How to teach a billion transistor chip a new trick

Dr. ir. Bart Kienhuis
University Leiden, LIACS
Computer Systems Group
http://www.liacs.nl

ICT-Kenniscongres 10 April 2006

Outline

- Focus is programming ICs for Stream based Applications
  - FPGAs
- Stream Based Applications
  - Multi-media
  - Imaging
  - Bio-informatics
  - Classical DSP
- Increasing demand for high-performance, embedded compute performance
- We know how to build billion transistor FPGAs, but cannot program efficiently

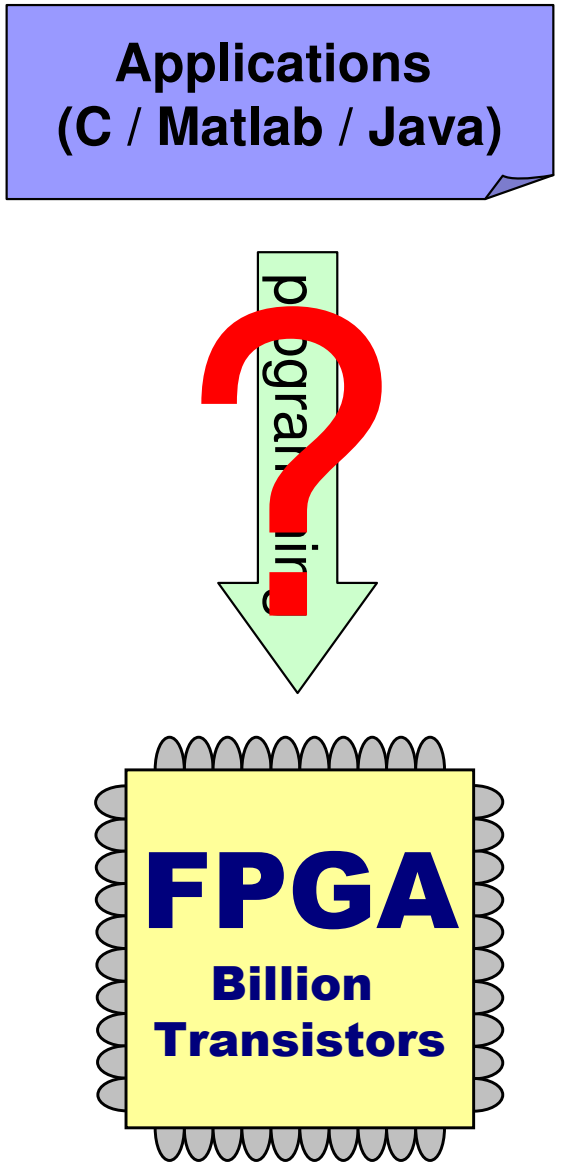

Stream Based Applications

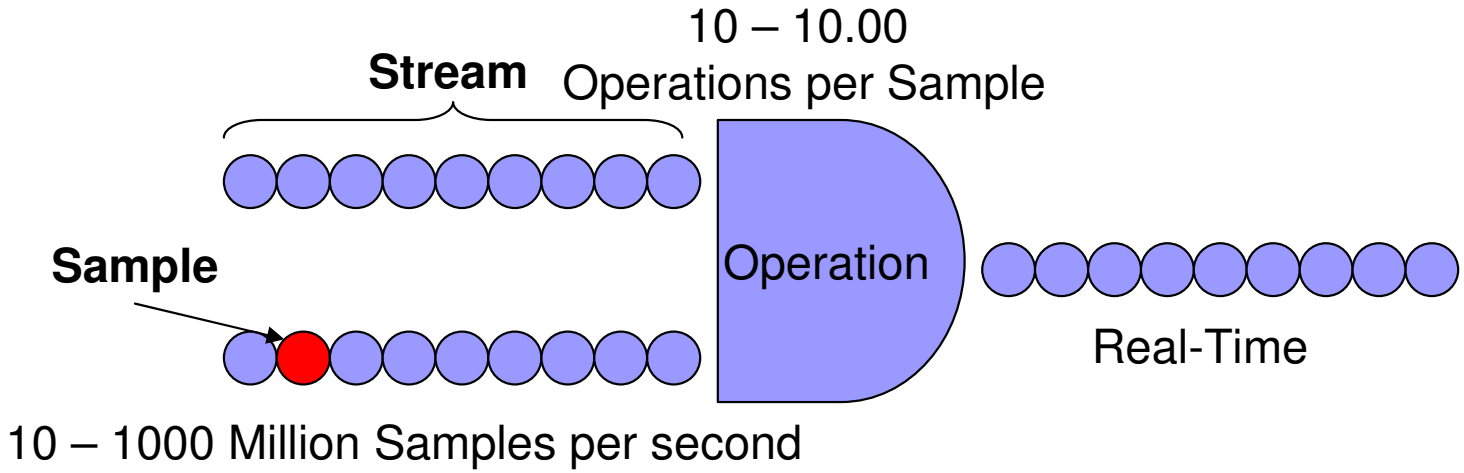

- Imaging
  - ***HDTV*** 1080x720 @ 100 Hz
    - 77 Million samples per seconds
    - Suppose 100 operations per sample
    - **8 Billion** Operations per second
- Classical DSP
  - ***LOFAR*** One antenna
    - Sampling rate 200*2 = 400 Million Samples per second
    - 100 Operations per second (First stage)
    - **40 Billion** Operations per second
    - There are 10.000 Antennas expected = 40*10^13 operations per second
    - Next beamforming is performed, statistics obtained, but on decimated signal

Compute Requirements ?

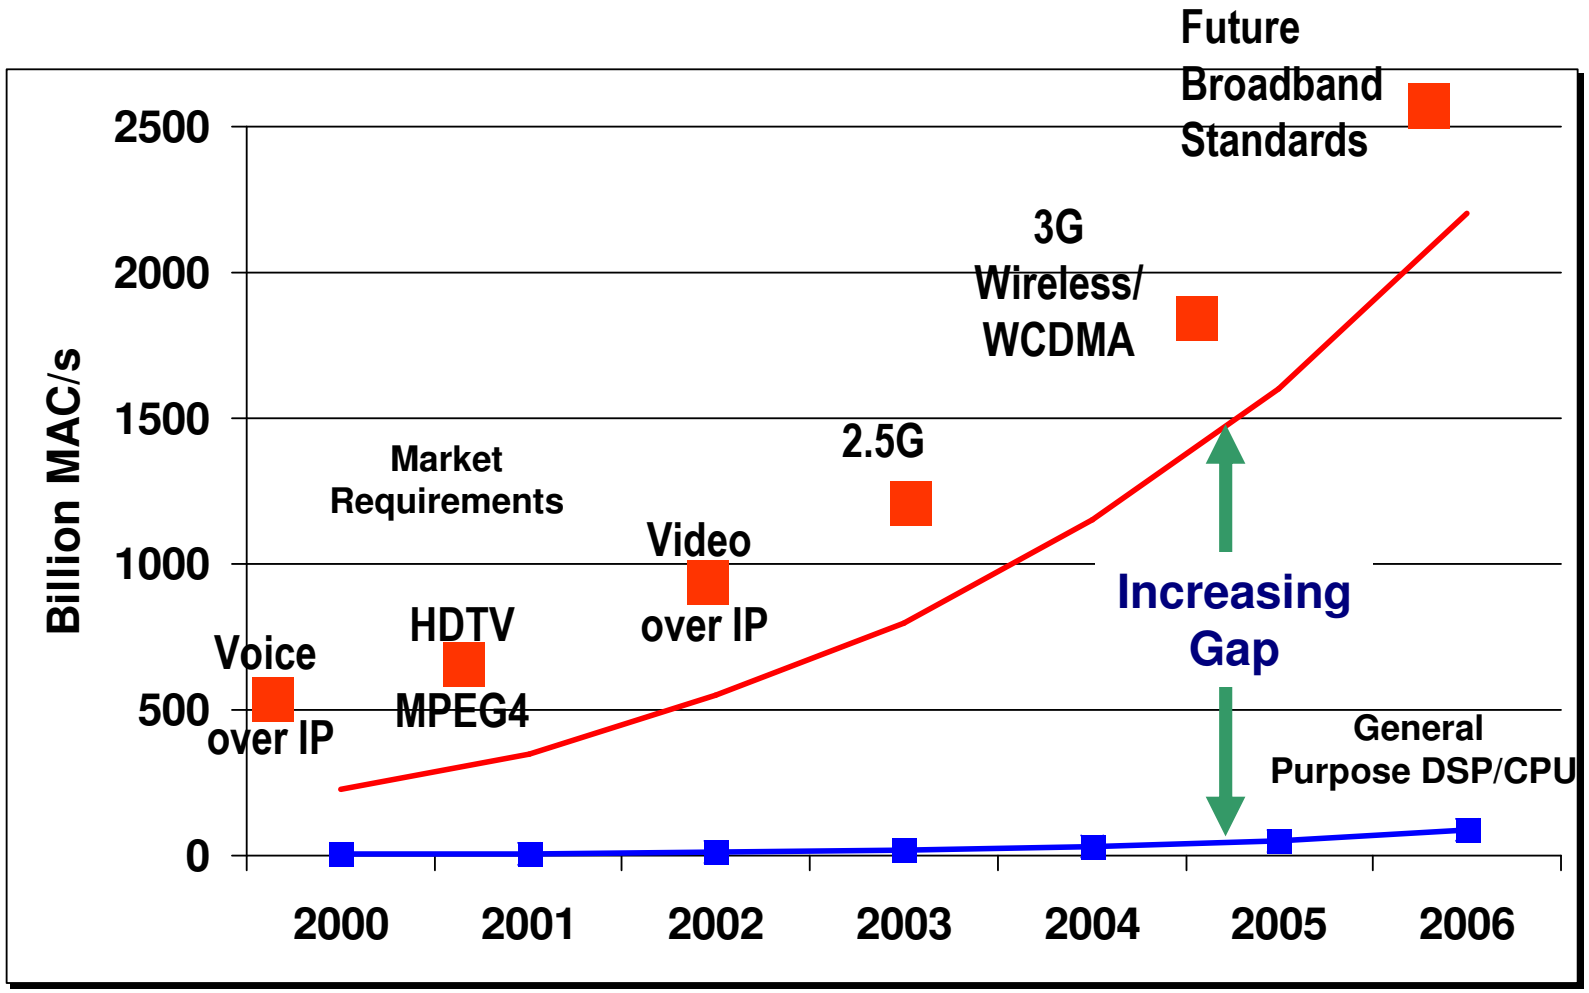

Source: TI, Xilinx – 1 MAC = 8 bit Multiply-Accumulate

Trend: ferocious appetite for more *embedded* compute power

Intrinsic Computational Efficiency (ICE)

