

# Welcome to MCS 572

- 1 About the Course
  - content and organization
  - expectations of the course
- 2 Supercomputing
  - definition and classification
- 3 Measuring Performance
  - speedup and efficiency
  - Amdahl's Law
  - Gustafson's Law
  - quality up

MCS 572 Lecture 1  
Introduction to Supercomputing  
Jan Verschelde, 26 August 2024

# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing

- definition and classification

## 3 Measuring Performance

- speedup and efficiency
- Amdahl's Law
- Gustafson's Law
- quality up

# Catalog Description

Introduction to supercomputing on vector and parallel processors; architectural comparisons, parallel algorithms, vectorization techniques, parallelization techniques, actual implementation on real machines.

Prerequisites:

- 1 Numerical Analysis (MCS 471) or Numerical Analysis of Partial Differential Equations (MCS 571).
- 2 Sufficient computer literacy.  
Scientific software (MCS 507) prepares for this course, but is not a prerequisite.
- 3 Or (if in doubt about the above) consent of the instructor.

MCS 572 is one of the courses on the computational science prelim.

# Content of the Course

Two recommended text books:

**Barry Wilkinson and Michael Allen:** *Parallel Programming. Techniques and Applications Using Networked Workstations and Parallel Computers.* 2nd Edition. Prentice-Hall 2005.

**David B. Kirk and Wen-mei W. Hwu:** *Programming Massively Parallel Processors. A Hands-on Approach.* Elsevier 2010.  
Fourth Edition, 2023, with Izzat El Hajj as 3rd author.

Parallel programming goals:

- 1 design and analysis of parallel programs;
- 2 implementation using MPI, OpenMP, and threads;
- 3 application to scientific problems.

This is a *computational*, not a programming course.

Although the course is hands-on, there are many theoretical aspects.

# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing

- definition and classification

## 3 Measuring Performance

- speedup and efficiency
- Amdahl's Law
- Gustafson's Law
- quality up

# Organization and Expectations

Three different types of parallel algorithms:

- 1 using Message-Passing Interface (MPI) for clusters,
- 2 for shared memory: pthreads and OpenMP,
- 3 programming Graphics Processing Units (GPUs) using CUDA (Compute Unified Device Architecture) of NVIDIA.

Activities throughout the semester:

- several homework collections,
- the midterm exam (could be replaced by one homework),
- three computer projects.

The first two computer projects will be on prescribed topics and may be solved in pairs. The third project must be done individually and could form the basis for a project presentation at the end.

# on the use of programming languages

This is not a programming, but a *computational* course.

- 1 Python allows for high level parallel programming.  
`mpi4py` enables distributed memory parallel programming.
- 2 Julia has a MATLAB-like syntax and offers support for
  - 1 message passing via `MPI.jl`,
  - 2 multithreading and multitasking,
  - 3 vendor agnostic GPU acceleration.
- 3 C++ achieves performance portability.

Observe on the use: we are using programming languages to run short programs, or code from software libraries.

# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing

- definition and classification

## 3 Measuring Performance

- speedup and efficiency
- Amdahl's Law
- Gustafson's Law
- quality up



# what is a supercomputer?

Supercomputing = use of a supercomputer (also called high performance computing).

## Definition

A *supercomputer* is a computing system (hardware, system & application software) that provides close to the best currently achievable sustained performance on demanding computational problems.

Classification at [www.top500.org](http://www.top500.org).

A *flop* is a floating point operation. Performance is often measured in number of flops per second.

If two flops can be done per clock cycle, then a processor at 3GHz can theoretically perform 6 billion flops (6 gigaflops) per second.

All computers in the top 10 achieve more than 1 petaflop per second.

