

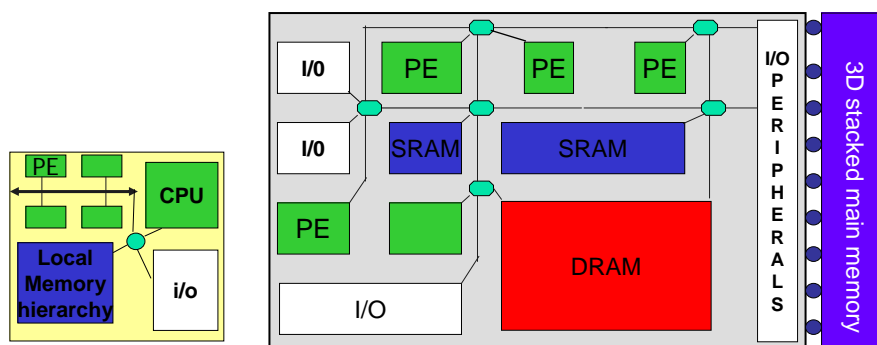
Luca Benini – DEIS Università di Bologna
lbenini@deis.unibo.it

MPSoCs – Hardware platforms

- Why MPSoCs?
- MSoC architectures
- Case studies

Luca Benini ARTIST2 / UNU IIST 2007

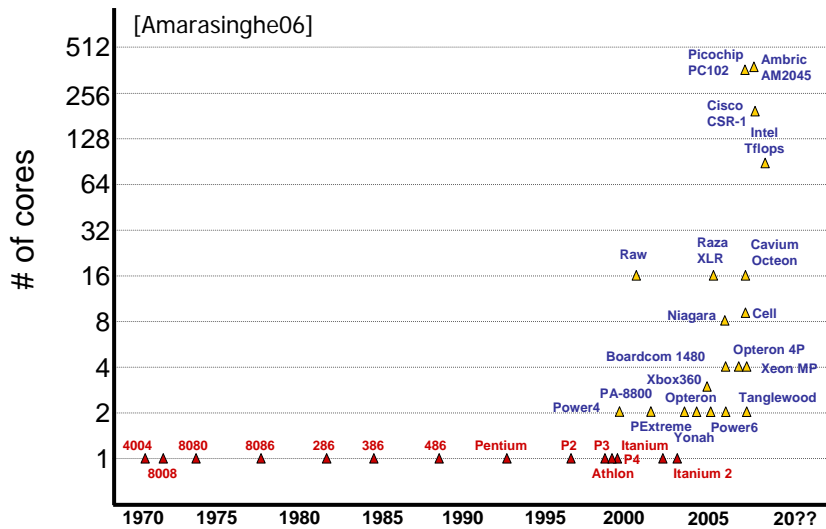
Architecture Evolution



- Roadmap continues: 90→65→45 nm
- “Traditional” Bus based SoCs fit in one tile !!
- Communication demand is staggering, but **unevenly distributed**, because of architectural heterogeneity


Luca Benini ARTIST2 / UNU IIST 2007

Multicores Are Here!



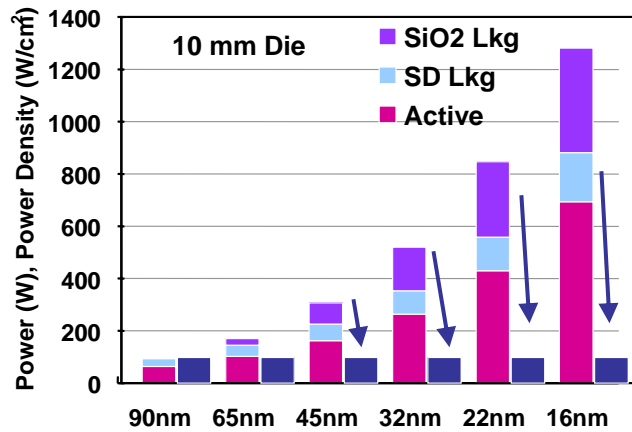
Luca Benini ARTIST2 / UNU IIST 2007

MPSoC – 2005 ITRS roadmap



Luca Benini ARTIST2 / UNU IIST 2007

Power is the Challenge!



Technology, Circuits, and Architecture to constrain the power


Luca Benini ARTIST2 / UNU IIST 2007

Near Term Solutions


- Move away from Frequency alone to deliver performance
- More on-die memory
- Multi-everywhere
 - Multi-threading
 - Chip level multi-processing
- Throughput oriented designs
- Performance by higher level of integration

Luca Benini ARTIST2 / UNU IIST 2007

μArchitecture Techniques





Multi-threading



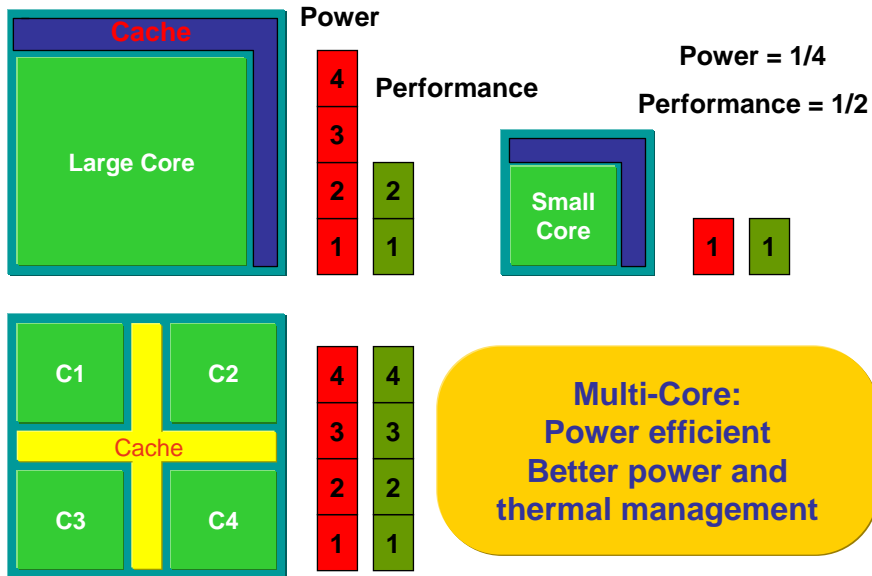
Improved performance, no impact on thermals & power delivery

Chip Multi-processing



Luca Benini ARTIST2 / UNU IIST 2007

Multi-Core



Multi-Core:
Power efficient
Better power and thermal management

Luca Benini ARTIST2 / UNU IIST 2007

Embedded vs. General Purpose

Embedded Applications

- *Asymmetric Multi-Processing*
 - Differentiated Processors
- Specific tasks known early
 - Mapped to dedicated processors
- Configurable and extensible processors: performance, power efficiency
- Communication
 - Coherent memory
 - Shared local memories
 - HW FIFOs, other direct connections
- Dataflow programming models
- Classical example – Smart mobile – RISC + DSP + Media processors

Server Applications

- *Symmetric Multi-Processing*
 - Homogeneous cores
- General tasks known late
 - Tasks run on any core
- High-performance, high-speed microprocessors
- Communication
 - large coherent memory space on multi-core die or bus
- SMT programming models (Simultaneous Multi-Threading)
- Examples: large server chips (eg Sun Niagara 8x4 threads), scientific multi-processors

Luca Benini ARTIST2 / UNU IIST 2007

MPSoC architectures

Luca Benini ARTIST2 / UNU IIST 2007

Example system platforms

- Generic
- Automotive
- Wireless
- Multimedia

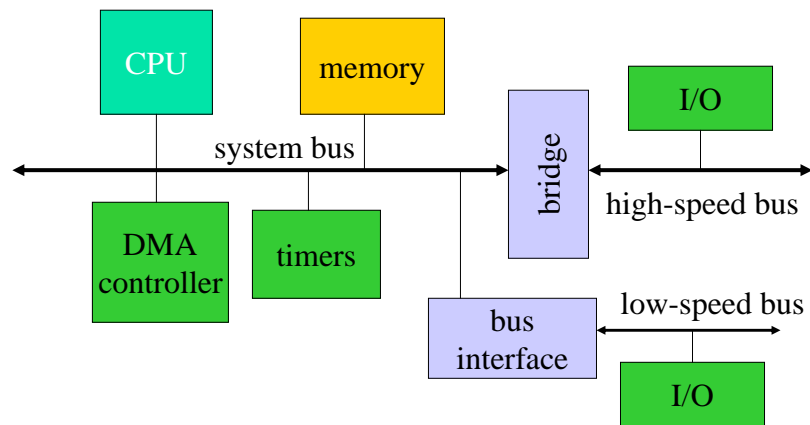
Luca Benini ARTIST2 / UNU IIST 2007

PC-based platform

- Basic hardware components:
 - CPU;
 - memory;
 - timers;
 - DMA;
 - minimal I/O devices.
- Basic software:
 - BIOS.

Luca Benini ARTIST2 / UNU IIST 2007

PC-style hardware architecture



Luca Benini ARTIST2 / UNU IIST 2007

Strong ARM

- StrongARM system includes:
 - CPU chip (3.686 MHz clock)
 - system control module (32.768 kHz clock).
 - Real-time clock;
 - operating system timer
 - general-purpose I/O;
 - interrupt controller;
 - power manager controller;
 - reset controller.

Luca Benini ARTIST2 / UNU IIST 2007


Pros and cons

- Plentiful hardware options.
- Simple programming semantics.
- Good software development environments.
- Performance-limited.

