



Introduction to Parallel Computing

Alexandre David

B2-206

<http://www.cs.aau.dk/~adavid/teaching/MTP-07/>

Welcome.

Web page: schedule, book, exercises, slides, everything about the course.
Also accessible from my personal web page.

E-mail: adavid@cs.aau.dk, also from my web page.



Presentation of the Course

- Parallel Computing
 - Little on parallel hardware
 - Mostly on parallel algorithms and design
- Models for Parallelism (PRAM...)
- Tools for Parallelism (MPI, pthreads...)
- 15 lectures, 3x30 min + exercises
- TAs: Morten Kühnrich & Simon Holm Thøgersen – {mokyhn,odie}@cs.aau.dk

16-02-2007

Alexandre David, MVP'07

2

Updated from last year, changed content according to feedback and discussions with other lecturers.


No overlap with other courses.

Only place in the curriculum where you have a chance to learn about parallel programming.

Meaning of the course:

- Models for parallel machines and programs, programming paradigms, etc...
- Tools for parallelism: standard API s.a. MPI, pthreads, OpenMP (not in the course but you have materials in the book).

The Course Book



- *Introduction to Parallel Computing*
- Covers many aspects on parallel computing.
- Both basic and advanced topics.
- I will follow the book but not cover it all.

16-02-2007

Alexandre David, MVP'07

3

Claims to be modular and suitable for a wide variety of undergraduate and graduate level courses.

Covers traditional algorithms (sorting, graph, searching) and scientific computing algorithms (matrix, FFT).

Practical, code examples.

MPI, pthreads, also OpenMP (not in this course): 3 most widely used standards for writing portable parallel programs.

Recent (2003) and solid book.

Show the plan of the chapters (Figures/chap1.pdf).



Course & Assignments

- Lectures will be alternated between theory & practice.
 - Course will be reorganized from last year.
- Assignments will be half theory, half practice.
 - 7 assignments, all compulsory.
 - Model: complete them until they are good.
 - Careful: Do not accumulate delay.



Goals of the Course

- Design, analysis, and implementation of parallel algorithms.
 - Principles of parallel algorithm design.
 - Analytical modeling of parallel programs.
 - Tools such as MPI & pthreads.
 - Some examples.

16-02-2007

Alexandre David, MVP'07

5

Main goal (design, analysis, implementation) that needs important notions (sub-goals).

Short motivation:

•Hardware: Parallel computing has changed from tightly scalable message passing platforms to today's inexpensive clusters and multiprocessor machines.

•Software: Programming models have evolved from custom to standard APIs. MPI = standard message passing library, pthreads = thread library, OpenMP = directive based models.

Impact on process of **design, analysis, and implementation** of parallel algorithms: What this course is about.



Do We Need Parallelism? ?

- ? ■ Complexity of parallel programs: How to *specify* and *coordinate* **concurrent** tasks? *race/deadlock/livelock*
- ? ■ Question of portable algorithm standards?
- ? ■ Need to accelerate applications?

16-02-2007 Alexandre David, MVP'07 6

There is another course dealing with the question of *concurrency*, in particular problems of deadlocks, livelocks, synchronization, condition variables, etc... I won't spend much time on the needed concepts, only what I need for this course (more pragmatic and practical approach).

- Specify: How to decompose a problem (most often sequential) into a set of parallel tasks to execute?
- Coordinate: Efficiency (control overhead in extra communication) and correctness issues (race conditions, deadlocks, livelocks).

Recall of race condition (several execution orderings may yield different results with the same program), deadlock (system not responding or doing anything), livelock (infinite loop without progress).

Algorithm standards: Actually in terms of API, now there are some.

Need to accelerate: Think of spending 2 years of development when the platform is going to be obsolete by the time you are finished.



Trends in Hardware

- Everything points towards parallelism from multi-core, hyper-threading, multi-threads, superscalar, ... technologies.
- Because
 - Limits to continue to increment performance with single processors.
 - Other constraints like heat, complexity, yields, etc...

