



Systems and Technology Group

# Power Efficient Processor Design and the Cell Processor

H. Peter Hofstee, Ph. D.

[hofstee@us.ibm.com](mailto:hofstee@us.ibm.com)

Architect, Cell Synergistic Processor Element

IBM Systems and Technology Group

Austin, Texas

# Agenda

- **Power Efficient Processor Architecture**
- **System Trends**
- **Cell Processor Overview**

# Power Efficient Architecture


## Limiters to Processor Performance

- **Power wall**
- **Memory wall**
- **Frequency wall**


# Power Wall (Voltage Wall)

## Power components:

- Active power
- Passive power
  - Gate leakage
  - Sub-threshold leakage (source-drain leakage)



Gate dielectric approaching a fundamental limit (a few atomic layers)



**NET: INCREASING PERFORMANCE  
REQUIRES INCREASING EFFICIENCY**

# Memory wall

- **Main memory now nearly 1000 cycles from the processor**
  - Situation worse with (on-chip) SMP
- **Memory latency penalties drive inefficiency in the design**
  - Expensive and sophisticated hardware to try and deal with it
  - Programmers that try to gain control of cache content, but are hindered by the hardware mechanisms
- **Latency induced bandwidth limitations**
  - Much of the bandwidth to memory in systems can only be used speculatively
  - Diminishing returns from added bandwidth on traditional systems

## Frequency wall

- **Increasing frequencies and deeper pipelines have reached diminishing returns on performance**
- **Returns negative if power is taken into account**
- **Results of studies depend on issue width of processor**
  - The wider the processor the slower it wants to be
  - Simultaneous Multithreading helps to use issue slots efficiently
- **Results depend on number of architected registers and workload**
  - More registers tolerate deeper pipeline
  - Fewer random branches in application tolerates deeper pipelines

# Microprocessor Efficiency

- **Recent History:**
  - Gelsinger's law
    - 1.4x more performance for 2x more transistors
  - Hofstee's corollary
    - 1/1.4x efficiency loss in every generation
    - Examples: Cache size, OoO, Superscalar, etc. etc.
  
- **Re-examine microarchitecture with performance per transistor as metric**
  - Pipelining is last clear win




## Attacking the Performance Walls

- **Multi-Core Non-Homogeneous Architecture**
  - Control Plane vs. Data Plane processors
  - Attacks **Power Wall**
- **3-level Model of Memory**
  - Main Memory, Local Store, Registers
  - Attacks **Memory Wall**
- **Large Shared Register File & SW Controlled Branching**
  - Allows deeper pipelines (11FO4 ... helps power!)
  - Attacks **Frequency Wall**

