

A parallel algorithm for dot product over word-size finite field using floating-point arithmetic

Stef Graillat

Joint work with Jérémy Jean

LIP6/PEQUAN - Université Pierre et Marie Curie (Paris 6) - CNRS

SYNASC 2010

12th International Symposium on Symbolic and Numeric Algorithms for
Scientific Computing

Timisoara, Romania, September 23-26, 2010



- Dot products: key tool in numerical linear algebra
- Fast algorithms in scientific computing
- Cryptology
- Error-correcting codes
- Computer algebra

Floating-point numbers

Normalized floating-point numbers $\mathbb{F} \subset \mathbb{R}$:

$$x = \pm \underbrace{x_0.x_1 \dots x_{M-1}}_{\text{mantissa}} \times b^e, \quad 0 \leq x_i \leq b-1, \quad x_0 \neq 0$$

b : basis, M : precision, e : exponent such that $e_{\min} \leq e \leq e_{\max}$

Approximation of \mathbb{R} by \mathbb{F} with rounding $\mathbf{fl} : \mathbb{R} \rightarrow \mathbb{F}$.

Let $x \in \mathbb{R}$ then

$$\mathbf{fl}(x) = x(1 + \delta), \quad |\delta| \leq \mathbf{u}$$

Unit rounding $\mathbf{u} = b^{1-M}$ for rounding toward zero

Standard model of floating-point arithmetic

Let $x, y \in \mathbb{F}$ and $\circ \in \{+, -, \cdot, /\}$.

The result $x \circ y$ is not in general a floating-point number

$$\mathbf{fl}(x \circ y) = (x \circ y)(1 + \delta), \quad |\delta| \leq \mathbf{u}$$

IEEE 754 standard (1985)

Type	Size	Mantissa	Exponent	Unit rounding	Interval
Single	32 bits	23+1 bits	8 bits	$\mathbf{u} = 2^{1-24} \approx 1,92 \times 10^{-7}$	$\approx 10^{\pm 38}$
Double	64 bits	52+1 bits	11 bits	$\mathbf{u} = 2^{1-53} \approx 2,22 \times 10^{-16}$	$\approx 10^{\pm 308}$

Finite field \mathbb{F}_p (p prime)

$\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z} = GF(p) = \{0, 1, \dots, p-1\}$ is a finite field with characteristic p

Finite field \mathbb{F}_p (p prime)

$\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z} = GF(p) = \{0, 1, \dots, p-1\}$ is a finite field with characteristic p

Operations in the field, for $a, b \in \mathbb{Z}/p\mathbb{Z}$:

- Addition: $a + b \in \{0, \dots, 2(p-1)\} \rightarrow a + b \pmod{p} \in \mathbb{Z}/p\mathbb{Z}$
- Multiplication: $ab \in \{0, \dots, (p-1)^2\} \rightarrow ab \pmod{p} \in \mathbb{Z}/p\mathbb{Z}$

Finite field \mathbb{F}_p (p prime)

$\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z} = GF(p) = \{0, 1, \dots, p-1\}$ is a finite field with characteristic p

Operations in the field, for $a, b \in \mathbb{Z}/p\mathbb{Z}$:

- Addition: $a + b \in \{0, \dots, 2(p-1)\} \rightarrow a + b \pmod{p} \in \mathbb{Z}/p\mathbb{Z}$
- Multiplication: $ab \in \{0, \dots, (p-1)^2\} \rightarrow ab \pmod{p} \in \mathbb{Z}/p\mathbb{Z}$

Reduction modulo p for $a \in \mathbb{Z}/p\mathbb{Z}$:

$$a \pmod{p} = a - \left\lfloor \frac{a}{p} \right\rfloor p = a - \lfloor a \cdot \text{inv}P \rfloor p$$

