Welcome to the PDC Summer School
Introduction to High Performance Computing

PDC Center for
High Performance Computing

Erwin Laure
Director PDC-HPC
Co-Chair PDC Summer School
19. Aug. 2019

Goal of the School

PDC Center for
High Performance Computing

- Introduce you to high-performance computing
- Give you practical knowledge to apply to your own work
- Discuss practical applications and look into the future

What is expected of you?

- Attend the classes
  - Sign up attendance sheet
- Read the handouts and material on the web
- Get lab attendance form signed every lab session
- Submit a project abstract before you leave
  - Then complete the project for credits
- Fill in the online feedback form available on the agenda page:
  - https://agenda.albanova.se/conferenceTimeTable.py?confId=6743

PDC Center for High Performance Computing

3

PDC HPC Summer School DD3258/DD2358 2019
7.5 ETCS

Attendance, two weeks Lectures and labs:
Get Lab attendance sheet signed !

Project, finished Fall '19:

Grade: Grad.: P,F Undergrad. : E... A

Support:
Lab assistant, Project advisor, Examiner

Project:
For some application and HPC architecture of your choice:
- Develop efficient program for non-trivial problem
- Demonstrate and report how efficient it is.

Expected work on the project is **3 weeks** of work ***incl.*** report writing
Deadline for reports: Nov 15, 2019.

The project is ***not*** about:

- *Substantial* development of *new* code.
- *Scientific results* obtained with code

So:

Prioritize measurements and analysis/interpretation!
Demonstrate use of tools (profiling, …) ,
and simple performance model.
NO TIME for development of new **significant** *code.*

Examples:

* Parallelize a code you know and/or work with;
choose interesting part.
* Write a simple code for key algorithm of
bigger solution process
* Write a simple code for a simple problem

Now – during lab-afternoons

- Discuss with instructors & course participants, form groups of size *G*.
- *Define project* and choose *tutor*: Thor, Stefano, Xin, Niclas, ...
- Write very short synopsis, check with supervisor !
- Submit synopsis to *summer-info@pdc.kth.se* ***before end of the course***

Later -

- Start the work ***ASAP:***
- Finish the work; Get in touch with tutor !!
- Submit report to ***tutor***.

The report will be graded and sent back with comments; you may have to complete some parts and hand in again.

- KTH students: LADOK
- Other students: Certificate will be sent to you

1. Develop initial version of program;
2. Develop approximate Performance model = theoretical prediction:
   time = *f*(problem size *N*, #processors *P*, problem partitioning parameters, …)
   Try to assess the *communication* and *computation* times separately.
3. *Measure* performance, e.g. $t = f(N,P,...)$, for different problem sizes, if relevant
   x = wall clock time start to finish, (*not* CPUtime), …

Size \ # proc	1	2	4 …	*n*
N_1	x	x	x	x
N_2	x	x	x	x
…				
N_M	x	x	x	x

4. If suitable, plot "speedup" and/or "efficiency", MFLOPS?, …
   - Make several measurements to discover variations – discuss sources of variability. (interactive nodes, dedicated,…)
   - Compare w. prediction; Interpret: ***Why these numbers?***
   - Identify "bottlenecks" by profiling tools; find remedy & make changes
   - Check improvement by measurements
   - Write report with description of problem, ***algorithm***, and design decisions, pertinent graphs of measurements and profiling, "before and after".

Single processor performance

Algorithm:
    BLAS etc. library
Memory hierarchy
    Disk - main - cache - register;
Organization of loops
    data layout (cache misses)
    index strides (-"-)
    "unrolling"
Compiler directives ("-O2")

Multi-processor performance

Algorithm: Communication !
    Latency vs. bandwidth
    # messages vs. size

Problem partitioning
    Load balancing

Other

- Group size G: $G = 2$ recommended.
- "Standard" grade C. A requires exceptional work
  Requirements for grade $\geq$ C increase with G.
- Proposed schedule
  - < 19-09-30 First iteration: status report, quick feedback from advisors
  - < 19-10-31 Second (final ?) iteration, results, quick feedback/grading
  - -------------- 19-11-15 --------------
  - > 19-12-9 … evaluation may take a while
  - ***> 2020-01-01 evaluation turnaround time may be very long***

Report:

- Background, formulas, relevant problem sizes, …:
- Algorithm, parallelization principle,...
- "Embarrassingly parallel" OK
- Performance model and measurements.
- Graphs, and textual description of what the graphs show, what we learn from them
- ***Interpretation: WHY these numbers?***

Various Information

- Certificates will be issued to all successful students
- Tutors will be available for lab sessions
  - Ask them questions
  - But they will also ask you!
- Labs in groups of 2-3 people
- Wireless
  - Eduroam
  - If you don't have eduroam you can use KTHOPEN
    - Passwords will be distributed as needed
- All material available via the course homepage https://agenda.albanova.se/conferenceTimeTable.py?confId=6743

PDC Center for High Performance Computing

10

Accounts

- If you do not have a PDC account yet, we will fix this tomorrow during the lab.
- **Tell us NOW!**
- **Have an ID handy**
- Bring your laptop to the session tomorrow afternoon so we can fix any last minute problems!!

