

Can a light typing discipline be compatible with efficient implementation of finite field inversion ?

Daniele Canavese¹, Emanuele Cesena², Rachid Ouchary¹, Marco Pedicini³, and Luca Roversi⁴

¹ Politecnico di Torino, Dip. di Automatica e Informatica, Torino, Italy
daniele.canavese@polito.it, rachid.ouchary@polito.it

² Theneeds Inc., San Francisco, CA
ec@theneeds.com

³ Università degli Studi Roma Tre, Dipartimento di Matematica e Fisica, Roma, Italy
pedicini@mat.uniroma3.it

⁴ Università degli Studi di Torino, Dip. di Informatica, Torino, Italy
roversi@di.unito.it

Abstract. We show that an algorithm implementing the Binary-Field Arithmetic operation of *multiplicative inversion* exists as a purely functional term which is typeable in *Dual Light Affine Logic* (DLAL). As a consequence, the set Λ^{DLAL} of functional terms typeable in DLAL is large enough to program the whole set of arithmetic operations. Second, and most important, we show that Λ^{DLAL} can be seen a domain specific language that forces the programmer to think about algorithms under a non standard mental pattern which may results in more essential descriptions of known algorithms which, also, may be more efficient.

1 Introduction

In this paper we address the question if a functional programming approach can be of broader interest when implementing efficient arithmetics. The challenge is posed by a double front of constraints:

1. efficient arithmetics implementation is generally done by programming at architectural level even by keeping in account the running architecture,
2. algorithms are in the feasible range of the complexity bounds (*i.e.*, FPTIME) and even the polynomial degree in the known bounds is subject to full consideration.

The arithmetic over binary extension fields has many important applications in the domains of theory of codes and in cryptography. Finite fields' arithmetic operations include: addition, subtraction, multiplication, squaring, square root, multiplicative inverse, division and exponentiation.

Declarative programming, by its nature, does not permit a tight control on complexity parameters, however the scenario is changed in the last twenty years with the introduction of type systems with implicit complexity bounds where implicit means that the inclusion in the complexity class is certified by the type system: if the program is typeable then the type guarantees on execution time. The point in so far considered type systems is that the restrictions on programming schemes hardly permit to specify any algorithm even if it belongs to the right complexity class. Therefore a certain number of new type systems have been introduced in the last a few years with the declared objective to capture a broader class of polynomial algorithms with respect to the one which was shown to be in the previous systems.

Our pragmatic workplan is to make fully operational a declarative framework where feasible arithmetics can be programmed in a certified environment ensuring that complexity is polynomial, for this reason we introduced in [1] a variant of the system DLAL [2], that we called TFA (typeable functional assembly) having in mind what kind of programming patterns should be used in arithmetics. In fact, we would like to have an even improved control on these systems in order to certify polynomial computations up to a certain exponent, maybe as a development on the quantitative approach introduced in [3].

We build on our previous paper where we introduced basic materials in order to make arithmetics in finite fields by using a declarative language. Principal algorithms are known to be polynomial in

complexity nevertheless it was not an easy task to show their typeability in the formal system: and this is an obstacle in the use of such systems. On the other hand it is maybe because the unusual programming pattern that all these difficulties arise. It is because in these systems iteration is forced to add a marker in the type which blocks nesting of iterations. On the road we discovered and put in practice several patterns derived by classical Map or MapThread terms, and the result contained in this paper completes the work by proving that also the term for inversion in finite fields is in TFA.

The efficient implementation of inversion we started from is known as BEA (binary euclidean algorithm) in the realization given by Fong in [4], see Figure 1. It declines perfectly by the imperative programming toolbox: direct assignments of variables in memory, control flow in the form of a double nested iteration: a goto-statement which creates a loop and an internal while statement. It is obvious that in a declarative programming language there is no goto-statement and even while-loops are to be realised by structural iterations on some data type (typically Church integers). This was a first step in order to write the algorithm in a declarative style, the second one was to simulate those operations with direct access to data structures: this forced us to have and use a reverse of the binary sequence representing the number to invert and then to control the access to the head of the sequence. But the most challenging step was to cope with type constraints on variable duplications which oblige to a parsimonious attitude while programming, in the constant trying to approximate at the best, linear types: in fact, the point is to think like if terms would be linear terms (any variable is used once), and then very carefully relax to have non-linear variables. This is our main result: in Section 3 we

INPUT: $a \in \mathbb{F}_{2^m}, a \neq 0$.
OUTPUT: $a^{-1} \bmod f$.

1. $u \leftarrow a, v \leftarrow f, g_1 \leftarrow 1, g_2 \leftarrow 0$.
2. While z divides u do:
 - (a) $u \leftarrow u/z$.
 - (b) If z divides g_1 then $g_1 \leftarrow g_1/z$ else $g_1 \leftarrow (g_1 + f)/z$.
3. If $u = 1$ then return(g_1).
4. If $\deg(u) < \deg(v)$ then $u \leftrightarrow v, g_1 \leftrightarrow g_2$.
5. $u \leftarrow u + v, g_1 \leftarrow g_1 + g_2$.
6. Goto Step 2.

Fig. 1: Binay-Field inversion as in **Algorithm 2.2** at page 1048 in [4].

show a typeable multiplicative inversion in TFA. Then in Section 4 we go back to imperative: we give a new implementation of inversion which we called DCEA, the imperative version of the typeable term which implements inversion in TFA. DCEA is the imperative version of the control flow we had to realise in order to comply with typing in TFA. Note that in principle we could perform less efficiently because of the operation rearrangements. In fact, our experimental data show that even with our potentially longer loops we do not loose too much. On the other hand, and this is the main reason to compare the two algorithms in the imperative setting: the regular loops should be easier to be optimised by compiler and easier to be executed by speculative prediction of conditional branches. In Section 5, we have measures of the behaviour of our algorithm DCEA, where it is clear that the new version do not sacrifice efficiency for polynomial certification, since it follows very closely timings of BEA. Moreover, by taking timings on different architectures we obtained compatible results in all the cases and we outperformed execution time of BEA while using a machine with UltraSPARC architecture.

2 Typeable Functional Assembly

This work keeps developing the project started in [1] which introduces Typeable Functional Assembly (TFA). TFA is the type assignment for λ -terms in Figure 2. TFA is DLAL [2] whose set of formulas is quotiented by a specific recursive equation we shall recall from [1] in a few. Here we reintroduce the strictly necessary notions about TFA.

