(Re-)Configurable Solutions for the high-volume ASSP Market

Fabio Campi,
STMicroelectronics
Technology R&D,
Configurable Logics

Reconfigurable Architectures Workshop

Keynote presentations 2003-2008 (I)

- 2003
  - R.Hartenstein, Univ Kaiserslautern, "Are we really ready for a breakthrough?"
  - M.Vorbach, PACT XPP: "Reconfigurable Processor for Mobile phones"
- 2004
  - P.Lysaght, Xilinx: "Of Wires and Gates"
  - D.Bouldin, UTennessee: "Impacting Education using FPGAs"
- 2005
  - I.Bolsen, Xilinx: "FPGAs the heart of embedded systems"
  - M.Hutton, Altera: "Old and new challenges in FPGA design"

Reconfigurable Architectures Workshop

Keynote presentations 2003-2008 (II)

- 2006
  - M.Gokhale, Los Alamos Labs: "The Outer Limits: reconfigurable computing in space and orbit"
  - R.Hartenstein, Univ Kaiserslautern: "New horizons of high performance"
- 2007
  - G.Bremer, Xilinx: "Problem oriented Configurable and reconfigurable architectures"
  - S.Singh, Microsoft: "Programming models for reconfigurable systems"
- 2008
  - D.Sciuto, Polimi: "Design Methodology for partial dynamic reconfiguration"

Reconfigurable Architectures Workshop

Keynote presentations 2003-2008:

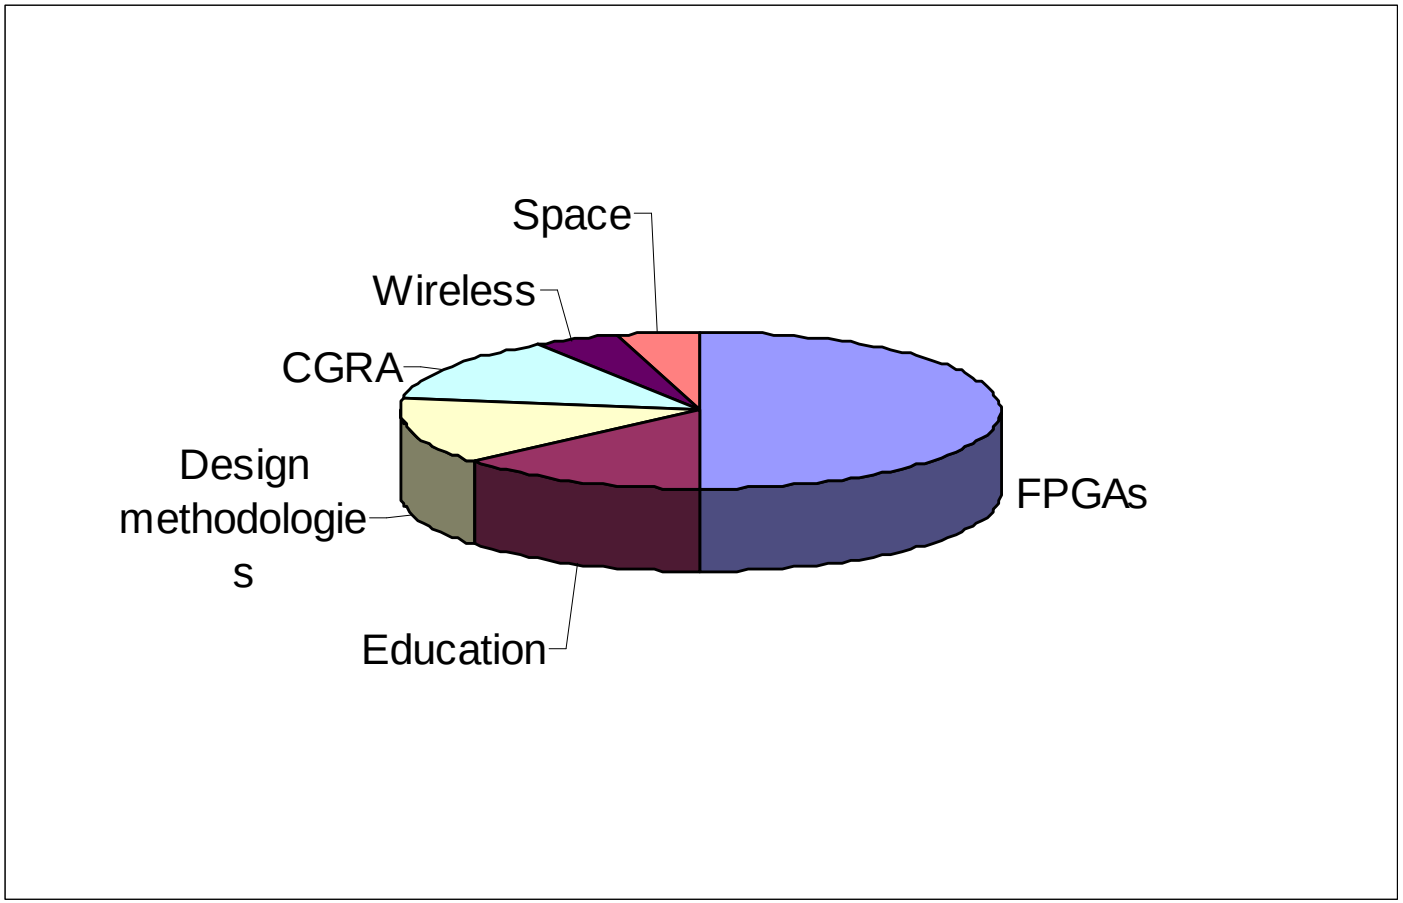

Reconfigurable Architectures Workshop

Keynote presentations 2003-2008:

- FPGAs: Impressive success story, clearly established as state-of-art signal processing platform
- CGRAs: Intriguing story, valuable alternative to FPGAs, it remains to be investigated the real impact on the silicon market
- Embedded reconfigurable architectures. Can this be a factor in ASSP Design ??? We thought so……

Application Fields: Application Specific Standard Products

- High end, embedded systems for signal processing:
  - Wireless Mobile
    - Cellular Infrastructure
    - Connectivity
    - Multimedia processing
  - Home Entertainment
    - HD Image processing
  - Computer Peripherals
    - Data Storage
    - Printers

High end Signal processing Solutions

"Up there where processors do not dare.."

- FPGAs: Very general purpose, good performances for low volumes, some issues in programming productivity
- CGRAs: Recore Montium, Pact XPP, many others. Loosely Domain oriented, still to be established as solid alternatives
- Multi-Core Architectures: PicoChip, Tilera, Cell.

- ASSP: Very domain oriented, high performance high volumes. Exploit specific hardware acceleration.User friendliness is not so relevant, post-fab programmability is achieved lowering the grain of asic acceleration

ASSP Trends (1)

Wayne Wolf, DAC 2004 :

- "MP-SoC make the most sense in high-volume markets that have strict performance, power and cost goals. Communications, multi-media and networking are examples of markets that meet these requirements"

ASSP Trends (2)