PDC Center for High Performance Computing

11

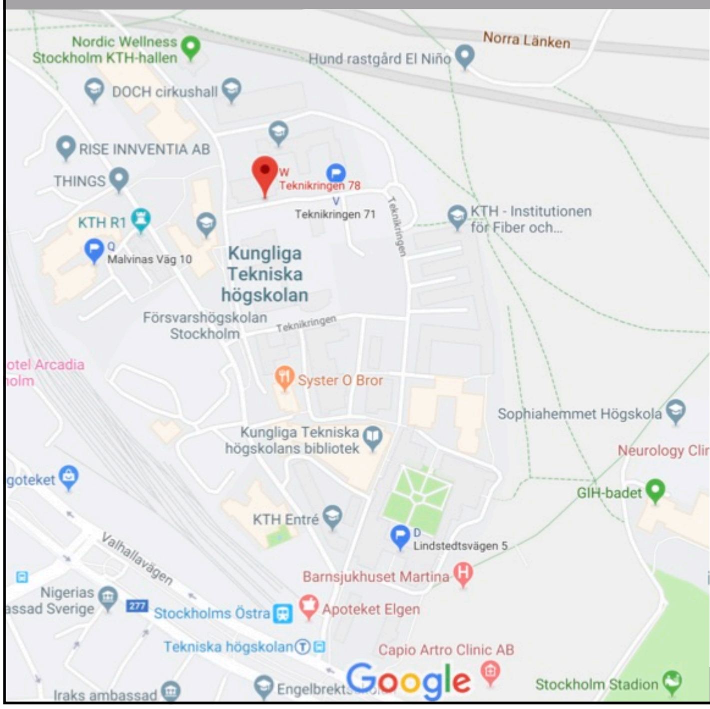

Rooms

- First week: Q2 and Q33/31
- Second week:
  - Q,
  - D (Lindstedtsvägen 5)
  - V (Teknikringen 71)
  - W (Teknikringen 78)
  - Check agenda page!

12

Social Program

- Get-together picnic this lunch at 12.15
- Possibility to see the PDC computer room tomorrow afternoon
- Summing-up dinner on Thursday, August 29, at 18:00, Syster o Bro (http://www.systerobror.se), Drottning Kristinas väg 24

PDC Center for
High Performance Computing

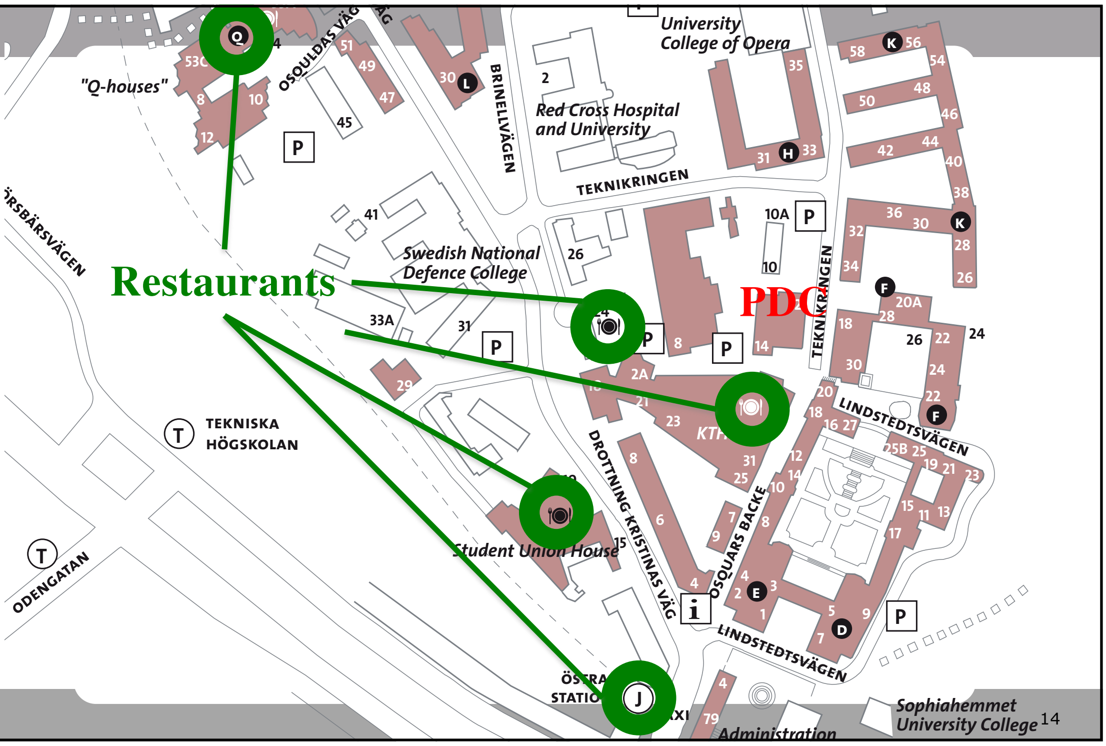

PDC Center for
High Performance Computing

What is PDC

PDC Center for High Performance Computing

- A major Swedish Supercomputing Center supported by SNIC (Swedish National Infrastructure for Computing)
  - Supporting a wide variety of sciences
- Located at the Royal Institute of Technology (KTH), Stockholm
  - Associated to School of Electrical Engineering and Computer Science and communication (EECS)
- Summerschool run jointly with EECS
  - Since 1996