Luca Benini ARTIST2 / UNU IIST 2007

TI Open Wireless Multimedia Applications Platform

- Dual-processor shared memory system:

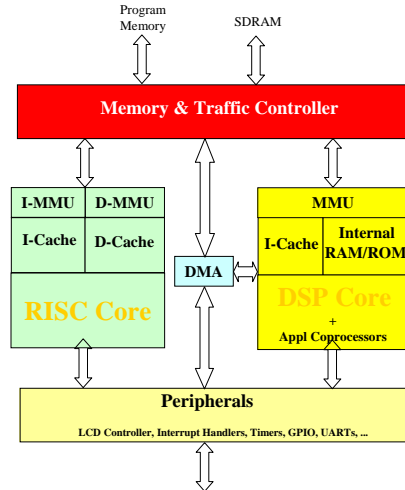


<http://www.ti.com/sc/docs/apps/wireless/omap/overview.htm>

Luca Benini ARTIST2 / UNU IIST 2007

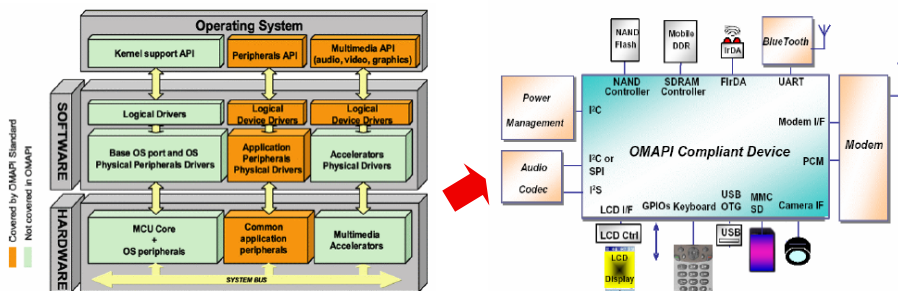
TI OMAP™ Hardware platform

- ARM9 core
 - 16KB I cache
 - 8KB D cache
 - 2 way set associative
 - 150 MHz
- C55x DSP core
 - 16KB I cache
 - 8KB RAM set
 - 2 way set associative
 - 200 MHz



Luca Benini ARTIST2 / UNU IIST 2007

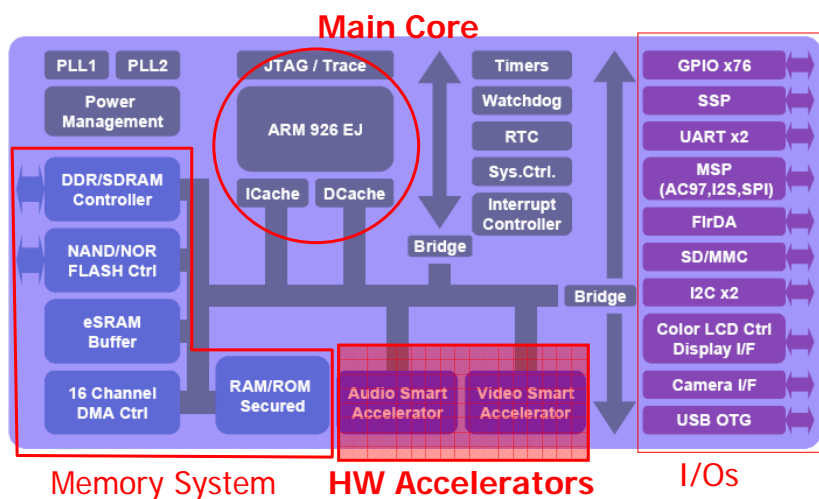
OMAPI Standard (ST/TI)



- Goal: standardize the interfaces between application processor and peripheral devices in a mobile product
- Provide standard services (APIs) in the OS that can be used by application developers

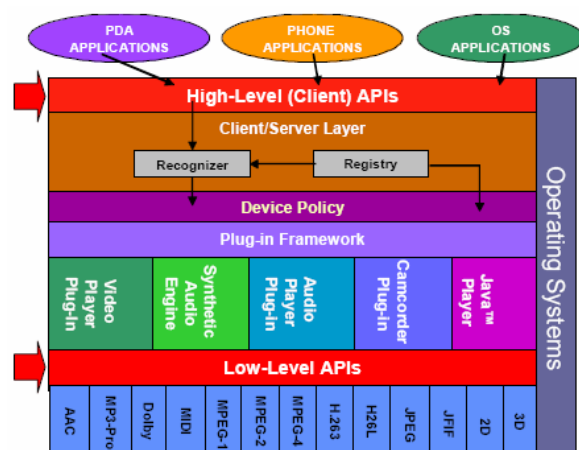
Luca Benini ARTIST2 / UNU IIST 2007

STMicro Nomadik platform



Luca Benini ARTIST2 / UNU IIST 2007

Nomadik SW platform



- Compliant with OMAPI standard

Luca Benini ARTIST2 / UNU IIST 2007

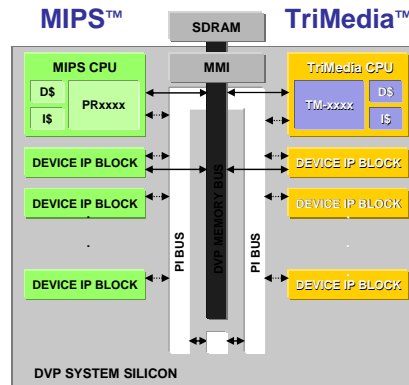
Philips Digital Video Nexperia Platform

General-purpose Scalable RISC Processor

- 50 to 300+ MHz
- 32-bit or 64-bit

Library of Device IP Blocks

- Image coprocessors
- DSPs
- UART
- 1394
- USB
- ...and more



Scalable VLIW Media Processor:

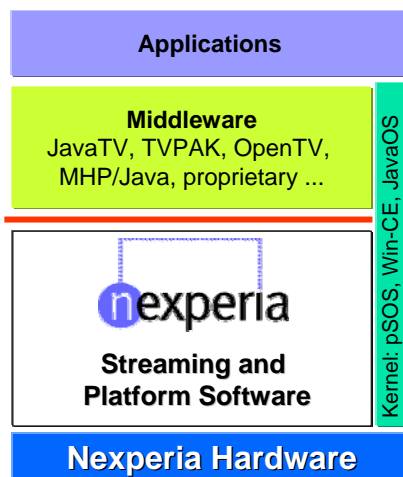
- 100 to 300+ MHz
- 32-bit or 64-bit

Nexperia™ System Buses

- 32-128 bit

Luca Benini ARTIST2 / UNU IIST 2007

Nexperia-DVP Software



■ Nexperia™ -DVP Software Architecture

- Supports multiple OSs and middleware software
- Abstracts platform functionality via consistent APIs

■ Nexperia™ -DVP Streaming Software

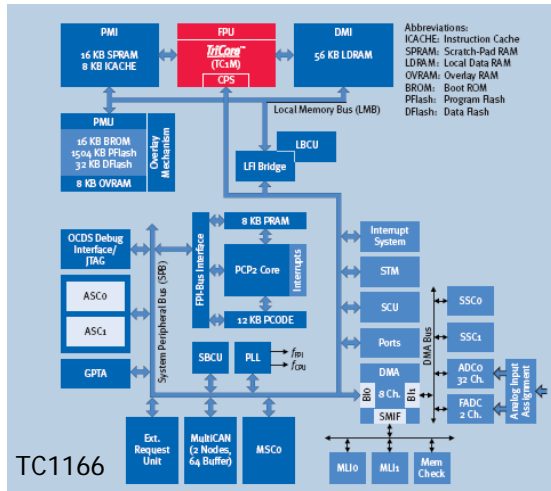
- Encapsulates implementation of streaming media components (hardware and software)

■ Nexperia™ Platform Software

- OS independent device drivers for on-chip and off-chip devices

Luca Benini ARTIST2 / UNU IIST 2007

Infineon Automotive Platform



Applications

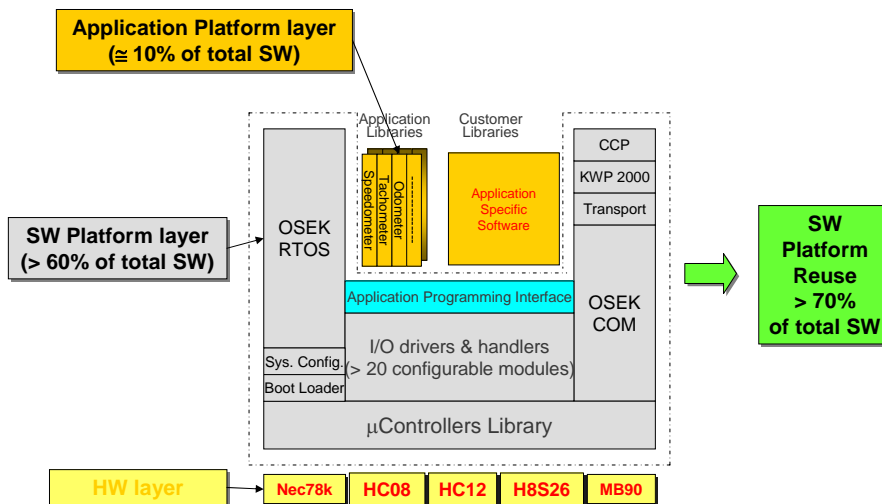
- High Performance drives / servo drives,
- Industrial control Robotics

