The Good, the Bad and the Lucky

Jürgen Schnack

Department of Physics – University of Bielefeld – Germany

http://obelix.physik.uni-bielefeld.de/∼schnack/

ScaleMP Workshop

RWTH Aachen, Aachen, October 16, 2009

The Good

Good I – Lanczos scaling

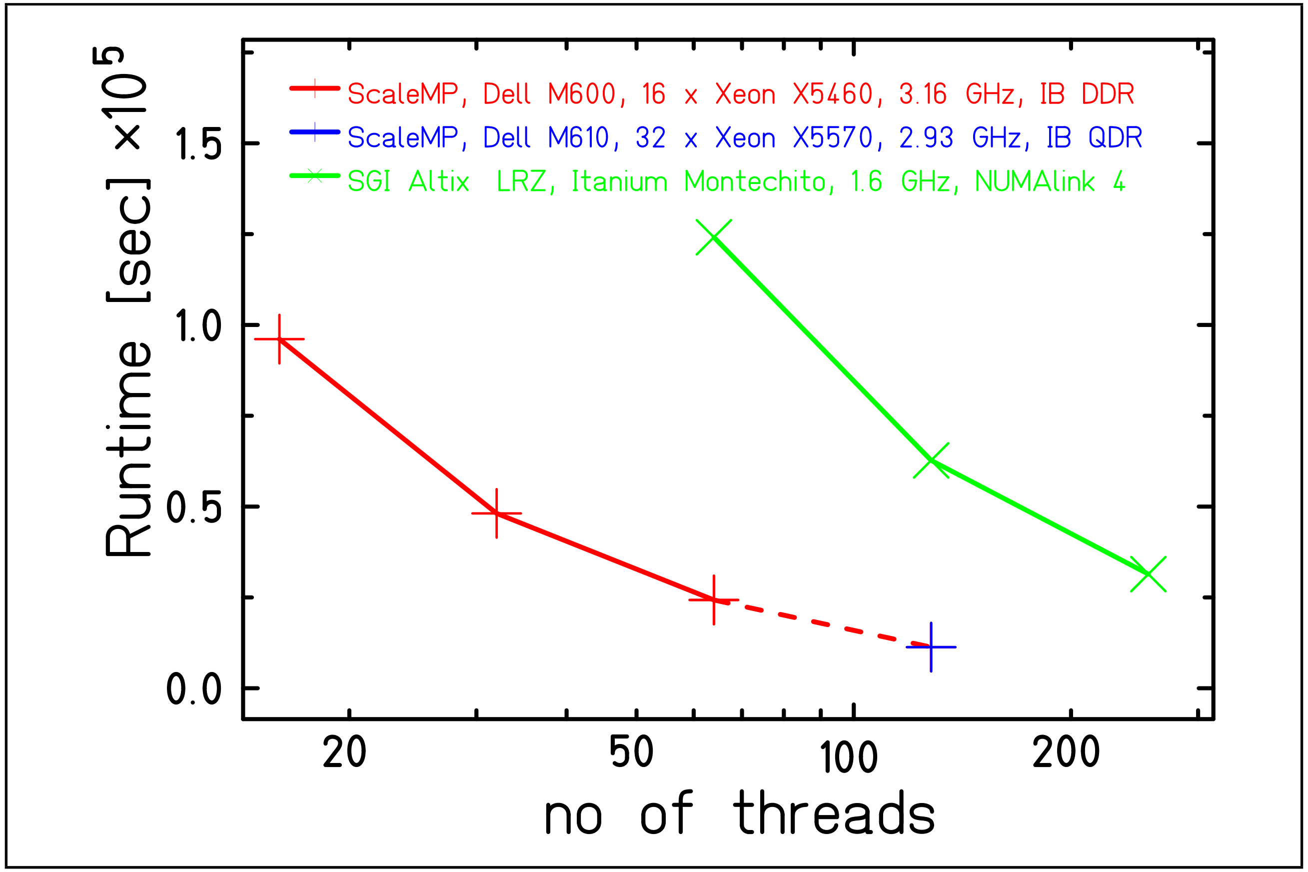

Almost perfect scaling of openMP parallelized Lanczos code: the SGI Altix (LRZ) is compared to new DELL machines using ScaleMP. The dimension of the underlying Hilbert space, i.e. the length of the used Lanczos vectors is 601,080,390.

