
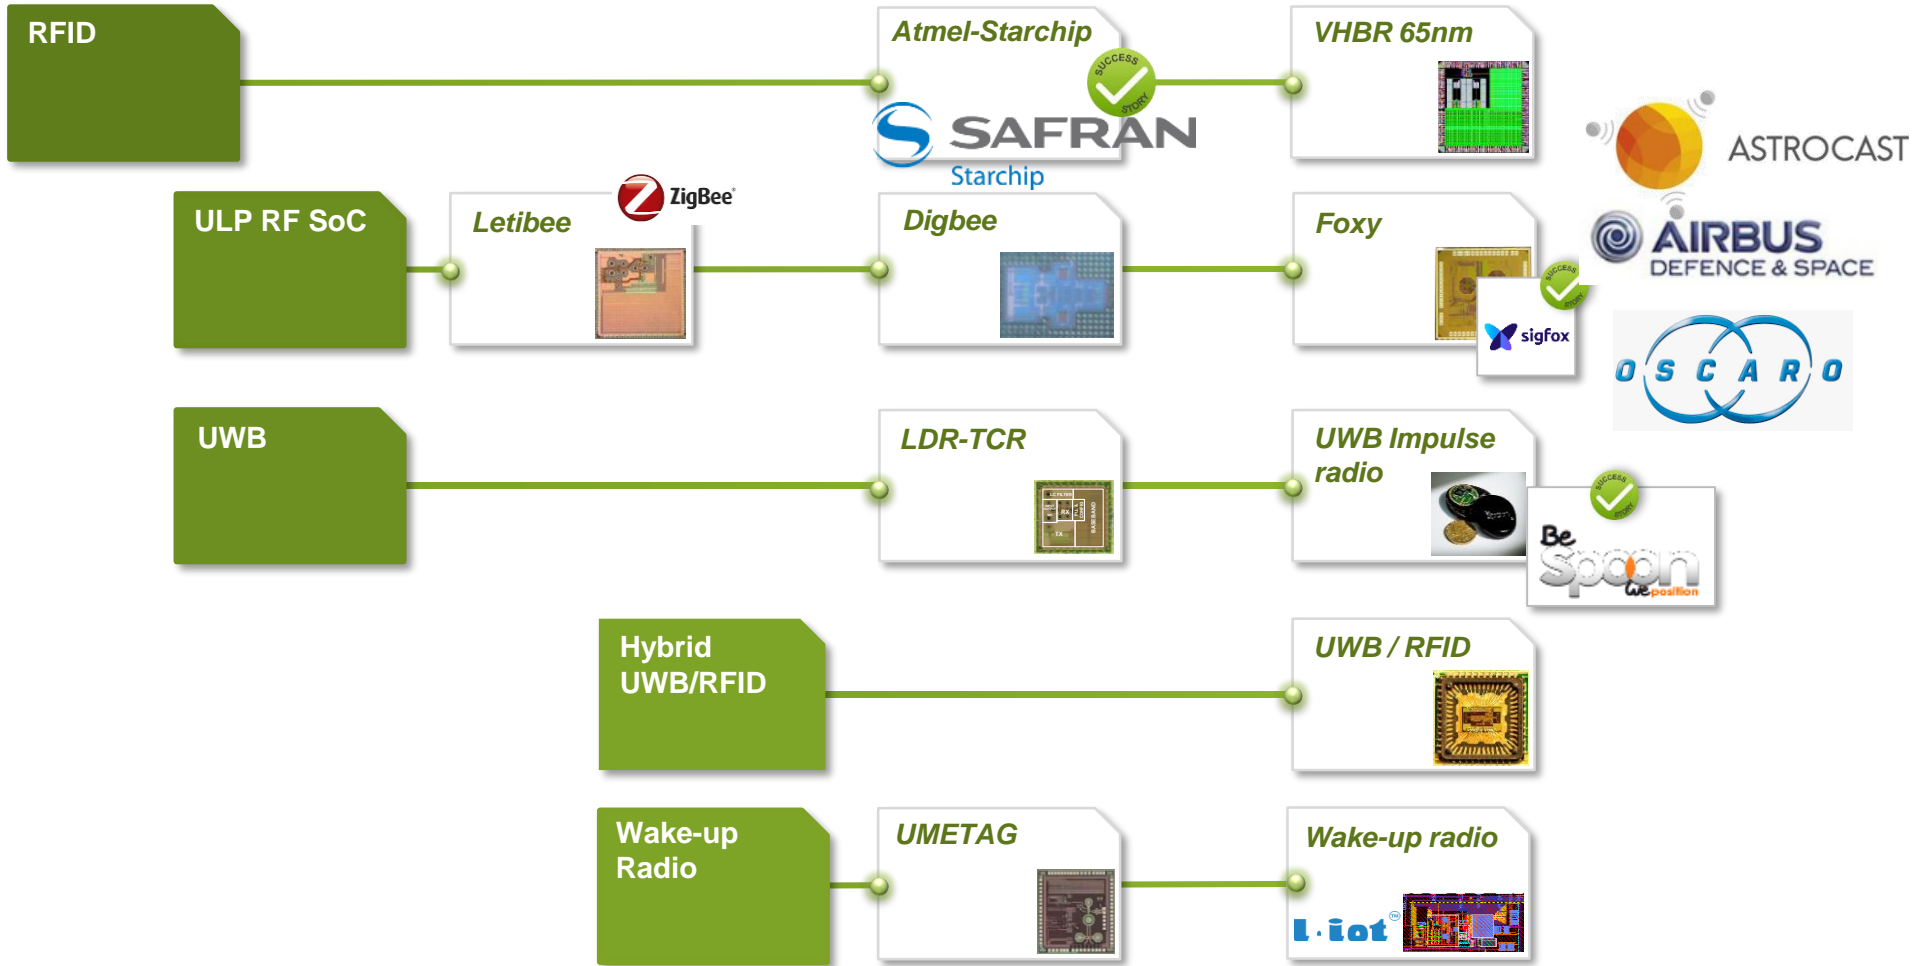


# A RISC-V ISA EXTENSION FOR ULTRA-LOW POWER IOT WIRELESS SIGNAL PROCESSING

Carolynn Bernier, Hela Belhadj Amor, Zdeněk Přikryl

Oct 1, 2019








## Motivation : A **software-defined** “Smart” wireless transceiver for IoT

- **PHY-agnostic solution for LPWA-IOT**



- Address « multi-mode » markets and lower hardware bug fix costs
- **Offer future-proofed designs to our clients**
  - Our clients’ advanced prototypes have evolving needs : satellite-IoT, Ultra-wide band localization, LPWA-IoT.
- **A new experimental platform**
  - Design new “RF software sensors”
  - Use light-weight ML algorithms to extract information from the RF signal



