SIMULATING QUANTUM MATERIALS

Goals, challenges and perspectives in the era of pre-exascale computing

May 29, 2018 | Edoardo Di Napoli |

OUTLINE

Quantum Materials in a nutshell

Perspectives and trends

Challenges and goals

The SL Quantum Materials, its goals and challenges

OUTLINE

Quantum Materials in a nutshell

Perspectives and trends

Challenges and goals

The SL Quantum Materials, its goals and challenges

IN THE BEGINNING . . .

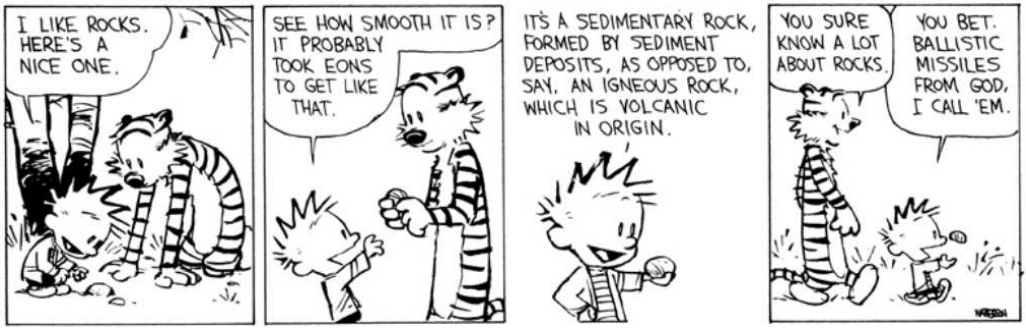

. . . THEN P.A.M. DIRAC

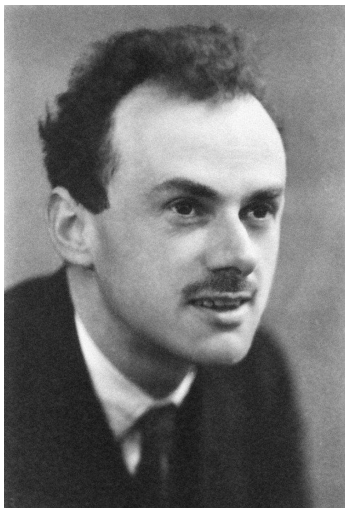

"The underlying **physical laws** *necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus* **completely known**, *and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that* **approximate practical methods** *of applying quantum mechanics should be developed, which can lead to an explanation of the main features of complex atomic systems* **without too much computation**".

THE SCHRÖDINGER EQUATION

general form and properties

$$\hat{H}\ \psi(\mathbf{r}) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) \right) \psi(\mathbf{r}) = E\ \psi(\mathbf{r})$$

THE SCHRÖDINGER EQUATION

general form and properties

$$\hat{H}\ \psi(\mathbf{r}) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) \right) \psi(\mathbf{r}) = E\ \psi(\mathbf{r})$$

$\hat{H} =$ Hamiltonian of a physical system:

- Hermitian operator
- In stationary cases it represents the (constant) energy of the system
- It is composed of a Kinetic ($-\frac{\hbar^2}{2m}\nabla^2$) and a Potential ($V(\mathbf{r})$) energy term.

THE SCHRÖDINGER EQUATION

general form and properties

$$\hat{H}\ \psi(\mathbf{r}) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) \right) \psi(\mathbf{r}) = E\ \psi(\mathbf{r})$$

$\hat{H} =$ Hamiltonian of a physical system:

- Hermitian operator
- In stationary cases it represents the (constant) energy of the system
- It is composed of a Kinetic ($-\frac{\hbar^2}{2m}\nabla^2$) and a Potential ($V(\mathbf{r})$) energy term.

$\psi(\mathbf{r}) =$ Wavefunction:

- one-dimensional complex valued function of many variables
- It is interpreted as a probability amplitude so that its square is a density distribution $\Rightarrow \int |\psi(\mathbf{r})|^2 d\mathbf{r} = 1$
- It generally depends on a set of discrete countable parameters

WHAT ABOUT JUST ELECTRONS?

