Parallel Computing
- a Status Report

Sverker Holmgren

Dept. of Information Technology, Uppsala University

- Director for Swedish National Infrastructure for Computing (SNIC)
- Chairman of the Board for the National Graduate School for Scientific Computing (NGSSC)

- Co-Director for the Strategic Research Area effort eSSENCE
- Prof. in Scientific Computing and head of the Research Program in Computational Science at Uppsala University

In this Presentation...

- Some history
- e-Science
- e-Infrastructure
  - SNIC
  - Current hardware
  - Parallel computing everywhere...
- Parallel computing examples
  - Quantum Molecular Dynamics
  - Genetic Analysis

More than 25 Years !

Parallel computing has been an every-day tool for solving computational problems in science and technology since the mid-80´s

- Parallel algorithms
- Tools for parallel programming
- Parallel codes
  - Message passing (MPI) in the 80´s
  - Shared memory (OpenMP) in the 90´s
- Parallel performance analysis
- Parallel computer systems
  - For a long time: 10-1000 processors

e-Science

”Research where advanced computational tools are developed and used”

- Simulation
- Analysis of data
- Collaborative and/or distributed projects

A new mode of research, complementing theory and experiments !

- (Parallel) algorithms and (parallel) computing tools
- e-Infrastructure

e-Science - Internationally

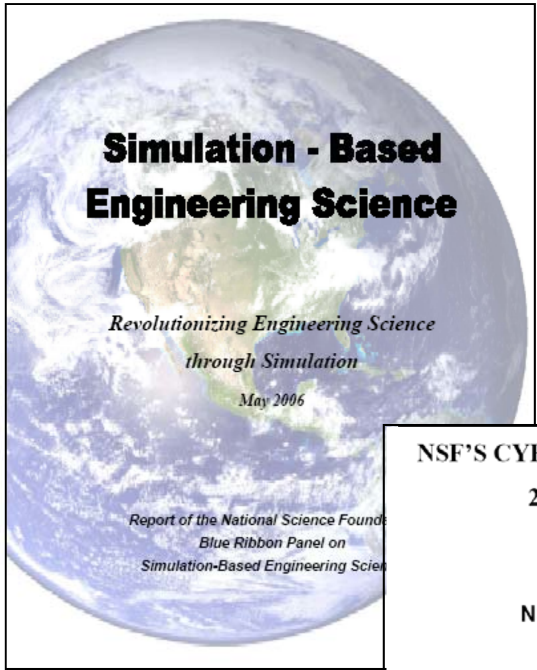

www.nature.com/nature/focus/futurecomputing/
2020 Computing: Everything, everywhere
2020 Computing: Champing at the bits
2020 Computing: Exceeding human limits
2020 Computing: The creativity machine
2020 Computing: Science in an exponential world
2020 Computing: Can computers help to explain biology?
2020 Computing: A two-way street to science's future

