High-Performance Scientific Computing

INFO 0939, Fall 2023

http://people.montefiore.ulg.ac.be/geuzaine/INFO0939/

Based in part on material from Victor Eijkhout's SSC 335/394 course at the Texas Advanced Computing Center

Introduction

Mathematics & Science

- In science, we use mathematics to understand physical systems.

- Different fields of science explore different 'domains' of the universe, and have their own sets of equations, encapsulated in theories.

- Determining the theories and governing equations requires observation or experimentation, and testing hypotheses.

THE GRAND CHALLENGE EQUATIONS

$$B_i A_i = E_i A_i + \rho_i \sum_j B_j A_j F_{ji}$$

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t}$$

$$\vec{F} = m\vec{a} + \frac{dm}{dt}\vec{v}$$

$$dU = \left(\frac{\partial U}{\partial S}\right)_V dS + \left(\frac{\partial U}{\partial V}\right)_S dV$$

$$\nabla \cdot \vec{D} = \rho$$

$$Z = \sum_j g_j e^{-E_j/kT}$$

$$\nabla \times \vec{H} = \frac{\partial \vec{D}}{\partial t} + \vec{J}$$

$$F_j = \sum_{k=0}^{N-1} f_k e^{2\pi i jk/N}$$

$$\nabla^2 u = \frac{\partial u}{\partial t}$$

$$P(t) = \frac{\sum_i W_i B_i(t) P_i}{\sum_i W_i B_i(t)}$$

$$\nabla \cdot \vec{B} = 0$$

$$p_{n+1} = r\, p_n (1 - p_n)$$

$$-\frac{h^2}{8\pi^2 m}\nabla^2 \Psi(r,t) + V\,\Psi(r,t) = -\frac{h}{2\pi i}\frac{\partial \Psi(r,t)}{\partial t}$$

$$-\nabla^2 u + \lambda u = f$$

$$\frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla)\vec{u} = -\frac{1}{\rho}\nabla p + \gamma \nabla^2 \vec{u} + \frac{1}{\rho}\vec{F}$$

$$\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2} = f$$

• NEWTON'S EQUATIONS • SCHROEDINGER EQUATION (TIME DEPENDENT) • NAVIER-STOKES EQUATION •
• POISSON EQUATION • HEAT EQUATION • HELMHOLTZ EQUATION • DISCRETE FOURIER TRANSFORM •
• MAXWELL'S EQUATIONS • PARTITION FUNCTION • POPULATION DYNAMICS •
• COMBINED 1ST AND 2ND LAWS OF THERMODYNAMICS • RADIOSITY • RATIONAL B-SPLINE •

[Courtesy of San Diego Supercomputer Center]

Scientific Computing

Why should we care about scientific computing?

- Computational research has emerged to complement experimental methods in basic research, design, optimization, and discovery in all facets of engineering and science
- In certain cases, computational simulations are the only possible approach to analyze a problem:
  - Experiments may be cost prohibitive (eg. *flight testing a 1,000 fuselage/wing-body configurations for a modern fighter aircraft*)
  - Experiments may be impossible (eg. *interaction effects between the International Space Station and Shuttle during docking*)
- Simulation capabilities rely heavily on the underlying compute power (e.g. amount of memory, total compute processors, and processor performance)
  - Fostered the introduction and development of *super-computers* starting in the 1960's
  - Large-scale compute power is tracked around the world via the *Top500 List* (more on that later)

Scientific Computing: a definition

"The efficient computation of constructive methods in applied mathematics"

- Applied math: getting results out of application areas
- Numerical analysis: results need to be correctly and efficiently computable
- Computing: the algorithms need to be implemented on modern hardware

Examples of Scientific Computing

(it really is everywhere)

Material Science

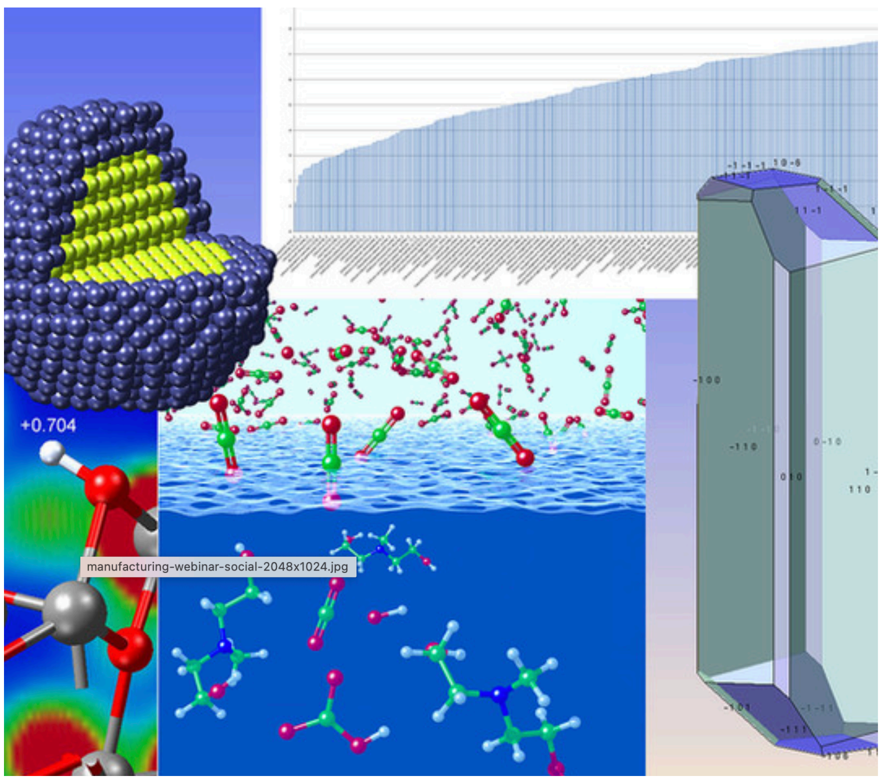

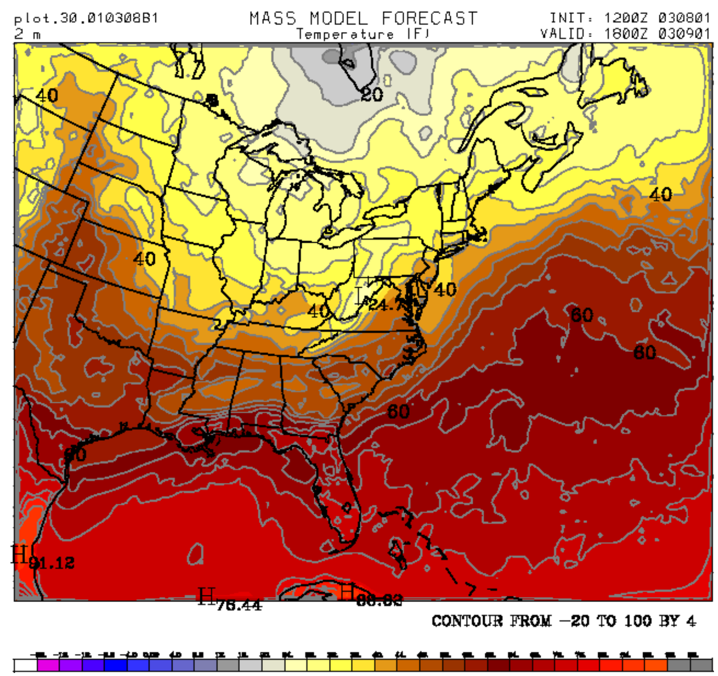

Weather Forecasting

Examples of Scientific Computing
(it really is everywhere)