**Software-Defined Transceiver**



- Bottleneck : Existing software-defined radio (SDR) solutions are NOT ULP !**

TABLE IX  
COMPARISON OF EXISTING SDR PLATFORMS

	Programmability	Flexibility	Portability	Modularity	Computing Power	Energy Efficiency	Soft Core	FPGA	Cost (USD)
Imagine-based [151]	✓	×	×	×	Medium	Low	Imagine Stream Processor	N/A	N/A
USRP X300 [17]	✓	✓	×	✓	High	Low	PC	Xilinx Kintex-7	~ 4 - 5K Total
USRP E310 [17]	✓	✓	✓	✓	High	High	Dual-core ARM Cortex-A9	Xilinx Artix-4	~ 3K Total
KUAR [34]	✓	×	×	×	Medium	Low	Pc + 2x PowerPC cores	Xilinx Virtex II Pro	N/A
LimeSDR [146]	✓	✓	×	✓	High	Low	PC	Intel Cyclone IV	~ 300 Board Only
Ziria [147]	✓	✓	×	×	High	Low	PC	Depends on App	N/A
Sora [18]	✓	✓	×	×	High	Low	PC	Xilinx Virtex-5	~ 900 Board Only
SODA [67]	✓	✓	✓	×	High	High	ARM Cortex-M3 + Processing Elements	N/A	N/A
Iris [148]	✓	✓	✓	✓	High	High	Dual-core ARM Cortex-A9	Xilinx Kintex-4	~ 1.2K Total
Atomix [19]	✓	✓	✓	✓	High	Medium	TI 6670 DSP	N/A	~ 200 DSP Only
BeagleBoard-X15 [159]	✓	✓	✓	✓	High	Medium	2x TI C66x DSPs + 2x ARM Cortex-A15 & 2x M4	N/A	~ 270 Board Only
Airblue [20]	✓	✓	✓	✓	High	High	N/A	Intel Cyclone IV	~ 1.3K Board Only
WARP v3 [21]	✓	×	✓	✓	High	High	2x Xilinx MicroBlaze cores	Xilinx Virtex-6	~ 7K Total
PSoC 5LP [164]	✓	×	✓	✓	Low	High	ARM Cortex-M3	N/A	10 Board Only
Zynq-based [166]	✓	✓	✓	✓	High	High	Dual-core ARM Cortex-A9	Xilinx Kintex-4	~ 1.2K Total

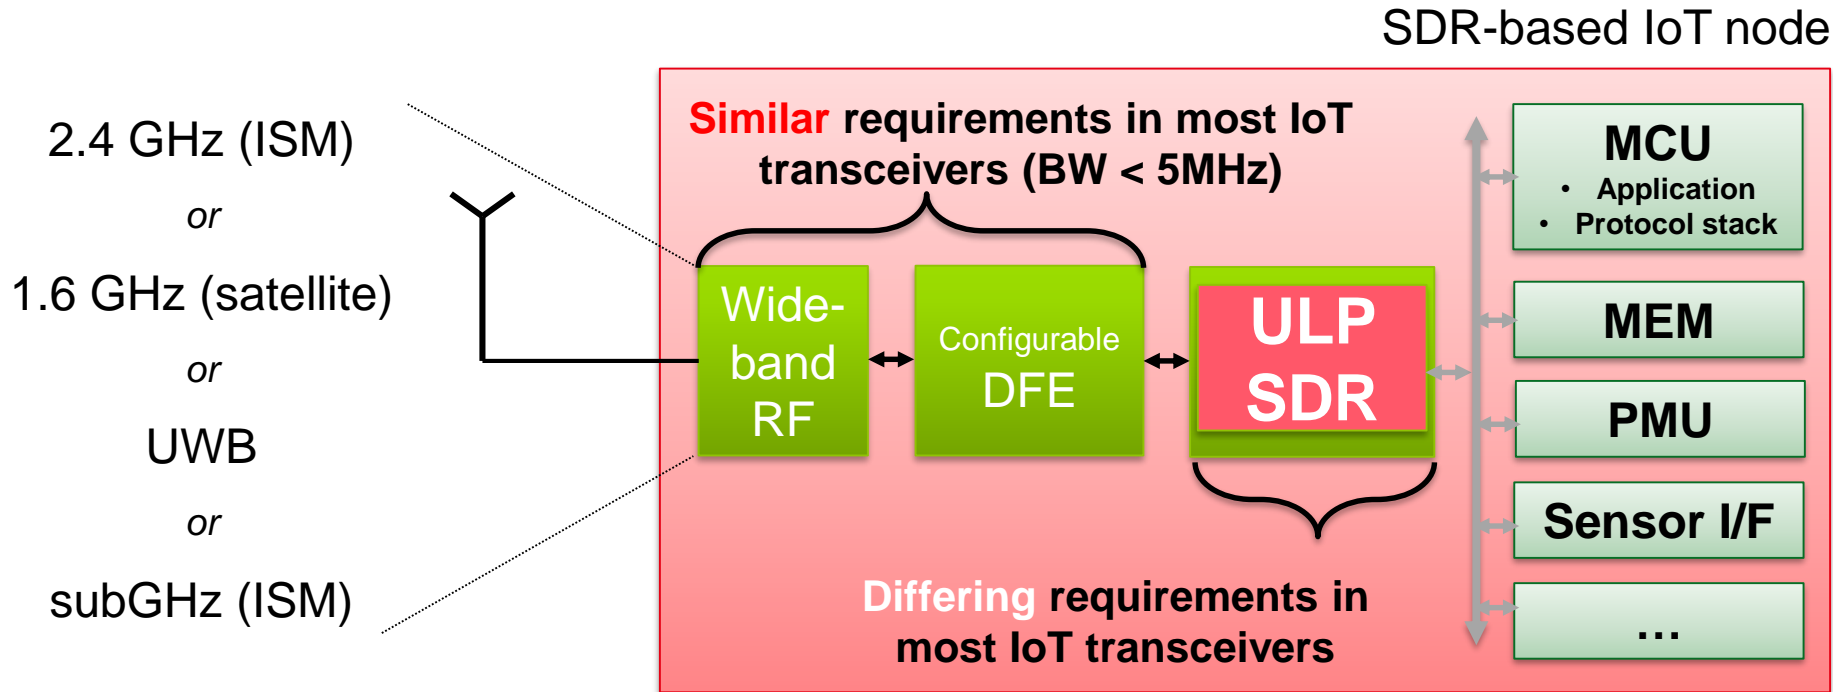


High cost (200 - 5K USD)  
General purpose → High power

[Akeela, 2018]



**Solution : Design of ULP-SDR**



Heterogeneous multi-core platform

Challenge:

Target mW-level power consumption



# SYSTEM REQUIREMENTS

- **Target Architecture**

- A very small and fast **core** (signoff ~300 MHz) associated to a TCPM and TCDM

- **Software DSP limited to decimated sample streams**

- DFE includes easily configurable and common HW operators : FIR filters, down-converters, AGC...