15

PDC's Mission

PDC Center for High Performance Computing

Research

Conduct world-class research and education in parallel and distributed computing methodologies and tools as part of CSC's HPCViz department

Infrastructure (PDC-HPC)

Operation of a world-class ICT infrastructure for Swedish research, including HPC and data services, with associated user support and training

16

PDC HPC Infrastructure

System	System/Processor	TPP (TF)	Cores
Beskow	Cray XC40 Intel Haswell	2,430	67,456
Tegner	SuperMicro Intel Ivy Bridge & Haswell Nvidia K420 & K80	65 + GPU	1,800 + GPU

PDC Center for High Performance Computing

Heat Reuse Project

- Background: today around 1.3 MW used at PDC
- Project started 2009 to re-use this energy
- Goals:
  - -Save cooling water for PDC
  - -Save heating costs for KTH
  - -Save the environment
- Use district cooling pipes for heating when no cooling is required
- Now integrated in a KTH-wide system

PDC Center for High Performance Computing

Storage

PDC Center for
High Performance Computing

- ~20 TB disk
  - Accessible via AFS
- ~5000 TB disk
  - Currently attached to individual systems
  - Lustre parallel file system
- ~700 TB LCG disk
- IBM tape robot
  (~2900 slots, ~2.3 PB)
  - Accessible via HSM, TSM, and dCache

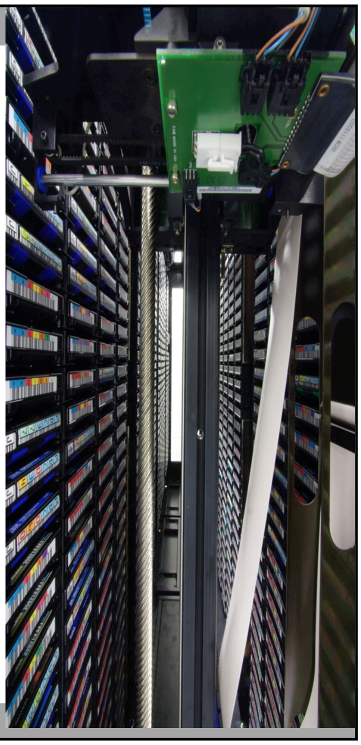

PDC Supporting Swedish Research (Data from 2011)

- About 2/3 of the total SNIC computational resources are operated by PDC
- In 2011 some 380 Million core hours have been used on PDC's SNIC machines, that equals to a usage of over 90% of the nominally available compute power

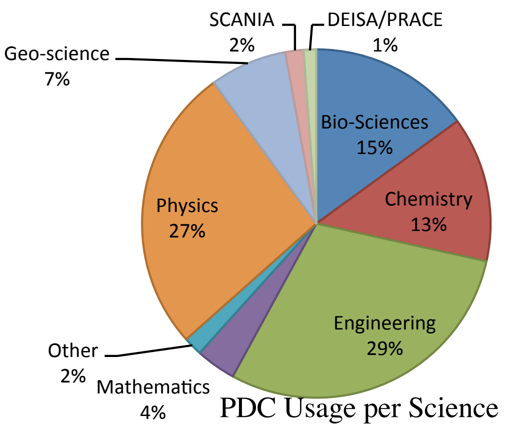

PDC Usage per Science
01/2011-11/2011

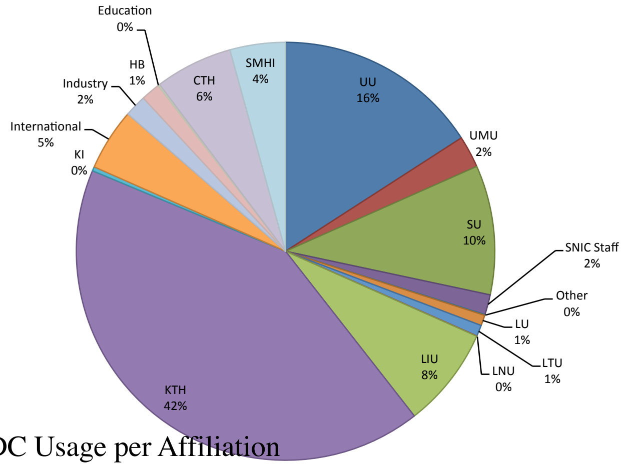

PDC Usage per Affiliation
01/2011-11/2011

20

PDC Center for
High Performance Computing

Software

- Commercial/Community codes
  - Gaussian
  - Jaguar
  - Gromacs
  - VASP
  - Fluent
  - Blast
  - Edge
  - Starcd
  - Dalton
  - Charmm
  - Numerical libraries
- User codes
  - CFD
  - Ab initio
  - Monte Carlo
  - Bioinformatics
- Programming
  - MPI, OpenMP
  - Fortran/C/C++ compilers
  - Tools
  - – acumem
  - – totalview
  - – papi, papiex
  - – jumpshot
  - – hpctoolkit
  - – Allinea tools
  - – CrayPAT