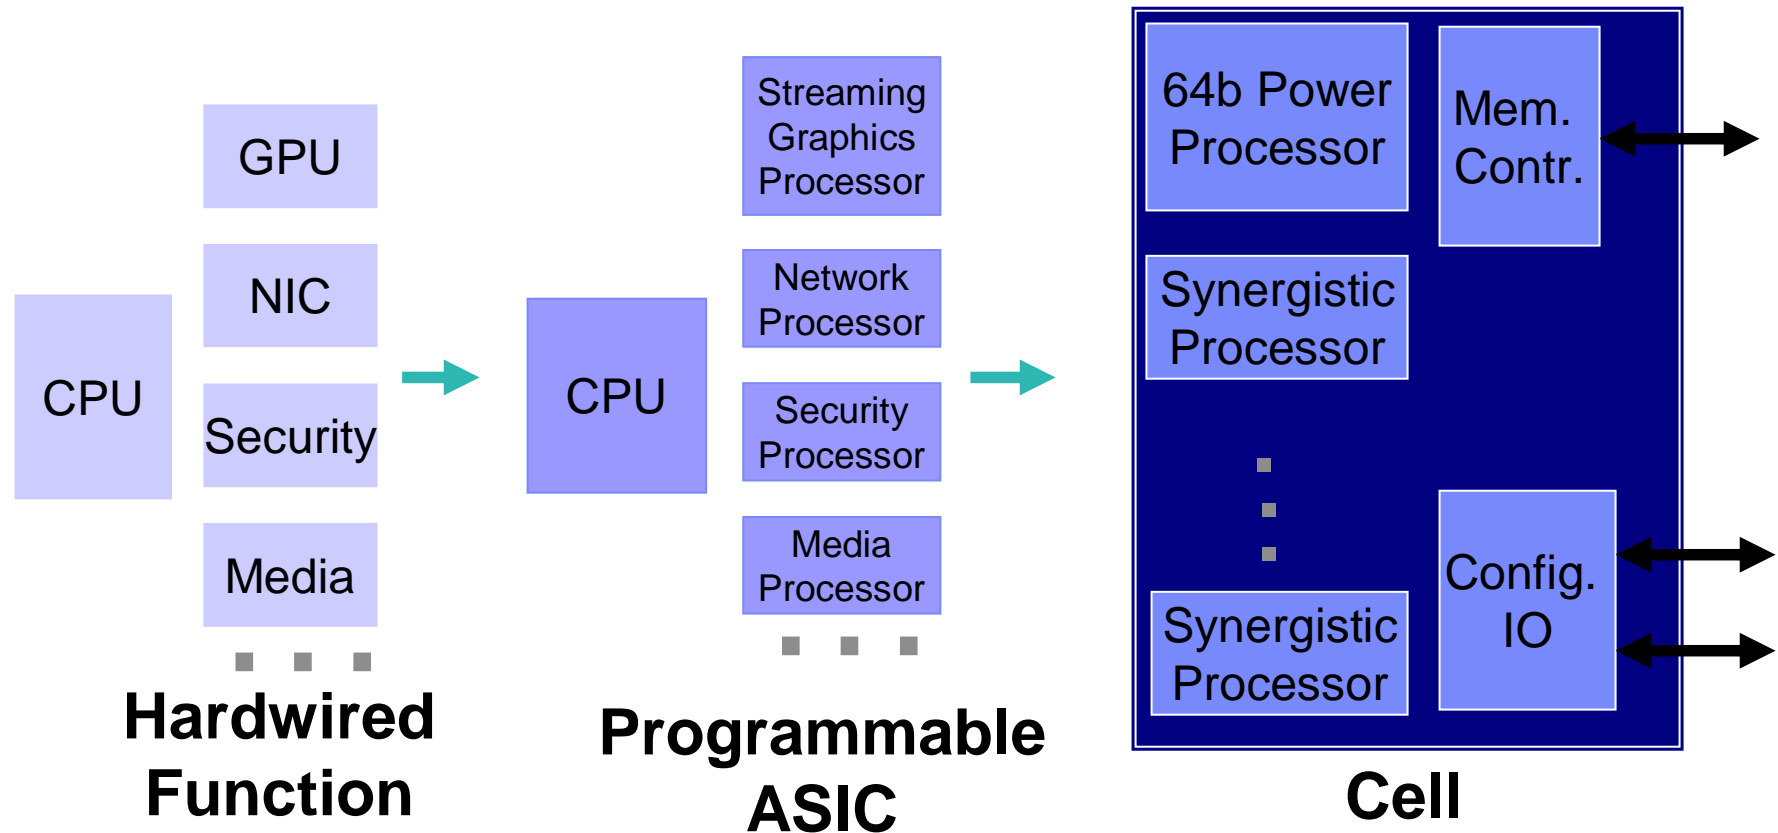
# System Trends

## System Trends toward Integration



- **Increased integration is driving processors to take on many functions typically associated with systems**
  - Integration forces processor developers to address off-load and acceleration in the design of the processor
  - Integration of bridge chip functionality
- **Virtualization technology is used to support non-homogeneous environments**

# Next Generation Processors address Programming Complexity and Trend Towards Programmable Offload Engines with a Simpler System Alternative



## “Outward Facing” Aspects of Cell


- **Cell is designed to be responsive**
- **.. to human user**
  - Real-time response
  - Supports rich visual interfaces
- **.. to network**
  - Flexible, can support new standards
  - High-bandwidth
  - Content protection, privacy & security
- **Contrast to traditional processors which evolved from “batch processing” mentality (inward focused).**

# Cell Overview

# Key Attributes of Cell


- **Cell is Multi-Core**
  - Contains 64-bit Power Architecture <sup>TM</sup>
  - Contains 8 Synergistic Processor Elements (SPE)
- **Cell is a Flexible Architecture**
  - Multi-OS support (including Linux) with Virtualization technology
  - Path for OS, legacy apps, and software development
- **Cell is a Broadband Architecture**
  - SPE is RISC architecture with SIMD organization and Local Store
  - 128+ concurrent transactions to memory per processor
- **Cell is a Real-Time Architecture**
  - Resource allocation (for Bandwidth Measurement)
  - Locking Caches (via Replacement Management Tables)
- **Cell is a Security Enabled Architecture**
  - SPE dynamically reconfigurable as secure processors

# Cell Chip Block Diagram





# Cell Prototype Die (Pham et al, ISSCC 2005)




## Cell Highlights

- Observed clock speed
  - > **4 GHz**
- Peak performance (single precision)
  - > **256 GFlops**
- Peak performance (double precision)
  - >**26 GFlops**
- Area 221 mm<sup>2</sup>
- Technology 90nm SOI
- Total # of transistors 234M


