ab initio Electronic Structure Calculations

New scalability frontiers using the BG/L Supercomputer

C. Bekas, A. Curioni and W. Andreoni
IBM, Zurich Research Laboratory
Rueschlikon 8803, Switzerland

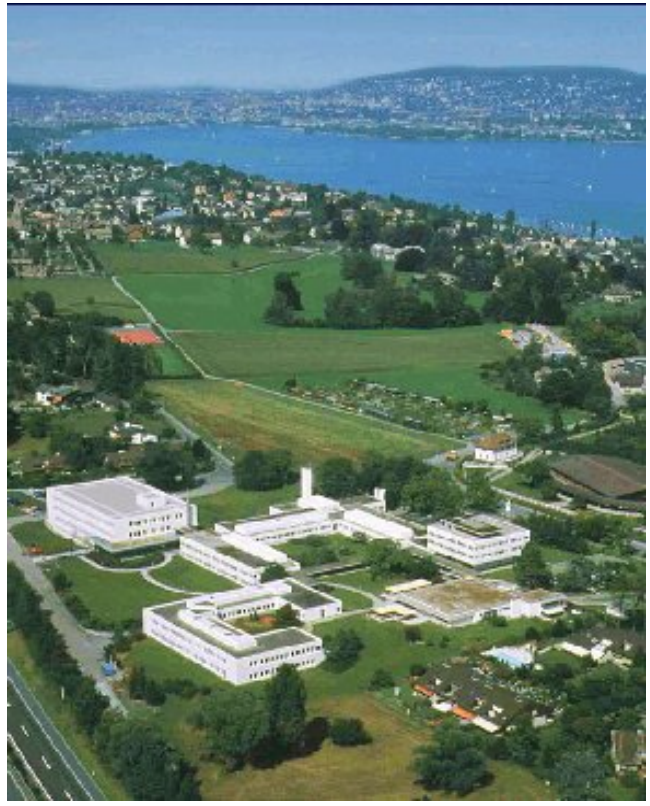

ab initio Electronic Structure Calculations: Why?

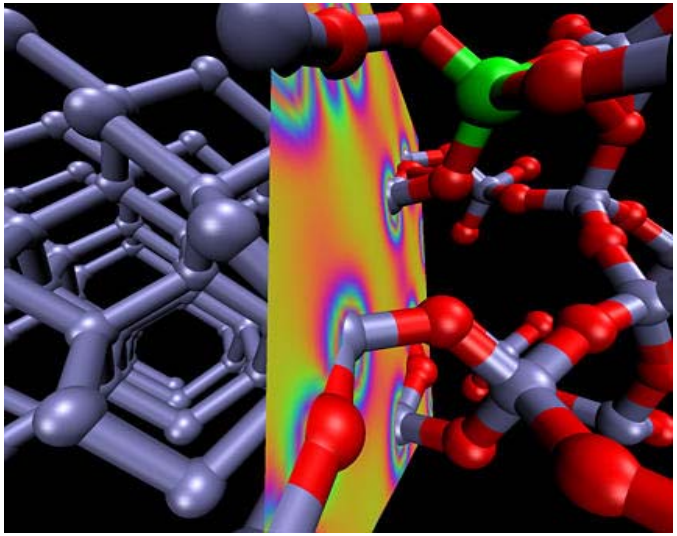

Semiconductors

Drug-protein interaction

***ab initio* electronic structure calculations in Molecular Dynamics (MD):**

- **Simulate the behaviour of materials at the atomic level by applying the fundamental laws of Quantum Mechanics**

Why is it different than classical MD?

- **The interaction between electrons is NOT modelled by empirical force fields but described quantum-mechanically**
- **Greatly improved accuracy (vs classical MD) at the nano level! …though it comes at a cost…**

In this talk

- Introduction to the mathematical formulation of *ab initio* calculations…
- in particular…the Density Functional Theory (DFT)…formulation and
- The parallel plane wave implementation of *ab initio* calculations … in which we
- identify the computationally intensive "spots"… 3D *parallel FFT transforms*
- *Task Groups strategy for parallel 3D FFTs*
- Showcase code: *CPV*, included in www.quantum-espresso.org

Computational platform: Blue Gene /L Supercomputer

- Massively parallel system with key features particularly suitable for *ab initio* calculations:
  - ✓ Dual core processors, delivering close to peak performance 2.8 Gflop/s (5.6 per node) for key linear alg. kernels (i.e. DGEMM)
  - ✓ Very fast-scalable interconnection network, facilitating both 3D Torus and Tree network, ideal for parallel 3D FFT kernels
  - ✓ Demonstrated performance for *ab initio* calculations: 110.4 Tflop/s for the CPMD code (www.cpmd.org)

Formulation: The Wave Function

We seek to find the *steady state* of the *electron distribution*

- **All (N) particles are described by a complicated *wave function* ψ …**
- **ψ is a function of a 3N dimensional space …in particular it is determined by**
  - **The position r_k of all (N) particles (including nuclei and electrons)**
  - **It is normalized in such a way that**

$$\int |\psi|^2 dr_1 dr_2 dr_3 ... = N$$

Max Born's probabilistic interpretation: Considering a region *D in r*, then

$$Pr(\psi, D) = \int_D |\psi|^2 dr_1 dr_2 dr_3 ...$$

…describes the probability of all particles being in region *D*. Thus: the distribution of electrons e_i in space is defined by the wave function ψ

Mathematical Modeling: The Hamiltonian

Steady state of the electron distribution:

- **it is such that it minimizes the total energy of the molecular system…(energy due to dynamic interaction of all the particles involved because of the forces that act upon them)**

***Hamiltonian H* of the molecular system:**

- ***Operator* that governs the interaction of the involved particles…**
- **Considering all forces between nuclei and electrons we have…**

$$H = H_{nucl} + H_e + U_{nucl} + V_{ext} + U_{ee}$$

H_{nucl}	*Kinetic energy of the nuclei*
H_e	*Kinetic energy of electrons*
U_{nucl}	*Interaction energy of nuclei (Coulombic repulsion)*
V_{ext}	*Nuclei electrostatic potential with which electrons interact*
U_{ee}	*Electrostatic repulsion between electrons*

Density Functional Theory (DFT)

High complexity is mainly due to the many-electron formulation of *ab initio* calculations…is there a way to come up with an one-electron formulation?

Key Theory

DFT: Density Functional Theory (Hohenberg,Kohn,Sham)

- ✓ **The total ground energy of a molecular system is a functional of the ground state electronic density…(number of electrons in a cubic unit)**
- ✓**The energy of a system of electrons is at a minimum if it is an exact density of the ground state!**

- **This is an existence theorem…the density functional always exists**
- **…but the theorem does not prescribe a way to compute it…**
- **This energy functional is highly complicated…**
- **Thus approximations are considered…concerning:**
  - **Kinetic energy and**
  - **Exchange-Correlation energies of the system of electrons**

One electron Schrödinger's equation

Let the columns of Ψ:

$$\Psi = [\psi_1, \psi_2, \ldots, \psi_N]$$

hold the wave functions corresponding to the electrons…Then it holds that

$$H\Psi = \epsilon\Psi$$

- **This is an eigenvalue problem…that becomes a usual…**
- **"algebraic" eigenvalue problem when we discretize ψ_i in a suitable basis**
- **Difficult and nonlinear problem…since:**
- **Hamiltonian and wave functions depend upon each other through the unknown density functional**

Variational Principle in the DFT framework (in simple terms!)

Minimal energy and the corresponding electron distribution amounts to calculating the M smallest eigenvalues/eigenvectors of the Schrödinger equation, where M is the number of electrons of the system

Density functional Theory: Formulation (1/2)

Equivalent eigenproblem (r is space):

$$[H_{e_i} + V_{tot}(r;\rho)]\psi_i(r) = \epsilon_i\psi_i(r)$$

ψ_i	One electron wave function
ϵ_i	Energy of the i-th state of the system
H_{e_i}	Kinetic energy of electron e_i
$V_{tot}(r;\rho)$	Total potential that acts on e_i at position r
$\rho(r) = \sum_{occ} \lvert\psi_i(r)\rvert^2$	Charge density at position r

Density functional Theory: Formulation (2/2)

Furthermore:

$$V_{tot}(r;\rho) = V_{ion} + V_H + Vxc$$

V_{ion} **Potential due to nuclei and core electrons**

V_H **Coulomb potential from valence electrons** $\nabla^2 V_H = -4\pi\rho$

V_{xc} **Exchange-Correlation potential…function of the charge density ρ**

Non-linearity: ***The Hamiltonian depends upon the charge density ρ while ρ itself depends upon the wave functions (eigenvectors) ψ_i***

Some short of iteration is required until convergence is achieved!

Fourier Space discretization: Plane Waves (PW)

Using a plane wave basis to express the wave functions has a number of advantages:

- ✓ **Natural basis for periodic systems like crystals**
- ✓ **Isolated systems can be readily represented using periodic cells**
- ✓ **Most parts of the Hamiltonian become diagonal matrices in the PW basis**
- ✓ **Maximum cut off frequency is determined by the kinetic energy operators (Laplacian) allowing for easy systematic control of accuracy**

Quantities are calculated either in real or Fourier space depending to where they are expressed locally. Efficient parallel 3D FFT transforms become crucial in PW *ab initio* calculations

Example: Charge density: $\rho(r) = \sum_{occ} |\psi_i(r)|^2$

Requires double summation over plane-waves but only a short loop in real space. Major cost kernel. Opportunity for hierarchical parallelism.