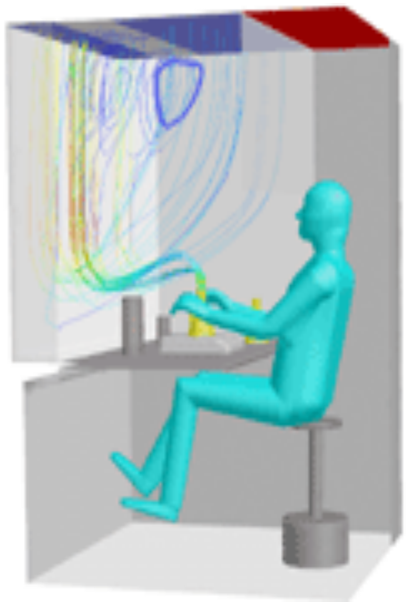

Streamlines for workstation ventilation

Heating, ventilation, and air conditioning

Aerospace

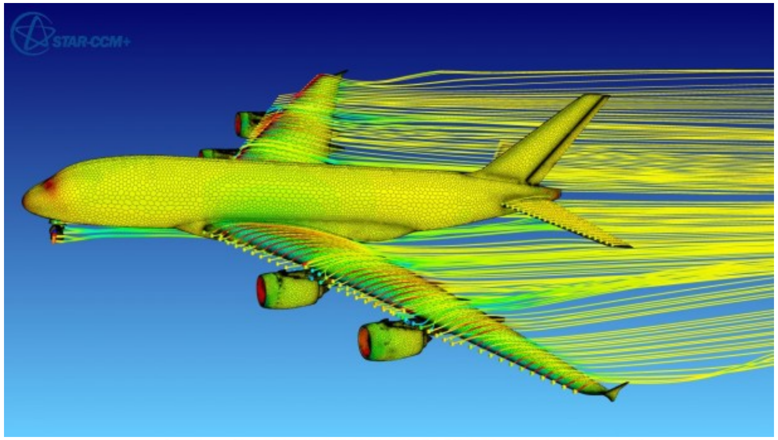

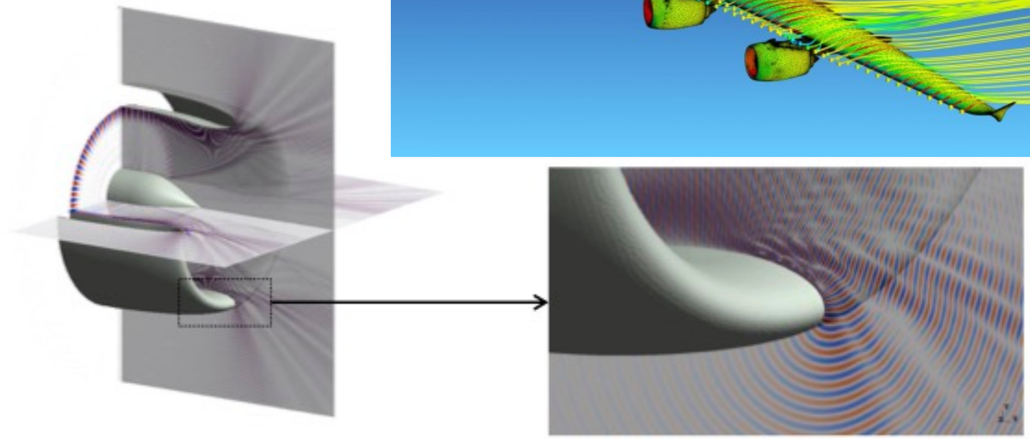

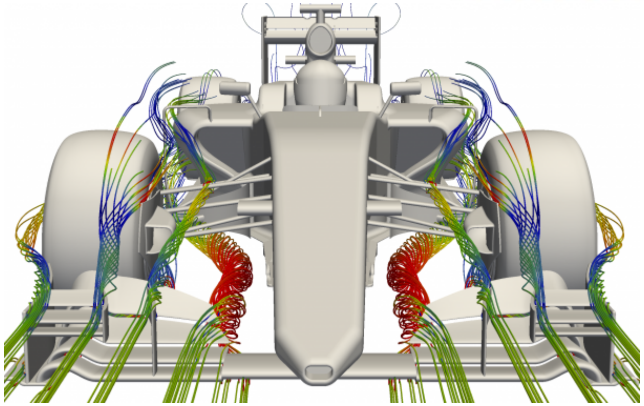

Automotive

Examples of Scientific Computing

(*it really is everywhere*)

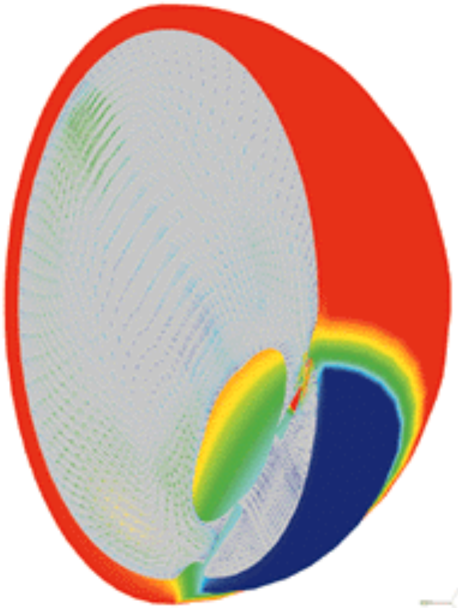

Temperature and natural convection currents in the eye following laser heating.

Biomedical engineering

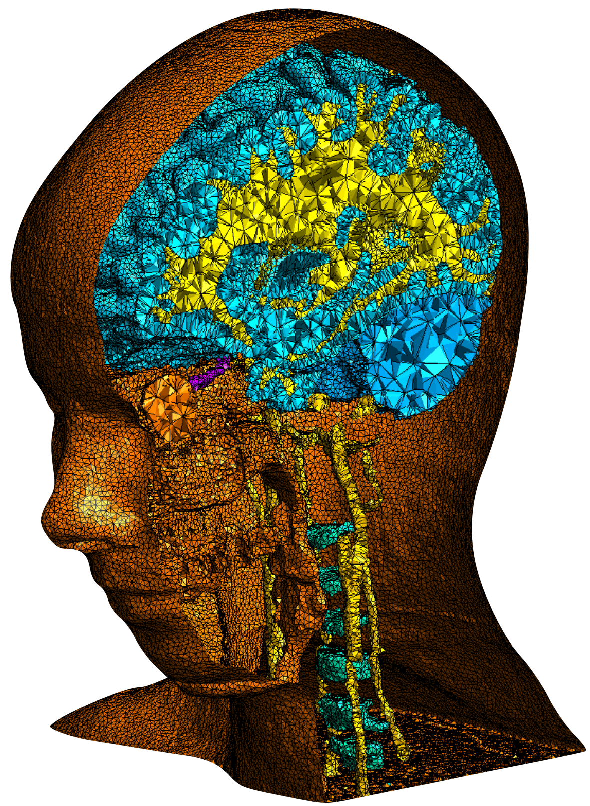

Fields induced in human body close to power lines

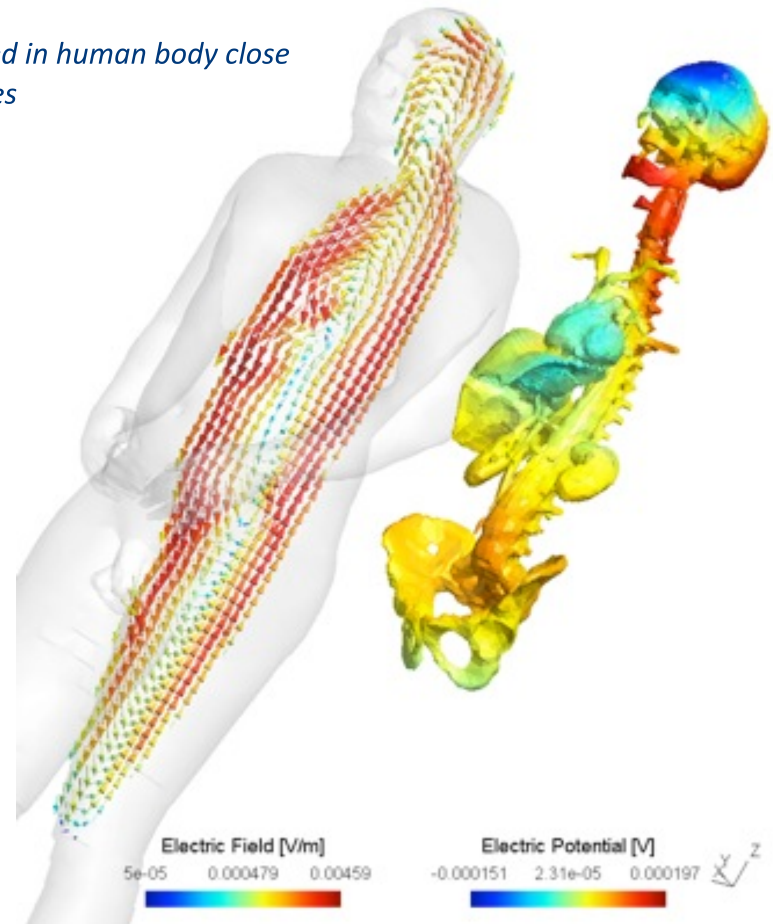

Examples of Scientific Computing

(*it really is everywhere*)