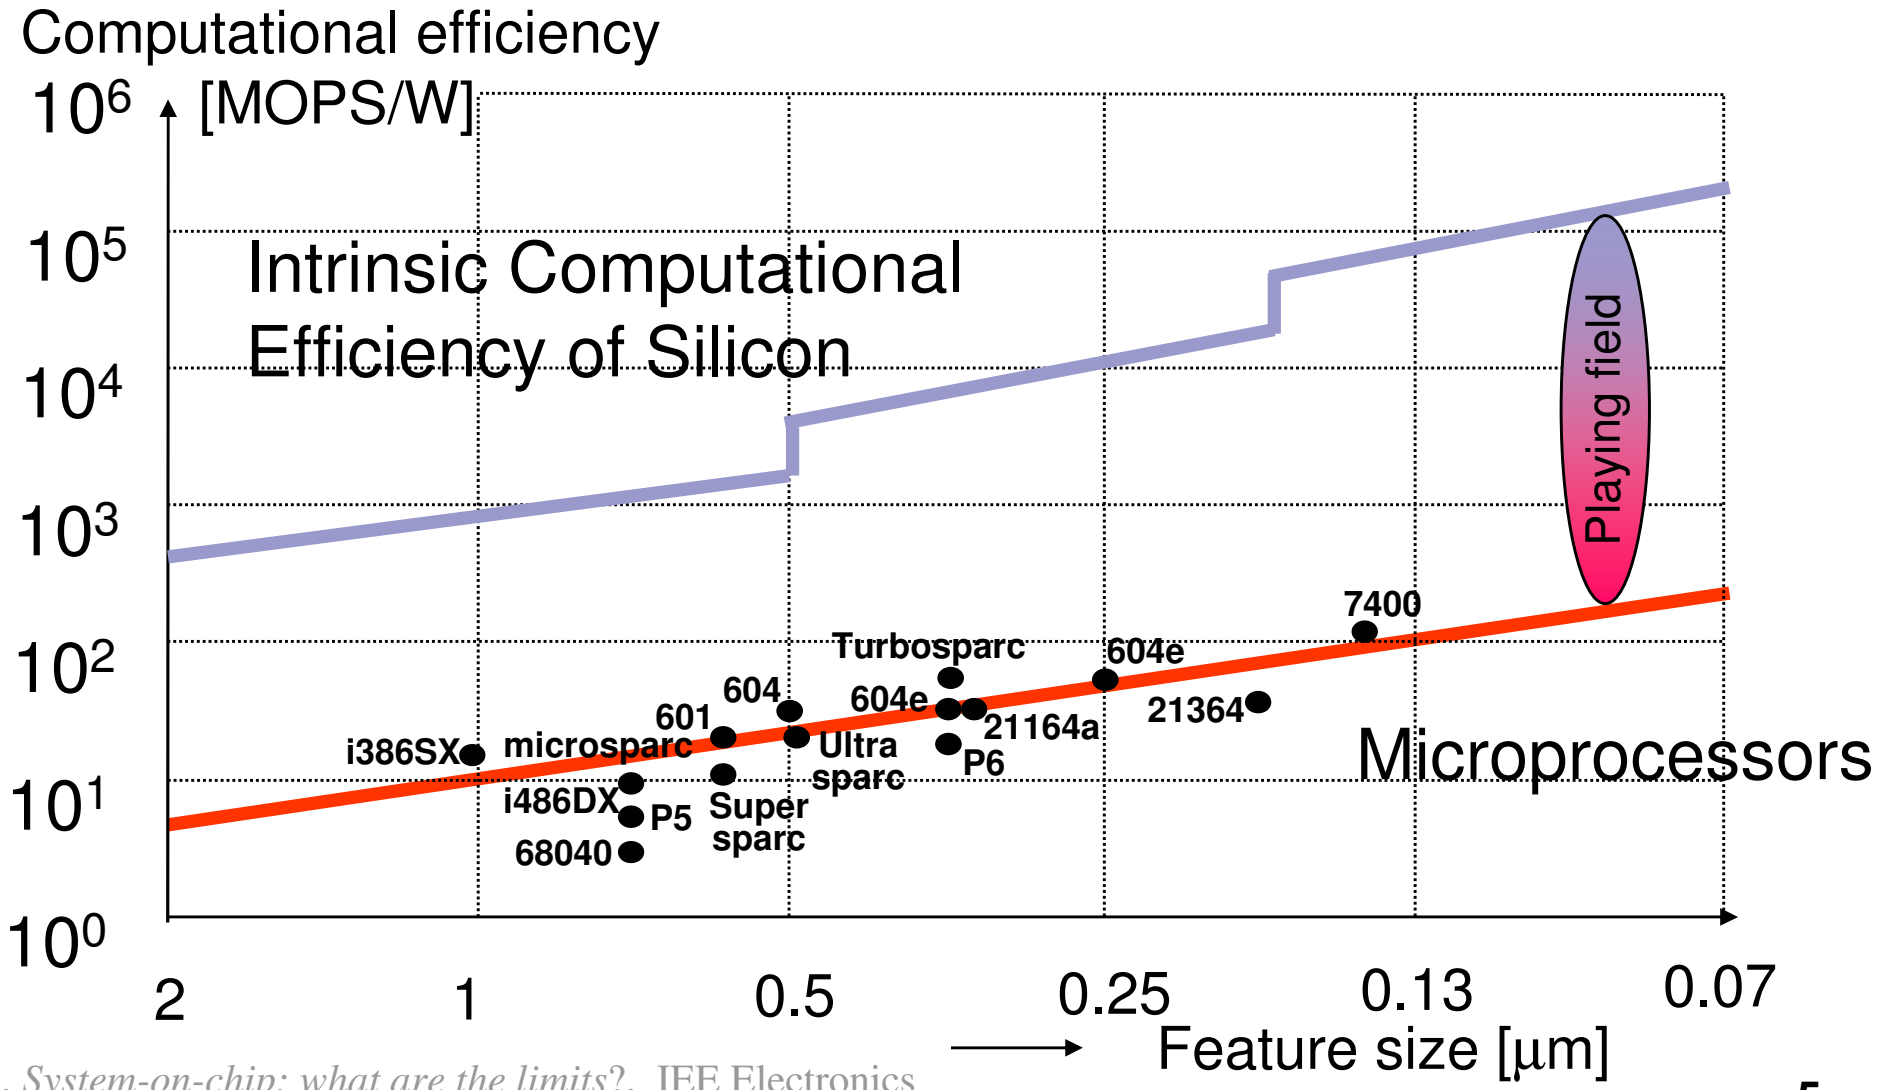

E. Roza, *System-on-chip: what are the limits*?, IEE Electronics Communication Engineering Journal, vol13, No6, Dec 2001, pp 249-255.

Xilinx Virtex II Pro

- Heterogeneous Architecture
  - Multi CPUs
  - Distributed Memory Blocks
  - Programmable Logic (IP cores)
- Commercially Available Platform
  - Virtex-4 is released and available
  - Virtex-5 is on the drawing board
- FPGAs are becoming more important in Embedded System Design
  - Flexible since they can easily be re-programmed
  - More functionality at lower cost
  - replacing a larger part of the ASICs market

