

# Cryptographic Engineering

## Cryptography in Software – Optimization Basics



---

Ruben Niederhagen

(based on content by Peter Schwabe)

Department of Mathematics and Computer Science (IMADA)

- Servers, main frames, ...
- PC, notebooks, ...
- Tablets, smartphones, ...
- Embedded devices, ...
- Smartcards, passports, ...



## **Embedded microcontrollers:**

- This is what you're looking at in the software assignment.
- Typically very tight size constraints (ROM and RAM).
- Different optimization targets: size, speed, . . .
- No (or very little) parallel computation capabilities.

## **Servers, workstations, laptops, smartphone:**

- No serious size constraints for crypto.
- Optimization target: speed (high throughput or *low latency*).
- Various different levels of parallelism.

- Some software makes extensive use of batching.
- Faster for many computations, if those are performed “together”.
- Example: McBits software (Bernstein, Chou, Schwabe, 2013):
  - 15486208 cycles on Intel Ivy Bridge for 256 decryptions
  - **NOT:**  $15486208/256 = 60493$  cycles for one decryption.
  - Software needs to wait until enough inputs are available.
  - Delay from input to output is delay of 256 decryptions.
- Highly parallel architectures (e.g., GPUs) focus on throughput.
- This can be a problem for, e.g., low-latency network communication.

- Tools like `time` or `time.h` have too low resolution.
- For serious optimization need to count CPU cycles.
- Use CPU's built-in cycle counter, e.g.,

on AMD64:

```
static long long cpucycles(void)
{
    unsigned long long result;
    asm volatile("rdtsc;"
                "shlq $32,%%rdx;"
                "orq %%rdx,%%rax"
                : "=a" (RES) :
                : "%rdx");
    return result;
}
```

on STM32F407:

```
#include <libopencm3/cm3/scs.h>
#include <libopencm3/cm3/dwt.h>

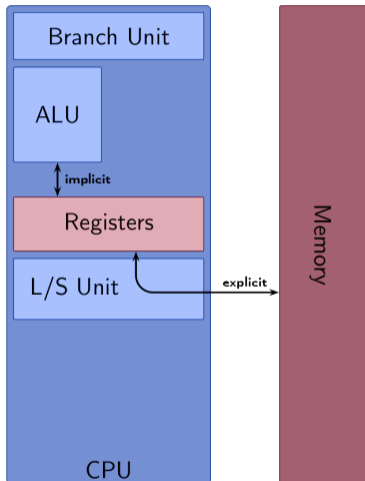
// active cycle counter
SCS_DEMCR |= SCS_DEMCR_TRCENA;
DWT_CYCCNT = 0;
DWT_CTRL |= DWT_CTRL_CYCCNTENA;

// read cycle counter
uint32_t clk_cnt = DWT_CYCCNT;
```

1. Your program is not running exclusively on the CPU, there may be interrupts.  
**Solution:** Measure many times, take the median (not average!).  
**Remark:** Also report quartiles.
2. The `rdtsc` instruction reports *reference* cycles, your CPU may run at a different speed.  
**Solution:** Switch off frequency scaling and TurboBoost/TurboCore.
3. Hyperthreading may run another process on the same physical core as your program.  
**Solution:** Switch off hyperthreading.
4. Getting reproducible, publicly verifiable benchmarks is hard.  
**Solution:** Use public benchmarking framework SUPERCOP by Bernstein and Lange:  
<http://bench.cr.yp.to>  
**Remark:** Submit cryptographic software to eBACS!

# Computers and Computer Programs

A highly simplified view



- A program is a sequence of *instructions*.
- Load/Store instructions move data between memory and registers (processed by the L/S unit).
- Branch instructions (conditionally) jump to a position in the program.
- Arithmetic instructions perform simple operations on values in registers (processed by the ALU).
- Registers are fast (fixed-size) storage units, addressed “by name”.

# A First Program

## Adding up 1000 integers

1. Set register R1 to zero.
2. Set register R2 to zero.
3. Load 32-bits from address  $\text{START} + \text{R2}$  into register R3.
4. Add 32-bit integers in R1 and R3, write the result in R1.
5. Increase value in register R2 by 4.
6. Compare value in register R2 to 4000.
7. Goto line 3 if R2 was smaller than 4000.



