Synergistic Processing In Cell's Multicore Architecture

Michael Gschwind, et al.

Presented by: Jia Zou

CS258

3/5/08

Goal for Cell

- Increase processor efficiency for most performance per area
- Reduce area per core, have more core in a given chip are
- Take advantage of the application parallelism
  - Aimd at data-processing intensive applications

Cell Architecture

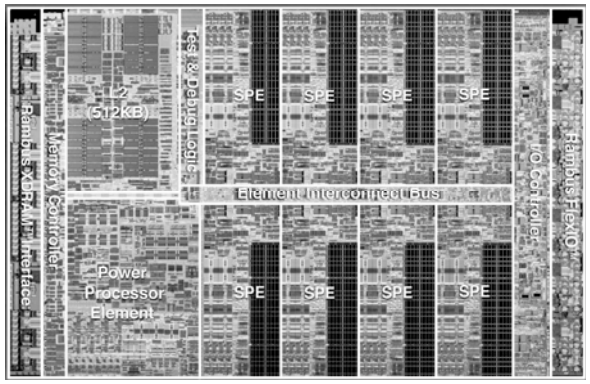

Design Philosophy

- Simple cores, lots of them
  - Any complexity reduction directly translates into increased performance
  - Exploiting the compiler to eliminate hardware complexity
- PPE serves as controller, SPE provides performance
  - PPE and SPEs share address translation and virtual memory architecture

Synergic Processing Unit

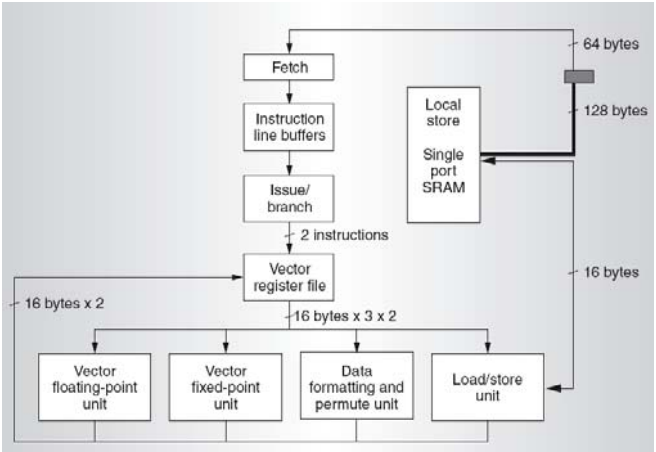

Data alignment for Scalar and Vector Processing

- SPU has no separate support for scalar processing
  - Unified scalar/SIMD register
  - Unified execution unit
  - Simpler control unit
    - Software-controlled data-alignment approach
  - Simplifies scalar data extraction, insertion, sharing between scalar and vector data
    - Increases compiler efficiency