$\frac{}{\emptyset \mid \mathbf{x} : A \vdash \mathbf{x} : A} \text{ a}$	$\frac{\Delta \mid \Gamma \vdash \mathbf{M} : A}{\Delta, \Delta' \mid \Gamma, \Gamma' \vdash \mathbf{M} : A} \text{ w}$	$\frac{\Delta, \mathbf{x} : A, \mathbf{y} : A \mid \Gamma \vdash \mathbf{M} : B}{\Delta, \mathbf{z} : A \mid \Gamma \vdash \mathbf{M}\{z/x, z/y\} : B} \text{ c}$
$\frac{\Delta \mid \Gamma, \mathbf{x} : A \vdash \mathbf{M} : B}{\Delta \mid \Gamma \vdash (\lambda \mathbf{x}. \mathbf{M}) : A \multimap B} \multimap \text{I}$	$\frac{\Delta \mid \Gamma \vdash \mathbf{M} : A \multimap B \quad \Delta' \mid \Gamma' \vdash \mathbf{N} : A}{\Delta, \Delta' \mid \Gamma, \Gamma' \vdash (\mathbf{M}\mathbf{N}) : B} \multimap \text{E}$	
$\frac{\Delta, \mathbf{x} : A \mid \Gamma \vdash \mathbf{M} : B}{\Delta \mid \Gamma \vdash (\lambda \mathbf{x}. \mathbf{M}) : !A \multimap B} \Rightarrow \text{I}$	$\frac{\Delta \mid \Gamma \vdash \mathbf{M} : !A \multimap B \quad \emptyset \mid \Delta' \vdash \mathbf{N} : A \quad \Delta' \leq 1}{\Delta, \Delta' \mid \Gamma \vdash (\mathbf{M}\mathbf{N}) : B} \Rightarrow \text{E}$	
$\frac{\emptyset \mid \Delta, \Gamma \vdash \mathbf{M} : A}{\Delta \mid \S \Gamma \vdash \mathbf{M} : \S A} \S \text{I}$	$\frac{\Delta \mid \Gamma \vdash \mathbf{N} : \S A \quad \Delta' \mid \mathbf{x} : \S A, \Gamma' \vdash \mathbf{M} : B}{\Delta, \Delta' \mid \Gamma, \Gamma' \vdash \mathbf{M}\{N/x\} : B} \S \text{E}$	
$\frac{\Delta \mid \Gamma \vdash \mathbf{M} : A \quad \alpha \notin \text{fv}(\Delta, \Gamma)}{\Delta \mid \Gamma \vdash \mathbf{M} : \forall \alpha. A} \forall \text{I}$	$\frac{\Delta \mid \Gamma \vdash \mathbf{M} : \forall \alpha. A}{\Delta \mid \Gamma \vdash \mathbf{M} : A[B/\alpha]} \forall \text{E}$	

Fig. 2: Type assignment system TFA

Every judgment $\Delta \mid \Gamma \vdash \mathbf{M} : A$ comes with two different kinds of contexts. The formula A is assigned as type to the λ -term \mathbf{M} with hypothesis from the *polynomial context* Δ and the *linear context* Γ .

In TFA we define every single ground data type.

Let \mathcal{G} be a countable set of *variables*. We range over \mathcal{G} by *lowercase Greek letters*. Any type A belongs to the quotient \mathcal{F}/\approx of the following language \mathcal{F} of formulas:

$$F ::= \mathcal{G} \mid F \multimap F \mid !F \multimap F \mid \forall \mathcal{G}. F \mid \S F$$

We define the quotient on \mathcal{F} when introducing the *Sequences of booleans* here below. *Uppercase Latin letters* A, B, C, D will range over \mathcal{F}/\approx . *Modal formulas* $!A$ can occur in negative positions only. The notation $A[B/\alpha]$ is the clash free substitution of B for every free occurrence of α in A . Clash-free means that occurrences of free variables of B are not bound in $A[B/\alpha]$.

The λ -term \mathbf{M} belongs to Λ , the λ -calculus given by:

$$\mathbf{M} ::= \mathcal{V} \mid (\lambda \mathbf{x}. \mathbf{M}) \mid (\mathbf{M}\mathbf{M}) \quad (1)$$

with \mathcal{V} the set of variables we range over by *any lowercase Latin Teletype letter*. *Uppercase Teletype Latin letters* $\mathbf{M}, \mathbf{N}, \mathbf{P}, \mathbf{Q}, \mathbf{R}$ will range over Λ . We shall write $\lambda \mathbf{x}. \mathbf{M}$ in place of $(\lambda \mathbf{x}. \mathbf{M})$ in absence of ambiguity. Application $(\mathbf{M}_1 \mathbf{M}_2) \dots \mathbf{M}_n$ is left associative. We shall tend to shorten it as $\mathbf{M}_1 \mathbf{M}_2 \dots \mathbf{M}_n$. The set of free variables of \mathbf{M} is $\text{fv}(\mathbf{M})$. The standard β -reduction \rightarrow^* on λ -terms is the reflexive, transitive, and contextual closure of:

$$(\lambda \mathbf{x}. \mathbf{M}) \mathbf{N} \rightarrow \mathbf{M}\{N/x\} \quad (2)$$

Both polynomial and linear contexts are maps $\{\mathbf{x}_1 : A_1, \dots, \mathbf{x}_n : A_n\}$ from the domain of variables \mathcal{V} , to the co-domain of formulas \mathcal{F} . The difference between the two kinds of context is that variables in the polynomial context may occur an arbitrary number of times in the *subject* \mathbf{M} of $\Delta \mid \Gamma \vdash \mathbf{M} : A$. Every variable in the linear context must occur at most once in \mathbf{M} . Every pair $\mathbf{x} : A$ of any kind of context is a *type assignment for a variable*. The notation $\S \Gamma$ is a shorthand for $\{\mathbf{x}_1 : \S A_1, \dots, \mathbf{x}_n : \S A_n\}$, if Γ is $\{\mathbf{x}_1 : A_1, \dots, \mathbf{x}_n : A_n\}$.

Finite types are functions that project one argument out of the many they have:

$$\mathbb{B}_n \equiv \forall \alpha. \mathbb{B}_n[\alpha] \text{ with } \mathbb{B}_n[\alpha] \equiv \overbrace{\alpha \multimap \dots \multimap \alpha}^{n+1} \multimap \alpha .$$

Finite types with $n = 2$ are *lifted booleans* we denote by \mathbb{B}_2 . Their canonical representatives are:

$$1 \equiv \lambda \mathbf{x}. \lambda \mathbf{y}. \lambda \mathbf{z}. \mathbf{x} : \mathbb{B}_2 \quad 0 \equiv \lambda \mathbf{x}. \lambda \mathbf{y}. \lambda \mathbf{z}. \mathbf{y} : \mathbb{B}_2 \quad \perp \equiv \lambda \mathbf{x}. \lambda \mathbf{y}. \lambda \mathbf{z}. \mathbf{z} : \mathbb{B}_2$$

We need \perp in order to simplify the definition of the functions we want to program.

The type *Tuples* is:

$$(A_1 \dots \otimes A_n) \equiv \forall \alpha. (A_1 \otimes \dots \otimes A_n)[\alpha] \multimap \alpha \text{ with } (A_1 \otimes \dots \otimes A_n)[\alpha] \equiv A_1 \multimap \dots \multimap A_n \multimap \alpha$$

If we extend Definition (1) as follows:

$$\mathbf{M} ::= \dots \mid \langle \mathbf{M}, \dots, \mathbf{M} \rangle \mid \langle \mathcal{V}, \dots, \mathcal{V} \rangle. \mathbf{M} .$$

and we extend β -reduction:

$$(\lambda \langle \mathbf{x}_1, \dots, \mathbf{x}_n \rangle. \mathbf{M}) \langle \mathbf{N}_1, \dots, \mathbf{N}_n \rangle \rightarrow \mathbf{M}^{\langle \mathbf{N}_1 / \mathbf{x}_1, \dots, \mathbf{N}_n / \mathbf{x}_n \rangle} .$$

we can type the new terms with the following rules derivable in TFA:

$$\frac{\Delta_1 \mid \Gamma_1 \vdash \mathbf{M}_1 : A_1 \quad \dots \quad \Delta_n \mid \Gamma_n \vdash \mathbf{M}_n : A_n}{\Delta_1, \dots, \Delta_n \mid \Gamma_1, \dots, \Gamma_n \vdash \langle \mathbf{M}_1, \dots, \mathbf{M}_n \rangle : (A_1 \otimes \dots \otimes A_n)} \otimes_I \quad \frac{\Delta \mid \Gamma, \mathbf{x}_1 : A_1, \dots, \mathbf{x}_n : A_n \vdash \mathbf{M} : B}{\Delta \mid \Gamma \vdash \lambda \langle \mathbf{x}_1, \dots, \mathbf{x}_n \rangle. \mathbf{M} : (A_1 \otimes \dots \otimes A_n) \multimap B} \multimap_{I \otimes}$$