Aim

Let $p \geq 3$ a prime number and $(a_i)_i, (b_i)_i$ two vectors of N scalars in $\mathbb{Z}/p\mathbb{Z}$. We want to compute the dot product of a and b in $\mathbb{Z}/p\mathbb{Z}$:

$$a \cdot b = \sum_{i=1}^N a_i b_i \pmod{p}$$

Let $p \geq 3$ a prime number and $(a_i)_i, (b_i)_i$ two vectors of N scalars in $\mathbb{Z}/p\mathbb{Z}$. We want to compute the dot product of a and b in $\mathbb{Z}/p\mathbb{Z}$:

$$a \cdot b = \sum_{i=1}^N a_i b_i \pmod{p}$$

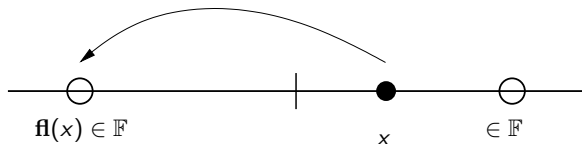
Assumptions:

- The integers are stored as floating-point numbers $\rightarrow \mathbb{F} \cap \mathbb{N}$
- The prime p satisfies $p - 1 < 2^{M-1}$
- The numbers are assumed to be **nonnegative**
- The rounding mode is **rounding toward zero**

Rounding toward zero in \mathbb{R}^+

Let $x \in \mathbb{R}^+$ $\mathbf{fl}(x)$ be the rounding toward zero of x in \mathbb{F}

- Equivalent to a truncation

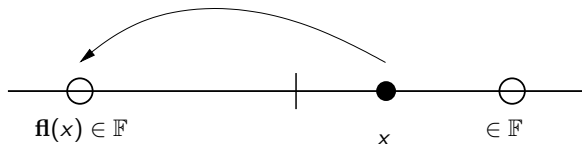


Rounding toward zero in \mathbb{R}^+

Let $x \in \mathbb{R}^+$ $\mathbf{fl}(x)$ be the rounding toward zero of x in \mathbb{F}

- Equivalent to a truncation
- The rounding is less or equal to the exact number:

$$\forall x \in \mathbb{R}^+, \mathbf{fl}(x) \leq x$$



Rounding toward zero in \mathbb{R}^+

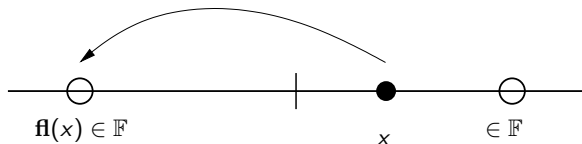
Let $x \in \mathbb{R}^+$ $\mathbf{fl}(x)$ be the rounding toward zero of x in \mathbb{F}

- Equivalent to a truncation
- The rounding is less or equal to the exact number:

$$\forall x \in \mathbb{R}^+, \mathbf{fl}(x) \leq x$$

- The rounding error is nonnegative:

$$\forall x \in \mathbb{R}^+, x - \mathbf{fl}(x) \geq 0$$



Error-free Transformations (EFT)

Problem : the result of a floating-point operation is generally not representable by a floating-point numbers.

Solution: *Error-free transformations*

- non-evaluated sum of two floating-point numbers
 - the floating-point result of the operation
 - the rounding error (which is representable in \mathbb{F} in our cases)
- For $a, b \in \mathbb{F} \cap \mathbb{N}$ and $\circ \in \{+, \times\}$,

$$a \circ b = \mathbf{fl}(a \circ b) + e, \text{ with } e \in \mathbb{F},$$

which is mathematically true.

Error-free Transformations for the product (1/2)

For $a, b, c \in \mathbb{F}$,

- $\text{FMA}(a, b, c)$ is the rounding of $a \cdot b + c$

Algorithm 1 (EFT for the product of two floating-point numbers)

```
function  $[x, y] = \text{TwoProductFMA}(a, b)$ 
```

```
   $x = \text{fl}(a \cdot b)$ 
```

```
   $y = \text{FMA}(a, b, -x)$ 
```

The FMA is now included in the IEEE 754-2008 standard

Error-free Transformations for the product (2/2)