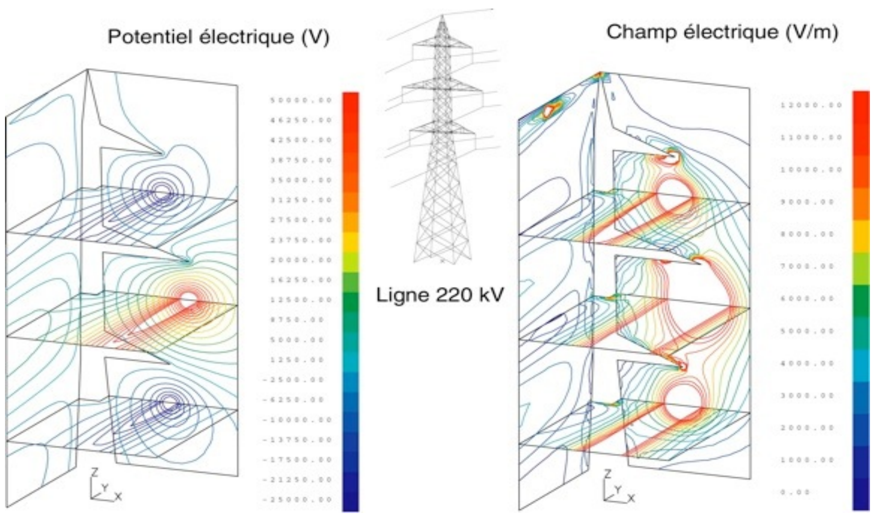

Electrical Engineering

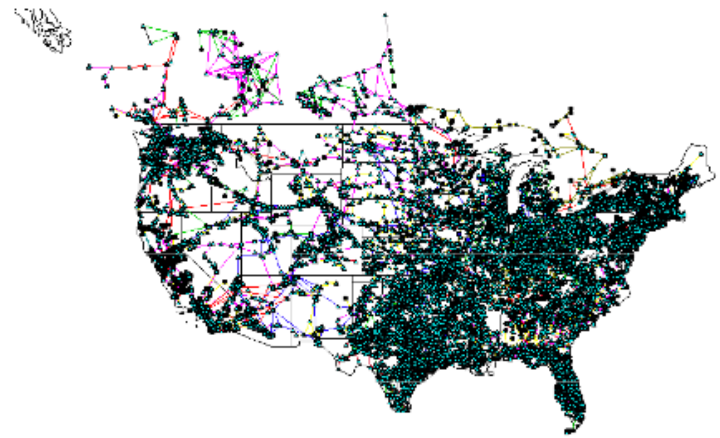


The New York Times
Thursday, September 4, 2008

Report on Blackout Is Said To Describe Failure to React

By MATTHEW L. WALD
Published: November 12, 2003

A report on the Aug. 14 blackout identifies specific lapses by various parties, including FirstEnergy's failure to react properly to the loss of a transmission line, people who have seen drafts of it say.

A working group of experts from eight states and Canada will meet in private on Wednesday to evaluate the report, people involved in the

E-MAIL
PRINT
SINGLE-PAGE
REPRINTS
SAVE
SHARE

Examples of Scientific Computing
(*it really is everywhere*)

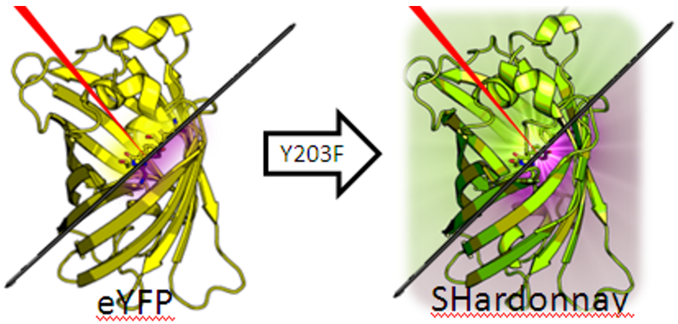

Fluorescent proteins for medical diagnostics

From the nanoscale...

Astronomy and planet habitability

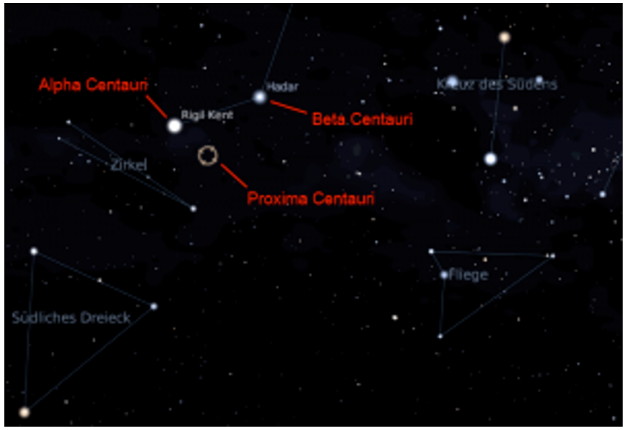

... to the infinitely large

... to day-to-day living ...

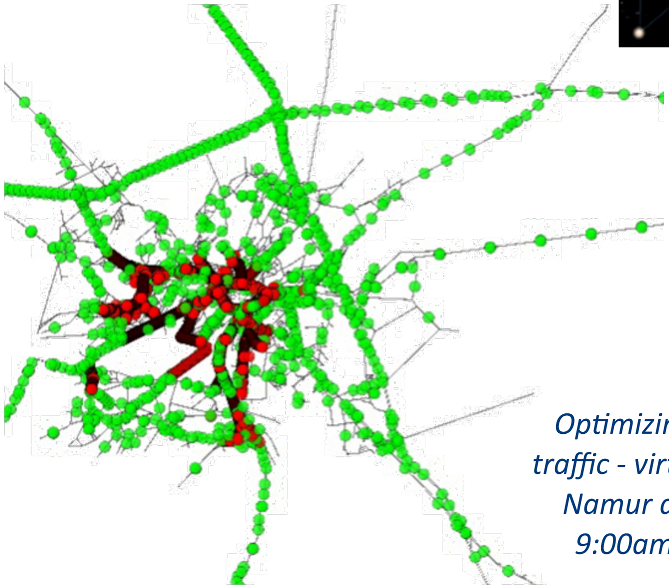

Optimizing traffic - virtual Namur at 9:00am

The Top500 List

- http://www.top500.org
- Owner submitted benchmark performance since 1993
  - based on a dense linear system solve
  - http://www.netlib.org/benchmark/hpl/

