Linear-scaling ab initio study of surface defects in metal oxide and carbon nanostructures

Rubén Pérez

SPM Theory & Nanomechanics Group

Departamento de Física Teórica de la Materia Condensada &

Condensed Matter Physics Center (IFIMAC) @ UAM

http://www.uam.es/spmth

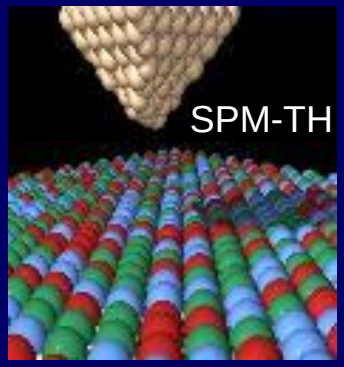

UAM
UNIVERSIDAD AUTONOMA
DE MADRID

IFIMAC

7th RES Users Meeting, Barcelona, 13/09/2013

Outline

1. SPMTH Group: ab initio Forces and currents on nanostructures

2. PRACE Tier-0 Project Access
   - What are Tier-0 machines and PRACE?
   - Why to apply for Tier-0 resources?
   - How?: Preparatory Access: code scalability
   - Problems during the testing stage
   - Project Access / Problems during the running stage

3. What are we doing with these resources?:
   - O(N) vs diagonalization for total energy calculations
   - CuO: 3D STM/AFM Maps: Cuantification of forces near defects only seen on the STM channel
   - Vacancies on Graphene: electronic, magnetic and mechanical properties

SPM-TH

Scanning Probe Microscopy Theory & Nanomechanics Group
Forces and Transport in Nanostructures

Home New

About us

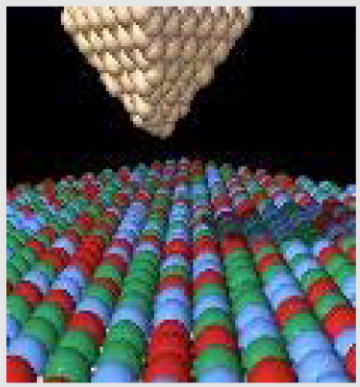

Our work is focused on ab initio r
Materials Science and Nanotech
currents at the atomic scale. Ir
and develop new capacities for th
the Scanning Probe Microsco
use currents and forces to visuali
nanoscale.

We work, in close collaborat
worldwide, in topics such as:

1. Scanning Tunneling (STM) and Atomic Force (AFM) energy dissipation mechanisms and nanomanipulation.
2. Nanomechanics: Fracture. Friction and Wear at the Transport properties of nanocontacts.
3. Adsorption, reactivity, and self-organization of fullerer on surfaces.

We use a suite of total-energy methods based in Density F very efficient codes based on local orbitals (like FIREBALL a