Features

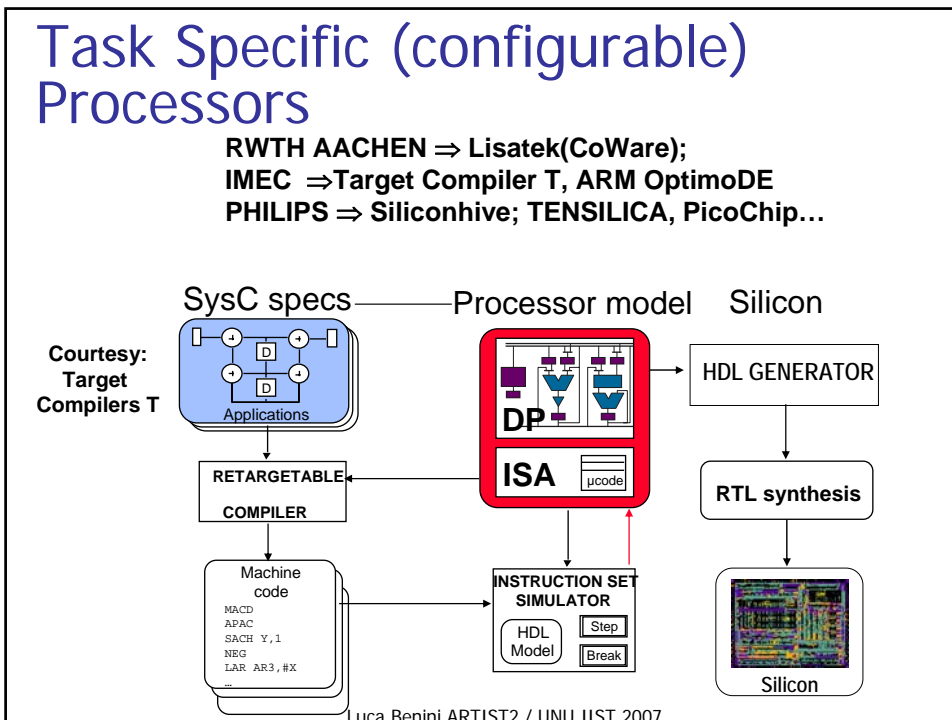
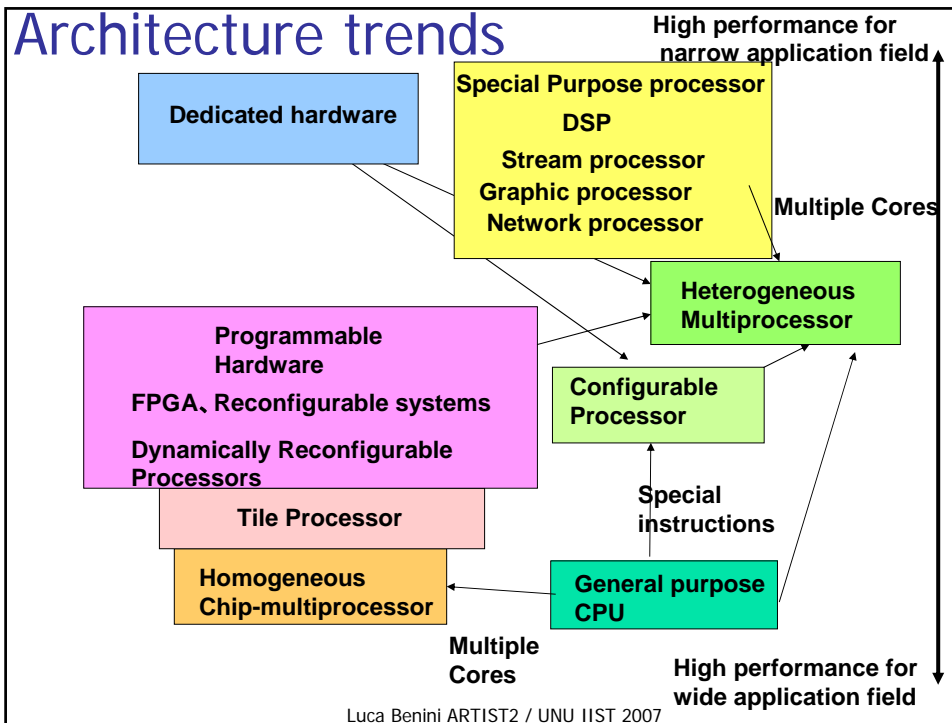
- 32-bit super-scalar TriCore™ V1.3 CPU, 4 stage pipeline
 - Fully integrated DSP capabilities
 - Single precision floating point unit (FPU)
 - 80 MHz at full industrial temperature range
- 32-bit peripheral control processor with single cycle instruction (PCP2)
- Memories
 - 1.5 MByte embedded progr. flash with ECC
 - 32 KByte data flash - EEPROM emulation
 - 56 KBSRAM, 8 KB I\$, 16 KB Imem
- 8-channel DMA controller
- Interrupt system with 2 x 255 hardware priority arbitration levels serviced by CPU and PCP2 Coprocessor
- Triple bus structure: 64-bit local memory buses to internal flash and data memory, 32-bit system peripheral bus, 32-bit remote peripheral bus

Luca Benini ARTIST2 / UNU IIST 2007

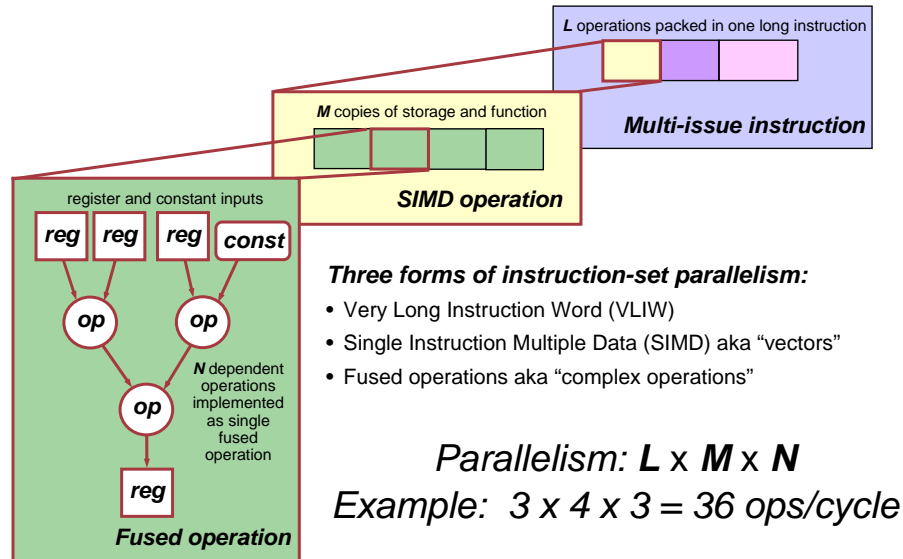
MOSAIC SW Architecture & Components for Automotive Dashboard and Body Control



Luca Benini ARTIST2 / UNU IIST 2007

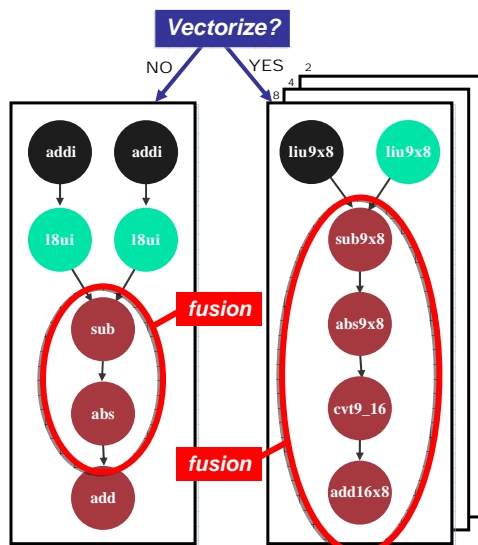


Parallelism at Three Levels in Extensible Instructions



Luca Benini ARTIST2 / UNU IIST 2007

Example: SAD (sum of absolute differences)




Original C Code

```
short total=0;
char *p1, *p2;
for i=1,m
  for j=1,n
    total += abs(*p1++ - *p2++)
```

Sample Software Pipelined Schedule Vector + Fusion + FLIX Configuration

```
loop j=1,n/8 by 2:
  liu9x8[j]; liu9x8[j]; fusion[j-2]
  liu9x8[j+1]; liu9x8[j+1]; fusion[j-1]
```



Luca Benini ARTIST2 / UNU IIST 2007

Dynamically Reconfigurable Processors


- Reconfigurable systems → Previous lesson
 - Flexible but It takes 10⁷ s milliseconds for dynamic reconfiguration.
- Dynamically Reconfigurable Processors
 - Improves area efficiency by changing hardware structure.
 - IPs used in various SoCs.
 - History
 - Reconfigurable Co-processor Garp(1997), CHIMAERA(2000)
 - Multicontext reconfigurable devices WASMII(1992), Time-multiplexing FPGA(1997), PipeRench(1998), DRL(1998)
 - Functional-level synthesis
 - Various commercial products are available since 2000
 - IPFlex DAPDNA-2, NEC electronics DRP-1, PACT Xpp, Elixent DFabrix
 - SONY's VME(Virtual Mobile Engine) is embedded in Network Workman and PSP
 - Recently, many Japanese vendors start to develop commercial products
 - Fujitsu
 - Hitachi
 - Lucent
 - Sanyo
 - Toshiba (Mep+D-Fabrix)

Luca Benini ARTIST2 / UNU IIST 2007

Processing Element

- Specialized for media/stream processing
Coarse grain ↔ Fine grain: LUT of FPGAs
- Components
 - ALU
 - Shifter+Mask unit
 - Multiplexers
 - Registers
- Operations and interconnection between components are changeable
- No instruction fetch mechanism : A part of large datapath

Luca Benini ARTIST2 / UNU IIST 2007



Putting it all together

	2004	2006	2008	2010	2012
Technology Node (nm)	90	65	45	32	22
Loosely coupled Sub-Systems	2	4	6	8	12
General Purpose CPU	Single	➔ Multiple			
Hardware Accelerator	Hardwired	➔ Reconfigurable			

- Constant SoC Die Size
- Slow evolution of peripherals (area decrease)
- GP CPU sub system complexity 2x each node (constant area),
- Embedded Memory capacity 2x at each node (constant area)
- Loosely coupled DSP sub system complexity increase by 30% at each node (30% area decrease)

Luca Benini ARTIST2 / UNU IIST 2007

Main trends

- Host CPU evolving toward multi-core architecture to meet the performance increase requirements
- HW acceleration mapped on reconfigurable arrays
 - Performances close to dedicated HW in many areas
 - Good fit with regular design constraints imposed by 45nm process and beyond
 - Excellent structure for best optimized power management
 - And ... FLEXIBILITY ...