The *Sequences of booleans*, or simply Sequences, is the following recursive type:

$$\mathbb{S} \approx \forall \alpha. \mathbb{S}[\alpha] \text{ with } \mathbb{S}[\alpha] \equiv (\mathbb{B}_2 \multimap \alpha) \multimap ((\mathbb{B}_2 \otimes \mathbb{S}) \multimap \alpha) \multimap \alpha . \quad (3)$$

The recursive definition of \approx in (3) determines the equivalence relation \mathcal{F}/\approx among formulas of \mathcal{F} . As we said we take \mathcal{F} up to such a relation. I.e., if \mathbf{M} has type \mathbb{S} , then we can equivalently use any of the “unfolded forms” of \mathbb{S} as type of \mathbf{M} . The canonical values of type \mathbb{S} are:

$$[\varepsilon] \equiv \lambda \mathbf{t}. \lambda \mathbf{c}. \mathbf{t} \perp : \mathbb{S} \\ [\mathbf{b}_{n-1} \dots \mathbf{b}_0] \equiv \lambda \mathbf{t}. \lambda \mathbf{c}. \mathbf{c} \langle \mathbf{b}_{n-1}, [\mathbf{b}_{n-2} \dots \mathbf{b}_0] \rangle : \mathbb{S} . \quad (4)$$

In accordance with (3), the Sequence $[\mathbf{b}_{n-1} \dots \mathbf{b}_0]$ that occurs in (4) is a function that takes two constructors as inputs and yields a Sequence. Only the second constructor is used in (4) to build a Sequence out of a pair whose first element is \mathbf{b}_{n-1} , and whose second element is — recursively! — another Sequence $[\mathbf{b}_{n-2} \dots \mathbf{b}_0]$. It is well known that adding recursive equations among the formulas of DLAL is harmless as far as Polynomial time soundness is concerned [5, 2].

2.1 Basic types and (meta-)combinators

Church numerals. Their type is $\mathbb{U} \equiv \forall \alpha. \mathbb{U}[\alpha]$ where $\mathbb{U}[\alpha] \equiv !(\alpha \multimap \alpha) \multimap \mathbb{S}(\alpha \multimap \alpha)$ with canonical representatives:

$$\mathbf{u}\varepsilon \equiv \lambda \mathbf{f}. \lambda \mathbf{x}. \mathbf{x} : \mathbb{U} \quad \bar{n} \equiv \lambda \mathbf{f}. \lambda \mathbf{x}. \mathbf{f} (\dots (\mathbf{f} \mathbf{x}) \dots) : \mathbb{U} \text{ with } n \text{ occurrences of } \mathbf{f}$$

Lists. Their type is $\mathbb{L}(A) \equiv \forall \alpha. \mathbb{L}(A)[\alpha]$ where $\mathbb{L}(A)[\alpha] \equiv !(A \multimap \alpha \multimap \alpha) \multimap \mathbb{S}(\alpha \multimap \alpha)$ with canonical representatives:

$$\{\varepsilon\} \equiv \lambda \mathbf{f}. \lambda \mathbf{x}. \mathbf{x} : \mathbb{L}(A) \\ \{\mathbf{M}_{n-1} \dots \mathbf{M}_0\} \equiv \lambda \mathbf{f}. \lambda \mathbf{x}. \mathbf{f} \mathbf{M}_{n-1} (\dots (\mathbf{f} \mathbf{M}_0 \mathbf{x}) \dots) : \mathbb{L}(A) \text{ with } n \text{ occurrences of } \mathbf{f}$$

Church words $\{\mathbf{b}_{n-1} \dots \mathbf{b}_0\}$, with all \mathbf{b}_i s booleans, are the typical instance of lists with type $\mathbb{L}_2 \equiv \mathbb{L}(\mathbb{B}_2)$ we need. In every Church word $\{\mathbf{b}_{n-1} \dots \mathbf{b}_0\}$, or simply *word*, the *least significant bit* (l.s.b.) is \mathbf{b}_0 , while the *most significant bit* (m.s.b.) is \mathbf{b}_{n-1} . The same convention holds for every Sequence $[\mathbf{b}_{n-1} \dots \mathbf{b}_0]$.

The combinator Xor. It has type $\mathbb{B}_2 \multimap \mathbb{B}_2 \multimap \mathbb{B}_2$ and extends the *exclusive or* on lifted booleans:

$$\begin{array}{ll} \text{Xor } \mathbf{0} \ \mathbf{0} \rightarrow^* \ \mathbf{0} & \text{Xor } \mathbf{1} \ \mathbf{1} \rightarrow^* \ \mathbf{0} \\ \text{Xor } \mathbf{0} \ \mathbf{1} \rightarrow^* \ \mathbf{1} & \text{Xor } \mathbf{1} \ \mathbf{0} \rightarrow^* \ \mathbf{1} \\ \text{Xor } \perp \ \mathbf{b} \rightarrow^* \ \mathbf{b} & \text{Xor } \mathbf{b} \ \perp \rightarrow^* \ \mathbf{b} \end{array} \quad (\text{where } \mathbf{b} : \mathbb{B}_2).$$

Whenever one argument is \perp then it gives back the other argument. This is an application oriented choice [1].

The combinator wRev. It has type $\mathbb{L}_2 \multimap \mathbb{L}_2$ and *reverses the bits* of a word:

$$\mathbf{wRev} \{\mathbf{b}_{n-1} \dots \mathbf{b}_0\} \rightarrow^* \{\mathbf{b}_0 \dots \mathbf{b}_{n-1}\} .$$

Meta-combinator MapState[·]. Let $\mathbf{F} : (A \otimes S) \multimap (B \otimes S)$ be a closed term. Then, $\text{MapState}[\mathbf{F}] : \mathbb{L}(A) \multimap S \multimap \mathbb{L}(B)$ applies \mathbf{F} to the elements of the input list, keeping track of a *state* of type S during the iteration. If $(\mathbf{F}) \langle \mathbf{b}_i, \mathbf{s}_i \rangle \rightarrow^* \langle \mathbf{b}'_i, \mathbf{s}_{i+1} \rangle$, for every $0 \leq i \leq n-1$:

$$\text{MapState}[\mathbf{F}] \{\mathbf{b}_{n-1} \dots \mathbf{b}_0\} \mathbf{s}_0 \rightarrow^* \{\mathbf{b}'_{n-1} \dots \mathbf{b}'_0\} .$$