Theorem 1

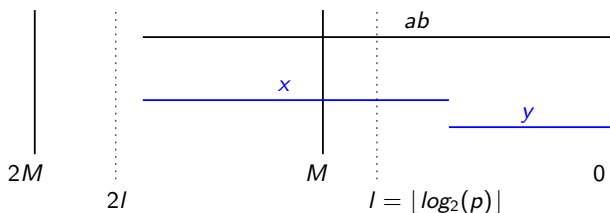
Let $a, b \in \mathbb{F} \cap \mathbb{N}$ and $x, y \in \mathbb{F}$ such that

$$[x, y] \leftarrow \text{TwoProductFMA}(a, b)$$

Then

$$ab = x + y, \quad x = \mathbf{fl}(ab), \quad 0 \leq y < \mathbf{u.ufp}(x), \quad 0 \leq x \leq ab$$

Algorithm TwoProductFMA requires 2 flops.



Binary euclidean division (1/2)

For $a, d \in \mathbb{F} \cap \mathbb{N}$, $d \neq 0$, the euclidean division of a by d is

$$a = qd + r, \quad 0 \leq r < d$$

For $a \in \mathbb{F} \cap \mathbb{N}$ and $\sigma = 2^k$, $\sigma \geq a$, one defines

Algorithm 2 (Split of a floating-point numbers)

```
function  $[x, y] = \text{ExtractScalar}(\sigma, a)$ 
```

```
     $q = \mathbf{fl}(\sigma + a)$ 
```

```
     $x = \mathbf{fl}(q - \sigma)$ 
```

```
     $y = \mathbf{fl}(x - a)$ 
```

\mathbf{fl} is rounding toward zero

Algorithm first proposed by Rump, Ogita and Oishi in rounding to the nearest

Binary euclidean division (2/2)

Theorem 2

Let $a \in \mathbb{F} \cap \mathbb{N}$, $\sigma = 2^k$, $\sigma \geq a$ and $x, y \in \mathbb{F}$ such that

$$[x, y] \leftarrow \text{ExtractScalar}(\sigma, a)$$

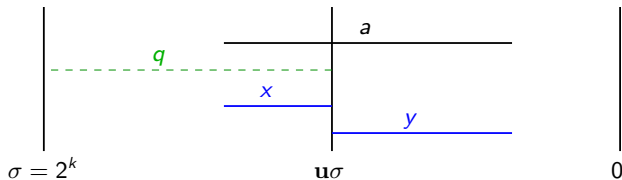
Then

$$a = x + y, \quad 0 \leq y < \mathbf{u}\sigma, \quad 0 \leq x \leq a, \quad x \in \mathbf{u}\sigma\mathbb{N}$$

Algorithm `ExtractScalar` requires 3 flops.

Remark:

$$a = x + y = x' \mathbf{u}\sigma + r, \quad x' \in \mathbb{N}, \quad 0 \leq r < \mathbf{u}\sigma$$



Computation of dot products

Assumption : $p - 1 < 2^{M-1}$ and $N < 2^{M/2}$

Computation of dot products

Assumption : $p - 1 < 2^{M-1}$ and $N < 2^{M/2}$

Idea :

- Split the number with a representation with only half the mantissa
- Sum them without error
- Reduction modulo p only at the end

Computation of dot products

Assumption : $p - 1 < 2^{M-1}$ and $N < 2^{M/2}$

Idea :

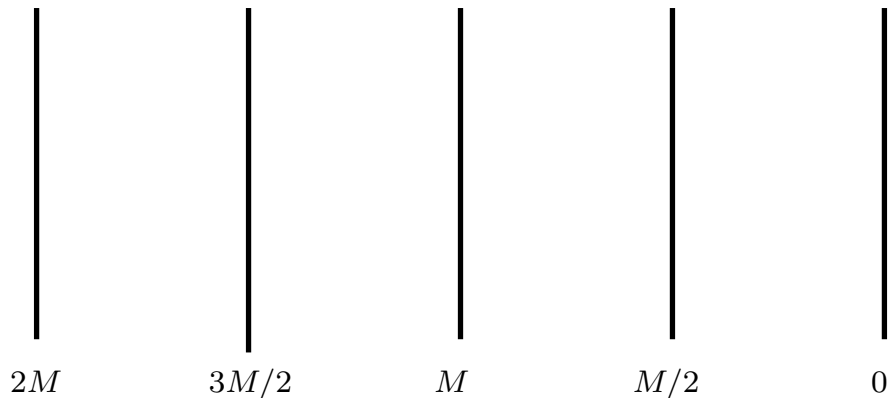
- Split the number with a representation with only half the mantissa
- Sum them without error
- Reduction modulo p only at the end