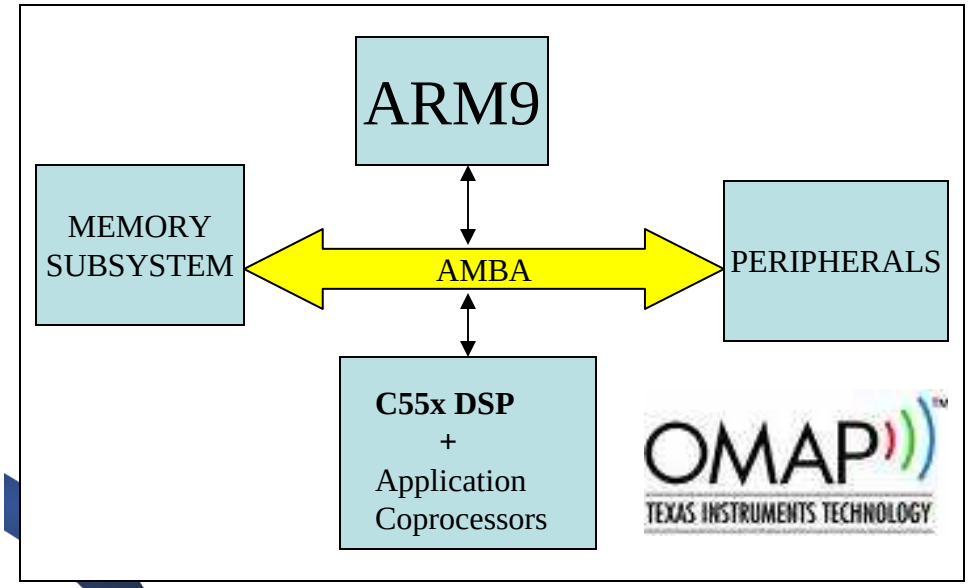

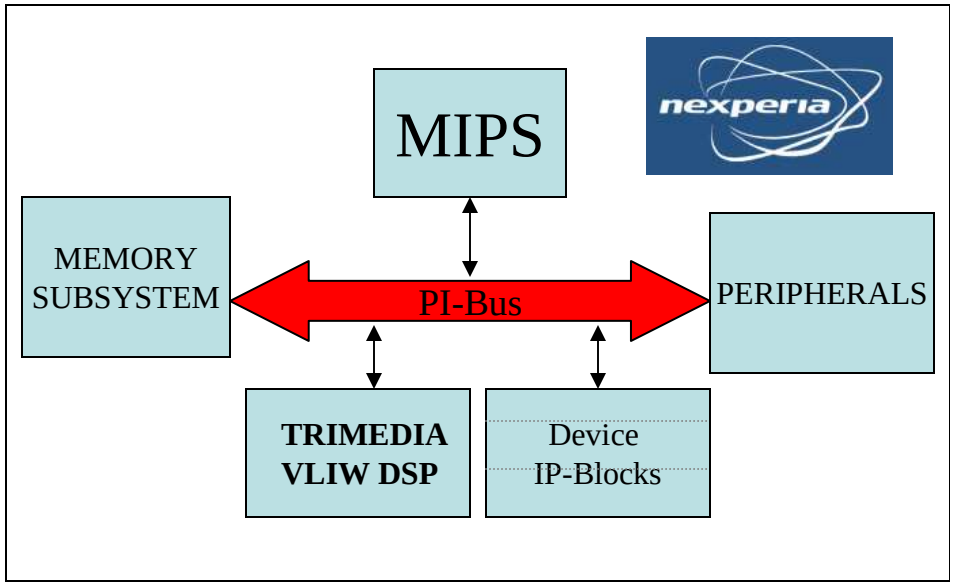

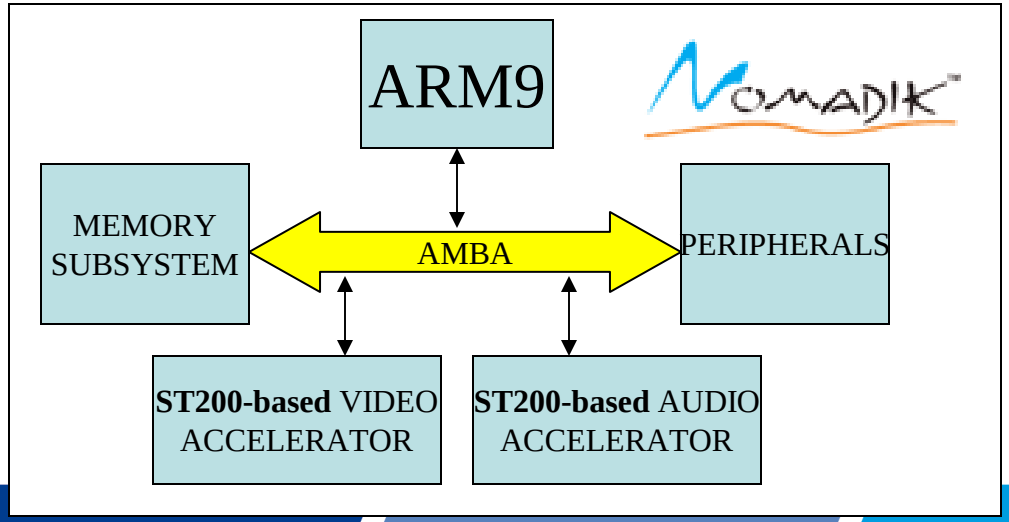

ASSP Trends: Task specific (Configurable) Processor

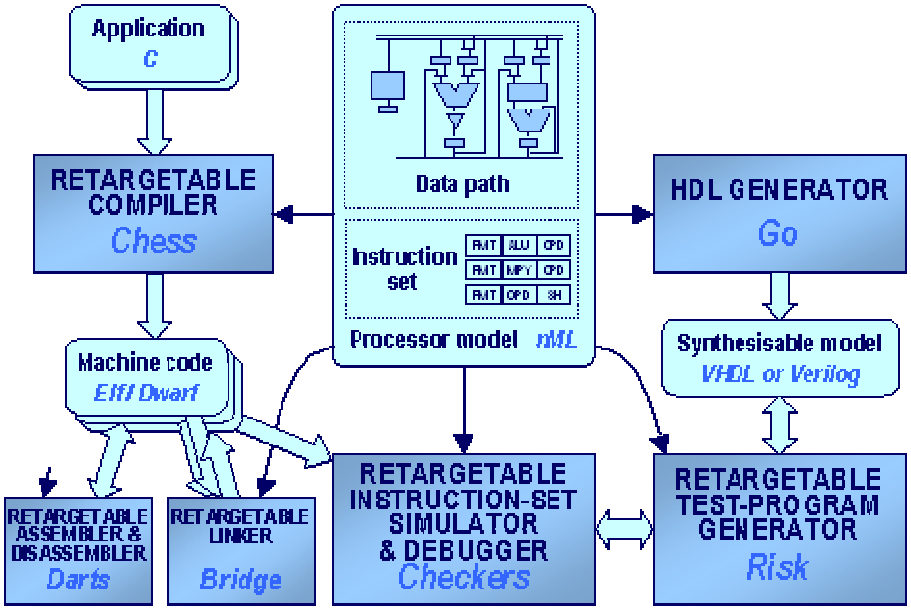

- RWTH Aachen => Lisatek (Coware)
- IMEC => Target Compiler (TGT)
- ARM => OptimoDE
- Philips => Silicon Hive
- Tensilica => XTensa, XPress

ASSP Trends: Case Study (1)

- Cell Phones: Key market driver

	1990	2000	2005	2010
MUnits	<100	400	700	900
Features	Voice	Voice & Data	Multimedia	Global Convergence

Source: A.Artieri, ST, DAC2006

ASSP Trends: Case study (2)

	2004	2006	2008	2010	2012
Technology Node	90	65	45	32	22
Loosely coupled subsystems	2	4	6	8	12
General Purpose CPU	Single ----------→ Multiple				
Hardware Accelerator	Hardwired ---------→ Reconfigurable				

- Constant die Size
- Slow evolution of peripherals (area decrease)
- CPU subsystem complexity doubles each node
- Embedded memory capacity doubles each node
- Loosely coupled DSP subsystem complexity increase 30% each node (30% area decrease)

Source A.Artieri, ST, DAC2006

Run Time Reconfigurable Architectures

- Can run time reconfigurable architectures [eFPGAs and/or CGRAs] be used as embedded computation cores in a MPSoC ASSP?

ASSP and Run-time reconfigurable hardware (1)

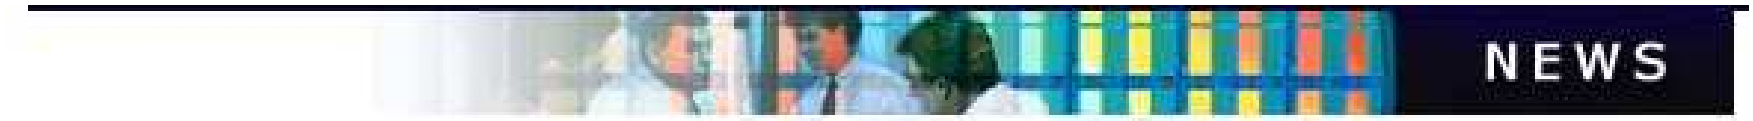

STMicroelectronics Introduces Highly Integrated Microcontroller with Embedded Programmable Logic for Wireless Infrastructure Applications

GreenFIELD-STW21000 significantly reduces external component count therefore reducing total cost and power consumption

San Francisco, Ca., Embedded Systems Conference, March 9, 2005 – STMicroelectronics (NYSE: STM), a leading developer of system-on-chip (SOC) solutions for the wireless market, has today disclosed details of a multi-purpose microcontroller that targets applications in wireless infrastructure equipment. Under the development name 'GreenFIELD' and part number, STW21000, the new chip combines an ARM926EJ-S 330 MIPS RISC processor core, 16-Mbit of on-chip eDRAM memory, an eFPGA (embedded Field-Programmable Gate Array) block, and a wide range of analog and digital peripherals. The GreenFIELD-STW21000 is the second advanced System-on-Chip (SoC) product to be announced from ST's wireless infrastructure division.

ASSP and Run-time reconfigurable hardware (2)

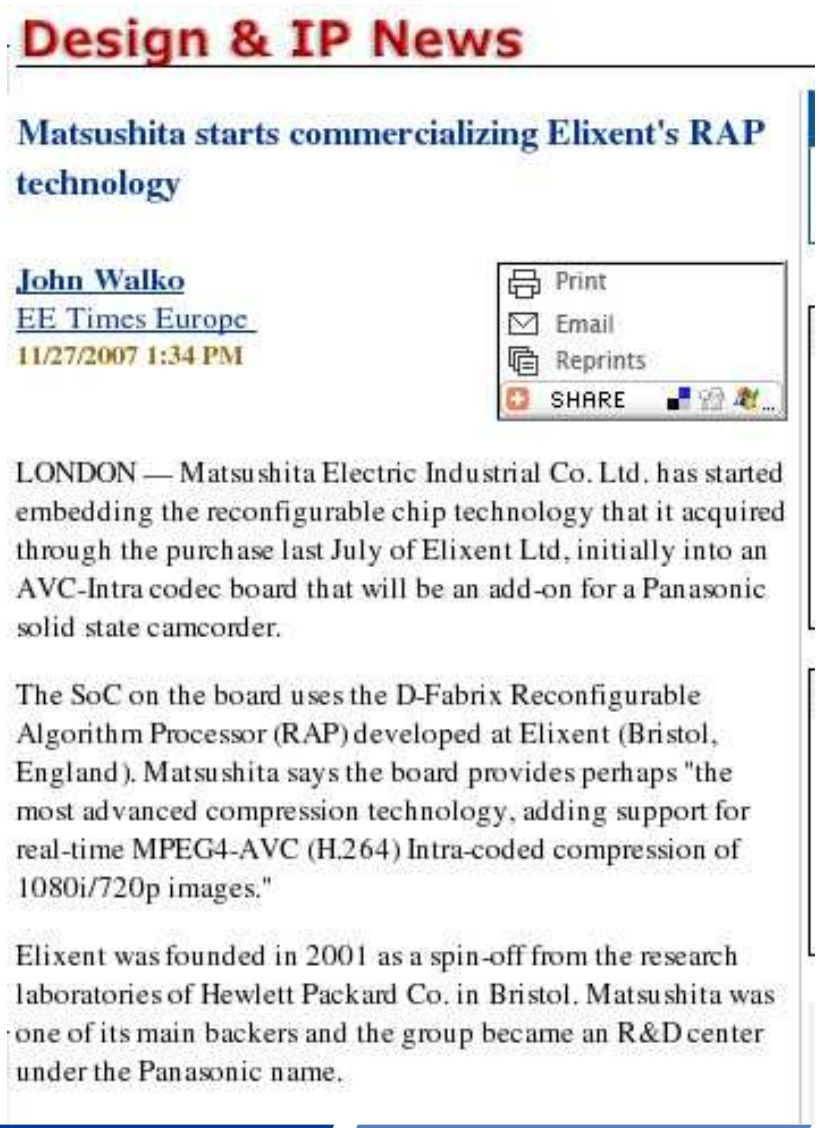

Design & IP News

Matsushita starts commercializing Elixent's RAP technology

John Walko
EE Times Europe
11/27/2007 1:34 PM

Print
Email
Reprints
SHARE

LONDON — Matsushita Electric Industrial Co. Ltd. has started embedding the reconfigurable chip technology that it acquired through the purchase last July of Elixent Ltd, initially into an AVC-Intra codec board that will be an add-on for a Panasonic solid state camcorder.

The SoC on the board uses the D-Fabrix Reconfigurable Algorithm Processor (RAP) developed at Elixent (Bristol, England). Matsushita says the board provides perhaps "the most advanced compression technology, adding support for real-time MPEG4-AVC (H.264) Intra-coded compression of 1080i/720p images."

Elixent was founded in 2001 as a spin-off from the research laboratories of Hewlett Packard Co. in Bristol. Matsushita was one of its main backers and the group became an R&D center under the Panasonic name.

Case Study: Morpheus (1)

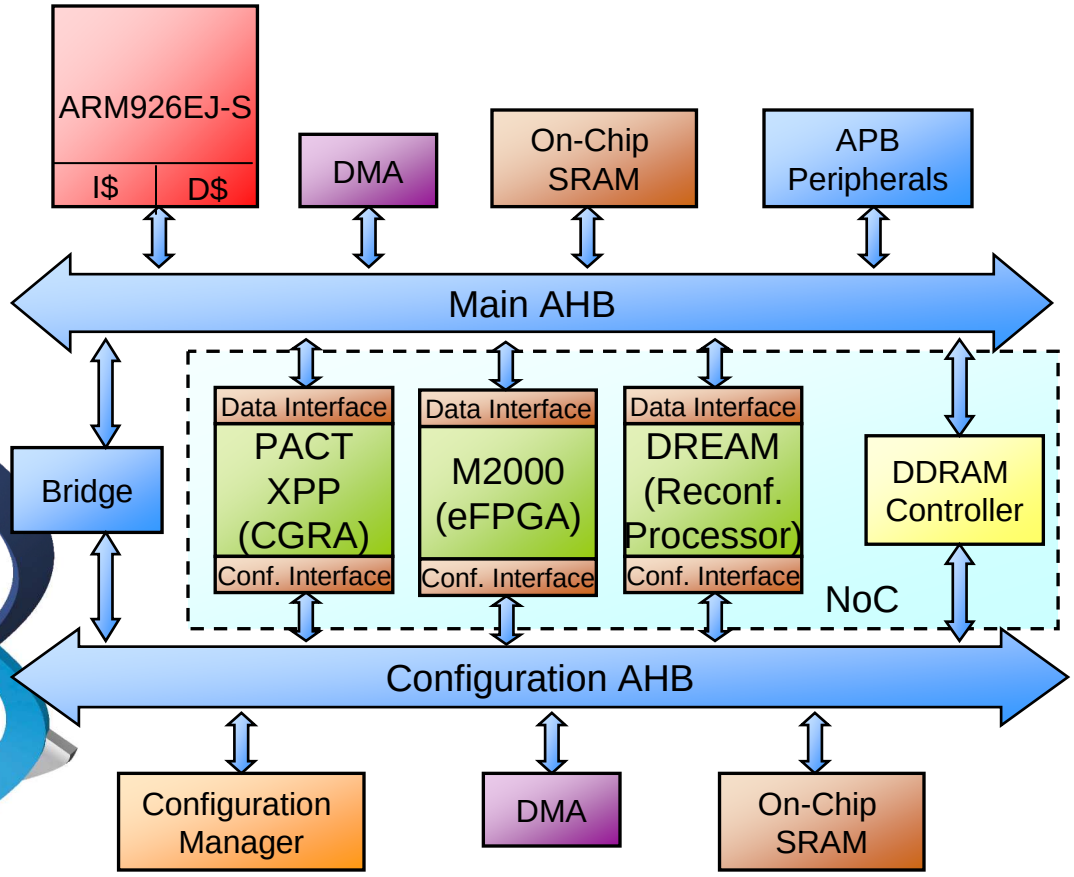

- Heterogeneous architecture for heterogeneous applications
- M2000 eFPGA for IO and bit-wise computation
- DREAM Adaptive processor for control-intensive tasks
- Pact XPP CGRA for Data-intensive computation
- NoC infrastructure for massive data communication
- Transport triggered computation
- Integrated, library based toolset providing friendly user-interface