21

PDC Center for
High Performance Computing

SeRC - Swedish eScience Research Center

- Collaboration of KTH, SU, KI, and LiU

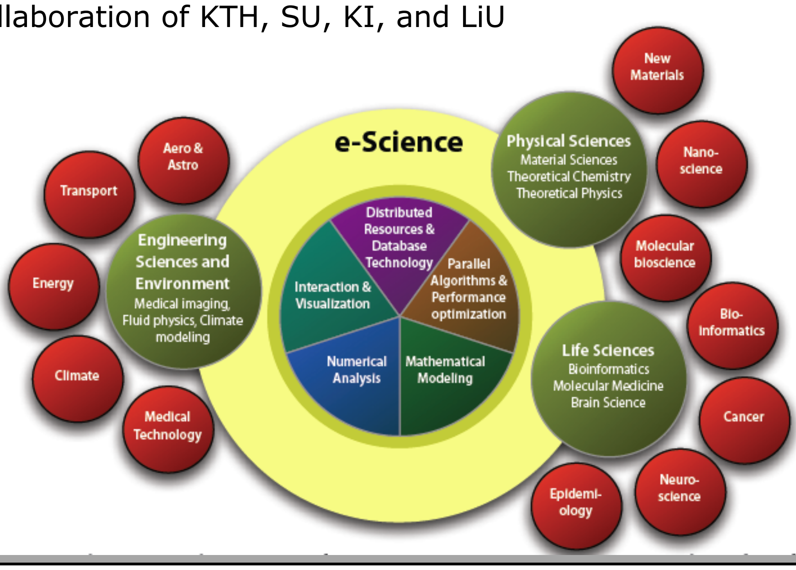

Other Activities

PDC Center for High Performance Computing

- PRACE - Partnership for Advanced Computing in Europe
  Design of future European supercomputing landscape
  - Prototype
    - Focus on energy efficiency and high density packing
  - Tier-1 system
    - Beskow
- EUDAT – European Data Infrastructure
- Advanced software support
  - SeRC
  - EPiGRAM-HS Exascale Programming Models
  - ExaFLOW Fet-HPC Project
  - BioExcel Centre of Excellence
  - Excellerat Centre of Excellence
- Visit program
  - HPC-Europa3

23

The PDC Summer School is part of

PDC Center for High Performance Computing

SeSE

Swedish e-Science Education

24

eScience Education Within SeRC

Arvind Kumar
Computational Science and Technology
EECS, KTH Royal Institute of Technology
Stockholm, Sweden
https://www.kth.se/profile/arvindku/
http://sese.nu/

SeRC Swedish e-Science Research Centre

essence THE e-SCIENCE COLLABORATION

History of e-Science education in Sweden

1996

NGSSC
FORSKARSKOLA I TEKNISK-VETENSKAPLIGA BERÄKNINGAR
NATIONAL GRADUATE SCHOOL IN SCIENTIFIC COMPUTING

2004

KCSE
Workshops/Seminars
Graduate Program
Multidisciplinary Research
Industrial Focal Point
International Contacts
HPC

KTH Computational Science and Engineering Center

2013

SeSE

Hast et al. 2015 11th IEEE Conf. on e-Science

Collaboration between two research initiative

Equal Financing from SeRC and eSSENCE
Typical budget: ~150K SEK

Swedish e-Science Education (SeSe)
Graduate School

Pavlin Mitev
Uppsala
University
Director
Administration

Arvind Kumar
KTH Royal Inst. Of Tech.
Vice Director
Coordination

SeSe Coordination group:
Anna Delin, KTH
Elias Jarlebring, KTH

Gunilla Svensson, Stockholm Uni.

Anders Hast, Uppsala Uni.
Gunilla Kreiss, Uppsala Uni.
Ingela Nyström, Uppsala Uni.
Ola Spjuth, Uppsala Uni.

Jonas Lindemann, Lund Uni.
...

SeSe Mission

- To provide education in fields where the use of **e-Science** is emerging
- Identify areas where courses within **e-Science** are needed
- A meeting place for graduate students using **e-Science** tools and techniques

e-Science Training/Education

Training

Help users to use the tools & resources efficiently = less user support

Field specific & general courses

Education

Theory of HPC & Neighbouring courses

HPC Software & Hardware

PRACE

SNIC

SeSE

SeSE course format:

- 1 week preparation week
- 1 week on campus
- 1 week project
- Equivalent of 3 to 5 credits

Who can attend SeSE courses?