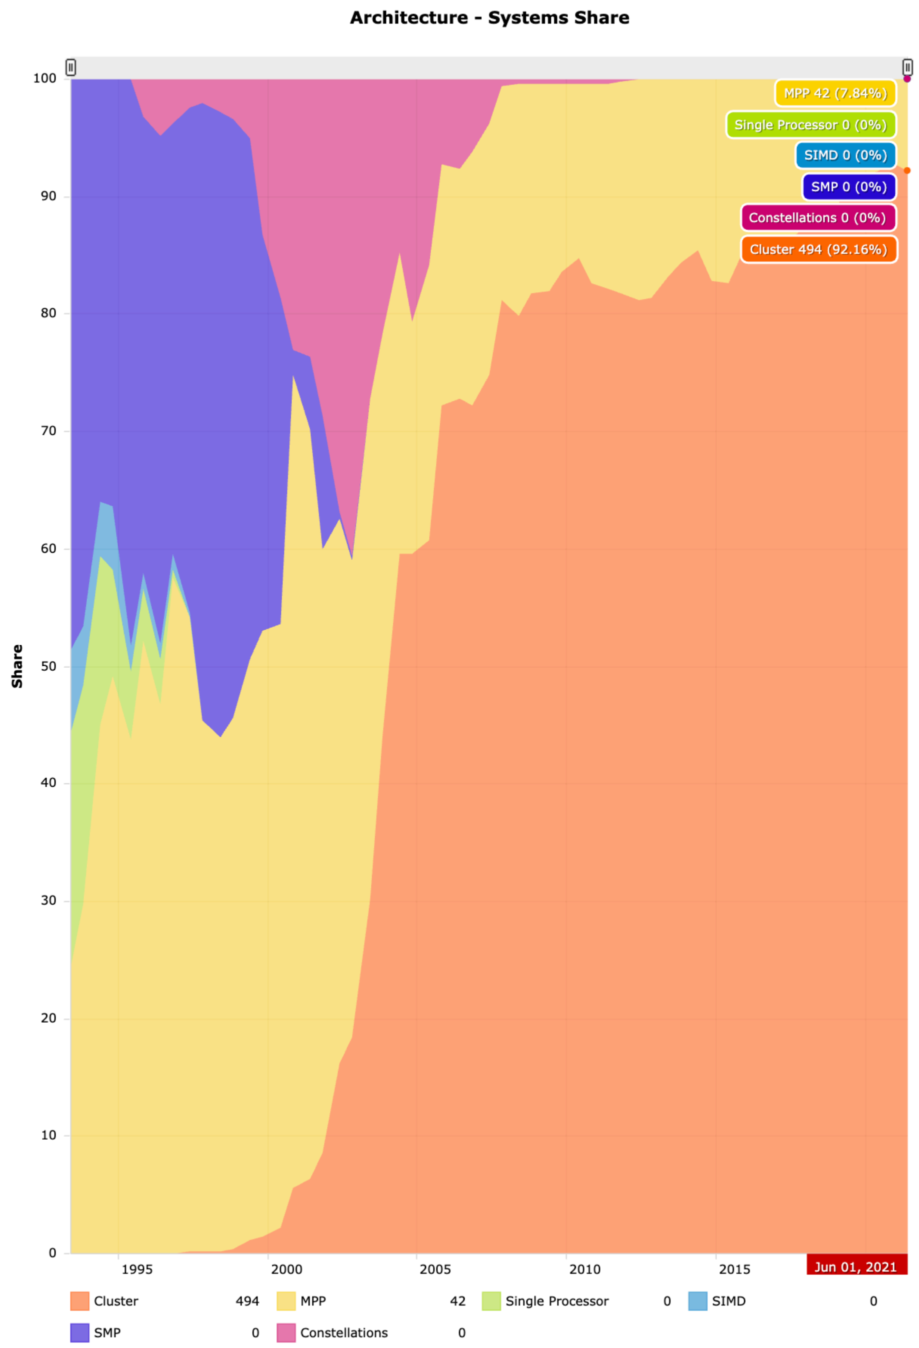

Architecture - Systems Share
Share
MPP 42 (7.84%)
Single Processor 0 (0%)
SIMD 0 (0%)
SMP 0 (0%)
Constellations 0 (0%)
Cluster 494 (92.16%)
Jun 01, 2021
Cluster 494
MPP 42
Single Processor 0
SIMD 0
SMP 0
Constellations 0

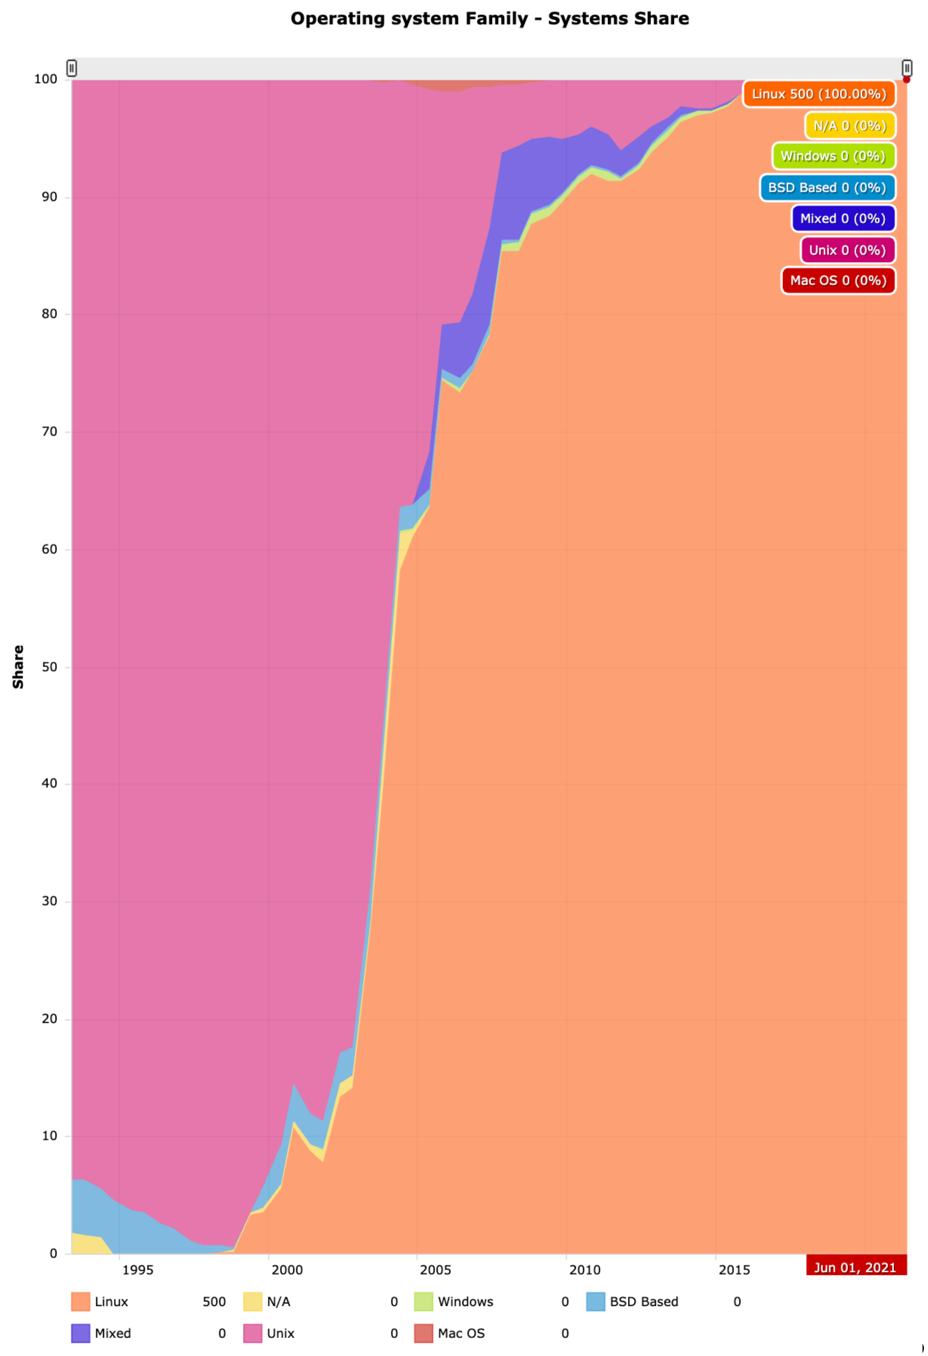

Operating system Family - Systems Share
Share
Linux 500 (100.00%)
N/A 0 (0%)
Windows 0 (0%)
BSD Based 0 (0%)
Mixed 0 (0%)
Unix 0 (0%)
Mac OS 0 (0%)
Jun 01, 2021
Linux 500
N/A 0
Windows 0
BSD Based 0
Mixed 0
Unix 0
Mac OS 0