Electronic Hamiltonian (neglecting spin and relativistic terms)

Atoms, molecules and solids are composed of many interacting electrons

$$\hat{H}_{\mathrm{el}} = \sum_i \frac{\mathbf{p}_i^2}{2m} - \sum_{i,I} \frac{z_I\, e^2}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{e^2}{|\mathbf{r}_i - \mathbf{r}_j|}$$

WHAT ABOUT JUST ELECTRONS?

Electronic Hamiltonian (neglecting spin and relativistic terms)

Atoms, molecules and solids are composed of many interacting electrons

$$\hat{H}_{\mathrm{el}} = \sum_i \frac{\mathbf{p}_i^2}{2m} - \sum_{i,I} \frac{z_I\, e^2}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{e^2}{|\mathbf{r}_i - \mathbf{r}_j|}$$

- 1st term represents the kinetic energy of electrons ($\mathbf{p} = i\hbar\nabla$)
- 2nd term describes the Coulomb potential well experienced by electrons due to the atoms nuclei
- 3rd describes repulsion (due to same charge) between electrons

Note: Dynamics of electrons and nuclei are somewhat separated if it was not for the presence of the Coulomb potential

INTRO TO MANY-ELECTRON PROBLEM

Exact solution

Electronic Schrödinger Equation

$$\hat{H}_{\text{el}}\ \Psi(\mathbf{r}_1;\sigma_1,\mathbf{r}_2;\sigma_2,\ldots,\mathbf{r}_N;\sigma_N) = E\,\Psi(\mathbf{r}_1;\sigma_1,\mathbf{r}_2;\sigma_2,\ldots,\mathbf{r}_N;\sigma_N)$$

$\Psi : \left(\mathbb{R}^3 \times \{\pm\frac{1}{2}\}\right)^N \longrightarrow \mathbb{C}$ high-dimensional anti-symmetric function.

INTRO TO MANY-ELECTRON PROBLEM

Exact solution

Electronic Schrödinger Equation

$$\hat{H}_{\text{el}}\ \Psi(\mathbf{r}_1;\sigma_1,\mathbf{r}_2;\sigma_2,\ldots,\mathbf{r}_N;\sigma_N) = E\Psi(\mathbf{r}_1;\sigma_1,\mathbf{r}_2;\sigma_2,\ldots,\mathbf{r}_N;\sigma_N)$$

$\Psi : \left(\mathbb{R}^3 \times \{\pm\frac{1}{2}\}\right)^N \longrightarrow \mathbb{C}$ high-dimensional anti-symmetric function.

The Exponential Wall

Assume $N = 10^3$ electrons (100 molecules of water)

Assume $q = 3$ bits per variable (extremely coarse grid)

$$\text{Time to solution} : \frac{2^N}{\text{CPUclock}} \sim \frac{2^{10^3}}{10^9} \sim 10^{290}\text{sec} \qquad \text{Age Universe } \sim 10^{23}\text{sec}$$

JÜLICH Forschungszentrum

INTRO TO MANY-ELECTRON PROBLEM

Exact solution

Electronic Schrödinger Equation

$$\hat{H}_{\text{el}}\ \Psi(\mathbf{r}_1;\sigma_1,\mathbf{r}_2;\sigma_2,\ldots,\mathbf{r}_N;\sigma_N) = E\Psi(\mathbf{r}_1;\sigma_1,\mathbf{r}_2;\sigma_2,\ldots,\mathbf{r}_N;\sigma_N)$$

$\Psi : \left(\mathbb{R}^3 \times \{\pm\frac{1}{2}\}\right)^N \longrightarrow \mathbb{C}$ high-dimensional anti-symmetric function.

The Exponential Wall

Assume $N = 10^3$ electrons (100 molecules of water)

Assume $q = 3$ bits per variable (extremely coarse grid)

$$\text{Time to solution} : \frac{2^N}{\text{CPUclock}} \sim \frac{2^{10^3}}{10^9} \sim 10^{290}\text{sec} \qquad \text{Age Universe } \sim 10^{23}\text{sec}$$

$$\text{Storing } \Psi : \text{ Bits} = q^{3N} \sim 10^{1500} \qquad \text{Baryons Universe} = 10^{80}$$

YOU MAY FEEL LIKE CALVIN . . .