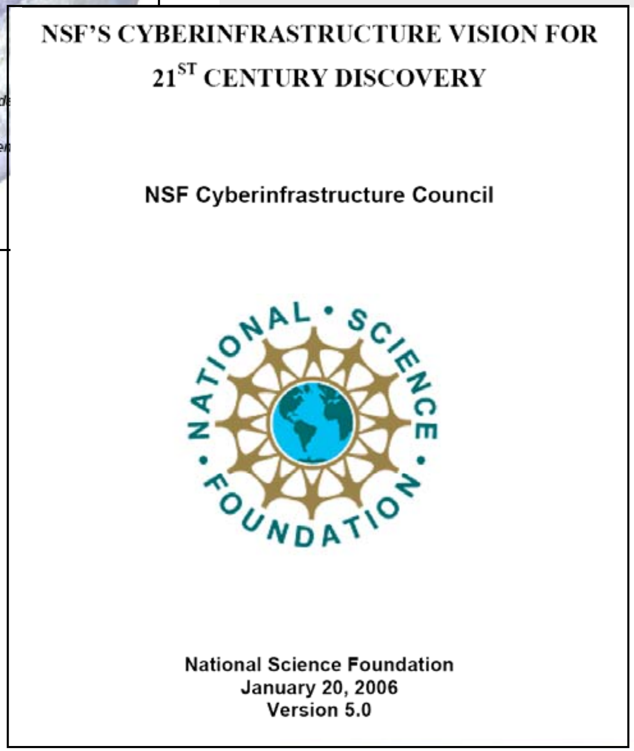

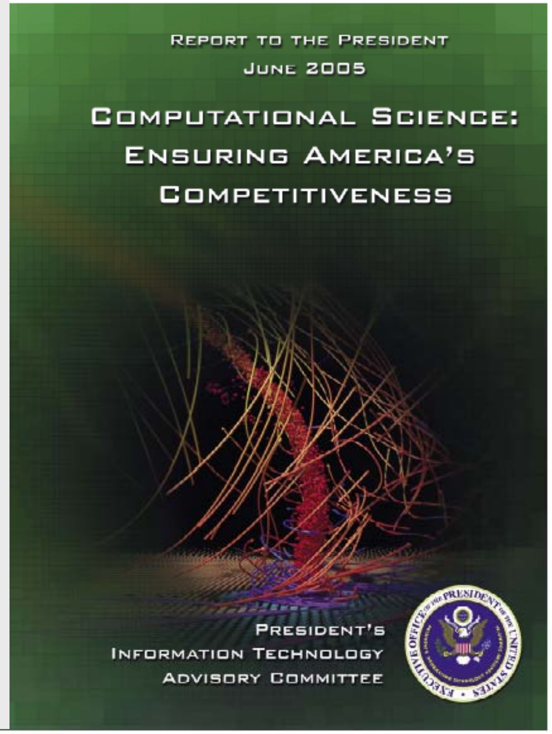

e-Science - Internationally

Computational science is now indispensable to the solution of complex problems in every sector, from traditional science and engineering domains to such key areas as national security, public health, and economic innovation. Advances in computing and connectivity make it possible to develop computational models and capture and analyze unprecedented amounts of experimental and observational data to address problems previously deemed intractable or beyond imagination. [2005 report from the US President's Information technology Advisory Committee (PITAC)]

eScience, promises to revolutionise the scientific discovery process, as did the 'Scientific Renaissance' in laying the basis for modern science. [2009 communication from the European Commission to the European Parliament on ICT Infrastructures for eScience]

e-Science – in Sweden

- Swedish e-Infrastructure:
  - SNIC and the six SNIC Centers
  - SUNET
  - SNDx

- Sweden has strong e-Science research!

- e-Science was one of the 22 Strategic Research Areas pointed out in the Government Bill on Research and Innovation in 2008
  - Extra funding, 70 Mkr/year
  - Two research efforts + e-Infrastructure:
    - eSSENCE (UU, LU, UmU)
    - SERC (KTH, LiU, SU, KI)
    - Additional funding to SNIC (e-Infrastructure)

SNIC: e-Infrastructure for Swedish Research

- Super-scale computing
  - 10^5-10^6 cores
  - EU-level systems (PRACE)
- National systems
  - 10^4-10^5 cores
  - A few, at some centers
- Foundation level systems
  - 10^3-10^4 cores
  - At each SNIC center

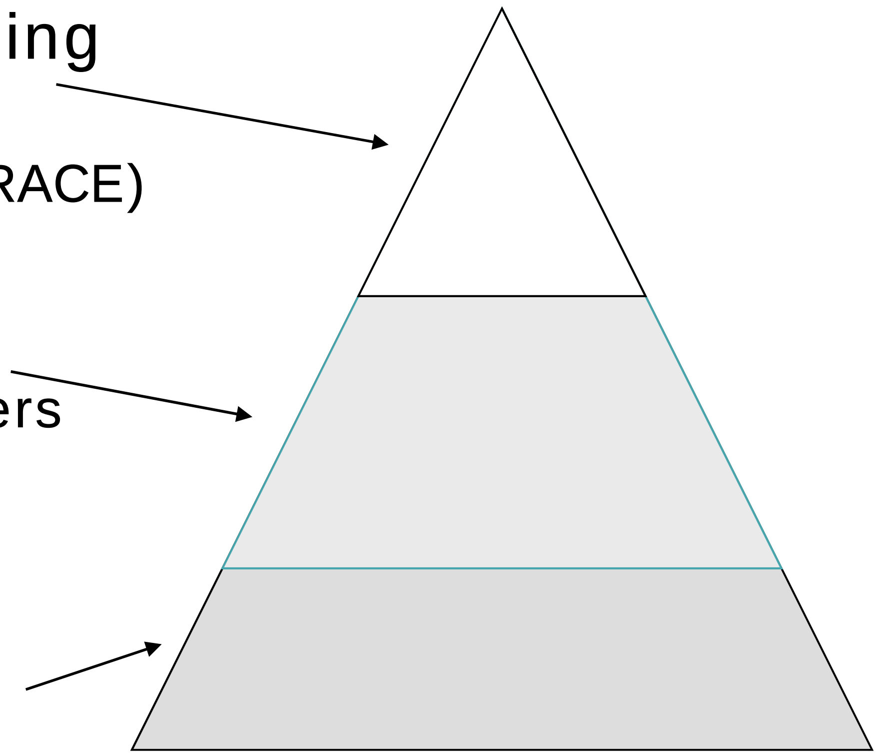

- Different modes of access (login, grid, ...)

Most Recent SNIC systems

- Large-scale computer resources
  - Neolith - 6440 cores (NSC, Sep 2007)
  - Akka - 5376 cores (HPC2N, Jul 2008)
  - Ekman/ Vagn - 10144 cores (PDC/NSC, Sep 2008) - KAW
  - Lindgren - xxxx cores (PDC, Aug 2010) – upgrade autumn 2010
  - New cluster with fat nodes (HPC2N, Spring 2011)
  - Neolith replacement (NSC, 2012)
- Foundation level resources
  - Ferlin - 5376 cores (PDC, 2008)
  - Four clusters at UPPMAX, NSC, C3SE and LUNARC (2010): In total xxxx cores
  - New clusters at UPPMAX, C3SE, LUNARC (2011)

  - SweGrid
  - SweStore

Recent Systems in the World

Cray XT Jaguar, 224162 Opteron cores, 2,331 Pflop peak performance (Jun 2010). Located at ORNL, USA

IBM BlueGene/P, 294912 PowerPC cores,1 Pflop peak performance (Jun 2010). Located at Jülich, Germany. The first PRACE Tier-0 system

Parallel Computing Everywhere...