The Blue Gene/L Supercomputer

System

64 Racks, 64x32x32

180/360 TF/s
32 TB

Rack

2.8/5.6 TF/s
512 GB

(32 chips 4x4x2)
16 compute, 0-2 IO cards

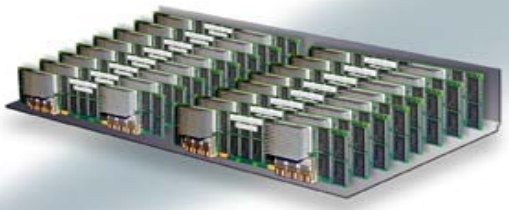

90/180 GF/s
16 GB

2 chips, 1x2x1

5.6/11.2 GF/s
1.0 GB

2 processors

2.8/5.6 GF/s
4 MB

BG/L Interconnection Network

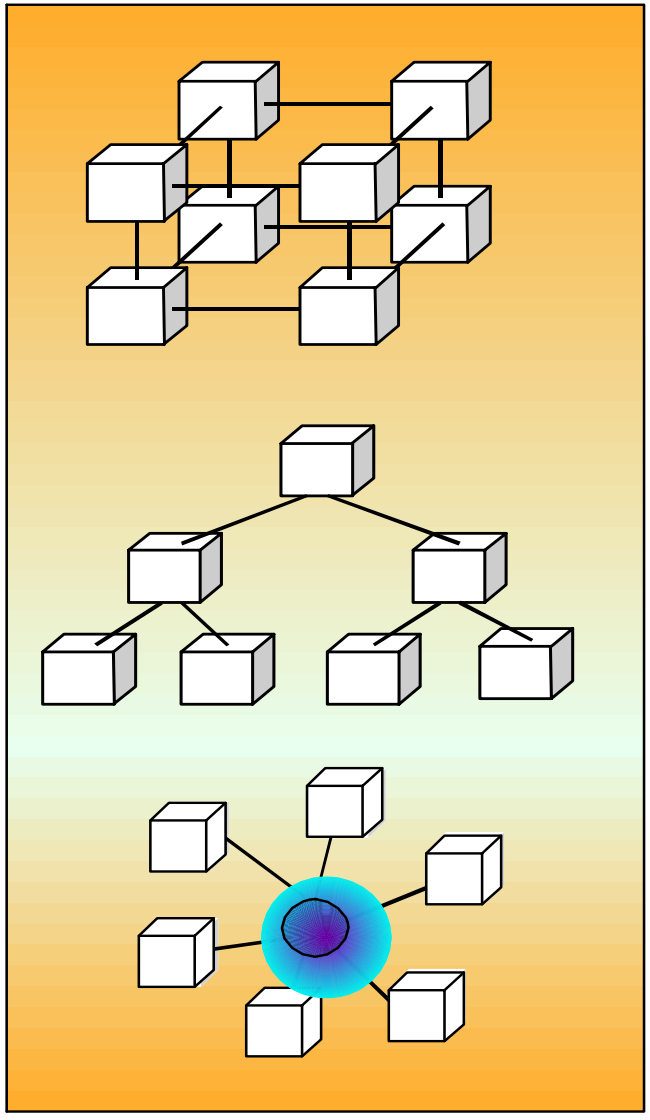

3 Dimensional Torus

- Interconnects all compute nodes (65,536)
- Virtual cut-through hardware routing
- 1.4Gb/s on all 12 node links (2.1 GB/s per node)
- Communications backbone for computations
- 0.7/1.4 TB/s bisection bandwidth, 67TB/s total bandwidth

Global Collective Network

- One-to-all broadcast functionality
- Reduction operations functionality
- 2.8 Gb/s of bandwidth per link; Latency of tree traversal 2.5 µs
- ~23TB/s total binary tree bandwidth (64k machine)
- Interconnects all compute and I/O nodes (1024)

Low Latency Global Barrier and Interrupt

- Round trip latency 1.3 µs

Control Network

- Boot, monitoring and diagnostics

Ethernet

- Incorporated into every node ASIC
- Active in the I/O nodes (1:64)
- All external comm. (file I/O, control, user interaction, etc.)

Distributed Memory Implementation in CPV

Distribute PWs and real space mesh following such that:

- Each processor has the same number of plane waves
- A processor hosts full planes of real-space grid points
- All plane waves with common y and z components are on the same processor
- The number of different (y, z) pairs of plane-wave components is the same on each processor
- The number of real-space planes is the same on each processor

Distributed Memory 3D FFT

For each wave function: Distribute its coefficients over the G-vectors across the z-direction, thus forming "pencils"

