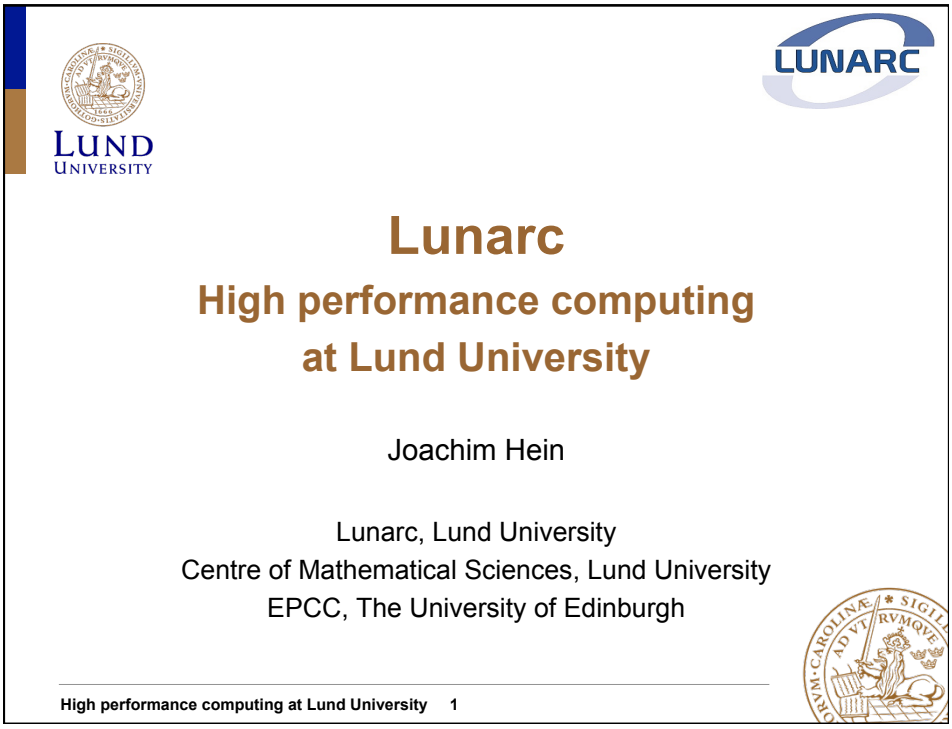
LUND
UNIVERSITY
LUNARC
Lunarc
High performance computing
at Lund University
Joachim Hein
Lunarc, Lund University
Centre of Mathematical Sciences, Lund University
EPCC, The University of Edinburgh
High performance computing at Lund University 1

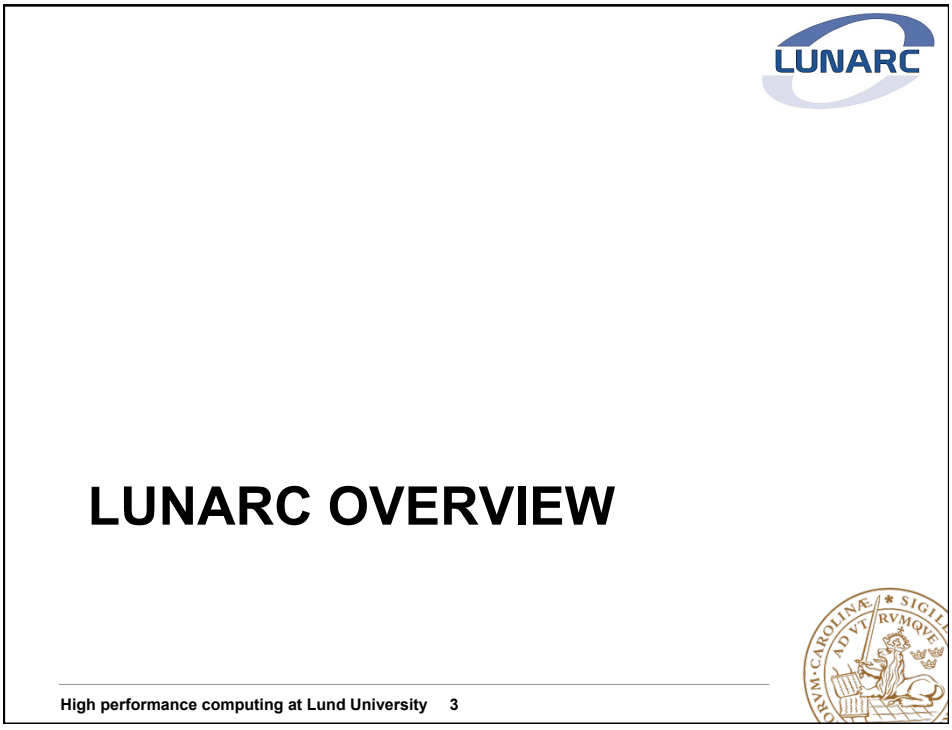
LUNARC
LUNARC OVERVIEW
High performance computing at Lund University 3