# A First Program

Adding up 1000 integers in readable syntax

```
int32 result
int32 tmp
int32 ctr

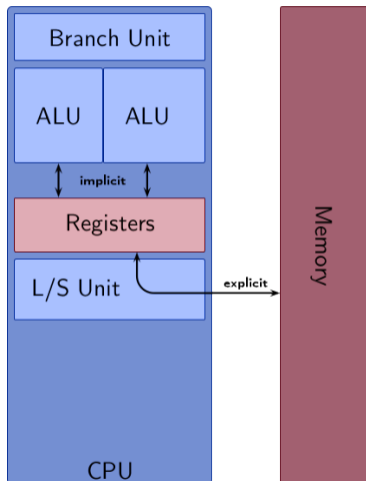
result = 0
ctr = 0
looptop :
tmp = mem32 [ START + ctr ]
result += tmp
ctr += 4
unsigned <? ctr - 4000
goto looptop if unsigned <
```

- Easy approach: Per cycle execute one instruction, then go for the next.
- Cycles needs to be long enough to finish the most complex supported instruction.
- Other approach: Chop instructions into smaller tasks, e.g. for addition:
  1. Fetch instruction
  2. Decode instruction
  3. Fetch register arguments
  4. Execute (actual addition)
  5. Write back to register
- Overlap instructions (while one instruction is in step 2, the next one can do step 1).
- This is called pipelined execution (many more stages possible).
- Advantage: Cycles can be much shorter (higher clock speed).
- Requirement for overlapping execution: Instructions must be independent.

- While the ALU is executing an instruction the L/S and branch units are idle.
- Idea: Duplicate fetch and decode, handle two or three instructions per cycle.
- While we're at it: Why not deploy two ALUs.
- This concept is called *superscalar* execution.
- Number of independent instructions of one type per cycle:  
**throughput.**
- Number of cycles that need to pass before the result can be used:  
**latency.**

# An Example Computer

Still highly simplified



## Latencies and throughputs:

- At most 4 instructions per cycle.
- At most 1 Load/Store instruction per cycle.
- At most 2 arithmetic instructions per cycle.
- Arithmetic latency: 2 cycles.
- Load latency: 3 cycles.
- Branches have to be last instruction in a cycle.

- Need at least 1000 load instructions:  $\geq 1000$  cycles.
- Need at least 999 addition instructions:  $\geq 500$  cycles.
- At least 1999 instructions:  $\geq 500$  cycles.
- **Lower bound:** 1000 cycles.

## Latencies and throughputs:

- At most 4 instructions per cycle.
- At most 1 Load/Store instruction per cycle.
- At most 2 arithmetic instructions per cycle.
- Arithmetic latency: 2 cycles.
- Load latency: 3 cycles.
- Branches have to be last instruction in a cycle.

# How about our program?

```
int32 result
int32 tmp
int32 ctr
```

```
result = 0
```

```
ctr = 0
```

```
looptop :
```

```
tmp = mem32[START + ctr]
```

```
# wait 2 cycles for tmp
```

```
result += tmp
```

```
ctr += 4
```

```
# wait 1 cycle for ctr
```

```
unsigned <? ctr - 4000
```

```
# wait 1 cycle for unsigned <
```

```
goto looptop if unsigned <
```

- Addition has to wait for load.
- Comparison has to wait for addition.
- Each iteration of the loop takes 8 cycles.
- Total: > 8000 cycles.
- **This program is very inefficient!**

# Making the Program Fast

## Step 1 — Unrolling

```
result = 0
tmp = mem32[START + 0]
# wait 2 cycles for tmp
result += tmp
tmp = mem32[START + 4]
# wait 2 cycles for tmp
result += tmp
tmp = mem32[START + 8]
# wait 2 cycles for tmp
result += tmp
...
tmp = mem32[START + 3996]
# wait 2 cycles for tmp
result += tmp
```

- Remove all the loop control: *unrolling*
- Each load-and-add now takes 3 cycles.
- Total:  $\approx 3000$  cycles
- Better, but still too slow.

# Making the Program Fast

## Step 2 — Instruction Scheduling

```
result = mem32[START + 0]
tmp0 = mem32[START + 4]
tmp1 = mem32[START + 8]
tmp2 = mem32[START + 12]

result += tmp0
tmp0 = mem32 [START + 16]
# wait 1 cycle for result
result += tmp1
tmp1 = mem32 [START + 20]
# wait 1 cycle for result
...
result += tmp2
tmp2 = mem32[START + 3996]
# wait 1 cycle for result
result += tmp0
# wait 1 cycle for result
result += tmp1
# wait 1 cycle for result
result += tmp2
```

