

Backtracking-assisted multiplication

Houda Ferradi¹ · Rémi Géraud¹  · Diana Maimuț¹ ·
David Naccache¹ · Hang Zhou¹

Received: 25 November 2016 / Accepted: 13 August 2017 / Published online: 5 September 2017
© Springer Science+Business Media, LLC 2017

Abstract This paper describes a new multiplication algorithm, particularly suited to lightweight microprocessors when one of the operands is known in advance. The method uses backtracking to find a multiplication-friendly encoding of the operand known in advance. A 68HC05 microprocessor implementation shows that the new algorithm indeed yields a twofold speed improvement over classical multiplication for 128-byte numbers.

Keywords Multiplication · Integer arithmetics · Backtracking

Mathematics Subject Classification (2010) 90-08 · 90C39 · 90C90

1 Introduction

A number of applications require performing long multiplications in performance-restricted environments. Indeed, low-end devices such as the 68HC05 or the 80C51 microprocessors

This article is part of the Topical Collection on *Recent Trends in Cryptography*

✉ Rémi Géraud
remi.geraud@ens.fr

Houda Ferradi
houda.ferradi@ens.fr

Diana Maimuț
diana.maimut@ens.fr

David Naccache
david.naccache@ens.fr

Hang Zhou
hang.zhou@ens.fr

¹ Département d'informatique de l'ÉNS, École normale supérieure, PSL Research University, Paris, France

have a very limited instruction-set, very limited memory, and operations such as multiplication are rather slow: a mul instruction typically claims 10 to 20 cycles.

General multiplication has been studied extensively, and there exist algorithms with very good asymptotic complexity such as the Schönhage-Strassen algorithm [18] which runs in time $O(n \log n \log \log n)$ or the more recent Fürer algorithm [13], some variants of which achieve the slightly better $O(2^{3 \log^* n} n \log n)$ complexity [14]. Such algorithms are interesting when dealing with extremely large integers, where these asymptotics prove faster than more naive approaches.

In many cryptographic contexts however, multiplication is performed between a variable and a *pre-determined constant*:

- During Diffie-Hellman key exchange [9] or ElGamal [10] a constant g must be repeatedly multiplied by itself to compute $g^x \bmod p$.
- The essential computational effort of a Fiat-Shamir prover [11, 12] is the multiplication of a subset of fixed keys (denoted s_i in [11]).
- A number of modular reduction algorithms use as a building-block multiplications (in \mathbb{N}) by a constant depending on the modulus. This is for instance the case of Barrett’s algorithm [2] or Montgomery’s algorithm [17].

The main strategy to exploit the fact that one operand is constant consists in finding a decomposition of the multiplication into simpler operations (additions, subtractions, bit-shifts) that are hardware-friendly [3]. The problem of finding the decomposition with the least number of operations is known as “single constant multiplication” (SCM) problem. $\text{SCM} \in \text{NP-complete}$ as shown in [4], even if fairly good approaches exist [1, 7, 8, 20] for small numbers. For larger numbers, performance is unsatisfactory unless the constant operand has a predetermined format allowing for ad hoc simplifications.

In this paper, we propose a completely different approach: the constant operand is encoded in a computation-friendly way, which makes multiplication faster. This encoding is based on linear relationships detected amongst the constant’s digits (or, more generally, subwords), and can be performed offline in a reasonable time for 1024-bit numbers and 8-bit microprocessors. We use a graph-based backtracking algorithm [16] to discover these linear relationships, using recursion to keep the encoder as short and simple as possible.

2 Multiplication algorithms

We now provide a short overview of popular multiplication methods. This summary will serve as a baseline to evaluate the new algorithm’s performance.

Multiplication algorithms usually fall in two broad categories: general divide-and-conquer algorithms such as Toom-Cook [6, 19] and Karatsuba [15]; and the generation of integer multiplications by compilers, where one of the arguments is statically known. We are interested in the case where small-scale optimizations such as Bernstein’s [3] are impractical, but general purpose multiplication algorithms à la Toom-Cook are not yet interesting.