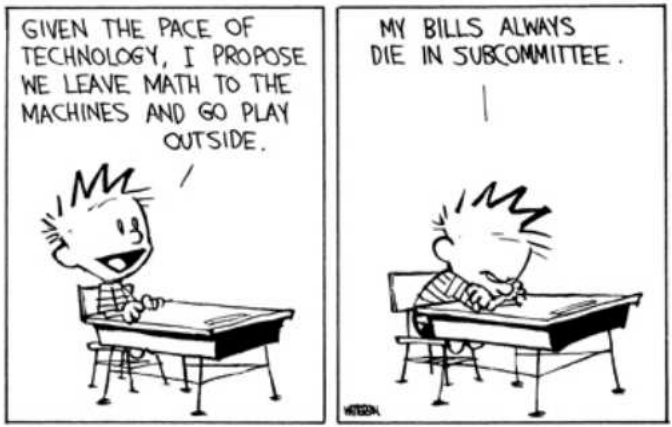

OUTLINE

Quantum Materials in a nutshell

Perspectives and trends

Challenges and goals

The SL Quantum Materials, its goals and challenges

APPROXIMATE SOLUTIONS

A plethora of methods

Definition

Ab initio is a Latin term meaning “from the beginning” and is derived from the Latin ab (“from”) + initio, ablative singular of initium (“beginning”).

Ab initio molecular orbitals methods:

- Hartree-Fock – # ops scales as $\sim N^4$
- Møller-Plesset
Configuration Interaction (CI)
Coupled Clusters (CC)
etc.
} – # ops scales as $N^4 \div N^7$

Ab initio electron density methods:

- Density Functional Theory (DFT) – # ops scales as $N \ln N \div N^3$
- Car-Parrinello Molecular Dynamics (MD)

TYPICAL APPLICATIONS

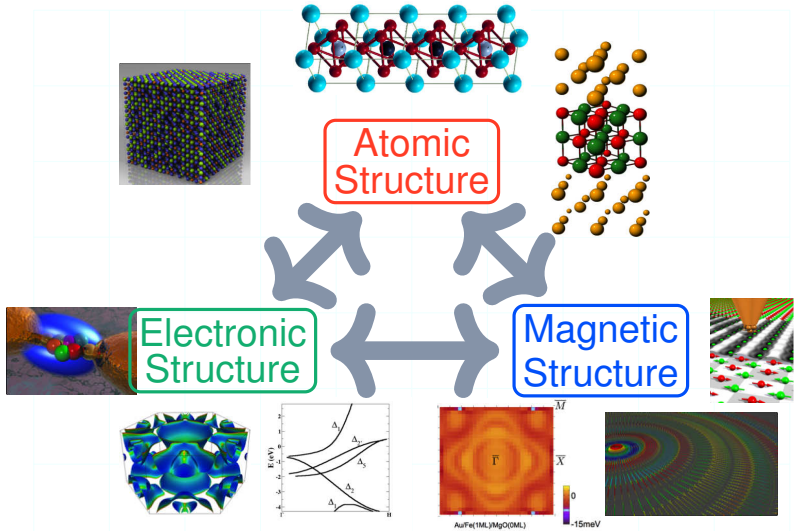

TRENDS IN ACADEMIA AND INDUSTRY

Private sector:

1. **Target (an example)**: the restructuring of the energy system by investigating materials for energy conversion processes and storage technologies.
2. **Method**: Using HTC for a large number of small scale simulations aimed at low impact materials screening
3. **Path**: Stable and verified simulation software (license preferred).

TRENDS IN ACADEMIA AND INDUSTRY

Private sector:

1. **Target (an example)**: the restructuring of the energy system by investigating materials for energy conversion processes and storage technologies.
2. **Method**: Using HTC for a large number of small scale simulations aimed at low impact materials screening
3. **Path**: Stable and verified simulation software (license preferred).