Case Study: Morpheus (2)

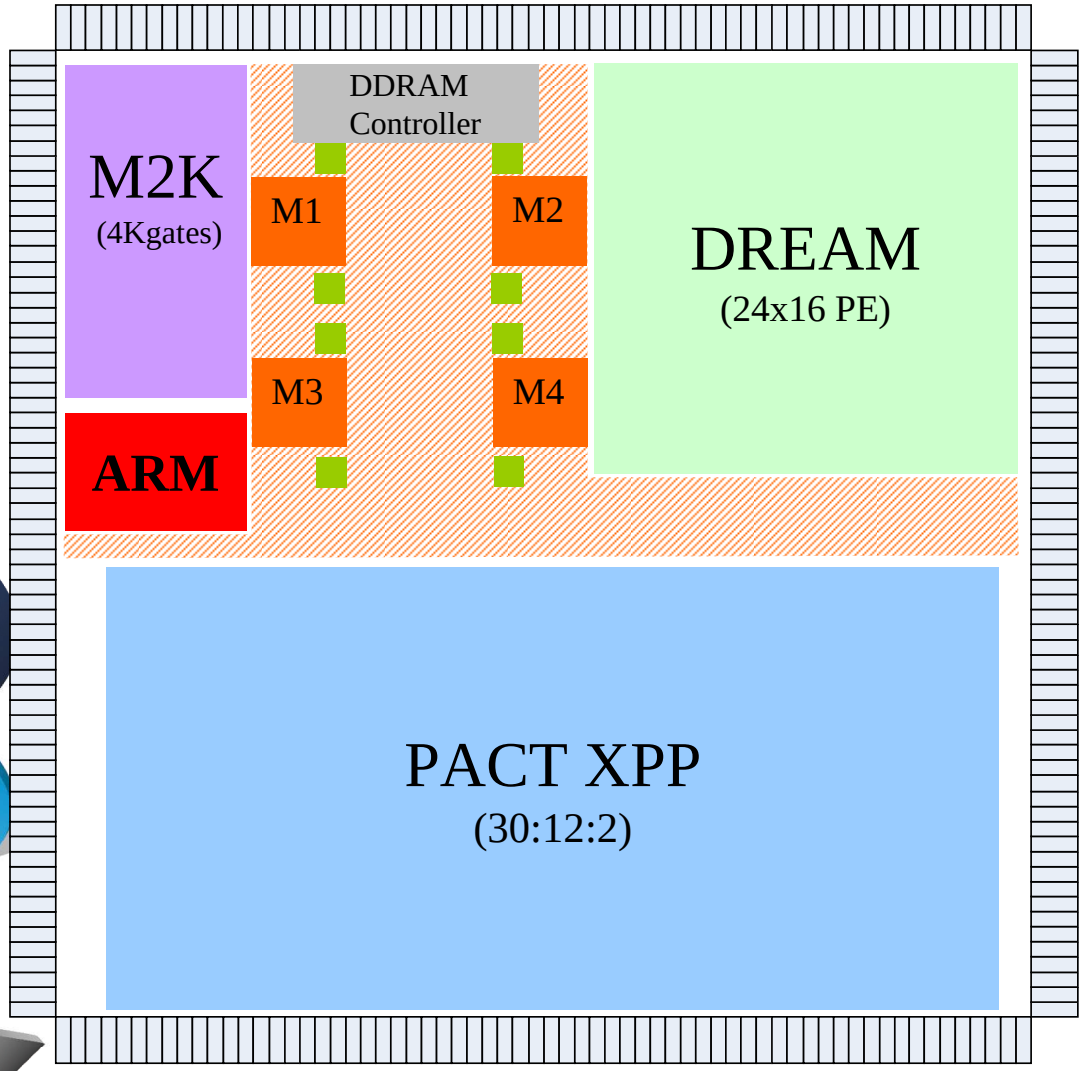

- Appealing signal processing solution for "domain-oriented computation"
- Good tradeoff between usability and performance results
- Reconfigurable cores feature area ~5 to 10 times higher than hardware-oriented accelerators, and require dynamic reconfiguration.
- This is justified by the broad application range

ASSP and Run time reconfigurable hardware

- Reconfigurable Architectures (eFPGA, CGRAs) are being heavily investigated as embedded solutions for heterogeneous MPSoC

- Drawbacks
  - Area
  - Configuration/Bitstream handling (area, communication and synchronization overhead)
- The average core area is roughly 5 to 10x wrt ASIC current solutions, and requires time multiplexing (dynamic reconfiguration)

The Reconfiguration Trade-off

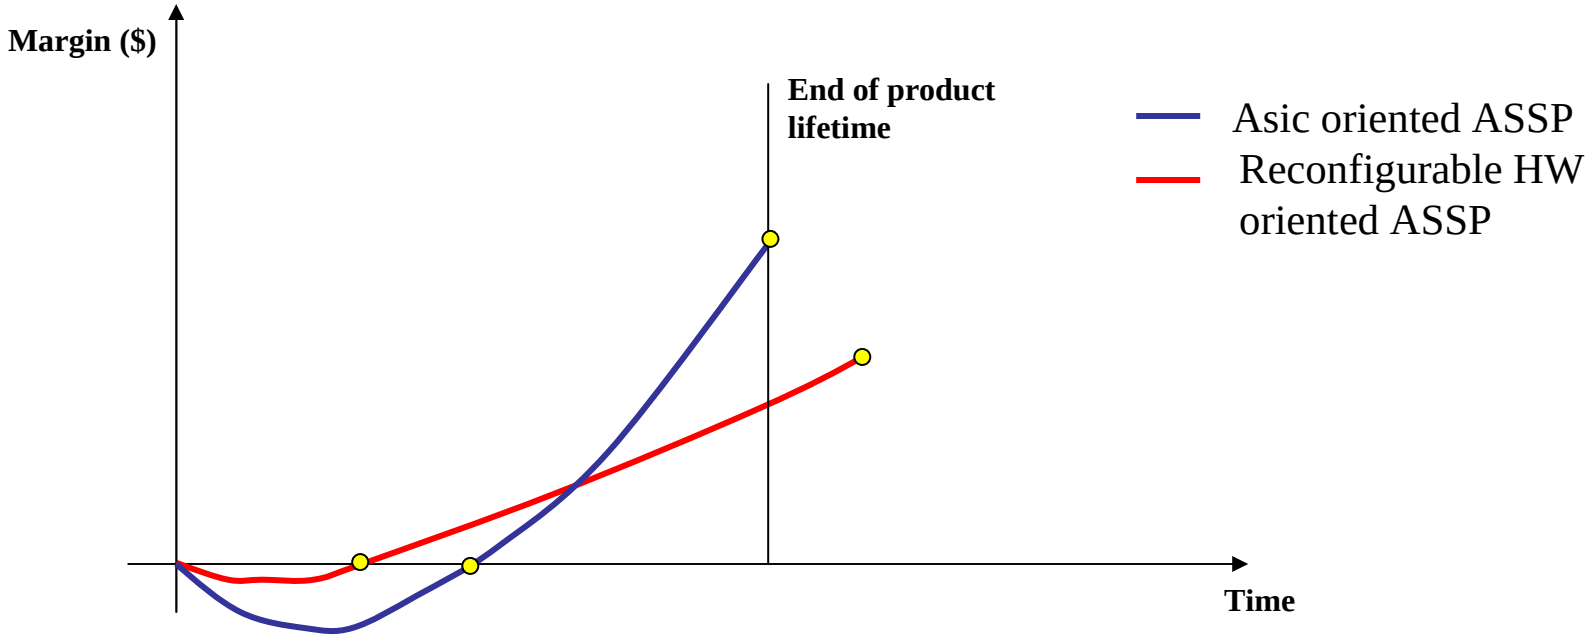

- Reconfigurable hardware *does enlarge* the volumes of the ASSP, thus reducing NREs
- Reconfigurable hardware *may* enlarge the product lifetime
- Reconfigurable hardware *heavily affects area*, reducing margins

Fast, Cheap and Under Control (1)

- The semiconductor industry is caught on two horns of the economics dilemma:
  - Deep Sub-wavelength lithography mask costs and yield issues
  - Design turnaround time and portability of design effort

- StdCells have taken us to 90 nm, we need a next generation fabric to bring us back into "fast, cheap and under control" implementation flows

Source: I.Bolsen, Xilinx; C.Hamlin LSI Logics; P.Magarshack, ST; Z.OrBach, eASIC; L.Pileggi, CMU

DAC 2003

Fast, Cheap and Under Control (2)

- The Research community has invested a lot at architecture level, there is a lot of space at circuit/implementation level
- We must add new design parameters between design time and run-time configuration. Success is measured on following metrics:
  - Fast design time
  - Low NRE Costs (Design/Mask)
  - Manufacturability
- The goal is to replicate in MPSoC ASSP the success of discrete FPGAs

