

HPC Curriculum Seminar

HPC Curriculum at the Niels Bohr Institute

- High performance parallel computing 7.5 ECTS
 - OpenCL, OpenMP, OpenACC, MPI
 - Applications: Astrophysics, other physics
- Scientific Computing 7.5 ECST
- Programming Massively-Parallel Hardware 7.5 ECTS
 - Intersection of advanced architectures and compilers
 - CUDA, OpenMP and Futhark
- Data Parallel Programming 7.5 ECTS
 - Flattening nested loops - Futhark

Curriculum at DTU

- Key concepts:
 - OpenMP, MPI
 - Quantum
 - Accelerators
 - Cache optimization
 - Parallel vs concurrent
- Core HPC courses of 5 ECTS:
 - HPC
 - MPI
 - Large scale modelling
- Difficulties:
 - Very mixed backgrounds of students
 - It is not always called HPC

Copenhagen Business School

- On the way, but no courses on HPC
- Very few HPC researchers
- But there is an MSc in Business Administration and Data Science. The curriculum:
 - Python programming
 - Linear algebra
 - Data mining, ML and DL
 - NLP and text analytics
- Electives in advanced statistics, time series analysis etc.
- Uses DeiC Interactive HPC/UCloud
- Social data in finance research center open in 2023 - need of more computing facilities

Pedro Petersen Moura Trancoso from Chalmers

- What is HPC?

- High-performance computing
- High-performance computing systems
- Trojan horse?
- Mistake?
- A holistic view of things
- Course formats
 - Lectures, labs, projects
 - A little project at the end of each course
- Pedro teaches sustainable computing
 - efficiency
 - Each week they have to read one of three papers
 - Students meet with the ones who read the same paper
 - Then presents for the rest of the class

The Curriculum at ITU

- From Pinar Tözün
- Cluster managed by Slurm
- Lottie Greenwood manages the cluster
- Teaching students how to work with remote hardware
- Courses
 - Operating systems and c
 - Parallel programming
 - Computer system performance
 - Advanced data systems
 - Large scale data analysis
 - Big data management

Aalborg University

- Thomas Arildsen
- CLAAUDIA offer it infrastructure and guidance to researchers
 - In-house gpu cluster
 - Virtual machines
- Overview of curriculum (examples)
 - High Performance Programming
 - Numerical Scientific Computing
 - Computational Modelling
 - Computational Statistics and Applications
 - Scientific Computing using Python
- Missing good, centralized hpc overview internally
- Dask for Python is more useful than OpenMPI for C
- Trend towards higher-level tools
 - Requires less awareness of communication
 - Interfacing to GPU

Roskilde University

- Thomas Schröder
 - Head of upcoming HPC center
 - Professor in physics
- There is no specific HPC track
- New masters on the way
- Will build their own HPC center soon
- Has a big GPU cluster
- Bachelor courses
 - Empirical data
 - Essential computing
 - Scientific computing
 - Data science and visualization
 - Field Course and Data Analysis
- MSc courses
 - AI/DL
 - Scientific computing and data science
 - Statistical Physics and Scientific Programming
 - Applied Data Science and Data Visualization
 - Data Analysis and Modelling in Environmental Science

TALK: Master Programme in HPC (MPHPC) @ Chalmers: Walking up the Steps

- Core faculty is computer architecture
- Develop the machines rather than utilizing them
- HPC is not only huge systems. It could also be small systems with huge number of constraints. A drone for example. Or a self driving car. But also data centers and IoT infra structure.
- Goal: Running advanced algorithms on advanced hardware.
 - Hardware-software co-design
 - Domain specific components
- Core knowledge:
 - Computer architecture
 - High-performance Parallel Computing
 - Sustainable computing
- ppedro@chalmers.se

Aalborg University - A Slightly Different Perspective

- Jimmy J Nielsen
- Department of electronic systems
- 3x5 ECTS blocks
- Course in numerical scientific computing
- Goal of the course:
 - Toolbox for exploiting hardware
 - Numpy (vectorized) Numba
 - Multi-core cluster with Dask
 - Mandelbrot set computations

- GPU programming in OpenCL
- Free book:
 - "Introduction to High Performance scientific Computing" by Victor Eijkhout
- When compiling with numba the effect of level 1 and 2 cache can be seen

Sven's Presentation: DTUs Computing Center (DCC)

- Interactive HPC and batch jobs
- Hosting services and consulting services

Cosmin Oancea: Research Bases Teaching in HPC

- cosmin.oancea@diku.dk
- Concerns backed up by former students
- Working on Futhark - functional parallel language

Rasmus Wridt Larsen

- Internship at Livermore with Olga Pierce.

Steffen Holst Larsen

- Developing SYCL at Intel in Edinburgh

Day II - Friday the 26th of August 2022

Data Science Lab - Erik Dam

- HPC is just a tool for people who don't study computer science.
 - They do not know of shell scripts
 - Or scheduling
 - Or even ssh
- They typically need a bigger machine
- Or guidance on how to reduce the complexity of their program
- Perhaps they need access to accelerators for deep learning
- Some cases end up in proper data science and publications

Data Science Lab - Stefan Oehmcke - ML Methods for Geosciences

- Example: Counting trees in Sahara
- Giving out simple deep learning framework
- Later the students can dive deeper into deep learning to make more complicated models

Danfoss Drives

- Christian Graulund
- Manager of Research and University Collaboration
- Danfoss Group

- Danfoss are "Engineering Tomorrow"
- Three core segments
 - Power solutions - construction, hydraulics
 - Climate solutions - heating for instance
 - Danfoss Drives - Brewery, HVAC, automotive, etc
- Global megatrends
 - Digitalization
 - Climate change
 - Electrification
 - Food supply
- Danfoss Drives
 - AC Drives /adjustable frequency drives / frequency converters
 - Electric motor control
 - Pump and pressure control
 - Food and beverages - conveyers - has to be very fast and very precise
 - Still quite a few more hardware than software developers
- Research areas of interest
 - ML and AI - do not need very big and heavy models - heterogeneous hardware
 - Simulations - digital twins - integration in factory scale simulations - surrogate models - it is hard to do hard fidelity in real time
 - Classical HPC is not necessarily very important to Danfoss
- christian.graulund@danfoss.com
- Looking more and more into accelerated hardware
- Is there an increase of local decision making on the edge?
 - All decisions happen at the edge
 - IF you need real-time control, it has to be done at the edge
 - Preprocessing on the edge
 - The intelligence is going downwards and up
 - In some cases no kind of connection to the cloud is allowed The cloud is important, but turning off the cloud is also important.
- Do you need "performance engineering"?
 - More functionality can be fitted in an old platform
 - A lot of what is being done at Danfoss is is embedded software engineering
 - We need to understand heterogeneous computing platforms
 - Offloading to GPUs and FPGAs - what are the tradeoffs
 - Need of broad students who knows about both cloud computing and offloading - knowledge of working in a distributed network - what happens when half of the network drops out
- What is your wishlist of what should be included in courses over the next 5 to 7 years?
 - Leveraging heterogeneous computing platforms - not only relevant to Donfoss but entire industry
 - Small devices are also like a computing cluster when you think of all the sensors and accelerators
 - Some of that is networking, some of that is edge computing, some of that is cloud computing
 - Interfacing - what can be done with an edge node - understanding connections
- How are FPGAs important to Danfoss?
 - To get realtime systems running - that is the big application inhouse
 - Can other applications leverage the FPGAs if the devices comes with it anyways

- What kind of ML can you do an FPGA?