16-02-2007

Alexandre David, MVP'07

7

Put a product name on all of these technologies to make it more concrete:

- Multi-core: X2 Athlon64, Intel's dual-core (P4 & new mobile CPU).
- Hyper-threading: Intel's technology to utilize the CPU better (switch thread instead of staying stalled on data).
- Multi-thread: Microsoft Xbox CPU, multithread triple core.
- Superscalar: Every modern GPU and CPU, several instructions in the same clock cycle (pipeline, different execution units).

Single processors have implicit lack of parallelism and have bottlenecks such as critical data paths and limited memory sub-system.

Example: multi-core recent design adopt simpler architectures replicated several times (no out-of-order execution), why? Complexity, efficiency/watt, die-size, price transistor/OPS (operation per second).



Arguments for Parallelism

- Computational power:
 - Moore's law.
 - Translating transistors into useful OPS.
- Memory/disk speed:
 - Performance/yr: CPU +40%, DRAM +10%.
 - How to feed data?
- Parallel platforms: larger aggregate cache+bandwidth+IPC...

16-02-2007

Alexandre David, MVP'07

8

•Computational power: The demand for higher and higher computation power always grows, Moore's law is only the technological answer to that demand. *We want to pursue Moore's law, that's why it still holds.* The question is: How to continue from now on? (with the problems mentioned before).

•Memory/disk speed: Overall system performance is defined by CPU speed and the ability of the system to feed data to it. We have cheated so far by bridging the speed gap with caches that work thanks to the **data locality** property of almost all programs. Still we have both problems of **latency** and **bandwidth**. The same applies to disks (you are familiar with RAID technology).

•Parallel platforms: Linear increase in the number of processors of cache, bandwidth, etc... in total. The question is: How to use the increased resource such that the performance has a linear increase as well?

•Other arguments: Data communication (SETI@home a.k.a. search for extraterrestrial intelligence), constraints on location of data & resources that require distributed/parallel algorithms.



Scope of Parallel Algorithms

- Engineering & design.
- Scientific computing.
- Commercial (web) applications.
- Embedded systems.
- Gaming industry.

16-02-2007

Alexandre David, MVP'07

9

- Engineering & design: complex physical processes, geometric & mathematical modeling in context of parallel computers.
- Scientific applications: human genome sequencing, computational physics & chemistry.
- Commercial applications: multiprocessor & cluster machines for web & database servers.
- Embedded systems: cars, planes, etc... have many computer systems communicating via some network. 90% of computer systems are embedded systems.
- Gaming industry: Xbox 360 (triple core CPU + general multi vertex/pixel shader engine), PS3 with Cell processor (8 simple computational units + 1 G5 on one core).

Parallel Programming Platforms





Outline

- Implicit Parallelism (2.1)
- Limitations of Memory System Performance (2.2)
- Dichotomy of Parallel Computing Platforms (2.3)

16-02-2007

Alexandre David, MVP'07

11

I will follow the book, skipping sections from time to time but the order will be respected. References to sections are given on the web and reminded during the lectures (plan may change). Lectures are intended to teach you and help you read the book.



Motivations

- Bottlenecks in computers:
 - Processor
 - Memory
 - Datapath
- Addressed with multiplicity.
- Parallelization not solution to everything
 - Sub-optimal serial code bad
 - Optimize serial first (similar characteristics)

16-02-2007

Alexandre David, MVP'07

12

Bottlenecks of different kinds.

- Processor: less and less, though can be depending on some bad behaviors (branch miss prediction in large pipelines).
- Memory: more and more considering the speed gap between processor and memory, the problem being how to feed the processor with data so that it does not stay idle.
- Datapath: depends on programs and architecture, linked to previous ones.

Motivations for optimizing serial programs and why we talk about implicit parallelism: Sub-optimal serial code exhibits **unreliable** and **misleading** behaviors. Undesirable effects coming from poor cache utilization, bad branch prediction, etc ... that may become even worse in a parallel context (distribute data, synchronize, etc ...).

Similar characteristics in serial programs with intrinsic parallelism of modern processors (pipelines). Understanding architecture is the first step to good programming.