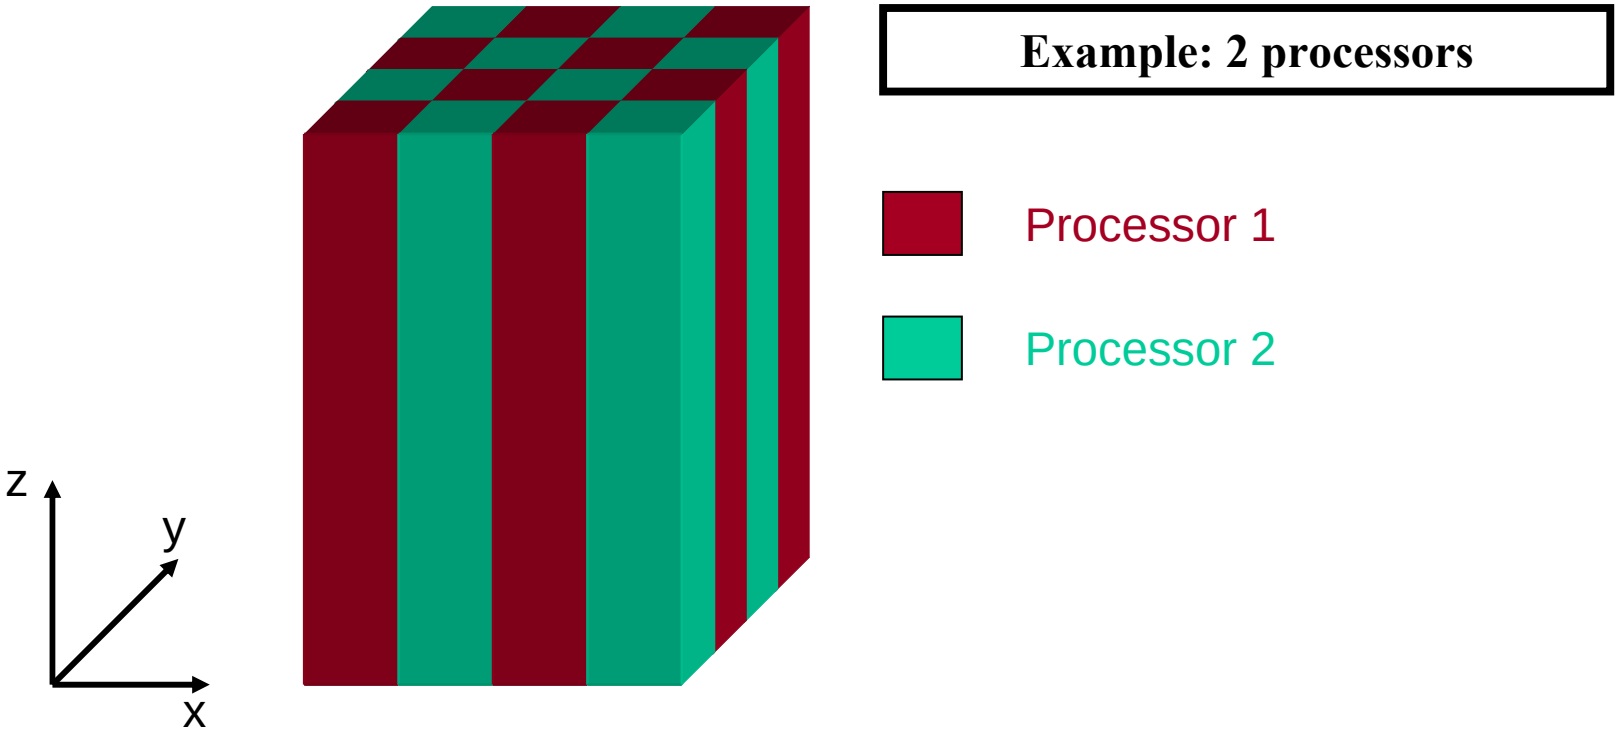

Distributed Memory Implementation in CPV

3D FFT: can be computed in 3 steps

1D FFT across Z ⇒ 1D FFT across Y ⇒ 1D FFT across X

…or 3D FFT in two steps

1D FFT across Z ⇒ 2D FFT across X-Y planes

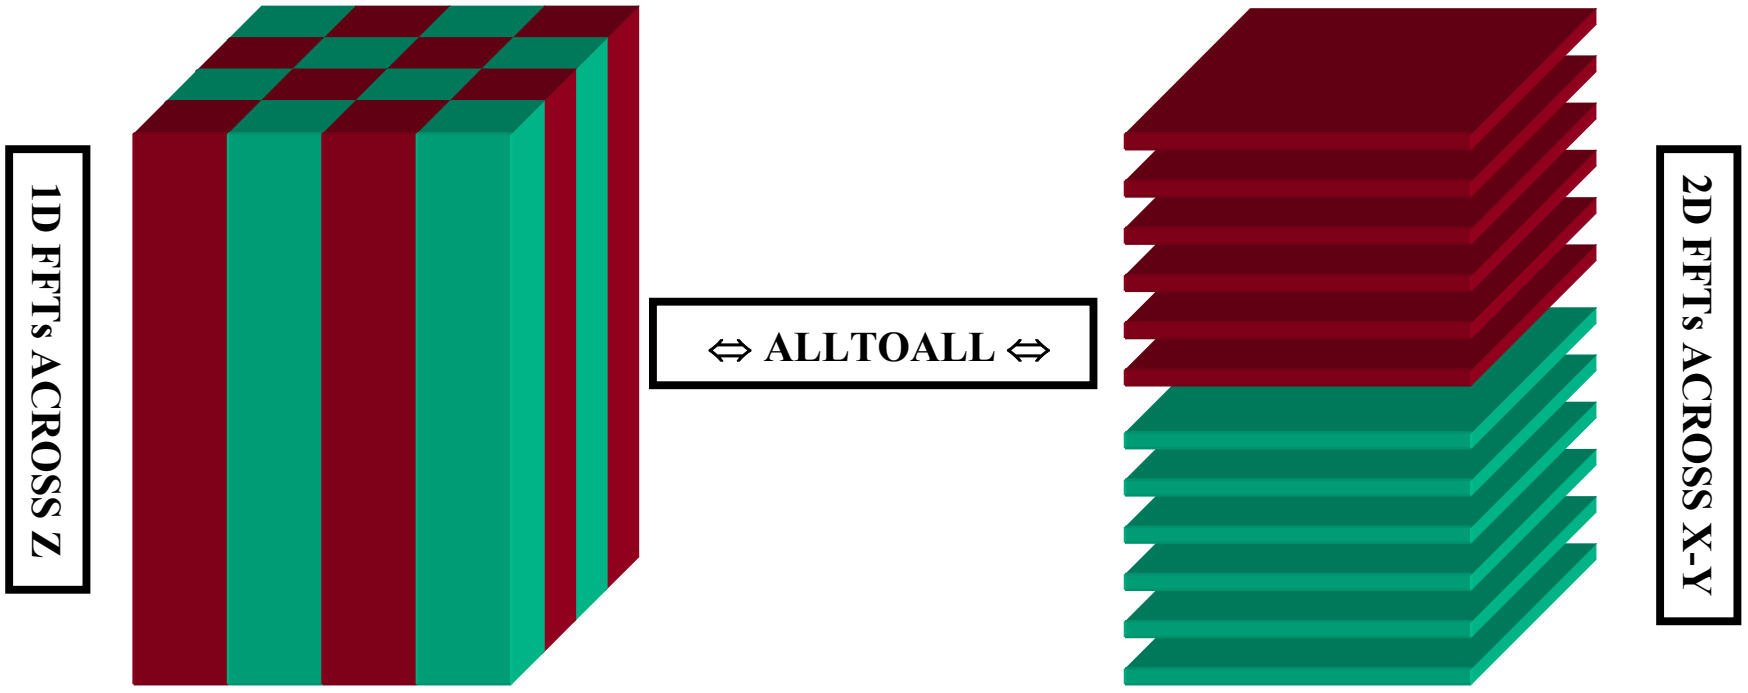

Limited Scalability of Standard 3D FFT

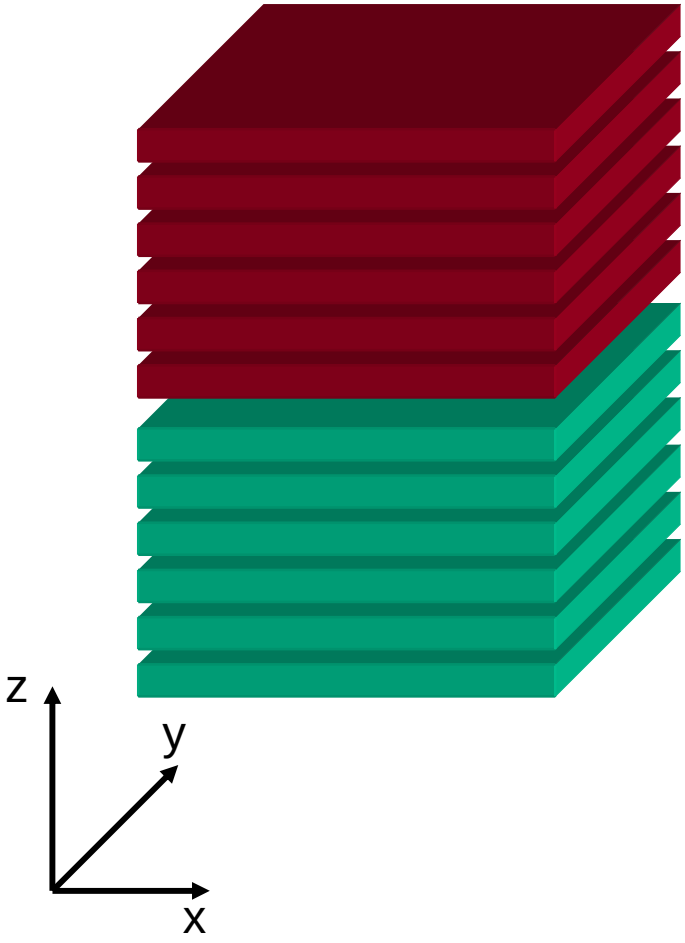

Each processor takes a number of whole planes…

Very good scheme for small – medium sized computational platforms…but

Observe that scalability is limited by the number of planes across the Z-direction! ... Which is in the order of a few hundreds…O(100)…

Thus: not appropriate for the BG/L system

3D FFT Using Task Groups

$$\rho(r) = \sum_{occ} |\psi_i(r)|^2$$

Loop across the number of electrons. Each iteration requires 1 3D FFT.

Hierarchical parallelism*: Assign to each Task Group a number of iterations.