Meta-combinator $\text{MapThread}[\cdot]$. Let $F : \mathbb{B}_2 \multimap \mathbb{B}_2 \multimap A$ be a closed term. Then, $\text{MapThread}[F] : \mathbb{L}_2 \multimap \mathbb{L}_2 \multimap \mathbb{L}(A)$ applies F to the elements of the input list. If $((F) a_i) b_i \rightarrow^* c_i$, for every $0 \leq i \leq n - 1$:

$$\text{MapThread}[F] \{a_{n-1} \dots a_0\} \{b_{n-1} \dots b_0\} \rightarrow^* \{c_{n-1} \dots c_0\} .$$

In particular, $\text{MapThread}[\backslash a. \backslash b. \langle a, b \rangle] : \mathbb{L}_2 \multimap \mathbb{L}_2 \multimap \mathbb{L}(\mathbb{B}_2^2)$ is such that:

$$\text{MapThread}[\backslash a. \backslash b. \langle a, b \rangle] \{a_{n-1} \dots a_0\} \{b_{n-1} \dots b_0\} \rightarrow^* \{\langle a_{n-1}, b_{n-1} \rangle \dots \langle a_0, b_0 \rangle\} .$$

3 Multiplicative inversion as a term of TFA

Preliminaries. We introduce $\text{wHeadTail}[L, S]$ with parameters L and S . $\text{wHeadTail}[L, S]$ embodies the core mechanism of predecessor for Church numerals [6, 5] inside typing systems like TFA.

$$\begin{aligned} \text{wHeadTail}[L, S] &= \backslash w. \backslash f. \backslash x. L (w (\text{wHTStep } f) (\text{wHTBase } x)) \\ \text{wHTStep} &= \backslash f. \backslash e. \langle \text{ft}, \text{et}, \text{t} \rangle. (S f e \text{ft et t}) \\ \text{wHTBase} &= \backslash x. \langle \backslash e. \backslash 1. 1, \text{DummyElement}, x \rangle \end{aligned} \quad (5)$$

Both L and S must be closed λ -terms that play the role of parameters. S stands for “step” and L for “last”. For any A and α , let $X \equiv (A \multimap \alpha \multimap \alpha \otimes A \otimes \alpha)$. For any B , if $L: X \multimap B$ and $S: (A \multimap \alpha \multimap \alpha) \multimap A \multimap \alpha \multimap X$, then $\text{wHeadTail}: \mathbb{L}(A) \multimap \mathbb{L}(B)$. For example, wHeadTail can develop the following computation on a list $\backslash f. \backslash x. f b (f a x)$:

$$\begin{aligned} &\text{wHeadTail} (\backslash f. \backslash x. f b (f a x)) \\ &\rightarrow \backslash f. \backslash x. L ((\text{wHTStep } f) b ((\text{wHTStep } f) a (\text{wHTBase } x))) \\ &\rightarrow \backslash f. \backslash x. L (\langle \backslash \text{ft}, \text{et}, \text{t} \rangle. S f b \text{ft et t}) (S f a (\backslash e. \backslash 1. 1) \text{DummyElement } x) \end{aligned} \quad (6)$$

It is an iteration of wHTStep from wHTBase on the input. If DummyElement is different from any possible element of the list, the rightmost occurrence of S in (6) knows that the iteration is at its step zero and it can operate on a as consequence of this fact. In general, S can identify a sequence of iteration steps of predetermined length, say l . Then, S can operate on the first l elements of the list in a specific way. The distinguishing invariant of the computation pattern that wHeadTail develops is that S can have simultaneous stepwise access to two consecutive elements in the list. For example, S in (6) can use a and DummyElement at step zero. At step one it has access to b and et and the latter may contain a or some element derived from it. This invariant is crucial to implement a bitwise forwarding mechanism of the state in the term of TFA that implements the multiplication inverse.

L is the last operation we can perform on the final triple that the iteration of wHeadTail yields. For example, let:

$$\begin{aligned} L &= \backslash \langle _, _ \rangle. l \\ S &= \backslash f. \backslash e. \langle \backslash \text{ft}. \backslash e. \backslash \text{t}. \langle f, e, \text{ft et t} \rangle \end{aligned} \quad (7)$$

We implement a λ -term that pops the last element out of the input list. Developing from (5) we get:

$$\begin{aligned} &\dots \rightarrow \backslash f. \backslash x. L (\langle \backslash \text{ft}, \text{et}, \text{t} \rangle. S f b \text{ft et t}) (S f a (\backslash e. \backslash 1. 1) \text{DummyElement } x)) \\ &\rightarrow \backslash f. \backslash x. L (\langle \backslash \text{ft}, \text{et}, \text{t} \rangle. S f b \text{ft et t}) \langle f, a, (\backslash e. \backslash 1. 1) \text{DummyElement } x \rangle) \\ &\rightarrow \backslash f. \backslash x. L (\langle \backslash \text{ft}, \text{et}, \text{t} \rangle. S f b \text{ft et t}) \langle f, a, x \rangle) \\ &\rightarrow \backslash f. \backslash x. L (S f b f a x) \\ &\rightarrow \backslash f. \backslash x. L \langle f, b, f a x \rangle) \\ &\rightarrow \backslash f. \backslash x. f a x \end{aligned} \quad (8)$$

Main details about wInv. We reformulate BEA in Fig. 1 as a λ -term wInv of TFA as in Fig. 3. wInv starts building a list which it obtains by means of MapThread applied to eleven lists. For example, let $u = z^2$ and $v = z^3 + z + 1$ and $g_1 = 1$ and $g_2 = 0$ be an input of BEA. We can represent the polynomials as words:

$$\begin{aligned} U &= \backslash f. \backslash x. f 0 (f 1 (f 0 (f 0 x))) & V &= \backslash f. \backslash x. f 1 (f 0 (f 1 (f 1 x))) \\ G1 &= \backslash f. \backslash x. f 0 (f 0 (f 0 (f 1 x))) & G2 &= \backslash f. \backslash x. f 0 (f 0 (f 0 (f 0 x))) \end{aligned} \quad (9)$$

```

wInv = \U. # Word in input.
(wProj # Extract the bits of G1 from the resulting threaded word.
(D # Parameter of wInv. It is a Church numeral. Its value is
# the square of the degree n of the binary field.
(\tw.wRevInit (BkwVst (wRev (FwdVst tw)))) # Step function of D.
) (MapThread[\u.\v.\g1.\g2.\m.\stop.\sn.\rs.\fwdv.\fwdg2.\fwdm.
<u,v,g1,g2,m,stop,sn,rs,fwdv,fwdg2,fwdm>]
U [m_{n-1}...m_1 1] # V is a copy of the modulus.
[ 0... 0 1] # G1 with n components.
[ 0... 0 0] # G2 " " "
[m_{n-1}...m_1 1] # M is a copy of the modulus.
[ 0... 0 0] # Stop with n components.
[ B... B B] # StpNmbr " " "
[ B... B B] # RghtShft " " "
[ 0... 0 0] # Fwfv " " "
[ 0... 0 0] # FwdG2 " " "
[ 0... 0 0] # FwdM " " "
) # Base function of D.
)
#
# LEGENDA
# Meaning | Text abbreviation | Name of variable
# -----
# Step number | StpNmbr | sn
# Right shift | RghtShft | rs
# Forwarding of V | FwdV | fwdv
# Forwarding of G2 | FwdG2 | fwdg2
# Forwarding of F | FwdM | fwdm

```

Fig. 3: Definition of wInv.