Today, all computers are parallel computers !

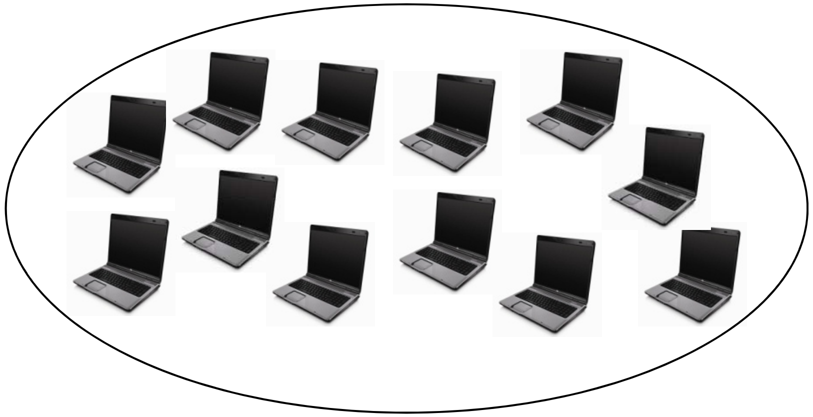

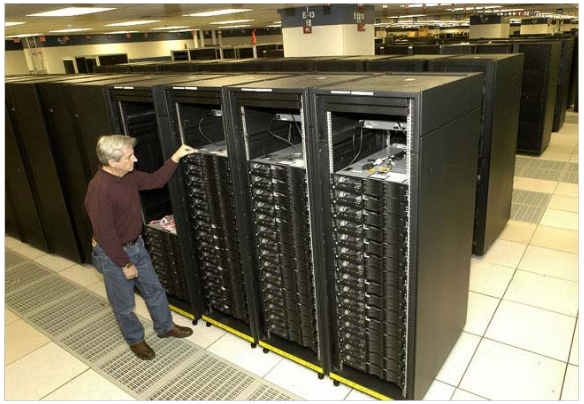

Multicore Processors

The Multicore Era is here!

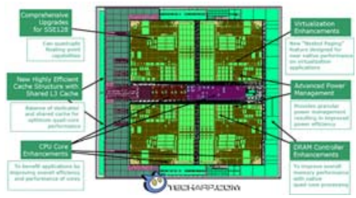

The introduction of multicore processors has a profound impact on parallel computing

Individual processor [core] costs will drop so quickly that in the near future the world can view processors as a nearly free commodity [IDC presentation on HPC at International Supercomputing Conference, ICS07, Dresden, June 26-29 2007]