- SeSE courses is open for PhD students enrolled in a graduate programme at a Swedish university (or Nordic university)
- Travel and accommodation funding available

Course Financing

Course of 10 students minimum		
Developing the course 80 000 SEK	Giving the course 80 000 SEK	Travel grants 60,000 SEK Max 10K/student

SeSe courses

CORE topics

- Scientific computing
- Software tools
- Visualization
- Computational data science

- Focus on general skills for eScience
- The objective is to provide a broad coverage
- Offered more regularly

For a list of previous course visit http://ww.sese.nu

Advanced/specialized topics

- Climate modelling
- Matrix computations for statistics
- Optimization tools for simulations
- Multiscale modelling
- Fluid dynamics

- Focus on specific applications of eScience
- Advanced topics in a specific area
- For experts
- Usually not offered regularly

Course statistics

Total courses	48
Unique courses	26
Core courses	13
Advanced courses	13

Course selection

- Is the topic ‘hot’ (e.g. AI, Climate modelling)?
- Is the course responsible person an influential scientists or teacher?
- Does the course fit in the 3-week format?
- Does the course covers an advanced topic usually not covered in available programs at the universities?
- Will the course attract at least 10 students?

e-Science courses at the universities

- At least 30 second cycle and 12 third cycle ‘e-science’ courses exist at KTH and LU
- There is an overlap between available II/III cycle courses and SESE courses
- In general there is a lack of course on computational biology

- So we are working on developing new courses

SG2212 Computational fluid dynamics
DD2360 Applied GPU Programming
MAI0122 Numerical Solution of Initial Boundary Value Problems
Introduction to Visualization and Computer Graphics
BB2490 Analysis of Data from High-Throughput Molecular Biology Experiments
(TFYA75) Applied Physics - Bachelor Project
93FY51 Fysik fördjupning
SF2520 Applied Numerical Methods
SF2522 Computational Methods for Stochastic Differential Eq.
SF2524 Matrix computations for large-scale system
SG2221 Wave motions and stability
DD2370 Computational Methods for Electromagnetics
MAI0129: Stochastic Galerkin Methods for Partial Differential Eq.
TDDD56 Multicore and GPU Programming
DD2257 Visualization
732A54/TDDE31 Big data analytics
SF2521 Numerical Solutions of Differential Eqs
Numerical algorithms for data-intensive science

SG2224 Applied Computational Fluid Dynamics
DD2356 Methods in High Performance Computing
DD1354 Models and Simulation
TDDD43 Advanced Data Models and Databases
Topological methods for visualization
DD3354 Probabilistic Graphical Models
Deep Learning, Advanced
Machine Learning
2990 Multivariate prediction modelling (medicine)
SF3565 Program Construction in C++ for Scientific Computing
SF3580 Numerical linear algebra
SG3112 Turbulence
DD3258 Introduction to High Performance Computing
DF21500 Multicore Computing
Advanced Data Models and Databases
2992 Biostatistics III: Survival analysis for epidemiologists
SF3561 The Finite Element Method
SF3584 Preconditioning for Linear Systems
Programming Frameworks for Deep Learning

Data Collected by Elias Jarlebring

SeSe Courses in 2019

Spring Semester

- Maximizing performance in practical HPC applications
- Introduction to climate modelling
- Advanced Molecular Dynamics

To sign up for the courses go to http://ww.sese.nu

Fall Semester

- Computational python
- Semantic web technologies
- Introduction to high performance computing
- Numerical solutions of initial boundary value problems
- Matrix computations in statistics with applications
- SeRC visualization school

For suggestions and feedback contact:
Arvind Kumar (arvkumar@kth.se) or Pavlin Mitev (pavlin.mitev@kemi.uu.se)

What is Parallel Computing/ High Performance Computing?

PDC Center for High Performance Computing

What is Parallel Computing

- Traditionally, software has been written for ***serial*** computation:
  - To be run on a single computer having a single Central Processing Unit (CPU);
  - A problem is broken into a discrete series of instructions.
  - Instructions are executed one after another.
  - Only one instruction may execute at any moment in time.

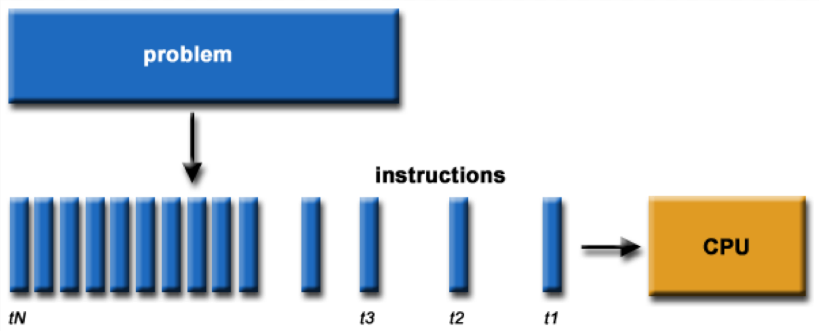

39

Parallel Computing

- In the simplest sense, ***parallel computing*** is the simultaneous use of multiple compute resources to solve a computational problem:
  - A problem is broken into discrete parts that can be solved concurrently
  - Each part is further broken down to a series of instructions

