### Paradigmas de Computação Paralela

### **Optimising performance (MPI)**

João Luís Ferreira Sobral Departamento do Informática Universidade do Minho

21 Nov 2017



#### **Performance models**

• Make it possible to compare algorithms, scalability and the identification of bottlenecks before a considerable time is invested in implementation

#### What is the definition of performance?

- There are multiple alternatives:
  - Execution time, efficiency, scalability, memory requirements, throughput, latency, project costs / development costs, portability, reuse potential
  - The importance of each one depends on the concrete application
- Most common measure in parallel applications (*speed-up*): tseq/tpar

#### Amdahl law

- The sequential component of an application limits the maximum speed-up
  - If **s** is the sequential faction of an algorithm then the maximum possible gain is **1**/**s**.
- Reinforces the idea that we should prefer algorithms suitable for parallel execution: *think parallel*.



#### **Performance models**

- Should explain observations and predict behaviour
  - Defined as a function of the problem dimension, number of processors, number of tasks, etc.
- Execution time
  - Time measured since the first process (thread) starts execution until the last process (thread) terminates
  - Texec = Tcomp + Tcomm + Tfree
    - Computation time time spent in computations, excluding communication/ synchronization and free time.
      - The sequential version can be used to estimate Tcomp.
    - Free time when a processor becomes starved (without work)
      - Can be complex to measure since it depends on the order of tasks
      - Can be minimized with adequate load distribution and/or overlapping computation and communication



#### **Performance models(cont.)**

- **Communication time** time that processes spend sending/receiving data.
  - Computed using communication latency (ts) and throughput (1/tw):
    - Tmens = ts + twL

ts and tw can be obtained experimentally, by a ping-pong test and a linear regression.

Tamnho	MPI (lat)	MPI (bw)	MPI-MX (lat)	MPI-MX (bw)	MPP (lat)	MPP (bw)	KaRMI (lat)	KaRMI (bw)
1	31	0,3	3,4	2,3	46	0,2	63	0,1
8	31	2,6	3,4	23,7	46	1,7	75	1,1
100	31	25,6	4,7	168,7	48	16,6	75	10,7
1000	55	146,3	8,1	983,9	94	106,4	200	40,0
10000	258	310,3	18,4	4355,9	279	286,0	387	206,0
100000	1136	703,8	125,5	6373,0	1017	786,5	1137	703,0
1E+06	9859	811,4	970,2	8246,0	8282	953,2	9787	817,0
ts (us)	31		3		46		63	
tw (us)	0,010		0,001		0,008		0,010	





#### **Performance models – Example Jacobi Method**

• Iterative method, at each iteration the new matrix value is computed as the average neighbour values

$$X_{i,j}^{(t+1)} = aX_{i,j}^{(t)} + b(X_{i-1,j}^{(t)} + X_{i,j-1}^{(t)} + X_{i+1,j}^{(t)} + X_{i,j+1}^{(t)})$$



Computação Paralela



#### **Performance models – Example Jacobi Method**

 Execution time for each iteration in Jacobi method, for a NxN matrix, on P processors, a partition by columns and N/P columns per processor.

**T***comp* = operations per element *x* num. elements per processor *x* t*c* 

= 6 x (N x N/P) x tc (tc = time for a single operation)

 $= 6tcN^2/P$ 

**T***comm* = messages per processor *x* time required for each message

= 2 x (ts + twN)

Tfree = 0, since in this problem the workload is well distributed

#### Texec = Tcomp + Tcomm + Tfree

 $= 6tcN^{2}/P + 2ts + 2twN$ 

 $= O(N^2/P+N)$ 



#### **Performance models – Example (cont)**

- In certain cases, execution time may not be the most adequate performance measure.
- Speed-up and efficiency are two related metrics.
- **Speed-up (G)** indicates the reduction in execution time attained in P processors
  - Ratio between the best sequential algorithm and the execution time of the parallel version

speed-up = Tseq / Tpar,

• **Efficiency** (E) gives the faction of time that processors perform useful work:

E = Tseq / (P x Tpar)

• Jacobi case:

$$G = \frac{6t_{c}N^{2}P}{6t_{c}N^{2} + 2Pt_{s} + 2Pt_{w}N} \qquad E = \frac{6t_{c}N^{2}}{6t_{c}N^{2} + 2Pt_{s} + 2Pt_{w}N}$$

Computação Paralela



#### **Scalability analysis**

- Execution time, speed-up and efficiency can be used for quantitative analysis of performance
- Jacobi example:
  - Execution time decreases when P increases, but it is limited by the time to exchange two lines
- $Texe = 6tcN^2/P + 2ts + 2twN$
- Execution time increases with N, tc, ts e tw
- Efficiency decreases when P, ts e tw increase
- Efficiency increases with N and Tc;

$$E = \frac{6t_c N^2}{6t_c N^2 + 2Pt_s + 2Pt_w N}$$

- Scalability for problems with fixed size (strong scalability)
  - Analysis of Texec and E when P increases
  - In general, *E* decreases. Texec can increase if it has a positive power of P.
- Scalability for problems with variable size (weak scalability)
  - In some cases, more processors are used to solve larger problems, keeping the same efficiency levels
  - **Isoefficiency** indicates what is the required increase in the problem dimension, to keep the same efficiency, when the number of processors increases

















### **Measuring time in MPI**

- Time functions in MPI
  - double MPI\_Wtime() returns the wall time (high resolution)
  - double MPI\_Wtick() returns the clock resolution (in seconds)
- Wall time can differ from process to process
  - There is no notion of "global time"
    - Each machine provides a local wall time
  - Application execution time should be wall time of the slowest process
  - <u>Note</u>: in some parallel algorithms process termination is not trivial



#### **Experimental study and evaluation of implementations**

- Parallel computing has a strong experimental component
  - Many problems are too complex for a realization only based on models
  - Performance model can be calibrated with experimental data (e.g., Tc)
- How to ensure that result are precise and reproducible?
  - Perform multiple experiments and verify clock resolution
  - Results should not change among in small difference: less than 2-3%
- Execution profile:
  - Gather several performance data: number of messages, data volume transmited
  - Can be implemented by specific tools or by directly instrumenting the code
    - There is always an overhead introduced in the base application
- Speed-up anomalies
  - superlinear (superior to the number of processors) in most cases it is due the cache effect





#### Technique to measure the application time-profile (profiling)

- **Polling**: the application is periodically interrupted to collect performance data
- **Instrumentation**: code is introduced (by the programmer or by tools) to collect performance data about useful events
- Instrumentation tends do produce better results but also produces more interference (e.g., overhead)
- Exemple: vampir





#### Distributed memory (MPI) vs Shared memory (OpenMP) optimisation

- Distributed memory vs shared memory
  - Data placement is explicit (vs implicit)
  - Static scheduling is preferred (vs dynamic)
  - Synchronization is costly (only performed by global barriers & message send)
- Improve scalability on distributed memory
  - Minimise communication among processes
    - Eventually duplicating computation
  - Minimise idle (free) time with a good load distribution
- Practical advise
  - Measure communication overhead
  - Measure load balance
  - Avoid centralised control