Good II – Lanczos background

$$\langle\, i \,|\, \psi_2 \,\rangle \;=\; \sum_{j=1}^{\text{dim}} \langle\, i \,|\, \underset{\sim}{H} \,|\, j \,\rangle \langle\, j \,|\, \psi_1 \,\rangle$$

1. Matrix $\langle\, i \,|\, \underset{\sim}{H} \,|\, j \,\rangle$ sparse. Recursive algorithm. Only few eigenvalues targeted.

2. openMP-parallelization of loops
   - BAD: $j \Rightarrow$ writing conflicts in i.
   - GOOD: $i \Rightarrow$ independent reading possible.

3. Analytical coding of states $|\, j \,\rangle$, i.e. physical meaning $\Leftrightarrow$ integer lables.

4. Evaluation of non-zero $\langle\, i \,|\, \underset{\sim}{H} \,|\, j \,\rangle$ *on the fly*.

J. Schnack, P. Hage, H.-J. Schmidt, *Efficient implementation of the Lanczos method for magnetic systems*, J. Comput. Phys. **227** (2008) 4512

The Bad

Bad I – Full matrix diagonalization

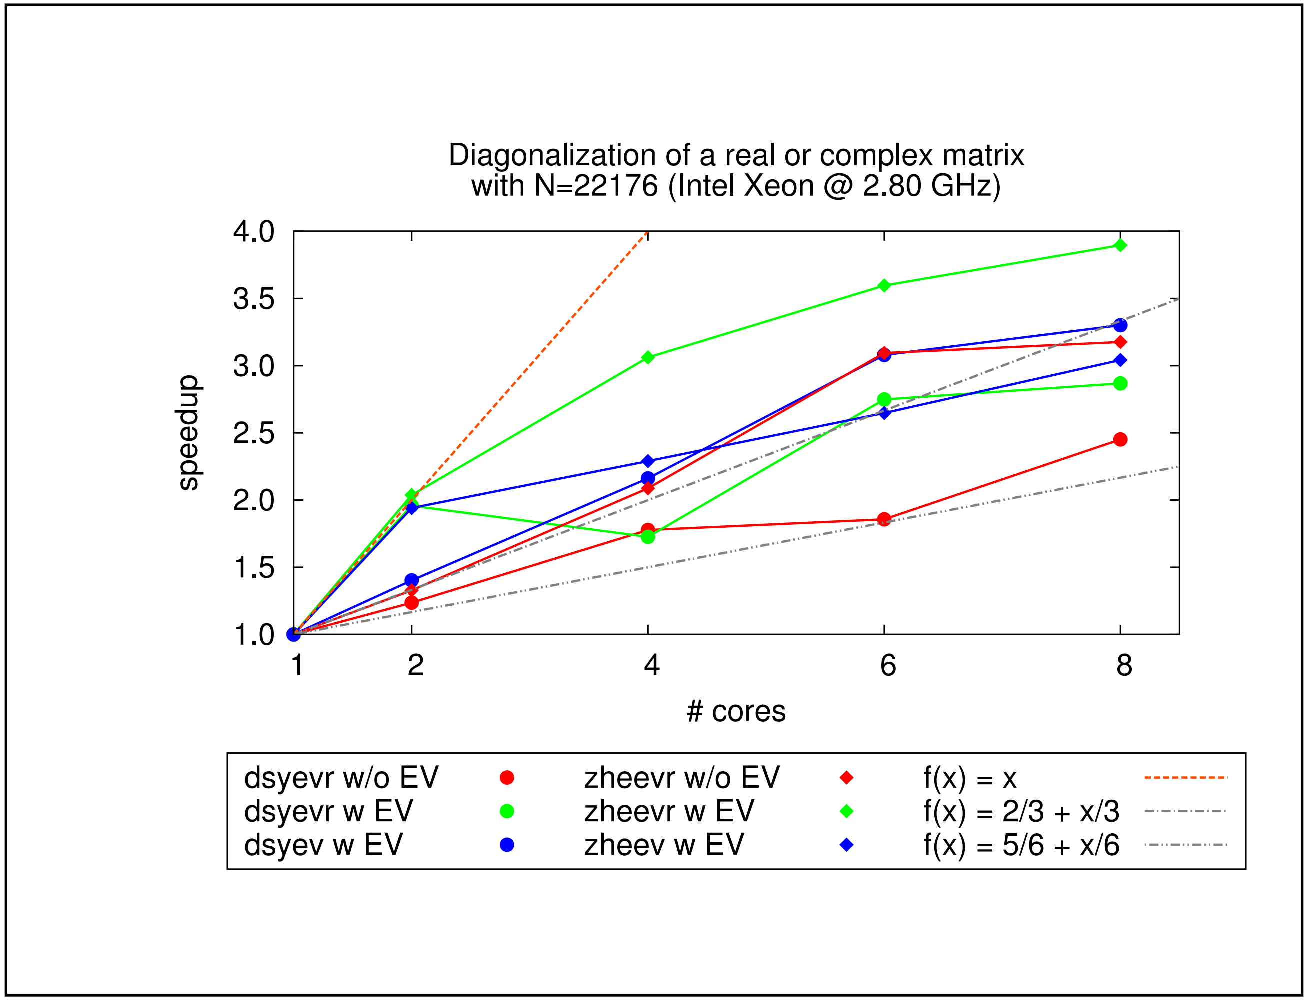

Scaling of full matrix diagonalization (Martin Höck).