Parallel Computing Example 1

Quantum Molecular Dynamics

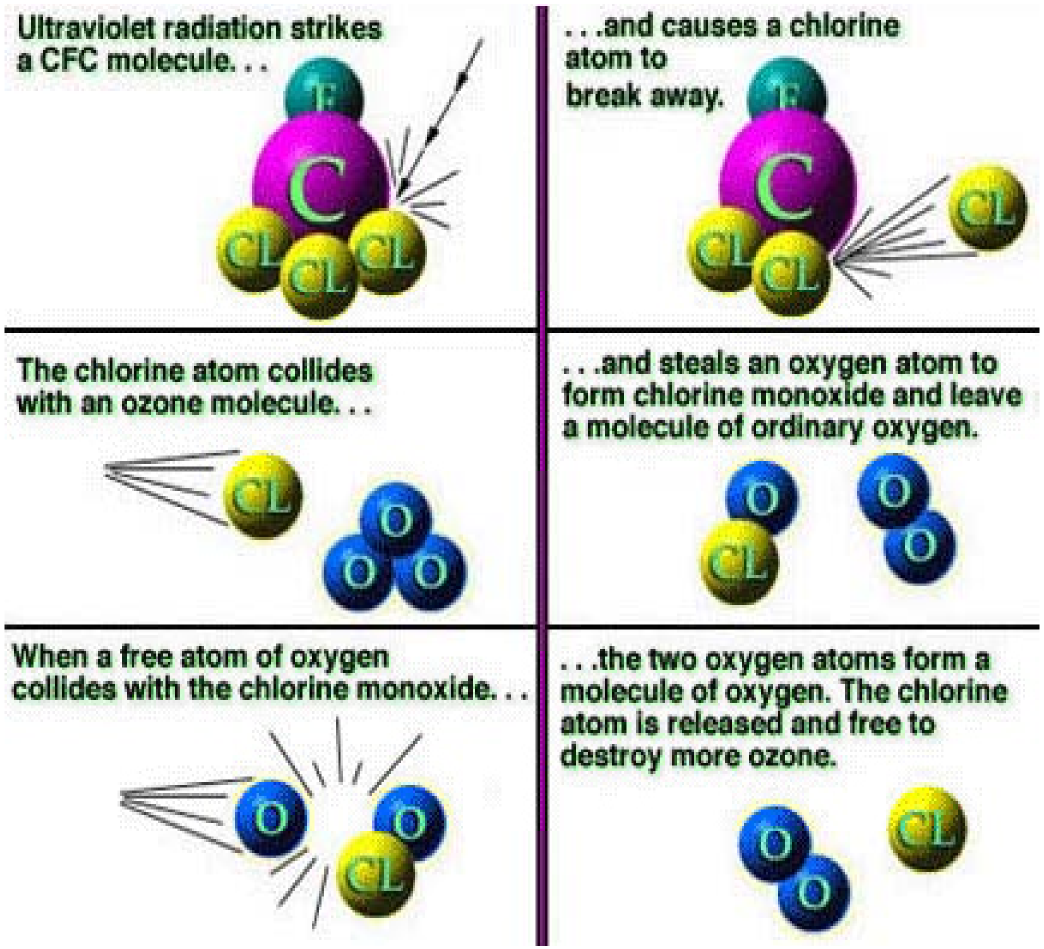

The Time-Dependent Schrödinger Equation

Dynamics of a quantum system is described by the (linear) PDE

$$i\frac{\partial\Phi}{\partial t} = H(x,t)\Phi$$

- H: Sum of kinetic and potential energy operators.
- $|\Phi(x, t)|^2$ is the probability of finding the system at (x, t).
- Example: CCl_3F (Freon-11): 5 nuclei, 66 electrons. x is 211-dimensional (!!)

- Born-Oppenheimer approximation: Wave packets representing the degrees of freedom for the nuclei evolve on potential surfaces.