- **Real-time processing of complex samples**

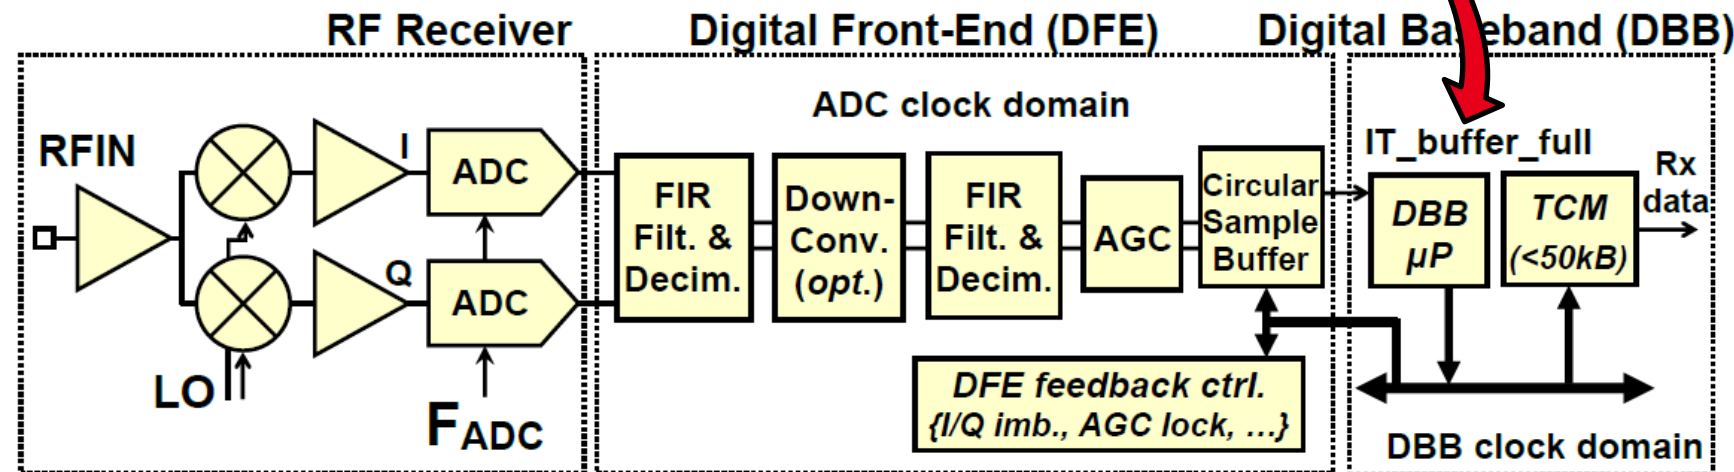
- Samples are temporarily stored in sample buffer and processed in blocks
- Integer processing only

- **Limit size of memory → big impact on power → configurable in size**

- TCPM (high speed non volatile)
- TCDM (stack usage !)
- Sample buffer
- Limit read/write to TCM


- **Single-cycle sleep**

- Wait for next block of samples
- Radio = OFF/ON





- **Wireless DSP requires linearity and low distortion**
  - Operators MUST NOT saturate
  - Operators MUST NOT overflow → but checking for overflows is too costly
- **Wireless DSP must conserve dynamic range (DR)**
  - The useful signal is often contained in the least significant bits
  - Beware of quantification noise → take care when rescaling the signal !
- **Most wireless signals are complex :  $i(t) + j*q(t)$** 
  - Frequent use of MUL, ADD, SUB, MAG, SHIFT, ... instructions on 8/16/32 bit **complex** data
- **Demodulation/compensation algorithms are mostly based on correlations**
  - i.e. multiplication
- **Input signal stream is typically  $\leq 8$  bits**
  - i.e. data streams are typically 8 / 16 / 32 bits
  - → fits well on a 32-bit machine

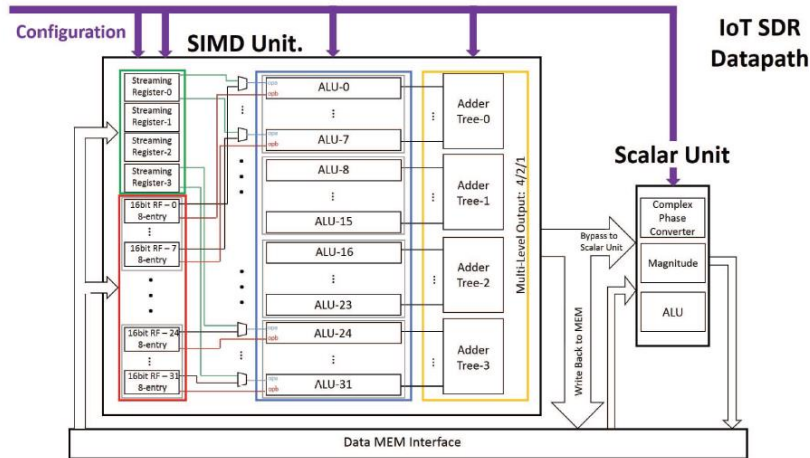




# WHICH PROCESSOR FOR OUR SDR ?

## Academic: Dedicated processors

### Custom SIMD [Chen, HPCA16]



✓ Promising power consumption

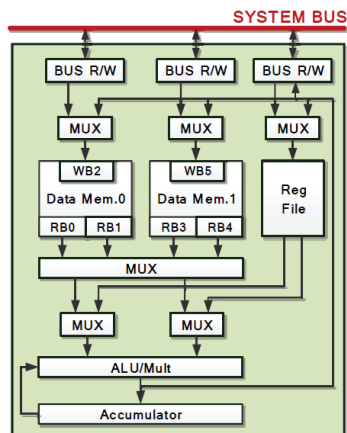
✗ Dedicated architectures → difficult to program

✗ No software tool-chains

✗ Low frequency clock → Large surface overheads

✗ Inefficient use of advanced CMOS nodes

### Custom MCU [Wu, GlobalSIP16]



## Commercial: GP processors, DSP



Previous work:

M3/M0+ vs. RISCY



[Belhadj, DATE19]

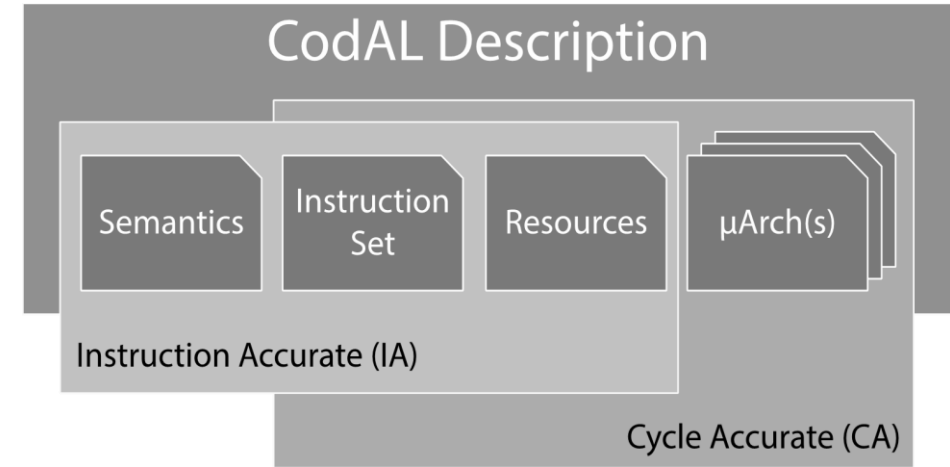
→ Lessons learned: GP processor can rival dedicated SoA processor architectures (with additional benefits)