Technology fabrics for Digital Signal Processing

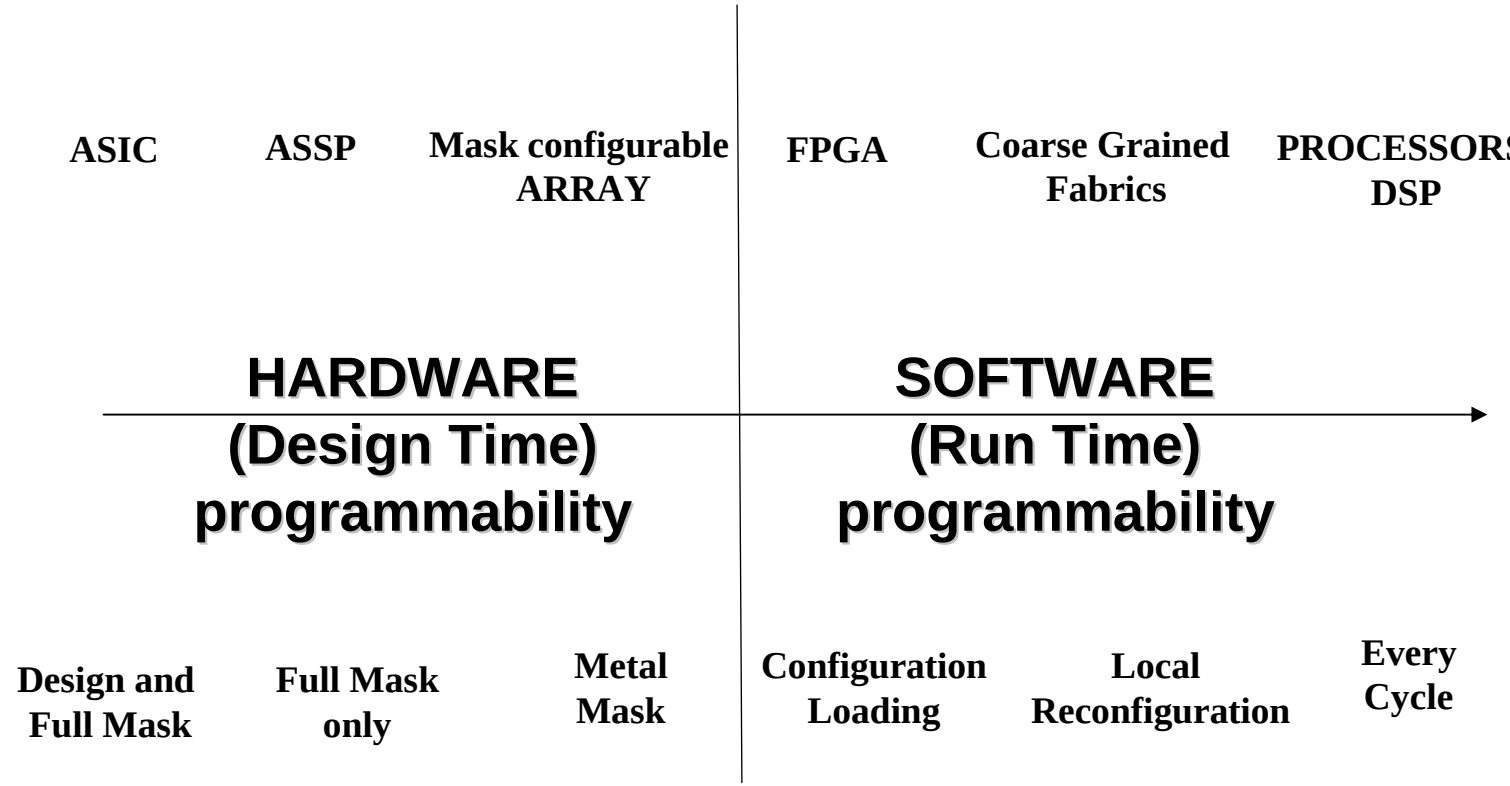

Outline

- Mask Programmable hardware: an introduction
- Technology orientations
- Exploitation of Mask Programmable hardware in the ASSP domain
- Why addressing the reconfigurable computing community?

Mask Programmable Logics (1)

We define *mask programmable logic*:

- A layout region of a device organized in a logical structure composed of elementary, uncommitted logic modules.
- These modules are designed over the majority, but not the totality of the mask layers that compose the layout
- The region can be tuned to perform a given functionality by defining the remaining reduced subset of the mask layers

Mask Programmable Logics (2)

- Depending on the solution DFM rules, clock distribution, test structures, some signal paths may be taken care of in the design fabric, *amortizing design/verification costs*
- Application specific customizations only affect higher mask sets, *amortizing mask costs*

Metal- or Via- Programmable Logics

- Mask-programmable technology are classified according to the nature of their customization and the granularity of their basic elements

- *Metal-programmable logics* allow for better performance/resource utilization but induce higher costs
- *Via-programmable logics* utilize pre-fabricated and pre-buffered metal layers.
  - *Printability issues are taken into account in the fabric design*
  - *Similarly to FPGAs, regular interconnects lead to safe and predictable timing analysis*

Fine Grained Fabrics

- Fine Grained: Similar to classic SOG: transistors and low level metallization to ease composition of standard blocks.
- User customization is dual:
  - A library of coarser functionalities is built above the fabric, typically by the provider
  - A synthesis step is performed using such library on the user application
- This solution is very flexible, utilizes standard tools but the gate-density may be low

Medium Grained Fabrics

- Medium Grained: Basic elements can be coarser blocks such as mux, nand, nor.
- Again, more complex functionalities are built grouping basic gates. This solution is less homogeneous, require specific allocation tools, but offers higher gate density with respect to fine grains

Y.Ran et al, Designing via-configurable Logic blocks for regular fabrics", TVLSI 2006

PLA-based Fabrics

- PLA-based: PLAs may be used for their inherent regularity and predictability
- They offer significant density, at the cost of higher power consumption and the need for specific synthesis and mapping tools

N.Jayakumar et al, “A metal and via Mask-programmable VLSI design methodology”, DAC04

LUT-based Fabrics

- LUT-based: Porting the concept of FPGA to mask-programmable contexts. SRAMs are collapsed to an array of connection to GND/VDD rails, while routing switches become via- or metal- programmable
- Standard FPGA synthesis and P&R tools can be exploited, and designs can be prototypized on SRAM fpgas before implementation providing a fast and safe deployment methodology.
- Collapsing SRAMs, design metrics are altered and area overhead may be introduced.
  - To limit this alternative interconnect schemes have been evaluated in literature (VPGA)

L.Pileggi at al, “Exploring regular fabrics to optimize performance-cost tradeoffs”, DAC2003

Coarse-Grained Fabrics

- Coarse-Grained: The array module can be composed of a regularly repeated set of pre-placed standard cells.

- In this case, typically performance is better while area occupation may suffer from some redundancy
- The layout of std-cells do not offer the design regularity for DFM available in most previous solutions

LightSpeed: US Patents 6014038 , 6943415, 6690194

SRAM-based Fabrics

- SRAM-based: One vendor provides an array of SRAM-based LUTs connected by a mask-programmable routing architecture.

- This minimizes routing overhead of FPGAs while retaining higher design flexibility.
- The obvious drawback of this solution is the area, power and control overhead induced by the configuration