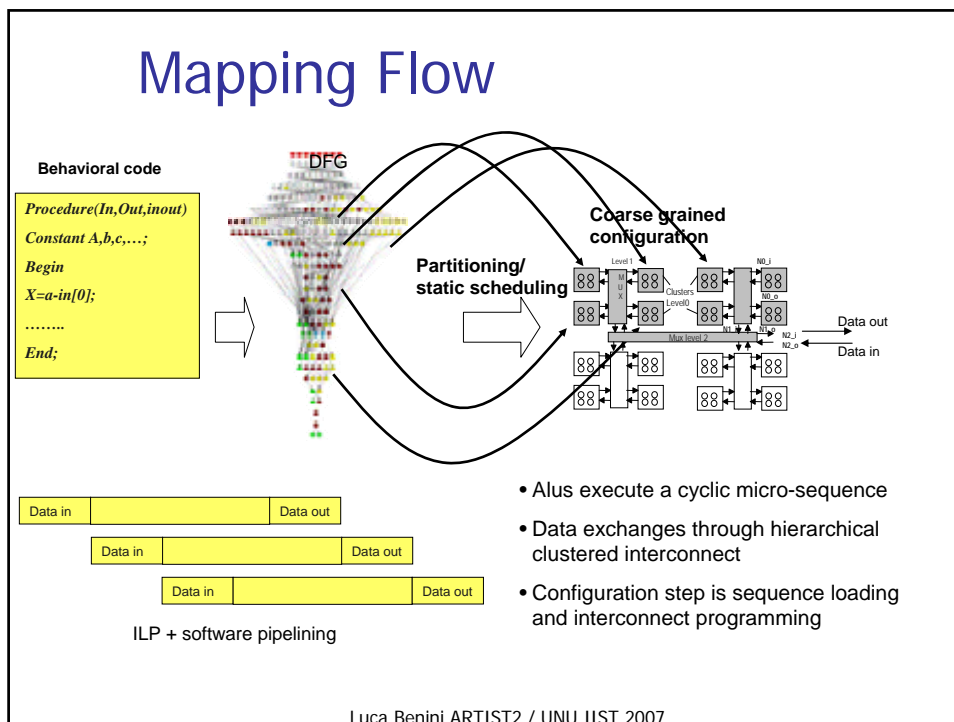
Luca Benini ARTIST2 / UNU IIST 2007

Reconfigurable HW (DSP fabric)

- Target signal processing and arithmetic intensive applications
- Reconfigurable array of simple DSP core (CNode)
- Low power architecture
 - Hierarchical clock gating
 - Distributed leakage control (fine grain power gating)
- Programmable DMA engine
- Reconfigurable at run time, multi task

Luca Benini ARTIST2 / UNU IIST 2007

Mapping Flow

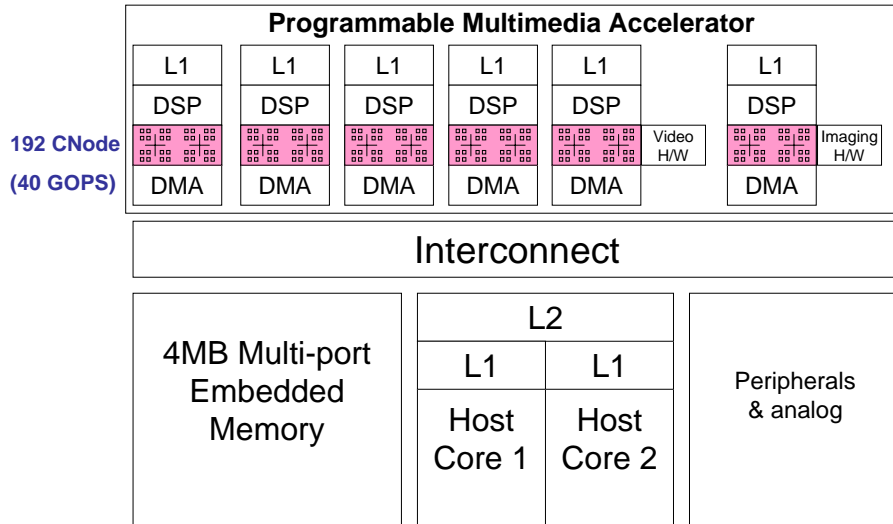


Mapping Flow

- 3D optimization problem (place/route/schedule)
- Traditional scheduling techniques for VLIW or clustered VLIW don't apply
 - The solution don't take into account the spatial dimension of the problem
- Traditional P&R used in FPGA don't apply neither because they don't consider the time dimension

Luca Benini ARTIST2 / UNU IIST 2007

What can fit in 45mm² in 45nm



Luca Benini ARTIST2 / UNU IIST 2007

Case Study: GPUs

Luca Benini ARTIST2 / UNU IIST 2007

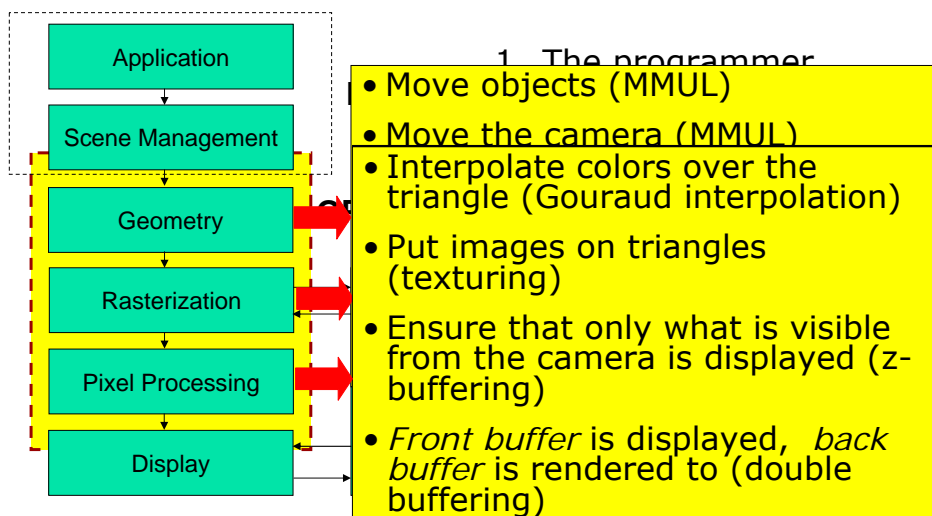
Mobile graphics platforms



300-400 million mobile phones with graphics hardware (OpenGL ES) by 2009

Luca Benini ARTIST2 / UNU IIST 2007


The 3D Graphics Pipeline



Luca Benini ARTIST2 / UNU IIST 2007

Why is it hard with 3D graphics on mobile devices?

- Small amount of memory
- Limited instruction set
- Low clock frequency
 - 100-200 MHz ARM9–400-600 MHz ARM11
- Small area on the chip for CG
- Must be cheap and physically small
- Powered by batteries!
 - A memory access is one of the most expensive operations
 - Battery growth: 9% per year
 - Performance growth: 40% per year
- Small display, but very close to the eye
 - Avg. Eye-to-pixel angle 1-4x larger than for desktop



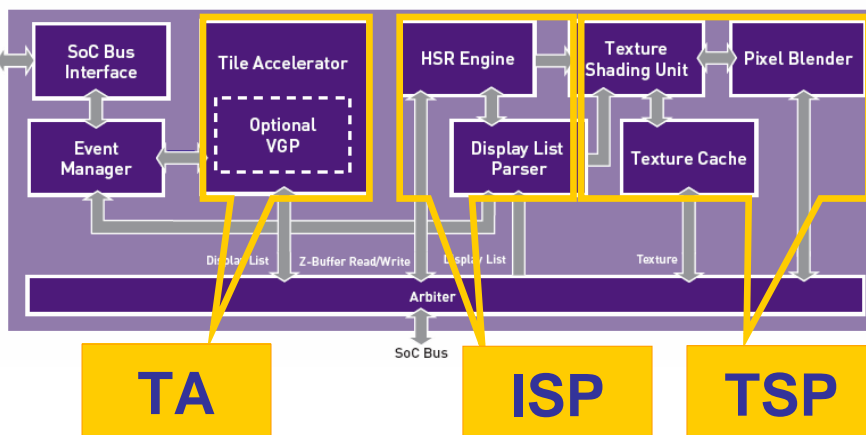
Limited resources, but high quality rendering!

Luca Benini ARTIST2 / UNU IIST 2007

PowerVR MBX low-power GPU

Features


- Tile-based rendering
- Tile accelerator
- ITC™: PowerVR internal true color: color ops on-chip at 32-bpp
- Image synthesis processor
- ESAA4Free™: full screen anti-aliasing for realism at mobile display resolutions
- Texture and shading processor
- PVR-TC™: texture compression for small memory footprints.