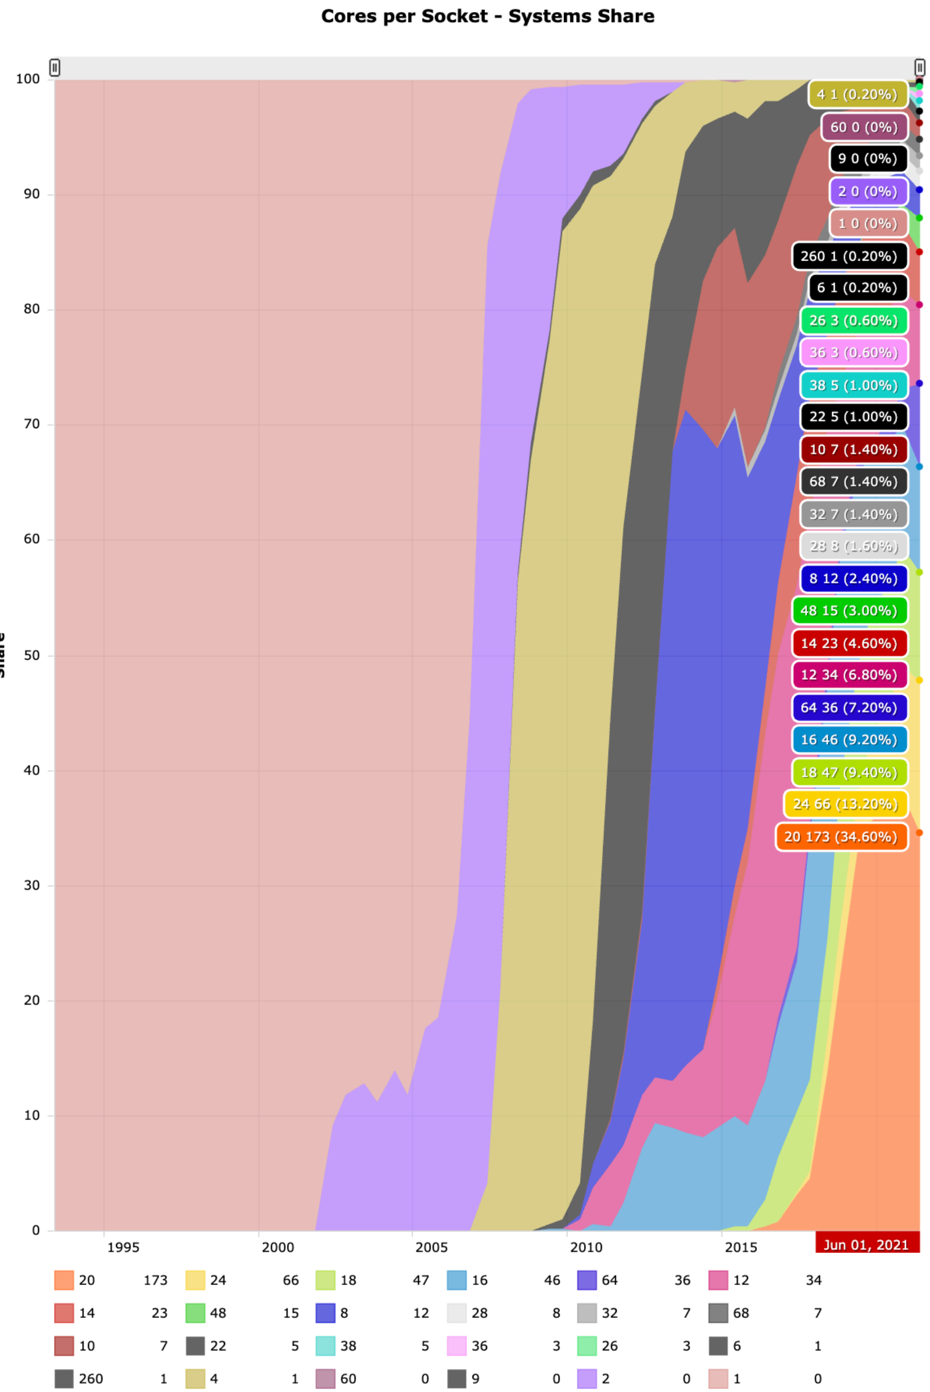
Cores per Socket - Systems Share
Share
4 1 (0.20%)
60 0 (0%)
9 0 (0%)
2 0 (0%)
1 0 (0%)
260 1 (0.20%)
6 1 (0.20%)
26 3 (0.60%)
36 3 (0.60%)
38 5 (1.00%)
22 5 (1.00%)
10 7 (1.40%)
68 7 (1.40%)
32 7 (1.40%)
28 8 (1.60%)
8 12 (2.40%)
48 15 (3.00%)
14 23 (4.60%)
12 34 (6.80%)
64 36 (7.20%)
16 46 (9.20%)
18 47 (9.40%)
24 66 (13.20%)
20 173 (34.60%)
Jun 01, 2021
20 173
24 66
18 47
16 46
64 36
12 34
14 23
48 15
8 12
28 8
32 7
68 7
10 7
22 5
38 5
36 3
26 3
6 1
260 1
4 1
60 0
9 0
2 0
1 0

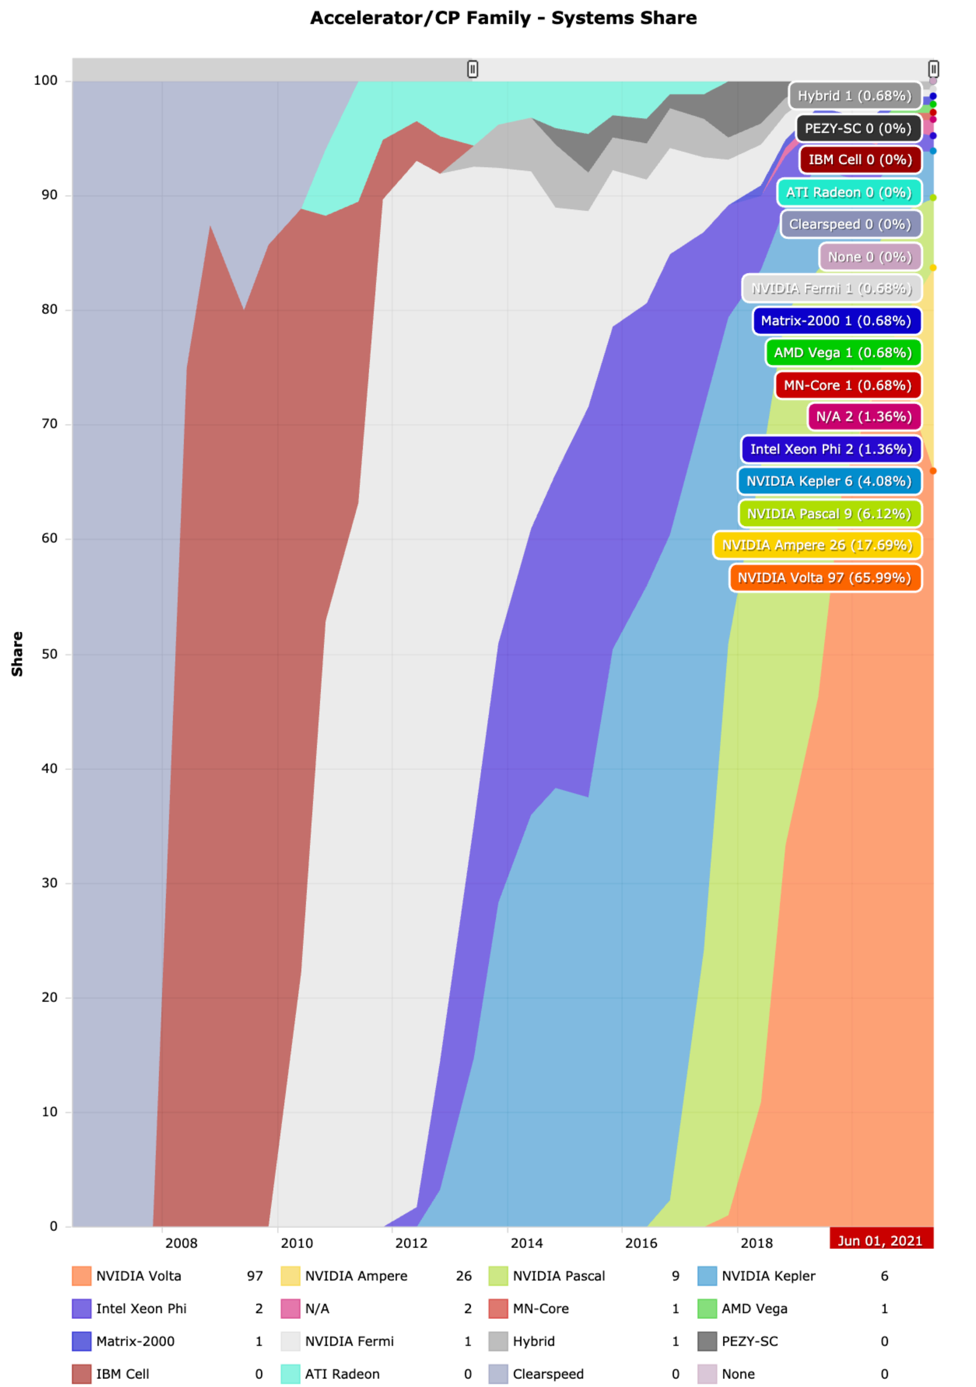
Accelerator/CP Family - Systems Share
Share
Hybrid 1 (0.68%)
PEZY-SC 0 (0%)
IBM Cell 0 (0%)
ATI Radeon 0 (0%)
Clearspeed 0 (0%)
None 0 (0%)
NVIDIA Fermi 1 (0.68%)
Matrix-2000 1 (0.68%)
AMD Vega 1 (0.68%)
MN-Core 1 (0.68%)
N/A 2 (1.36%)
Intel Xeon Phi 2 (1.36%)
NVIDIA Kepler 6 (4.08%)
NVIDIA Pascal 9 (6.12%)
NVIDIA Ampere 26 (17.69%)
NVIDIA Volta 97 (65.99%)
Jun 01, 2021
NVIDIA Volta 97
NVIDIA Ampere 26
NVIDIA Pascal 9
NVIDIA Kepler 6
Intel Xeon Phi 2
N/A 2
MN-Core 1
AMD Vega 1
Matrix-2000 1
NVIDIA Fermi 1
Hybrid 1
PEZY-SC 0
IBM Cell 0
ATI Radeon 0
Clearspeed 0
None 0

