

# **Review of recent processor evolutions and trends**

**Sverre Jarp**

**CERN openlab CTO**

**CERN openlab monthly meeting**

**18 September 2007**





# Agenda

- **Introduction**
- **Intel Core 2 (Clovertown → Harpertown)**
- **AMD Quad-core Opteron (Barcelona)**
- **SUN Ultrasparc T2 (Niagara 2)**
- **Some software recommendations**
- **Concluding remarks**



# Introduction

- **In the past, Moore's law was exploited to create a single processor (with a single core) running at maximum speed!**
  - Pentium 4 went “quickly” from 1.5 GHz to 3.8 GHz
- **Little emphasis on multithreading and performance tuning**
  - “Let's just wait for the next speed increase in 6 months!”

**Trend brutally stopped by leakage current (i.e. heat)**




# The new game

- **Now, the industry is searching for a “new” sweet spot:**
  - Wider execution (via “instruction-level parallelism”)
  - Multi-core (sharing nothing or sharing cache)
  - Many-core (sharing fast interconnect)
  - Hardware multithreading
    - N execution states share each execution engine
  - Specialized execution engines
    - Hash/Cipher, Graphics
  - Thinner execution
  - In-order execution
- **Redefine sharing**
- **All – at a reasonable frequency**
  - Between 1 and 3 GHz



# Hierarchy of execution engines

- A processor contains cores
- A core contains hardware threads
- A thread owns/shares execution units