Luca Benini ARTIST2 / UNU IIST 2007

Tiled (sort middle) architecture

- Apply geometry transf. (incl. projection) to vertices



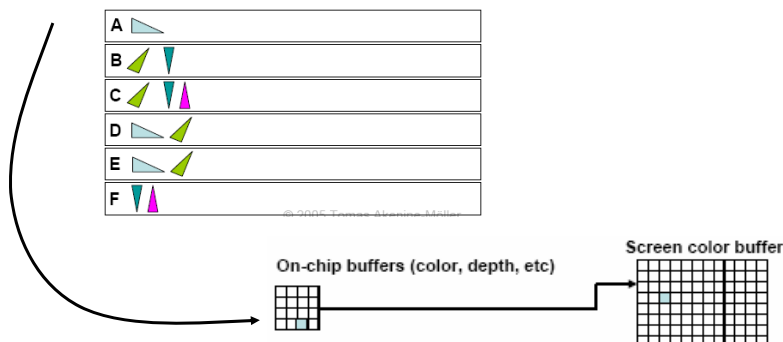
- Create a triangle list for each tile
 - Holds pointers to all triangles overlapping a tile



© 2005 Tomas Akenine-Möller
Luca Benini ARTIST2 / UNU IIST 2007


Tiled processing

- Process one tile at a time, and rasterize triangles in list
- Work on local (on-chip) tile buffers
 - Color, depth, stencil
- Copy color tile buffer to off-chip display buffer
 - may need to copy depth buffer as well



Luca Benini ARTIST2 / UNU IIST 2007

P, TSP



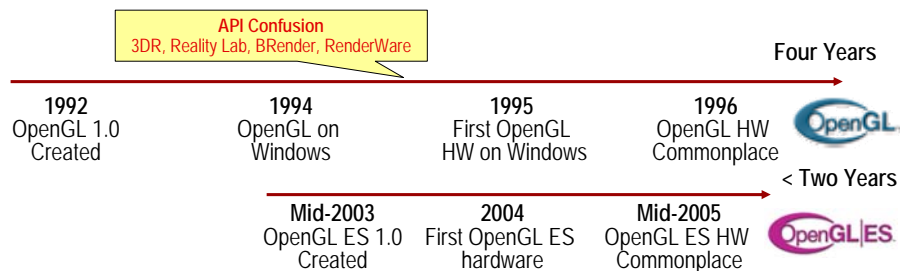
- CPU sends triangle data to MBX
- **Tile Accelerator (TA)**: sorts triangles, and creates a list of triangle pointers for each tile
 - Needs an entire scene before ISP and TSP blocks can start
 - So TA works on the next image, while ISP and TSP work on the current image (i.e., they work in a pipelined fashion)
- **Image synthesis processor (ISP)**: implements Z-buffer, color buffer, stencil buffer for tile
 - Depth testing: test 32 pixels at a time against Z-buffer
 - Records which pixels are visible
 - Groups pixels with same texture and sends to TSP
 - These are guaranteed to be visible, so we only texture each pixel once (deferred texturing)
- **Texture and Shading Processor (TSP)**: Handles texturing and shading interpolation
 - Uses texture compression
 - Performs over-sampling



Luca Benini ARTIST2 / UNU IIST 2007

Mobile 3D API


- The Mobile 3D industry is embryonic - and moving fast!
- **We are where PC graphics were in 1996 - but evolving 2-3 times faster!**
 - Just nine months since OpenGL ES 1.0 released
 - Compliant graphics acceleration already on the market
- OpenGL ES has become the industry standard for embedded graphics
 - We avoided two years of API indecision that occurred on the PC



Luca Benini ARTIST2 / UNU IIST 2007

The case for a higher abstraction


- A game is much more than just 3D rendering
 - Objects, properties, relations (scene graph)
 - Key frame and other animations
 - Etc. (game logic, sounds, ...)
- If everything else but rendering is in Java
 - A very large percentage of the processing is in slow Java
 - Even if rendering was 100% in HW, total acceleration remains limited
- A higher level API could help
 - More of the functionality could be implemented in native (=faster) code
 - Only the game logic must remain in Java
- M3G (JSR-184), a new API
 - Nodes and scene graph
 - Extensive animation support
 - Binary file format and loader




Luca Benini ARTIST2 / UNU IIST 2007

Freescale iMX31

- System
- CPU: A
- VFP - V
- Co-pro
- Image I
- MPEG-
- HW Po
- DVFS
- GPU: M
- Power

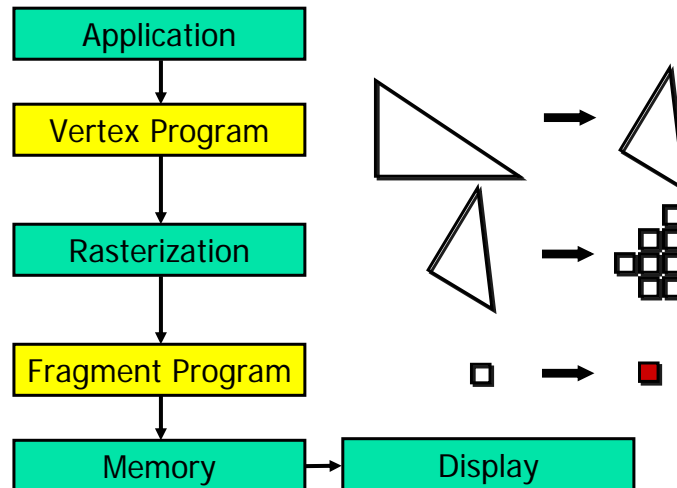


ients



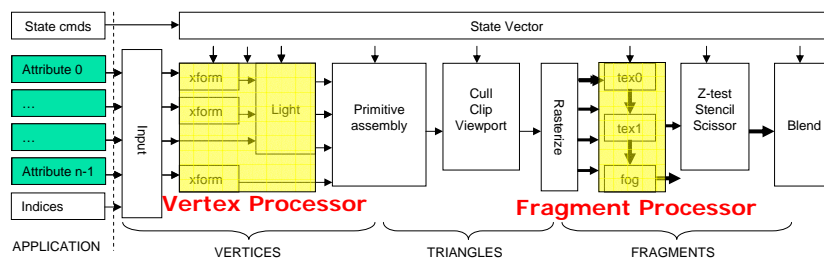
Luca Benini ARTIST2 / UNU IIST 2007

Programmable GPU Model



Luca Benini ARTIST2 / UNU IIST 2007

The OpenGL ES 2.0 Pipeline

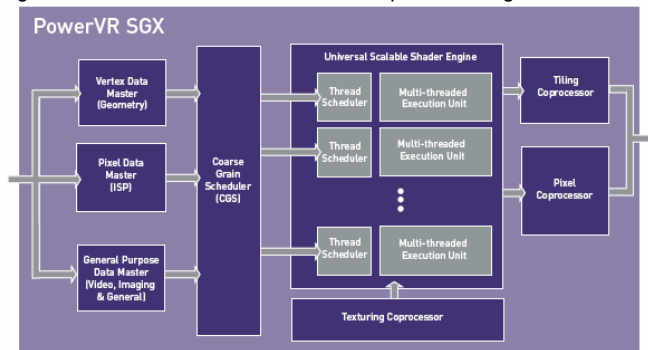


- What Changes From ES 1.1 to ES 2.0?
 - General-purpose attributes replace fixed input arrays
 - Vertex shader programs replace transform and lighting
 - General-purpose uniforms replace fixed lighting & texture state
 - General-purpose varyings replace fixed fragment attributes
 - Fragment shader programs replace texture / fog / alpha test

Luca Benini ARTIST2 / UNU IIST 2007


PowerVR SGX

- Advanced shader-based GPU (OpenGL ES 2.0 compliant)
 - USSE: scalable programmable, multi-threaded engine for graphics, video, imaging and other mathematically-intensive tasks.
 - Tasks are automatically broken down into processing packets which are then scheduled across a number of multi-threaded execution units
 - Coprocessors (texture, pixel and tiling accelerators) assist the MT EUs
 - Latency tolerant architecture
 - geometry and rasterisation are decoupled using tile-based rendering, enabling on-chip processing hidden-surface removal and deferred pixel shading



A high-end system: PS3

A look into the future...




Luca Benini ARTIST2 / UNU IIST 2007

Cell: Single-Chip Embedded Multiprocessor

Highlights (3.2 GHz)

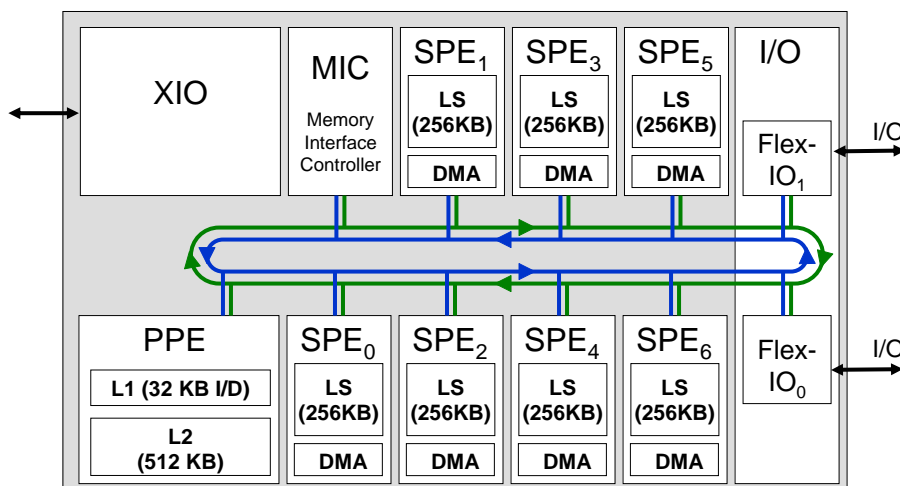
- 241M transistors
- 235mm²
- 9 cores, 10 threads
- >200 GFlops (SP)
- >20 GFlops (DP)
- Up to 25 GB/s memory B/W
- Up to 75 GB/s I/O B/W
- >300 GB/s EIB
- Top frequency >4GHz (observed in lab)