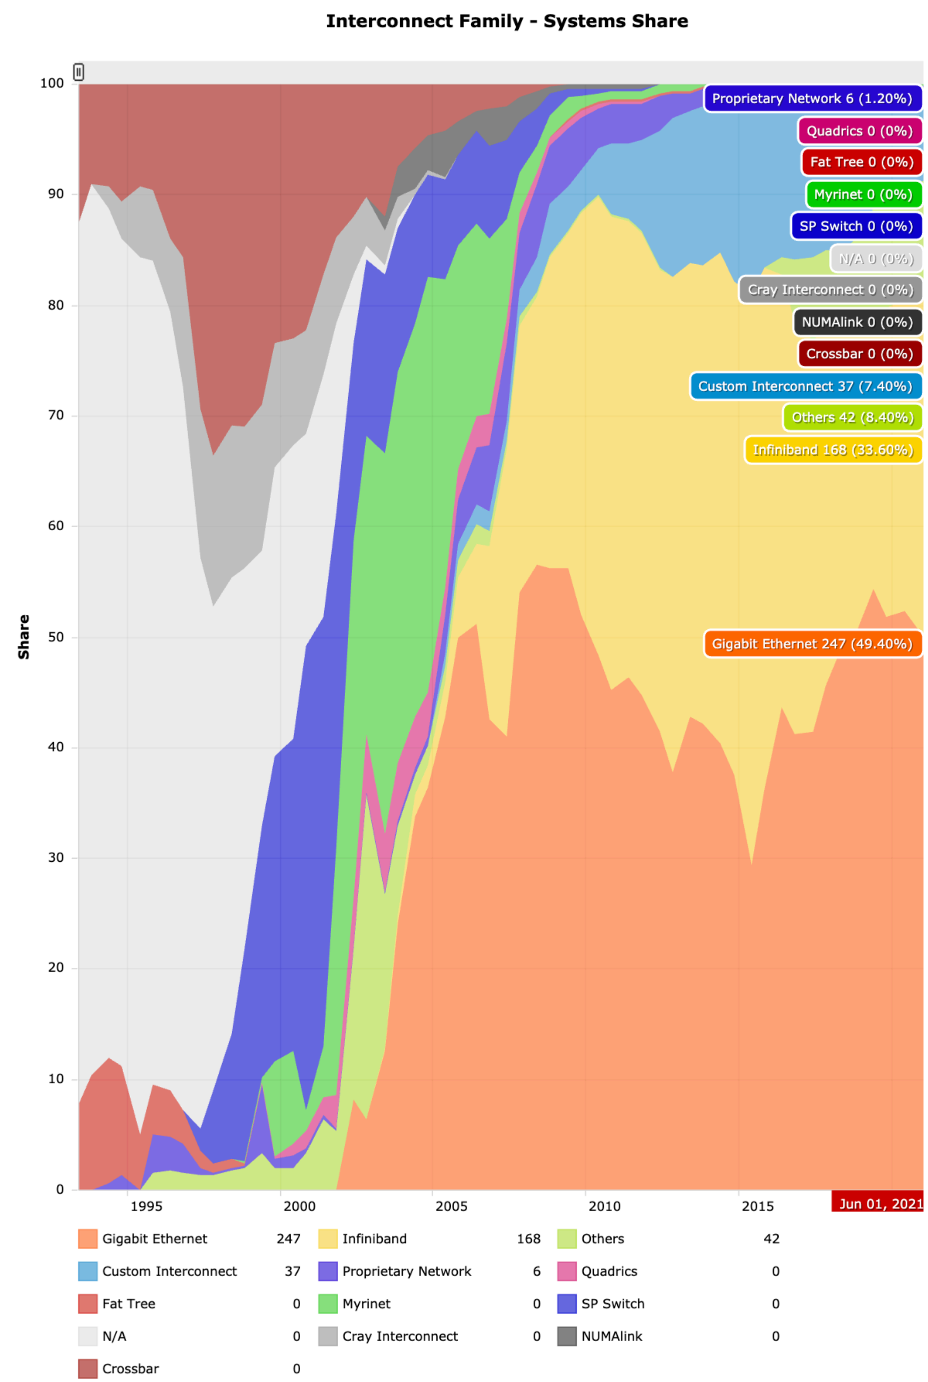

Interconnect Family - Systems Share
Share
Proprietary Network 6 (1.20%)
Quadrics 0 (0%)
Fat Tree 0 (0%)
Myrinet 0 (0%)
SP Switch 0 (0%)
N/A 0 (0%)
Cray Interconnect 0 (0%)
NUMAlink 0 (0%)
Crossbar 0 (0%)
Custom Interconnect 37 (7.40%)
Others 42 (8.40%)
Infiniband 168 (33.60%)
Gigabit Ethernet 247 (49.40%)
Jun 01, 2021
Gigabit Ethernet 247
Infiniband 168
Others 42
Custom Interconnect 37
Proprietary Network 6
Quadrics 0
Fat Tree 0
Myrinet 0
SP Switch 0
N/A 0
Cray Interconnect 0
NUMAlink 0
Crossbar 0

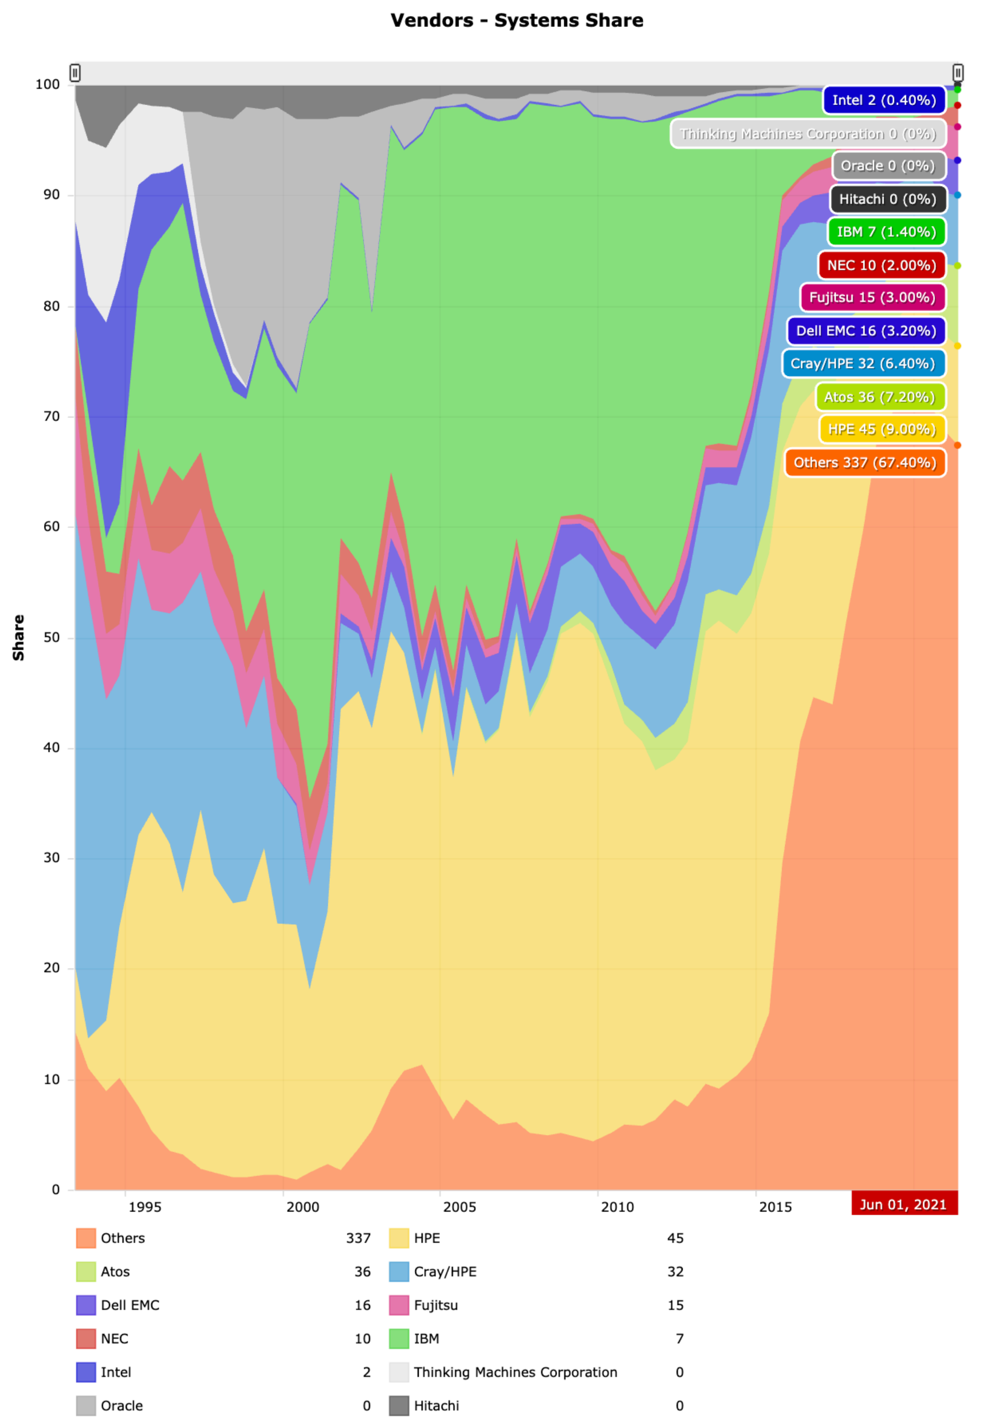

Vendors - Systems Share
Share
Intel 2 (0.40%)
Thinking Machines Corporation 0 (0%)
Oracle 0 (0%)
Hitachi 0 (0%)
IBM 7 (1.40%)
NEC 10 (2.00%)
Fujitsu 15 (3.00%)
Dell EMC 16 (3.20%)
Cray/HPE 32 (6.40%)
Atos 36 (7.20%)
HPE 45 (9.00%)
Others 337 (67.40%)
Jun 01, 2021
Others 337
HPE 45
Atos 36
Cray/HPE 32
Dell EMC 16
Fujitsu 15
NEC 10
IBM 7
Intel 2
Thinking Machines Corporation 0
Oracle 0
Hitachi 0