Public sector across topics and disciplines

1. **Target**: Fundamental research in method and functionality development
2. **Method**: Using large scale computations to simulate real world materials
3. **Path**: Code implementation without concern for abstractions and efficiency

TRENDS IN THE SUPERCOMPUTING INDUSTRY

Commodity Cluster

Blue Waters (NCSA)

- Cray: 22640 XE6 and 4224 XK7 nodes
- 64 GB memory (XE6) and 32 GB (XK7)
- XE6: 2 AMD Interlagos chips (8 cores each) per node
- XE7: 1 AMD Interlagos and 1 Kepler accelerator
- Peak performance: 13+ PFLOPs

Specialized Cluster

Sunway TaihuLight (Wuxi)

- 40,960 Chinese-designed SW26010 manycore 64-bit RISC processors (256 cores each)
- 10,649,600 CPU cores across the entire system
- Own operating system, Sunway RaiseOS 2.0.5
- Own customized implementation of OpenACC 2.0
- Peak performance: 125 PFLOPs

TRENDS IN THE SUPERCOMPUTING INDUSTRY

Modular Supercomputer

JURECA cluster

- 1872 compute nodes
- Two Intel Xeon E5-2680 v3 Haswell CPUs per node
- 75 compute nodes equipped with two NVIDIA K80 GPUs
- Peak performance: 2.2 PFLOPs

JURECA booster

- 1640 compute nodes
- one Intel Xeon Phi 7250-F Knights Landing CPUs per node
- ntel Omni-Path Architecture high-speed network with non-blocking fat tree
- Peak performance: 5 PFLOPs

OUTLINE

Quantum Materials in a nutshell

Perspectives and trends

Challenges and goals

The SL Quantum Materials, its goals and challenges

1ST CHALLENGE: SYSTEM SIZE

Simulating Ideal periodic crystals

- Simple cell
- Small number of atoms $\mathcal{O}(10)$
- Many $\mathbf{k}$-points

Real world materials have defects

- For these systems size matters
- Large number of atoms needed $\mathcal{O}(10^3)$
- Cubic (N^3) methods struggle

REAL WORLD MATERIALS HAVE DEFECTS

$SrTiO_3$ with a 7.6 degree bounday ($N = 2950$)