Throughout the paper we will assume unsigned integers, and denote by w the word size (typically, $w = 8$), a_i , b_i and r_i the binary digits of a , b and r respectively:

$$a = \sum_{i=0}^{n-1} 2^{wi} a_i, \quad b = \sum_{i=0}^{n-1} 2^{wi} b_i, \quad \text{and} \quad r = a \times b = \sum_{i=0}^{2n-1} 2^{wi} r_i.$$

2.1 Textbook multiplication

A direct way to implement long multiplication consists in extending textbook multiplication to several words. This is often done by using a MAD¹ routine.

A MAD routine takes as input four n -bit words $\{x, y, c, \rho\}$, and returns the two n -bit words c', ρ' such that $2^n c' + \rho' = x \times y + c + \rho$. We write

$$\{c', \rho'\} \leftarrow \text{MAD}(x, y, c, \rho).$$

If such a routine is available then multiplication can be performed in n^2 MAD calls using Algorithm 1. The MIRACL big number library [5] provides such a functionality.

Algorithm 1 MAD-based multiplication algorithm

Input: $a, b \in \mathbb{N}$.

Output: $r \in \mathbb{N}$ such that $r = a \times b$.

1. for $i \leftarrow 0$ to $2n - 1$
 2. $r_i \leftarrow 0$
 3. for $i \leftarrow 0$ to $n - 1$
 4. $c \leftarrow 0$
 5. for $j \leftarrow 0$ to $n - 1$
 6. $\{c, r_{i+j}\} \leftarrow \text{MAD}(a_i, b_j, c, r_{i+j})$
 7. $r_{i+n} \leftarrow c$
 8. return r
-

This approach is unsatisfactory: it performs more computation than often needed. Assuming a constant-time MAD instruction, Algorithm 1 runs in time $O(n^2)$.

2.2 Karatsuba’s algorithm

Karatsuba [15] proposed an ingenious divide-and-conquer multiplication algorithm, where the operands a and b are split as follows:

$$r = a \times b = (2^L \bar{a} + \underline{a}) \times (2^L \bar{b} + \underline{b}),$$

where typically $L = nw/2$. Instead of computing a multiplication between long integers, Karatsuba performs multiplications between shorter integers, and (virtually costless) multiplication by powers of 2. Karatsuba’s algorithm is described in Algorithm 2.

Algorithm 2 Karatsuba multiplication algorithm

Input: $a, b \in \mathbb{Z}$.

Output: $r \in \mathbb{Z}$ such that $r = a \times b$.

1. $u \leftarrow \bar{a} \times \bar{b}$
 2. $v \leftarrow \underline{a} \times \underline{b}$
 3. $w \leftarrow (\bar{a} + \underline{a})(\bar{b} + \underline{b}) - u - v$
 4. $r \leftarrow 2^{2L} \times u + 2^L \times w + v$
 5. return r
-

¹An acronym standing for “Multiply Add Divide”.

Table 1 An example showing how linear relationships between individual words are encoded and interpreted. All the a_i are computed from a_3 and a_7 only

step	a_0	a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	a_9	meaning	reg1	reg2	reg3
0				1	=			1			$a_5 \leftarrow a_3 + a_7$	a_5	a_3	a_7
1					=	2					$a_4 \leftarrow a_5 + a_5$	a_5	a_4	a_7
2	=				-1			1			$a_0 \leftarrow a_7 - a_4$	a_0	a_4	a_7
3					2					=	$a_9 \leftarrow a_4 + a_4$	a_0	a_4	a_9
4	1	=								-1	$a_1 \leftarrow a_0 - a_9$	a_0	a_1	a_9
5		1							=	1	$a_8 \leftarrow a_1 + a_9$	a_8	a_1	a_9
6		1	=							1	$a_2 \leftarrow a_1 + a_8$	a_8	a_1	a_2
7			2				=				$a_6 \leftarrow a_2 + a_2$	a_8	a_6	a_2

This approach is much faster than naive multiplication – on which it still relies for multiplication between short integers – and runs² in $\Theta(n^{\log_2 3})$.