Scalar Layering

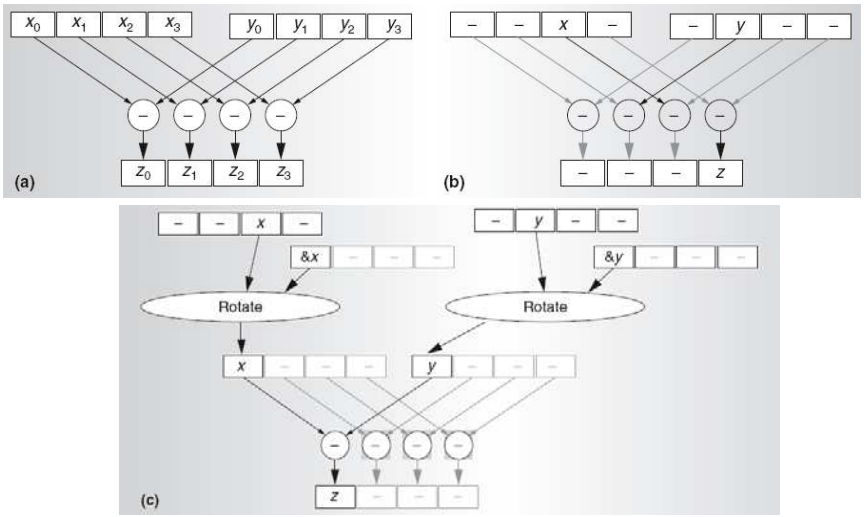

Data-Parallel Conditional Execution

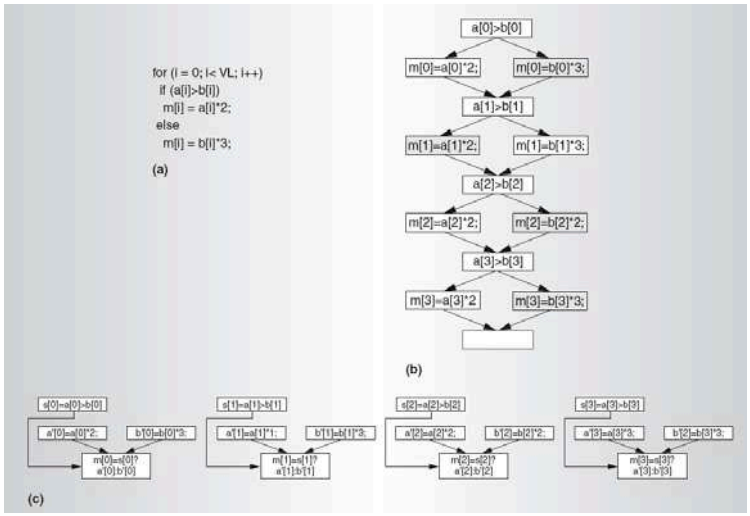

Deterministic Data Delivery

- SPE has local stores
  - 4Kb – 4Gb address range
  - Stores both instruction and data
  - All memory operations that the SPU executes refer to address space of this local store
  - Different from cache memory by:
    - No cache coherency problem
    - Offers low and deterministic access latency

Statically Scheduled ILP

- Instruction fetches are scheduled statically
- Delivery up to two instructions per cycle
  - One to each complex
- Static branch prediction: prepare-to-branch instruction => initiate instruction prefetch

SPE Microarchitecture

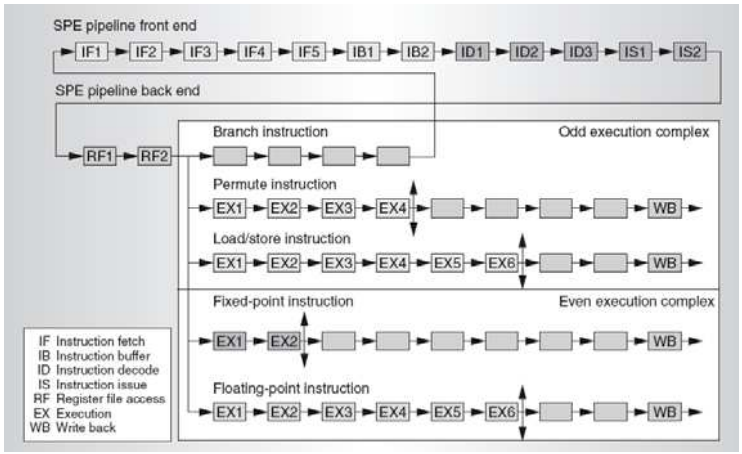
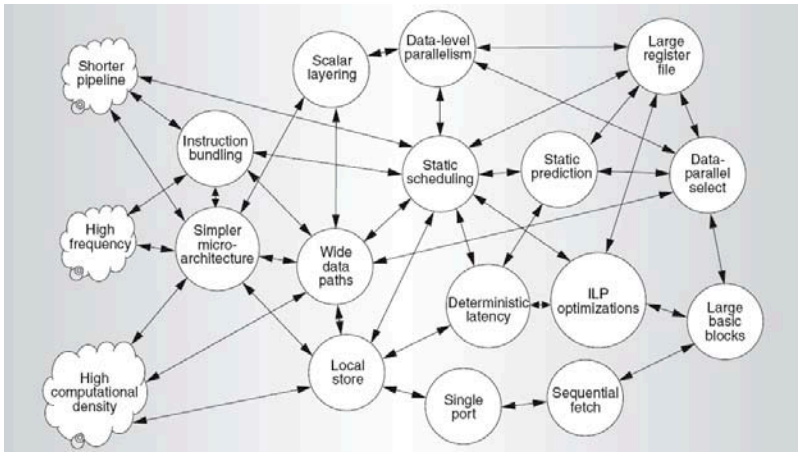
Design Goals and Decisions