# top 10 of November 2011

Rank	Site	Computer/Year Vendor	Cores	R <sub>max</sub>	R <sub>peak</sub>	Power
1	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIx 2.0GHz, Tofu interconnect / 2011 Fujitsu	705024	10510.00	11280.38	12659.9
2	National Supercomputing Center in Tianjin China	NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 / 2010 NUDT	186368	2566.00	4701.00	4040.0
3	DOE/SC/Oak Ridge National Laboratory United States	Cray XT5-HE Opteron 6-core 2.6 GHz / 2009 Cray Inc.	224162	1759.00	2331.00	6950.0
4	National Supercomputing Centre in Shenzhen (NSCS) China	Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 / 2010 Dawning	120640	1271.00	2984.30	2580.0
5	GSIC Center, Tokyo Institute of Technology Japan	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows / 2010 NEC/HP	73278	1192.00	2287.63	1398.6
6	DOE/NNSA/LANL/SNL United States	Cray XE6, Opteron 6136 8C 2.40GHz, Custom / 2011 Cray Inc.	142272	1110.00	1365.81	3980.0
7	NASA/Ames Research Center/NAS United States	SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon 5570/5670 2.93 Ghz, Infiniband / 2011 SGI	111104	1088.00	1315.33	4102.0
8	DOE/SC/LBNL/NERSC United States	Cray XE6, Opteron 6172 12C 2.10GHz, Custom / 2010 Cray Inc.	153408	1054.00	1288.63	2910.0
9	Commissariat a l'Energie Atomique (CEA) France	Bull bulx super-node S6010/S6030 / 2010 Bull	138368	1050.00	1254.55	4590.0
10	DOE/NNSA/LANL United States	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband / 2009 IBM	122400	1042.00	1375.78	2345.0

# top 10 of November 2013

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Super Computer Center in Guangzhou China	<b>Tianhe-2 (MilkyWay-2)</b> - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory United States	<b>Titan</b> - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL United States	<b>Sequoia</b> - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	<b>K computer</b> , SPARC64 Villifx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory United States	<b>Mira</b> - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
6	Swiss National Supercomputing Centre (CSCS) Switzerland	<b>Piz Daint</b> - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,325
7	Texas Advanced Computing Center/Univ. of Texas United States	<b>Stampede</b> - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462,462	5,168.1	8,520.1	4,510
8	Forschungszentrum Juelich (FZJ) Germany	<b>JUQUEEN</b> - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	458,752	5,008.9	5,872.0	2,301
9	DOE/NNSA/LLNL United States	<b>Vulcan</b> - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	393,216	4,293.3	5,033.2	1,972
10	Leibniz Rechenzentrum Germany	<b>SuperMUC</b> - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM	147,456	2,897.0	3,185.1	3,423

# top 10 of June 2016

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Supercomputing Center in Wuxi China	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
2	National Super Computer Center in Guangzhou China	Tianhe-2 (MilkyWay-2) - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 3151P NUDT	3,120,000	33,862.7	54,902.4	17,808
3	DOE/SC/Oak Ridge National Laboratory United States	Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
4	DOE/NNSA/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
5	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,660
6	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
7	DOE/NNSA/LANL/SNL United States	Trinity - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	301,056	8,100.9	11,078.9	
8	Swiss National Supercomputing Centre (CSCS) Switzerland	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,325
9	HLRS - Höchstleistungsrechenzentrum Stuttgart Germany	Hazel Hen - Cray XC40, Xeon E5-2680v3 12C 2.5GHz, Aries interconnect Cray Inc.	185,088	5,640.2	7,403.5	
10	King Abdullah University of Science and Technology Saudi Arabia	Shaheen II - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	196,608	5,537.0	7,235.2	2,834

# top 10 of November 2020

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DDE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	<b>Sierra</b> - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
4	<b>Sunway TaihuLight</b> - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
5	<b>Selene</b> - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63,460.0	79,215.0	2,646
6	<b>Tianhe-2A</b> - TH-IVB-FEP Cluster, Intel Xeon ES-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
7	<b>JUWELS Booster Module</b> - Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100, Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite, Atos Forschungszentrum Juelich (FZJ) Germany	449,280	44,120.0	70,980.0	1,764
8	<b>HPC5</b> - PowerEdge C6140, Xeon Gold 6252 24C 2.1GHz, NVIDIA Tesla V100, Mellanox HDR Infiniband, Dell EMC Eni S.p.A. Italy	669,760	35,450.0	51,720.8	2,252
9	<b>Frontiera</b> - Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR, Dell EMC Texas Advanced Computing Center/Univ. of Texas United States	448,448	23,516.4	38,745.9	
10	<b>Dammam-7</b> - Cray CS-Storm, Xeon Gold 6248 20C 2.5GHz, NVIDIA Tesla V100 SXM2, InfiniBand HDR 100, HPE Saudi Aramco Saudi Arabia	672,520	22,400.0	55,423.6	