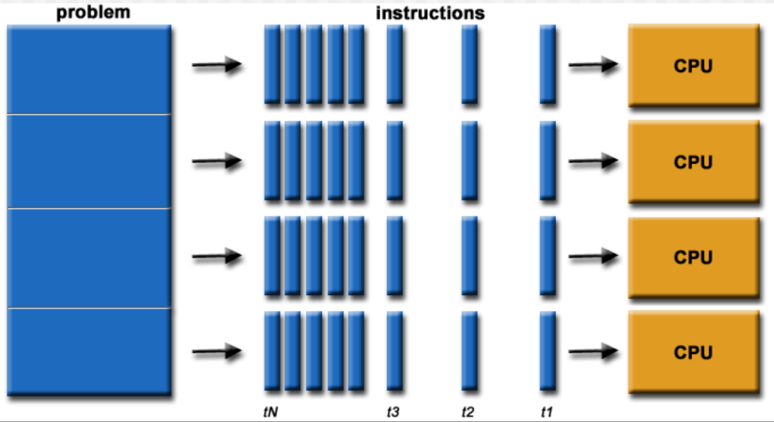

40

Solve a given problem in less time

- Increase the productivity, save money

41

Solve a larger problem

- Many problems require more resources than available on a single computer (particularly memory)
  - “Grand Challenge" (en.wikipedia.org/wiki/Grand_Challenge) problems requiring PetaFLOPS and PetaBytes of computing resources.
  - Web search engines/databases processing millions of transactions per second

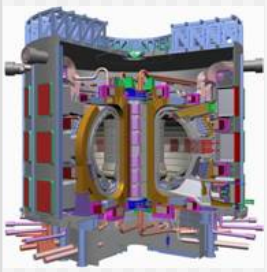

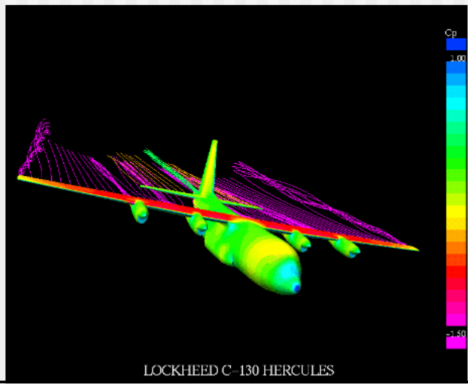

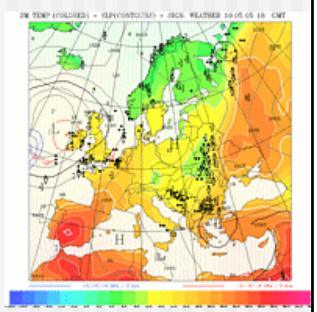

Parallelism on different levels

- CPU
  - Instruction level parallelism, pipelining
  - Vector unit
  - Multiple cores
    - Multiple threads or processes
- Computer
  - Multiple CPUs
  - Co-processors (GPUs, FPGAs, …)
- Network
  - Tightly integrated network of computers (supercomputer)
  - Loosely integrated network of computers (distributed computing)

43

FLOPS

- FLoating Point Operations per Second
- Most commonly used performance indicator for parallel computers
- Typically measured using the Linpack benchmark
- Most useful for scientific applications
- Other benchmarks include
  SPEC, NAS, stream (memory)

Name	Flops
Yotta	10^{24}
Zetta	10^{21}
Exa	10^{18}
Peta	10^{15}
Tera	10^{12}
Giga	10^{9}
Mega	10^{6}

44

Uses for Parallel Computing

- Historically "the high end of computing”
  - Atmosphere, Earth, Environment
  - Physics - applied, nuclear, particle, condensed matter, high pressure, fusion, photonics
  - Bioscience, Biotechnology, Genetics
  - Chemistry, Molecular Sciences
- Geology, Seismology
- Mechanical Engineering - from prosthetics to spacecraft
- Electrical Engineering, Circuit Design, Microelectronics
- Computer Science, Mathematics

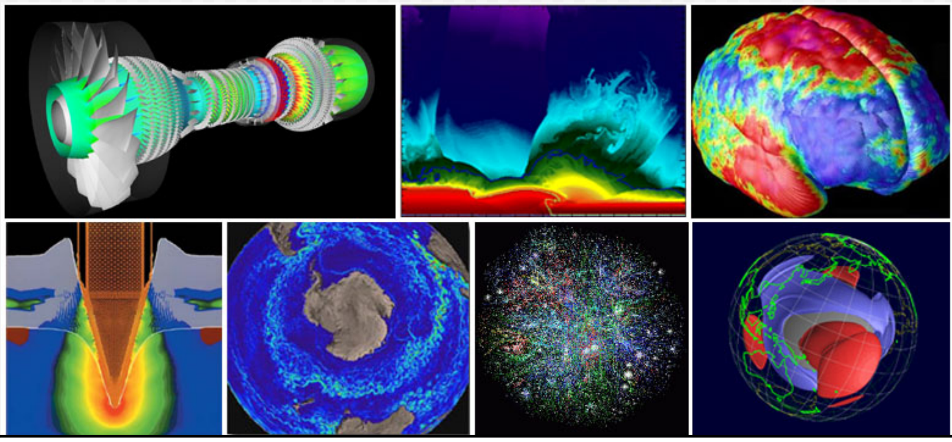

45

Uses for Parallel Computing 2

- Today, commercial applications provide an equal or greater driving force; require processing of large amounts of data in sophisticated ways
  - Databases, data mining
  - Oil exploration
  - Web search engines, web based business services
  - Medical imaging and diagnosis
- Pharmaceutical design
- Management of national and multi-national corporations
- Financial and economic modeling
- Advanced graphics and virtual reality, particularly in the entertainment industry
- Networked video and multi-media technologies
- Collaborative work environments

