

The Cell Processor

Computing of tomorrow or yesterday?

Torsten Höfler
Department of Computer Science
TU Chemnitz

24.05.2005



CHEMNITZ UNIVERSITY
OF TECHNOLOGY

Outline

- 1 The Vision
- 2 The Cell Architecture
 - Introduction
 - The PowerPC Processing Element (PPE)
 - The Synergistic Processing Element (SPE)
 - The Element Interconnect Bus (EIB)
 - The I/O Interconnect - FlexIO
 - The Memory Interface Controller (MIC)
- 3 Cell Programming
- 4 Summary & Conclusion



The Researchers

- started in mid 2000 by Sony Toshiba and IBM
- Sony has PS2 arch - needs chip for PS3
- Toshiba has memory experience - needs Chips for HDTV
- IBM has technical knowledge in processor manufacturing
- billions of dollars have been invested
- ⇒ high throughput multi purpose processor



Known Problems to Overcome

- memory latency gap
- instruction latency gap (control logic)
- old fashioned x86 arch
- hardware architecture other than ISA
- implicit mem hierarchy (caching)



Main Changes compared to x86

- RISC design
- in order CPU (no OOO)
- SoC design (different cores)
- heterogeneous Multiprocessing




Outline

- 1 The Vision
- 2 **The Cell Architecture**
 - Introduction
 - The PowerPC Processing Element (PPE)
 - The Synergistic Processing Element (SPE)
 - The Element Interconnect Bus (EIB)
 - The I/O Interconnect - FlexIO
 - The Memory Interface Controller (MIC)
- 3 Cell Programming
- 4 Summary & Conclusion




1st Patent - Ken Kutaragi (Sony,1999)


- cell = software or hardware cell
- software cell = program + data
- hardware cell = execution logic + local memory
- hardware cells process software cells
- no fixed architecture (network distribution)
- cell computer is created ad-hoc (multiple devices)



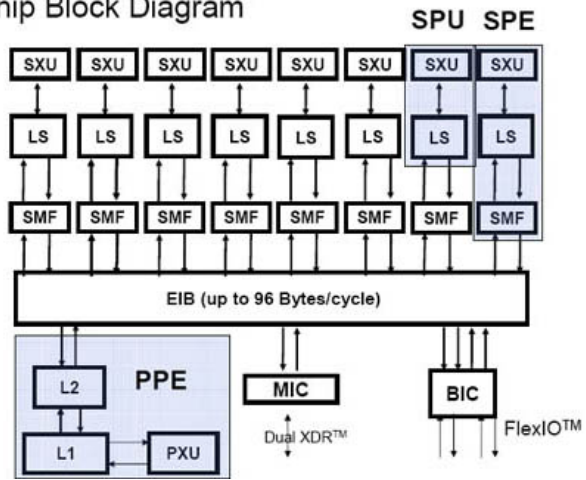
- Heterogeneous Multiprocessing - 9 Core Processor
- SPU - Synergistic Processing Unit
- MFC - Memory Flow Controller
- SPE - Synergistic Processing Element (=SPU+MFC)
- PPE - PowerPC Processing Element
- 1PPE + 8 SPEs = Cell Processor
- 9 full blown CPUs



- connected via EIB - Element Interconnection Bus
- FlexIO - IO/Processor Interconnect (BIC)
- MIC - dual port XDR Memory Interface
- hardware DRM :-)
- interrupts controller routes only to PPE
- effective NoC (Network on Chip)
- ⇒ single precision 256 GFlops peak




Cell Chip Block Diagram



source & copyright: IBM

First Prototype


- 90nm SOI, 8 copper layers
- 234 Mill. Transistors
- 60-80 W (prototype)
- only 6-7 SPEs enabled (manufacturing errors)
- IBMs virtualization technology
- 1.1V, 4GHz