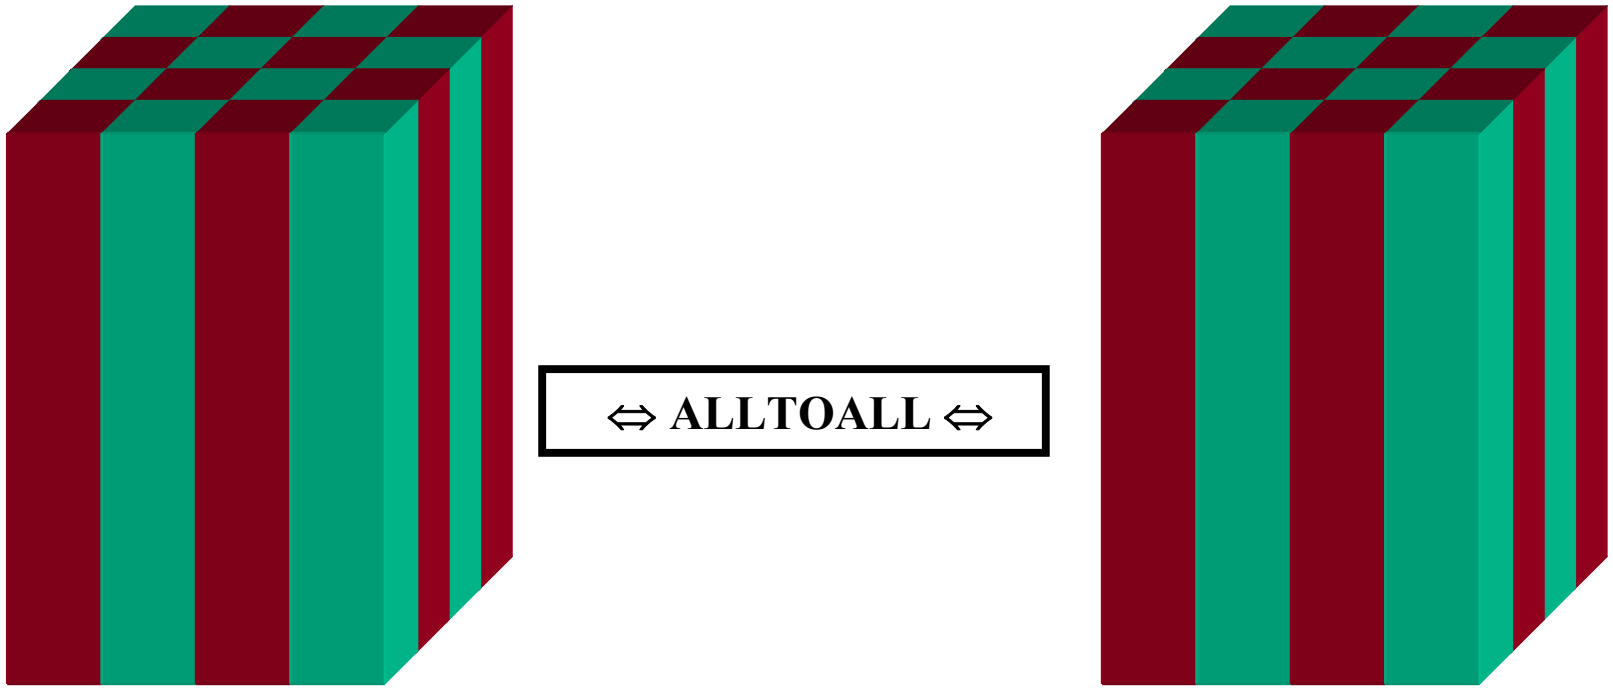

* J. Hutter and A. Curioni, Parallel Computing (31) 1, 2005

3D FFT Using Task Groups

- **The Task Groups of processors will work on different eigenstates concurrently**
- **Number of processors per group: Ideally the one that achieves the best scalability for the original parallel 3D FFT scheme**