Bad II – Full matrix diagonalization

$$\sum_{j=1}^{\text{dim}} \langle\, i \,|\, \underset{\sim}{H} \,|\, j \,\rangle\langle\, j \,|\, \psi_n \,\rangle \quad = \quad E_n \,\langle\, i \,|\, \psi_n \,\rangle$$

1. Goal: All eigenvalues E_n and eigenvectors $\langle\, i \,|\, \psi_n \,\rangle$.
   Essential for quantum mechanics and thermodynamics.

2. MKL routines, e.g. zheev(d/r), already openMP parallelized.
   - Scaling VERY poor, i.e. up to 4 . . . 8 acceptable.
   - More than 8 cores run at a 100 %, but need longer.

3. New approaches needed (1,2).

(1) A. Honecker, J. Schüle, in C. Bischof, M. Bücker, P. Gibbon, G. Joubert, T. Lippert, B. Mohr, F. Peters, editors, *Advances in Parallel Computing*, volume 15, IOS Press (2008)
(2) Paolo Bientinesi, Inderjit S. Dhillon, Robert A. van de Geijn, *A Parallel Eigensolver for Dense Symmetric Matrices Based on Multiple Relatively Robust Representations*, SIAM Journal on Scientific Computing, Vol. 27, No. 1, 2005.

