

# A Short Primer on (HPC) Clusters

Henrik Thostrup Jensen

April 18th 2006

# What is a Cluster Anyway

- It does NOT do any of the following:
  - Use magic
  - Automatically make your program run faster
  - Provide a single OS image of all resources

# So what is it

- So what is a cluster really
  - A set of closely connected computers
    - Usually homogeneous
    - Connectivity: GB Ethernet, SCI, Infiniband
    - They run some form of \*nix
  - Purpose
    - To solve problems faster than on single machine
    - To solve problems that cannot be solved on a single machine
- Performance is everything

# What Runs a Cluster

- An LRMS: Local Resource Management System
- Also called batch systems
  - Because that is what they are
  - A job gets a set of nodes/CPU's exclusively
- How does most LRMS work
  - Make a job description
  - Submit it to the LRMS
  - The LRMS will decide when to run it
- Torque/PBS, Sun Grid Engine, LoadLeveler, Condor
  - Some are open source, some are proprietary

# How does a program run on a cluster

- It is made explicitly parallel
  - This is difficult and hard work
- MPI is used for this
  - Many MPI solutions for different interconnect types
- This is the hard part
  - Usually it is more of a rewrite than a port
  - Setting up a cluster is the easy part
- Getting parallel program to perform well is immensely difficult
  - But challenges are fun!
  - Multi-core CPUs are becoming the norm
  - This is why this course is important