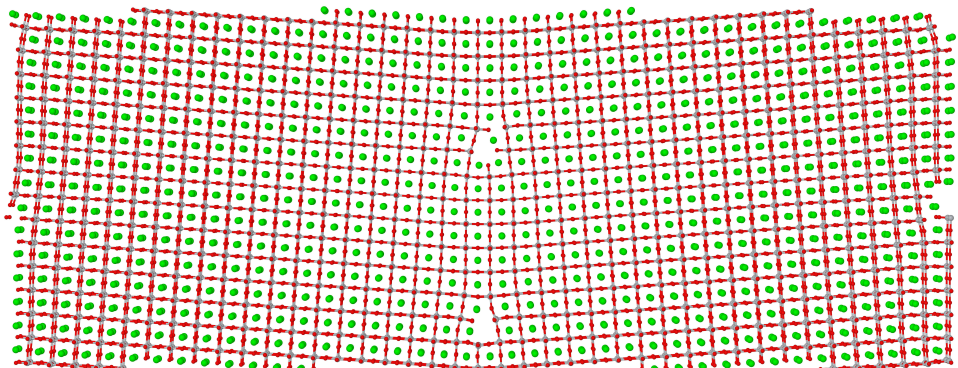

JÜLICH Forschungszentrum

LINEAR SCALABILITY

KKRnano

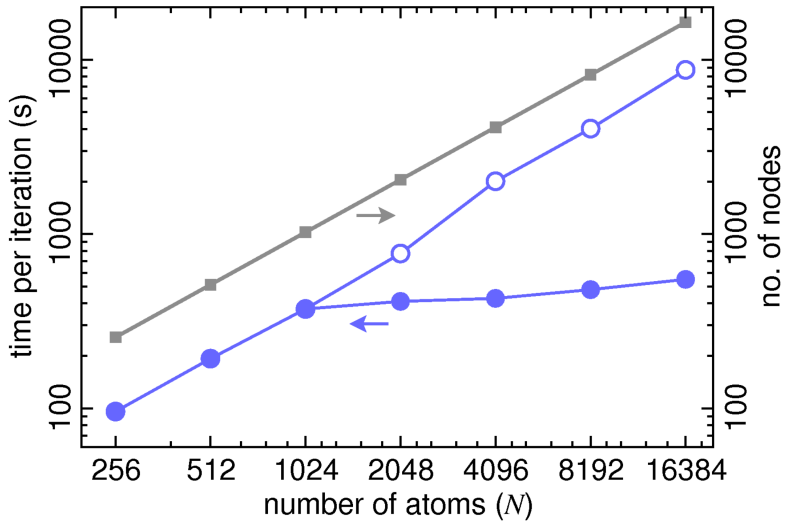

2ND CHALLENGE: CONVERGENCE

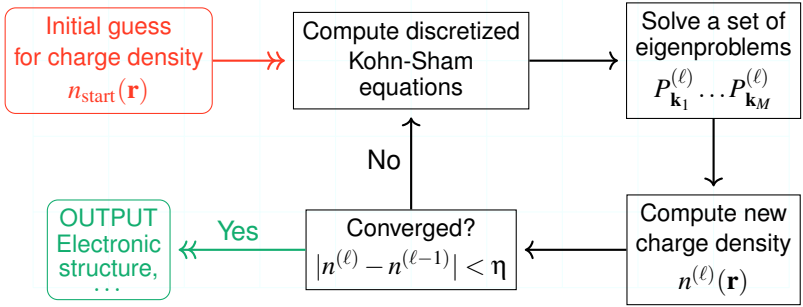

- ∃ problem of **"charge sloshing"**
- Often observed in metallic systems
- Dramatic slow-down of the SCF convergence
- Convergence worsen as size increases

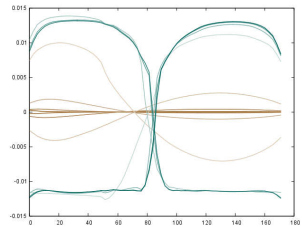

KERKER PRE-CONDITIONER: RESULTS FOR CU

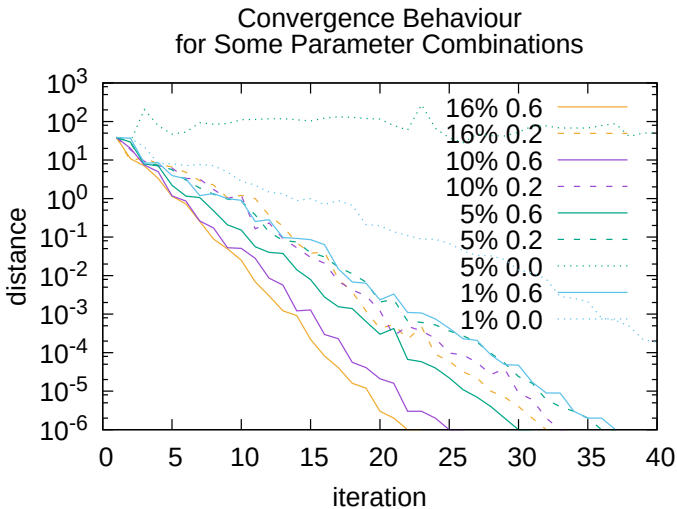

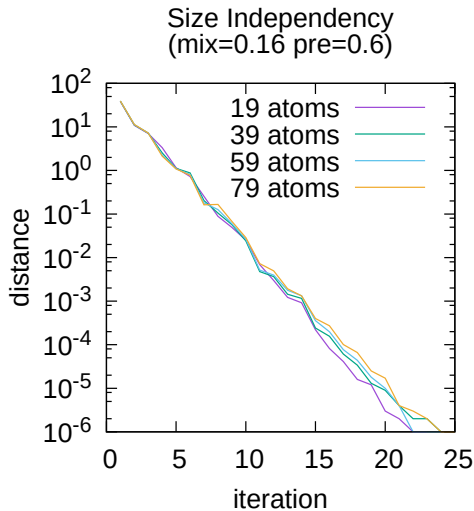

HIGH-PERFORMANCE COMPUTATIONS

Observations:

- Linear algebra algorithms ubiquitous in *Materials Science*
- Numerical libraries are used as **black boxes**.
- Domain-specific **knowledge** does not influence algorithm choice.
- Rigid legacy codes are hard to **modernize**.