Toshiba

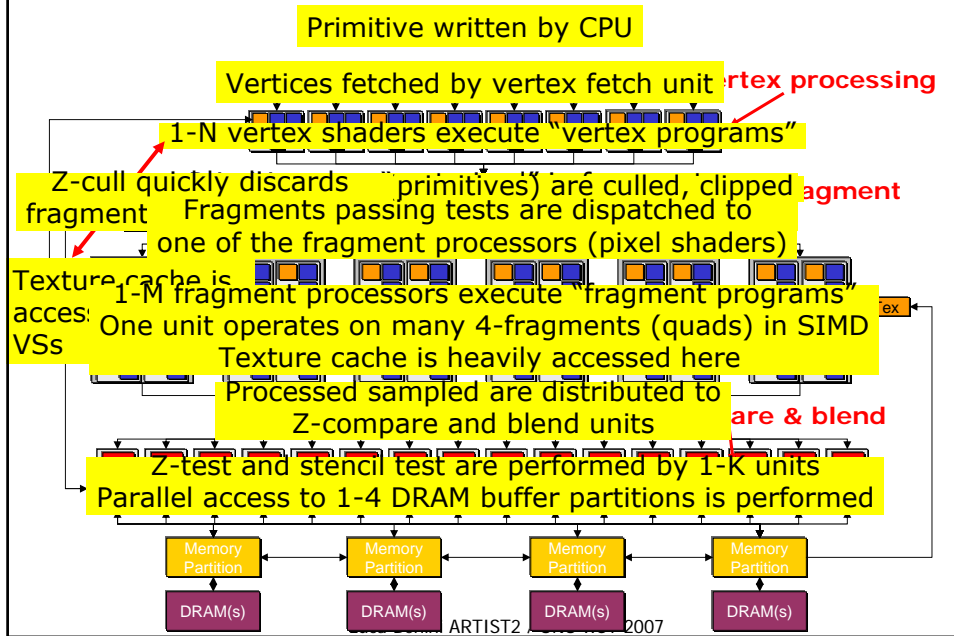
Luca Benini ARTIST2 / UNU IIST 2007

Cell Architecture




Luca Benini ARTIST2 / UNU IIST 2007

NVidia GeForce 7800 Architecture



GF8800 replaces the pipeline model

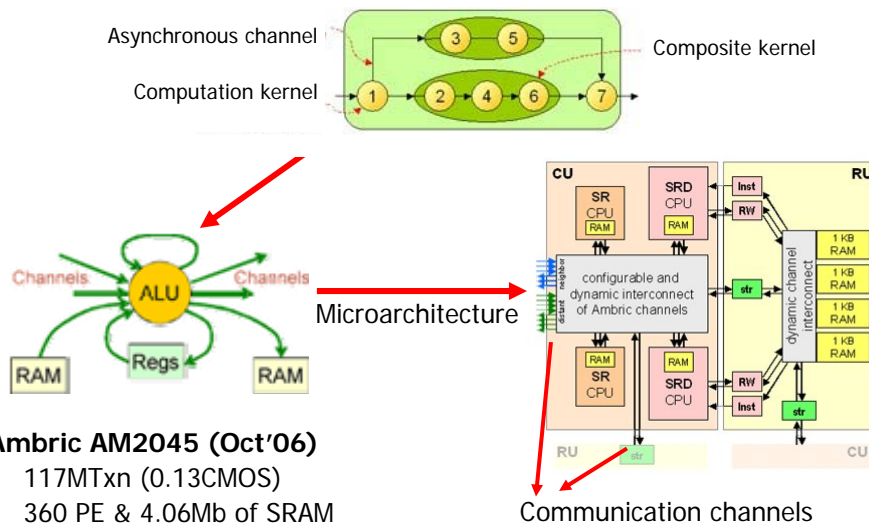
- The future of GPUs is programmable processing
- So – build the architecture around the processor



Case study: Polycores & NoCs

Luca Benini ARTIST2 / UNU IIST 2007

Embedded SoC Architecture Trends "Distributed" stream processors



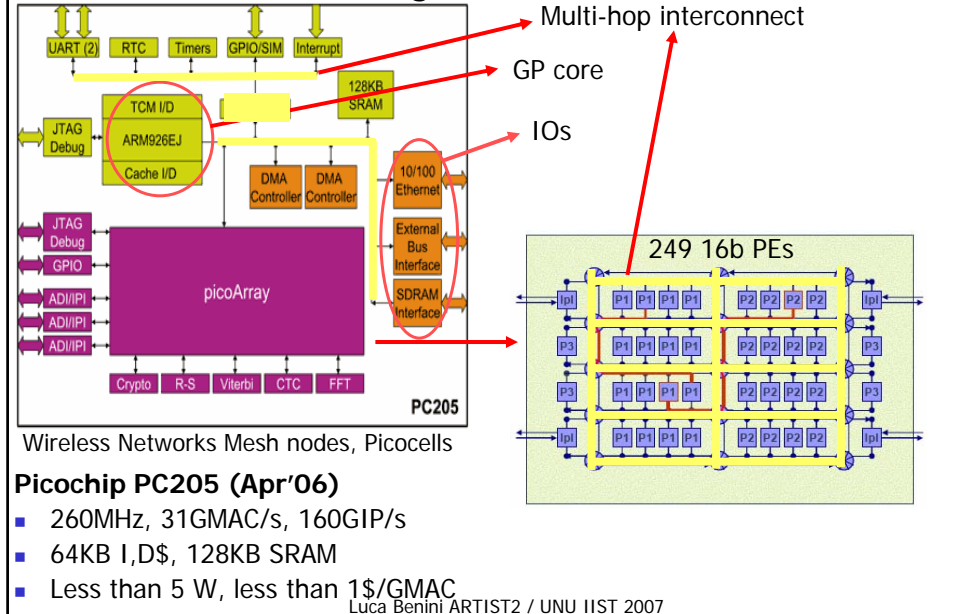
Ambric AM2045 (Oct'06)

- 117MTxn (0.13CMOS)
- 360 PE & 4.06Mb of SRAM
- 1.08 TOPS @ 333MHz (peak)
- 14 Watts → 77GOPS/Watt

Luca Benini ARTIST2 / UNU IIST 2007

Embedded SoC Architecture Trends

Heterogeneous clusters

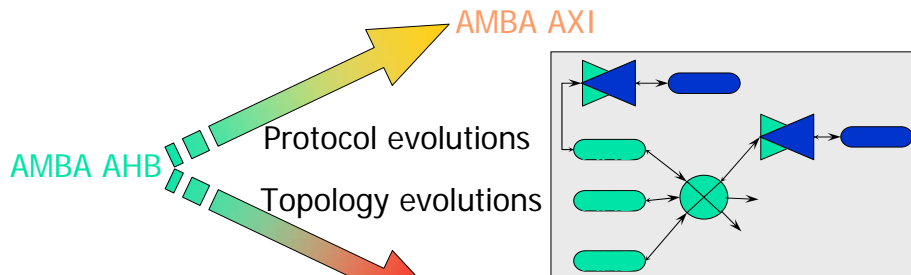


Vision: What Do We Need?

- Scalable
 - Don't want to change the way I design architecture even if requirements scale up exponentially
- Predictable
 - I want to know what to expect (latency, bandwidth), and I want to be able to negotiate it
- Robust
 - Keeps going and going... Even if something is broken inside
- Efficient
 - Silicon is expensive, power is precious
- Easy
 - To create, update, analyze, verify

Addressing Interconnect Issues

- High-end industrial solutions:
 - Evolutionary path from shared busses




- **Challenges**
 - **Complexity** (e.g. 4-SHB + 2XBar, 75 actors): how to analyze and verify "spaghetti interconnects"?
 - **Scalability**: bus is bandwidth-limited, Xbar is size-limited
 - **Predictability**: how to tie interconnects with floorplanning

Luca Benini ARTIST2 / UNU IIST 2007

The Network-on-Chip Paradigm

The "power of NoCs":

- **Clean separation** at session layer
 - Cores issue end-to-end transactions
 - Network deals with transport, network, link, physical
- **Modularity** at HW level: only 2 building blocks
 - Network interface
 - Switch (router)
- **Physical design aware** (floorplan global routing)




Scalability is supported from the ground up!