wInv builds an initial list by applying MapThread to the four words in (9) and to further seven words which play the role of a state of the computation. In our example, the whole initial list is:

```

wInvInput =
\f.\x.# |----- This is a state -----|
# v v
# U V G1 G2 M Stop StpNmb RghtShft FwdV FwdG2 FwdM
f <0,1, 0, 0,1, 0, B, B, 0, 0, 0> # msb
(f <1,0, 0, 0,0, 0, B, B, 0, 0, 0>
(f <0,1, 0, 0,1, 0, B, B, 0, 0, 0>
(f <0,1, 1, 0,1, 0 B, B, 0, 0, 0> # lsb
x)))

```

We call *threaded words* the initial list that wInv builds as its first step which is customary to fix some notation and terminology on. Every *element* of the threaded words is a tuple of booleans. Let $U[i]$ be the bit of column U in the i th element of wInvInput. Analogous notation on V, G1, etc. exists. We write $\langle V, \dots, M \rangle[i]$, or $\langle V[i], \dots, M[i] \rangle$ to denote the projection of the bits in column V, G1, G2 and M out of the i th element in wInvInput. Analogous notation holds for arbitrary sub-sequences we need to project out of U, ..., FwdM. The most significant bit msb of wInvInput is on top; its less significant bit lsb is at the bottom.

The term D is a Church numeral. It iterates $\backslash tw.wRevInit (BkwVst (wRev (FwdVst tw)))$, a step function, starting from wInvInput. The step function implements steps 2 through 5 of BEA in Fig. 1. The iteration by D implements the outermost loop that starts at step 2 and stops at step 6. FwdVst shortens *forward visit*. wRev reverses the threaded words it takes as input. BkwVst stands for *backward visit*. wRevInit reverses the threaded words it gets in input while reinitializing the bits in positions StpNmb, RghtShft, FwdV, FwdG2 and FwdM.

The implementation of `FwdVst` follows the pattern of `wHeadTail`. It can distinguish its step zero, and its last step. Yet, for every $0 < i \leq \text{msb}$, `FwdVst` has access to elements $\langle U, V, \dots, \text{FwdM} \rangle [i]$ and $\langle U, V, \dots, \text{FwdM} \rangle [i-1]$.

The identification of step zero allows `FwdVst` to simultaneously check which of the following mutually exclusive questions has a positive answer:

“Is $\text{Stop}[0]=1$?” (11)

“Does z divide both u and g_1 ?” (12)

“Does z divide u but not g_1 ?” (13)

“Neither of the previous questions has positive answer?” (14)

If (11) holds, `FwdVst` must behave as the identity. Such a situation is equivalent to saying that all the bits in position $G1$ contain the result.

Let us assume instead that (12) or (13) hold. Answering the first question requires to verify $U[0]=0$ and $G1[0]=0$. Answering the second one needs to check both $U[0]=0$ and $G1[0]=1$. Under our conditions, just after reading `wInvInput`, `FwdVst` generates the following first element of the whole threaded words it has to build:

$$\langle U[0], B, X, B, B, B, 0, Y, \text{FwdV}[0], \text{FwdG2}[0], \text{FwdM}[0] \rangle \quad (15)$$

If (12) holds, then X is $G[0]$ and Y is 1. If (13) holds, then X is $Xor\ G[0]\ M[0]$ and Y is 0. It is worth remarking that we record $V[0]$, $G2[0]$ and $M[0]$ in $\text{FwdV}[0]$, $\text{FwdG2}[0]$ and $\text{FwdM}[0]$. Finally, we replace them with B .

After the generation of the first element (15), for every $0 < i \leq \text{msb}$, the iteration that `FwdVst` implements proceeds as follows. Two elements constitute the input at step i :

$$\begin{aligned} &\langle U, V, G1, G2, M, \text{Stop}, \text{StpNmbr}, \text{RghtShft}, \text{FwdV}, \text{FwdG2}, \text{FwdM} \rangle [i] \\ &\langle U, V, G1, G2, M, \text{Stop}, \text{StpNmbr}, \text{RghtShft}, \text{FwdV}, \text{FwdG2}, \text{FwdM} \rangle [i-1] \end{aligned} \quad (16)$$

From them, `FwdVst` generates the following element which we see as being the output at step i :

$$\langle U[i], \text{FwdV}[i-1], X, \text{FwdG2}[i-1], \text{FwdM}[i-1], B, 0, Y, V[i], G2[i], M[i] \rangle \quad (17)$$

Yet, X and Y depend on u and g_1 being divisible by z . The element (17) becomes the $i-1$ th element in the succeeding step.

Finally, under the above condition that (12) or (13) hold, the last step of `FwdVst` adds two elements to the threaded words it builds. If 1 is the length of the threaded words that `FwdVst` starts from, the two elements are:

$$\begin{aligned} &\langle U[1], \text{FwdV}[1-1], X, \text{FwdG2}[1-1], \text{FwdM}[1-1], B, 0, Y, V[1], G2[1], M[1] \rangle \\ &\langle 0, \text{FwdV}[1], 0, \text{FwdG2}[1], \text{FwdM}[1], B, 0, Y, B, B \rangle \end{aligned} \quad (18)$$

where X and Y keeps depending on which between (12) or (13) hold.

Even though this might sound a bit paradoxically, the overall effect of iterating the process we have just described — the one which exploits the simultaneous access to two elements of a list to propagate some bits from step $i-1$ to step i — amounts to shifting the bits in positions V , $G2$ and M one step to their *left*. Leaving the bits of position U and $G1$ as they were will result in shifting them one step to their right. We shall clarify how and why when we describe `BkwdVst`. Roughly, only a correct concatenation of both `FwdVst` and `BkwdVst` shifts to the right every $U[i]$ and $G1[i]$, or $Xor\ G1[i]\ M[i]$ preserving the position of every other element.

The description of how `FwdVst` works concludes by assuming that neither (12) nor (13) hold. This occurs when $U[0]=1$. `FwdVst` must forcefully answer to: “Is u different from 1?”. Answering the question requires a complete visit of the threaded words that `FwdVst` takes in input. The visit serves to verify whether some $j > 0$ exists such that $U[j]=1$. The non existence of j implies that `FwdVst` sets $\text{Stop}[\text{msb}]=1$. This will impede any further change of any bit in any position of the threaded words

The threaded words (19) is the input of $wRev$ which yields the following instance of $wBkwdVstInput$:

```
\f.\x. # U V      G1 G2 M Stop StpNmb RghtShft FwdV FwdG2 FwdM
      f <0,B,Xor 1 1, 0,1, . B,      0,      0, 1, 0, 1> # original lsb.
      (f <0,1,Xor 0 1, 0,1, B,      0,      0, 1, 0, 1> # new lsb.
      (f <1,1,Xor 0 0, 0,0, B,      0,      0, 0, 0, 0>
      (f <0,0,Xor 0 0, 0,1, B,      0,      0, 1, 0, 1>
      (f <0,1,      0, 0,1, B,      B,      0, B, B, B> # msb
                                          x))))
```

$BkwdVst$ applies to (20). It finds that $Stop[0]=B$ and $RghtShft[0]=0$ which requires to shift all the bits of U and $G1$ one position to the their right. $BkwdVst$ commits the requirement by erasing the topmost element of (20). The result is:

```
\f.\x. # U V      G1 G2 M Stop StpNmb RghtShft FwdV FwdG2 FwdM
      f <0,1,Xor 0 1, 0,1, B,      0,      0, 1, 0, 1>
      (f <1,1,Xor 0 0, 0,0, B,      0,      0, 0, 0, 0>
      (f <0,0,Xor 0 0, 0,1, B,      0,      0, 1, 0, 1>
      (f <0,1,      0, 0,1, B,      B,      0, B, B, B> x)))
```