Sharing is possible at all levels:  
-Execution units,  
-HW threads,  
-Caches



# Processor Review



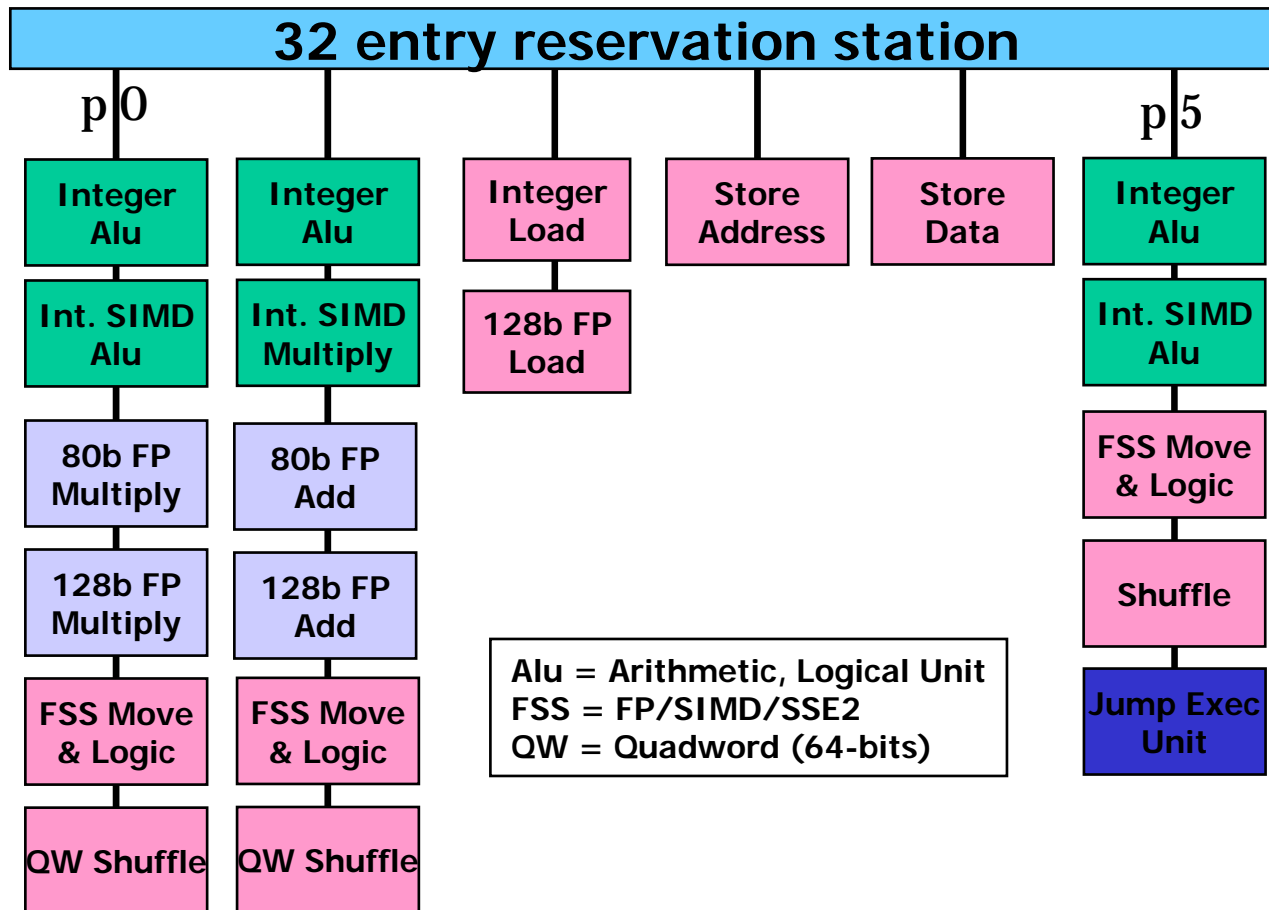
# Intel's current architecture

- **Intel Core 2 processor (Woodcrest, Clovertown):**
  - New, wide micro-architecture
    - *Netburst* completely replaced
  - 65 nm now, moving to 45 nm soon
    - “Penryn”-family expected in November; Core 3 a year later
  - Relatively high frequency
    - 3 GHz now, maybe somewhat higher in the future
  - Reasonable power consumption
    - Although FB-DIMMs add to the total power bill: 10 W/GB
  - Quad-core since last year
    - Although purists point out that this is two dual-cores bolted together
      - They work for us!



# Intel micro-architecture

- Execution ports in the Core 2 processor:



Issue ports in the Core 2 micro-architecture  
(from Intel Manual No. 248966-014)






# Intel CPU parameters

- Core 2 processor (Clovertown)

| Caches/TLB       | Size<br>(total / line) | Access<br>(cycles) | Porting | Associativity<br>(N-ways) |
|------------------|------------------------|--------------------|---------|---------------------------|
| L1I              | 32 KB / 64 B           | -                  |         | 8                         |
| L1D              | 32 KB / 64 B           | 3                  | dual    | 8                         |
| L2 (semi-shared) | 2 * 4 MB / 64 B        | 14                 |         | 16                        |
| ITLB0            | 128 e.                 | -                  |         |                           |
| DTLB0            | 16 e.                  | -                  |         | 4                         |
| DTLB1 (4K pages) | 256 e.                 | 2                  |         | 4                         |

|                   |                  |
|-------------------|------------------|
| Instruction issue | 4 * 4 $\mu$ -ops |
| CPU speed         | 3.0 GHz          |
| Bus speed         | 1333 * 8 B       |