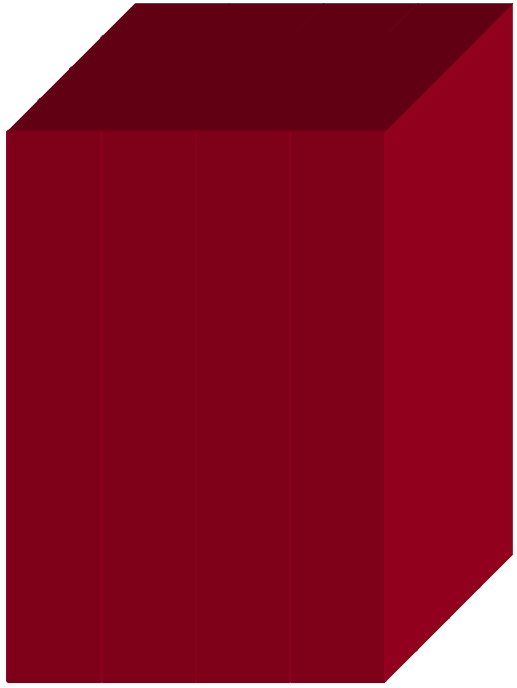

EIG 1: ONLY PROC 1

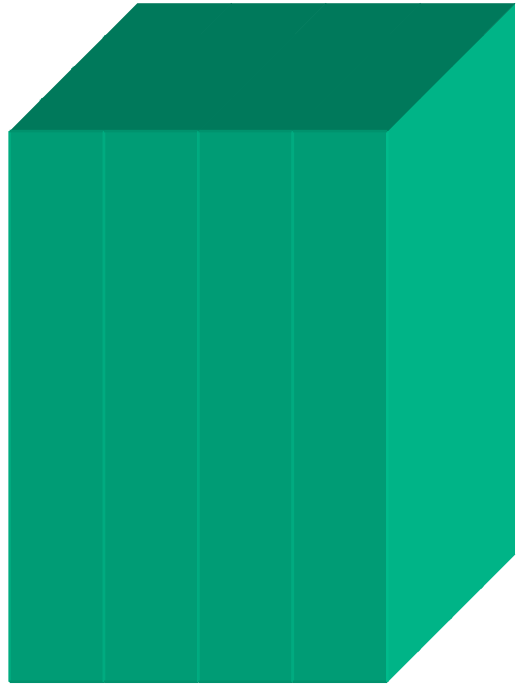

EIG 2: ONLY PROC 2

Numerical Example

80 H_2O molecules (A. Pasquarello)

- **240 atoms**
- **320 occupied states (i.e. loop length for density)**

FFT mesh z length: 128
Original parallel 3D FFT stops scaling at 128 nodes

TG parallel 3D FFTs continues to scale:
256 nodes…2TGs
512 nodes…4TGs

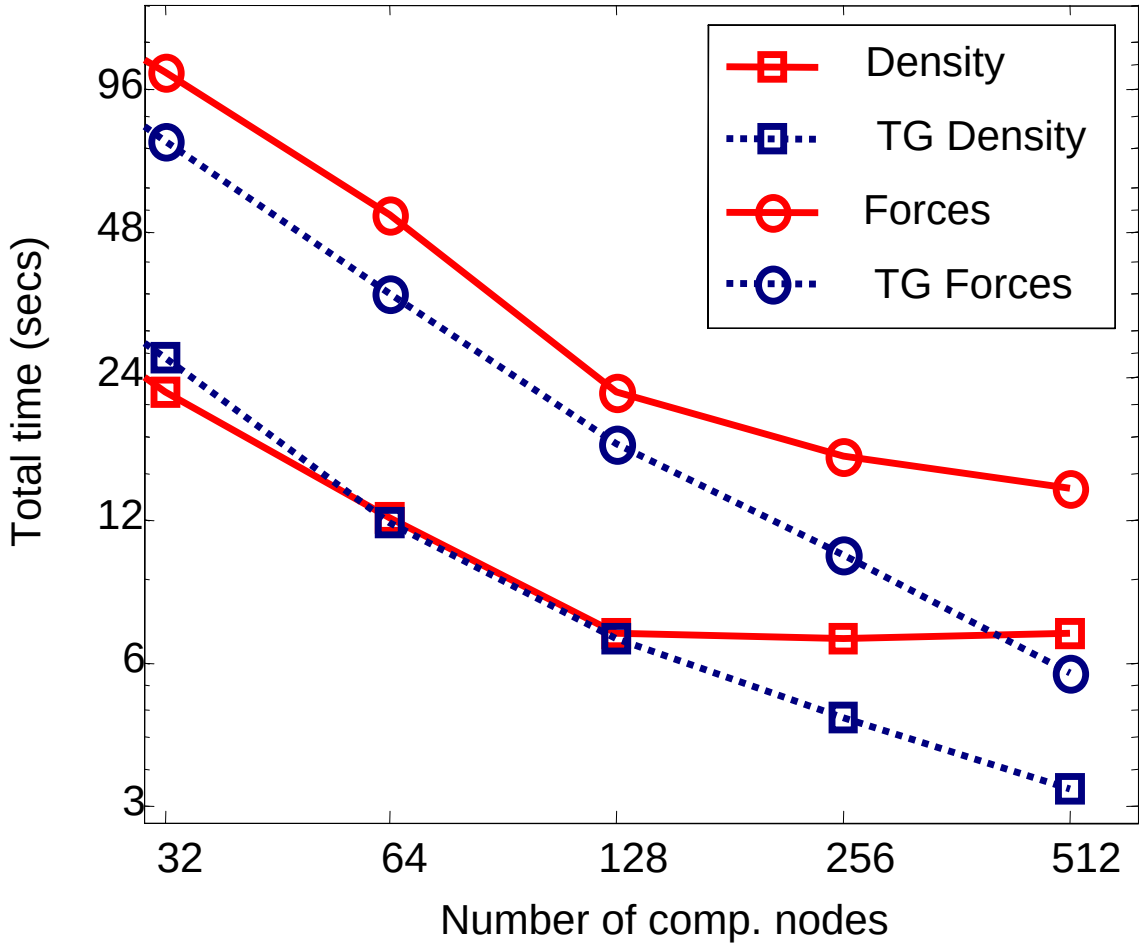

Comparison of Kernels

80 H_2O molecules (A. Pasquarello)
- **240 atoms**
- **320 occupied states (i.e. loop length for density)**

CPV Code: Uses Soft Pseudopotentials for the representation of core electrons. Requires much smaller cutoff for PWs.

FFT kernels are no longer dominant for large # CPUs, but Dense eigenproblems for orthogonalization stand out!