Finally, $wRevInit$ reverses (21), yielding:

```
\f.\x. # U V      G1 G2 M Stop StpNmb RghtShft FwdV FwdG2 FwdM
      f <0,1,      0, 0,1, B,      B,      B, 0, 0, 0>
      (f <0,0,Xor 0 0, 0,1, B,      B,      B, 0, 0, 0>
      (f <1,1,Xor 0 0, 0,0, B,      B,      B, 0, 0, 0>
      (f <0,1,Xor 0 1, 0,1, B,      B,      B, 0, 0, 0> x)))
```

Let us compare (22) and (19). All the bits of position U and $G1$ have been shifted while those ones of position V , $G2$ and M have not. Moreover, the bits of position $Stop$, \dots , $FwdM$ have been reinitialized so that (22) is a consistent input for $FwdVst$. We remark that the whole process of shifting the bits of positions U and $G1$ requires the concatenation of both $FwdVst$ and $BkwdVst$ up to some reverse. The first one shifts the bits of position V , $G2$ and M to the left while operates on those of position U and $G1$. The latter drops the correct element and fully realizes the shift to the right we need.

The code of $FwdVst$ and of $BkwdVst$. We recall that $FwdVst$ and $BkwdVst$ follow the programming pattern of $wHeadTail$. The step functions they relies on and their “last step functions”, implement branching. Choices depend on the values of the bits that belong to the state or on the values of bits in position U and $G1$.

Because of space constraints we include $LastStepFwdVst$ only. It is the last step of $FwdVst$. The aim is to give an idea of the branching structure of the terms that $wInv$ contains. We make use of simple syntactic sugar to introduce $LastStepFwdVst$. Let N be of type \mathbb{B} . Then $N M1 M0 MB$ is a λ -term which eventually chooses among $M1$, $M0$ and MB , depending on the normal form N evaluates to. For the sake of readability, we represent $N M1 M0 MB$ by means of a `switch`-structure:

```
switch (N) {
  case 1: ...
  case 0: ...
  case B: ... }
```

$LastStepFwdVst$ is in Fig 5. The name `stopt` recalls “Stop of the tail”, i.e. “Stop that comes from step $msb-1$ ”. Analogously `rst` is “RghtShft that comes from step $msb-1$ ”. Let us focus on the two branches with `stopt=B` and `rst=1` or `rst=0`. They take care of the situations that require the shift to the right of U and $G1$. I.e., if we think in general terms, they generate the two elements in (18). If we prefer to think in terms of our example, they generate the two topmost elements in (19). We remark that $LastStepFwdVst$ is completely linear. Branching after branching it yields a λ -abstraction that correctly builds required elements that complete the threaded words under construction.

```

LastStepFwdVst =
\f.
\<ft,et,t>. # Element from step i-1.
(\<ut,vt,g1t,g2t,mt,stopt,snt,rst,fwdvt,fwdg2t,fwdmt>.
 (switch (stopt) {
   case 1: # We checked U=1. The whole wInv must be the identity.
     \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
     (ft <ut,vt,g1t,g2t,mt,1,B,B,B,B> t)
   case 0: # So we have also RghtShft=B and U[0]=1.
     switch (snt) {
       case 1: # U is different from 1.
         \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
         (ft <ut,vt,g1t,g2t,mt,0,1,B,B,B,B> t )
       case 0: # Here we detect that U=1 and we set Step=1 !!!!
         \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
         (ft <ut,vt,g1t,g2t,mt,1,B,B,B,B> t )
       case B: # Can never occur.
         \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
         (ft <ut,vt,g1t,g2t,mt,0,B,B,B,B> t )
     }
   case B: # We have U[0]=0 and RghtShft=0 or RghtShft=1.
     switch (rst) {
       case 1: # U[0]=0 and G1[0]=0. We are shifting and we have to add
         # a new msb to the threaded words.
         \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
         (f <0,vt,0,g2t,mt,B,B,1,B,B,B >
          (ft <ut,fwdvt,g1t,fwdg2t,fwdmt,B,snt,1,B,B,B> t ))
       case 0: # U[0]=0 and G1[0]=1. We are shifting and we have to add
         # a new msb to the threaded words.
         \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
         (f <0,vt,0,g2t,mt,B,B,0,B,B,B >
          (ft <ut,fwdvt,g1t,fwdg2t,fwdmt,B,snt,0,B,B,B> t ))
       case B: # Can never occur.
         \f.\ft.\ut.\vt.\g1t.\g2t.\mt.\snt.\rst.\fwdvt.\fwdg2t.\fwdmt.\t.
         (ft <ut,vt,g1t,g2t,mt,B,B,B,B,B> t )
     }
   }
) f ft ut vt g1t g2t mt snt rst fwdvt fwdg2t fwdmt t
) et

```

Fig. 5: Definition of LastStepFwdVst.

4 An imperative version of $wInv$ and its correctness

In Section 3, we reformulate the imperative algorithm BEA in [4] as a functional program which TFA can give a type to. In this section we go in the opposite direction. From $wInv$ we synthesize an imperative program DCEA (DLAL Certified Euclidean Algorithm). $wInv$ suggests how to reorganize some steps of BEA so that DCEA becomes a slight variation of BEA.

In Table 1, we present DCEA and BEA side by side. The outermost loop exists in DCEA as well. The new variable *side* in DCEA takes Fwd and Bkwd as its values. When *side* is Fwd, DCEA mimics FwdVst as defined in $wInv$; it checks if z divides at least u and behaves accordingly. If true, *side* gets set to Bkwd. Then, the block B completes the shift to the right of u and g_1 exactly as BkwdVst would do in $wInv$. On the contrary, if *side* is Fwd and z does not divide u , then *side* eventually becomes Bkwd. Then DCEA compares the degrees of u and v exactly like BkwdVst in $wInv$. Finally, DCEA checks $u = 1$ just before switching from *side* = Fwd to *side* = Bkwd. We recall that this is what FwdVst does when z does not divide u and the computation is going to be controlled by BkwdVst.

INPUT: $a \in \mathbb{F}_{2^m}, a \neq 0$

OUTPUT: $a^{-1} \bmod f$

```

1:  $u \leftarrow a, v \leftarrow f, g_1 \leftarrow 1, g_2 \leftarrow 0, side \leftarrow \text{Fwd}$ 
2: if  $side = \text{Fwd}$  then
3:   if  $z$  divides  $u$  and  $z$  do not divides  $g_1$  then (A)
4:      $g_1 \leftarrow g_1 + f$ 
5:   end if
6:    $side \leftarrow \text{Bkwd}$ 
7: else
8:   if  $z$  divides  $u$  then (B)
9:      $u \leftarrow u/z; g_1 \leftarrow g_1/z$ 
10:  else
11:    if  $\deg(u) < \deg(v)$  then (C)
12:       $u \leftrightarrow v; g_1 \leftrightarrow g_2$ 
13:    end if
14:     $u \leftarrow u + v; g_1 \leftarrow g_1 + g_2$ 
15:  end if
16:   $side \leftarrow \text{Fwd}$ 
17: end if
18: if  $u = 1$  then (D)
19:   return  $g_1$ 
20: end if
21: go to 2

```

(a) DCEA.