Use `ExtractScalar` to get: $s_1 = \left\lfloor \frac{M}{2} \right\rfloor$

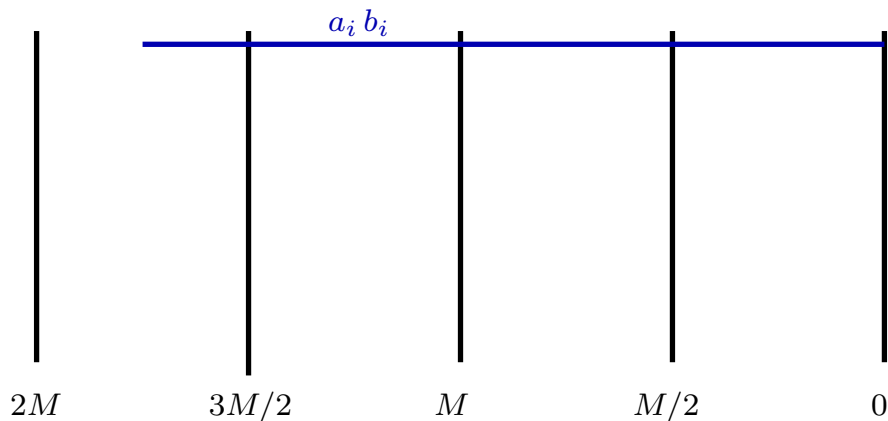
$$\forall i \in [1, N], \quad a_i b_i = \alpha_i + \beta_i + \gamma_i + \delta_i = A_i 2^{M+s_1} + B_i 2^M + C_i 2^{s_1} + D_i$$

$$a \cdot b = 2^{M+s_1} \sum_{i=1}^N A_i + 2^M \sum_{i=1}^N B_i + 2^{s_1} \sum_{i=1}^N C_i + \sum_{i=1}^N D_i \pmod{p}$$

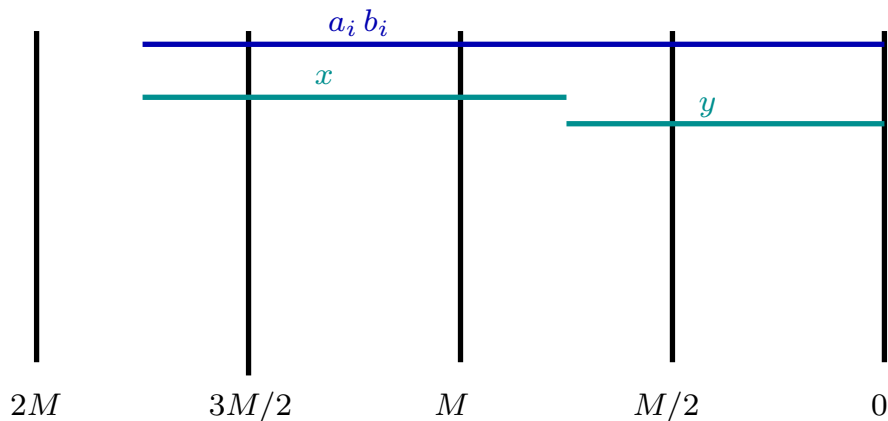
Principle of the splitting of $a_i b_i$ (1/2)