Heterogeneous Architecture

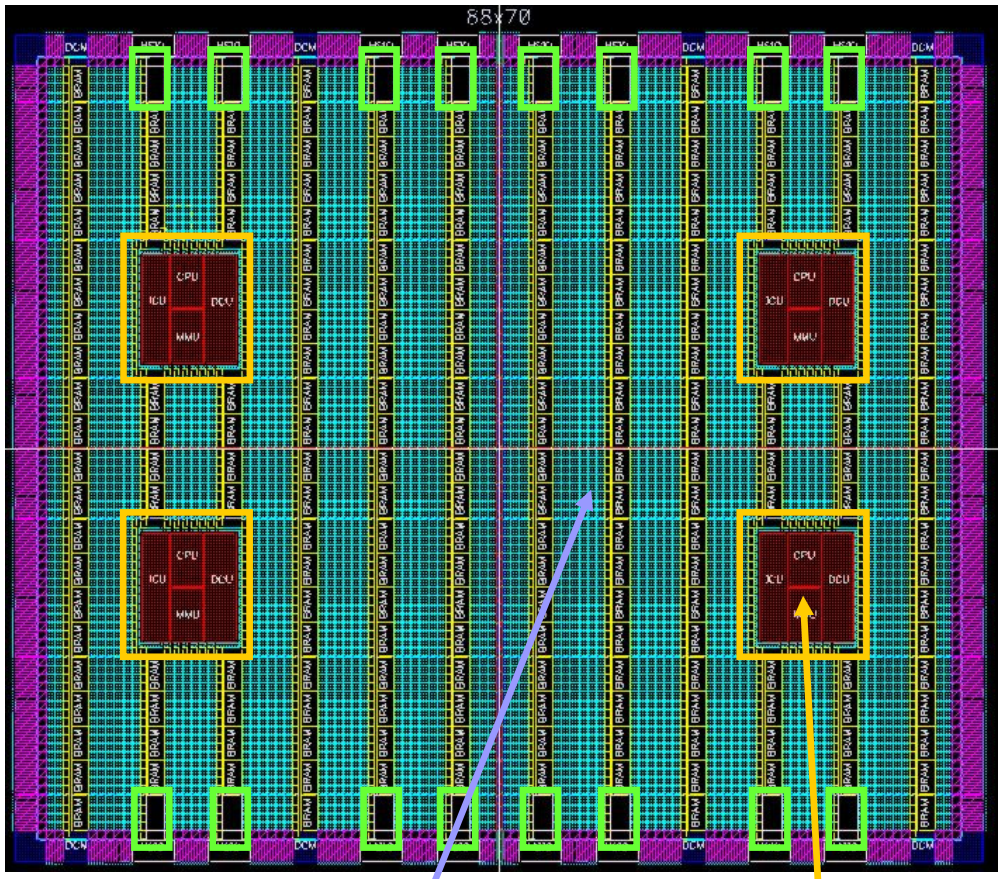

Building an FPGA: Logic First

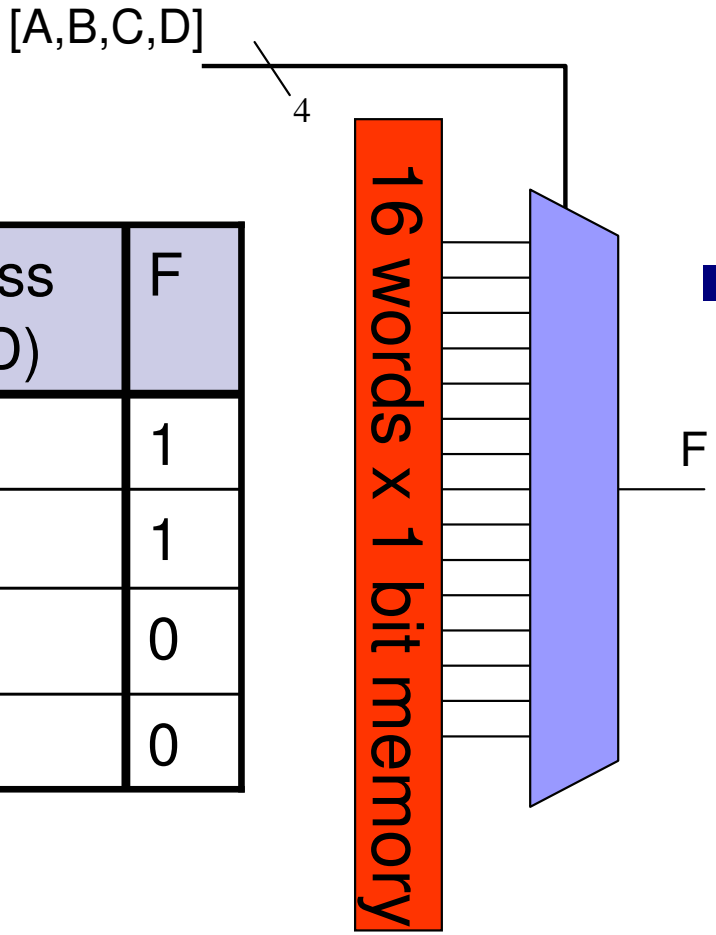

Address (ABCD)	F
0001	1
0010	1
1010	0
1011	0

- A 4-input lookup table (LUT) can implement any function of 4 inputs.
- For example, a 1-bit adder needs 2 LUTs:

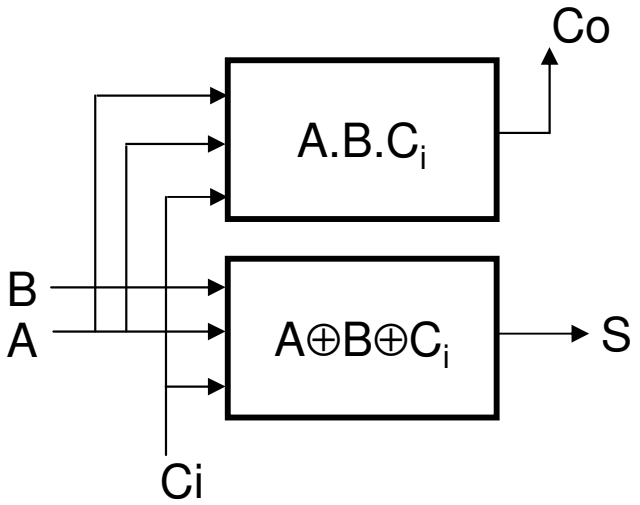

Arithmetic, Distributed RAM

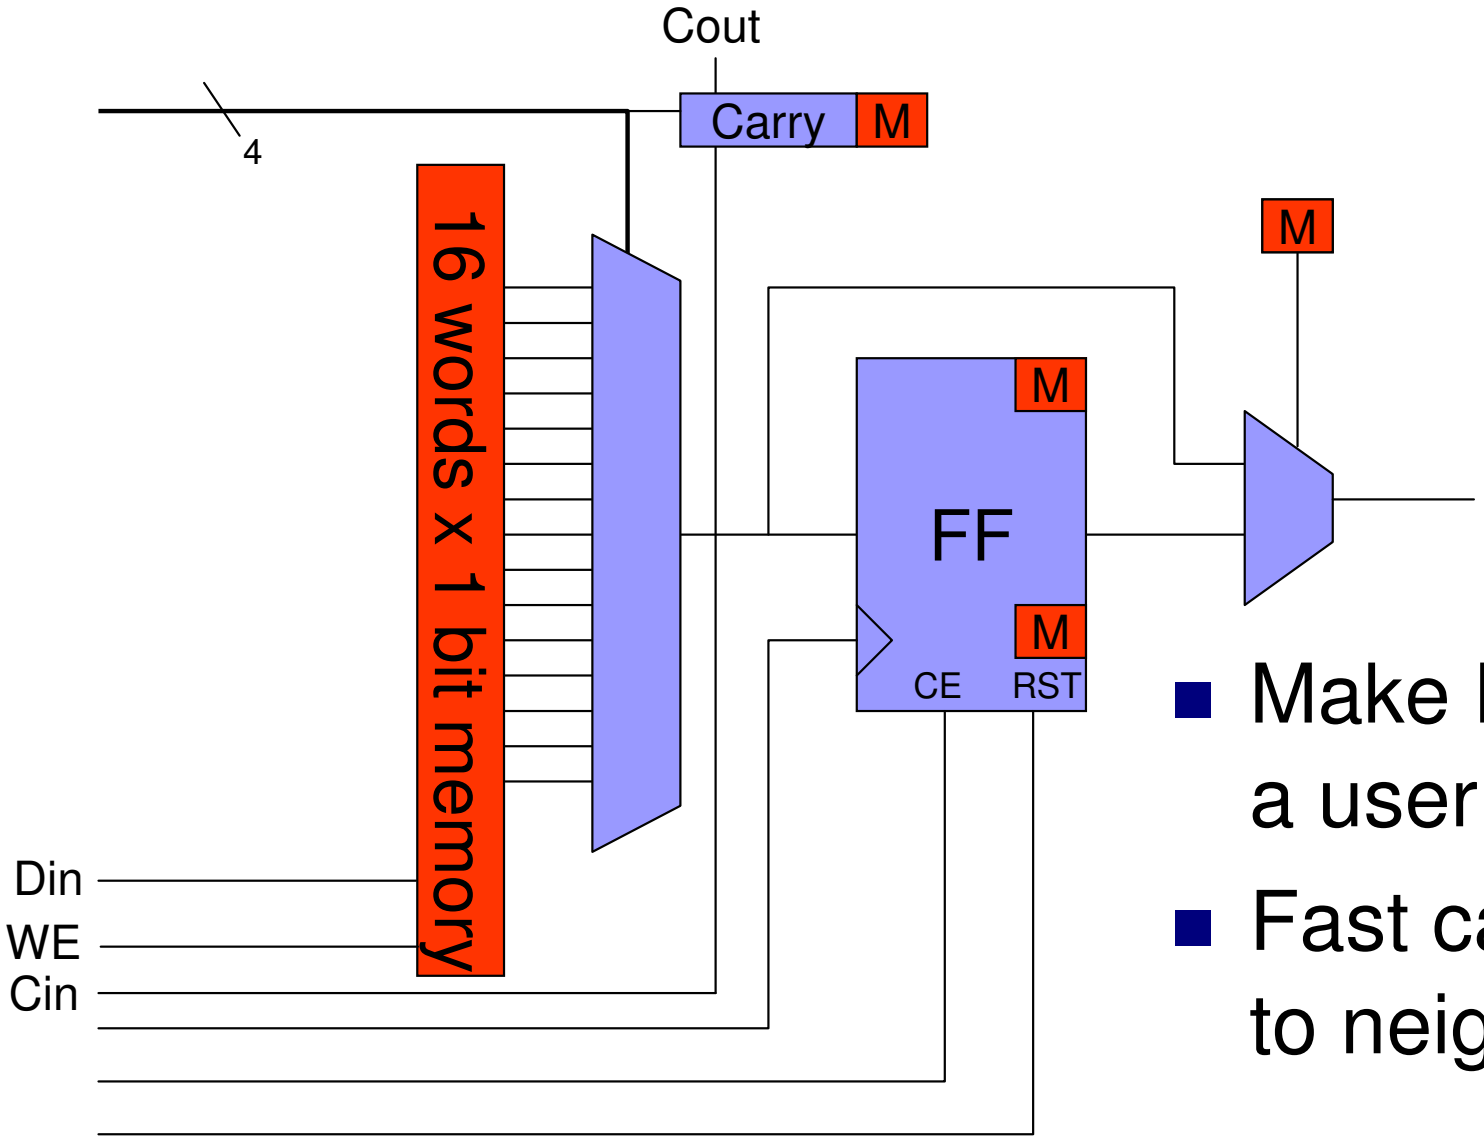

- Make LUT RAM a user resource.
- Fast carry ripple to neighbor.