Trends in Microprocessors

- Processor speed increase exponentially
- More and more transistors: How to use them wisely?
- Several instructions are issues in the same clock cycle (possibly on multiple functional units) : **superscalar** processors.
- How to select and execute instructions?

16-02-2007

Alexandre David, MVP'07

13

Already mentioned: Moore's law power x2 every 18 month, but for processor only. Global performance plagued by memory abysmal performance. Higher level of integration poses the question of how to best utilize transistors.

Functional units: MMU, FPU, etc usually part of the marketing buzz of microprocessor companies.

How to? That's the different architectures, not going into details but basically all processors are +/- RISC processors with a translation from assembler to microcode. Then most have branch prediction.



Pipelining and Superscalar Execution

- Pipeline idea: overlap stages in instruction execution.
- Example of car factory.
- The good: higher throughput.
- The bad: penalty of branch miss prediction.
- Multiple pipelines: several functional units.

16-02-2007

Alexandre David, MVP'07

14

Overlapping stages: cut instructions in small pieces, one per cycle, and try to occupy all the stages. But why after all? Not better to have a super powerful one stage-do-all? Car factory: Imagine a fast factory where it takes 12h to complete one car. If there is one unit doing all, then it will be busy all the time and throughput would be 1 car per 12h. How to improve? Buy 11 other full scale units? Super expensive! Cut the bit unit in 12 smaller parts, 1h per part, every car needs 12h but throughput is 1 car per hour. 12x faster, cost efficient.

Branch prediction: try to keep the pipeline busy by filling it ahead, but if did it wrong, then need to flush it (and loose all the computations). P4: 20 stages, miss prediction means loose 20 cycles.

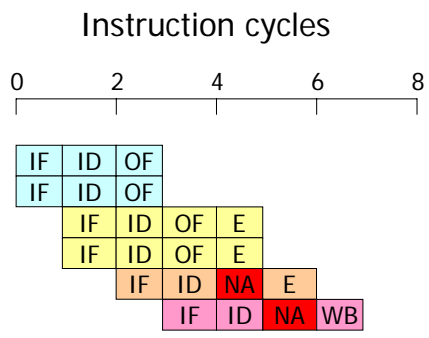
Pipelining and Superscalar Execution

Compiler

```
c=a+b+c+d
as
c=(a+b)+(c+d)
```

CPU

1. load R1,@1000
2. load R2,@1008
3. addR1,@1004
4. addR2,@100C
5. add R1,R2
6. store R1,@2000



2x IF, ID, OF, ... in the same cycle:
superscalar.

16-02-2007

Alexandre David, MVP'07

15

Dual issue or two-way superscalar execution.

IF: Instruction Fetch.

ID: Instruction Decode.

OF: Operand Fetch.

E: Instruction Execute.

WB: Write back.

NA: No Action.

Note: begin to execute 6th instruction at 4th clock cycle.

Compare the number of cycles without pipelining and/or superscalar.



Pipelining and Superscalar Execution

- Imagine another ordering (or factorization by the compiler): different performance.
- Resolve data dependency.
- Reordering by CPU possible (out-of-order execution).
- Resource dependency.

16-02-2007

Alexandre David, MVP'07

16

Data dependency: needs previous results in order to continue computations. $A=(B+C)*D$, we need $B+C$ before computing $*D$.

Resource dependency: needs functional units. $A=B*C+C*D+D*E+E*F+F*G$, obviously not all $*$ can be done in parallel because of lack of functional units.


Most processors are capable of out-of-order execution, not Xbox 360.

Limitations

- Bottleneck: slowest stage -> small stages to go fast -> long pipelines
 - BUT miss prediction gives big penalties
- How to keep busy the functional units?

Vertical waste:
no instruction on
execution unit. Here
no instruction on the
adder unit.