Lunarc – The e-Science hub

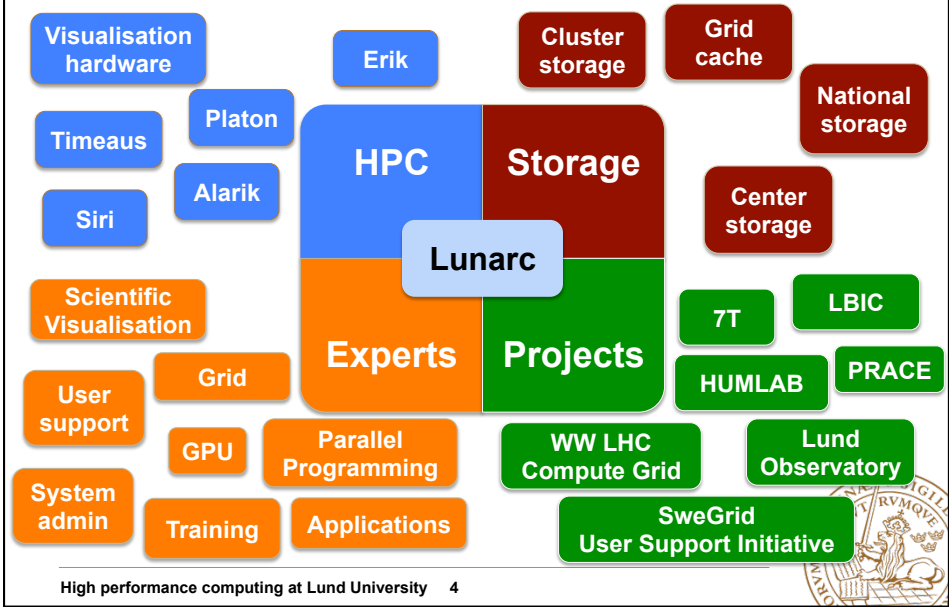

Lunarc computer room

Current Resources

208 x 16 cores
AMD Opteron 6220

32 – 64 GB
Infiniband 4x QDR
Direct login

216 x 8 cores
Intel Xeon e5520

24 GB
Infiniband QDR/DDR
Direct login

64 x 8 cores
Intel Xeon 5430

16 GB
GigE interconnect
Grid access

4 nodes, 8 core & 2 GPU each
Intel Xeon E5520
Nvidia Tesla (fermi) M2050

24 GB + 2 x 3 GB
Infiniband QDR/DDR
with Lund Observatory

GPU resource: Erik

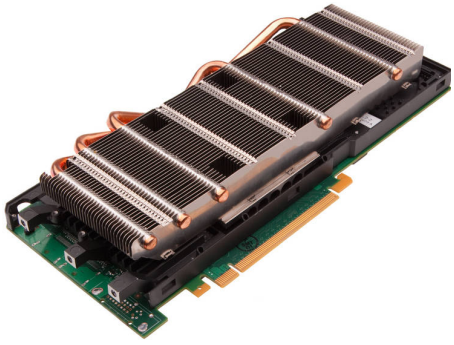

- SNIC funded
- CPU-GPU pilot system
- To arrive Summer 2012

- **68 Nvidea Tesla M2090 GPU**
- Parallel GPU programming
  - Multiple GPU per node
  - Multiple GPU nodes (message passing)

- GPU experts at SNIC HPC centres
- Expected: the largest Swedish GPU system

ALARIK
LUNARC'S NEWEST CPU SYSTEM

Alarik: Lunarc's newest CPU service
Hardware Overview

- **208 compute nodes**
- **3328 processing cores in total**

- Dual processor nodes: AMD Opteron 6220 "*Interlagos*"
  - 8 compute cores per processor, **16 per node**
  - 3 GHz clock frequency, highest AMD offers
- DDR3 1600 MHz main memory
  - 70 nodes with **64GB main memory**, 4GB per core
  - 138 nodes with **32GB main memory**, 2GB per core
- Fat tree Mellanox 4x QDR Infiniband switch network
  - Three head switches to improve bi-sectional bandwidth