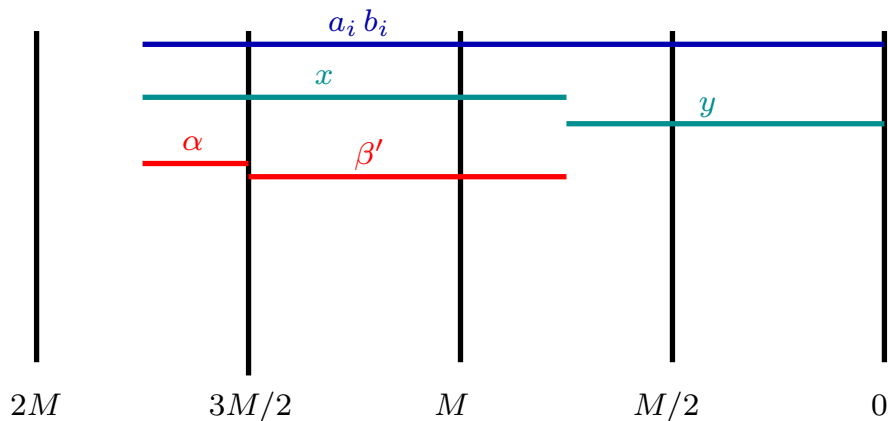
Principle of the splitting of $a_i b_i$ (1/2)



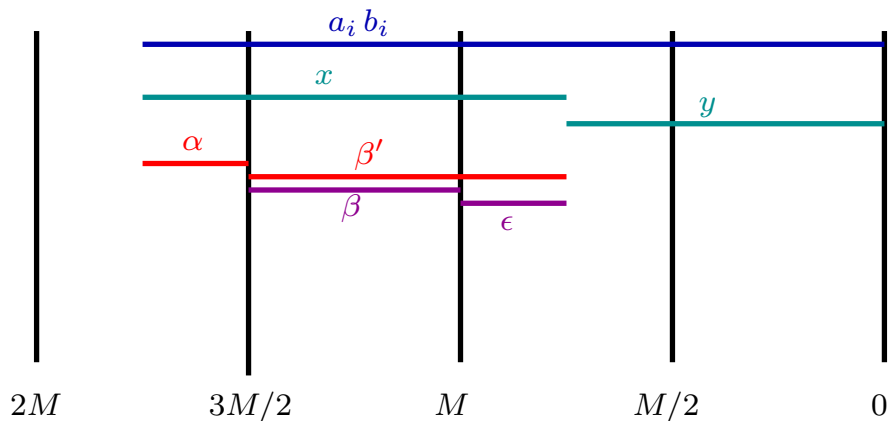
Principle of the splitting of $a_i b_i$ (1/2)



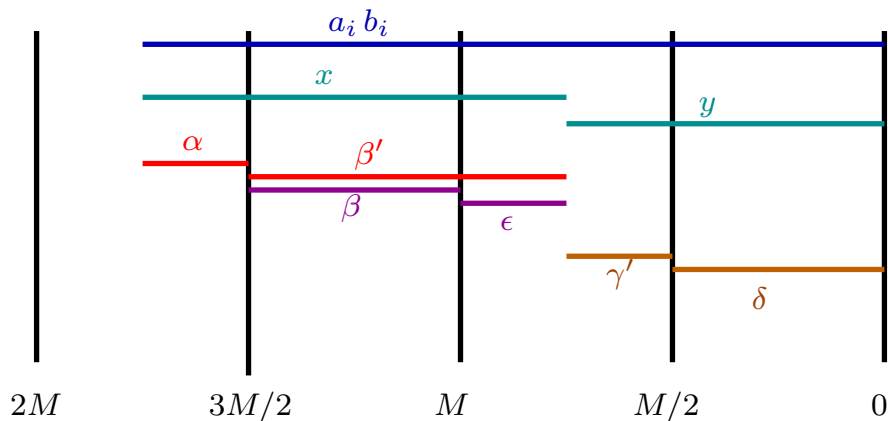
Principle of the splitting of $a_i b_i$ (1/2)