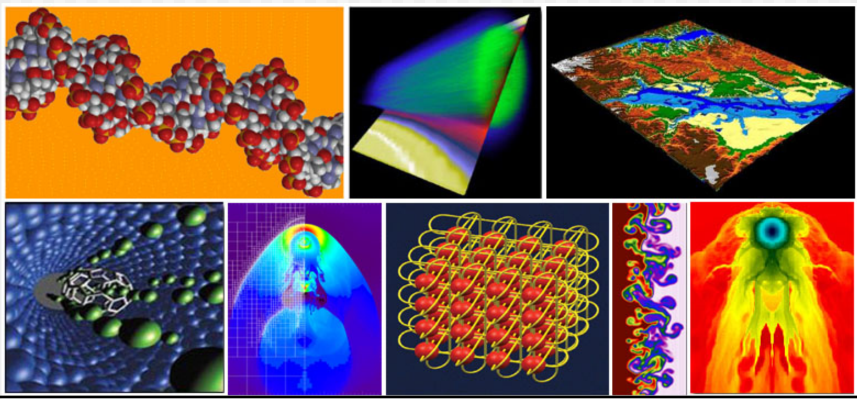

46

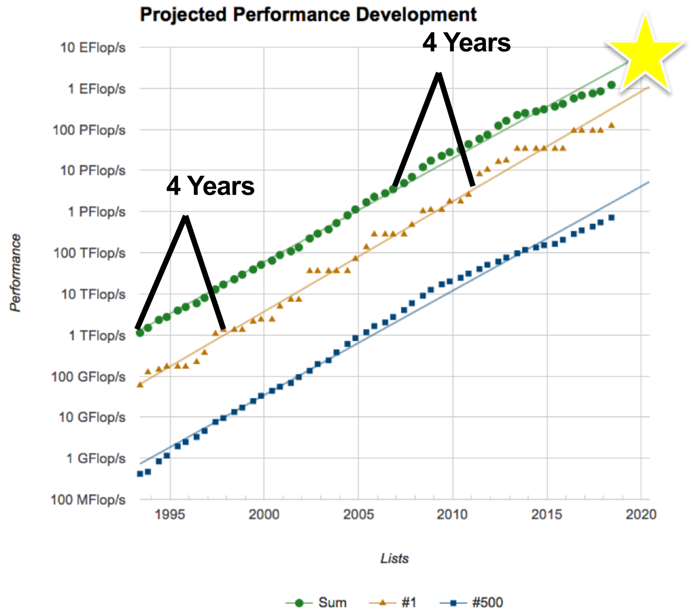

47

Top500 Nr 1: Summit

- Oak Ridge National Lab, USA
- IBM Power9, 22 cores, 3.07GHz & Nvidia Volta GV100
- 2,414,592 cores (2,211,840 GPU cores)
- 148.6 PF Linpack (200.79 PF theoretical peak)
- 10 MW

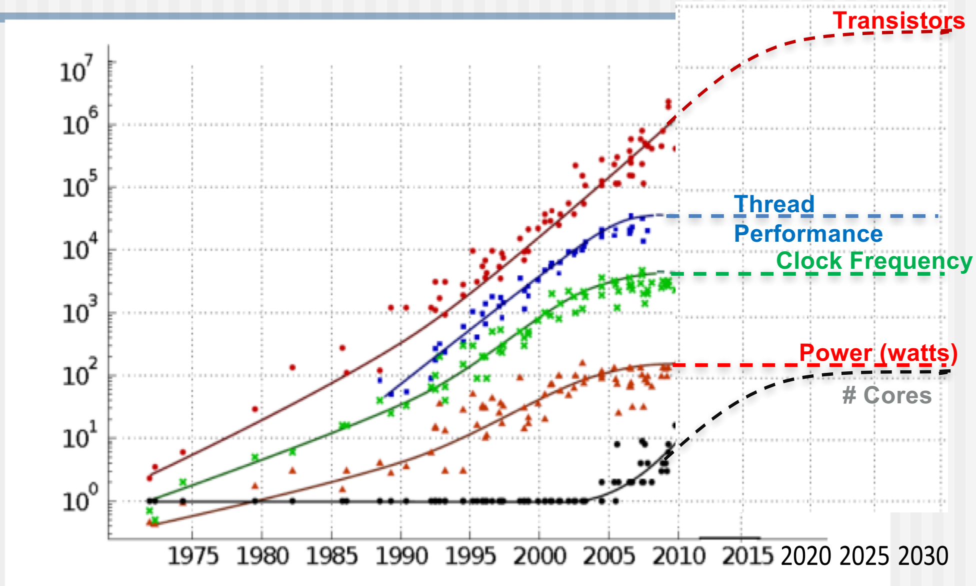

Figure courtesy of Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith

49

The New Computing World

- Increased performance only through parallelism
  - 10^9 FPU needed for exascale
  - Exploit parallelism at all levels of a code
    - Distributed memory, shared memory, vector/SIMD, multithreaded
- Data locality more important than reduction of operations
  - Complex memory hierarchies; HBM, stacked 3D-memory
- Computers are increasingly dynamic and heterogeneous
- Faults may become an increasing issue – or not