INPUT: $a \in \mathbb{F}_{2^m}, a \neq 0$

OUTPUT: $a^{-1} \bmod f$

```

1:  $u \leftarrow a, v \leftarrow f, g_1 \leftarrow 1, g_2 \leftarrow 0$ 
2: while  $z$  divides  $u$  do
3:    $u \leftarrow u/z$ 
4:   if  $z$  divides  $g_1$  then (E)
5:      $g_1 \leftarrow g_1/z$ 
6:   else
7:      $g_1 \leftarrow (g_1 + f)/z$ 
8:   end if
9: end while
10: if  $u = 1$  then (F)
11:   return  $g_1$ 
12: end if
13: if  $\deg(u) < \deg(v)$  then (G)
14:    $u \leftrightarrow v; g_1 \leftrightarrow g_2$ 
15: end if
16:  $u \leftarrow u + v; g_1 \leftarrow g_1 + g_2$ 
17: go to 2

```

(b) BEA.

Table 1: Algorithms that compute the multiplicative inversion for a given binary finite field.

Section 5 gives an account of experiments on the performance of DCEA, BEA and OpenSSL. We anticipate that, despite DCEA might look a trivial syntactical modification of BEA, it can behave better than BEA, and outperforms OpenSSL.

We now give evidence about the equivalence of DCEA and BEA. Proving the equivalence of DCEA and BEA amounts to show that they behave the same, starting from identical pre-conditions. Two significant cases exists.

- Let $z \mid u$, i.e. let z divide u . BEA iterates E in the loop that starts at line 2 until $z \nmid u$. Under the same pre-condition DCEA iterates the sequence A, D, B, D until $z \nmid u$, i.e. until z does not divide u . In both cases it should be evident that the post-conditions coincide.
- Let $u \neq 1$ and $z \nmid u$. BEA skips both E and F to execute G. DCEA skips A, D, and B to finally interpret C. The set of instructions in G and C coincide.

As soon as DCEA and BEA check $u = 1$ they are obviously equivalent because neither of the two alter any among u, v, g_1 and g_2 .

5 Experimental results

Calculating the multiplicative inverse of a finite field is typically the most time-consuming operation. A high-speed implementation is frequently much more than desirable, since its running-time can have a significant impact on several cryptographic algorithms and protocols. We present the results of a number of tests aiming to compare the execution time of DCEA, BEA and the implementation of the inversion that the OpenSSL library [7] supplies.

We performed our tests on four different hardware architectures whose specifications are given in Table 2. We remark that the “Core2”, “Pentium4” and “UltraSPARC” platforms are PC microprocessors while the “ATmega328” is a microcontroller. Note also that for the tests we used an Arduino Duemilanove⁵ prototype board which integrates an ATmega328.

CPU	Clock	Word size	Compiler	Optimization flags
Intel Core2 Duo P8600	2.40GHz	64 bit	gcc 4.7.3	-O4
Sun UltraSPARC IIe	500 MHz	64 bit	gcc 4.2.1	-O4
Intel Pentium4	3.00GHz	32 bit	gcc 4.4.5	-O4
ATmega328	16MHz	16 bit	avr-gcc 4.3.2	-Os

Table 2: Specifications of the test platforms.

We compare the performances of the aforementioned algorithms on the following binary fields:

- the field $\mathbb{F}_{2^{61}}[z]/(z^{61} + z^5 + z^2 + z + 1)$;
- the field $\mathbb{F}_{2^{163}}[z]/(z^{163} + z^7 + z^6 + z^3 + 1)$, recommended in the standard FIPS 186-3 [8];
- the field $\mathbb{F}_{2^{233}}[z]/(z^{233} + z^{74} + 1)$, also recommended in the standard FIPS 186-3 [8].

We recall that the field degrees 61, 163 and 233 are the dimension of the field, that is the length n of the threaded words we uniformly manipulate with `wInv` in Section 3.

5.1 Comparing the running-time of DCEA and BEA

In the following paragraphs, we briefly explain our findings regarding the performances of the DCEA with respect to the BEA.

In this context, the term “*generic* algorithm” refers to an implementation of a multiplicative inversion which can operate on any field $\mathbb{F}_{2^n}[z]/(m(z))$, with a generic degree n and modulus $m(z)$. On the other hand, the term “*specialized* algorithm” identifies a specific implementation that can only work on a given field $\mathbb{F}_{2^n}[z]/(m(z))$, with a fixed degree n and modulus $m(z)$. Note that fixing n and $m(z)$ *a-priori* usually allows to better optimize the code, so a specialized algorithm is typically faster than its generic counterpart.

For instance, the Listing 1.1 shows the C code used to test the DCEA algorithm over $\mathbb{F}_{2^{61}}$ on the Core2 platform. The snippet does not precisely reflect the pseudo-code listed in Table 1a since several ad-hoc optimizations were applied, such as:

- elimination of the selection on the *side* variable;
- elimination of the selection in the block A of Table 1a via a Shannon expansion (lines 31 – 32);
- the comparison between the polynomial degrees is computed using a simple and efficient integer comparison (line 37).

```

1  #define SHIFT(i) \
2    tmp1 = u[i] >> 1; \
3    tmp2 = (u[i + 1] & 1) << 63; \
4    tmp3 = g1[i] >> 1; \
5    tmp4 = (g1[i + 1] & 1) << 63; \
6    u[i] = tmp1 ^ tmp2; \
7    g1[i] = tmp3 ^ tmp4
8
9  #define SUM(i) \
10   u[i] ^= v[i]; \
11   g1[i] ^= g2[i]
12
13 #define SWAP_SUM(i) \
14   u[i] ^= v[i]; \
15   g1[i] ^= g2[i]; \
16   v[i] ^= u[i]; \
17   g2[i] ^= g1[i]
18
19 void dcea61(uint64_t r[], uint64_t a[])
20 {
21   uint64_t u[2] = { a[0] };
22   uint64_t v[2] = { (11 << 61) ^ (1 << 5) ^ (1 << 2) ^ (1 << 1) ^ 1 };
23   uint64_t g1[2] = { 1 };
24   uint64_t g2[2] = { 0 };
25   uint64_t x0;
26   uint64_t tmp1, tmp2, tmp3, tmp4;
27
28   do
29     if ((u[0] & 1) == 0)
30     {
31       x0 = ~((g1[0] & 11) - 11);
32       g1[0] ^= ((11 << 61) ^ (1 << 5) ^ (1 << 2) ^ (1 << 1) ^ 1) & x0;
33       SHIFT(0);
34       u[1] >>= 1;
35       g1[1] >>= 1;
36     }
37     else if (u[1] < v[1] || u[0] < v[0])
38     {
39       SWAP_SUM(0);
40       SWAP_SUM(1);
41     }
42     else
43     {
44       SUM(0);
45       SUM(1);
46     }
47   while (u[0] != 1 || u[1] != 0);
48
49   r[0] = g1[0];
50 }

```

Listing 1.1: The DCEA algorithm C code specialized for the field $\mathbb{F}_{2^{61}}$ and the Core2 architectures.