- Load values earlier.
- Load latencies are hidden.
- Use more registers for loaded values (tmp0, tmp1, tmp2).
- Get rid of one addition to zero.
- Now arithmetic latencies kick in.
- Total:  $\approx 2000$  cycles



# Making the Program Fast

## Step 3 — More Instruction Scheduling (two accumulators)

```
result0 = mem32[START + 0]
tmp0 = mem32[START + 8]
result1 = mem32[START + 4]
tmp1 = mem32[START + 12]
tmp2 = mem32[START + 16]

result0 += tmp0
tmp0 = mem32[START + 20]
result1 += tmp1
tmp1 = mem32[START + 24]
result0 += tmp2
tmp2 = mem32[START + 28]

...

result0 += tmp1
tmp1 = mem32[START + 3996]
result1 += tmp2
result0 += tmp0
result1 += tmp1
result0 += result1
```

- Use one more accumulator register (`result1`).
- All latencies hidden.
- Total: 1004 cycles
- Asymptotically  $n$  cycles for  $n$  additions.
- **This is program is efficient.**

## Summary of what we did.

- Analyze the algorithm in terms of machine instructions.
- Look at what the respective machine is able to do.
- Compute a lower bound of the cycles.
- Optimize until we (almost) reached the lower bound:
  - Unroll the loop.
  - Interleave independent instructions (instruction scheduling).
  - Resulting program is larger and requires more registers!
- Note: Good instruction scheduling typically requires more registers.
- Opposing requirements to register allocation (assigning registers to live variables, minimizing memory access).
- Both instruction scheduling and register allocation are NP hard.
- So is the joint problem.
- Many instances are efficiently solvable.

## What instructions and how many registers do we have?

- Instructions are defined by the **instruction set**.
- Supported register names are defined by the **set of architectural registers**.
- Instruction set and set of architectural registers together define the **architecture**.
- Examples for architectures: x86, AMD64, ARMv6, ARMv7, UltraSPARC.
- Sometimes base architectures are extended, e.g., MMX, SSE, NEON.

## What determines latencies etc?

- Different **microarchitectures** implement an architecture.
- Latencies and throughputs are specific to a microarchitecture.
- Example: Intel Core 2 Quad Q9550 implements the AMD64 architecture.

- Optimal instruction scheduling depends on the microarchitecture.
- Code optimized for one microarchitecture may run at very bad performance on another microarchitecture.
- Many software is shipped in binary form (cannot recompile).
- Idea: Let the processor reschedule instructions on the fly.
- Look ahead a few instructions, pick one that can be executed.
- This is called **out-of-order execution**.
- Typically requires more physical than architectural registers and **register renaming**.
- Harder for the (assembly) programmer to understand what exactly will happen with the code at runtime.
- Harder to come up with optimal scheduling.
- Harder to screw up completely.

- So far there was nothing crypto-specific in this lecture.
- Is optimizing crypto the same as optimizing any other software?
- No. Cryptographic software deals with secret data (e.g., keys).
- Information about secret data must not leak through side channels.
- Most critical for software implementations on “large” CPUs:  
Software must take constant time (independent of secret data).

- Consider the following piece of code:

**if  $s$  then**

$r \leftarrow A$

**else**

$r \leftarrow B$

**end if**

- General structure of any conditional branch.
- $A$  and  $B$  can be large computations,  $r$  can be a large state.
- This code takes different amount of time, depending on  $s$ .
- Obvious timing leak if  $s$  is secret.
- Even if  $A$  and  $B$  take the same amount of cycles this is generally not constant time!
- Reasons: Branch prediction, instruction-caches.
- **Never use secret-data-dependent branch conditions!**

- So, what do we do with this piece of code?

**if  $s$  then**

$r \leftarrow A$

**else**

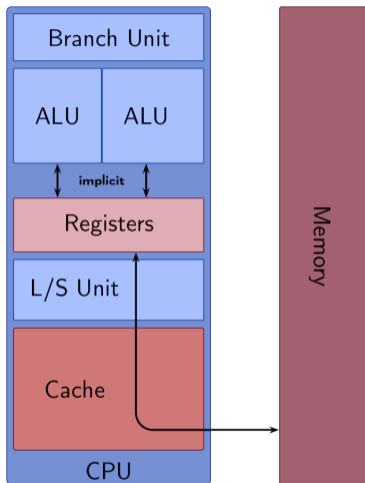
$r \leftarrow B$

**end if**

- Replace by

$r \leftarrow sA + (1 - s)B$

- Can expand  $s$  to all-one/all-zero mask and use XOR instead of addition, AND instead of multiplication.
- For very fast  $A$  and  $B$  this can even be faster than branching.