16-02-2007



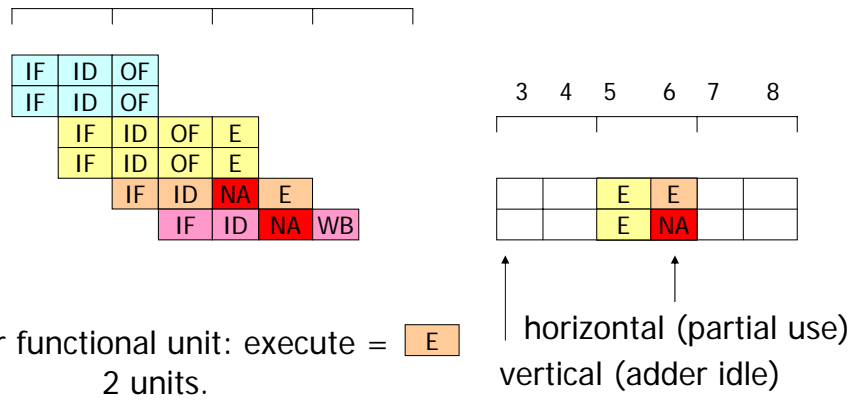
Horizontal waste:
parts of execution
units used.

17

Intrinsic parallelism: Pipeline (multiple stages) & multiple functional units (superscalar) implement parallelism.

Modern processors: 4 way superscalar, 10-20 stage pipeline.

Adder Utilization (fig 2.1)



16-02-2007

Alexandre David, MVP'07

18

More explanation for fig 2.1.



VLIWP ?

- Bundle instructions together to simplify the superscalar scheduler.
- IA64 (Itanium) is an example.
- Problems:
 - *Rely a lot* on the compiler.
 - Limited parallelism (not dynamic).

16-02-2007

Alexandre David, MVP'07

19

Very Long Instruction Word Processors!

Superscalar schedulers are complex and expensive (transistors). VLIW design idea is to rely on the compiler to bundle instructions together, so that the scheduler becomes very simple.



Limitations of Memory System Performance

- The **memory system** is most often the **bottleneck**.
- Performance captured by
 - latency and
 - bandwidth.
- Remark: In practice latency is complicated to define: CL2, CL3, 2-2-2-5,...

16-02-2007

Alexandre David, MVP'07

20

The problem is most often how to feed the processor with continuous data so that it does not **stall**.

Latency is the time from the issue of a memory request to the time the data becomes available to the processor.

Bandwidth is the rate at which data can be pumped to the processor.

Example: water hose. Latency: time before first drop of water comes out.

Bandwidth: rate flow of water.



Effect on Performance: An Example

- Processor @1GHz (1ns cycle), DRAM with 100ns latency, capable of executing 4 IPC.
- 4 IPC @1GHz -> 4GFLOPS *peak rating*.
- Processor must wait 100 cycles for every request.
 - Vector operations (dot product) @10MFLOPs.

16-02-2007

Alexandre David, MVP'07

21

No cache in this example to simplify. It is still general enough since we can consider first access to some memory and take cache miss into account.


ALU: arithmetic and logical unit.

FPU: floating point unit.

Here absolute worst case scenario but still we loose a factor 100 in performance.



Improving with Cache

- Note: Often “\$\$” on pictures (cash).
- Hierarchical memory architecture with several levels of cache (2 common).
- Instruction and data separate for L1.
- Low latency, high bandwidth, but small.
-  Why does it improve performance???

16-02-2007

Alexandre David, MVP'07

22

Common: Athlon 64 64K+64K L1, 1M L2. Pentium 4 more complicated NetBurst with execution trace cache (12K) and 16K L1, with 1M L2.

Now you have to think some time about why it helps. You know about cache hit ratio, cache miss, at least you've heard about it.



Why is \$\$ good?

- Temporal locality
 - Repeated access to the **same** data in a small window of time.
- Spatial locality
 - **Consecutive** data accessed by successive instructions.
- Vital assumptions, almost always hold.
- Very important for parallel computing.

16-02-2007

Alexandre David, MVP'07

23

REMEMBER these two! They are common to almost all programs and are vital to cache performance.

For parallel computing, even more important: apart from the aggregate higher amount of cache that must be used wisely, we have more penalty for moving data around processors (or processor nodes).