→ Lessons learned: size of register file has huge impact on cycle count

**RISC-V advantage !**

→ Lessons learned: post-increment, HW loop, SIMD → not important in our test benches (mix of DSP computing and control)



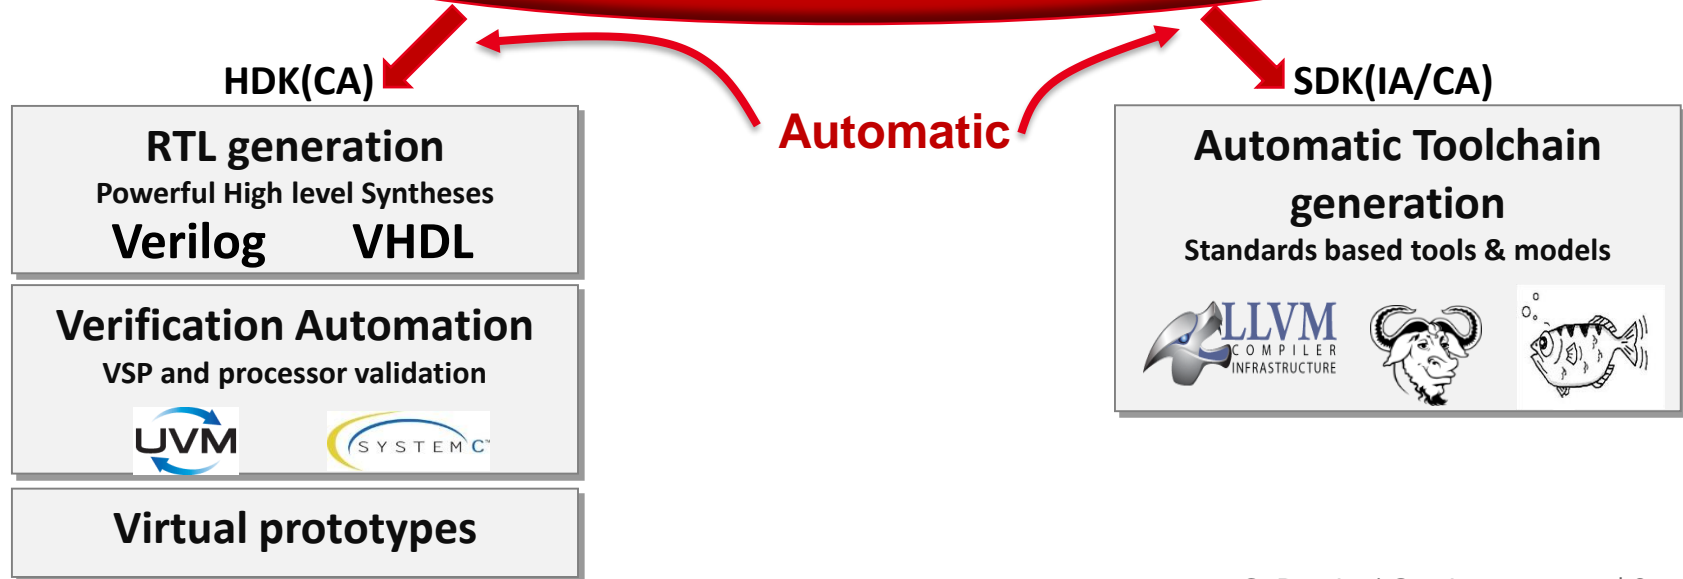


• RISC-V-based acceleration ?

• **Extend RISC-V ISA using dedicated instructions**

- Cudasip Studio : → An easy task ?
- Instruction Accurate (IA) model of new instructions
- Dedicated to RF DSP computing “zero cost” hardware implementation

**Cudasip Studio Toolset**



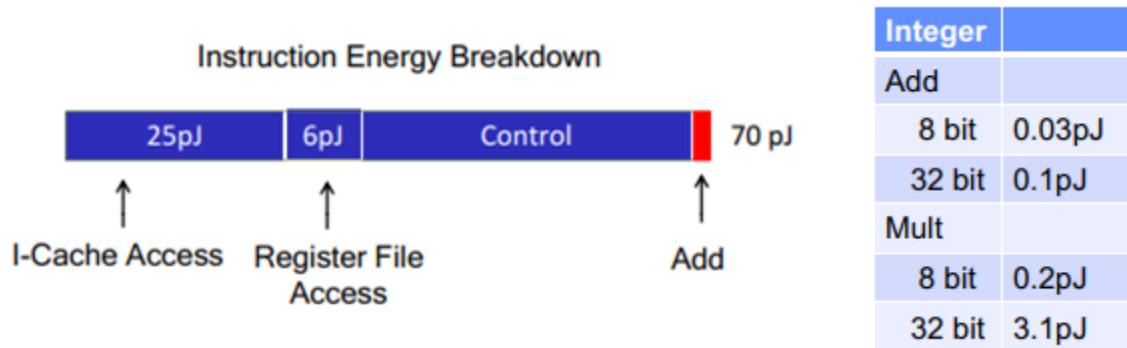


- **Wanted**

- Minimal set of USEFUL instructions.
- Only 32-bit opcodes for low decoding complexity.

- **Opportunities**

- Wide opcodes means up to 5 operands !
- First operation on 8-bit data is ALWAYS a complex multiplication
- Advanced CMOS allows single-cycle operators
- Tiny relative cost of ALU operators



45 nm, 0.9 V [M. Horowitz, ISSCC 2014]

## **REJECTED**

- **More general solution preferred :**
  - Halving variants (e.g. RADD)
- **Not clearly indispensable :**
  - CSMUL (complex-scalar multiply )
- **Useless :**
  - saturating instructions, MIN/MAX, 8 bit SIMD, CONJ