- Memory access goes through a **cache**.
- Small but fast transparent memory for re-used data.
- A load from memory places data also in the cache.
- Data remains in cache until replaced by other data.
- Loading data is fast if data is in the cache (**cache hit**).
- Loading data is slow if data is not in the cache (**cache miss**).



$T[0] \dots T[15]$
$T[16] \dots T[31]$
???
???
$T[64] \dots T[79]$
$T[80] \dots T[95]$
???
$T[112] \dots T[127]$
???
???
$T[160] \dots T[175]$
$T[176] \dots T[191]$
$T[192] \dots T[207]$
$T[208] \dots T[223]$
???
???



- Consider lookup table of 32-bit integers.
- Cache lines have 64 bytes.
- Crypto and the attacker's program run on the same CPU and thus share the cache.
- Tables are in cache.
- The attacker's program replaces some cache lines.
- Crypto continues, loads from table again.
- Attacker loads his data:
  - Fast: cache hit (crypto did not just load from this line).
  - Slow: cache miss (crypto just loaded from this line).

- This is only the most basic cache-timing attack.
- Non-secret cache lines are not enough for security.
- Load/Store addresses influence timing in many different ways.
- **Do not access memory at secret-data-dependent addresses.**
- Timing attacks are practical:  
Osvik, Tromer, Shamir, 2006: 65 ms to steal a 256-bit AES key used for Linux hard-disk encryption.
- *Remote* timing attacks are practical:  
Brumley, Tuveri, 2011: A few minutes to steal ECDSA signing key from OpenSSL implementation.

- Want to load item at (secret) position  $p$  from table of size  $n$ .
- Load all items, use arithmetic to pick the right one:

```
for  $i$  from 0 to  $n - 1$  do
```

```
   $d \leftarrow T[i]$ 
```

```
  if  $p = i$  then
```

```
     $r \leftarrow d$ 
```

```
  end if
```

```
end for
```

- Problem 1: if-statements are not constant time (see before).
- Problem 2: Comparisons are not constant time, replace by, e.g.:

```
static unsigned long long eq(uint32_t a, uint32_t b) {  
  unsigned long long t = a ^ b;  
  return 1 - ((-t) >> 63);  
}
```

## Lesson so far:

- Avoid all data flow from secrets to branch conditions and memory addresses.
- This can *always* be done; cost highly depends on the algorithm.
- Test this with `valgrind` and *uninitialized secret data* (see <https://www.post-apocalyptic-crypto.org/timecop/>).

*“In order for a function to be constant time, the branches taken and memory addresses accessed must be independent of any secret inputs. (That’s assuming that the fundamental processor instructions are constant time, but that’s true for all sane CPUs.)”*

— Langley, Apr. 2010

*“So the argument to the DIV instruction was smaller and DIV, on Intel, takes a variable amount of time depending on its arguments!”*

— Langley, Feb. 2013

- DIV, IDIV, FDIV on pretty much all Intel/AMD CPUs.
- Various math instructions on Intel/AMD CPUs (FSIN, FCOS, ...).
- MUL, MULHW, MULHWU on many PowerPC CPUs.
- UMULL, SMULL, UMLAL, and SMLAL on ARM Cortex-M3.

## Solution

- Avoid these instructions.
- Make sure that inputs to the instructions do not leak timing information.
- Requires assembler implementation or special compilers!

- Until early years 2000 each new processor generation had higher clock speeds.
- Nowadays, increase performance by number of cores:
  - A laptop has 4 physical (and 8 virtual) cores or more.
  - Smartphones typically have 4 or 6 cores or more.
  - Servers have 8, 16, ... cores.
  - Special-purpose hardware (e.g., GPUs) often comes with many more cores.
- Consequence: *“The free lunch is over”* (Herb Sutter, 2005)

*“As a result, system designers and software engineers can no longer rely on increasing clock speed to hide software bloat. Instead, they must somehow learn to make effective use of increasing parallelism.”*

— Maurice Herlihy: The Multicore Revolution, 2007

# Why multicore doesn't matter...

... for algorithm design in crypto.