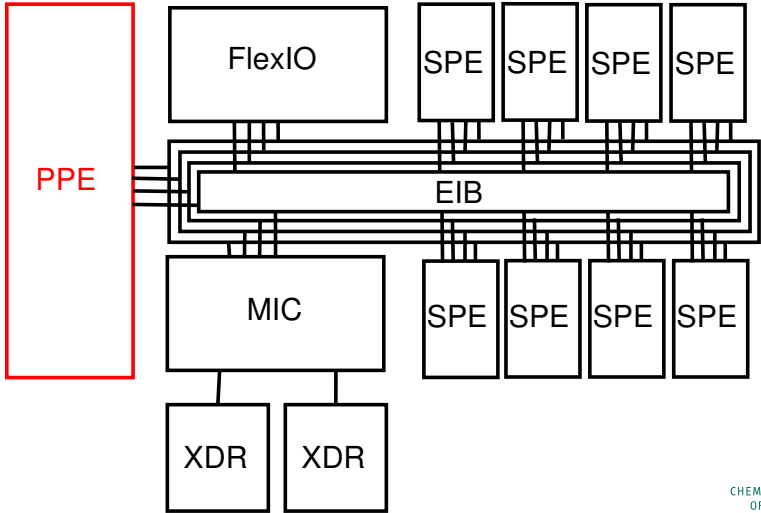


Outline

- 1 The Vision
- 2 **The Cell Architecture**
 - Introduction
 - **The PowerPC Processing Element (PPE)**
 - The Synergistic Processing Element (SPE)
 - The Element Interconnect Bus (EIB)
 - The I/O Interconnect - FlexIO
 - The Memory Interface Controller (MIC)
- 3 Cell Programming
- 4 Summary & Conclusion







- dual-threaded (SMT) 64 bit Power Architecture
- includes VMX (aka AltiVec) ISA
- simple architecture, only in-order execution
- super-scalar with deep 2-way pipeline (>20)
- 2 instructions issued per cycle
- 32kB+32kB L1, 512 kB L2 Cache
- supports virtualization \rightarrow logical partitioning (memory, I/O, time)
- simplified Power Architecture




Outline

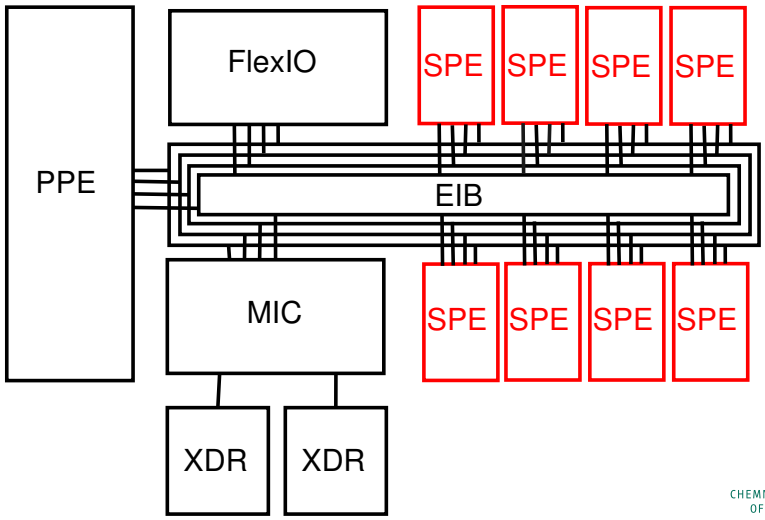
- 1 The Vision
- 2 **The Cell Architecture**
 - Introduction
 - The PowerPC Processing Element (PPE)
 - **The Synergistic Processing Element (SPE)**
 - The Element Interconnect Bus (EIB)
 - The I/O Interconnect - FlexIO
 - The Memory Interface Controller (MIC)
- 3 Cell Programming
- 4 Summary & Conclusion



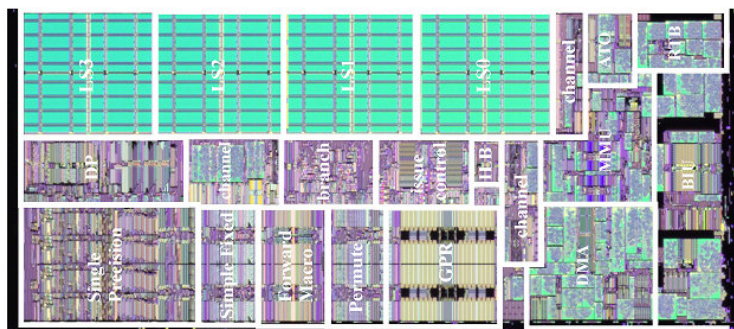
The Vision
The Cell Architecture
Cell Programming
Summary & Conclusion

Introduction
The PowerPC Processing Element (PPE)
The Synergistic Processing Element (SPE)
The Element Interconnect Bus (EIB)
The I/O Interconnect - FlexIO
The Memory Interface Controller (MIC)