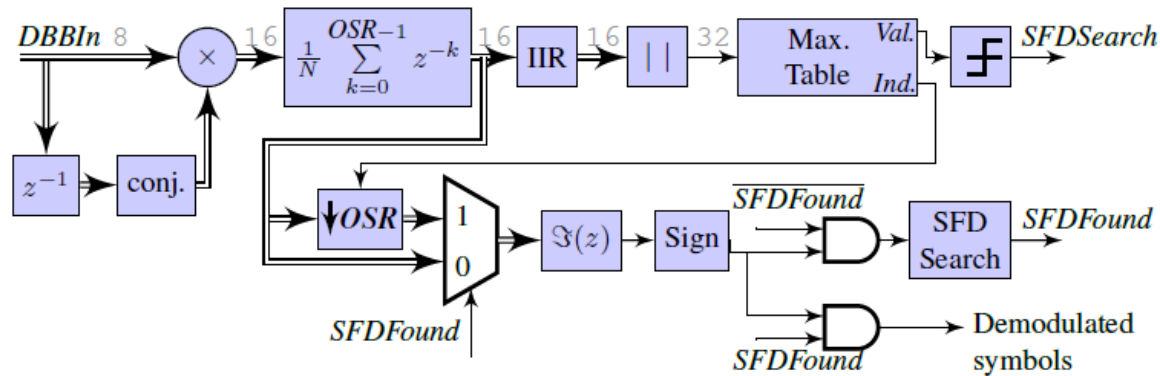
- 15 instructions using 3 major opcodes

Mnemonic	Instruction	Operation
<b>Zero-Cost Instructions</b>		
ADDC16 rd, rs1, rs2, imm	16-bit Addition & Shift Right Arithmetic Immediate	$rd.L = (rs1.L + rs2.L) \gg imm$ $rd.H = (rs1.H + rs2.H) \gg imm$
SUBC16 rd, rs1, rs2, imm	16-bit Subtraction	$rd.L = (rs1.L - rs2.L) \gg imm$ $rd.H = (rs1.H - rs2.H) \gg imm$
MUL2ADD16-32 rd, rs1, rs2, imm	Two "16x16" and Signed Addition	$rd = [(rs1.L * rs2.L) + (rs1.H * rs2.H)] \gg imm$
SRAC16 rd, rs1, imm	16-bit Shift Right Arithmetic Immediate	$rd.L = rs1.L \gg imm$ $rd.H = rs1.H \gg imm$
SLLC16 rd, rs1, imm	16-bit Shift Left Logical Immediate	$rd.L = rs1.L \ll imm$ $rd.H = rs1.H \ll imm$
CRASC16 rd, rs1, rs2, imm	16-bit Cross Add & Sub	$rd.L = (rs1.L + rs2.H) \gg imm$ $rd.H = (rs1.H - rs2.L) \gg imm$
CRSAC16 rd, rs1, rs2, imm	16-bit Cross Sub & Add	$rd.L = (rs1.L - rs2.H) \gg imm$ $rd.H = (rs1.H + rs2.L) \gg imm$
MULC8-16 rd, rs1, rs2, H1, H2, imm	Two "8x8" and Signed Subtraction  Two Crossed "8x8" and Signed Addition	if $Hx = 1$ , $\{ix, qx\} = \{rsx.B2, rsx.B3\}$ if $Hx = 0$ , $\{ix, qx\} = \{rsx.B0, rsx.B1\}$ $rd.L = ((i1 * i2) - (q1 * q2)) \gg imm$ $rd.H = ((i1 * q2) + (i2 * q1)) \gg imm$
MULC16 rd, rs1, rs2	Two "16x16" and Signed Subtraction Two Crossed "16x16" and Signed Addition	$rd.L = (rs1.L * rs2.L) \gg 16 - (rs1.H * rs2.H) \gg 16$ $rd.H = (rs1.H * rs2.L) \gg 16 + (rs1.L * rs2.H) \gg 16$
<b>Low-Cost Instructions</b>		
ADDC32 rd, rs1, rs2, imm	32-bit Addition & Shift Right Arithmetic Immediate	$rd = (rs1 + rs2) \gg imm$ $rd+1 = ((rs1+1) + (rs2+1)) \gg imm$
SUBC32 rd, rs1, rs2, imm	32-bit Subtraction	$rd = (rs1 - rs2) \gg imm$ $rd+1 = ((rs1+1) - (rs2+1)) \gg imm$
CRASC32 rd, rs1, rs2, imm	32-bit Cross Add & Sub	$rd = (rs1 + (rs2+1)) \gg imm$ $rd+1 = ((rs1+1) - rs2) \gg imm$
CRSAC32 rd, rs1, rs2, imm	32-bit Cross Sub & Add	$rd = (rs1 - (rs2+1)) \gg imm$ $rd+1 = ((rs1+1) + rs2) \gg imm$
MULC16-32 rd1, rd2, rs1, rs2, imm	Two "16x16" and Signed Subtraction Two Crossed "16x16" and Signed Addition	$rd1 = ((rs1.L * rs2.L) - (rs1.H * rs2.H)) \gg imm$ $rd2 = ((rs1.H * rs2.L) + (rs1.L * rs2.H)) \gg imm$
<b>Higher-Cost Instructions</b>		
MULC32 rd, rs1, rs2, rs3, rs4	Two "32x32" and Signed Subtraction Two Crossed "32x32" and Signed Addition	$rd = (rs1 * rs3) \gg 32 - (rs2 * rs4) \gg 32$ $rd+1 = (rs2 * rs3) \gg 32 + (rs1 * rs4) \gg 32$

- « Zero-cost »
  - Reconfigurable HW
  - Systematic output DR adjust
- « Low-cost »
  - 4 output / 2 input port register file
  - Duplicated ALU
- « Higher-cost »
  - 3 more 32-bit multipliers



## Testbench 1: FSK demodulation




## Testbench 3: 16 and 32-bit FFT

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N}nk} = \sum_{n=0}^{N-1} x_n W_N^{kn}$$

- Radix-4 decimation-in-frequency, complex FFT with bit-reversed outputs, N = 128, 2048
- Based on source code from a port of the ARM CMSIS DSP library to RISC-V

## Testbench 2: LoRa preamble synchronization

- Spreading Factor (SF) = 7, 11



## Testbench 4: CORDIC algorithm

$$x_{i+1} = x_i - s_i \cdot y_i \cdot 2^{-i}$$

$$y_{i+1} = y_i + s_i \cdot x_i \cdot 2^{-i}$$

$$z_{i+1} = z_i - s_i \cdot \text{atan}(2^{-i})$$

- 10 iteration CORDIC algorithm applied to 32-bit complex input data.



# RESULTS


Power Model	Baseline	+Extensions
All instr. except NOP and MUL	1	1.05
MUL	1.14	1.14
MULC16-32 / MULC16	-	1.3
MULC32	-	1.59

- Expect at least ~50% power reductions with reduced clock and VDD.