Add Interconnect

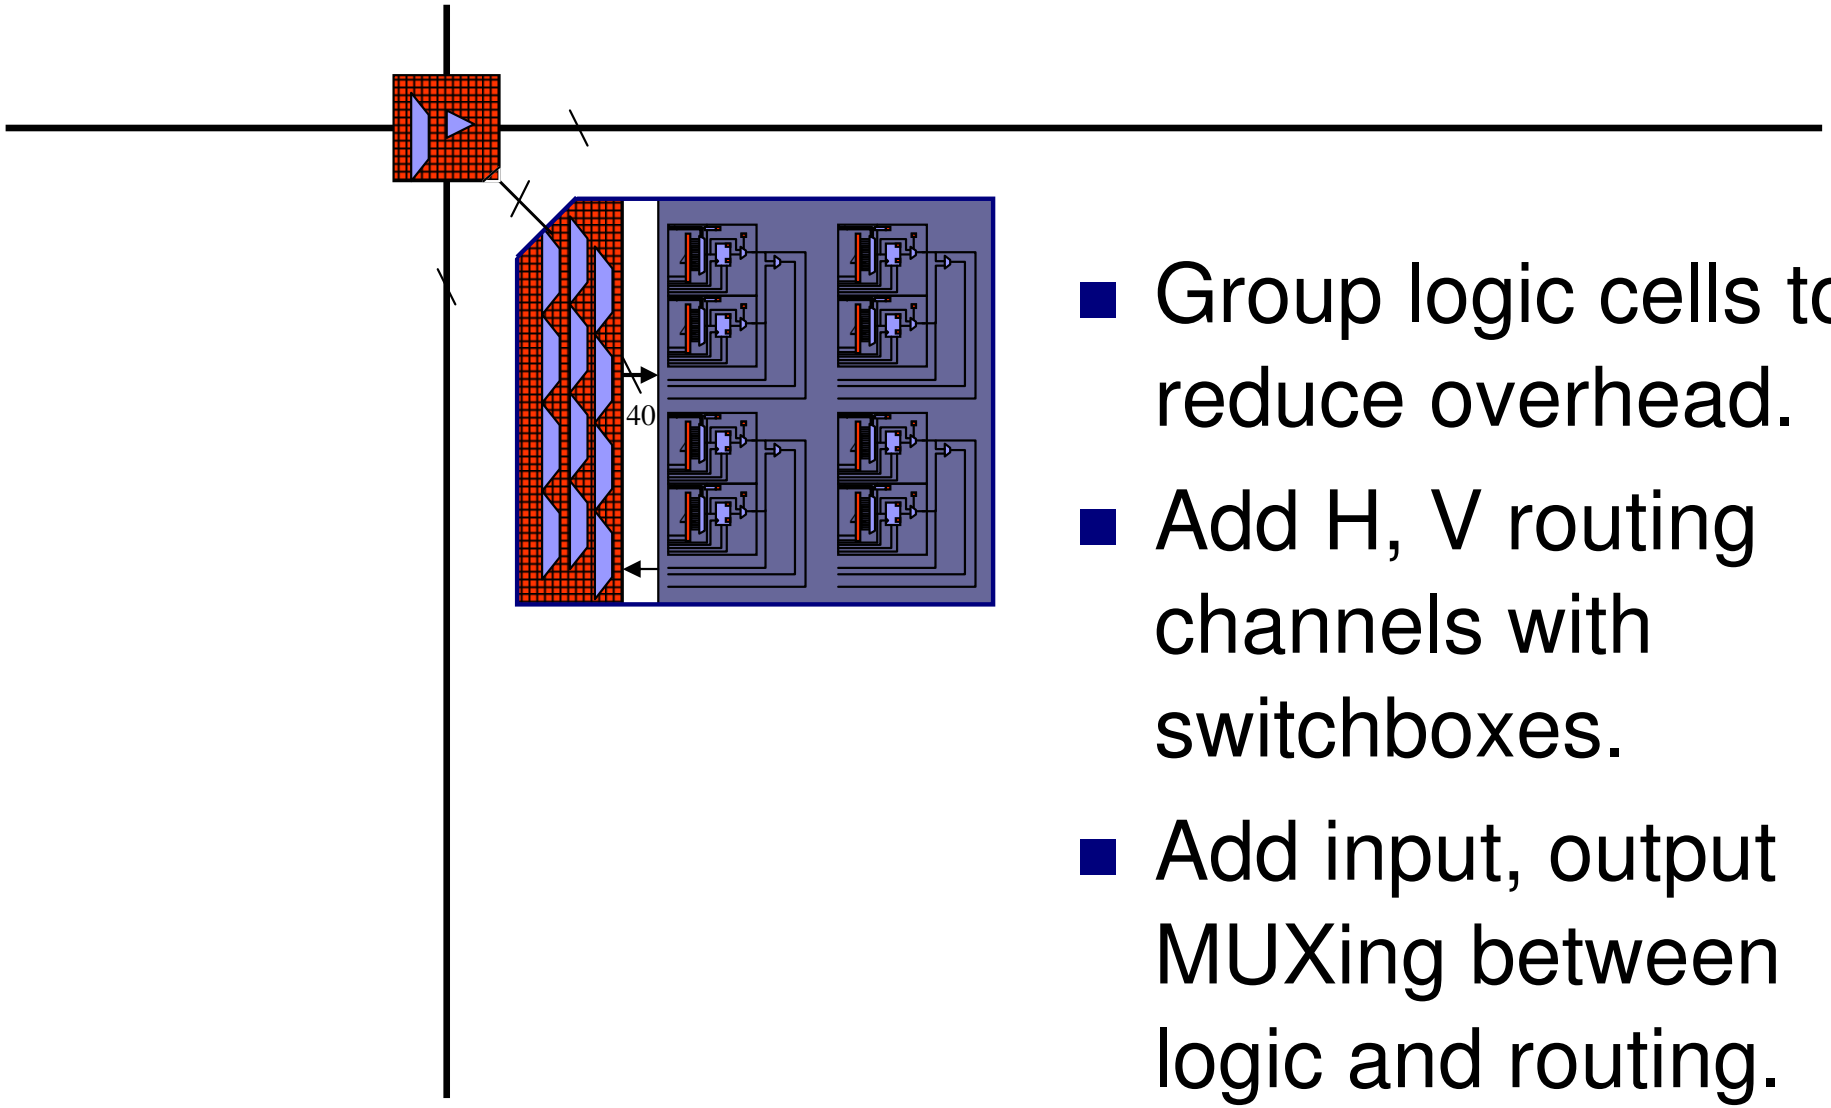

- Group logic cells to reduce overhead.
- Add H, V routing channels with switchboxes.
- Add input, output MUXing between logic and routing.