eASIC: US Patent 6331790

Mask Programmable Solutions Landscape

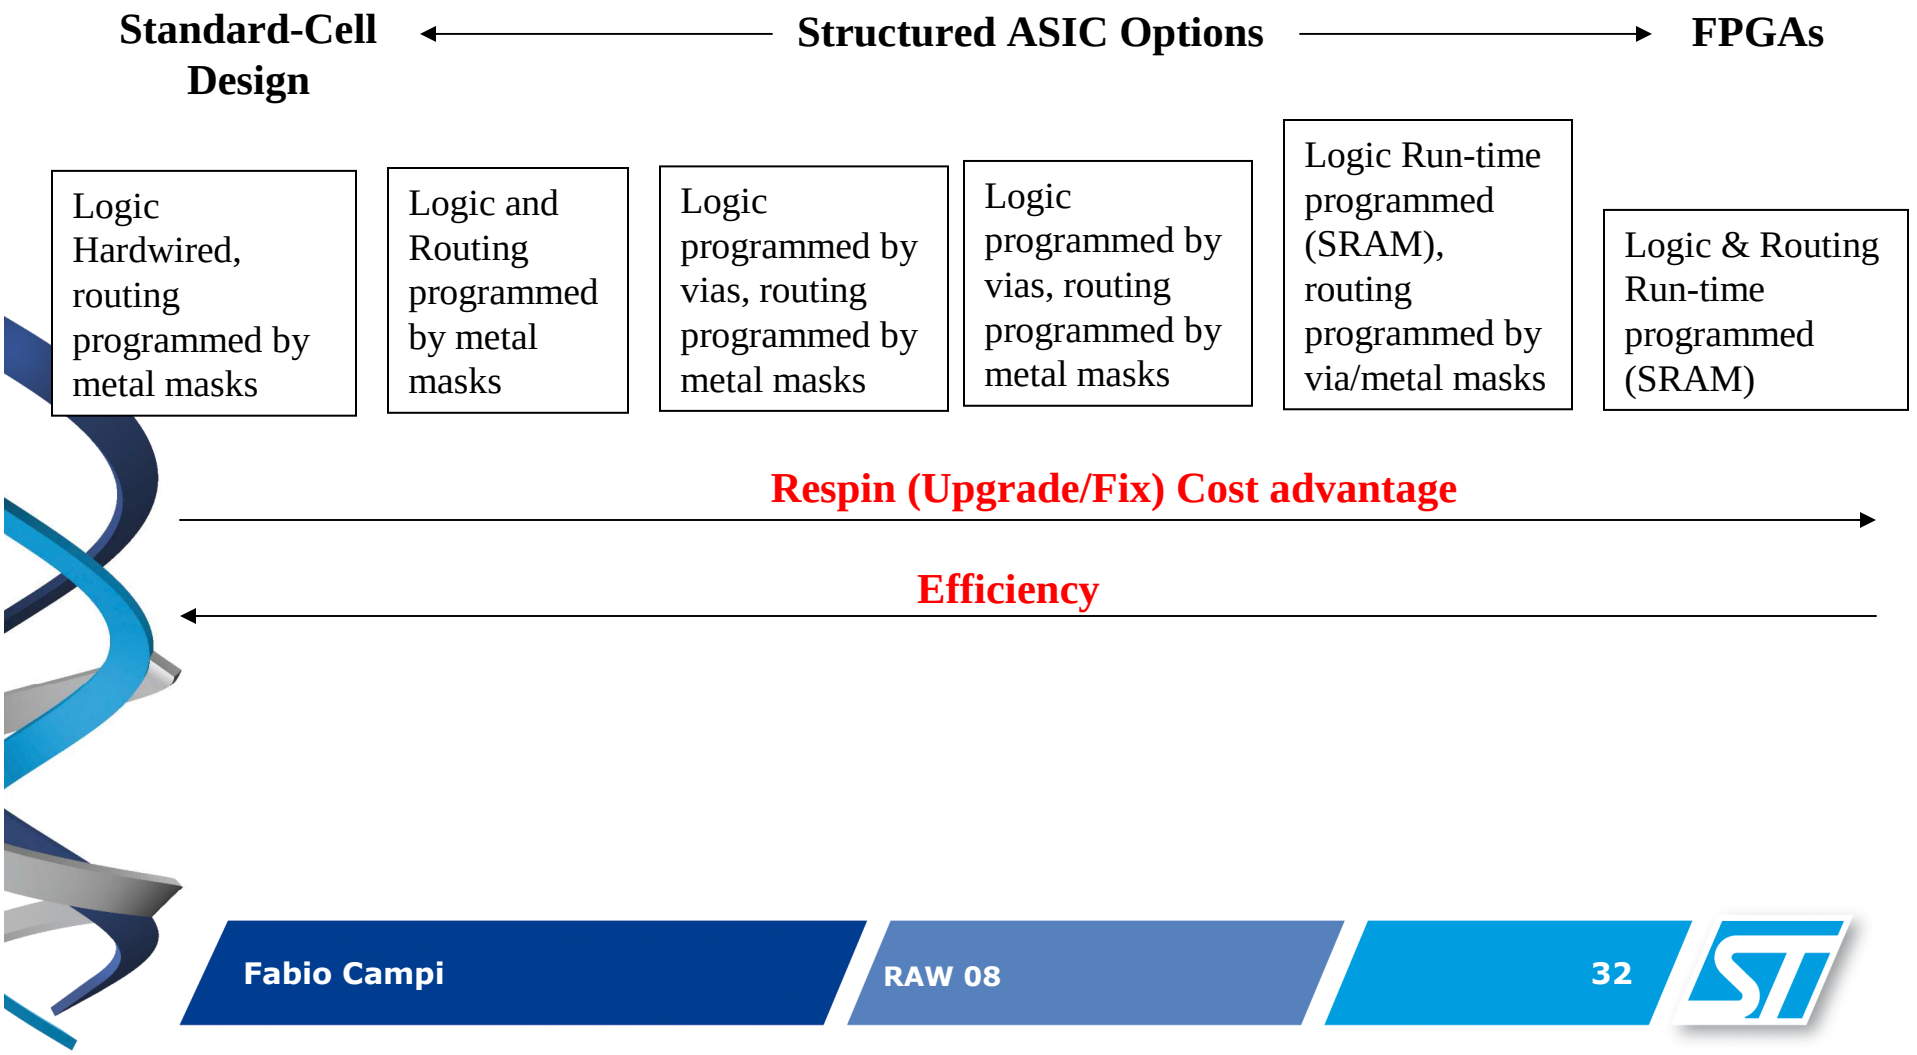

Altera Hardcopy

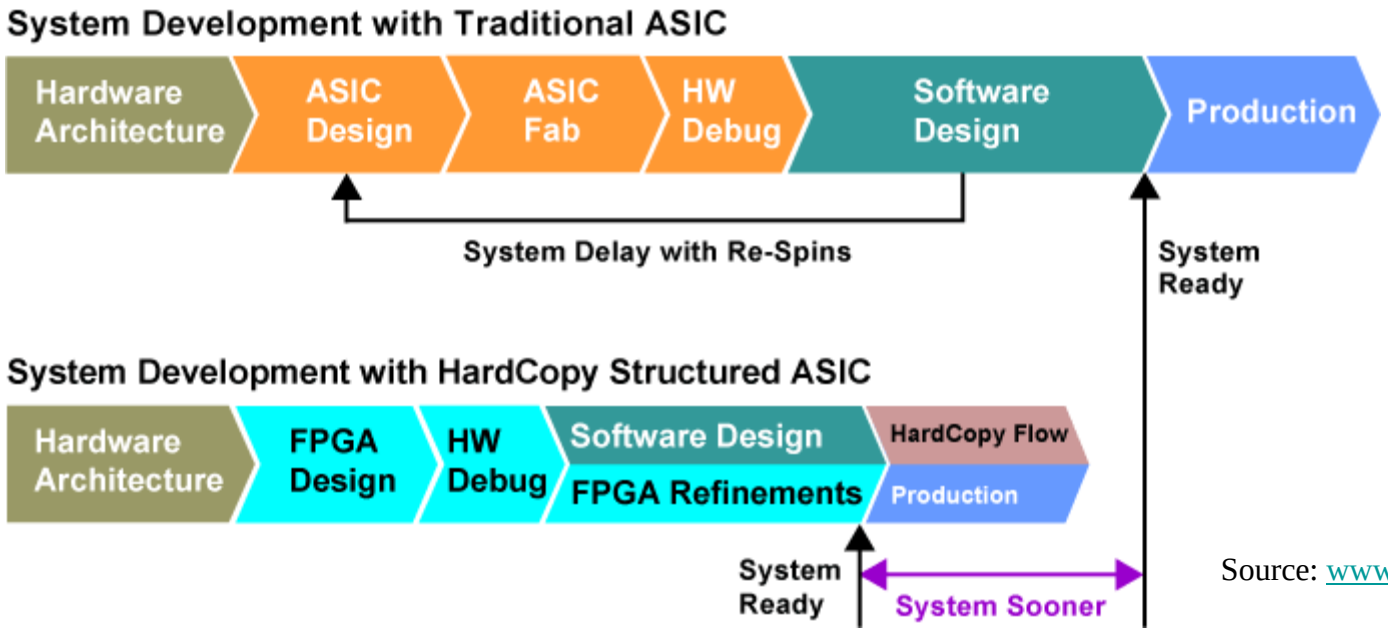

Source: www.altera.com, 08

- Hardcopy allows designers to prototype applications on the Stratix FPGA Family and utilize the final bit-stream as driver for an hardwired implementation

- Chips are built "hardwiring" the FPGA structure
- Full reuse of all Stratix IP blocks

Altera Hardcopy

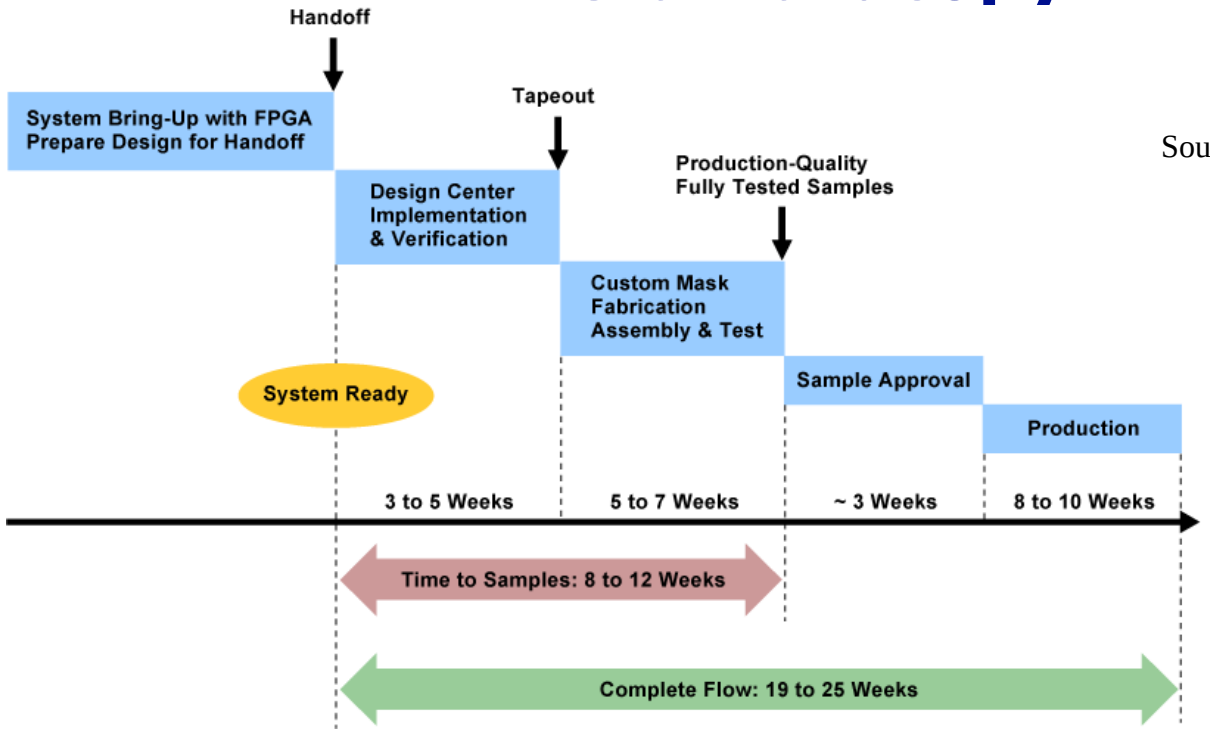

Source: www.altera.com, 08

- Claims:
  - 2.7 to 7M Asic gates
  - 4.2 to 16.3 Mbit on-chip memory
  - 50% power reduction with respect to FPGA solution

Platform ASIC (1)

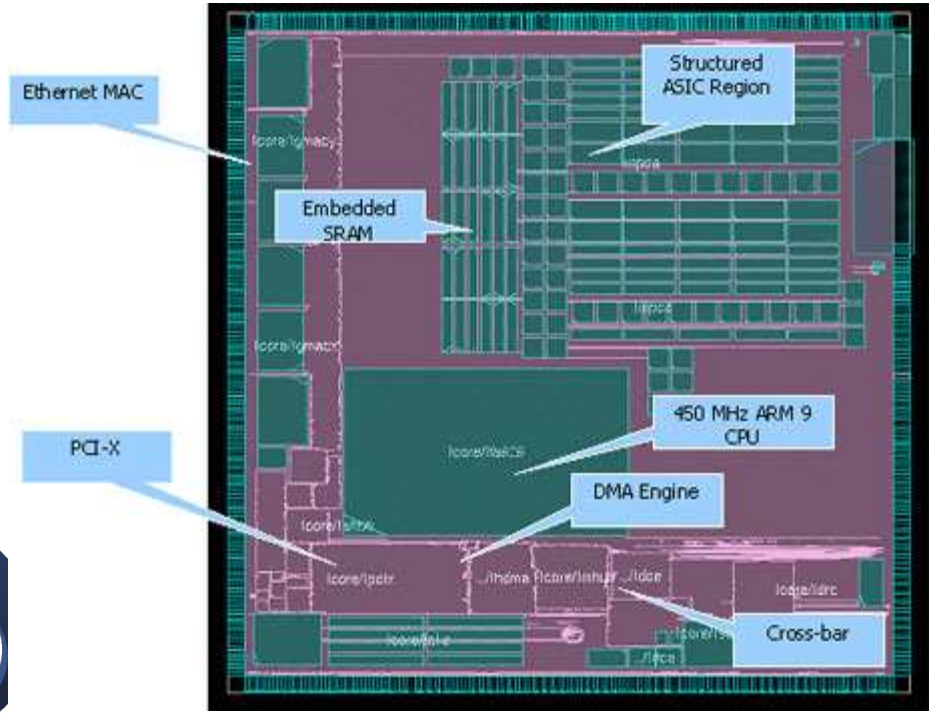

Source: www.faraday-tech.com, 08

- Several packaged solutions have emerged offering "Platform ASIC": fast Silicon solutions based on hardware IPs (Memory blocks, microcontrollers, PLLs, IO) where all glue logic and interconnect is mapped on Structured-ASIC

Platform ASIC (2)

Source: www.faraday-tech.com, 08

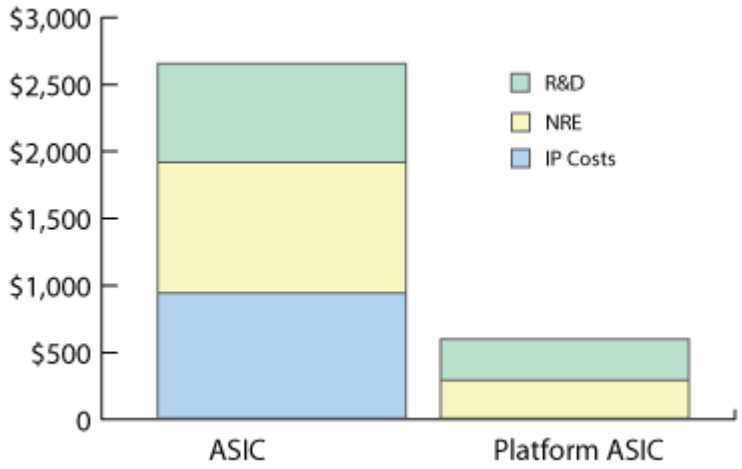

Source: elecrtonicdesign.com, 05

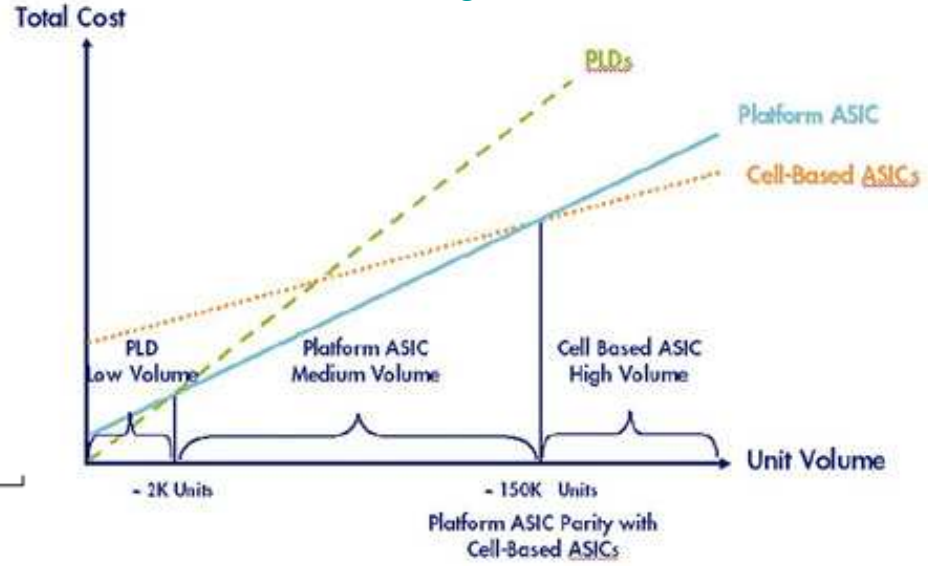

- This solution offer faster TTM, and lower risk for average-volume designs
- It is also possible to seamlessly port the design to ASIC in later stages of product life
- Offers: Faraday-tech, ARM, LSI Rapidchip, ChipX, others

Embedded Mask-Programmable Solutions

- Some vendors:
  - eASIC
  - ViAsic
  - Lightspeed
  - Mask-programmable versions of CGRAs
- Is there any added value embedding mask-programmable portions in a Multi-Core ASSP ?