Forthcoming hardware: Large memory and high core-count nodes

- 4 Nodes to be installed in spring 2012
- Four AMD Opteron processors "*Interlagos*"
  - 12 compute cores per processor, **48 cores per node**
  - 2.6 GHz clock frequency
- **128 GB of main memory** per node
- Ideal for:
  - **threaded applications** (e.g. Posix, OpenMP, …)
  - Applications requiring **large amounts of main memory**

MANAGING MEMORY TRAFFIC

Structure of an Alarik node

- 2 Processors
- 4 quad core groups (dyes)
- Memory attached to these groups (CC-NUMA architecture)
- Hypertransport
- 16 cores
- 8 FPUs
- A lot of structure that affects performance!

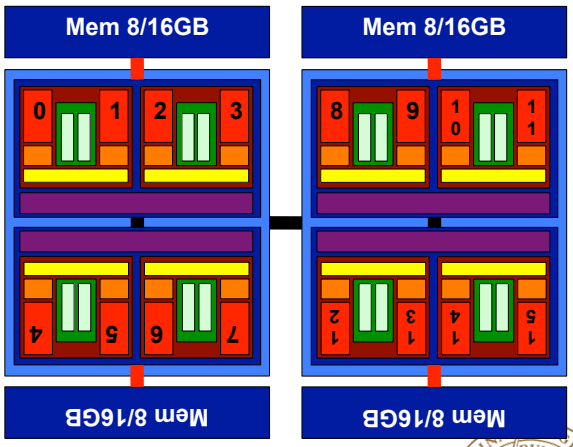

Investigating Memory bandwidth
How fast can one load and store data

Using all memory busses
- Binding, two orderings:
  - 0, 1, 2, 3, 4, 5, 6, 7, 8, …
  - 0, 8, 4, 12, 1, 5, 9, 13, …
- No Binding