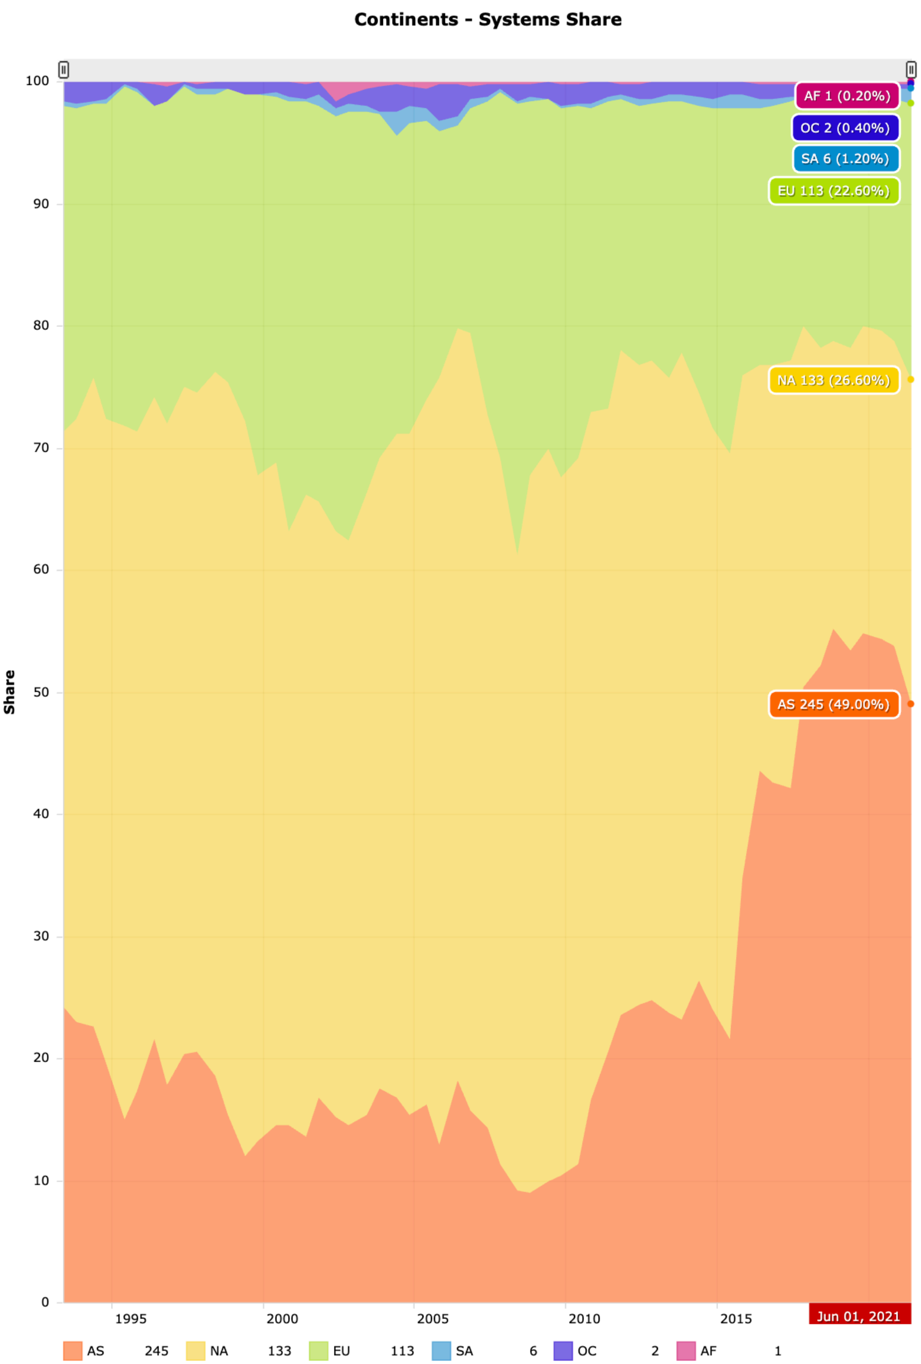
Continents - Systems Share
Share
AF 1 (0.20%)
OC 2 (0.40%)
SA 6 (1.20%)
EU 113 (22.60%)
NA 133 (26.60%)
AS 245 (49.00%)
Jun 01, 2021
1995
2000
2005
2010
2015
AS 245
NA 133
EU 113
SA 6
OC 2
AF 1

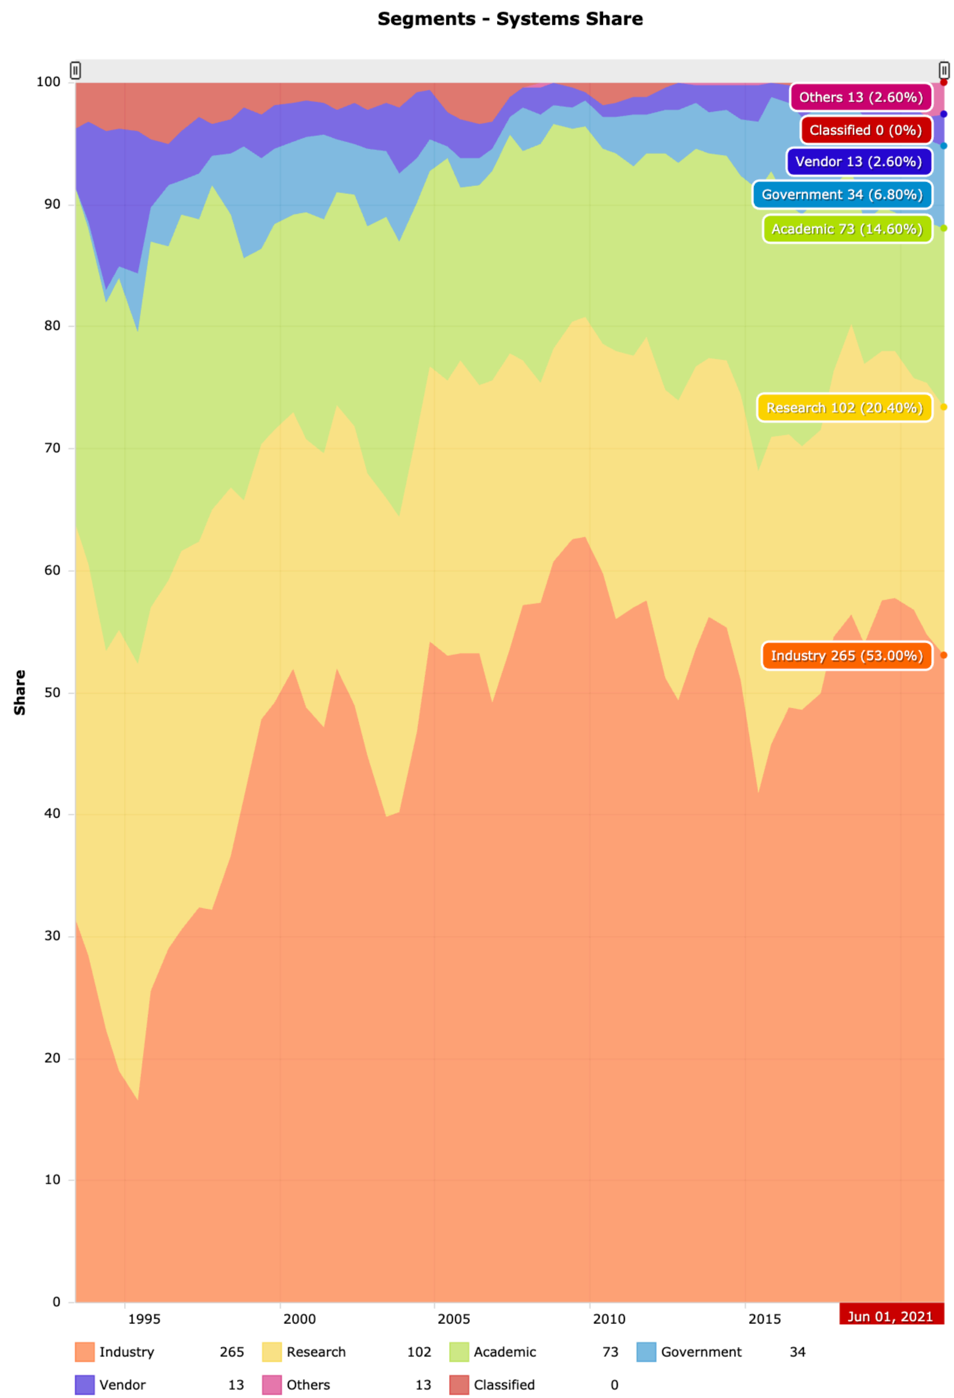
Segments - Systems Share
Share
Others 13 (2.60%)
Classified 0 (0%)
Vendor 13 (2.60%)
Government 34 (6.80%)
Academic 73 (14.60%)
Research 102 (20.40%)
Industry 265 (53.00%)
Jun 01, 2021
1995
2000
2005
2010
2015
Industry 265
Research 102
Academic 73
Government 34
Vendor 13
Others 13
Classified 0