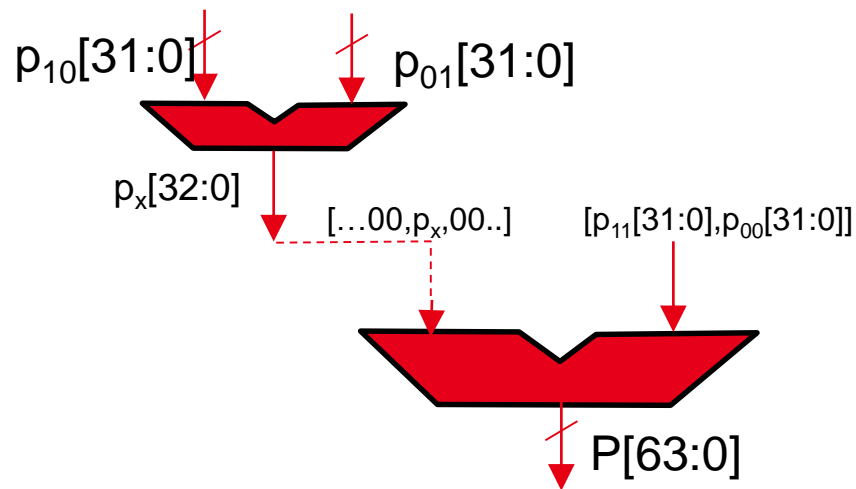


Testbench	Cycle count improvement (IA model)	Energy improvement (est.)
FSK Demod	22 %	
LoRa, SF=7	49 %	46 %
LoRa, SF=11	52 %	50 %
16-bit FFT, N=128	55 %	53 %
16-bit FFT, N=2048	57 %	55 %
32-bit FFT, N=128	34 %	32 %
32-bit FFT, N=2048	34 %	30 %
32-bit CORDIC, 10 iteration	28 %	

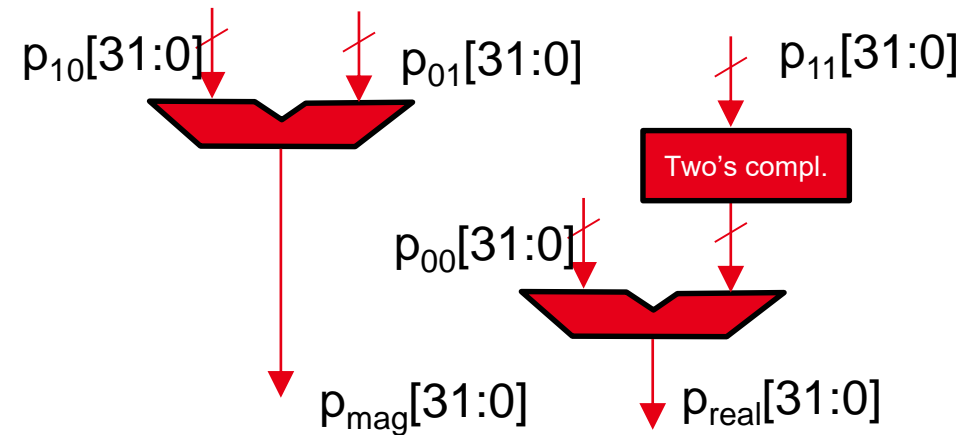
- **Finish CA model & run Power/Area analysis in 22 nm**
  - Reconfigurable hardware blocks designed in CodAL. Ex: 32-bit multiplication



### CASE : 32-bit integer multiplication



### CASE : 16-bit complex multiplication



## **Special thanks to :**

**Hela Belhadj Amor**

**Zdeněk Přikryl**

**Jerry Ardizzone**

**And**

**Ivan Miro Panades**

**Yves Durand**

**Henri-Pierre Charles**

**Simone Bacles-Min**

**Romain Lemaire**

**... and all of LISAN !**

Leti, technology research institute

Commissariat à l'énergie atomique et aux énergies alternatives

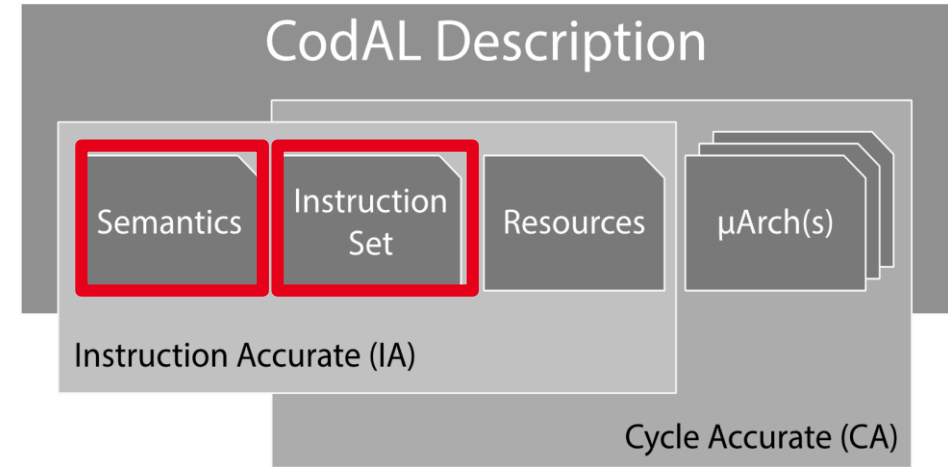
Minatec Campus | 17 rue des Martyrs | 38054 Grenoble Cedex | France

[www.leti.fr](http://www.leti.fr)





- Step 1 : ISA exploration using IA model



```

element opc_name
{
    use instance_data_type as name of
instances;

    assembler {textual form of the instruction};
    binary {The instructions's binary coding};

    semantics
    {
        The instruction's behavior is described
        using a subset of the ANSI C
        language.
    };
};

```

Used by IA and CA models

```

element i_load
{
    use opc_load as opc;
    use gpr_all as gpr_dst, gpr_src1;
    use simm12;

    assembler { opc gpr_dst ", " simm12 "(" gpr_src1 ")" };
    binary { simm12 gpr_src1 opc[OPC_FRAG1] gpr_dst opc[OPC_FRAG0] };

    semantics
    {
        ADDR_TYPE address;
        WORD_TYPE result;

        codasip_compiler_schedule_class(sc_load);

        address = rf_gpr_read(gpr_src1) + simm12;
        result = load(opc, address);
        rf_gpr_write(result, gpr_dst);
    };
};