Goal

Design and optimize linear algebra algorithms in order to:

- exploit available **knowledge**.
- increase the **parallelism** of complex tasks.
- facilitate **performance portability**

3RD CHALLENGE: HPC KNOWLEDGE TRANSFER

Legacy codes

- Conceived with a rigid abstraction model
- Developed by application scientists in isolation
- Carry the “curse of early optimization”
- Non-modular, often written in FORTRAN
- Many, many functionalities

Restructuring and modernization

- Back to mathematical model
- Decompose the numerical operations in terms of low-level kernels
- Separate data from algorithm
- Use architecture-aware kernels

HAMILTONIAN AND OVERLAP MATRIX GENERATION.

Matrix form

$$H = \sum_{a=1}^{N_A} \underbrace{A_a^H T_a^{[AA]} A_a}_{H_{AA}} + \underbrace{A_a^H T_a^{[AB]} B_a + B_a^H T_a^{[BA]} A_a + B_a^H T_a^{[BB]} B_a}_{H_{AB+BA+BB}}$$

$$S = \underbrace{\sum_{a=1}^{N_A} A_a^H A_a}_{S_{AA}} + \underbrace{\sum_{a=1}^{N_A} B_a^H \dot{U}_a^H \dot{U}_a B_a}_{S_{BB}}$$

A_a and $B_a \in \mathbb{C}^{N_L \times N_G}$ while $T_a^{\cdots} \in \mathbb{C}^{N_L \times N_L}$ and Hermitian.
$N_L = (l_{\max}+1)^2 \leq 121$, $N_G = 1,000 \div 50,000$, and $N_A =$ number of atoms.

PORTED TO HYBRID PLATFORMS (CPU+GPU)

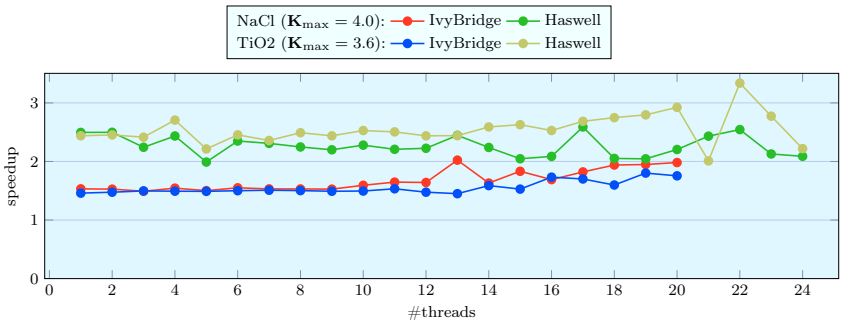

Figure 1: Speedup of our algorithm over FLEUR with $k_{max} = 4$ and increasing parallelism

PORTED TO HYBRID PLATFORMS (CPU+GPU)

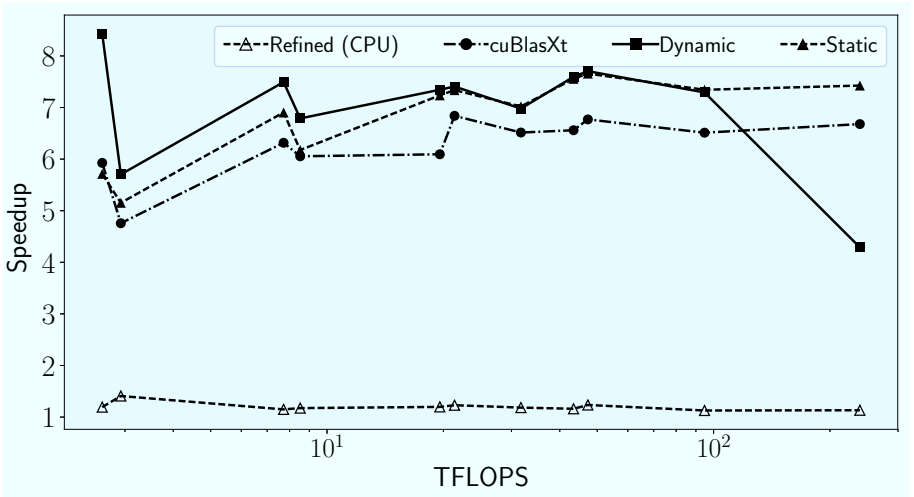

Figure: Speedup on RWTH-GPU for all implementations, relative to original code.

TAILORED NUMERICAL LIBRARY

ChASE library RELEASED.