## **Crypto is fast (single core of an old Intel Core i3-2310M):**

- > 50 RSA-4096 signatures per second
- > 8000 RSA-4096 signature verifications per second
- > 28000 Ed25519 signatures per second
- > 9000 Ed25519 signature verifications per second
  
- **If you perform only one crypto operation, you don't care.**
- **Many crypto operations are trivially parallel on multiple cores.**

## Scalar computation:

- Load 32-bit integer  $a$ .
- Load 32-bit integer  $b$ .
- Perform addition  
 $c \leftarrow a + b$ .
- Store 32-bit integer  $c$ .

## Vectorized computation:

- Load 4 consecutive 32-bit integers  $(a_0, a_1, a_2, a_3)$ .
  - Load 4 consecutive 32-bit integers  $(b_0, b_1, b_2, b_3)$ .
  - Perform addition  
 $(c_0, c_1, c_2, c_3) \leftarrow (a_0 + b_0, a_1 + b_1, a_2 + b_2, a_3 + b_3)$ .
  - Store 128-bit vector  $(c_0, c_1, c_2, c_3)$ .
- 
- Perform the same operations on independent data streams (SIMD).
  - Vector instructions available on most “large” processors.
  - Instructions for vectors of bytes, integers, floats, ...
  - Need to interleave data items (e.g., 32-bit integers) in memory.
  - Only limited help by compilers with vectorization.



- Imagine that:
  - vector addition is as fast as scalar addition and
  - vector loads are as fast as scalar loads.
- Need only 250 vector additions, 250 vector loads (plus adding up 4 partial sums).
- Lower bound of 250 cycles.
- Very straight-forward modification of the program.
- Fully unrolled loop needs only 1/4 of the space.

# Is it really that efficient?

- Consider the Intel Skylake processor with AVX2:
  - 32-bit load throughput: 2 per cycle
  - 32-bit add throughput: 4 per cycle
  - 32-bit store throughput: 1 per cycle
  
  - 256-bit load throughput: 2 per cycle
  - $8 \times$  32-bit add throughput: 3 per cycle
  - 256-bit store throughput: 1 per cycle
- **AVX2 vector instructions are almost as fast as scalar instructions but do  $8 \times$  the work.**
- Situation on other architectures/microarchitectures is similar.
- Reason: Cheap way to increase arithmetic throughput (less decoding, address computation, etc.).

- Data-dependent branches are expensive in SIMD.
  - Variably indexed loads (lookups) into vectors are expensive.
  - Need to rewrite algorithms to eliminate branches and lookups.
- Secret-data-dependent branches and secret branch conditions are the major sources of timing-attack vulnerabilities.
- ⇒ Synergies between speeding up code with vector instructions and protecting code!

## Carry handling:

- When adding two 32-bit integers, the result may have 33 bits (32-bit result + carry).
- Scalar additions keep the carry in a special *flag register*.
- Subsequent instructions can use this flag, e.g., “add with carry”.
- How about carries of vector additions?
  - Answer 1: Special “carry generate” instruction (e.g., CBE-SPU).
  - Answer 2: They’re lost, recomputation is expensive.
- Need to *avoid carries* instead of handling them.
- No problem for today’s lecture, but requires care for big-integer arithmetic.

## Removing instruction-level parallelism:

- If we do not vectorize, we perform multiple independent instructions.
- We turn data-level parallelism (DLP) into instruction-level parallelism (ILP).
- Pipelined and multiscalar execution need ILP.
- Vectorization removes ILP.
- Problematic for algorithms with, e.g., 4-way DLP.
- Good example to see this: ChaCha vs. Blake.
- Vectorization of ChaCha can resort to higher-level parallelism (multiple blocks).
- Harder for Blake: Each block depends on the previous one.

## Data shuffling:

Consider multiplication of 4-coefficient polynomials  $f = f_0 + f_1x + f_2x^2 + f_3x^3$  and  $g = g_0 + g_1x + g_2x^2 + g_3x^3$ :

$$r_0 = f_0g_0$$

$$r_1 = f_0g_1 + f_1g_0$$

$$r_2 = f_0g_2 + f_1g_1 + f_2g_0$$

$$r_3 = f_0g_3 + f_1g_2 + f_2g_1 + f_3g_0$$

$$r_4 = f_1g_3 + f_2g_2 + f_3g_1$$

$$r_5 = f_2g_3 + f_3g_2$$

$$r_6 = f_3g_3$$

- Ignore carries, overflows etc. for a moment.
- 16 multiplications, 9 additions.
- How to vectorize multiplications?
- Can easily load  $(f_0, f_1, f_2, f_3)$  and  $(g_0, g_1, g_2, g_3)$ .
- Multiply, obtain  $(f_0g_0, f_1g_1, f_2g_2, f_3g_3)$ .
- And now what?
- Answer:  
Need to shuffle data in input and output registers.
- Significant overhead, not clear that vectorization speeds up computation!