# top 10 of November 2022

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	<b>Frontier</b> - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,730,112	1,102.00	1,685.65	21,100
2	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
3	<b>LUMI</b> - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288	309.10	428.70	6,016
4	<b>Leonardo</b> - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,463,616	174.70	255.75	5,610
5	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.97GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096
6	<b>Sierra</b> - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94.64	125.71	7,438
7	<b>Sunway TaihuLight</b> - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93.01	125.44	15,371
8	<b>Pertinutter</b> - HPE Cray EX235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10, HPE DOE/SC/LBNL/NERSC United States	761,856	70.87	93.75	2,589
9	<b>Selene</b> - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63.46	79.22	2,646
10	<b>Tianhe-2A</b> - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61.44	100.68	18,482

# top 5 of June 2024

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	<b>Frontier</b> - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, <b>HPE</b> DOE/SC/Oak Ridge National Laboratory United States	8,699,904	1,206.00	1,714.81	22,786
2	<b>Aurora</b> - HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11, <b>Intel</b> DOE/SC/Argonne National Laboratory United States	9,264,128	1,012.00	1,980.01	38,698
3	<b>Eagle</b> - Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR, <b>Microsoft Azure</b> Microsoft Azure United States	2,073,600	561.20	846.84	
4	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, <b>Fujitsu</b> RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
5	<b>LUMI</b> - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, <b>HPE</b> EuroHPC/CSC Finland	2,752,704	379.70	531.51	7,107

## Performance Development





# system terms and architectures

**core** for a CPU: unit capable of executing a thread,  
for a GPU: a streaming multiprocessor.

$R_{\max}$  maximal performance achieved on the LINPACK  
benchmark (solving a dense linear system) for problem  
size  $N_{\max}$ , measured in Gflop/s, Tflop/s, Pflop/s.

$R_{\text{peak}}$  theoretical peak performance.

**Power** total power consumed by the system.

Types of architectures, using

- commodity leading edge microprocessors running at their maximal clock and power limits;
- special processor chips running at less than maximal power to achieve high physical packaging densities;
- mix of chip types and accelerators (GPUs).

# main conferences in supercomputing

- The International Conference for High Performance Computing, Networking, Storage, and Analysis. SC Conference Series.  
<https://supercomputing.org>.  
Established in 1988, draws more than 10,000 attendees.
- The IEEE International Parallel and Distributed Processing Symposium. [www.ipdps.org](http://www.ipdps.org).  
This international conference meets every year in Spring.
- SIAM Conferences on Parallel Processing for Scientific Computing and Computational Science cover High Performance Computing.  
<https://www.siam.org/conferences/>  
SIAM PP happens in even years, SIAM CSE in odd years.

# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing

- definition and classification

## 3 Measuring Performance

- **speedup and efficiency**
- Amdahl's Law
- Gustafson's Law
- quality up

# speedup and efficiency

By  $p$  we denote the number of processors.

$$\text{Speedup } S(p) = \frac{\text{sequential execution time}}{\text{parallel execution time}}.$$

Another measure for parallel performance:

$$\text{Efficiency } E(p) = \frac{\text{speedup}}{\text{\#processors}} = \frac{S(p)}{p} \times 100\%.$$

In the best case, we hope:  $S(p) = p$  and  $E(p) = 100\%$ .

If  $E = 50\%$ , then on average processors are idle for half of the time.


## superlinear speedup

While we hope for  $S(p) = p$ , we may achieve  $S(p) > p$ .

*Example.* Sequential search in unsorted list.

A parallel search by  $p$  processors divides the list evenly in  $p$  sublists.

$p = 3$  :



The sequential search time depends on position in list.

The parallel search time depends on position in sublist.

⇒ huge speedup if at first element of last sublist.

# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing

- definition and classification


## 3 Measuring Performance

- speedup and efficiency
- **Amdahl's Law**
- Gustafson's Law
- quality up

## predicting speedup

Consider a job that takes time  $t$  on one processor.