- ChASE is open source (BSD 2.0 license) and available at
- `https://github.com/SimLabQuantumMaterials/ChASE`
- `https://arxiv.org/abs/1805.10121` (Submitted to ACM TOMS)

Highlights

- Sequences of correlated eigenvalue problems
- Modern C++ interface: easy-to-integrate, performance portability
- Excellent strong- and weak-scale performance

EMERGING METHODS

BSE and NEGF methods

- Optoelectronics and semiconductor technology
  - lasers and light emitting diods
  - Transparent electronics
- Energy related applications
  - Photocatalytic water splitting
  - Photovoltaic cells

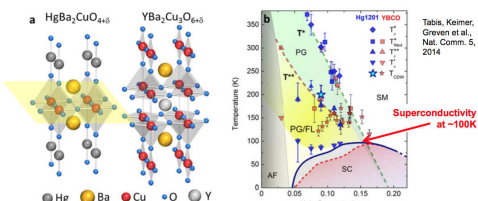

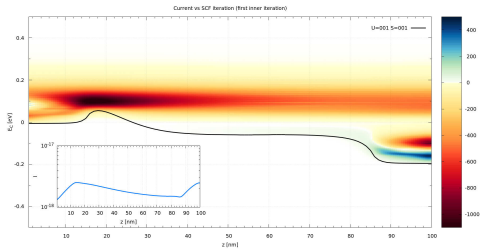

QMC and fRG methods

- Strong electron correlation
  - Energy scale
  - Topological effects
  - 2-D and 3-D phenomena
- Energy related applications
  - High-T_c superconductors
  - Anti-ferromagnetic effects

THE NEW FRONTIER: MACHINE LEARNING

Materials informatics

- HTC simulations
- Materials databases and data standardization
- Data-based properties prediction
- From simulations to production

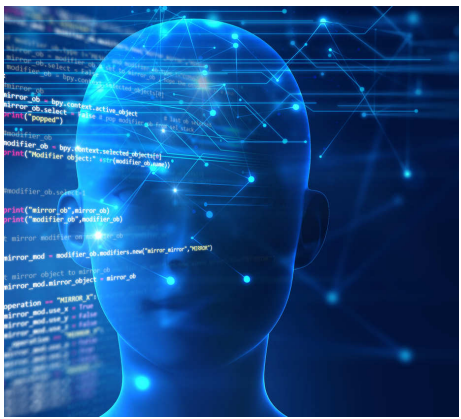

Machine Learning

- Materials classification
- Screening materials
- Inferring unknown materials properties
- Selecting DFT correlation potentials

OUTLINE

Quantum Materials in a nutshell

Perspectives and trends

Challenges and goals

The SL Quantum Materials, its goals and challenges

SIMULATION LABORATORY AS HPC ENABLER

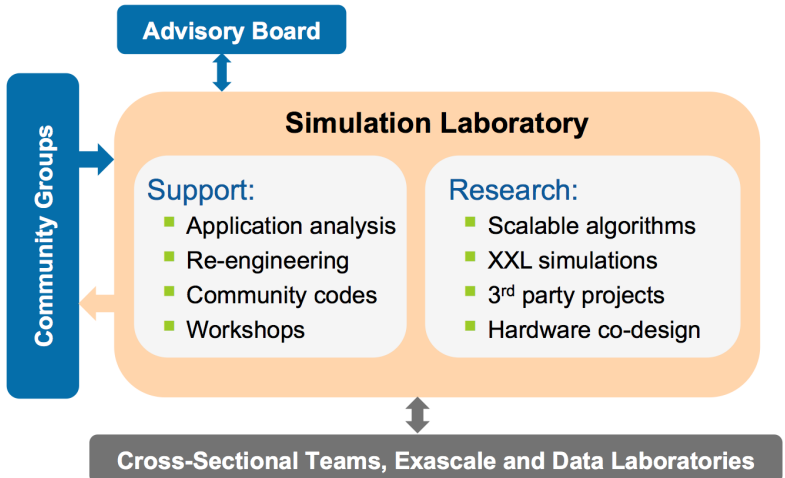

SIMLAB QUANTUM MATERIALS (SLQM)

SLQM area of activities

- Method development, algorithmic design, and code modernization
- HPC Knowledge transfer
- Programming models, Performance portability, Hybrid architectures, etc..