# Element Interconnect Bus

- **EIB data ring for internal communication**
  - Four 16 byte data rings, supporting multiple transfers
  - 96B/cycle peak bandwidth
  - Over 100 outstanding requests




# Power Processor Element

- **PPE handles operating system and control tasks**
  - 64-bit Power Architecture™ with VMX
  - In-order, 2-way hardware Multi-threading
  - Coherent Load/Store with 32KB I & D L1 and 512KB L2




# Synergistic Processor Element

- **SPE provides computational performance**
  - Dual issue, up to 16-way 128-bit SIMD
  - Dedicated resources: 128 128-bit RF, 256KB Local Store
  - Each can be dynamically configured to protect resources
  - Dedicated DMA engine: Up to 16 outstanding request



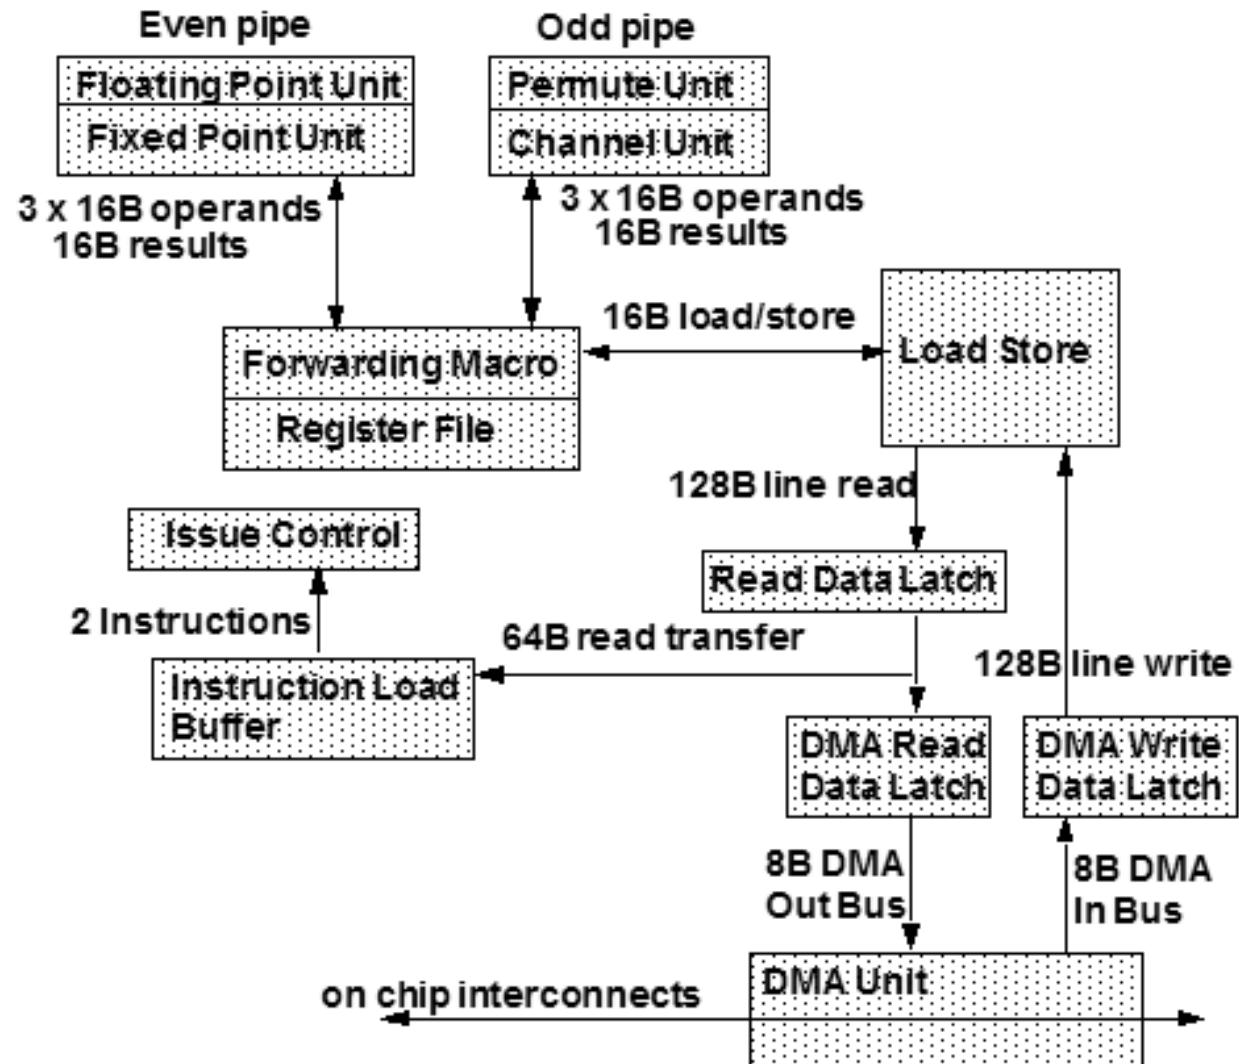
# SPE Highlights




**14.5mm<sup>2</sup> (90nm SOI)**

- **User-mode architecture**
  - No translation/protection within SPU
  - DMA is full Power Arch protect/x-late
- **Direct programmer control**
  - DMA/DMA-list
  - Branch hint
- **VMX-like SIMD dataflow**
  - Broad set of operations
  - Graphics SP-Float
  - IEEE DP-Float (BlueGene-like)
- **Unified register file**
  - 128 entry x 128 bit
- **256kB Local Store**
  - Combined I & D
  - 16B/cycle L/S bandwidth
  - 128B/cycle DMA bandwidth

# SPE Organization (Flachs et al, ISSCC 2005)




# SPE PIPELINE (Flachs et al, ISSCC 2005)






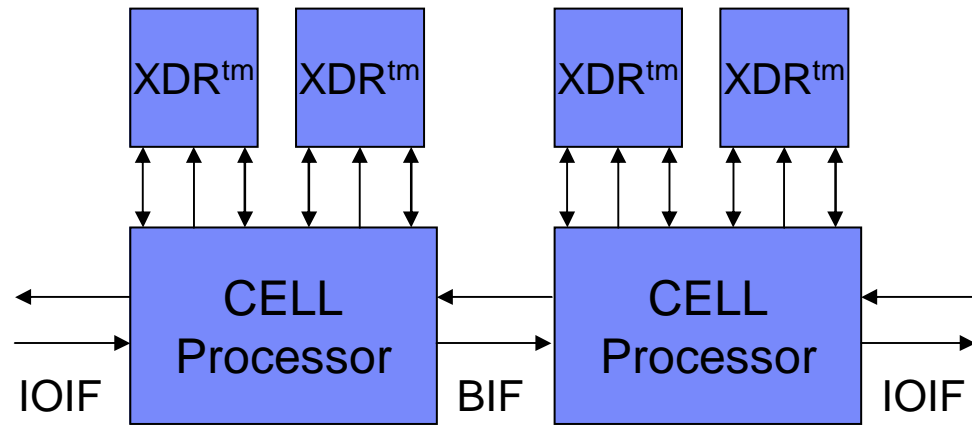

## I/O and Memory Interfaces

- **I/O Provides wide bandwidth**
  - Dual XDR™ controller (25.6GB/s @ 3.2Gbps)
  - Two configurable interfaces (76.8GB/s @6.4Gbps)
  - Flexible Bandwidth between interfaces
  - Allows for multiple system configurations



# Cell Processor Can Support Many Systems


- Game console systems
- Workstations (CPBW)
- HDTV