Luca Benini ARTIST2 / UNU IIST 2007

Building blocks: NI

- Session-layer interface with nodes
- Back-end manages interface with switches



Standardized node interface @ session layer.
Initiator vs. target distinction is blurred

1. Supported transactions (e.g. QoSread...)
2. Degree of parallelism
3. Session prot. control flow & negotiation

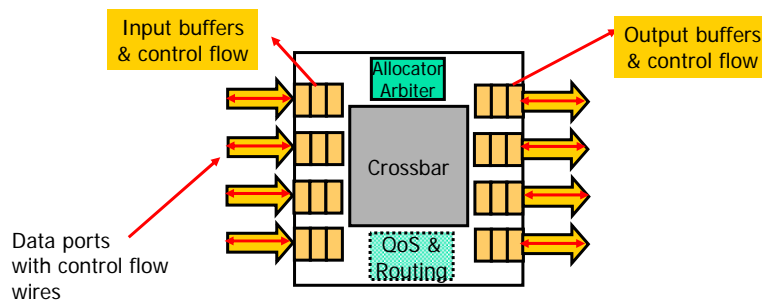
NoC specific backend (layers 1-4)

1. Physical channel interface
2. Link-level protocol
3. Network-layer (packetization)
4. Transport layer (routing)

Luca Benini ARTIST2 / UNU IIST 2007

Building blocks: Switch

- Router: receives and forwards packets
 - NOTE: Packet-based does not mean datagram!
- Level 3 or Level 4 routing
 - No consensus, but generally L4 support is limited (e.g. simple routing)




Luca Benini ARTIST2 / UNU IIST 2007

Æthereal: context

- **Consumer electronics**
 - reliability & predictability are essential
 - low cost is crucial
 - time to market must be reduced
- NoC offer **differentiated services**
 - to manage (and hence reduce) resources
 - to ease integration (and hence decrease TTM)

Luca Benini ARTIST2 / UNU IIST 2007

NoC services




- Request communication services using **connections**
 - opening & closing affect resource reservations **commitment**
 - With **properties**
 - data integrity (uncorrupted data transfer)
 - transaction ordering
 - un/ordered per slave/connection
 - transaction completion
 - flow control
 - data loss or not
 - delivery bounds
 - throughput, latency, jitter
- correctness
completion
bounds

Luca Benini ARTIST2 / UNU IIST 2007

Æthereal: features

- Conceptually, **two disjoint networks**
 - a network with throughput+latency guarantees (GT)
 - a network without those guarantees (best effort, BE)
- Several types of commitment in the network
 - **combine guaranteed worst case behaviour with good average resource usage**



Luca Benini ARTIST2 / UNU IIST 2007

Router architecture

- Best-effort router
 - Worm hole routing
 - Input queueing
 - Source routing
- Guaranteed throughput router
 - Contention-free routing
 - synchronous, using slot tables
 - time-division multiplexed circuits
 - Store and forward routing
 - Headerless packets
 - information is present in slot table

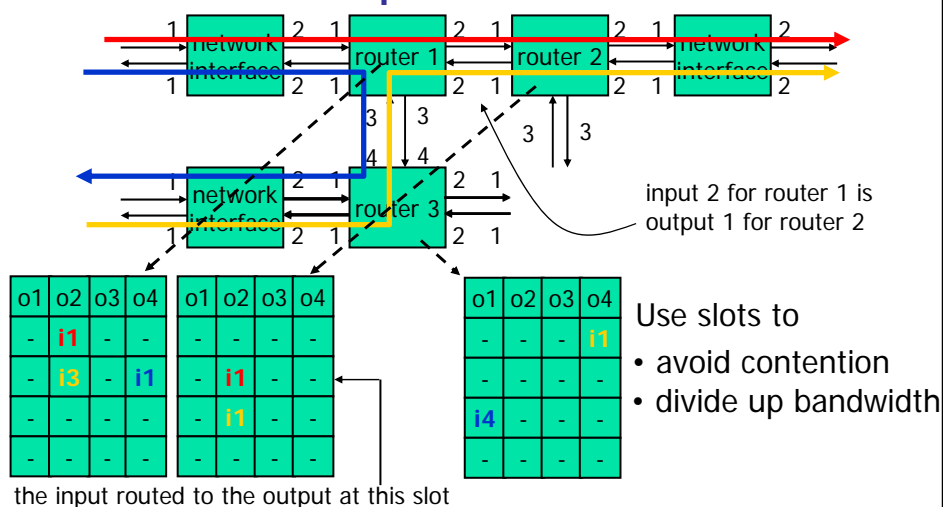
Luca Benini ARTIST2 / UNU IIST 2007

Contention-free routing

- Latency guarantees are easy in circuit switching
- Emulate circuits with packet switching
- **Schedule packet injection** in network such that they never contend for same link at same time
 - in space: disjoint paths
 - in time: time-division multiplexing
 - or a combination

Luca Benini ARTIST2 / UNU IIST 2007

CFR Example



Luca Benini ARTIST2 / UNU IIST 2007


CFR setup

- Use **best-effort packets to set up connections**
 - set up & tear down packets like in ATM (asynchronous transfer mode)
- Distributed, concurrent, pipelined
- Safe: always consistent
- Compute slot assignment compile time, run time, or combination
- **Connection opening is guaranteed to complete** (but without a latency guarantee) with commitment or rejection

Luca Benini ARTIST2 / UNU IIST 2007

Router implementation

- Memories (for packet storage)
 - Register based FIFOs are expensive
 - RAM based FIFOs are as expensive
 - 80% of router is memory
 - Special hardware FIFOs are very useful
 - 20% of router is memory
- Speed of memories
 - registers are fast enough
 - RAMs may be too slow
 - Hardware FIFOs are fast enough

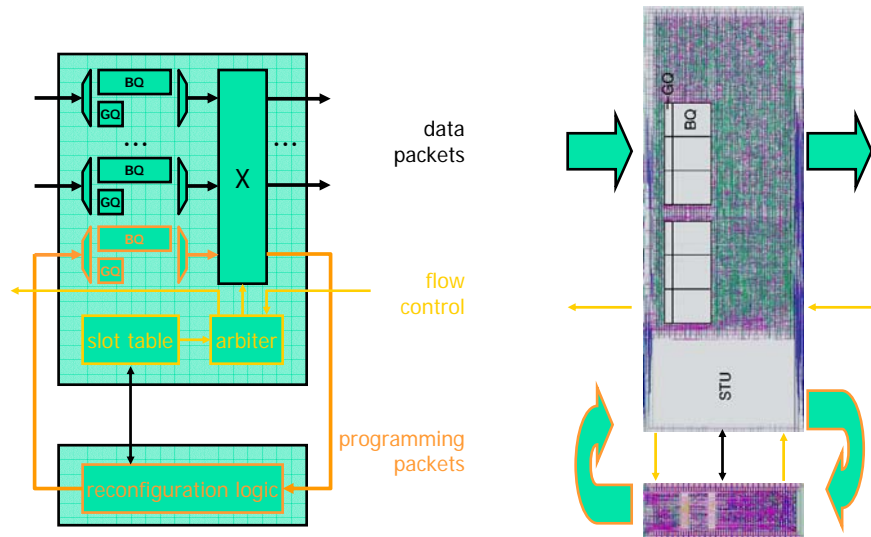


routers based on register-file and hardware fifos drawn to approximately same scale (1mm², 0.26mm²)




Luca Benini ARTIST2 / UNU IIST 2007

Layout



Results


- 5 input and 5 output ports (arity 5)
- 0.25 mm² CMOS12
- 500 MHz data path, 166 MHz control path
- flit size of 3 words of 32 bits
- $500 \times 32 = 16$ Gb/s throughput per link, in each direction
- 256 slots & 5x1 flit fifos for guaranteed throughput traffic
- 6x8 flit fifos for best effort traffic



Luca Benini ARTIST2 / UNU IIST 2007

xpipes: context


- Typical applications targeted by SoCs
 - Complex
 - Highly heterogeneous (component specialization)
 - Communication intensive
- xpipes is a synthesizable, heterogeneous NoC infrastructure
- Three year lifetime, mature research project
 - University of Bologna (architecture)
 - Stanford University (design technology)
 - University of Cagliari (design and backend)**



Heterogeneous topology

SoC *component specialization* leads to the integration of *heterogeneous cores*

Ex. MPEG4 Decoder



- Non-uniform block sizes
- SDRAM: communication bottleneck
- Many neighboring cores do not communicate


On a homogeneous fabric:

- Risk of under-utilizing many tiles and links
- Risk of localized congestion

Luca Benini ARTIST2 / UNU IIST 2007


Example: MPEG4 decoder

- Core graph representation with annotated average communication requirements




Luca Benini ARTIST2 / UNU IIST 2007


NoC Floorplans



General purpose: mesh



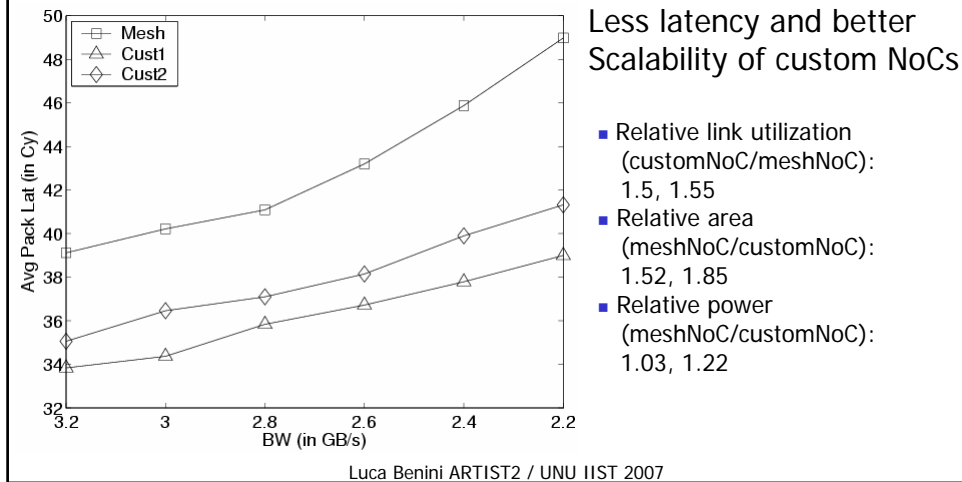
Application Specific NoC1 (centralized)



Application Specific NoC2 (distributed)