```

Used by IA model

Call to memory interface if\_ldst



# RECONFIGURABLE MULTIPLIER (8 BIT EXAMPLE HERE)

State 1 : the block performs 8-bit integer multiplication

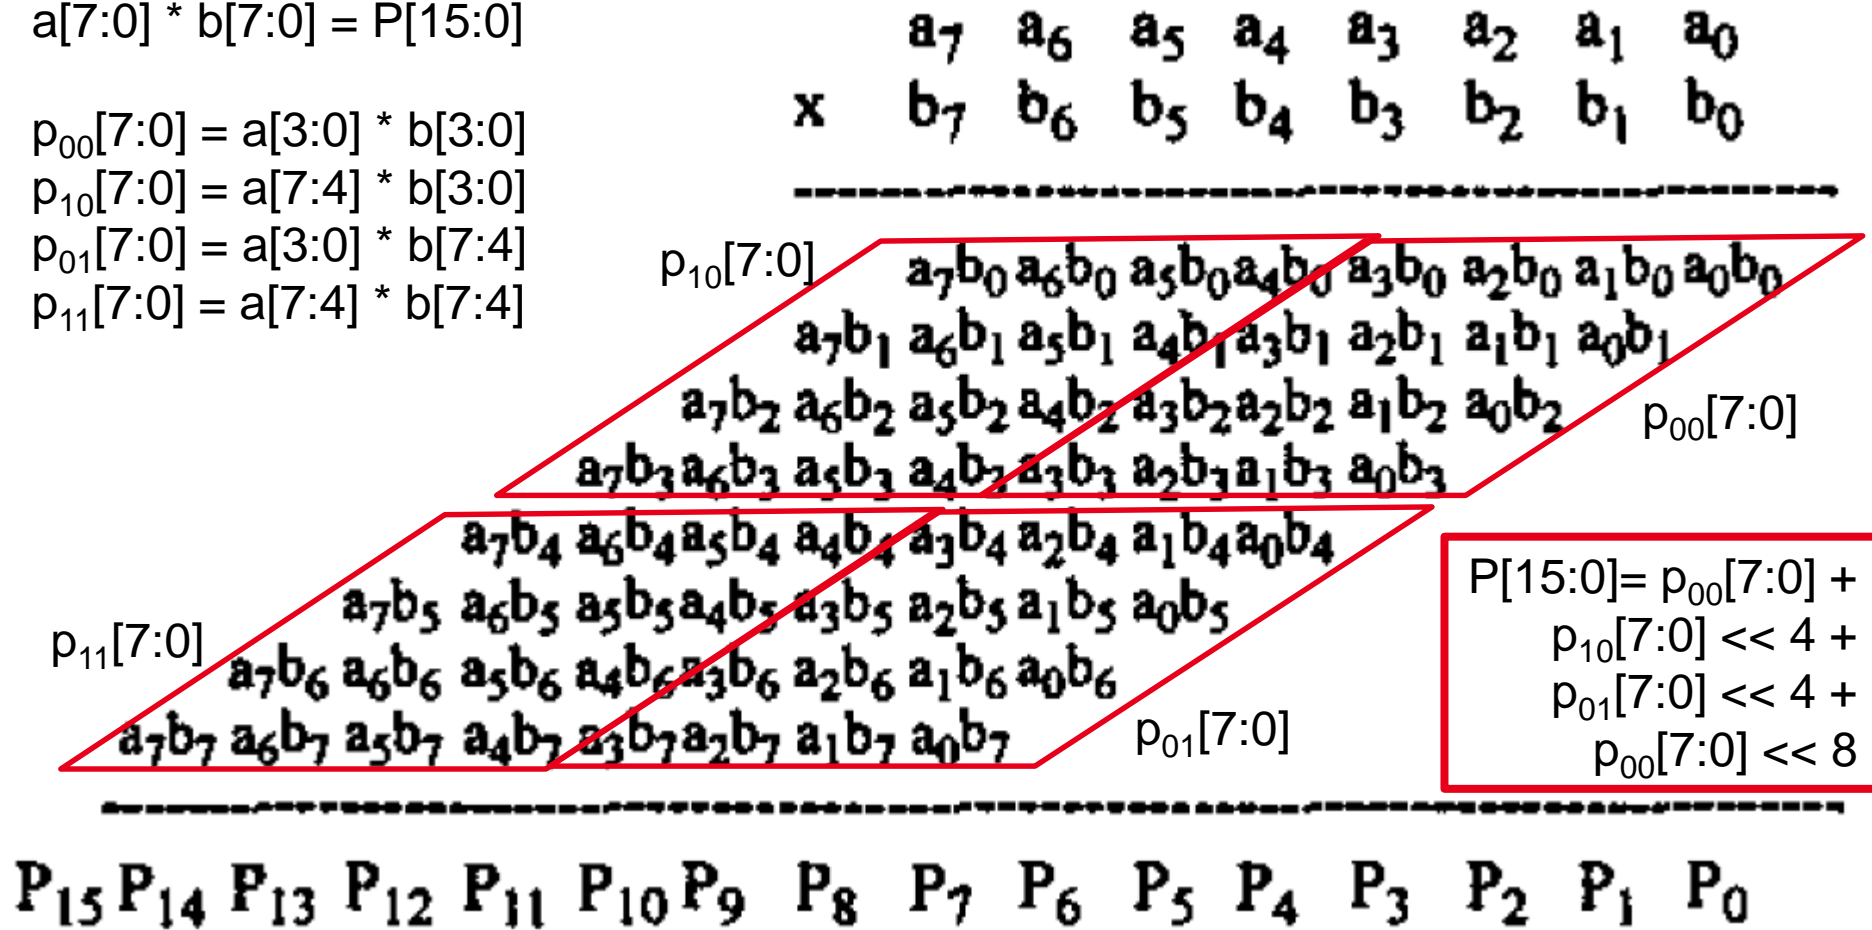
$$a[7:0] * b[7:0] = P[15:0]$$

$$p_{00}[7:0] = a[3:0] * b[3:0]$$

$$p_{10}[7:0] = a[7:4] * b[3:0]$$

$$p_{01}[7:0] = a[3:0] * b[7:4]$$

$$p_{11}[7:0] = a[7:4] * b[7:4]$$



## RECONFIGURABLE MULTIPLIER (8 BIT EXAMPLE HERE)

State 2 : the block performs a 4-bit **complex** integer multiplication :

$$(I_1 + j*Q_1) * (I_2 + j*Q_2) = P_{\text{real}} + j*P_{\text{imag}}$$

$Q_1$	$a_7$	$a_6$	$a_5$	$a_4$	$a_3$	$a_2$	$a_1$	$a_0$	$I_1$
$Q_2$	$b_7$	$b_6$	$b_5$	$b_4$	$b_3$	$b_2$	$b_1$	$b_0$	$I_2$