HPC and CECI

http://www.ceci-hpc.be

- CECI is the "Consortium des Equipements de Calcul Intensif" in Wallonia/Brussels
  - Once you create an account you can use all the CECI clusters
  - Funded by FNRS
  - UCLouvain, ULB, UNamur, UMons, ULiège
  - Single login for all clusters (more on that later)
- Machines in CECI grid:
  - NIC5 @ ULiège
  - Lemaitre3 @ UCLouvain
  - Vega @ ULB
  - Hercules2 @ FUNDP
  - Dragon2 @ UMons

NIC5 System Summary

- 73 compute nodes with two 32-cores AMD Epyc Rome 7542 CPUs at 2.9 GHz and 256 GB or 1TB of RAM
- Infiniband HDR interconnect
- 520 TB BeeGFS parallel filesystem
- The cluster is especially designed for massively parallel jobs (MPI) with many communications and/or a lot of parallel disk I/O, 2 days max.
- SSH to nic5.segi.ulg.ac.be with the appropriate login and id_rsa.ceci file (more on that later).

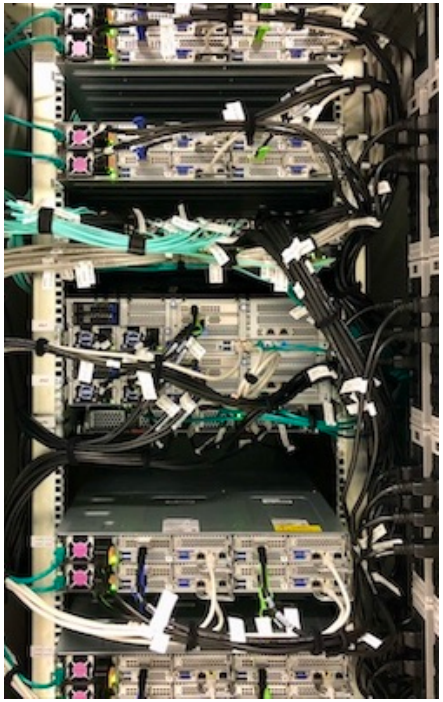

External Power and Cooling

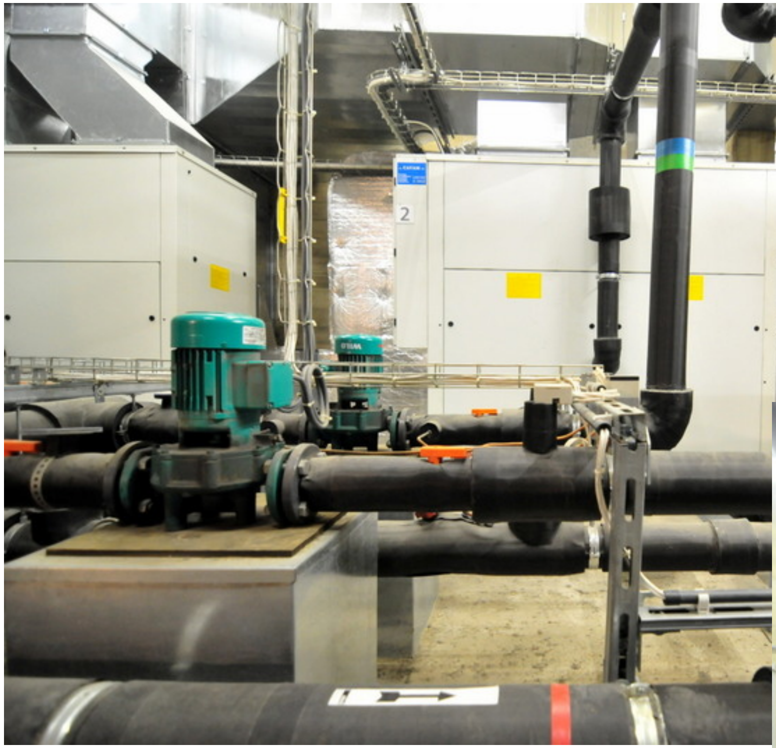

Class Goals/Topics

- Remember that definition "The efficient computation of constructive methods in applied mathematics"
  - Theory: numerical algorithms, (parallel) computation, and how to combine them
  - Practice: the tools of scientific computing, and in particular UNIX exposure (shells/command line, environment, compilers, profilers, debuggers, ...)
- Setup this year:
  - Classes on Tuesday @ 1:45pm, Montefiore R7
  - 1 project ; oral exam during January session

Computer Accounts

- CECI clusters
  - You should create an account now (https://login.ceci-hpc.be)
  - Only SSH access is allowed ; if not on campus, you must use the ULiège VPN:
- Jobs run in a managed environment
  - Login to the login node
  - Submit jobs to the scheduler
  - Wait
  - Collect results
- Production runs on the login node are forbidden
  - Avoid resource intensive tasks
  - Exceptions include compilers, “standard” UNIX commands (ls, mkdir, etc.)
- All required info is available on http://www.ceci-hpc.be
  - Read the FAQ!

CÉCI Clusters News Training FAQ Documentation Support Contact Create/Manage Account

CÉCI

Consortium des Équipements de Calcul Intensif

6 clusters, 10k cores, 1 login, 1 home directory

About

CÉCI is the 'Consortium des Équipements de Calcul Intensif'; a consortium of high-performance computing centers of UCLouvain, ULB, ULiège, UMons, and UNamur. The CÉCI is supported by the F.R.S-FNRS and the Walloon Region. Read more.

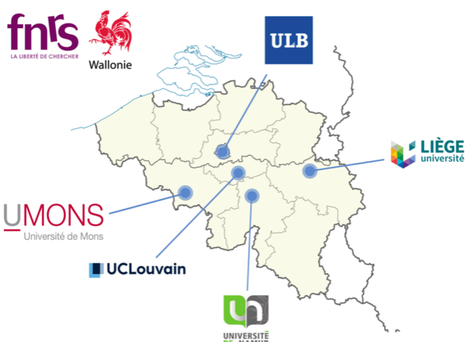

Quick links

- Connecting from a Windows computer
- Connecting from a UNIX/Linux or MacOS computer
- Slurm tutorial and quick start
- Slurm Frequently Asked Questions
- Tier-1 Zenobe quickstart
- Submission Script Generation Wizard

Videos of the training sessions

The training sessions Youtube channel is now ready.

Do not hesitate to have a look at the videos of the training sessions.

Latest News

MONDAY, 25 JANUARY 2021

LUMI Training - Practical course on the CUDA to HIP porting

CSC and the *LUMI User Support Team* organize their first training day that will focus on porting CUDA applications to HIP. This training aims to help the users to port their code to LUMI, the European pre-exascale supercomputer, that will achieve its high computing power thanks to a large number of nodes with AMD GPUs. This event will take place on **February 26**.

This session will provide an extensive dive into the Hipify tools, many examples and a hands-on session. If you're developing your code with CUDA and wish to harness the computing power of LUMI in the future this training is a fantastic opportunity.

You can find further details and the registration form here.

WEDNESDAY, 13 JANUARY 2021

NIC5 installed at ULiège

The new NIC5 HPC cluster is now available. It features 70 nodes with two 32 cores AMD EPYC Rome 7542 cpus at 2.9 GHz and 256 GB of RAM, 3 nodes with 1 TB of RAM, 520 TB of fast BeeGFS /scratch and a 100 Gbps Infiniband HDR interconnect (blocking factor 1,2:1),