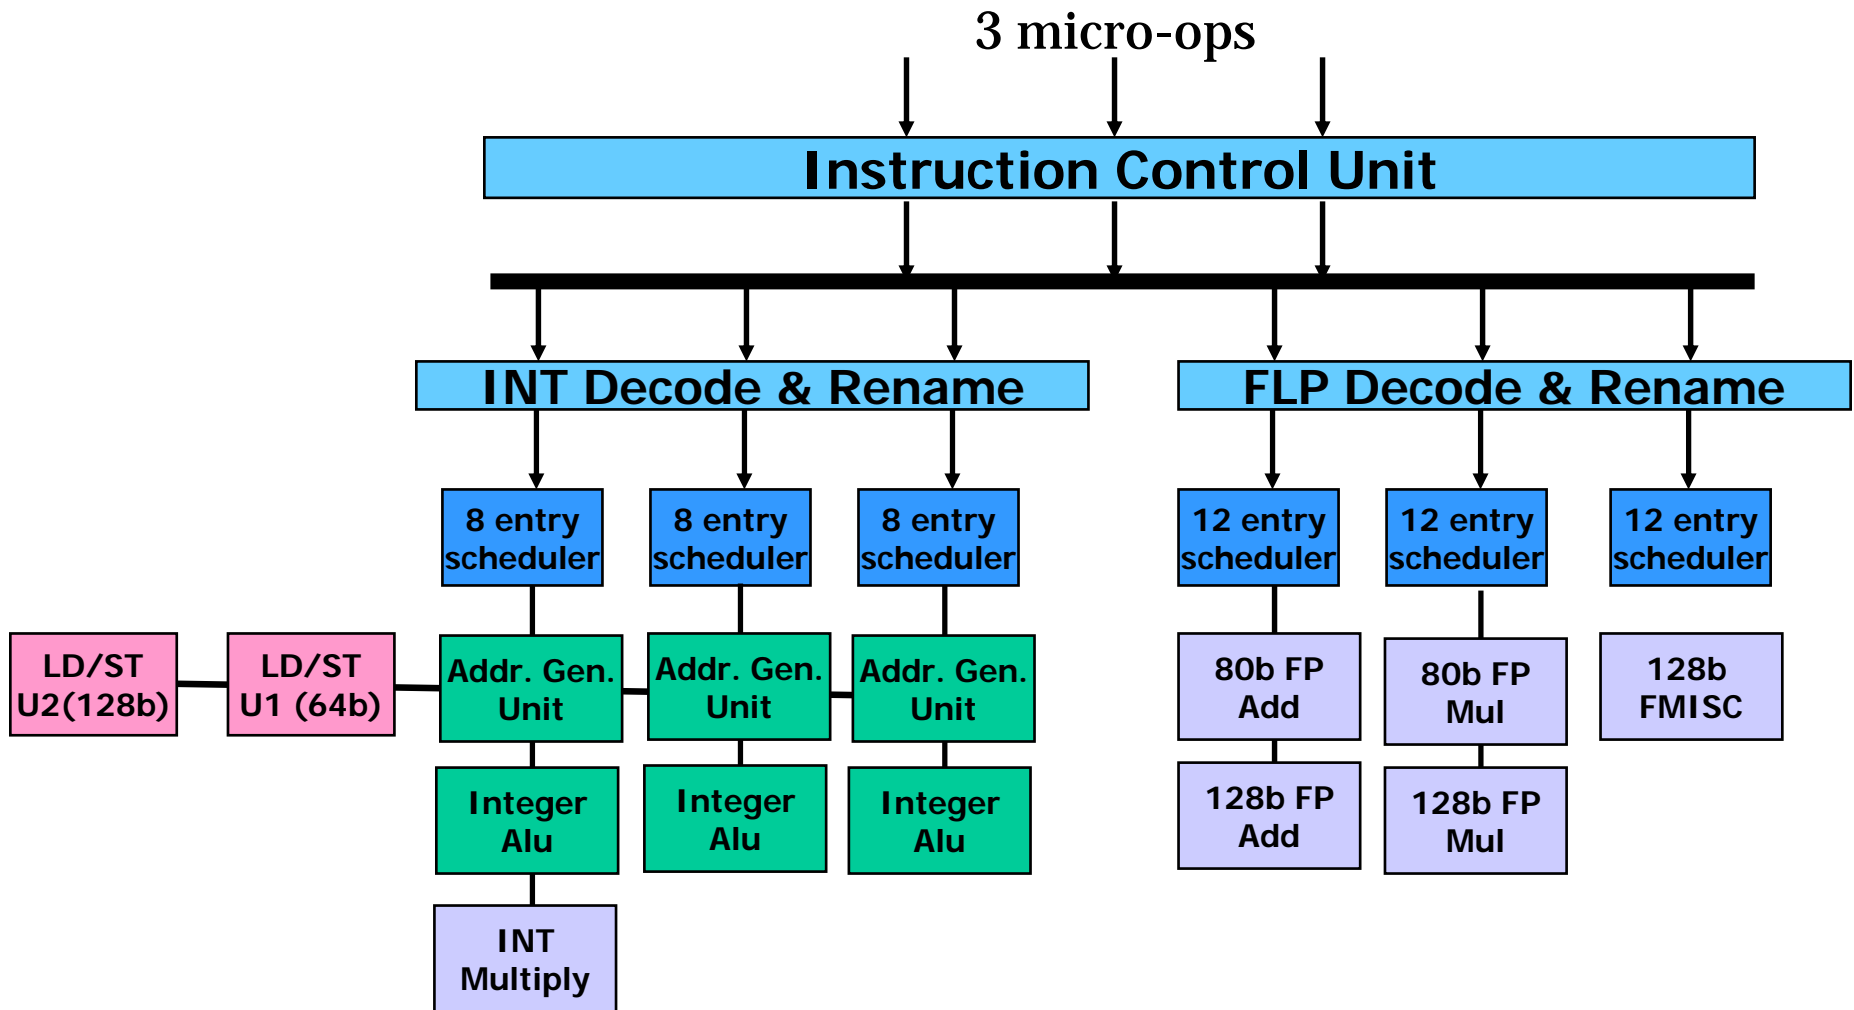
# AMD's current architecture

- **Barcelona processor:**
  - Opterons feature Integrated Memory Controller and HyperTransport since a long time
  - Native Quad core design
  - New K10 micro-architecture
  - Finally 65 nm technology
    - 45 nm around end-2008/beginning 2009
  - 95 W: not so cool; but RDIMMs are cooler than FB-DIMMs
  - Three cache levels
    - But L3 is surprisingly small
  - “Only” 2 GHz
    - Hoping to ramp up to 2.3/2.5 GHz by YE
- **Future processors:**
  - SSE5 extensions



# AMD micro-architecture

- Execution units in Barcelona processor:



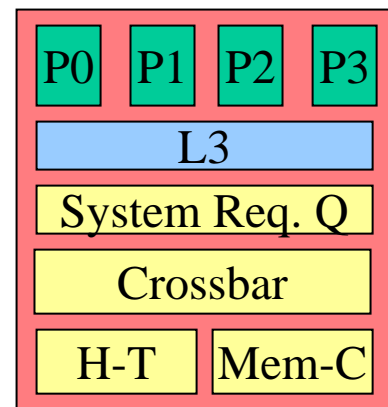


# AMD CPU parameters

- Barcelona processor:

| Caches/TLB  | Size<br>(total / line) | Access<br>(cycles) | Porting | Associativity<br>(N-ways) |
|-------------|------------------------|--------------------|---------|---------------------------|
| L1I         | 64 KB / 64 B           |                    |         | 2-way                     |
| L1D         | 64 KB                  | 3                  | dual    | 2-way                     |
| L2          | 512 KB                 | 12                 |         | -                         |
| L3 (shared) | 2 MB                   | <38                |         | -                         |
| L1 ITLB     | 32 e.                  |                    |         | fully                     |
| L1 DTLB     | 48 e.                  |                    |         | fully                     |

|                   |                  |
|-------------------|------------------|
| Instruction issue | 4 * 3 $\mu$ -ops |
| CPU speed         | 2.0 GHz          |
| Bus speed         | 2 * 8 * 667 MB/s |
| HyperTransport    | 2 * 8 * 2 GB/s   |