Luca Benini ARTIST2 / UNU IIST 2007

Performance, area and power



Xpipes: features

- Source based routing
 - Very high performance switch design
- Wormhole switching
 - Minimize buffering area while reducing latency
- Pipelined links
 - Link data introduction interval is not bound by wire delay
 - Link-latency (# of repeater stages) insensitive operation
- Parameterizable network building blocks
 - Plug-and-play composable for arbitrary network topology
 - Design time tunable buffer size, link width, virtual channels, # of switch I/Os
- Standard OCP interface

Luca Benini ARTIST2 / UNU IIST 2007

Link delay bottleneck

- Wire delay is serious concern for NoC Links

- If NoC "beat" is determined by worst case link delay, performance can be severely limited

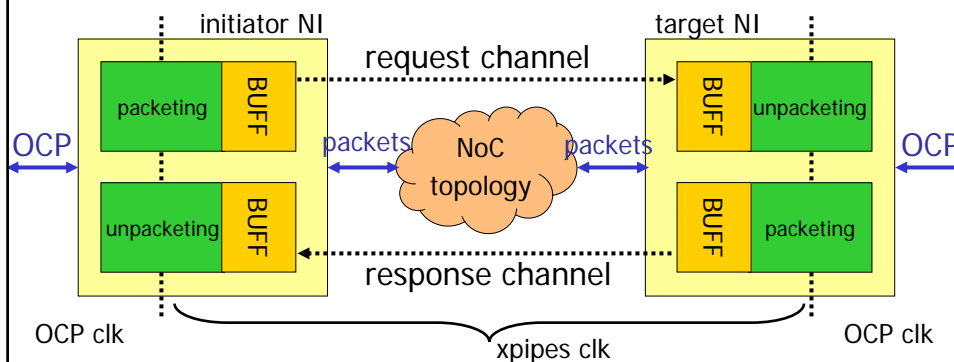
- ↳ Pipeline links

- Delay is transformed in Latency
 - Data introduction speed is not bound by link delay any longer!



Luca Benini ARTIST2 / UNU IIST 2007

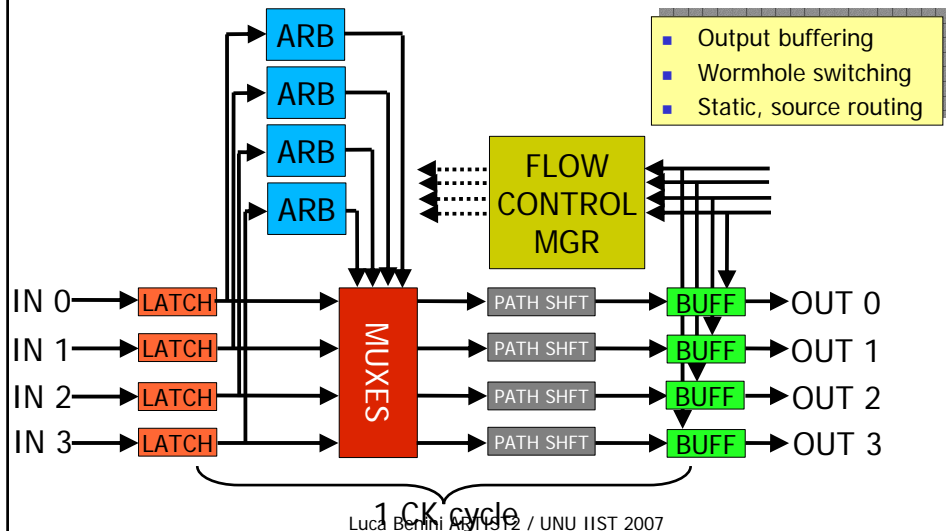
xpipes Architecture: the Network Interface



- OCP 2.0 protocol to connect to IP cores
- Performs packeting/unpacketing
- Handles routing via path lookup tables
- Dual clock operation

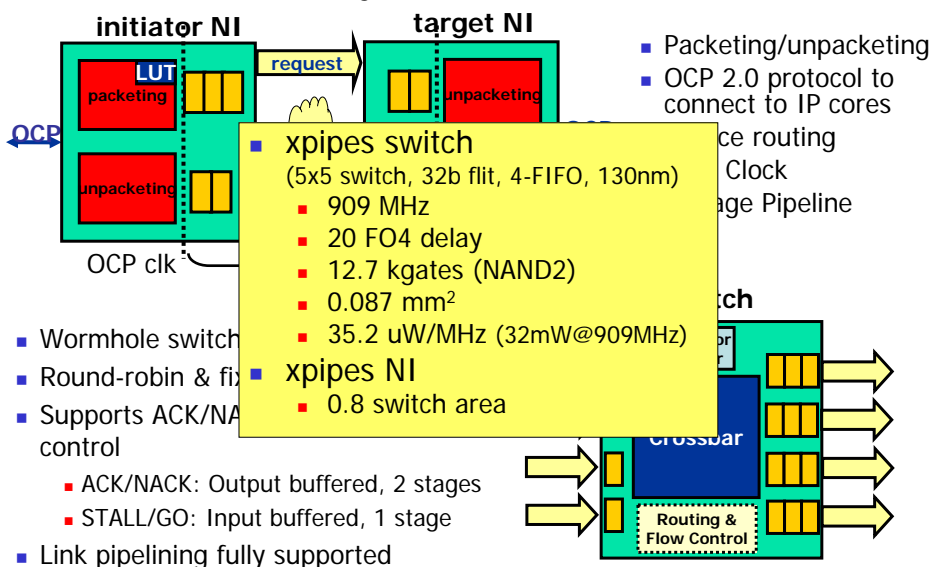
ST2 / UNU IIST 2007

xpipes Architecture: the Switch



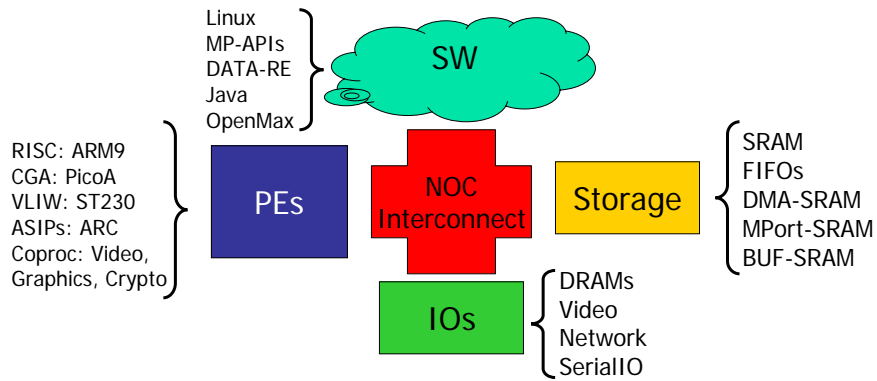
The xpipes NoC

- A soft macro library:



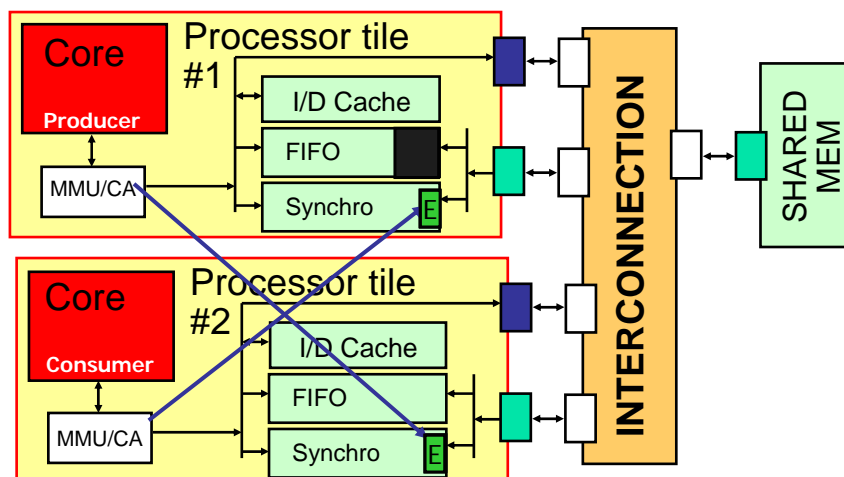
Communication-Centric Platforms

- A holistic approach to MPSoC architectural design (HW & SW) is needed!



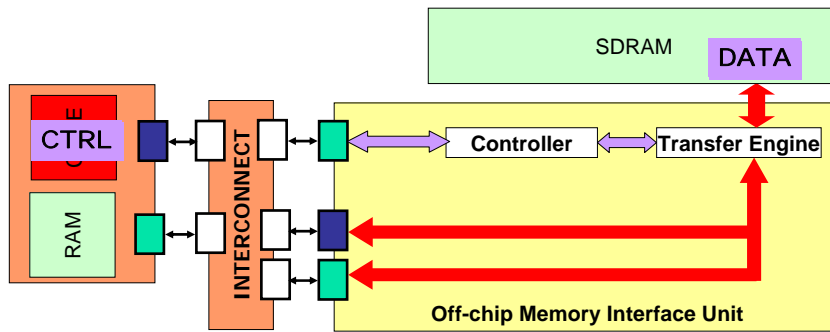
Luca Benini ARTIST2 / UNU IIST 2007

"NoC-friendly" stream processors



Luca Benini ARTIST2 / UNU IIST 2007

External Memory Bottleneck



- “Pull” memory channel
 - *Control Block* keeps programmable table of objects to be moved
 - Table entries can be programmed by different cores
 - *Transfer Engine* shuffles data among bus and Memory Controller
 - Triggers bus or SDRAM transactions
 - *Memory Controller* handles SDRAM accesses

Luca Benini ARTIST2 / UNU IIST 2007

Summary

- Why SoCs?
- SoC Platforms
- From SoC to MPSoC
- From MPSoC to NoC

Luca Benini ARTIST2 / UNU IIST 2007