(1) Mask Programmable CMOS library

- Logic cell library (~60 logic gates) with FE layers fixed (up to contact), common for all cells.
- Up to Metal-2 for logic cell definition
- Exploit Design regularity (DFM-friendly)
- Compatible to standard synthesis/P&R flows
- Drawbacks:
  - Gate Density (?)
  - Clock distribution, interconnect, scans must be handled by the user

(2) Structured ASIC Technology

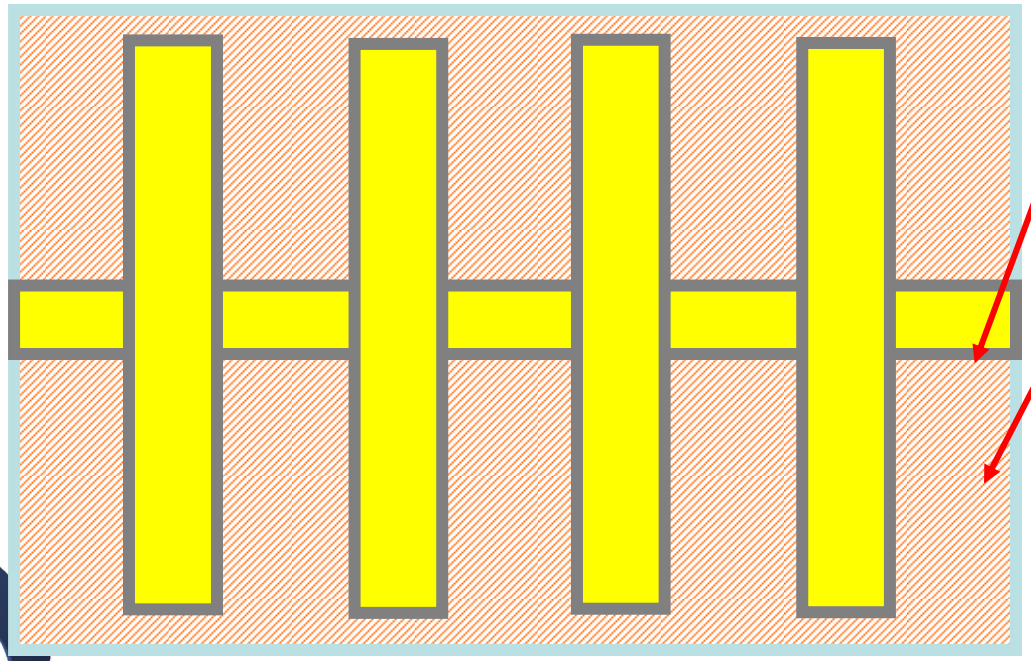

- Uncommitted Cells
- Regular Layout

- Mixed fabric composed of Std-Cell Flip-Flops mixed in a SOG-like array
- Compatible to standard synthesis/P&R flows
- Higher gate density wrt plain SoG
  - Placing of FFs impose constraints that may affect performance
  - Some issues for multiple clock domains

(3) Std-Cells based solution

- Tiles made up of regular standard cell library
- Standard cell mix: P&R Cells and Logic Cells
- Configurable area made up of metal-programmed regular repetition of tiles
  - Advantages: Density, Performance, Portability
  - Disadvantages: Non DFM Friendly, may be redundant

(4) Via-Configurable Datapath

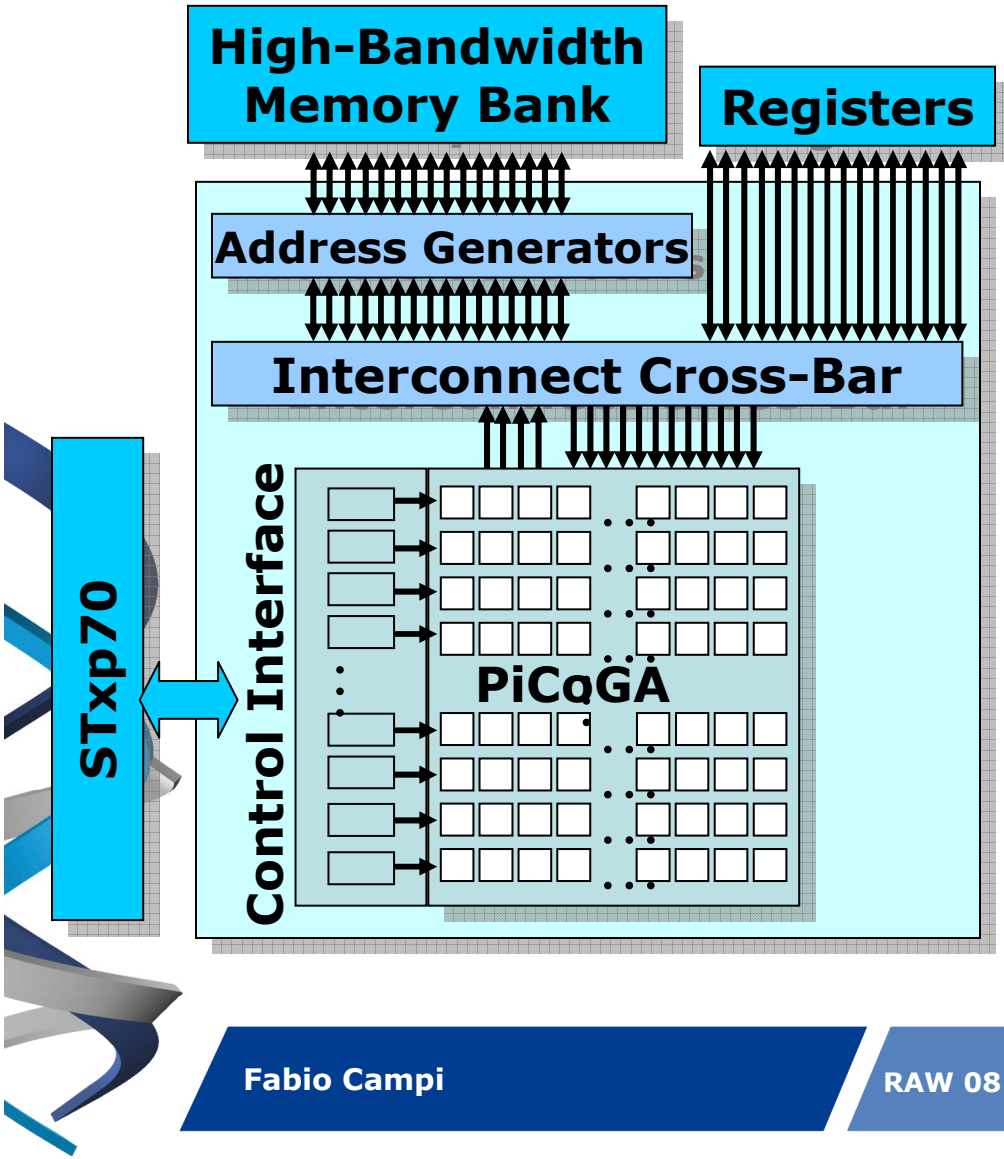

- MT-PiCoGA (Mask-Tuned LUT gate array)
- Mask programmable version of a Reconfigurable processor
- Performance boost (65 nm)
  - GA Area: 10 to 3 mm^2
  - GA Clk: 200 to 270 MHZ
- Only VIA masks affected by customization
- Same tool-flow as the run-time programmable version

SPEAr: "Pipecleaning" product line

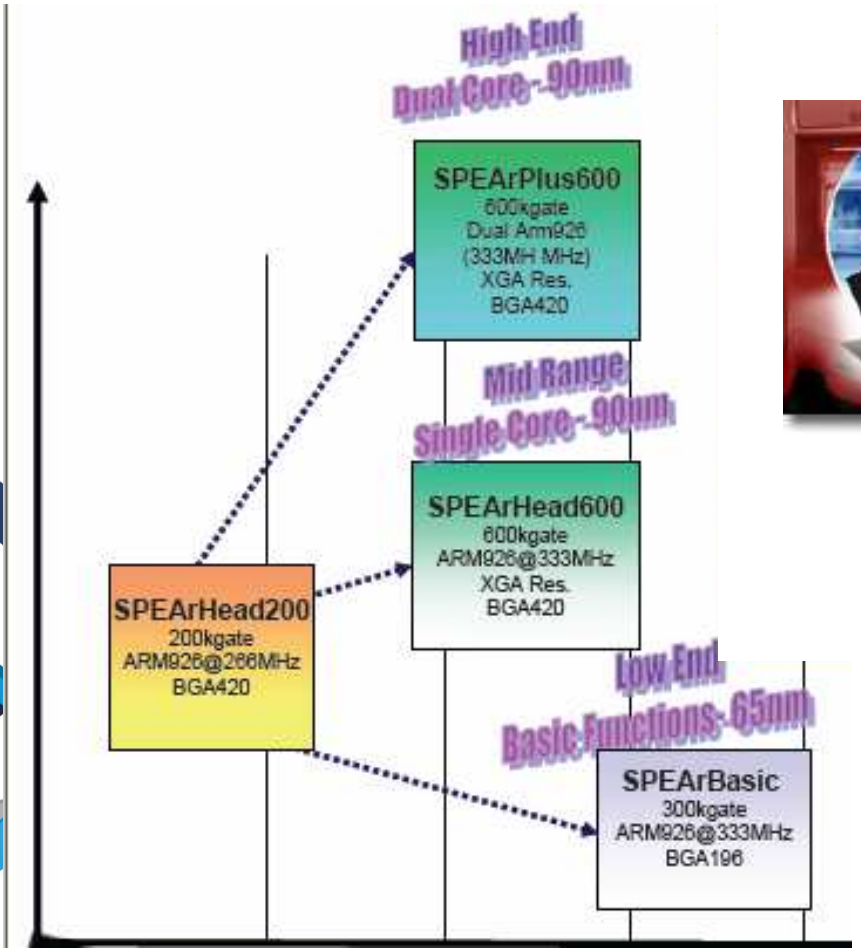

SPEAr Family:
"Structured Processor Enhanced Architecture"

ASSP for
Printer Applications

SPEAr: "Pipecleaning" product line (2)