2.3 Bernstein’s multiplication algorithm

When one of the operands is constant, different ways to optimize multiplication exist. Bernstein [3] provides a branch-and-bound algorithm based on a cost function.

The minimal cost, and an associated sequence, are found by exploring a tree, possibly using memoization to avoid redundant searches. More elaborate pruning heuristics exist to further speedup searching. The minimal cost path produces a list of operations which provide the result of multiplication.

Because of its exponential complexity, Bernstein’s algorithm is quickly overwhelmed when dealing with large integers. It is however often implemented by compilers for short (32 to 64-bit) constants.

3 The proposed algorithm

3.1 Intuitive idea

The proposed algorithm works with an alternative representation of the constant operand a . Namely, we wish to express some a_i as a linear combination of other a_j s with small coefficients. It is then easy to reconstruct the whole multiplication $b \times a$ from the values of the $b \times a_j$ only.

The more linear combinations we can find, the less multiplications we need to perform. Our algorithm therefore tries to find the longest sequence of linear relationships between the digits of a . We call this sequence’s length the *coverage* of a .

Yet another performance parameter is the number of registers used by the multiplier. Ideally at any point in time two registers holding intermediate values should be used. This is not always possible and depends on the digits of a .

As an example, consider the set of relations of Table 1. All words are expressed ultimately in terms of the values of a_3 and a_7 . In Table 1, we express a as a subset of words

²When repeated recursively.

$A \in \{a_0, \dots, a_{n-1}\}$ and build a sparse table U where $U_{i,j} \in \{-1, 0, 1, 2, =\}$, which encodes linear relationships between individual words. During multiplication, U describes how the different a_i can be derived from each other.

Hence it suffices to compute $b \times a_3$ and $b \times a_7$ to infer all other $b \times a_i$ by long integer additions. Note that the algorithm only needs to allocate three $(n + 1)$ -word registers `reg1`, `reg2` and `reg3` to store intermediate results.

The values allowed in U can easily be extended to include more complex relationships (larger coefficients, more variables, etc.) but this immediately impacts the algorithm's performance. Indeed, the corresponding search graph has correspondingly many more branches at each node.

Operations can be performed without overflowing (i.e., so that results fit in a word), or modulo the word size. In the latter case, it is necessary to subtract $b \ll w$ from the result, where w is the word size, to obtain the correct result. This incurs some additional cost.

As a concrete example, if the operations in Table 1 are understood modulo 256 (and the a_i are bytes), it provides the decomposition of the following 80-bit number:

$$a = 0x8d69f62\underline{a}9249e\underline{c}1f8d24$$

where we have underlined the values of a_3 and a_7 .

3.2 Backtracking algorithm

Algorithm 3 Step(u, w) macro

1. $(p_{d+1,0}, p_{d+1,1}) \leftarrow (u, w)$
 2. Backtrack($d + 1$)
-

Algorithm 4 EncodeDep(c, opcode) macro

1. if $c < 256$
 2. if $v_c = \text{False}$
 3. $(v_c, p_{d,2}, p_{d,3}) \leftarrow (\text{True}, c, \text{opcode})$
 4. Step(a, b)
 5. Step(a, c)
 6. Step(b, c)
 7. $v_c \leftarrow \text{False}$
-

Algorithm 5 Backtrack(d) macro

1. if $d > d^{\max}$ then $(d^{\max}, p^{\max}) \leftarrow (d, p)$
 2. $(a, b) \leftarrow (p_{d,0}, p_{d,1})$
 3. for all opcode in \mathcal{C}
 4. EncodeDep(opcode(a, b), opcode)
-

Algorithm 6 Main backtracking program**Input:** $A = \sum_{i=0}^{N-1} 256^i A_i$.**Output:** $U_{i,j}$ to be processed by the Virtual Machine.