Example: IBr oscillations

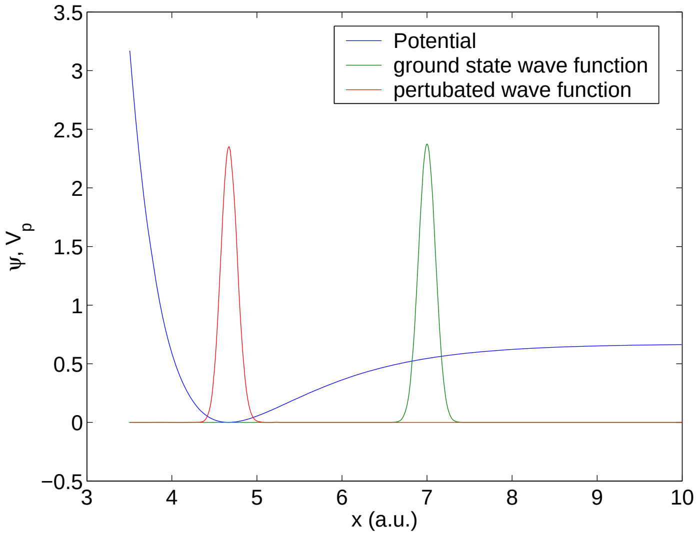

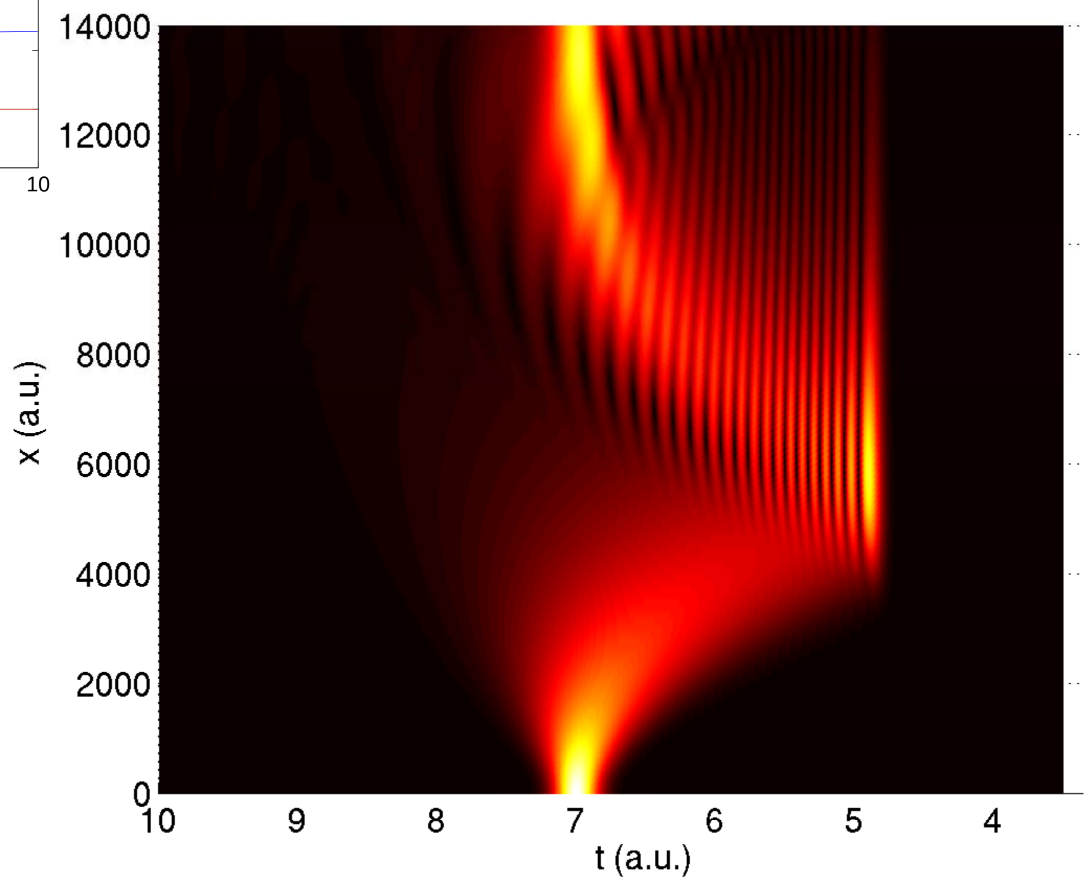

Spatial Discretization

- Block-structured grid in d dimensions
  - High-dimensional problem, d is approx. 3n for n nuclei
  - Currently using an equidistant, static grid
  - h-adaptivity with error control soon...
- High-order finite difference stencils
  - Large strides in data accesses for higher dimensions
  - Data dependencies on block boundaries (ghost cells)
  - High-dimensional ghost cell blocks

Target systems: Large-scale cluster of multicore nodes

- Message passing (MPI) between nodes
- OpenMP for worksharing within each node

Communication Pattern

Communication between MPI processes using “ghostpoints”

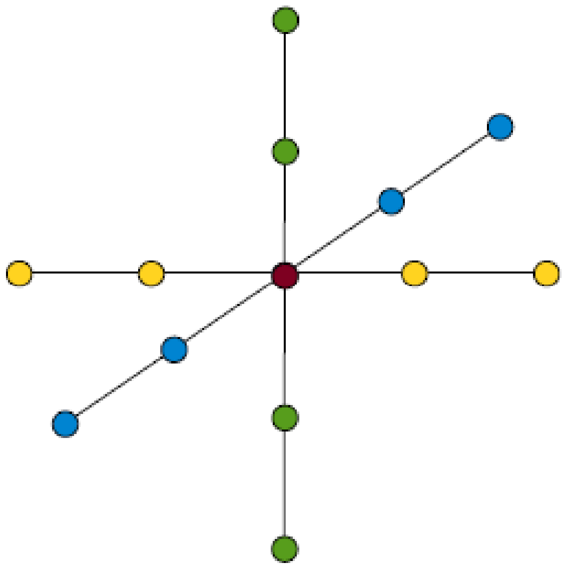

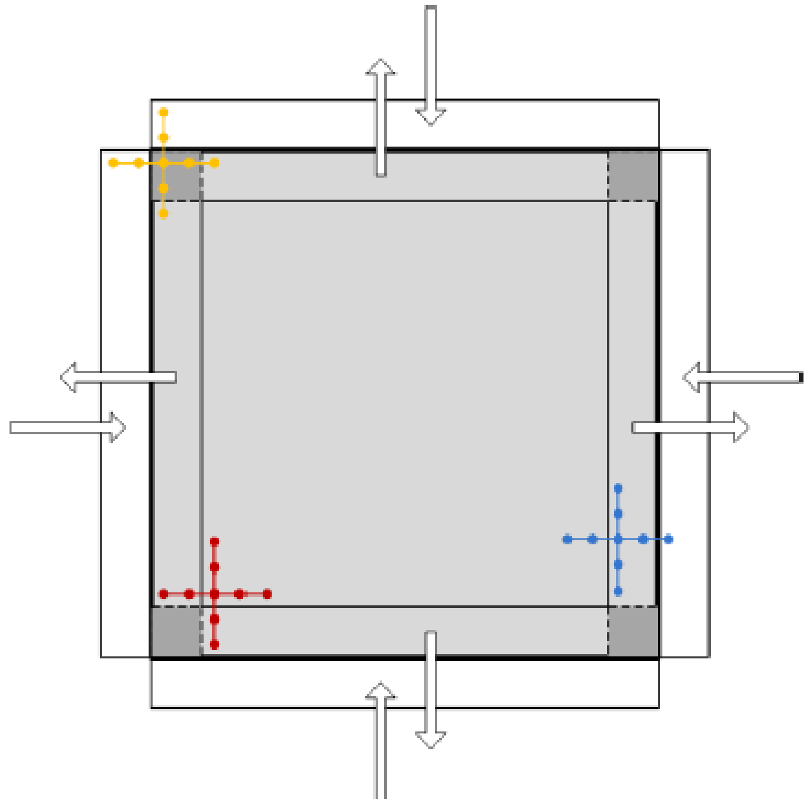

Temporal Discretization 1

- Exponential integrators
  - Allows for long time step
- The Short Iterative Lanczos algorithm
  - Evaluates exp*(-iH)* approx. at low cost
  - In each iteration:
    - Multiplication of wave vector with Hamiltonian matrix (Apply FD stencil, neighbor communication)
    - Two inner products (global synchronization points)

Temporal Discretization 2

- Modified Short Iterative Lanczos algorithm
  - Restructured Lanczos' algorithm brings the two inner products together. Eliminates one synchronization point
  - Combine s Lanczos iterations and compute the corresponding basis vectors at once. This effectively reduces the number of global synchronization points by a factor of s
    - Allows for temporal locality in the multiplication with the Hamiltonian matrix