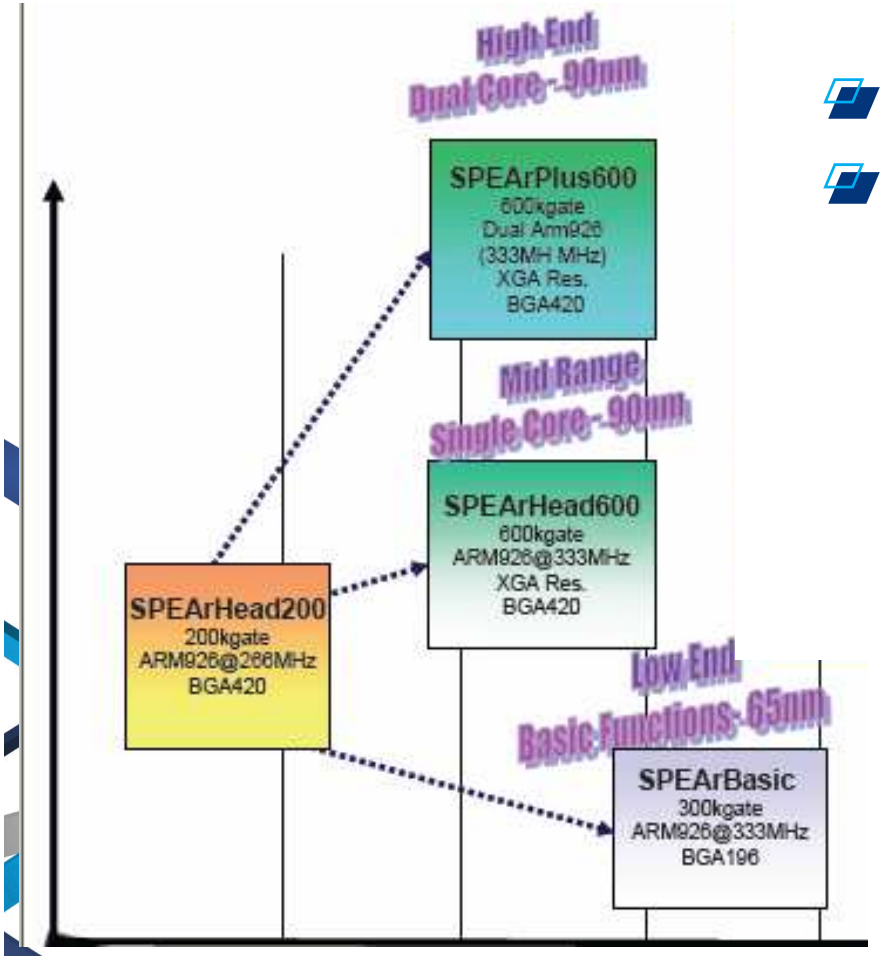

- 200K to 600K configurable gates
- Customization for new application domain:
  - 6 weeks vs 6 months for respin
  - 1/5 of full ASIC NRE
- Configurability advantage: Dual AIM:
  - high volume, low cost applications -> TTM
  - low volume, high added value -> NRE Costs

Exploitation of Mask programmability in ASSP (1)

- A gradual migration from FPGA to S-Asic to ASIC is not applicable to ASSPs, because integration is so high and costraint so tight that altering a portion constraints would require complete redesign
- Partition between different technology must be planned from early stages of the product -> S-ASIC accelerators can be part (or all) of a core in a MPSoC
- There is large space for architectural evaluation of the concept

Exploitation of Mask programmability in ASSP (2)

Selection of candidate blocks for Mask programmable based implementation:

- Key issues:
  - What parts are timing critical and must be hardwired?
  - What parts can be considered fixed over a large spectrum of applications, thus can be hardwired saving area/power/design cost ?
  - How much area overhead can be afforded in the design
  - How much configurable area may be needed in the future

Exploitation of mask programmability in ASSP (3)

- Metrics:
  - NRE Costs
  - TTM
  - Manufacturability

- The aim is to determine the best ratio between manufacturability versus NRE cost, that is
- Alter Design cost / manufacturability ratio in the direction of higher manufacturability per cost unit
- The opportunity is to invest high effort in heavily reused parts that can amortize NRE while providing solid configurable support for application specific customization

SPEAr: Loosely coupled coprocessor and Programmable IO

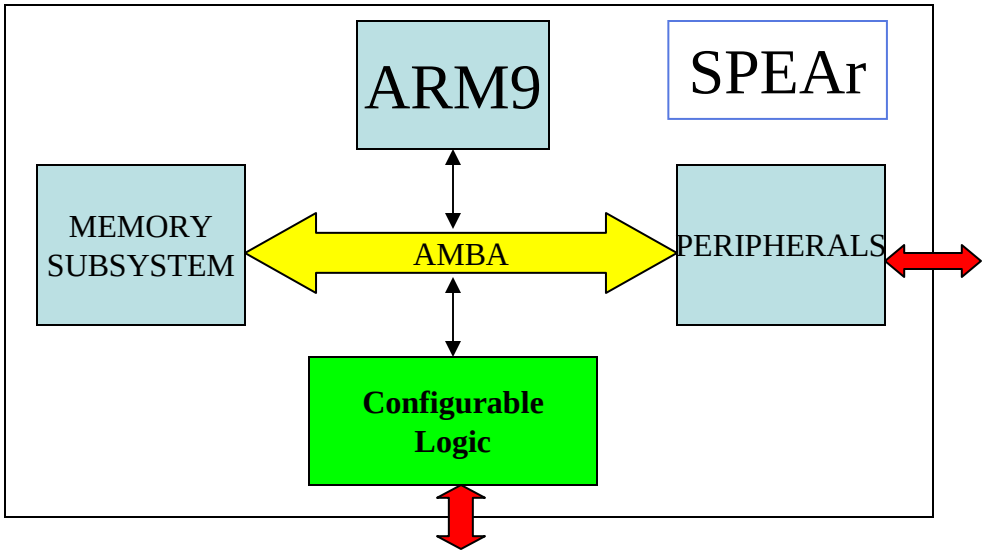

- First product line, efforts are focused on Technology aspects
- Single Clock Tree balanced between ASIC and Mask-Programmable parts
  - Mask programmable block drives overall speed
  - No real time constraints, few performance constraints
- Very few assumptions on the Customer utilization
- Configurable block also used for IO

Instruction Set Extension

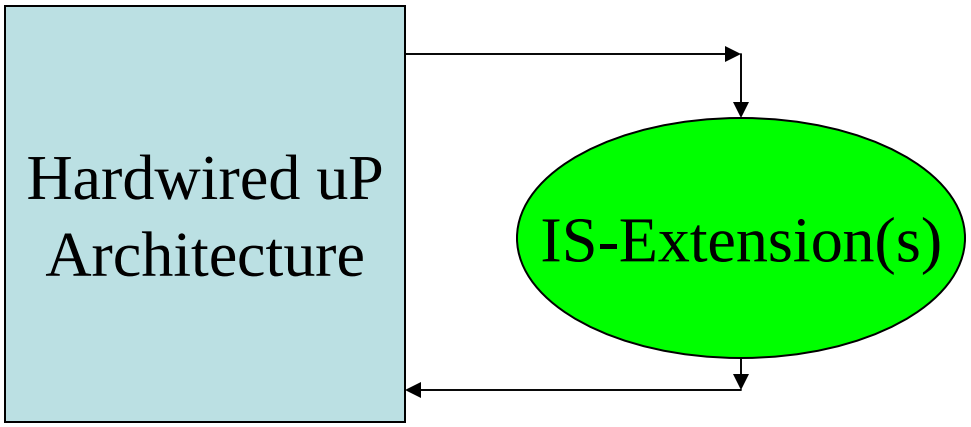

- Athanas/Silvermann, Instruction Set Metamorphosys
- Razdan/Smith, Prisc 1994

- Tensilica/Arc Model
- Not applied on eFPGA due to Array size and speed, can be exploited in the Mask-Programmable domain

Instruction Set Extension (2)

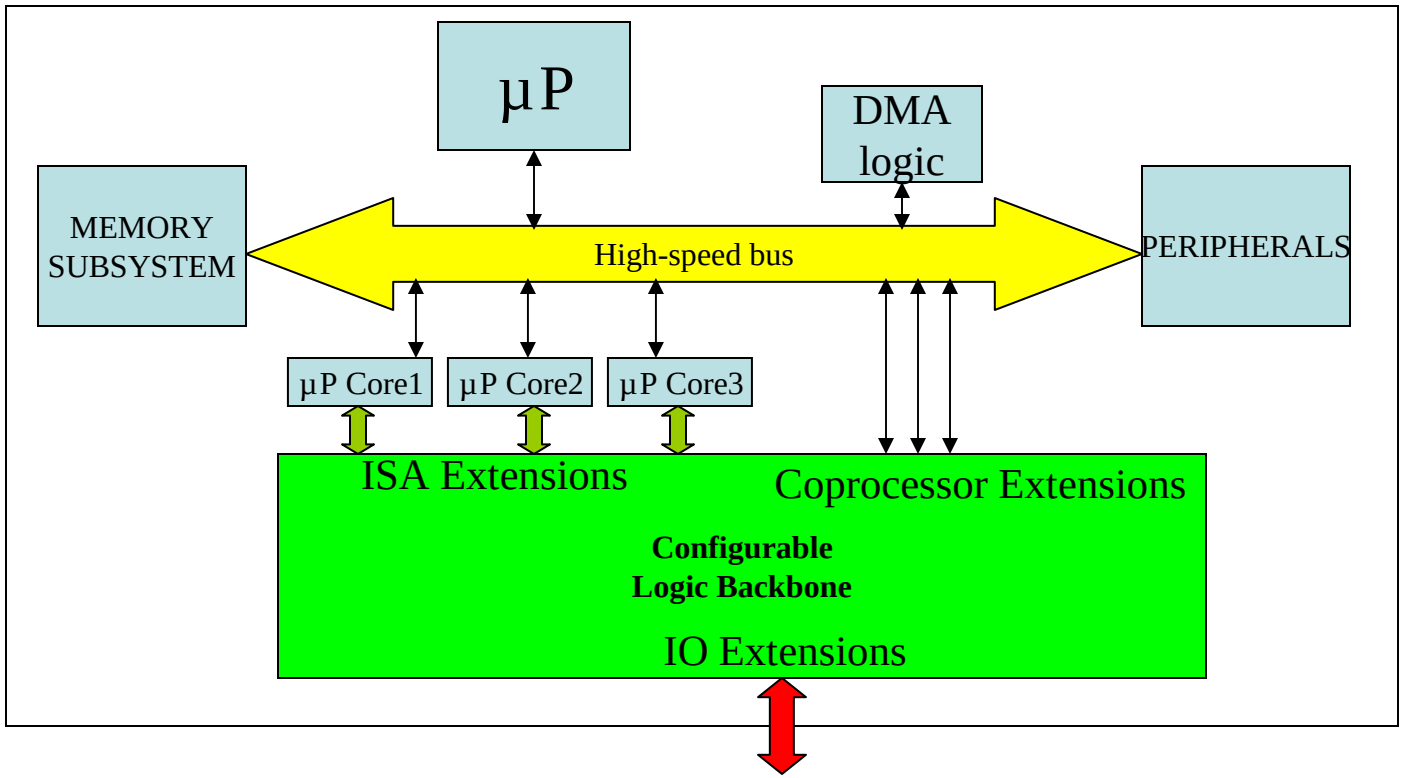

Notes on HDL coding style

- Porting the standard HDL code designed for ASIC is not a good idea
  - ASIC Designers should not be the ones coding for Mask-Programmable fabrics
- *Metrics change*. Often FFs are comparatively cheaper. Speed is an issue, on area there is an intrinsic redundancy. Long wires (top level connections should be avoided as much as possible
- Computation should be heavily pipelined