- Home media servers
- Supercomputers




## Cell Processor Based Workstation (CPBW) (Sony Group and IBM)

- **First Prototype “Powered On”**
- **16 Tera-flops in a rack (est.)**
  - ( equals 1 Peta-flop in 64 racks )
- **Optimized for Digital Content Creation, including**
  - Computer entertainment
  - Movies
  - Real-time rendering
  - Physics simulation

**CELL Processor  
Board**



**16 TFlop rack**



## Cell Processor Example Application Areas

- Cell is a processor that excels at processing of rich media content in the context of broad connectivity
  - **Digital content creation (games and movies)**
  - **Game playing and game serving**
  - **Distribution of (dynamic, media rich) content**
  - **Imaging and image processing**
  - **Image analysis (e.g. video surveillance)**
  - **Next-generation physics-based visualization**
  - **Video conferencing (3D?)**
  - **Streaming applications (codecs etc.)**
  - **Physical simulation & science**

## Summary

- **Cell ushers in a new era of leading edge processors optimized for digital media and entertainment**
- **Desire for realism is driving a convergence between supercomputing and entertainment**
- **New levels of performance and power efficiency beyond what is achieved by PC processors**
- **Responsiveness to the human user and the network are key drivers for Cell**
- **Cell will enable entirely new classes of applications, even beyond those we contemplate today**

## Acknowledgements

- **Cell is the result of a deep partnership between SCEI/Sony, Toshiba, and IBM**
- **Cell represents the work of more than 400 people starting in 2001**
- **More detailed papers on the Cell implementation and the SPE micro-architecture can be found in the ISSCC 2005 proceedings**