Build an Array

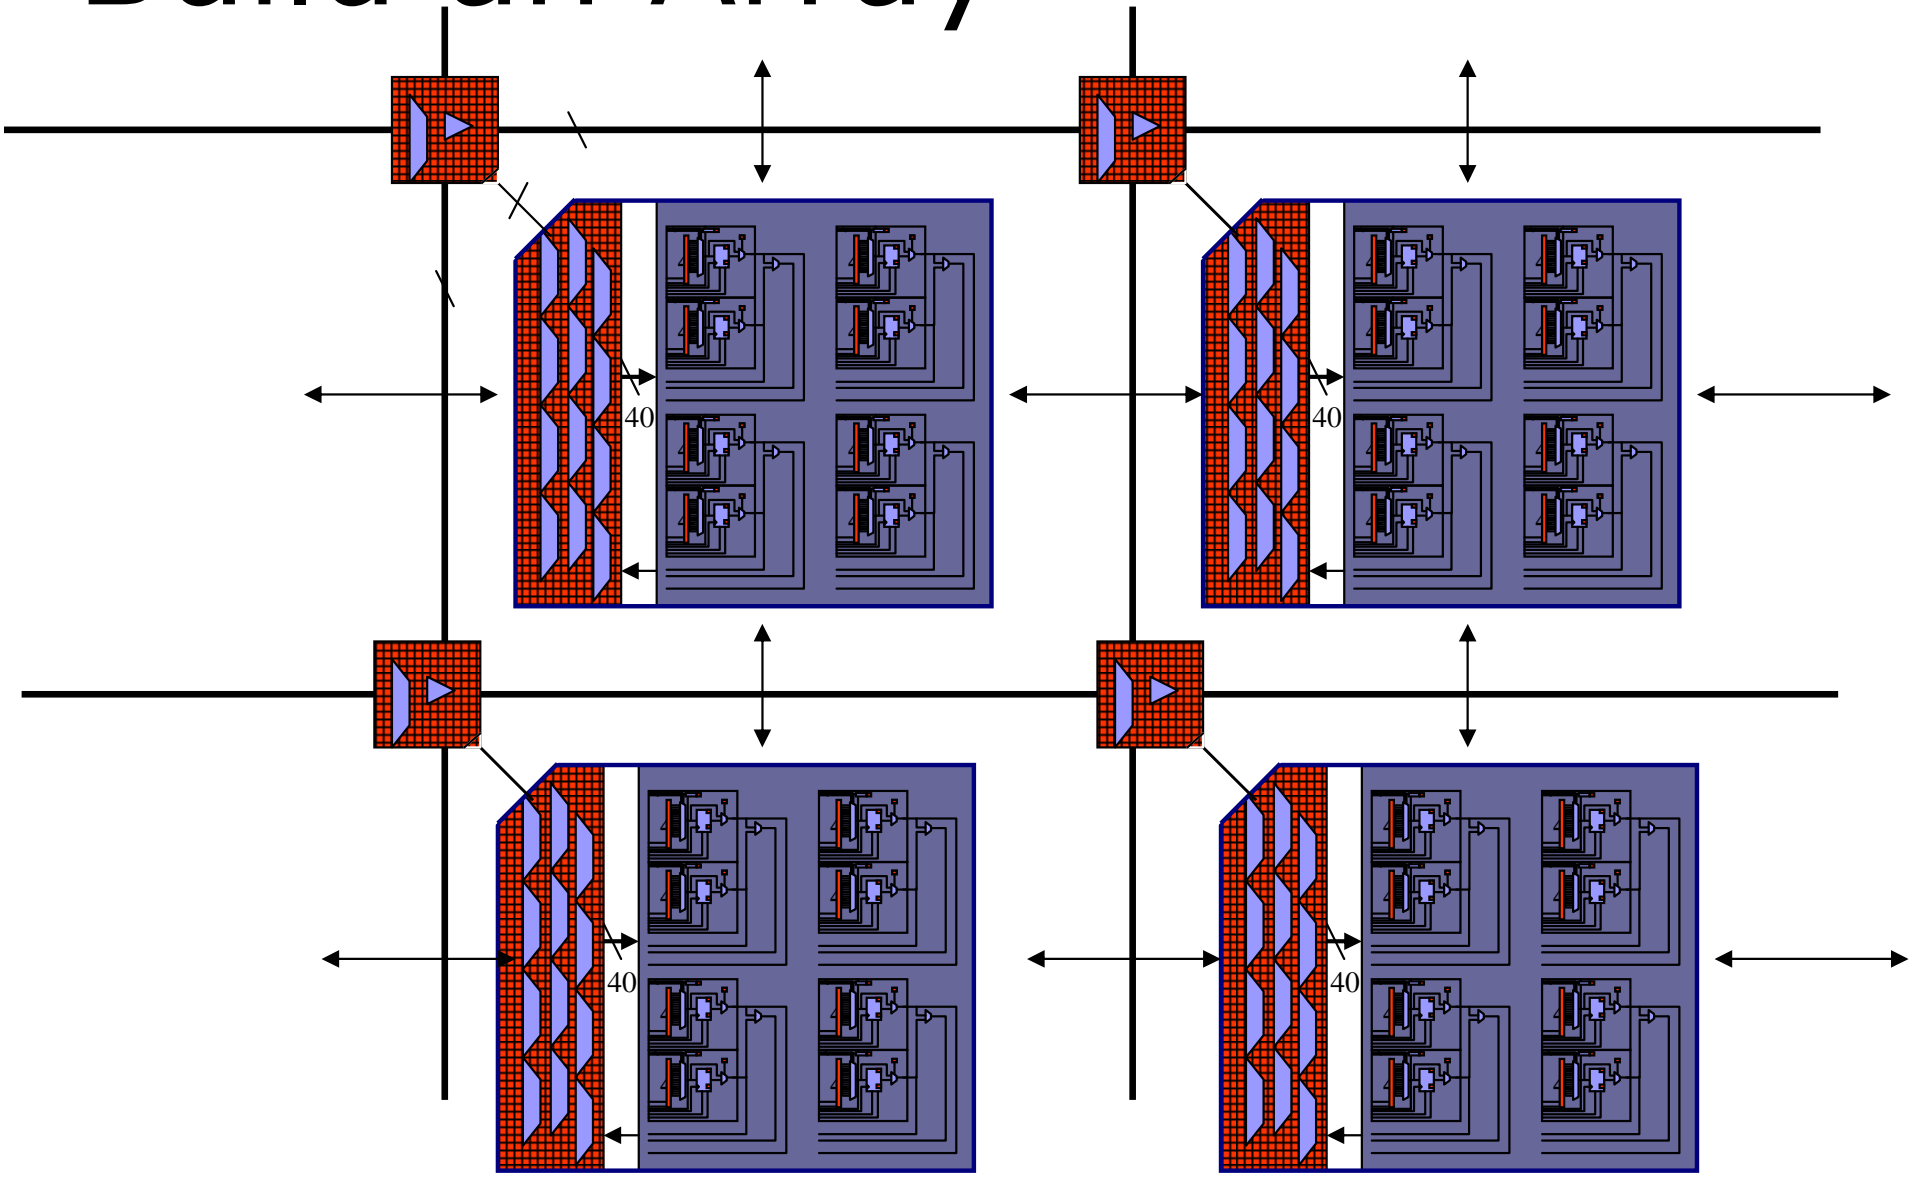

FPGA Technology

- Virtex-4 has already over a billion transistors!
- FPGA can take most advantage of Moore's law
  - More an more logic available on the same die
- CPUs in programmable logic
  - Hardcore = PowerPC
  - Softcore = MicroBlaze