Initialisation

1. for $i \leftarrow 0$ to 255, $v_i \leftarrow \text{True}$
2. for $i \leftarrow 0$ to $N - 1$, $v_{A_i} \leftarrow \text{False}$
3. $d^{\max} \leftarrow -1$

Backtracking

4. for $i \leftarrow 0$ to 255
5. for $j \leftarrow i + 1$ to 255
6. if $v_i = v_j = \text{False}$
7. $(p_{0,0}, p_{0,1}, v_i, v_j) \leftarrow (i, j, \text{True}, \text{True})$
8. Backtrack(0)
9. $(v_i, v_j) \leftarrow (\text{False}, \text{False})$

 U -matrix reconstruction

10. $U_{i,j} \leftarrow 0$
11. for $i \leftarrow 0$ to 255
12. $(\text{in}_1, \text{in}_2, \text{out}, \text{opcode}) \leftarrow p_i^{\max}$
13. $(U_{i,\text{in}_1}, U_{i,\text{in}_2}, U_{i,\text{out}}) \leftarrow (1, 1, \text{opcode})$
14. return $U_{i,j}$

Linear combinations amongst words of a are found by backtracking [16], the pseudocode of which is given in Algorithm 5. Our implementation focuses on linear dependencies amongst 8-bit words, as our main recommendation for applying the proposed multiplication algorithm is exactly an 8-bit microprocessor.

We take advantage of recursion and macro expansion (see Algorithms 3 to 5) to achieve a more compact code (see Appendix). In this implementation, p encodes the current depth's three registers of Table 1 as well as the current operation. With suitable listing, Algorithm 6 outputs a set of values being related, along with the corresponding relation. The dependencies that we take into account in our C code (provided in the full version or upon request) don't go beyond depth 2. Thus, the corresponding operations are $\mathcal{C} = \{+, -, \times 2\}$. We also add these operations performed modulo 256, to obtain more solutions. The alternative to this approach is to consider a bigger depth, which naturally leads to more possibilities.

Our program takes as an input an integer p that represents the percentage of a being covered (i.e., the coverage is $p/100$ times the length of the a). In a typical lightweight scenario, a 128-byte number is involved in the multiplication process. Our software attempts to backtrack over a coverage-related number of values out of 256. It follows immediately that at most a 50% coverage would be required for performing such a multiplication (as byte collisions are likely to happen).

The program takes as parameter the list of bytes of a . If some bytes appear multiple times, it is not necessary to re-generate each of them individually: generation is performed once, and the value is cloned and dispatched where needed.

Note that if precomputation takes too long, the list of a_i can be partitioned into several sub-lists on which backtrackings are run independently. This would entail as many initial multiplications by the online multiplier but still yield appreciable speed-ups.

3.3 Multiplication algorithm

Algorithm 7 Multiplication virtual machine

Input: b , $\text{instr} = (\text{opcode}, i, j, t, p)_k, R$.

Output: r .

1. $r \leftarrow 0$
 2. for each $(i, v) \in R$
 3. $\text{reg}[i] \leftarrow v \times b$
 4. $\text{PlaceAt}(v, \text{reg}[i])$
 5. for each $(\text{opcode}, i, j, t, p) \in \text{instr}$
 6. $\text{reg}[t] \leftarrow \text{opcode}(\text{reg}[i], \text{reg}[j])$
 7. $\text{PlaceAt}(p, t)$
 8. return r
-

With the encoding of a generated by Algorithm 6, it is now possible to implement multiplication efficiently.

To that end we make use of a specific-purpose multiplication virtual machine (VM) described in Algorithm 7. The VM is provided with instructions of the form

$$\text{opcode } \tau, i, j, p$$

that are extracted offline from U . Here, opcode is the operation to perform, i and j are the indices of the operands, τ is the index of the result, and $p \leftarrow w \times \tau$ is the position in r where to place the result, w being the word size. The value of p is pre-computed offline to allow for a more efficient implementation.

We store the result in a $2n$ -byte register initialized with zero. We also make use of a long addition procedure $\text{PlaceAt}(p, i)$ which “places” the contents of the $(n + 1)$ -byte register $\text{reg}[i]$ at position p in r . PlaceAt performs the addition of register $\text{reg}[i]$ starting from an offset p in r , propagating the carry as needed.