First memory bus only
- Binding, ordering:
  - 0, 1, 2, 3, 4, 5, 6, 7, 8, …

- Open64 compiler, OpenMP

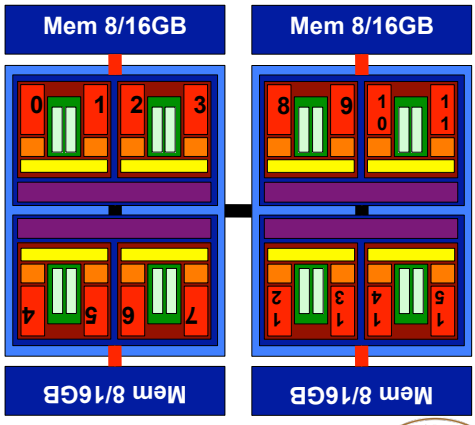

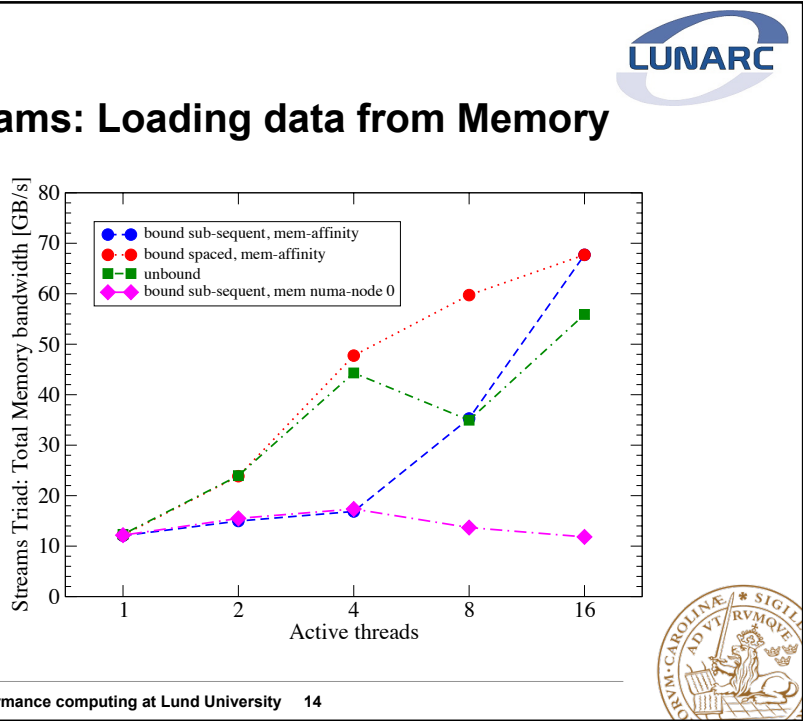

Streams: Loading data from Memory

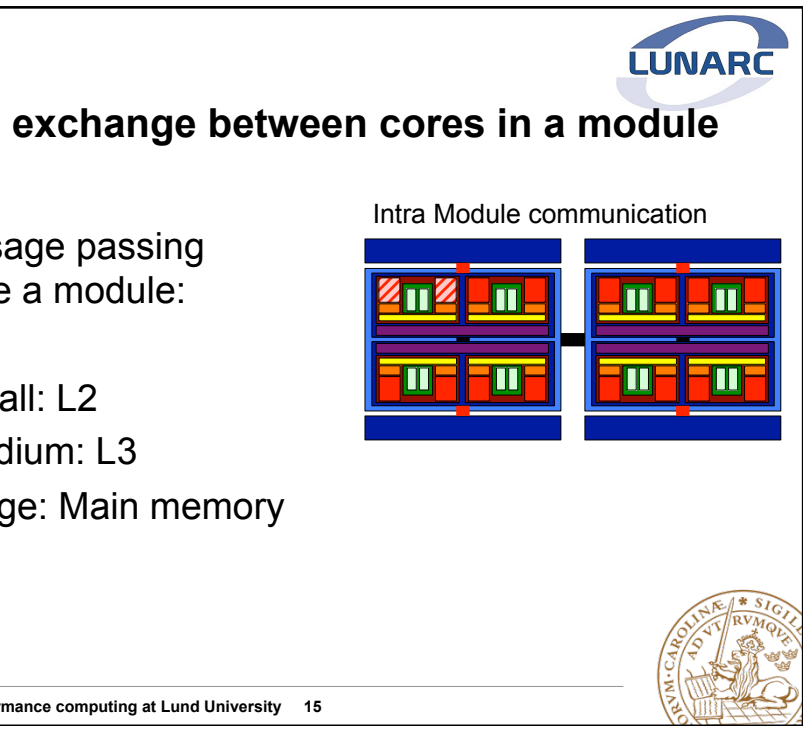

Data exchange between cores in a module

Message passing inside a module:

- Small: L2
- Medium: L3
- Large: Main memory

Intra Module communication

Data exchange between processors in a node

Message passing between processors inside a node:

- All sizes: via Hypertransport

Intra node communication

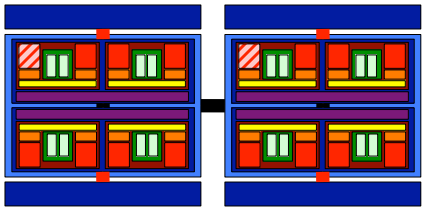

Data exchange between nodes

Message passing between nodes:

- All sizes: via IB-network
- RDMA, transfer doesn't involve cores

Inter node communication

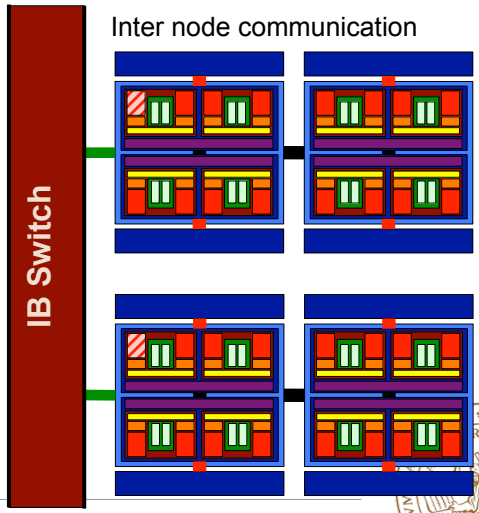

Ping-Pong bandwidth for a pair of tasks

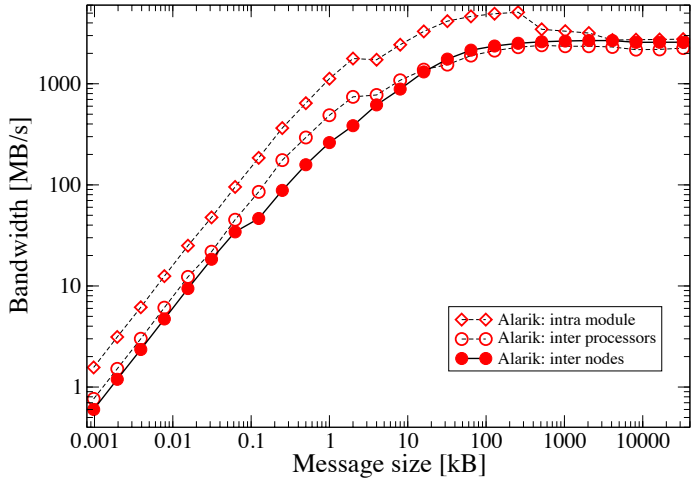

APPLICATION PERFORMANCE

Gromacs

- Molecular dynamics package
- Primarily for biochemical molecules (e.g. proteins, lipids)
- Assembler kernel (SSE instructions)
- Studying d.dppc benchmark, single precision, 122 katoms
  - Hard cut off, no PME