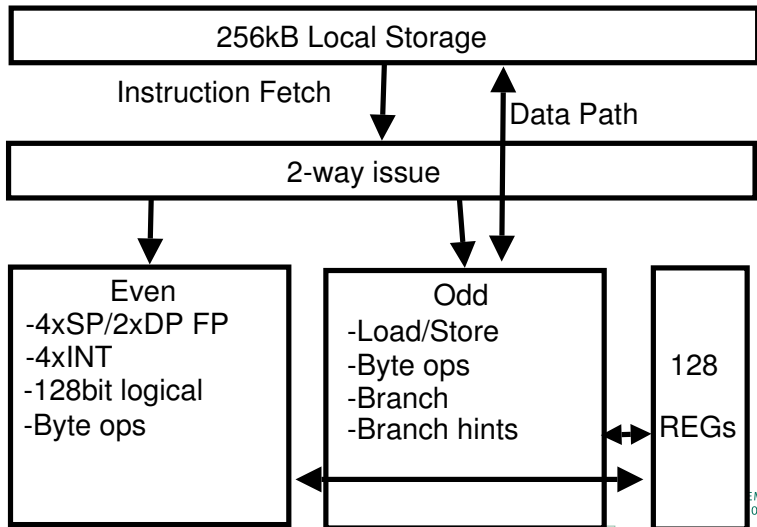


- fully blown vector CPUs with own RAM
- ISA: not VMX compatible!
- ISA: 32 bit fixed length
- 21 Mill transistors (14 Mill SRAM, 7 Mill logic)



- no branch prediction or scheduling logic (software)
- two independent short and simple pipes
- can issue two instructions in parallel
- one memory and one SIMD computation
- strictly in order
- instructions work with 128 bit compound data
- 4 SP FP units (not fully IEEE754 compliant - 32GFlops)
- slow DP arithmetic (fully IEEE754 - 3-4GFlops)
- 4 INT units (32 GFlops)

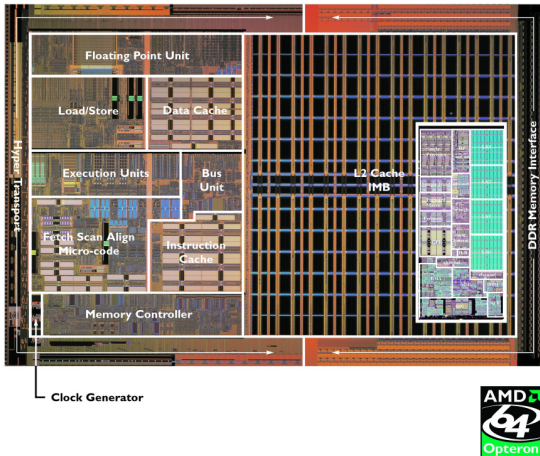




- 256kB local storage (LS) memory
- accessible in 128bit lines
- 128 128bit registers (2 cycles latency)
- registers are layered (hold all data types)
- no virtual memory, no coherency
- no processing in main memory
- DMA to move data between LS and main memory
- MFC connects to EIB, acts like MMU + synchronisation



Opteron: 3GHz, 3 FP units - SPE: 4GHz, 4 SP FP units




Outline


- 1 The Vision
- 2 The Cell Architecture**
 - Introduction
 - The PowerPC Processing Element (PPE)
 - The Synergistic Processing Element (SPE)
 - The Element Interconnect Bus (EIB)**
 - The I/O Interconnect - FlexIO
 - The Memory Interface Controller (MIC)
- 3 Cell Programming
- 4 Summary & Conclusion



The Vision
The Cell Architecture
Cell Programming
Summary & Conclusion

Introduction
The PowerPC Processing Element (PPE)
The Synergistic Processing Element (SPE)
The Element Interconnect Bus (EIB)
The I/O Interconnect - FlexIO
The Memory Interface Controller (MIC)





- four 128 bit wide concentric rings
- optimized for 1024 bit blocks
- 96 byte/cycle
- buffered point-to-point ring (cmp. SCI)
- SPE buffers and routes
- scalable (more SPEs increase latency)
- guaranteed bandwidth of $1/\#$ devices
→ Real Time capable




Outline

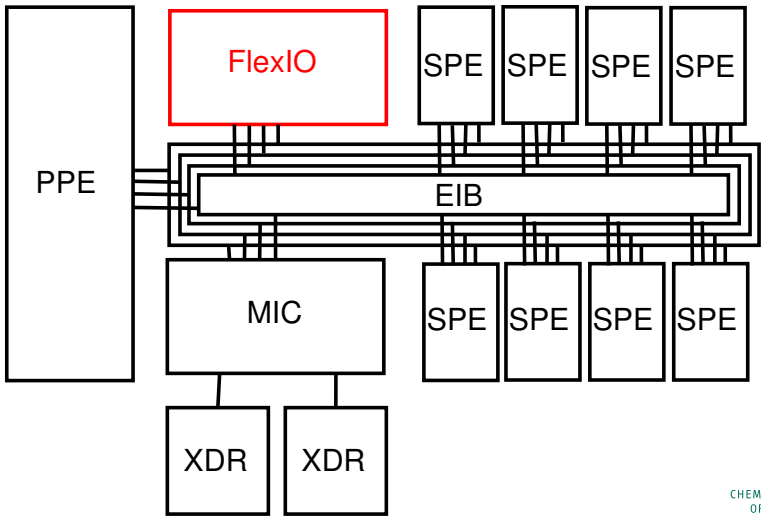
- 1 The Vision
- 2 **The Cell Architecture**
 - Introduction
 - The PowerPC Processing Element (PPE)
 - The Synergistic Processing Element (SPE)
 - The Element Interconnect Bus (EIB)
 - **The I/O Interconnect - FlexIO**
 - The Memory Interface Controller (MIC)
- 3 Cell Programming
- 4 Summary & Conclusion