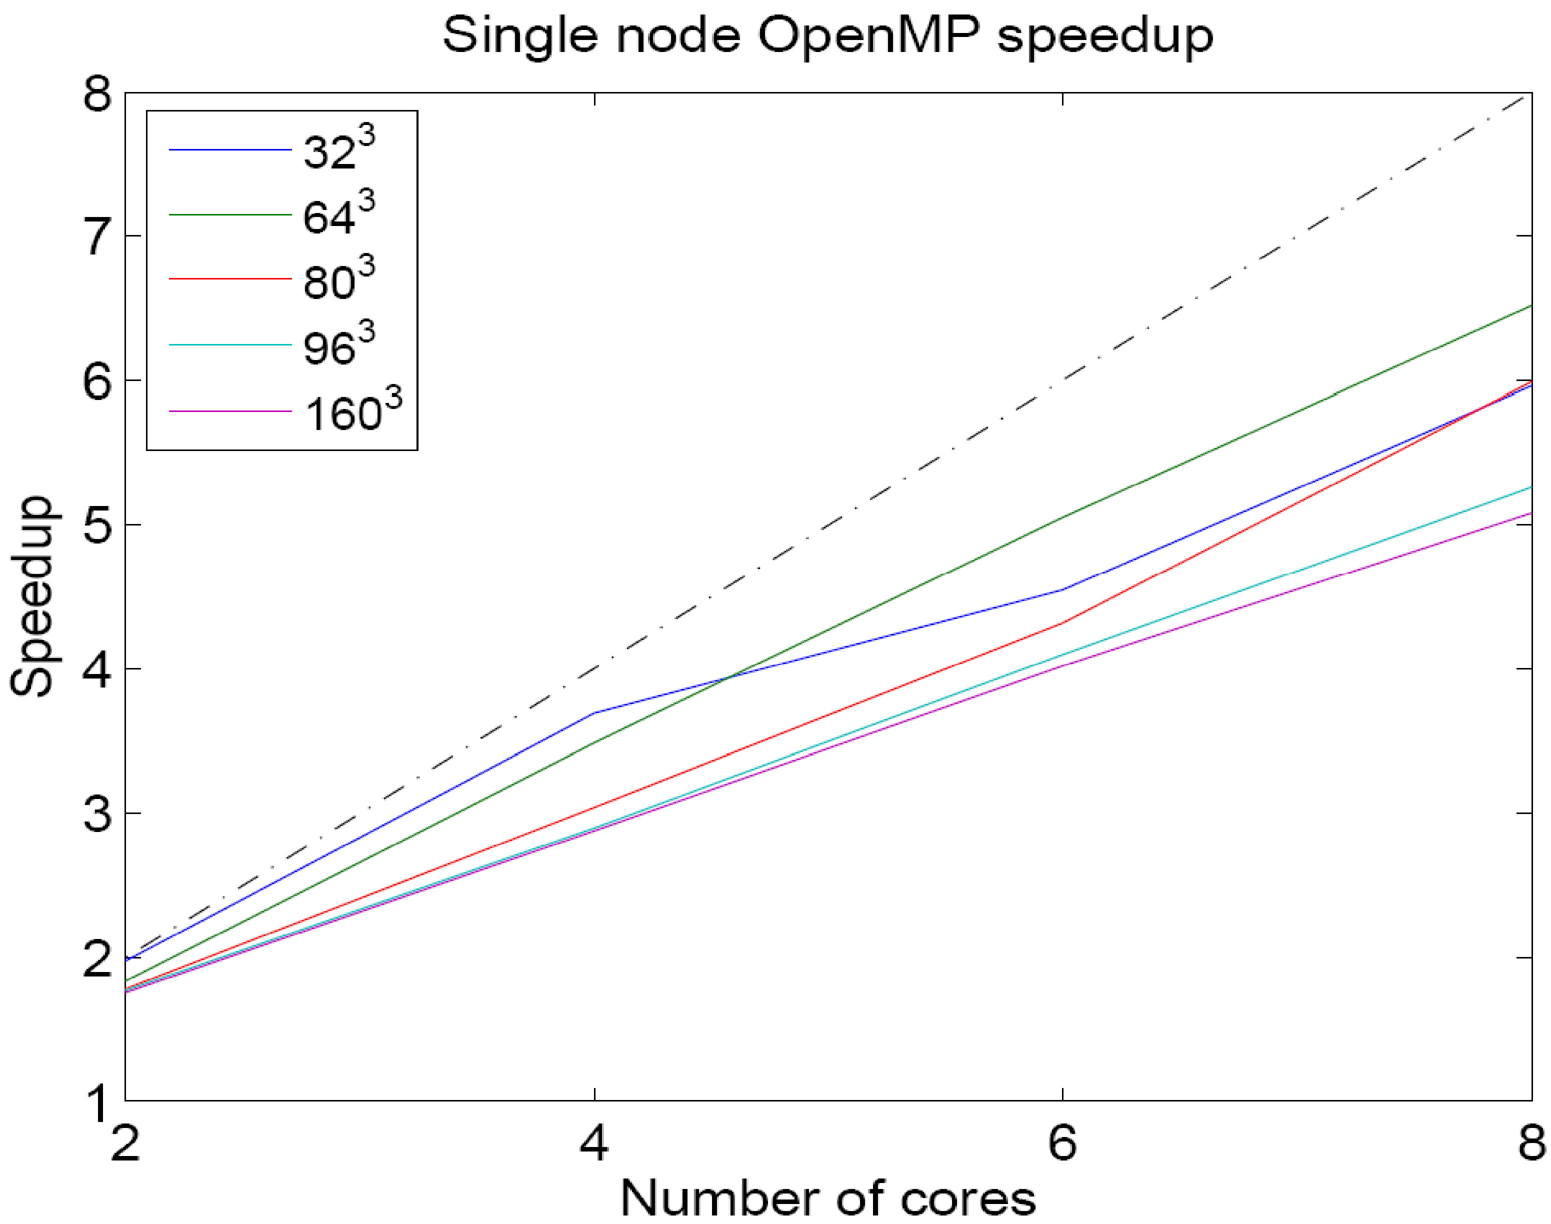
Single node OpenMP speedup
Speedup
Number of cores
32^3
64^3
80^3
96^3
160^3
8
7
6
5
4
3
2
1
2
4
6
8