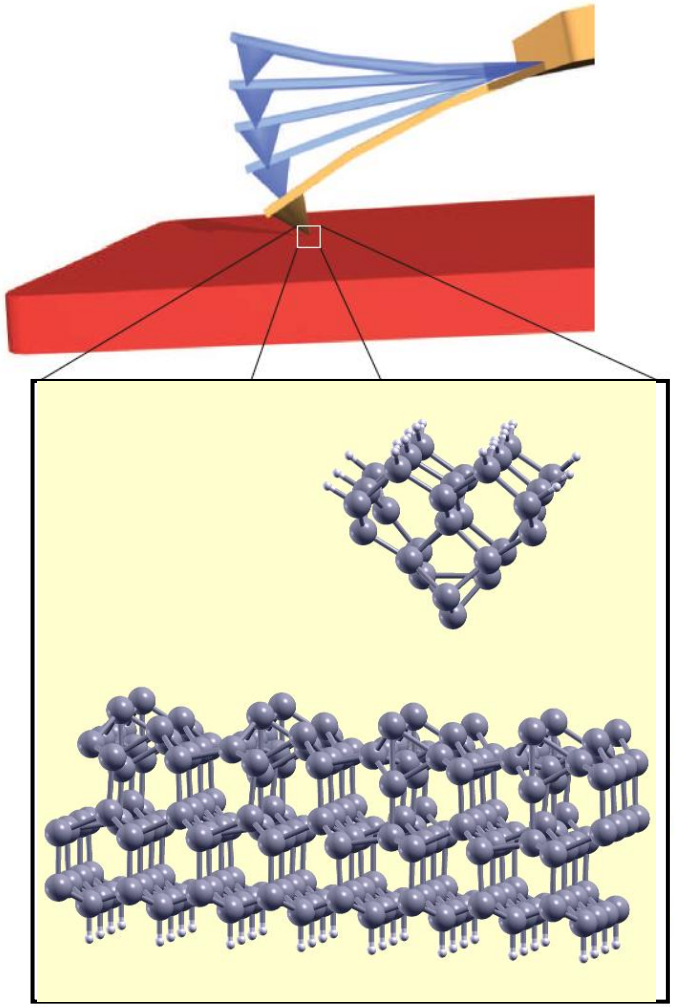

SR forces amenable to ab initio calculations

Methodology

Ab-initio total energy methods

(based in Density Functional Theory)

Non-equilibrium Green's Functions

both plane wave & local orbital basis: accuracy/efficiency balance

Linked with the local orbital description

Structure + electronic properties

Electronic transport

Energy vs tip height ⇒ Forces !!

OPENMX
VASP
+ vdW corrections:
semi-empirical (Grimme); vdW-DF (Klimes)

O(N) Method (Krylov subspace) !!

"The computer is a tool for clear thinking" Freeman J. Dyson

SPM-TH

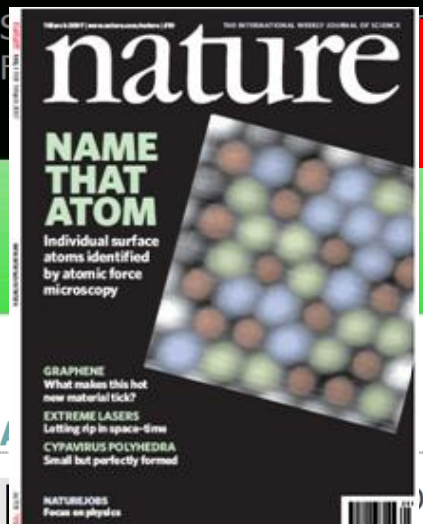

Chemical identification with FM-AFM

Nature 446, 64 (2007).

Single-atom manipulation with FM-AFM

Science 322, 413 (2008)

Nature Nanot. 4, 803 (2009)

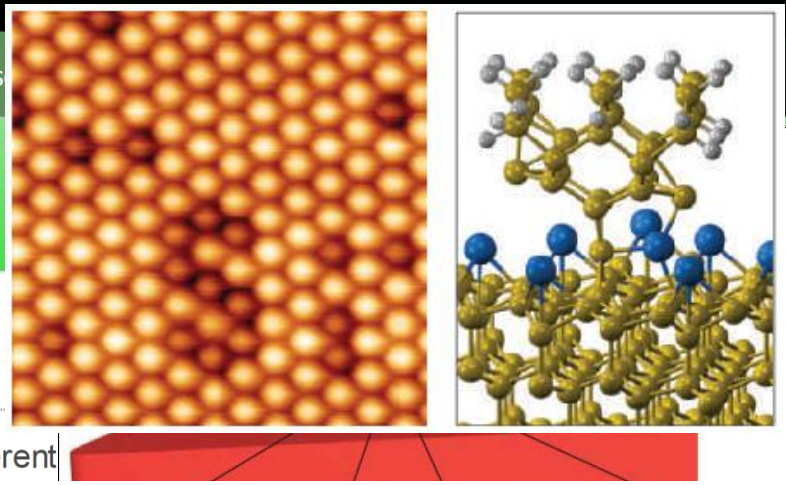

Selected for a Viewpoint in *Physics*

PRL **106,** 176101 (2011) PHYSICAL REVIEW LETTERS week ending 29 APRIL 2011

Forces and Currents in Carbon Nanostructures: Are We Imaging Atoms?