Gromacs performance

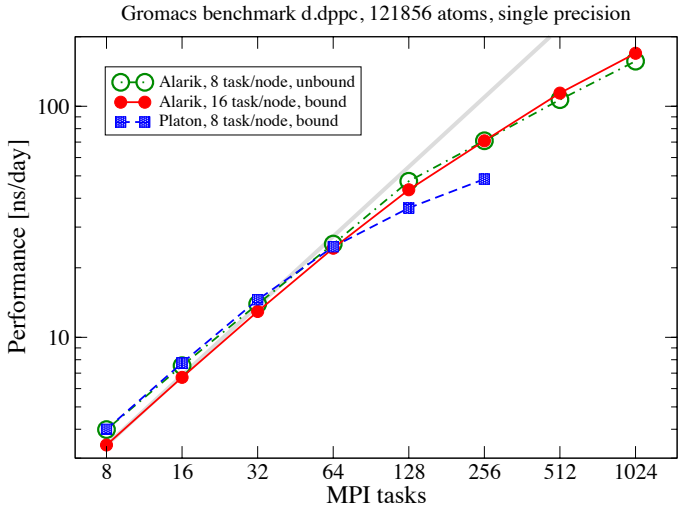

Alarik:

- **Gromacs 4.5.5**
- **GCC 4.6.2**
- **OpenMPI 1.4.4**

Platon:

- **Gromacs 4.5.4**
- **GCC 4.4.2**
- **OpenMPI 1.4.1**

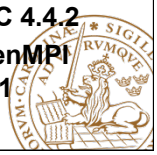

OpenFOAM

- Open source CFD software package
- Benchmark configuration supplied by: Johan Nilsson
  - Points: 3.9 M
  - Cells: 3.7 M
  - Faces: 11.3 M
  - Relevant for: simulation of acoustic fatigue in aircraft

OpenFOAM benchmark results: sonic

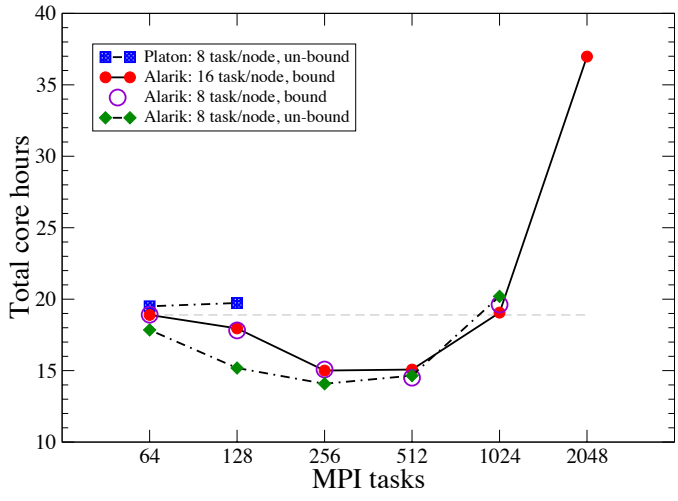

Courses and Training

- "Efficient programming of modern HPC architectures"
  - COMPUTE trainings course, HT1 2012
    - Fortran 95
    - Python scripting
    - Parallel programming: MPI and OpenMP
- "Message passing with MPI"
  - Lunarc Training course, 11-13 April 2012
  - Registration opened yesterday
- "Introduction to HPC"
  - Lunarc Training course, planed for May 2012

Acknowledgements

- Jonas Lindemann, Lunarc
- Anders Sjöström, Lunarc
- Magnus Ullner, Lunarc
- Anders Follin, Lunarc
- Johan Nilsson, Construction, Lund University

Summary

- Lunarc supports academic research
  - HPC cycles: CPU and GPU
  - Data storage
  - Expertise
  - e-Science projects
- Alarik – Lunarc's newest CPU service
  - Excellent parallel performance
- HPC training & consultancy
- If your PhD project has special e-Science needs
  - Speak to your supervisor and speak to us