- Most important question: Where does the parallelism come from?
- Easiest answer: Consider multiple batched encryptions, decryptions, signature computations, verifications, etc. (but that increases latency).
- Often: Can exploit lower-level parallelism.
- Rule of thumb: Parallelize on an as high as possible level.
- Vectorization is hard to do as “add-on” optimization.
- Reconsider algorithms and data structures, synergy with constant-time algorithms.

- Imagine registers that have only one bit.
- Perform arithmetic on those registers using XOR, AND, OR.
- Essentially the same as hardware implementations.
- But wait, registers are longer!
- Think of them as vectors of bits.
- This needs transposition of the “binary data matrix”.
- Perform the simulated hardware implementations on many independent data streams.
- Bitslicing works for every algorithm.
- Bitslicing is inherently protected against timing attacks.
- Efficient bitslicing needs a huge amount of data-level parallelism.



## 4-coefficient binary polynomial:

$(a_3x^3 + a_2x^2 + a_1x + a_0)$ , with  $a_i \in 0, 1$ .

## 4-coefficient bitsliced binary polynomials:

$$\begin{array}{l} p_0 = a_{0,3} \ a_{0,2} \ a_{0,1} \ a_{0,0} \\ p_1 = a_{1,3} \ a_{1,2} \ a_{1,1} \ a_{1,0} \\ p_2 = a_{2,3} \ a_{2,2} \ a_{2,1} \ a_{2,0} \\ p_3 = a_{3,3} \ a_{3,2} \ a_{3,1} \ a_{3,0} \\ p_4 = a_{4,3} \ a_{4,2} \ a_{4,1} \ a_{4,0} \\ \quad \dots \end{array} \quad \longrightarrow \quad \begin{array}{l} r_0 = \dots \ a_{4,0} \ a_{3,0} \ a_{2,0} \ a_{1,0} \ a_{0,0} \\ r_1 = \dots \ a_{4,1} \ a_{3,1} \ a_{2,1} \ a_{1,1} \ a_{0,1} \\ r_2 = \dots \ a_{4,2} \ a_{3,2} \ a_{2,2} \ a_{1,2} \ a_{0,2} \\ r_3 = \dots \ a_{4,3} \ a_{3,3} \ a_{2,3} \ a_{1,3} \ a_{0,3} \end{array}$$

## 4-coefficient binary polynomial:

$(a_3x^3 + a_2x^2 + a_1x + a_0)$ , with  $a_i \in 0, 1$ .

## 4-coefficient bitsliced binary polynomials:

```
typedef unsigned char poly4; /* 4 coefficients in the low 4 bits */
```

```
typedef unsigned long long poly4x64[4];
```

```
void poly4_bitslice(poly4x64 r, const poly4 p[64]) {  
    for (int i = 0; i < 4; i++) {  
        r[i] = 0;  
        for (int j = 0; j < 64; j++)  
            r[i] |= (unsigned long long)(1 & (p[j] >> i)) << j;  
    }  
}
```

```
typedef unsigned long long poly4x64[4];
typedef unsigned long long poly7x64[7];

void poly4x64_mul(poly7x64 r, const poly4x64 a, const poly4x64 b)
{
    r[0] = a[0] & b[0];
    r[1] = (a[0] & b[1]) ^ (a[1] & b[0]);
    r[2] = (a[0] & b[2]) ^ (a[1] & b[1]) ^ (a[2] & b[0]);
    r[3] = (a[0] & b[3]) ^ (a[1] & b[2]) ^ (a[2] & b[1]) ^ (a[3] & b[0]);
    r[4] = (a[1] & b[3]) ^ (a[2] & b[2]) ^ (a[3] & b[1]);
    r[5] = (a[2] & b[3]) ^ (a[3] & b[2]);
    r[6] = (a[3] & b[3]);
}
```

- XOR, AND, OR, etc. are usually fast (e.g., three 128-bit operations per cycle on Intel Core 2).
- Can be very fast for operations that are not natively supported (like arithmetic in binary fields).
- Active data set increases massively (e.g.,  $128\times$ ).
- For “normal” vector operations, register space is increased accordingly (e.g, 16 256-bit vector registers vs. 16 64-bit integer registers).
- For bitslicing: Need to fit more data into the same registers.
- Typical consequence: More loads and stores for register spills (that easily become the performance bottleneck).