That also explains the model numbering for AMD processors. ☺



Matrix Multiplication Example

- Common example, will be used many times in the course.
- $C=A*B$, where A ($n*m$), B ($p*n$), and C ($p*m$) are matrices.

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$

16-02-2007

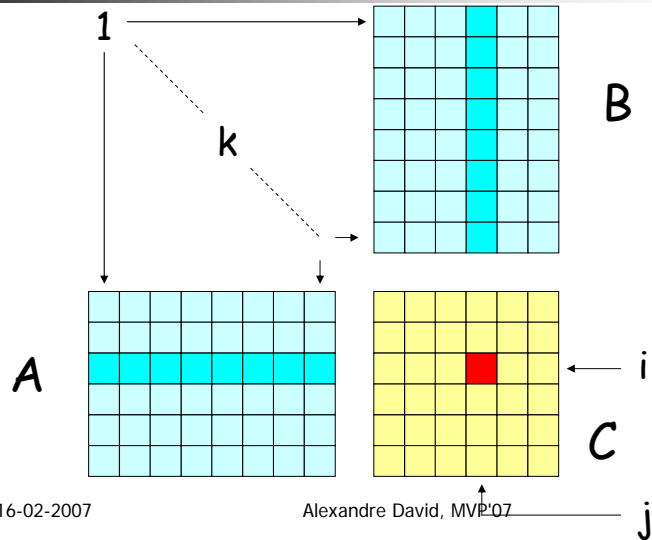
Alexandre David, MVP'07

24

Compatible dimensions required. In practice $n*n$.



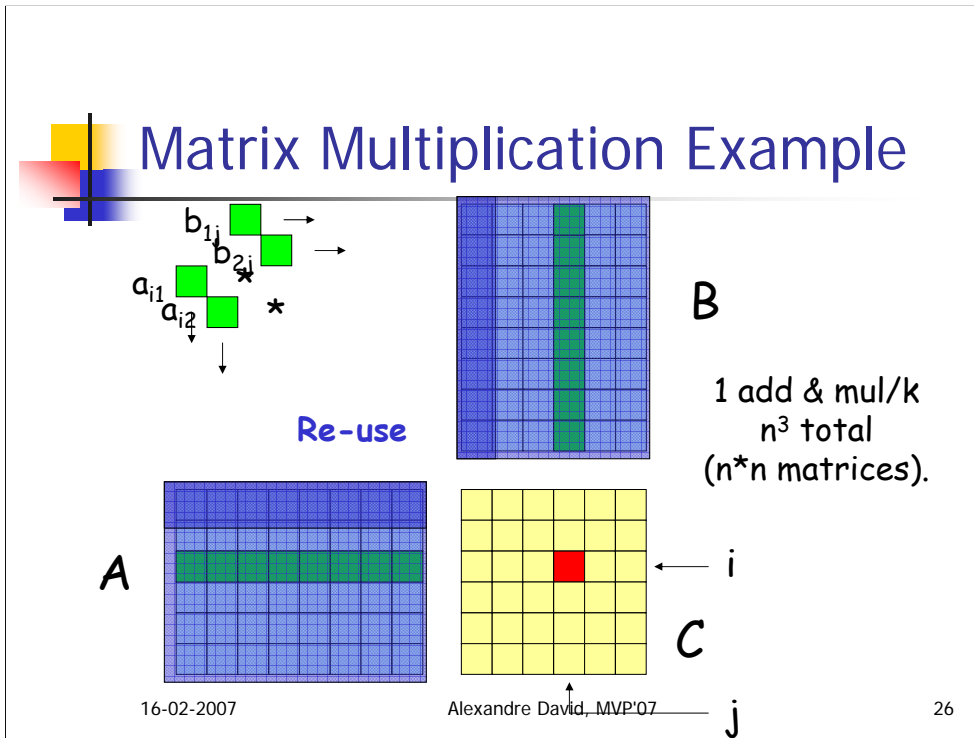
Matrix Multiplication Example



16-02-2007

Alexandre David, MVP'07

25



Re-use: spatial and temporal localities. Intuitively: n^3 accesses on $2 \times n^2$ matrices (if $n \times n$).