Finally, we assume that the list $R = (i, v)_k$ of root nodes (position and value) of U is provided.

After executing all the operations described in U , Algorithm 7 has computed $r \leftarrow a \times b$.

Remark 1 (Karatsuba multiplication) Using the notations of Algorithm 2 one can see that in settings where a is a constant, the numbers u, v, w all result from the multiplication of \bar{b}, \underline{b} and $\bar{b} + \underline{b}$ (which are variable) by \bar{a}, \underline{a} and $\bar{a} + \underline{a}$ (which are constant). Hence our approach can independently be combined with Karatsuba’s algorithm to yield further improvements.

4 Performance

The algorithm has an offline step (backtracking) and an online step (multiplication), which are implemented on different devices.

The offline step is naturally the longest; its performance is heavily dependent on the digit combination operations allowed and on how many numbers are being dealt with. More precisely, results are near-instant when dealing with 64 individual bytes and operations $\{+, -, \times 2\}$. It takes much longer if operations modulo 256 are considered as well, but this

Table 2 Performance on a 68HC05 clocked at 5 MHz

	Time	RAM	Code size
Usual algorithm	188 ms	395 bytes	1.1 kB
New algorithm	72 ms	663 bytes	1.7 kB

gives a better coverage of a , hence better *online* results. That being said, modulo 256 operations are slightly less efficient than operations over the integers ($\simeq 1.5$ more costly), since they require a subtraction of b afterwards.

Table 2 provides comparative performance data for a multiplication by the processed constant $\lfloor \pi 2^{1024} \rfloor$. Backtracking this constant took 85 days on an Altix UV1000 cluster.

As a final remark, note that one can also reverse the idea and generate a key by which multiplication is easy. This can be done by progressively picking VM operations until an operand (key) with sufficient entropy is obtained. While this is not equivalent to randomly selecting keys, the authors conjecture that, in practice, the existence of linear relations³ between key bytes should not significantly weaken public-key implementations.

Appendix: Source code

```
#include <stdio.h>
#include <string.h>
#include <math.h>
#include <stdlib.h>

int False = 0;
int True = 255;

#define step(u,v) path[dep+1][0]=u; path[dep+1][1]=v; backtracking(dep+1);
#define steps(x,n) c=x; if(c<256) {if (visited[c]==False) {visited[c]=True;\
    path[dep][2]=c;path[dep][3]=n;step(a,b);step(a,c);step(b,c);visited[c]=False;}}

int visited[2 * 256];
int maxdep,path[256][4],maxpath[256][4];
int size;

void backtracking(int dep){
    if (dep > maxdep){
        maxdep = dep;
        memcpy(maxpath, path, sizeof(path));
    }

    int a = path[dep][0];
    int b = path[dep][1];
    int c;

    steps(a+b, 1);
    steps(a<<1, 2);
    steps(b<<1, 3);
    steps((a<<1)%256, 4);
    steps((b<<1)%256, 5);
    steps((a+b)%256, 6);
    steps(abs(a-b), 7);
}
```

³These linear relations are unknown to the attacker.


```

int main() {
    int i,j;
    FILE *Fin, *Fout;
    int *A;

    Fin = fopen("input.txt", "r");
    fscanf(Fin, "%d", &size);
    A = (int*)malloc(sizeof(int)*size);
    for (i=0; i < size; ++i) {
        fscanf(Fin, "%d", &A[i]);
    }

    for (i=0; i<256; ++i) visited[i] = True;
    for (i=0; i<size; ++i) visited[A[i]] = False;

    free(A);

    maxdep = -1;
    for (i=0; i<256; ++i) {
        for (j=i+1; j<256; ++j) {
            if (visited[i]==True || visited[j]==True) continue;
            path[0][0] = i;
            path[0][1] = j;
            visited[i] = visited[j] = True;
            backtracking(0);
            visited[i] = visited[j] = False;
        }
    }

    Fout = fopen("output.txt", "w");
    for (i=0; i < maxdep; ++i)
        fprintf(Fout, "%3d %3d %3d %3d\n",
            maxpath[i][0], // a_i
            maxpath[i][1], // a_j
            maxpath[i][2], // a_p
            maxpath[i][3]); // op

    return EXIT_SUCCESS;
}