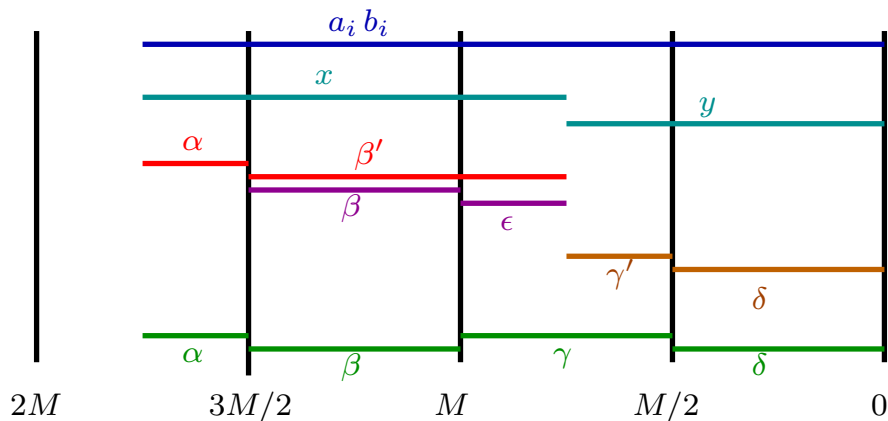
Principle of the splitting of $a_i b_i$ (1/2)



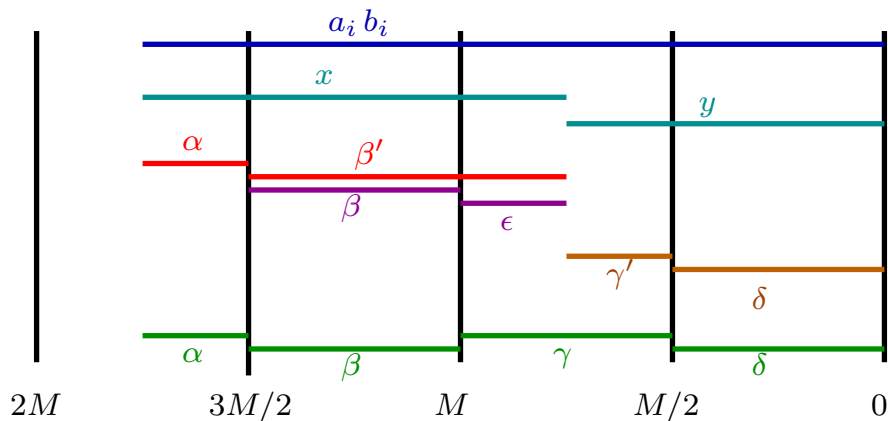
Principle of the splitting of $a_i b_i$ (1/2)



Principle of the splitting of $a_i b_i$ (1/2)



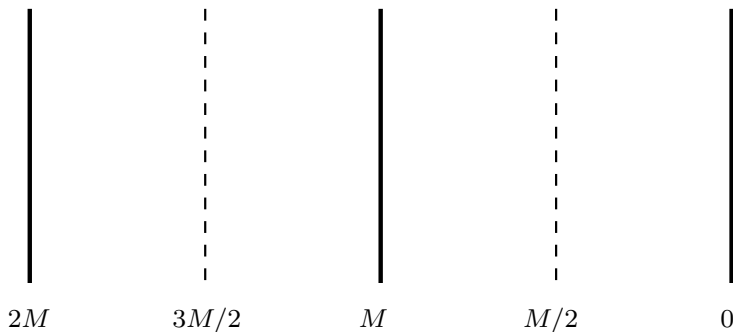
Principle of the splitting of $a_i b_i$ (1/2)



$$a_i b_i = \alpha + \beta + \gamma + \delta$$

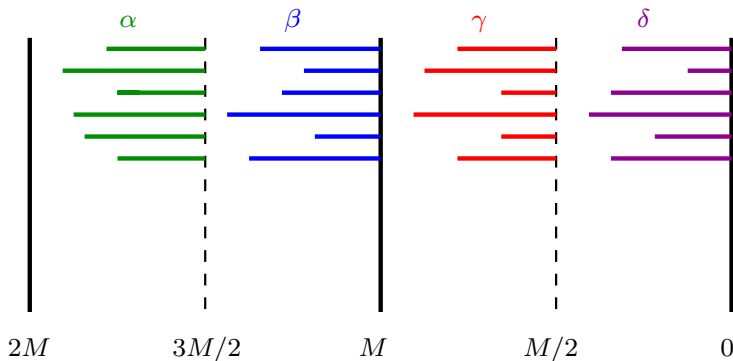
Principle of the splitting of $a_i b_i$ (2/2)