Algorithm	Core2		Pentium4		UltraSPARC		ATmega328	
	BEA	DCEA	BEA	DCEA	BEA	DCEA	BEA	DCEA
\mathbb{F}_{261} : generic algorithm	2.63	2.64	6.10	6.15	10.52	10.89	7596.13	7685.26
\mathbb{F}_{261} : specialized algorithm	0.70	0.78	1.45	1.47	2.99	2.58	1108.59	1159.87
\mathbb{F}_{2163} : generic algorithm	25.14	25.43	41.77	41.48	102.11	107.29	59755.80	60160.20
\mathbb{F}_{2163} : specialized algorithm	5.92	5.94	13.77	13.78	38.54	30.01	15250.45	15504.68
\mathbb{F}_{2233} : generic algorithm	46.34	46.83	91.22	91.23	199.23	205.62	122471.00	123576.85
\mathbb{F}_{2233} : specialized algorithm	10.83	10.88	30.30	30.75	86.34	82.39	36039.93	36695.22

Table 3: Average time in μs that BEA and DCEA take to compute an inversion.

In Table 3, we compare the average running-time in microseconds (μs) that DCEA and BEA need to calculate a single inversion on a given field. If we exclude the UltraSPARC architecture, DCEA can be effectively seen as a variant of BEA. Both programs perform essentially equally with a slight prevalence of BEA. However, on UltraSPARC, the specialized version of DCEA is significantly faster than the BEA. It achieves a 1.28 speedup factor over \mathbb{F}_{2163} .

5.2 Comparing the running-time of DCEA and OpenSSL

OpenSSL [7] is one of the most used and widespread cryptographic toolkits for securing a network communication, especially in the open-source world. Table 4 shows the average running-times of OpenSSL 1.0.1e and DCEA, both generic and specialized, to calculate a single inversion on the given field on an Intel Core2 CPU. The time is in microseconds (μs).

Field	OpenSSL	DCEA	
		Generic	Specialized
\mathbb{F}_{261}	3.97	2.64	0.78
\mathbb{F}_{2163}	13.05	25.43	5.90
\mathbb{F}_{2233}	20.58	46.83	10.88

Table 4: Average time in μs that OpenSSL algorithm and DCEA take to compute an inversion.

In Fig. 6, we graphically depicts the data in Table 4 to better emphasize the differences among the performances of OpenSSL 1.0.1e and DCEA.

We remark that OpenSSL implements MAIA (Modified Almost Inverse Algorithm) which is a generic version of BEA. The specialized versions of DCEA always outperform the OpenSSL, especially on the small fields. The generic version of DCEA seems faster only on \mathbb{F}_{261} . In this case it is twice as fast. This is an interesting result since several algebraic structures, such as hyperelliptic curves, can guarantee a high level of security using fields with a low order.

6 Conclusions and future work

We introduce the functional program `wInv`. It has type in TFA and it implements the multiplicative inverse in a given binary field of arbitrary, but fixed, degree. We complete a project started in [1]. The goal of [1] is to implement a library of potential real interest by using a language conceived in the ambit of Implicit Computational Complexity (ICC). We succeeded in spite of the widespread opinion

⁵ <http://arduino.cc/en/Main/ArduinoBoardDuemilanove>

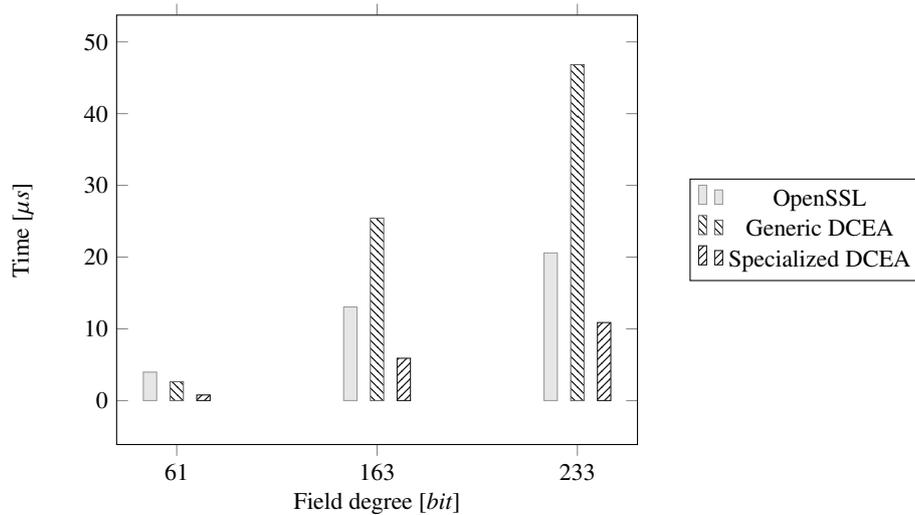


Fig. 6: Comparing the implementation of inversion in OpenSSL and DCEA.

that the expressivity of languages like the one we used is too weak to program anything interesting. By the way, we remark that the existence of $wInv$ gives an alternative proof that inversion has a polynomial cost.

We further strengthen the evidence that programming with a language with its roots inside ICC, like the λ -terms typed by TFA are, can be rewarding. From $wInv$ we synthesize the iterative algorithm DCEA which implements the inversion. Tests on typical hardware architectures show that DCEA is essentially as fast as BEA (DCEA is about 1% slower than BEA). Instead, tests on an UltraSPARC architecture show that DCEA performs better than BEA. One of our future goals is to certify that this is not a singularity and to obtain clues at why DCEA seems to perform better than BEA on UltraSPARC architectures.

A second goal is to assess to which extent we can exploit the functional programming patterns we use in the definition of DCEA to design a (specific domain) imperative language. The hope is to design one characterized by two seemingly contrasting features. It should be expressive enough in terms of programs we can write but its constructs would be limited enough to imply the application of good compiling techniques and optimizations able to fully exploit architecture like, for example, UltraSPARC.

References

1. Cesena, E., Pedicini, M., Roversi, L.: Typing a Core Binary-Field Arithmetic in a Light Logic. In Peña, R., van Eekelen, M., Shkaravska, O., eds.: Foundational and Practical Aspects of Resource Analysis (subtitle: 2nd International Workshop on Foundational and Practical Aspects of Resource Analysis, FOPARA 2011). Volume 7177 of LNCS., Springer (2012) 19 – 35
2. Baillot, P., Terui, K.: Light types for polynomial time computation in lambda calculus. *I&C* **207**(1) (2009) 41–62
3. Dal Lago, U.: Context semantics, linear logic, and computational complexity. *ACM Trans. Comput. Log.* **10**(4) (2009) Art. 25, 32
4. Fong, K., Hankerson, D., Lopez, J., Menezes, A.: Field inversion and point halving revisited. *IEEE Trans. Comput.* **53**(8) (2004) 1047–1059
5. Asperti, A., Roversi, L.: Intuitionistic light affine logic. *ACM ToCL* **3**(1) (2002) 1–39
6. Roversi, L.: A P-Time Completeness Proof for Light Logics. In: Ninth Annual Conference of the EACSL (CSL'99). Volume 1683 of Lecture Notes in Computer Science., Madrid (Spain), Springer-Verlag (September 1999) 469 – 483
7. The OpenSSL team: OpenSSL: The Open Source toolkit for SSL/TLS (2013)
8. National Institute of Standards and Technology: FIPS PUB 186-3 FEDERAL INFORMATION PROCESSING STANDARDS PUBLICATION Digital Signature Standard (DSS) (June 2009)