```

References

1. Avižienis, A.: Signed-digit number representations for fast parallel arithmetic. *IRE Trans. Electron. Comput.* **EC-10**(3):389–400. <https://doi.org/10.1109/TEC.1961.5219227> (1961)
2. Barrett, P.: Implementing the Rivest Shamir and Adleman public key encryption algorithm on a standard digital signal processor. In: Odlyzko, A.M. (ed.) *Advances in Cryptology—CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, Santa Barbara, CA, USA, August 1987, pp. 311–323. Springer, Heidelberg (1987)
3. Bernstein, R.: Multiplication by integer constants. *Softw. Pract. Exp.* **16**(7), 641–652 (1986)
4. Cappello, P.R., Steiglitz, K.: Some complexity issues in digital signal processing. *IEEE Trans. Acoust. Speech Signal Process.* **32**(5), 1037–1041 (1984)
5. Certivox. The MIRACL big number library. See <https://www.certivox.com/miracl>
6. Cook, S.A.: On the minimum computation time of functions. PhD thesis (1966)
7. Dempster, A.G., Macleod, M.D.: Constant integer multiplication using minimum adders. *IEE Proc.—Circ. Dev. Syst.* **141**(5), 407–413 (1994)
8. Dempster, A.G., Macleod, M.D.: Use of Multiplier Blocks to Reduce Filter Complexity. In: 1994 IEEE International Symposium on Circuits and Systems, ISCAS, 1994, pp. 263–266. London, England (1994). <https://doi.org/10.1109/ISCAS.1994.409247>
9. Diffie, W., Hellman, M.E.: New directions in cryptography. *IEEE Trans. Inf. Theory* **22**(6), 644–654 (1976)
10. ElGamal, T.: On computing logarithms over finite fields. In: Williams, H.C. (ed.) *Advances in Cryptology—CRYPTO’85*, volume 218 of *Lecture Notes in Computer Science*, Santa Barbara, CA, USA, August 18–22, 1986, pp. 396–402. Springer, Heidelberg (1986)

11. Feige, U., Fiat, A., Shamir, A.: Zero knowledge proofs of identity. In: Aho, A. (ed.) 19th Annual ACM Symposium on Theory of Computing, pp. 210–217, New York City, NY, USA, May 25–27, 1987. ACM Press (1987)
12. Feige, U., Fiat, A., Shamir, A.: Zero-knowledge proofs of identity. *J. Cryptol.* **1**(2), 77–94 (1988)
13. Fürer, M.: Faster integer multiplication. *SIAM J. Comput.* **39**(3), 979–1005 (2009)
14. Harvey, D., Van Der Hoeven, J., Lecerf, G.: Even faster integer multiplication. arXiv preprint [arXiv:1407.3360](https://arxiv.org/abs/1407.3360) (2014)
15. Karatsuba, A., Ofman, Y.: Multiplication of many-digital numbers by automatic computers. *Doklady Akad. Nauk SSSR* **145**, 293–294 (1962)
16. Knuth, D.: *The Art of Computer Programming* (1968)
17. Montgomery, P.L.: Modular multiplication without trial division. *Math. Comput.* **44**(170), 519–521 (1985)
18. Schönhage, A., Strassen, V.: Schnelle Multiplikation grosser Zahlen. *Computing* **7**(3–4), 281–292 (1971)
19. Toom, A.L.: The complexity of a scheme of functional elements realizing the multiplication of integers. *Soviet Math. Dokl.* **3**, 714–716 (1963)
20. Wu, H., Hasan, M.A.: Closed-form expression for the average weight of signed-digit representations. *IEEE Trans. Comput.* **48**(8), 848–851 (1999)