Sequenced DFG Computation: Garp, Piperench

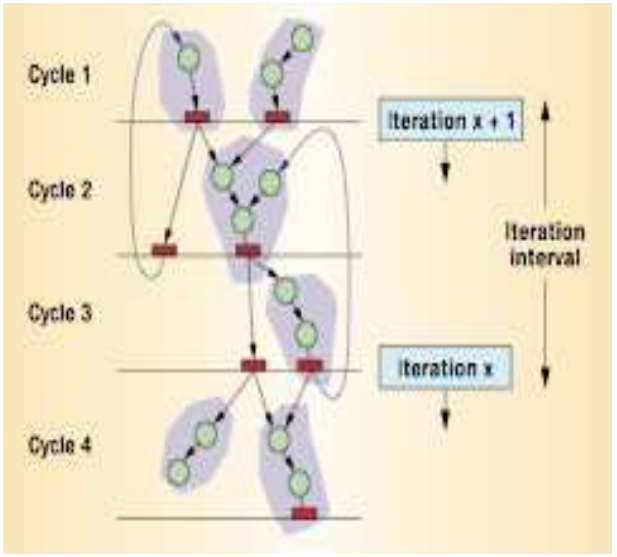

Callahan (UCB), “The Garp Architecture and C compiler”, IEEE Computer 2000

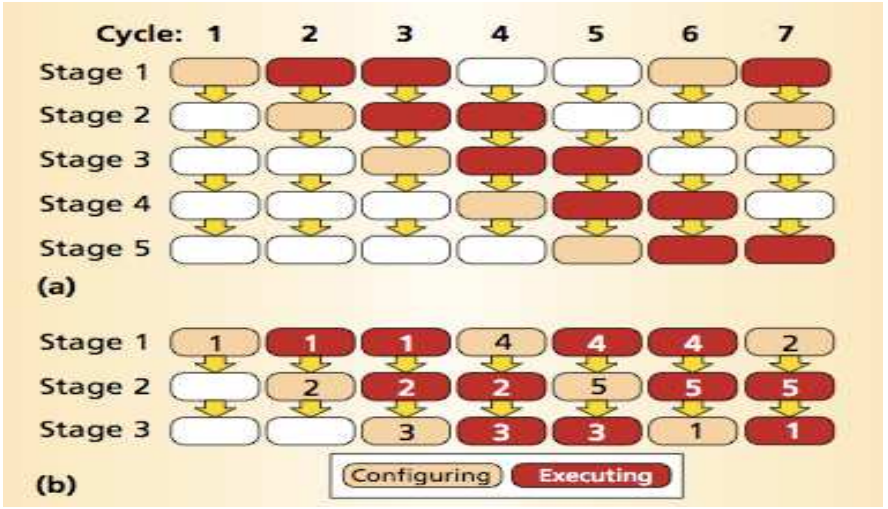

Source: Goldstein (CMU), “PipeRench: A Reconfigurable Architecture and compiler”, IEEE Computer 2000

- Computational kernels are mapped over a “striped” logic controlled by an hardwired sequencer
- Imperative description: DFGs are used as entry language, Single Assignment C is the preferred syntax (DIL, GriffyC)

Sequenced DFG computation (2)

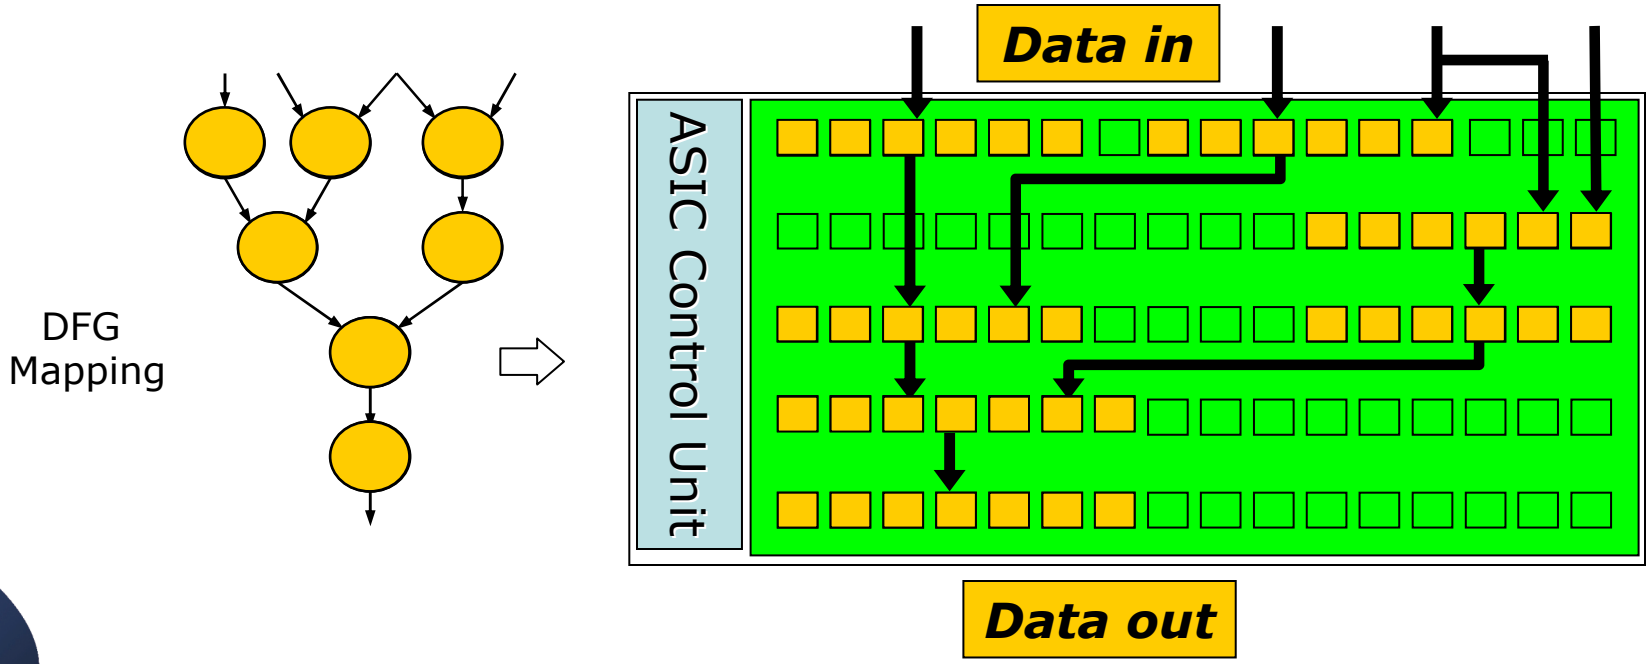

- Only operators are mapped on the fabric: synchronization and data consistency are implicit in the DFG and maintained by the hardwired sequencer
- Programming productivity significantly enhanced (~5x)

Notes on Physical implementation issues

- A strong limiting factor to "mixed-fabric" SoC is the need to balance clock trees over different technology supports, and be limited by the slowest fabric
- Asynchronous communication (Handshake or Asynchronous FIFOs) is a strong enabling factor for the deployment of such solutions

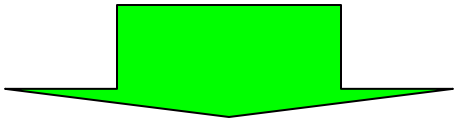

- High-end SoC feature heterogeneous multi-cores
- Heterogeneous Voltage/Freq domains are commonly deployed
- This "partitioned", design approach is an appealing environment for a third parameter: Heterogeneous technology fabrics (Run-Time/Mask-Programmable)

Tile-Oriented Computing (GP or ASSP)

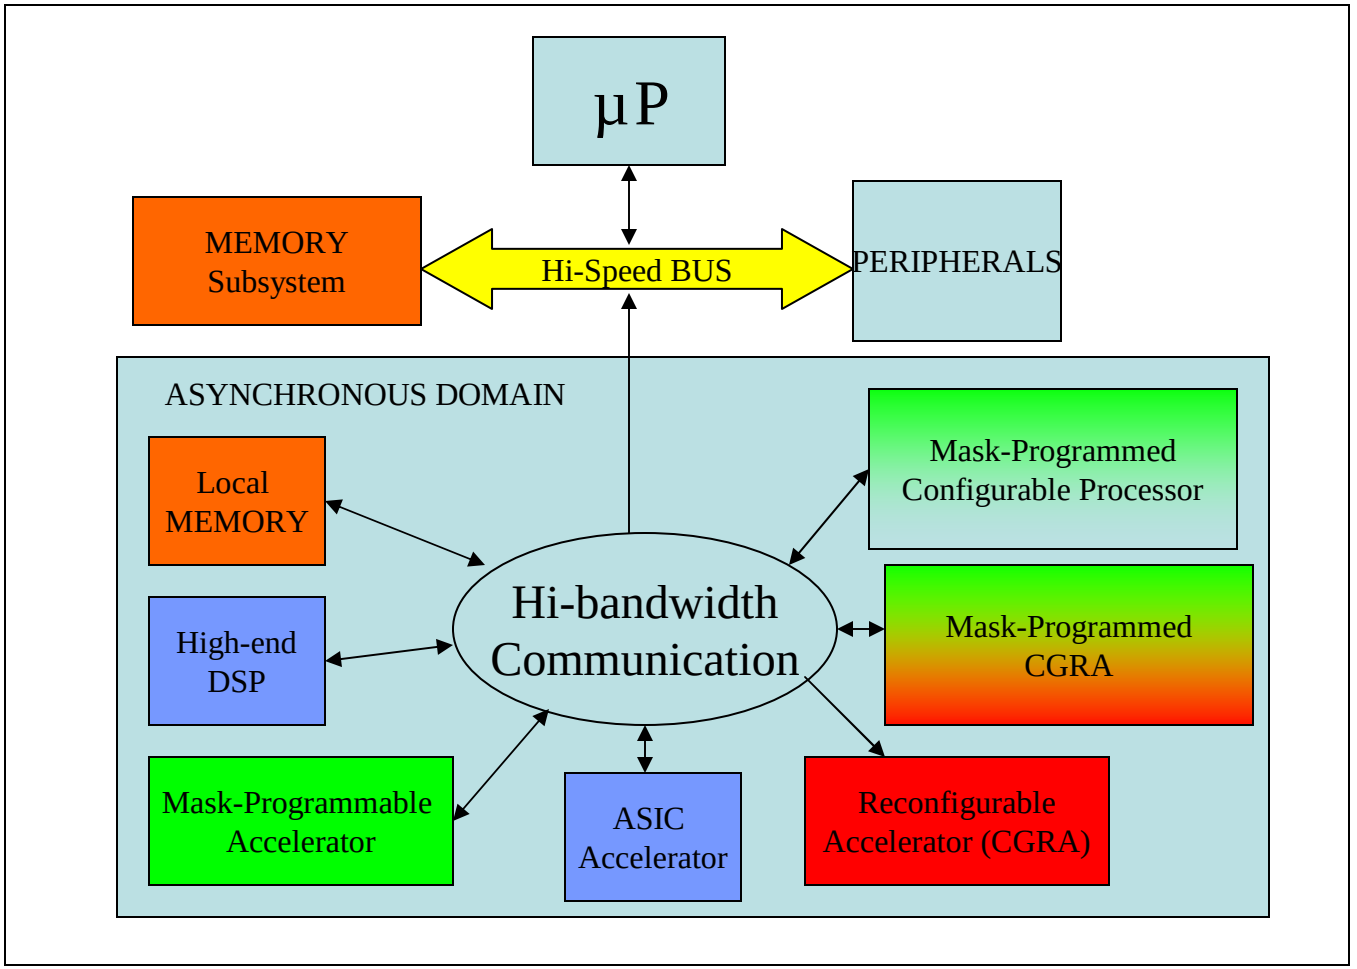

Conclusions

- Mask programmable technology is no breakthrough solution, borrowing concepts from SOG and FPGAs that were largely investigated in the 90s
- As compared to run-time programmability is largely less elegant and immediate
- From an economic perspective, it represents an useful way to retain most advantages of run-time reconfiguration while mitigating area costs
- Reconfigurable computing community has a portfolio of knowledge for the ideal exploitation of such concepts.
- Mask programmable fabrics should be seen as an opportunity rather than an alien or competitor technology, replicating to the ASSP domain the success of Altera Harcopy