Mission

The Simulation Laboratory Quantum Materials (SLQM) provides expertise in the field of quantum-based simulations in Materials Science with a special focus on high-performance computing. SLQM acts as a high-level support structure in dedicated projects and hosts research projects dealing with fundamental aspects of code development, algorithmic optimization, and performance improvement. The Lab acts as an enabler of large scale simulations on current HPC platforms as well as on future architectures by targeting domain-specific co-design processes.

THE SLQM TEAM

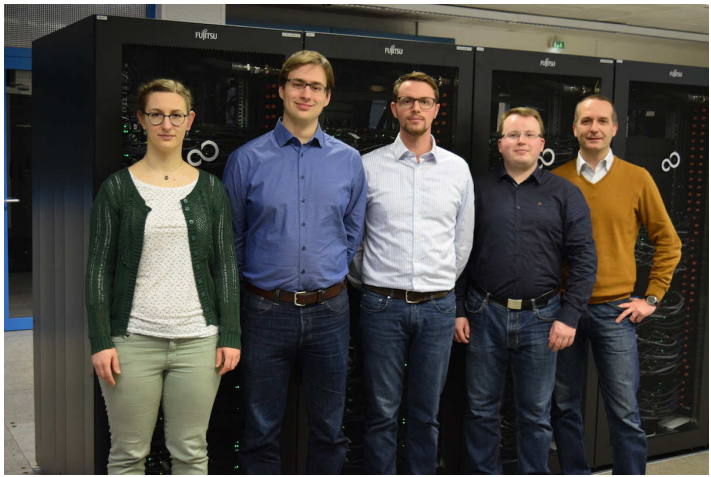

STRATEGIC PLAN

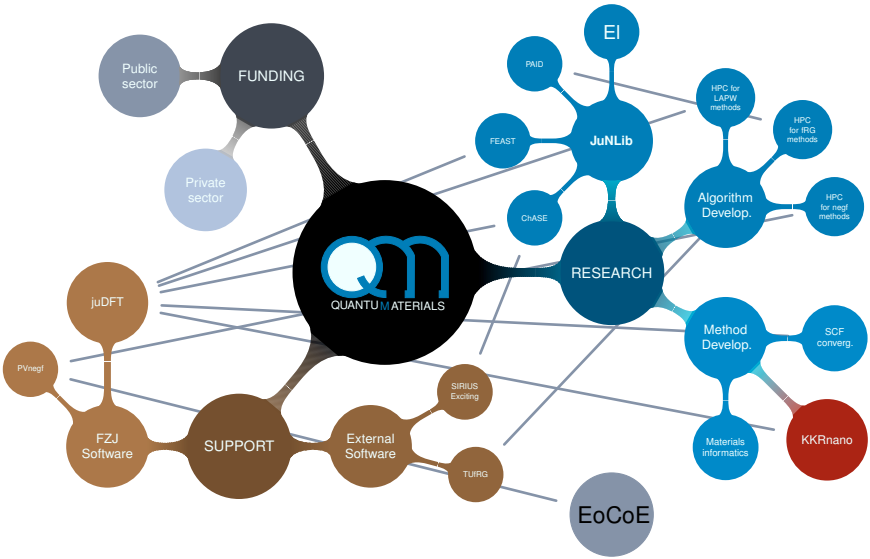

OUTLOOK

SLQM active areas of research

- **Adaptive integration** for functional Renormalization Group (fRG)
- HPC **Tensor** operations (contractions, generalized contractions, transpositions, etc.) with application in fRG, Quantum Chemistry, etc.
- Algebraic **eigensolvers** (sequences of dense eigenproblems, sparse eigenproblems) with applications in Density Functional Theory (DFT), Numerical RG, Excitonic Hamiltonians, etc.
- Articulated **linear algebra tasks** (Matrix initialization, etc.)
- Structured **Linear Systems** (Dyson, Poisson, and Keldysh equations)
- Self-consistent Field **preconditioners**
- Predicting material properties through **machine learning**

THANK YOU

For more information

e.di.napoli@fz-juelich.de

http://www.fz-juelich.de/ias/jsc/slqm