The Vision
The Cell Architecture
Cell Programming
Summary & Conclusion

Introduction
The PowerPC Processing Element (PPE)
The Synergistic Processing Element (SPE)
The Element Interconnect Bus (EIB)
The I/O Interconnect - FlexIO
The Memory Interface Controller (MIC)





- 12 unidirectional byte-wide lanes
- 96 pairs in the whole
- 6.4 Gb/s per lane (76.8 Gb/s in the whole)
- 7 lanes (44.8 GB/s) out, 5 lanes (32 GB/s) in
- coherent (cc-NUMA) and non coherent links
- also used as processor interconnect (cmp. HT)
- connect two processors glueless
- needs switch for more




Outline

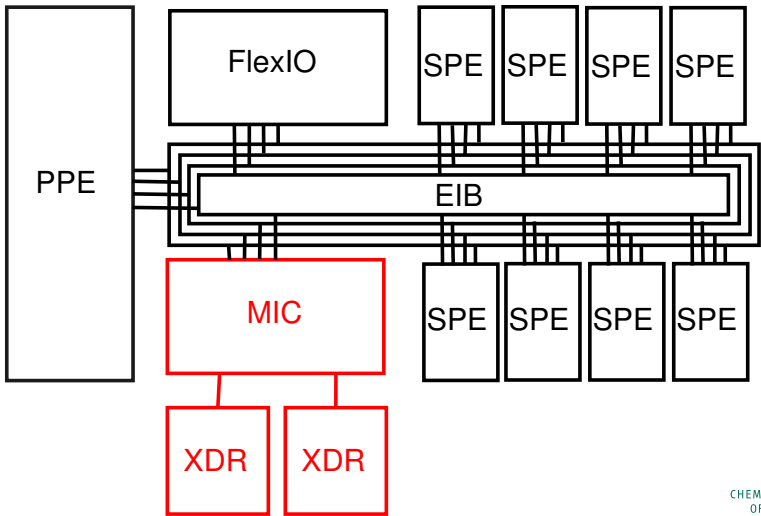
- 1 The Vision
- 2 **The Cell Architecture**
 - Introduction
 - The PowerPC Processing Element (PPE)
 - The Synergistic Processing Element (SPE)
 - The Element Interconnect Bus (EIB)
 - The I/O Interconnect - FlexIO
 - **The Memory Interface Controller (MIC)**
- 3 Cell Programming
- 4 Summary & Conclusion



The Vision
The Cell Architecture
Cell Programming
Summary & Conclusion

Introduction
The PowerPC Processing Element (PPE)
The Synergistic Processing Element (SPE)
The Element Interconnect Bus (EIB)
The I/O Interconnect - FlexIO
The Memory Interface Controller (MIC)





- two channel Rambus XDR Memory - 25.2 GB/s
- ECC protected (why?)
- PPE/SPE have conventional protection system to access main memory (cmp. MMU)
- PPE/SPE use virtual addresses
- no cache → moved to software



- no abstraction layer like x86
- difficult programming and optimization
- same problems as for Multi-Core or SMPs
- direct programming of SPEs, 256kb storage + 128 Regs
- no need for assembly - but it is available :)
- programmable in C/C++
- SPEs are allocated in software
- SPE virtualization by OS? (more SPE tasks as SPEs)
- PPE code must only be PPC970
→ Linux already running
- SPE code must be self-contained
- autoparallelizing compiler from IBM: Octopile




Programming Models:

- job queue
- self multitasking SPE
- Stream Processing
- Software Managed Cache
- MPI?



Let the Cells glow!



- general purpose Power CPU with 8 vector processors (SPEs)
- RISC approach - moving functionality to software
→ more complicated compiler (cmp. VLIW)
- very fast I/O System coupled with task distribution
- 560 Cells to get to the top of the top500 (theoretically)
- Realtime capabilities because of independent SPE
- intended to be the standard of tomorrow
- (hopefully) integrated in PS3 and Toshiba TVs (hdtv) in 2006



Resources / Additional Information:

- <http://www.research.scea.com/research/html/CellGDC05/>
- <http://www.blachford.info/computer/Cells/Cell0.html>
- <http://www.research.ibm.com/cell/>
- <http://www.cell-industries.com/>
- <http://www-306.ibm.com/chips/techlib/techlib.nsf/products/Cell>
- <http://www-128.ibm.com/developerworks/power/cell/>