Cache Characteristics

- Hit ratio (behavior): fraction of references satisfied by the cache.
- Cache line (= bus width): granularity.
- Associativity (architecture): “collision list” to reduce cache eviction.
- For the matrix: $2n^2$ fetches from memory to *populate the cache*, and then n^3 direct accesses at full speed.

16-02-2007

Alexandre David, MVP'07

27

Data re-use is the keyword. Cache line: word granularity is too expensive and bad for spatial locality. 4 words usually for L2 (access to system bus), and internally 256-bit data bus for L1 \leftrightarrow L2 (8 words).

The term *cache eviction* is not mentioned in the book and is missing. It is an important notion to know.




Impact on Memory Bandwidth (and Latency)

- Access to successive words much better than random access.
 - Higher bandwidth (whole cache line at once)
 - Better latency (successive words already in cache)

Example: Strided Access



(a) Column major data access



(b) Row major data access.

Figure 2.2 Multiplying a matrix with a vector: (a) multiplying column-by-column, keeping a running sum; (b) computing each element of the result as a dot product of a row of the matrix with the vector.

- a) Add vectors (temporary results) to get final result.
 - b) Compute final result incrementally.
- Strided access (a) yields poor performance.



Other Approaches to Hide Latency

- Prefetching
 - but may evict useful data because cache is small.
- Multi-threading
 - but needs higher bandwidth because all the threads share the same bus.

16-02-2007

Alexandre David, MVP'07

30

Prefetching is like fetching the next cache line when possible (hardware decides), or same effect by reordering instructions (hardware or compiler) to issue loads long before usage. Works if consecutive words are accessed by consecutive instructions (spatial locality).


Multi-thread: switch to another thread when a thread stalls for data and keep the processor busy.

In fact, both solutions address the latency problem and exacerbate the bandwidth problem. That was probably the design idea behind RAMBUS, though no multi-threading at the time to use it! + the fact that latency was way higher than other systems.


Multi-threading

1 thread/dot product

BUT: need more bandwidth!



B



A

i

C

16-02-2007

Alexandre David, MVP'07

31

Software: need to create the threads explicitly.

Hardware: need support for multi-thread.



Summary on Memory

- Exploit spatial and temporal locality in programs. For sequential and parallel programs!
- Operations/memory accesses is a good indicator of tolerance to memory bandwidth.
- Processing is cheap, memory is expensive.



Dichotomy of Parallel Computing Platforms

- Logical organization: programmer's view.
- Physical organization: actual hardware.
- Two critical components:
 - expressing parallel tasks (control structure)
 - specifying interaction between them (communication model).

16-02-2007

Alexandre David, MVP'07

33

The 2 critical components both through logical and physical organizations.



Control Structure

- Parallelism can be expressed at different levels of granularity
 - from instruction level parallelism
 - to processes.
- SIMD: single instruction stream, multiple data stream.
- MIMD: multiple instruction stream ...