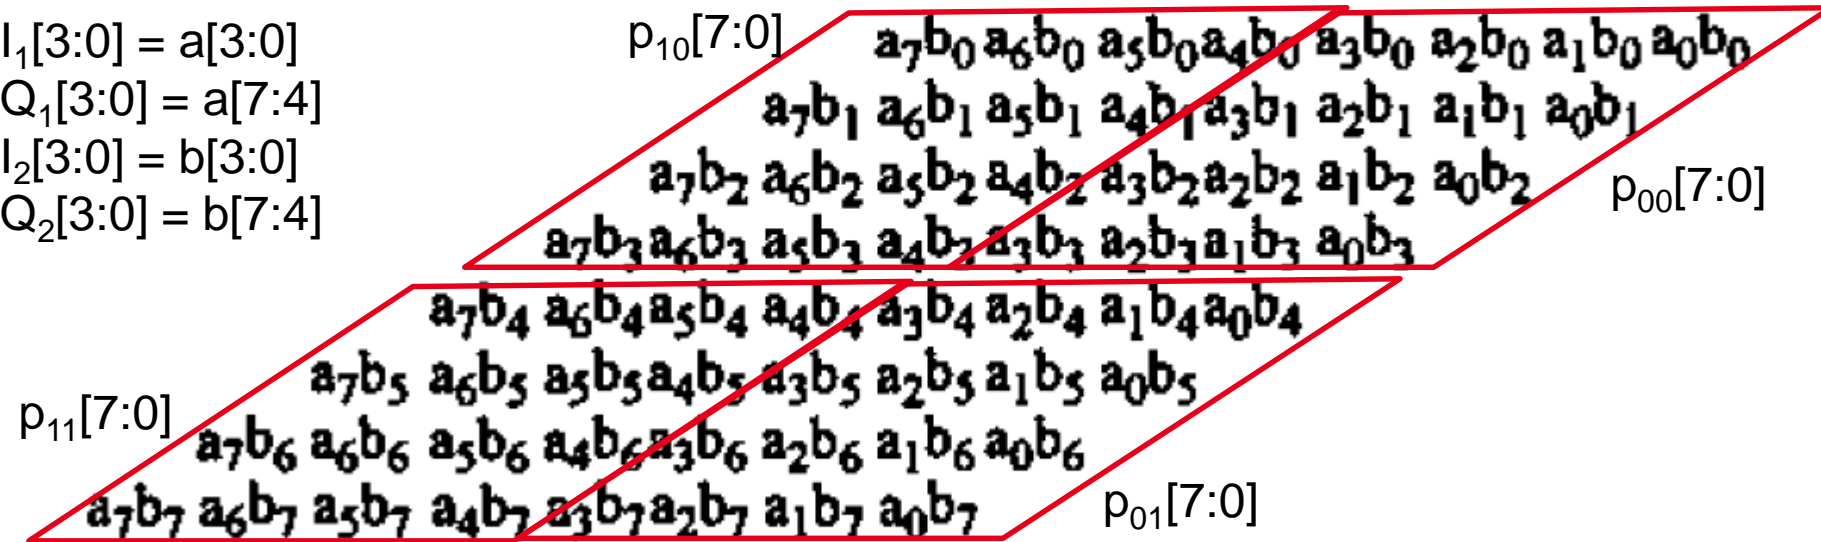
Input is redefined:

$$I_1[3:0] = a[3:0]$$

$$Q_1[3:0] = a[7:4]$$

$$I_2[3:0] = b[3:0]$$

$$Q_2[3:0] = b[7:4]$$




~~$P_{15} P_{14} P_{13} P_{12} P_{11} P_{10} P_9 P_8 P_7 P_6 P_5 P_4 P_3 P_2 P_1 P_0$~~

# RECONFIGURABLE MULTIPLIER (8 BIT EXAMPLE HERE)

State 2 : the block performs a 4-bit **complex** integer multiplication :

$$(I_1 + j*Q_1) * (I_2 + j*Q_2) = P_{\text{real}} + j*P_{\text{imag}}$$

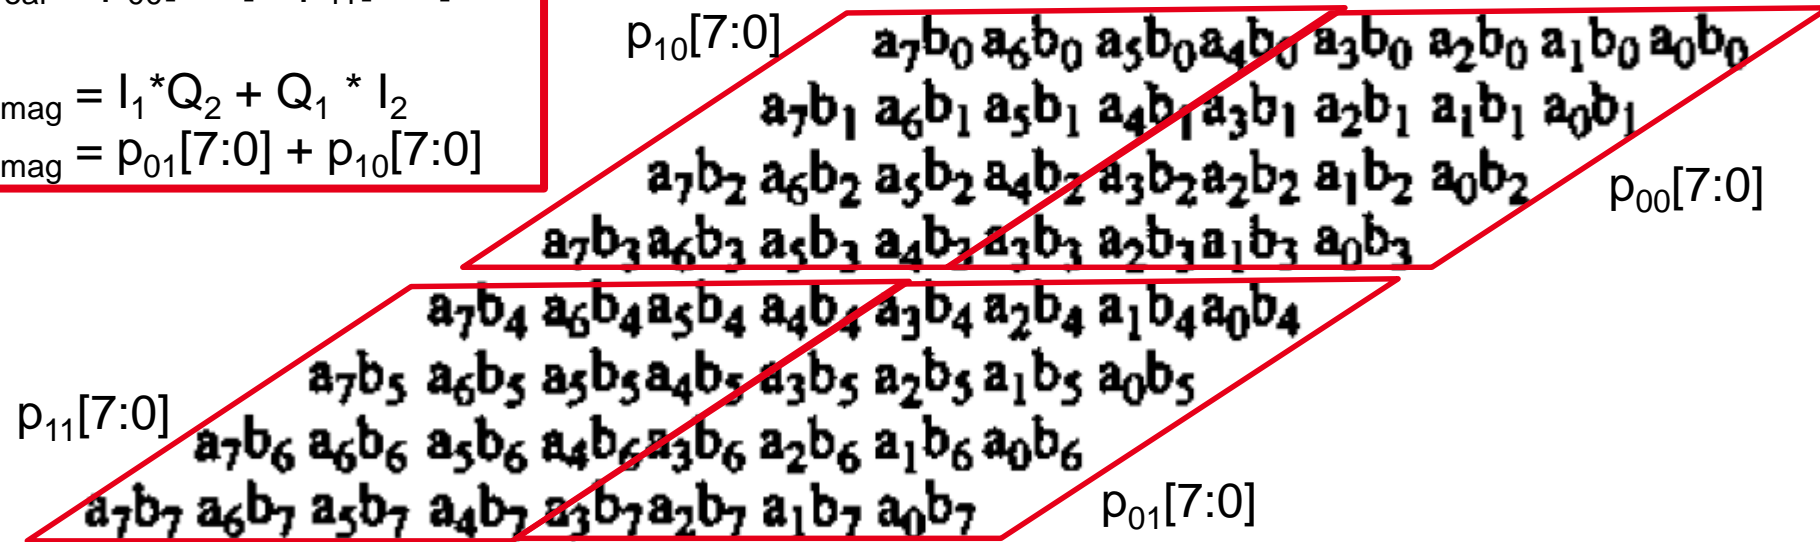


$$P_{\text{real}} = I_1 * I_2 - Q_1 * Q_2$$

$$P_{\text{real}} = p_{00}[7:0] - p_{11}[7:0]$$

$$P_{\text{imag}} = I_1 * Q_2 + Q_1 * I_2$$

$$P_{\text{imag}} = p_{01}[7:0] + p_{10}[7:0]$$



~~$P_{15} P_{14} P_{13} P_{12} P_{11} P_{10} P_9 P_8 P_7 P_6 P_5 P_4 P_3 P_2 P_1 P_0$~~