Let  $R$  be the fraction of  $t$  that must be done sequentially,  $R \in [0, 1]$ .



$$\text{Speedup on } p \text{ processors } S(p) \leq \frac{t}{Rt + \frac{(1-R)t}{p}} = \frac{1}{R + \frac{1-R}{p}} \leq \frac{1}{R}.$$

# Amdahl's Law

## Theorem (Amdahl (1967))

Let  $R$  be the fraction of the operations which cannot be done in parallel. The speedup with  $p$  processors is bounded by  $\frac{1}{R + \frac{1-R}{p}}$ .

*Corollary.*  $S(p) \leq \frac{1}{R}$  as  $p \rightarrow \infty$ .

*Example.* Suppose 90% of the operations in an algorithm can be executed in parallel. What is the best speedup with 8 processors? What is the best speedup with an unlimited amount of processors?

$$p = 8: \frac{1}{\frac{1}{10} + (1 - \frac{1}{10}) \frac{1}{8}} = \frac{80}{17} \approx 4.7 \qquad p = \infty: \frac{1}{1/10} = 10$$



# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing


- definition and classification

## 3 Measuring Performance

- speedup and efficiency
- Amdahl's Law
- **Gustafson's Law**
- quality up

## scaled speedup

Consider a job that took time  $t$  on  $p$  processors.  
Let  $s$  be the fraction of  $t$  that is done sequentially.



$$\text{Scaled speedup } S_s(p) \leq \frac{st + p(1-s)t}{t} = s + p(1-s) = p + (1-p)s.$$

# Gustafson's Law

***The problem size scales with the number of processors!***

## Theorem (Gustafson's Law (1988))

*If  $s$  is the fraction of serial operations in a parallel program run on  $p$  processors, then the scaled speedup is bounded by  $p + (1 - p)s$ .*

*Example.* Suppose benchmarking reveals that 5% of time on a 64-processor machine is spent on one single processor (e.g.: root node working while all other processors are idle). Compute the scaled speedup.

$$p = 64, s = 0.05: S_s(p) \leq 64 + (1 - 64)0.05 = 64 - 3.15 = 60.85.$$

# Welcome to MCS 572

## 1 About the Course

- content and organization
- expectations of the course

## 2 Supercomputing

- definition and classification

## 3 Measuring Performance

- speedup and efficiency
- Amdahl's Law
- Gustafson's Law
- **quality up**

# quality up

More processing power often leads to better results.

- finer granularity of a grid  
e.g.: discretization of space and/or time in a differential equation
- greater confidence of estimates  
e.g.: enlarged number of samples in a simulation
- compute with larger numbers (multiprecision arithmetic)  
e.g.: solve an ill-conditioned linear system

$$\text{quality up } Q(p) = \frac{\text{quality on } p \text{ processors}}{\text{quality on 1 processor}}$$

$Q(p)$  measures improvement in quality using  $p$  procesors, keeping the computational time fixed.

## summary and recommended reading

We defined supercomputing, speedup, and efficiency.  
Gustafson's Law reevaluates Amdahl's Law.

Available to UIC via the ACM digital library:

- Jeannette M. Wing: **computational thinking**.  
*Communications of the ACM* 49(3):33-35, 2006.
- Peter M. Kogge and Timothy J. Dysart: **Using the TOP500 to trace and project technology and architecture trends**.  
In *SC'11 Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*.  
ACM 2011.
- John L. Gustafson: **Reevaluating Amdahl's Law**.  
*Communications of the ACM* 31(5):532-533, 1988.

# Exercises

- 1 How many processors whose clock speed runs at 3.0GHz does one need to build a supercomputer which achieves a theoretical peak performance of at least 4 Tflop/s? Justify your answer.
- 2 Suppose we have a program where 2% of the operations must be executed sequentially. According to Amdahl's law, what is the maximum speedup which can be achieved using 64 processors? Assuming we have an unlimited number of processors, what is the maximal speedup possible?
- 3 Benchmarking of a program running on a 64-processor machine shows that 2% of the operations are done sequentially, i.e.: that 2% of the time only one single processor is working while the rest is idle. Use Gustafson's law to compute the scaled speedup.
- 4 Visit <https://acer.uic.edu/> and estimate the theoretical peak performance of UIC's supercomputer.