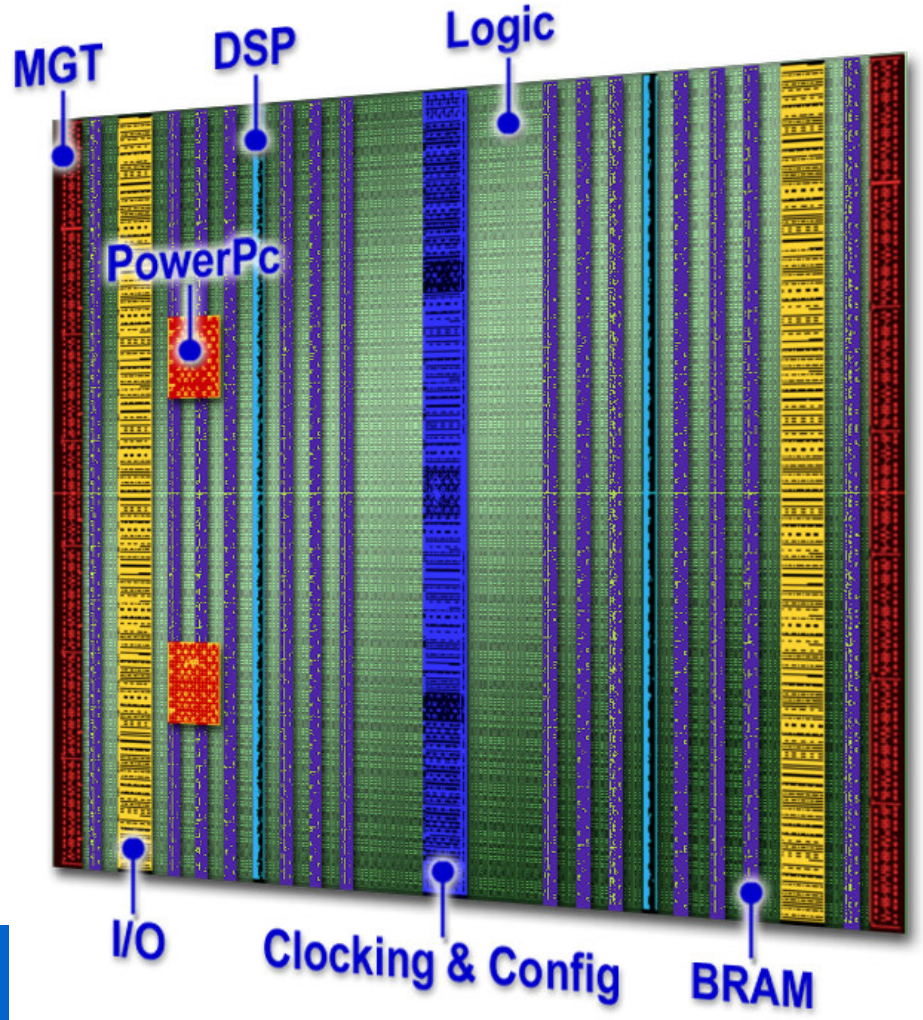

XILINX

ALTERA

FPGA, the last 10 years

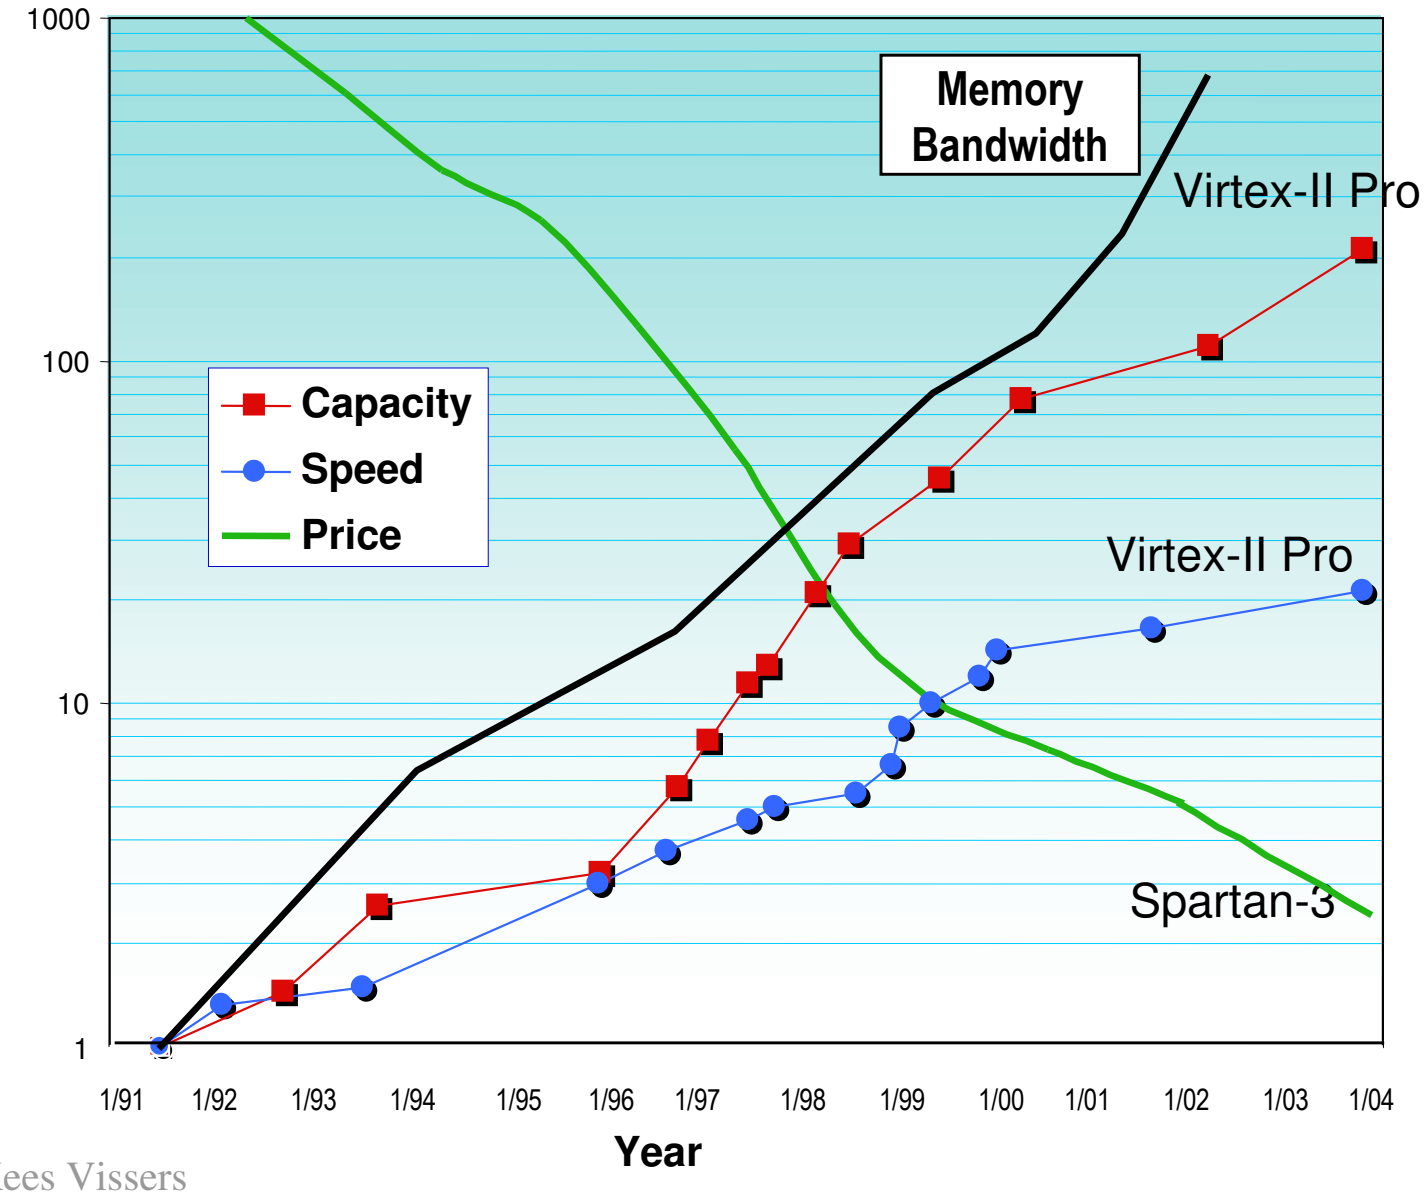

FPGA Programming

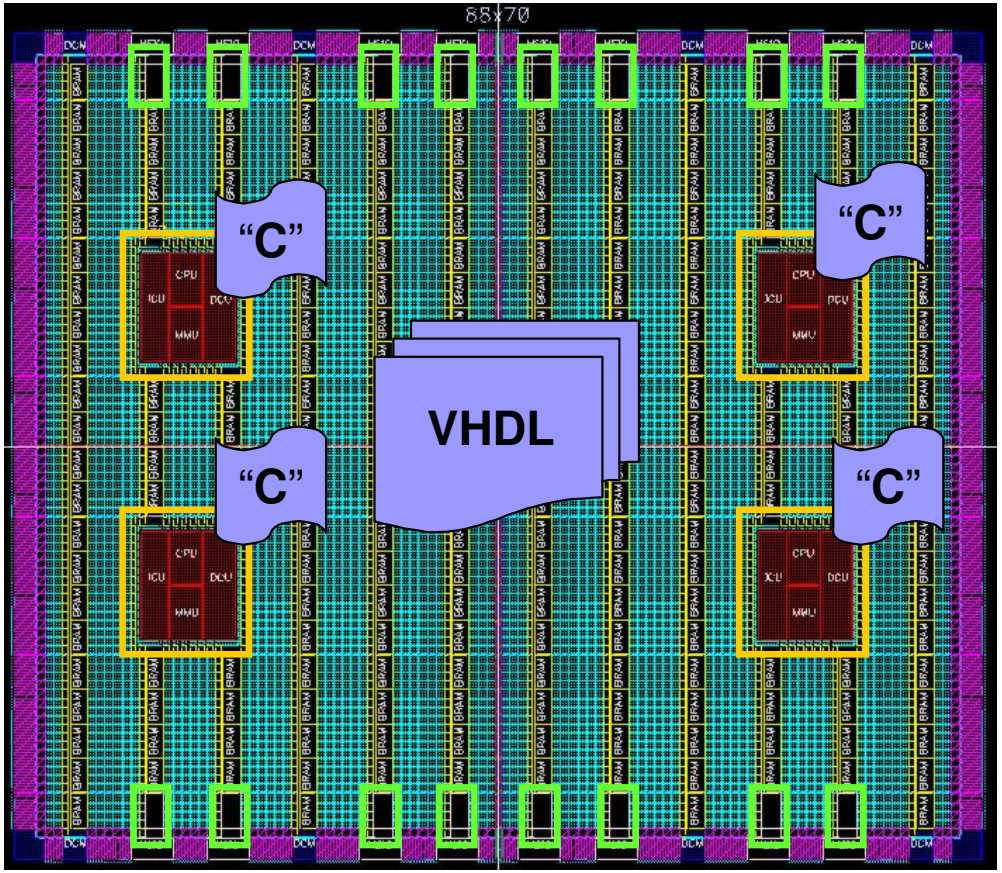

Virtex-II Pro FPGA

- FPGA is flexible in Hardware and Software…
  - How to teach a billion transistor chip a new trick
  - Solution Industry doesn't work
- Programming means:
  - Write "C" programs for the Micro processors
  - Write "VHDL" programs for the IP cores and interconnection network
- "C" programs and "VHDL" programs need to work together
  - High performance
  - Error proof
  - Nightmare to debug
- Programming is currently done manually with lots of engineers

Software Productivity

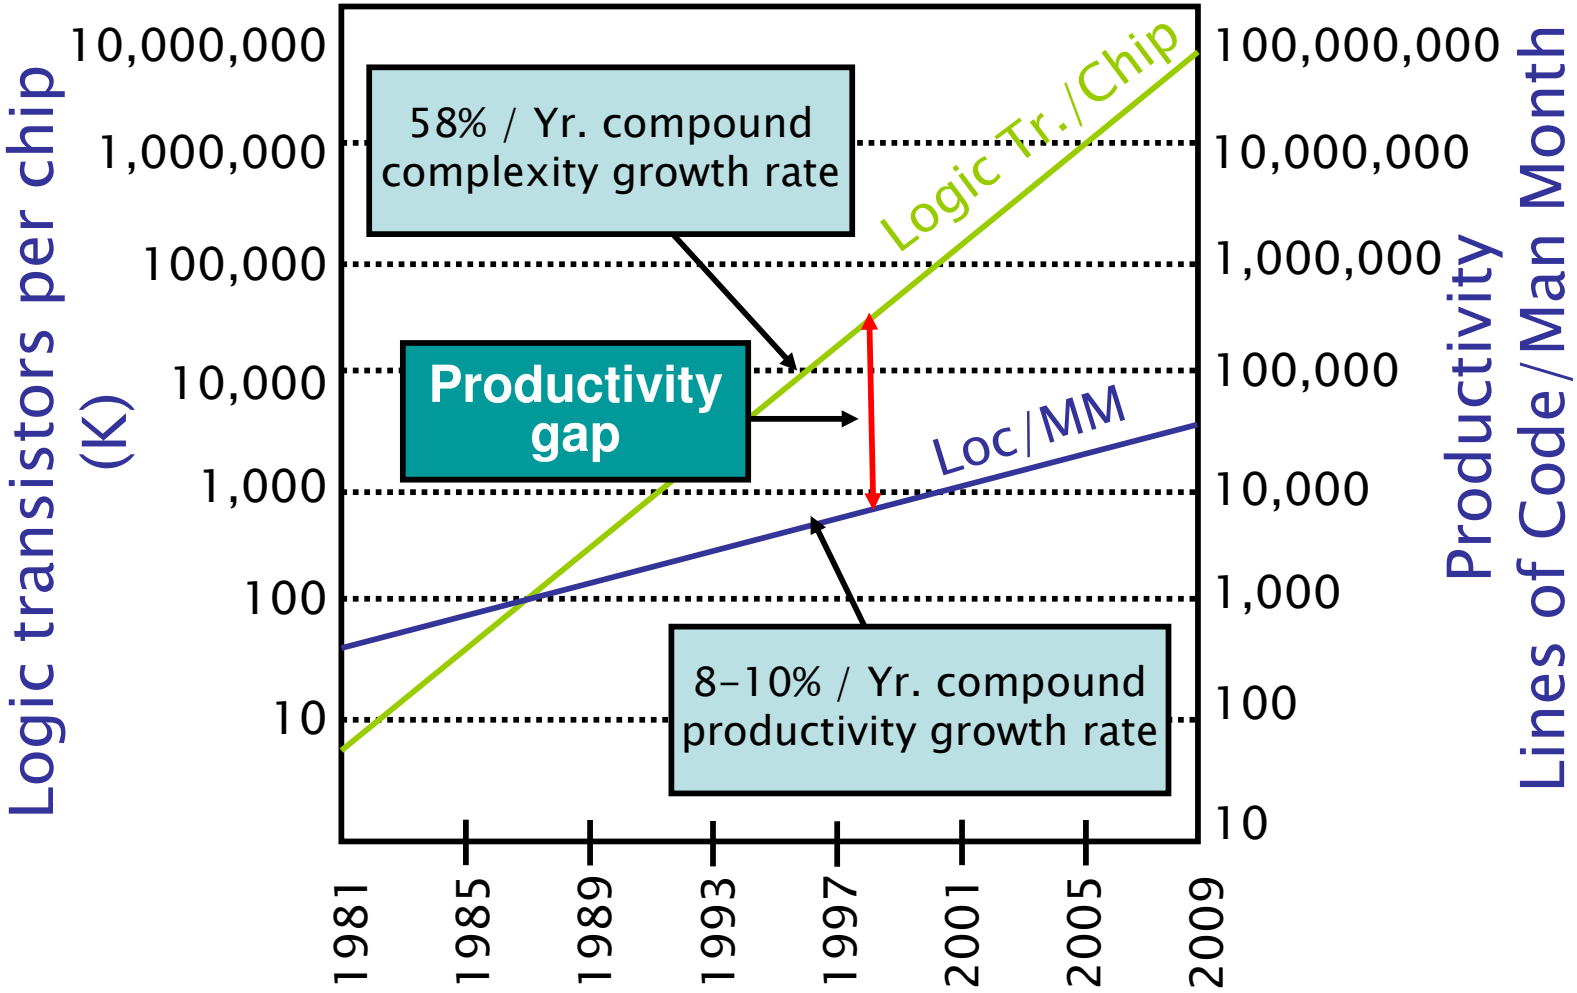

We know how to build Billion transistor chips, but do not know how to program them efficiently