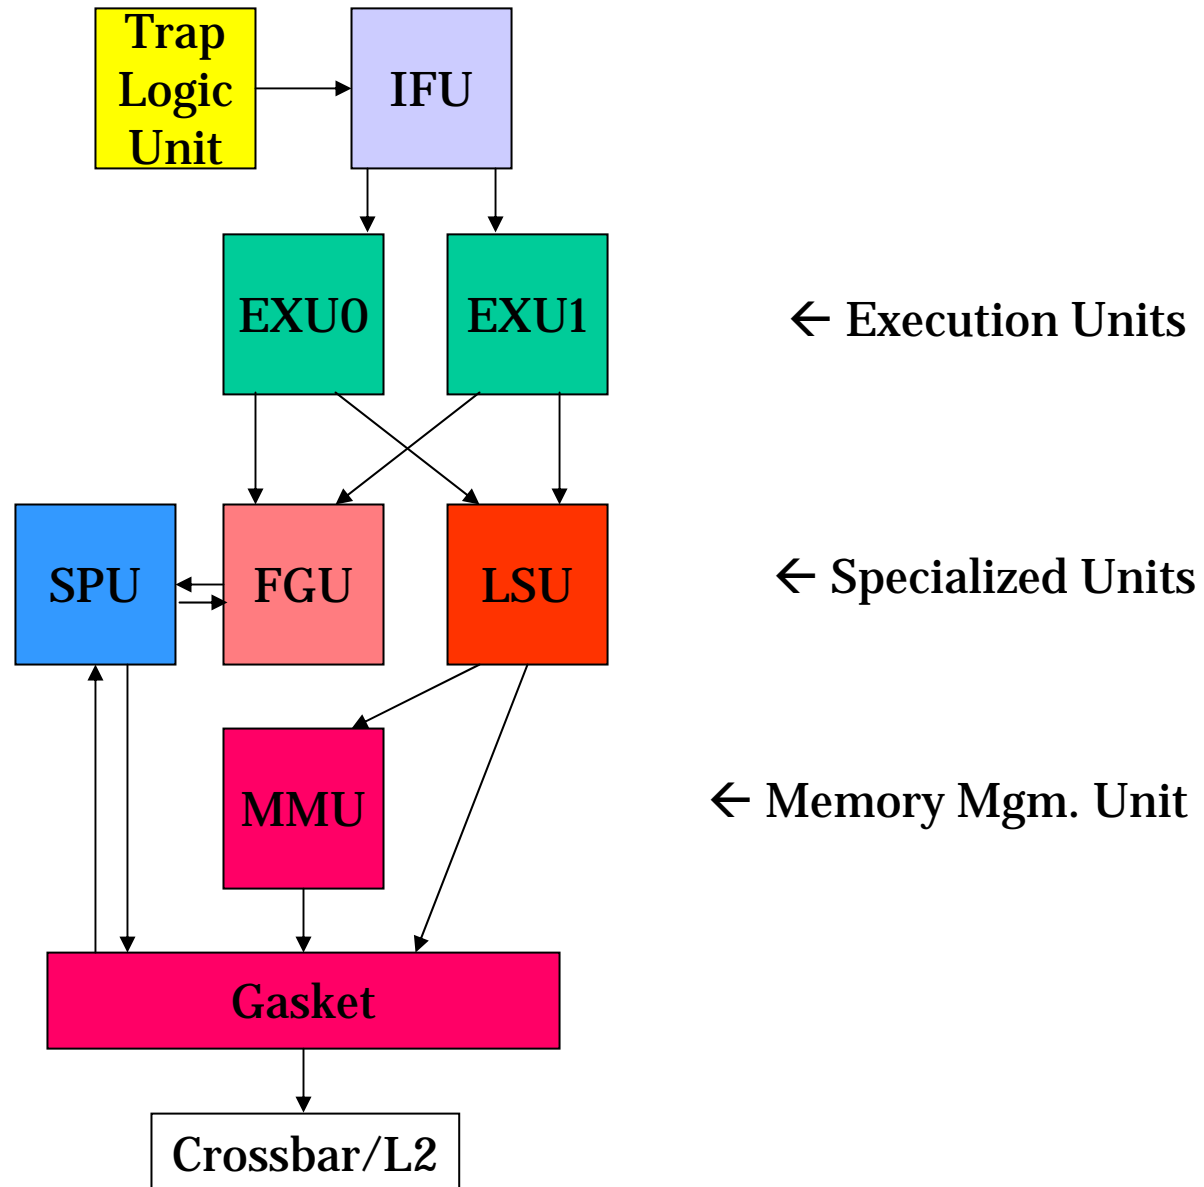
# Sun's “new” architectures

- **Niagara 2 processor:**
  - Evolution from Niagara 1
  - Aggressive throughput design
  - 65 nm technology (from TI)
  - Not yet multi-socket enabled
  - Multiple cores and multiple threads
    - 8 x 8
  - Power: 70 W
  - On-die cross-bar
  - Specialized engine
    - Stream Processing Unit
- **The “Rock” processor (2008) will also be focused on throughput:**
  - 4 cores / 4 blocks (with a FGU) / 2 threads