Some Results (ongoing work…)

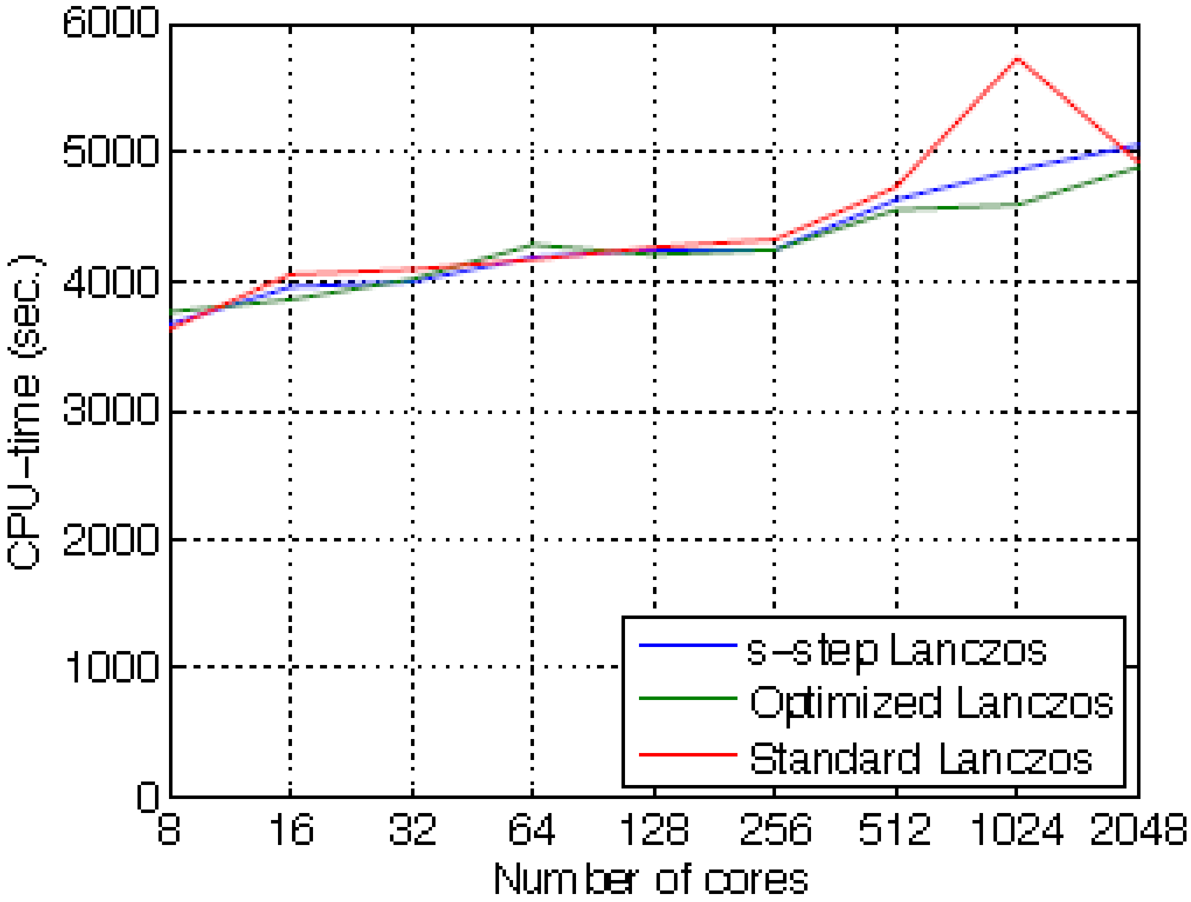

Scaled speedup on cluster, nodes have 8 cores

Parallel Computing Example 2

Genetic Analysis

Quantitative Trait Loci

QTL = A position in the genome affecting a quantitative trait

Quantitative trait = A trait that is measured on a countinous scale (e.g. Body weight, blood presure, crop yield, ...)

QTL model = Statistical model relating genotypes (genetic composition) to phenotypes (trait values) for a population.