Research at LIACS (COMPAAN)

Matlab/C/Java

```
for j = 1:1:N,
    [x(j)] = S1( );
end
for i = 1:1:K,
    [y(i)] = S2( );
end
for j = 1:1:N,
  for i = 1:1:K,
     [y(i), x(j)] = func(y(i), x(j) );
  end
end
for i = 1:1:K,
  [Out(i)] = Sink( y( I ) );
end
```

Parameterized Nested Loop Programs

Compaan

Kahn Process Network

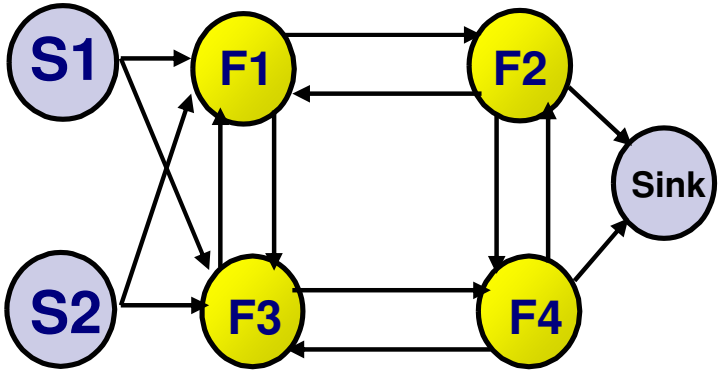

FPGA

Laura / Espam

How ICT Research at LIACS tackles the Productivity Gap (Leiden University, STW PROGRESS)

Toolbox

- Toolbox to change the number of processes
  - Unrolling (More Parallelism)
  - Merging (Less Parallelism)
- Better use of Resources
  - Re-timing
  - C-Slowing
- Exploration
- Track Technology changes

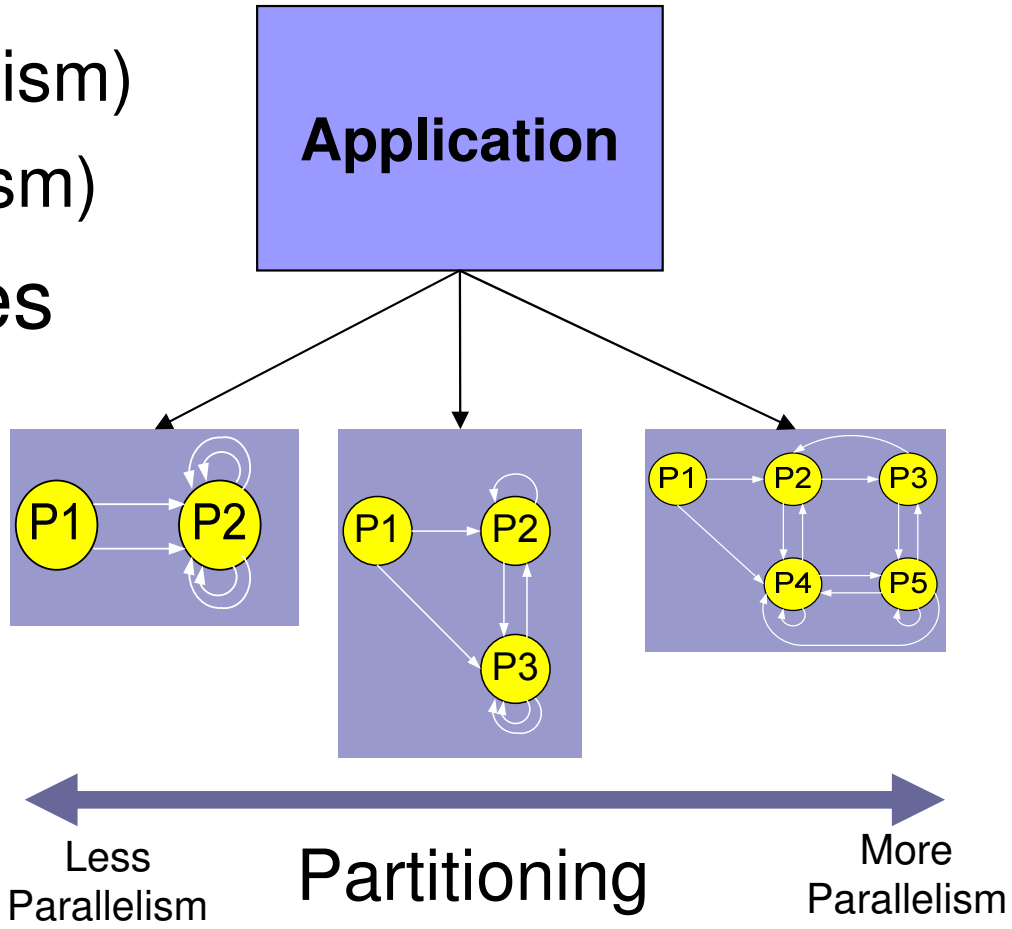

Results: Motion JPEG (1)

```
for k = 1:1:NumFrames,
  for j = 1:1:VNumBlocks,
    for i = 1:1:HNumBlocks,
      [ Block ] = VideoInMain();
      [ Block ] = DCT( Block );
      [ Block ] = Q( Block );
      [ Packets ] = VLE( Block );
      [  ] = VideoOut( Packets );
    end
  end
end
```

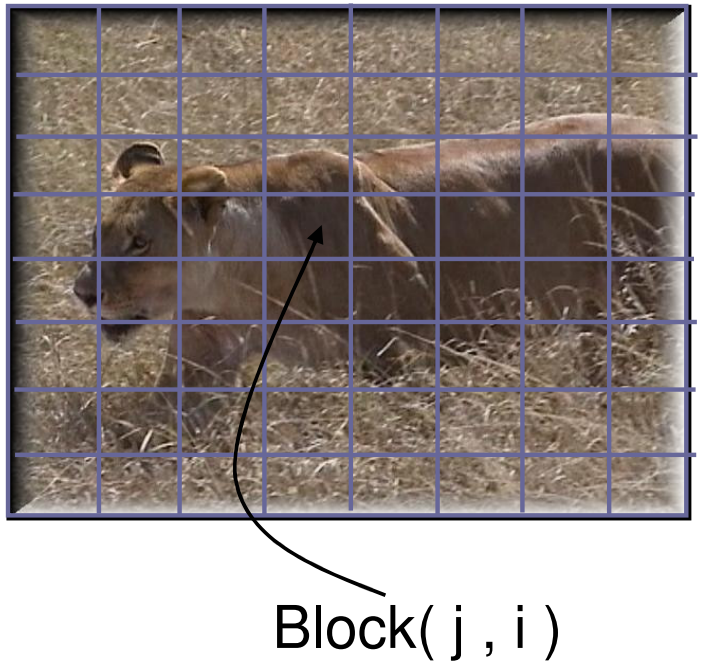

PHILIPS

PROGRESS

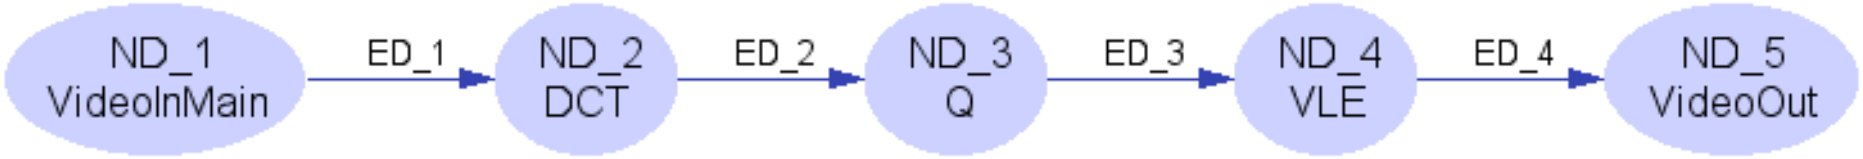

Generated Process Network

Motion JPEG (2)