The Lucky

Lucky I – Me

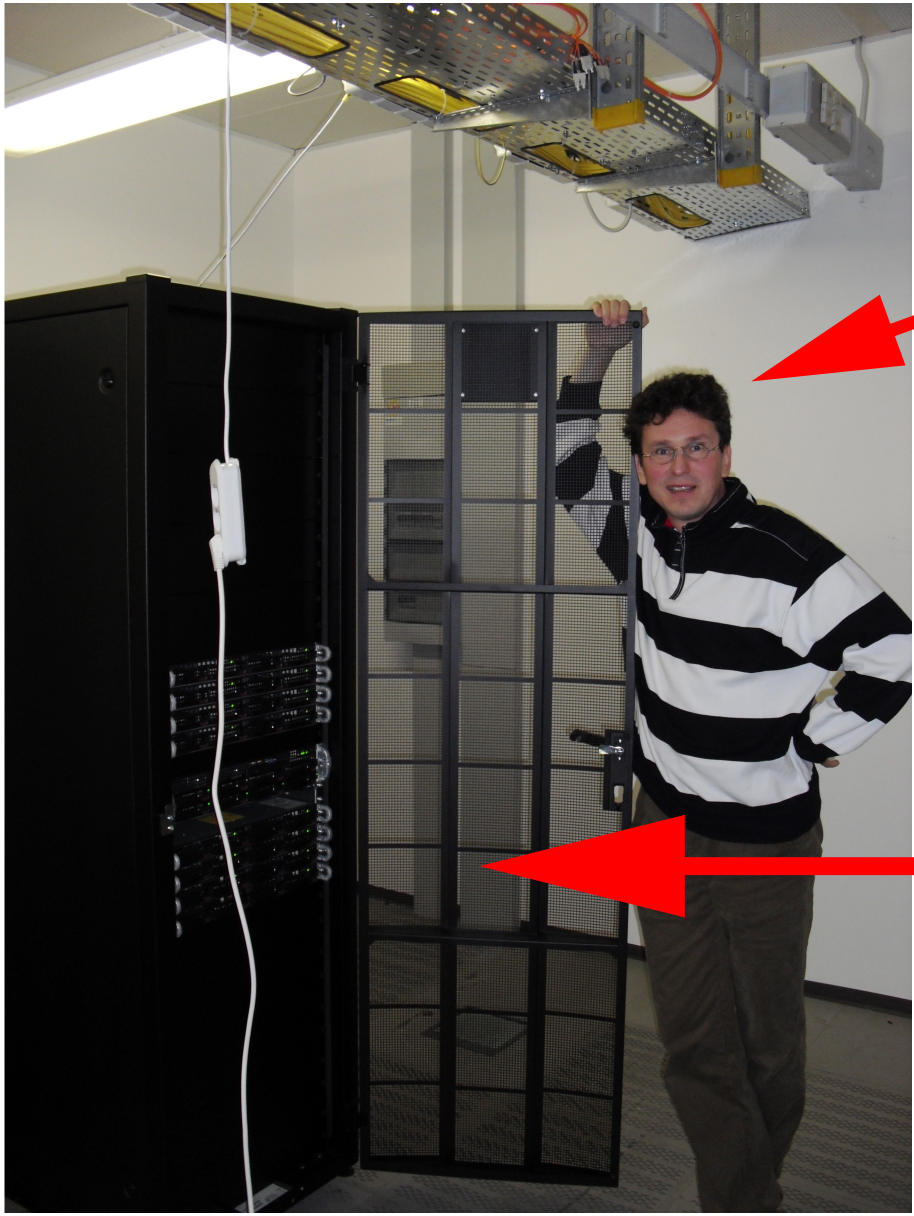

“cell professor”

128 cores, 384 GB RAM

Lucky II – BULL and ScaleMP

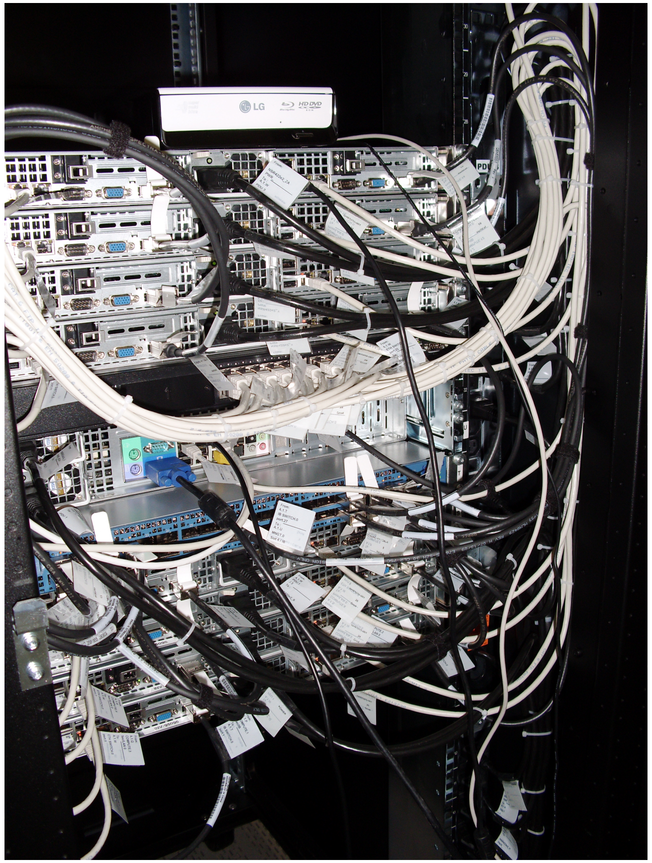

- 16 boards
- 2 Quad Nehalem + $2 \times 3 \times 4$ GB RAM per board
- IB, QDR
- login node
- almost up

Thank you very much for your attention.

German Molecular Magnetism Web

www.molmag.de

Highlights. Tutorials. Who is who. DFG SPP 1137.