16-02-2007

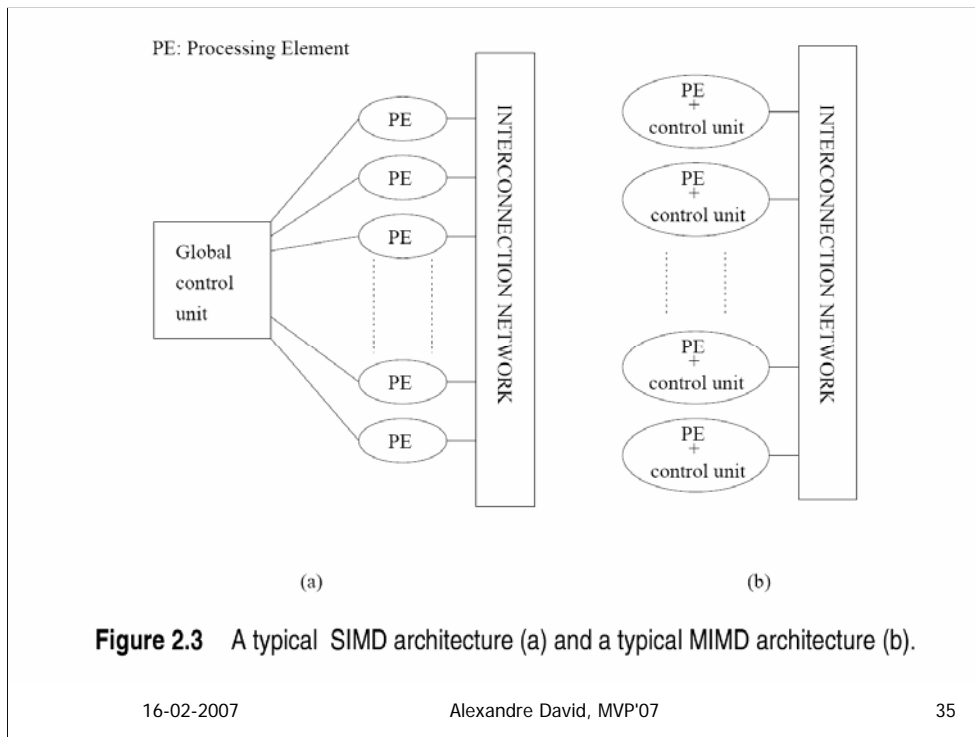
Alexandre David, MVP'07

34

Processing units in parallel computers either operate under the centralized control of a single control unit or work independently.

•SIMD: A single control unit dispatches the same instruction to various processors or functional units.

•MIMD: Each processor has its own control unit and can execute different instructions on different data items.



PE: processing element.

Typically, SIMD implemented as special instruction sets on processors (MMX, SSE, SSE2, 3DNow!), and MIMD implemented as multiprocessor machines or clusters.

SIMD relies on the regular structure of computations (such as those in image processing).

A variant of MIMD, called single program multiple data streams (SPMD) executes the same program on different processors, which is often the case in practice.



Communication Model: Shared Address Space

- Memory shared between several processors.
 - NUMA different access time
 - UMA same access time.
 - Cases with local cache considered UMA.
- Easier programming, one address space
 - but cache coherence mechanisms needed,
 - But need to solve contention (writes).

16-02-2007

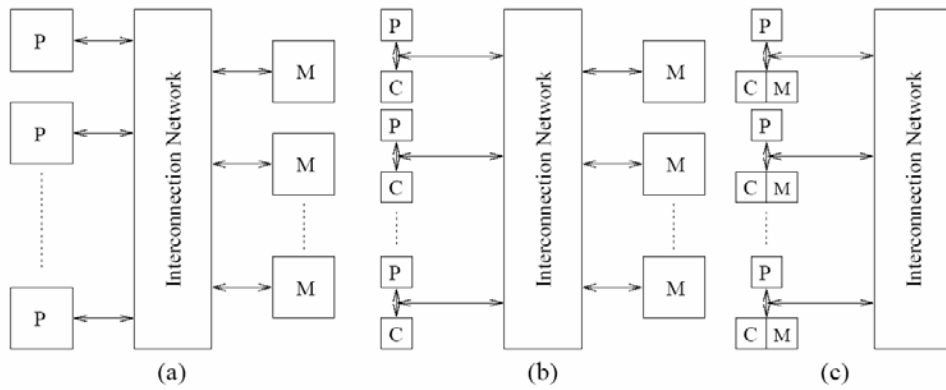
Alexandre David, MVP'07

36


NUMA: non-uniform memory access. Cheaper and easier to implement. Need locality to perform well.

UMA: uniform memory access. Performant uniform access expensive.

UMA vs. NUMA



Communication Model: Shared Address Space



Implemented as shared memory computers
or distributed memory computers.

16-02-2007

Alexandre David, MVP'07

38



Message-Passing Platforms

- Memory private to processors.
- Interaction via messages
 - Send/receive primitives.
 - MPI libraries.
- Hardware needed: good network interconnect.

16-02-2007

Alexandre David, MVP'07

39

Cheap and popular solution: cluster of MP machines connected via high bandwidth network.