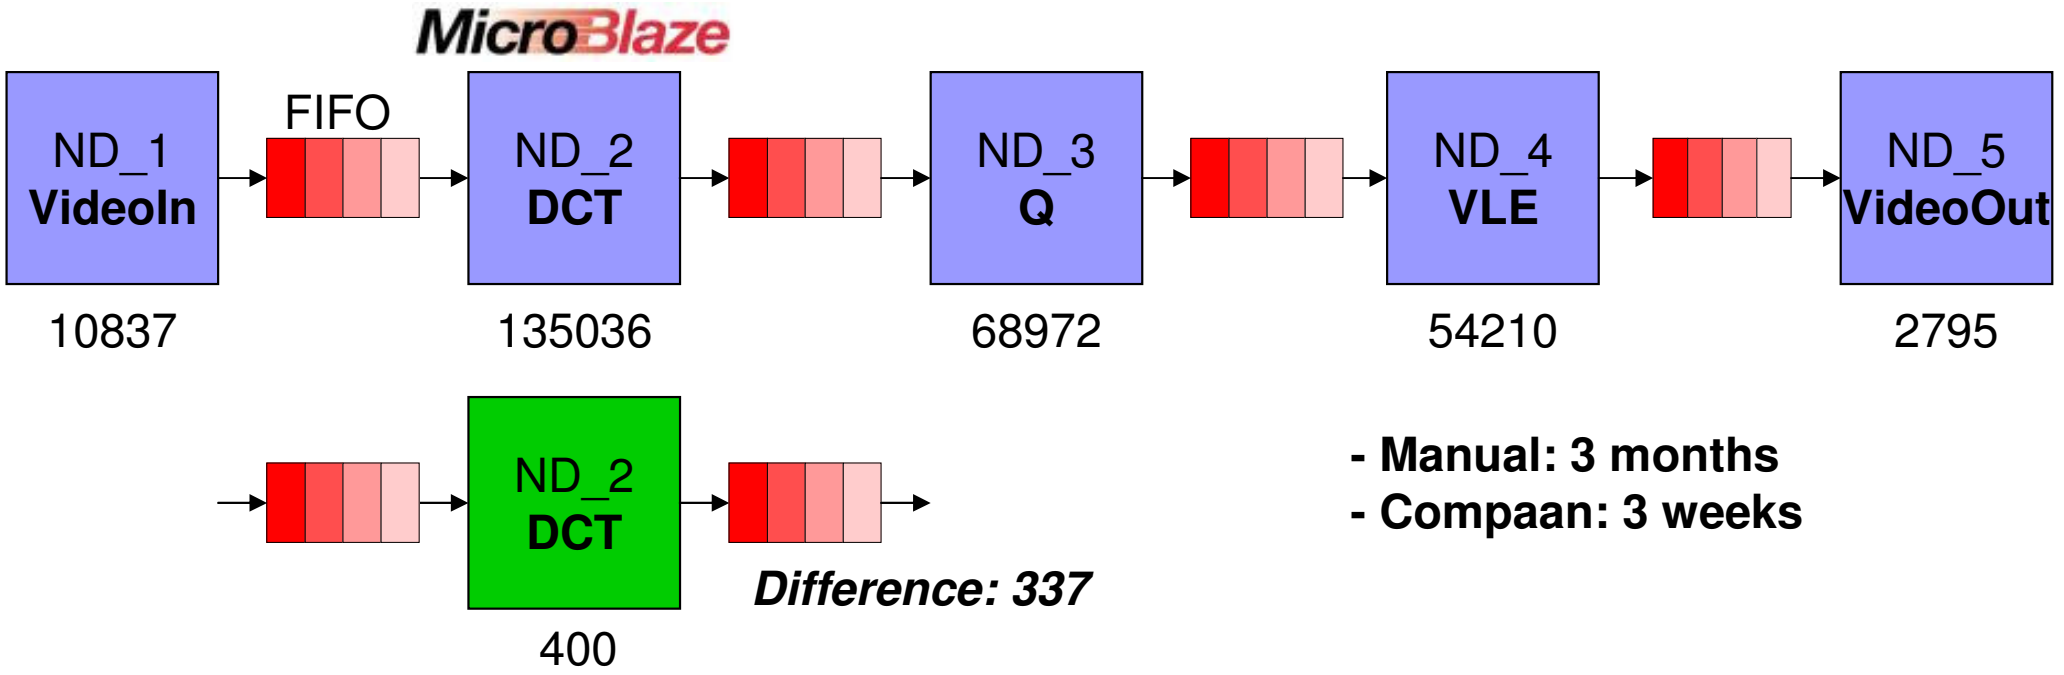

- Manual: 3 months
- Compaan: 3 weeks

- From Matlab to FPGA implementation
  - Automatically derived a multi processor solution
  - Free choice between Microprocessor (MicroBlaze) or Hardware (IPCore)
  - Automatically generated
    - "C" program for all Microprocessors
    - VHDL program for interconnections and IP Core integration
    - Correct by Construction
    - Seamless integration into FPGA Development environment (EDK)

Results: QR Case study

```
for k = 1:1:21,
  for j = 1:1:7,
    [r(j,j), rr(j,j), a, b, d(k)] =
            bcell( r(j,j), rr(j,j),x(k,j), d(k));
    for i = j+1:1:7,
      [r(j,i), x(k,i)] = icell( r(j,i), x(k,i), a, b );
    end
  end
end
```

QR decomposition
Beam former/Radar

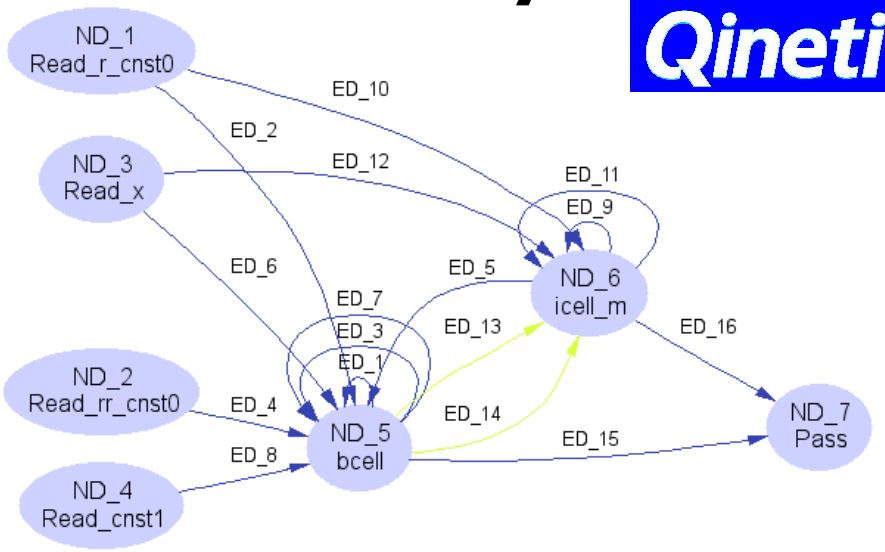

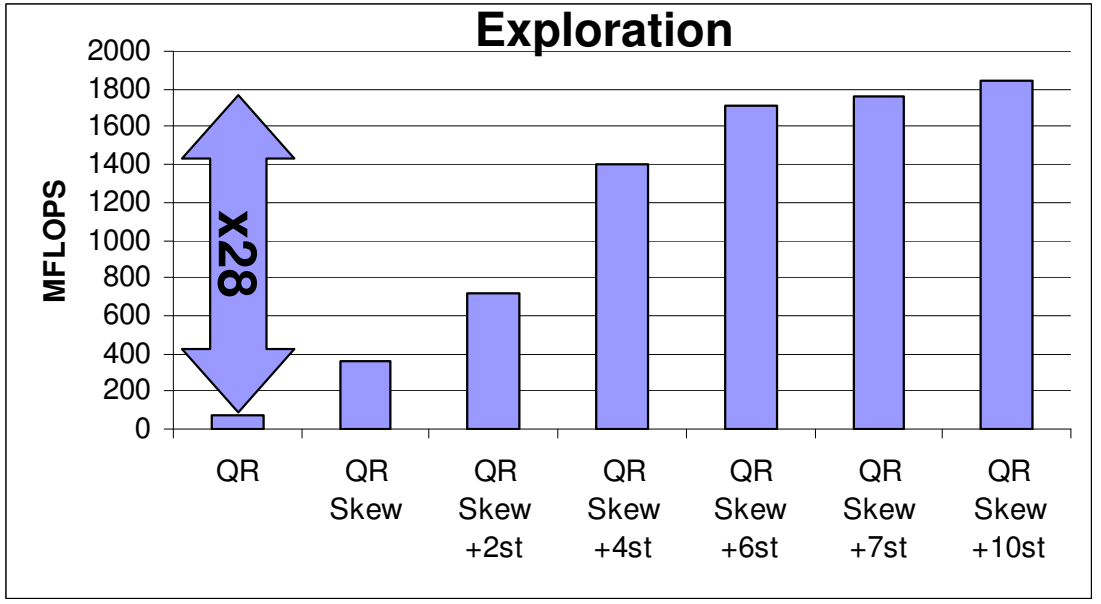

- 60MFlops out-of-the-box
- 2 GFlops, with Compaan/Laura

- Manual: 1 year
- Compaan: 3 days

Multidisciplinary Work

Mathematical Modeling (Polyhedral Mathematics)
Efficient Solvers (Parametric Integer Programming)

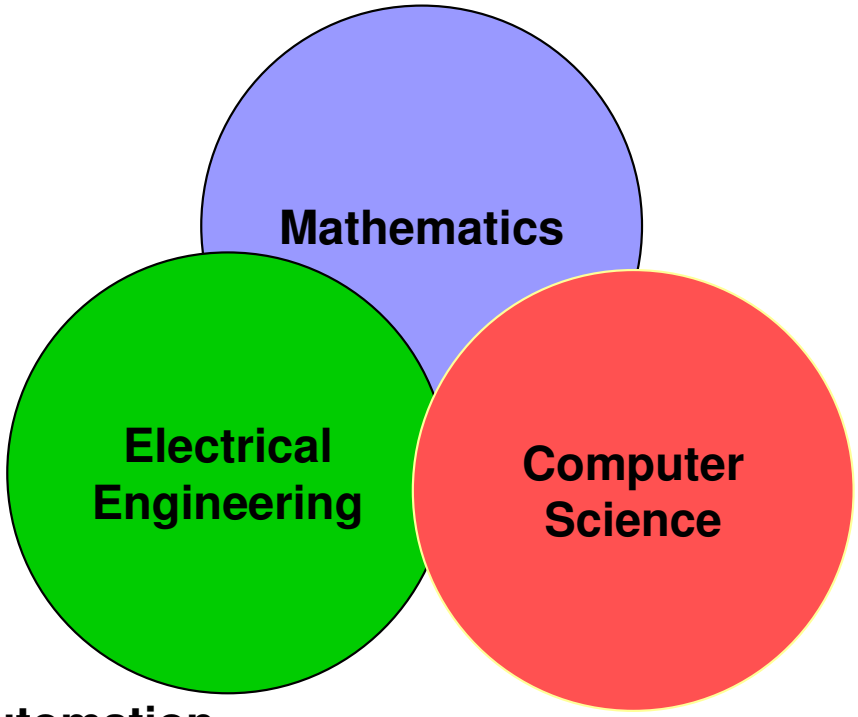

Processor Design
FPGA Technology
Electronics Design Automation

Models of Computations
Compiler Technology
Software Engineering

Compaan Project: 6 years, 6 PhDs, 2 Postdocs

Conclusions

- Stream Based Applications require Heterogeneous Multiprocessor architecture
  - Stream of 10 – 1000 million samples per second
  - 10 – 100 Giga operations per second
- Latest FPGAs are Heterogeneous Multiprocessor architectures
  - Over a Billion Transistor and still counting
- Problem is that we know how to build FPGAs, but programming them is very hard
- Programming/Compiler Technology is a strategic ICT component to reduce Productivity Gap
- Showed the Compaan Project at LIACS, which provides a solution for stream based applications.

Thanks to

- Prof. Ed Deprettere
- Dr. Todor Stefanov
- Dr. Sven Verdoolaege
- Alexandru Turjan
- Claudiu Zissulescu
- Hristo Nikolov
- Sjoerd Meijer
- Jerome Lemaitre
- Bin Jiang
- Wei Zhong

For more information http://www.liacs.nl/~kienhuis