Mapping of *d* QTL

D-dimensional global optimization problem

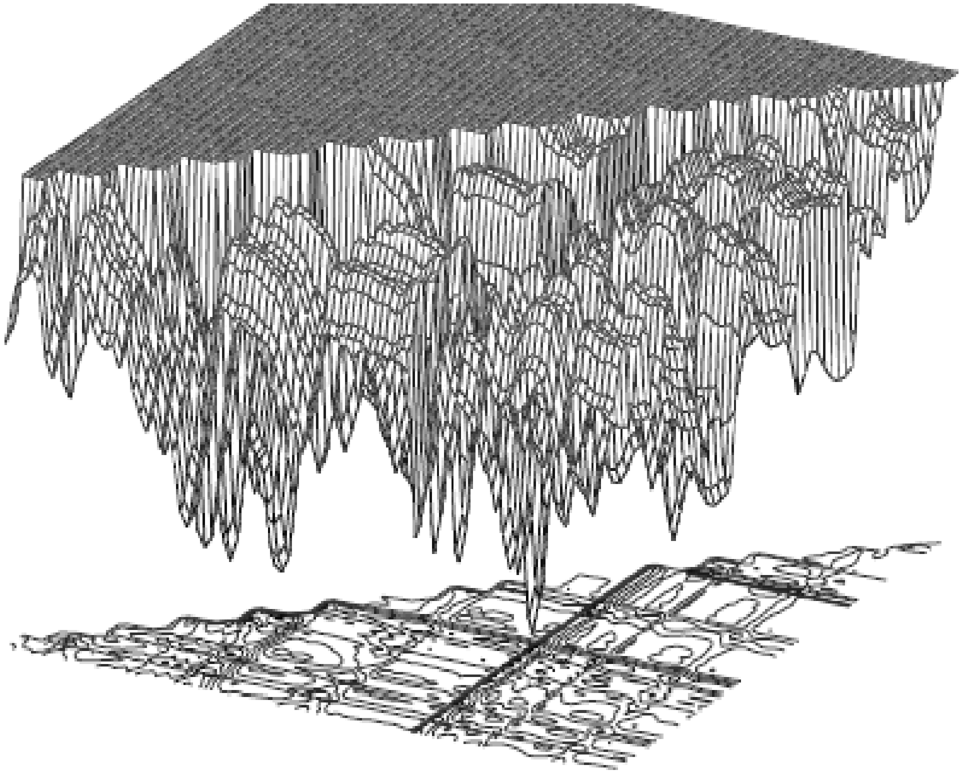

DIRECT

- DIvide search space in RECTangles and evaluate objective function at centers.

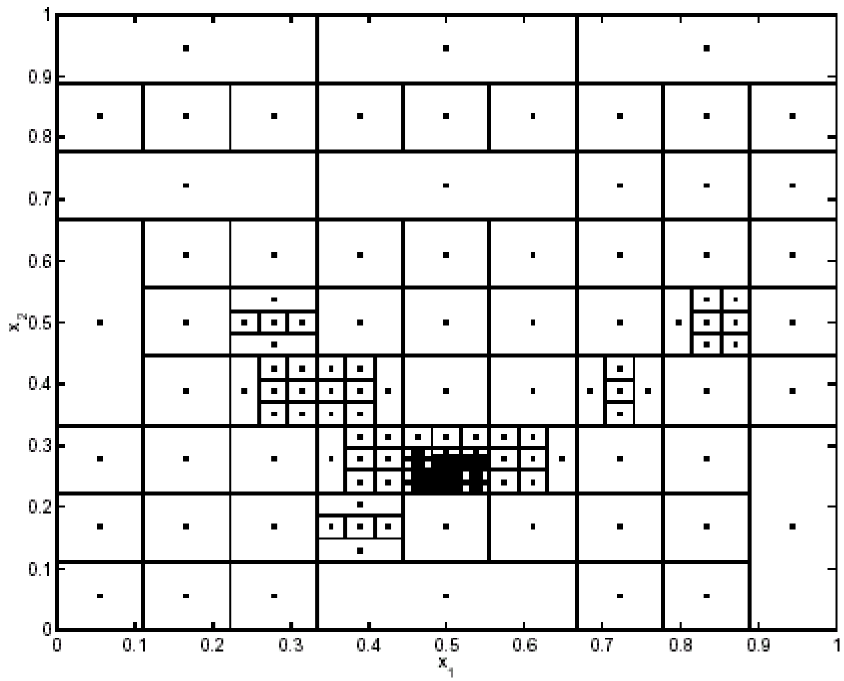

How to Parallelize DIRECT

- Domain decomposition
  - Parallelize over regions in the search space
- Convex hull
  - Parallelize over the rectangles divided in each iteration

Domain Decomposition

- Several instances of DIRECT in sets of boxes
- Cyclic distribution of initial boxes (load balance)
- NOTE: Communication and synchronization is removed from the original algorithm!

Implementation

Outer level (Domain decomposition)

- Execute one DIRECT job for each set of boxes
  (mpi, or grid submission script, or ...)
- Compare the local minima (very quick)

Inner level (Convex hull)

- OpenMP

Results on SweGrid (Six distributed clusters with 64-96 nodes each. Nodes have 2 quad-core Opterons each)