# Niagara 2 Execution Logic

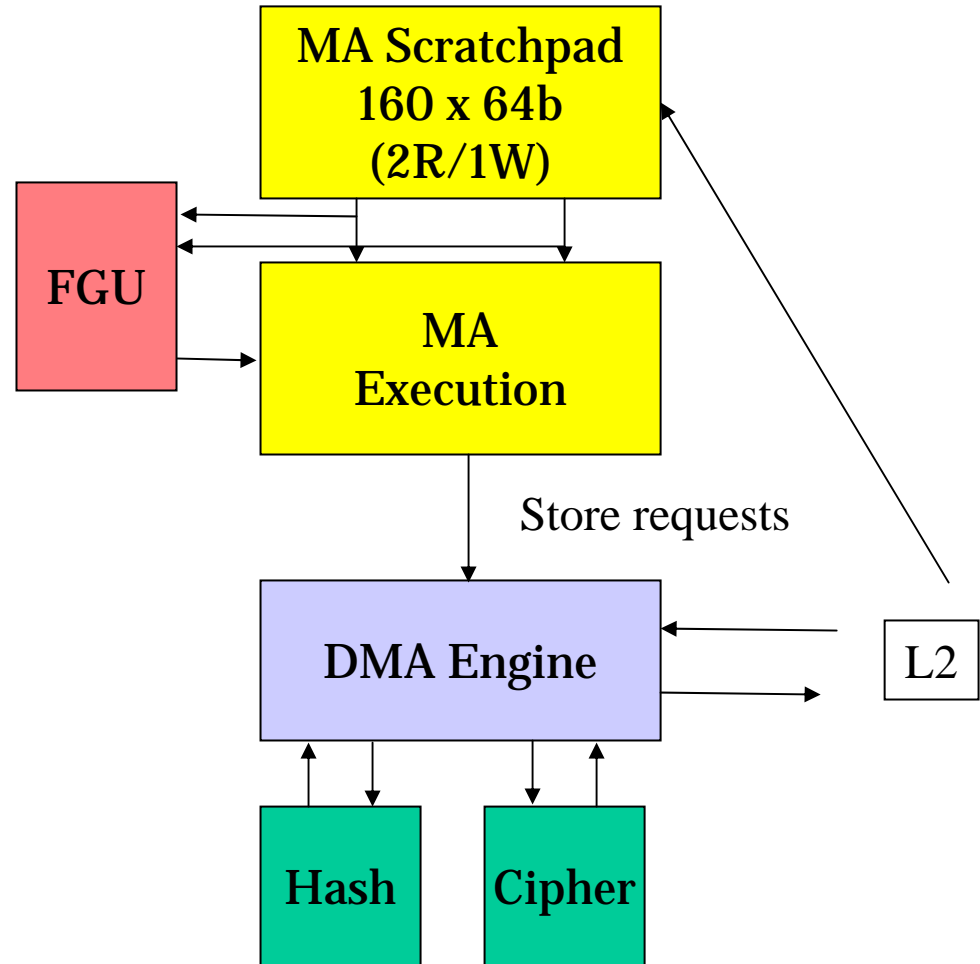
- **Simplified overview:**





# Stream Processing Unit

- **Cryptographic co-processor**
- **Modular Arithmetic Unit:**
  - RSA, elliptic curves, etc.
- **Cipher/Hash**
  - RC4, DES/3DES, AES
  - MD5, SHA





# Key CPU parameters

- SUN Ultrasparc T2 processor**

| <b>Caches/TLB</b>   | <b>Size<br/>(total / line)</b> | <b>Access<br/>(cycles)</b> | <b>Porting</b> | <b>Associativity<br/>(N-ways)</b> |
|---------------------|--------------------------------|----------------------------|----------------|-----------------------------------|
| <b>L1I</b>          | <b>16 KB / 32B</b>             |                            |                | <b>8</b>                          |
| <b>L1D</b>          | <b>8 KB / 16 B</b>             | <b>3</b>                   | <b>single</b>  | <b>4</b>                          |
| <b>L2 (8 banks)</b> | <b>4 MB / 64 B</b>             | <b>24 - 26</b>             |                | <b>16</b>                         |
| <b>ITLB</b>         | <b>64 e.</b>                   |                            |                | <b>fully</b>                      |
| <b>DTLB</b>         | <b>128 e.</b>                  |                            |                | <b>fully</b>                      |

|                          |                      |
|--------------------------|----------------------|
| <b>Instruction issue</b> | <b>8 * 2 ops</b>     |
| <b>CPU speed</b>         | <b>1.4 GHz</b>       |
| <b>Bus speed</b>         | <b>25.6 (?) GB/s</b> |