50

Summit Node Overview

Application Performance	200 PF
Number of Nodes	4,608
Node performance	42 TF
Memory per Node	512 GB DDR4 + 96 GB HBM2
NV memory per Node	1600 GB
Total System Memory	>10 PB DDR4 + HBM2 + Non-volatile
Processors	2 IBM POWER9™ 9,216 CPUs 6 NVIDIA Volta™ 27,648 GPUs
File System	250 PB, 2.5 TB/s, GPFS™
Power Consumption	13 MW
Interconnect	Mellanox EDR 100G InfiniBand
Operating System	Red Hat Enterprise Linux (RHEL) version 7.4

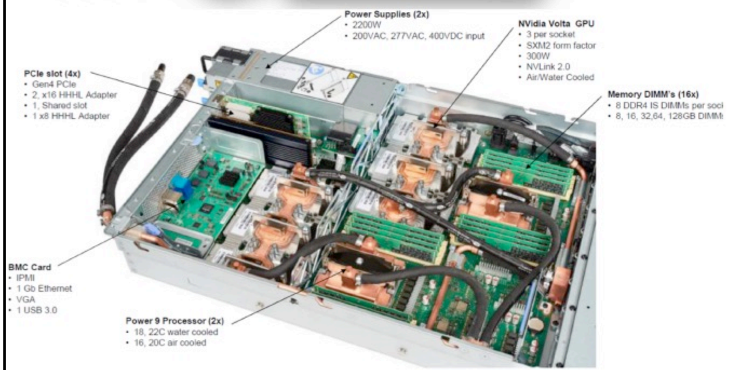

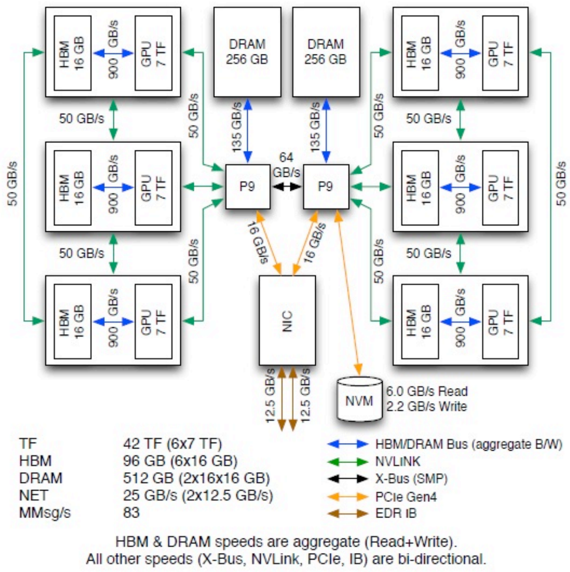

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

LUMI: Our vision of the EuroHPC pre-exascale system

- The LUMI system consists of **a pre-exascale (200+ Pflop/s) accelerated supercomputer,** which enables the convergence of **high-performance computing**, **artificial intelligence**, and **high-performance data analytics**.
- The system will be supplemented by a **number of supporting compute and storage resources,** maximizing its overall value.
  - These will include an additional CPU partition, data analysis partition and cloud resources.
- Maximized **user experience**: interactive supercomputing, easy user interfaces, containerization, pronounced support for data management.

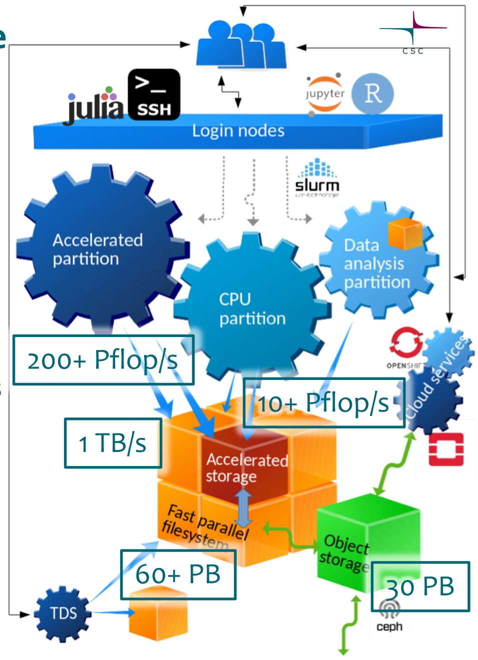

52

PDC Center for
High Performance Computing

Have interesting,
challenging,
fun
two weeks!