Split \longrightarrow 4 vectors of $N < 2^{M/2}$ elements with at most $M/2$ bits



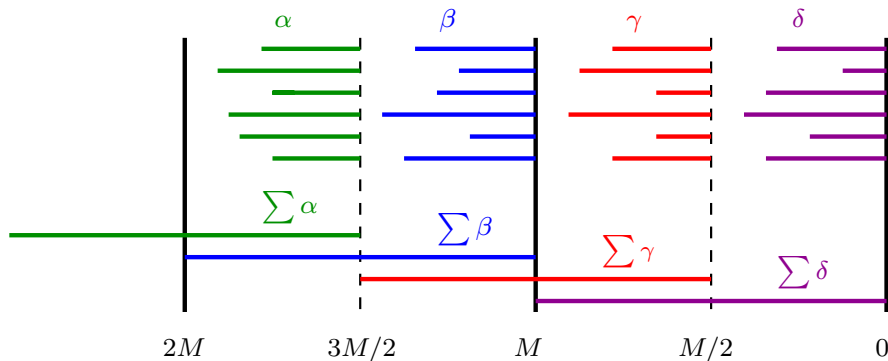
Principle of the splitting of $a_i b_i$ (2/2)

Split \rightarrow 4 vectors of $N < 2^{M/2}$ elements with at most $M/2$ bits



Principle of the splitting of $a_i b_i$ (2/2)

Split \rightarrow 4 vectors of $N < 2^{M/2}$ elements with at most $M/2$ bits



Final results:

$$a \cdot b = \sum_{i=1}^N \alpha_i + \sum_{i=1}^N \beta_i + \sum_{i=1}^N \gamma_i + \sum_{i=1}^N \delta_i \pmod{p}$$

Total cost: $16N + O(1)$ flops

Sequential algorithms

- Intel Itanium2 1.5GHz
- FMA instruction
- Double precision
- $p - 1 < 2^{53-1}$

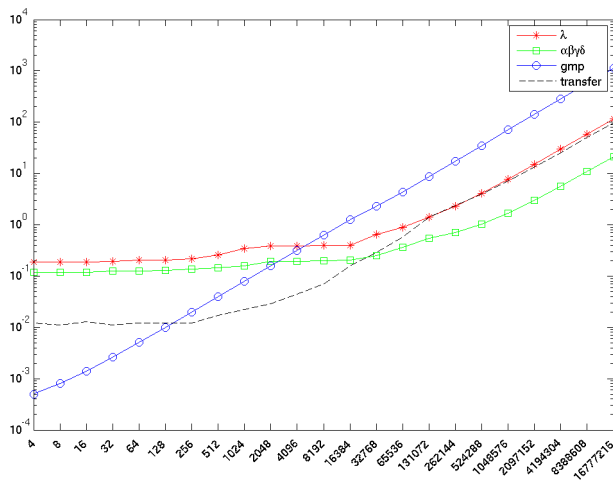
GPU algorithms

- Intel Core 2 Quad Processor Q8200 2.33GHz
- GPU: NVIDIA Tesla C1060
- FMA instruction
- Double precision
- $p - 1 < 2^{53-1}$

Comparison to a sequential GMP-based version.

GPU implementation

Timings for GPU implementations.



$$\rho = 2147483647 (\approx 2^{31})$$

Speedups:

- 10 for λ
- > 40 for $(\alpha, \beta, \gamma, \delta)$

Transfer time ignored.

Conclusion:

- An efficient algorithms using floating-point arithmetic well suited for parallelism
- Usage of error-free transformations when rounding toward zero

Future work:

- Port RNS algorithms to GPU
- Tests on new NVIDIA Fermi cards

Thank you for your attention