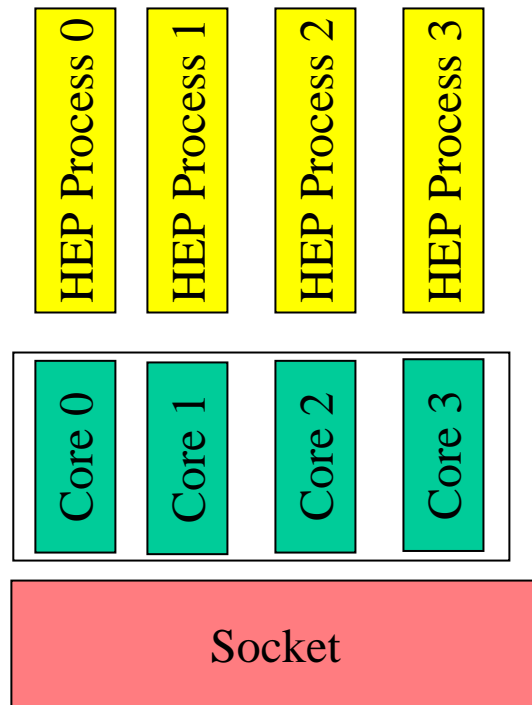


# **Software Issues and Recommendations**

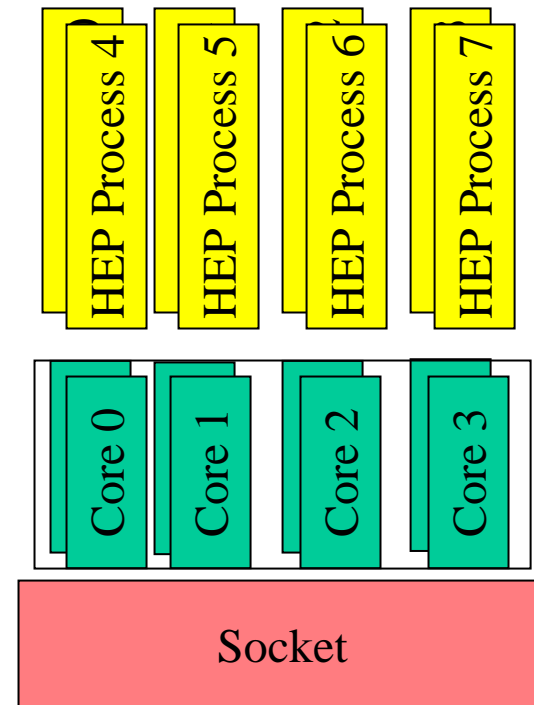


# Single threaded processes

- Simply illustrated:



Quad-core



Octo-core or Quad-core w/two-way HW Multithreading  
(seen by the OS as 8 independent CPUs)



# Our memory usage

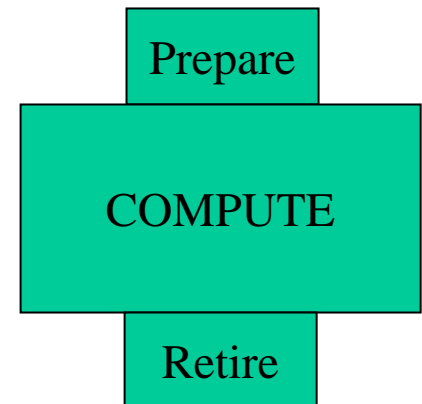
- **An initial preoccupation:**

- Today, we need 2 – 4 GB per single-threaded process.
- In other words, a dual-socket server needs at least:
  - Single core: 4 - 8 GB
  - Quad core: 16 - 32 GB
  - Future 16-way CPU: 64 – 128 GB (!)
  - Future 64-way CPU: 256 – 512 GB (!!)



# 1) Increased ILP


- Aim at creating richer “sections”, with especially the floating-point contents exposed
- Assist our C++ compilers in making these sections effective\*\*
  - Optimization in all important areas
    - Inlining of “tiny” methods
    - Disambiguation of data pointers/references
    - Minimization of if and switch statements
    - Etc.
  - Optimization of mathematical functions
    - Log, exp, sine, cosine, atan2, etc





## 2) Multithreading


- Explore new paradigms, for example:





# 3) Simplify/restructure code

- **Today, our frameworks are very complicated and heavy**
  - In one case, we observed 400+ shared libraries
- **Make a move à la BOOT?**
  - Test coverage of various applications has shown that frequently the 80/20 rule applies:
    - 20% of the code is enough to cover 80% of the (even complex) use cases
- **Having a more modular approach would be very beneficial**
  - For instance,
    - Quicker porting to assess new hardware
    - Quicker adoption of new paradigms





# Conclusions

- **Main issue (for manufacturers) is to deploy transistors intelligently**
- **Main issue (for consumers) is to benefit proportionally**
- **Some trends are emerging:**
  - Increase “uncore” sophistication
    - Crossbars, rings, etc. as interconnects
  - Preserve “fat” cores: 2, 4, 8 with limited multithreading
  - Push thin cores with substantial threading
  - Produce specialized units
    - Cryptography, Graphics, Networking, etc.

**The future will be exciting!**