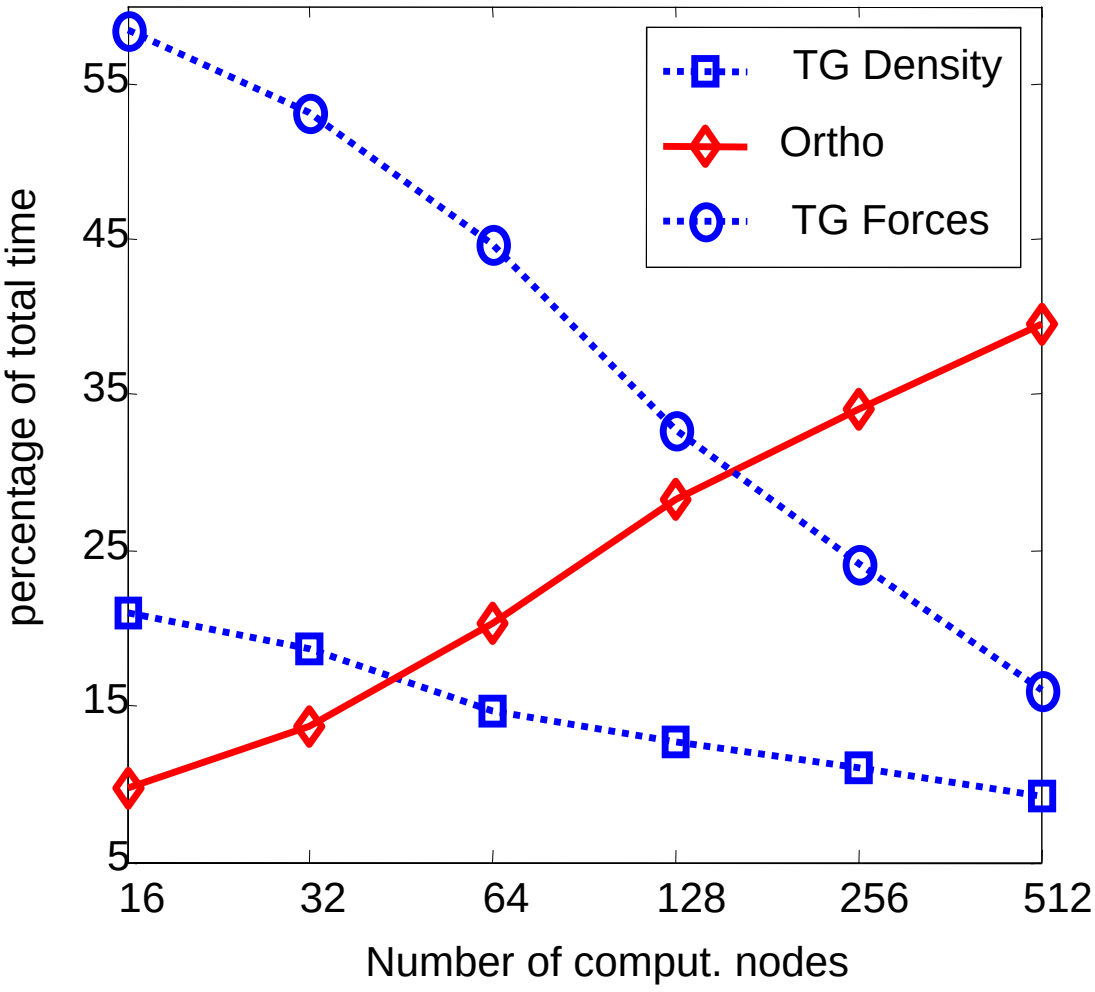

Comparison v.s. rival architectures

80 H_2O molecules
(A. Pasquarello)
• 240 atoms
• 320 occupied states (i.e. loop length for density)

Comparison vs. systems hosted at EPFL and new Cray XT3 at CSCS.

Employed TG strategy and DUAL core characteristics of the BG/L nodes:

- **DGEMM***
- **Double FPU directives / code per CPU code**

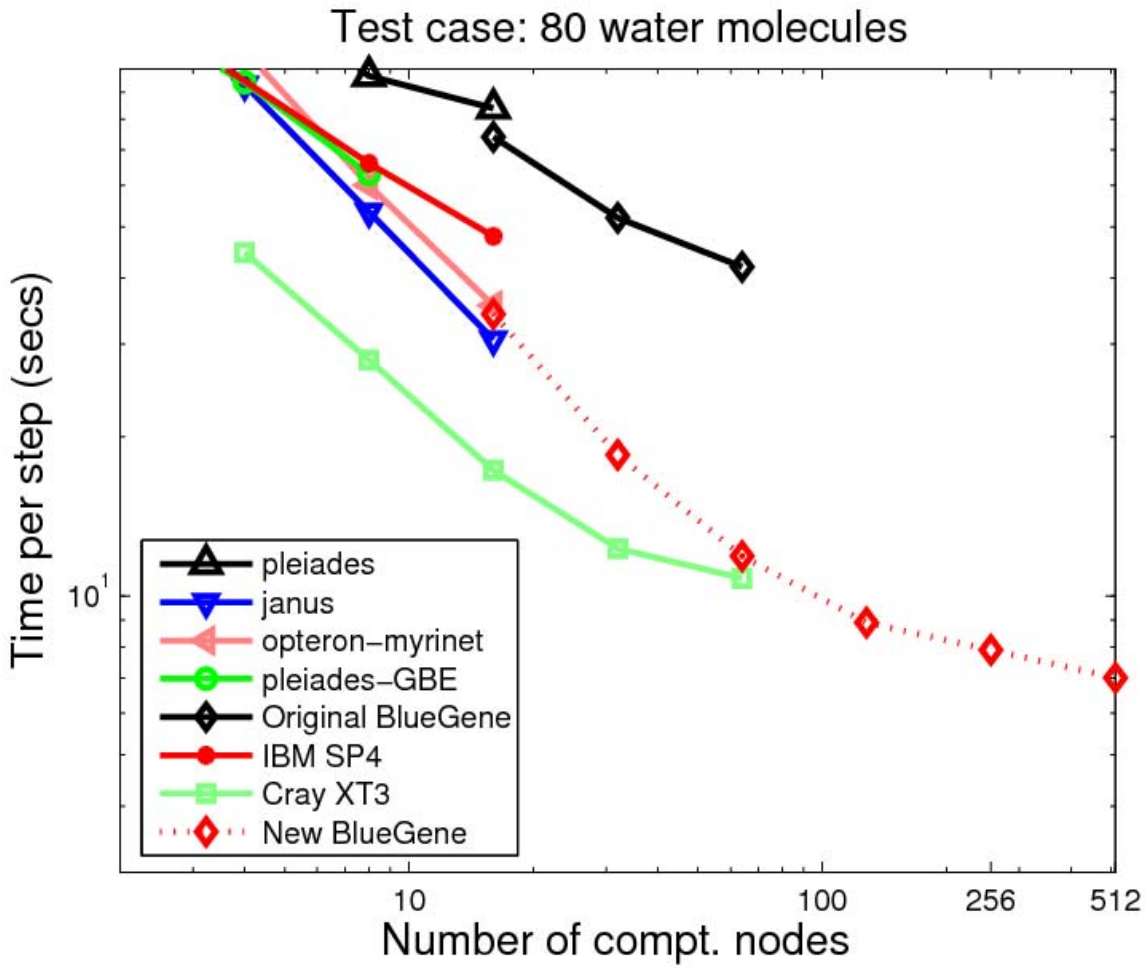

* J. Gunnels et al, Supercomputing, 2005 (G.B. prize)

Conclusions

- *ab initio* electronic structure calculations have much to offer in the coming years w.r.t. better understanding of physical processes at the nano scale
- Electronic structure codes are typical examples of large scale parallel computing applications
- The advent of the pioneering scalable BG/L system has triggered significant effort in scaling these codes to tens of thousands of processors. The largest TF count is achieved by such a code: CPMD
- FFT kernels are very important: Demonstrated scalability is here!
- Linear algebra kernels appear once more very important
- The synergy of new numerical algorithms as well as parallel computing practices will be crucial in crossing to the Peta Flop era