Martin Ondráček,[1] Pablo Pou,[2] Vít Rozsíval,[1] Cesar González,[3] Pavel Jelínek,[1] and Rubén Pérez[2,*]

Understanding Dissipative Tip–Molecule Interactions with Submolecular Resolution on an Organic Adsorbate

Gernot Langewisch, Wojciech Kamiński, Daniel-Alexander Braun, Rolf Möller, Harald Fuchs, André Schirmeisen,* and Rubén Pérez**

small NANO MICRO Vol. 8, 602 (2012)

We use a suite of total-energy methods based in Density Functional Theory (D very efficient codes based on local orbitals (like FIREBALL and OPENMX) to more accurate

Doctorate.

Outline

1. SPMTH Group: ab initio Forces and currents on nanostructures

2. PRACE Tier-0 Project Access
   - What are Tier-0 machines and PRACE?
   - Why to apply for Tier-0 resources?
   - How?: Preparatory Access: code scalability
   - Problems during the testing stage
   - Project Access / Problems during the running stage

3. What are we doing with these resources?:
   - O(N) vs diagonalization for total energy calculations
   - CuO: 3D STM/AFM Maps: Cuantification of forces near defects only seen on the STM channel
   - Vacancies on Graphene: electronic, magnetic and mechanical properties

What are Tier-0 machines and PRACE?

Tier-0: 50.000 – 460.000 cores, petaflop scale

System Name	Hosting Centre	Launch date	Architecture	Speed
CURIE	GENCI@CEA, France	Q2 2011	Bullx cluster	2 Petaflops/s
Hermit	GCS@HLRS, Germany	Q3 2011	Cray XE6	5 Petaflops/s
FERMI	CINECA, Italy	Q2 2012	BlueGene/Q	2 Petaflops/s
SuperMUC	GCS@LRZ, Germany	Q2 2012	iDataPlex	3 Petaflops/s
JUQUEEN	GCS@FZJ, Germany	Q4 2012	BlueGene/Q	5.87 Petaflops/s
MareNostrum	BSC, Spain	Q4 2012	iDataPlex	1 Petaflops/s

Partnership for Advanced Computing in Europe

PRACE (Partnership for Advanced Computing in Europe), the European research infrastructure for High Performance Computing (HPC), makes it possible for researchers from across Europe and the world to apply for access to Tier-0 HPC systems via a peer-review process.

Why Tier-0? (or better… Why bother?)

- ❑ It requires extra effort & different perspective from Tier-1 (typical job: from hundreds to thousands of processors)
- ❑ Many popular codes (e.g. VASP) do not have the required scalability
- ✓ Access to massive resources in terms of computer power and time

TGCC : PRACE Project 2012060986 on CURIE
PRACE Project : 2012060986
Project number : RA0986
16.250.000 core hours on standard nodes on Curie.
dates : 2012-11-01 to 2013-11-01

How?: Applying for a PRACE Tier-0 project

- Technical review
  - Scalability
- Scientific review

Project access vs Preparatory access
(different calls!!)

How?: Preparatory access projects

GOAL: to show the scalability of your code

1.1. Proposal ID
2010PA0638 **(Type A PA project intended for code scalability testing)**

1.3. Period of access to the PRACE facilities
01/11/2011 to 31/12/2011 **(just two months !!)**

1.4. Name of the PRACE facility assigned
Curie **(why?: large memory on Curie fat nodes)**

2.1. Project name
Large-scale O(N) DFT simulations of defects in metal oxides

2.2. Research field

Astrophysics
Earth Sciences and Environment
Medicine and Life Sciences
Chemistry and Materials
Fundamental Physics
Engineering and Energy
Mathematics and Computer Science
Finance and Social Science
Linguistics and Encryption

OpenMX scalability with hybrid MPI/OpenMP

of MPI process X # of threads (1-32) = # of cores < 4096 (limitation Prep. Acc.)

Table 1: 3400 atoms, 8 thread test

# cores	absolute timing (s)	speed-up	ideal speed-up	efficiency [%]
256	53.0	1.0000	1.0	100.0
512	31.3	1.6924	2.0	84.6
1024	15.7	3.3760	4.0	84.4
2048	9.4	5.6572	8.0	70.7
4096	7.3	7.3057	16.0	45.7

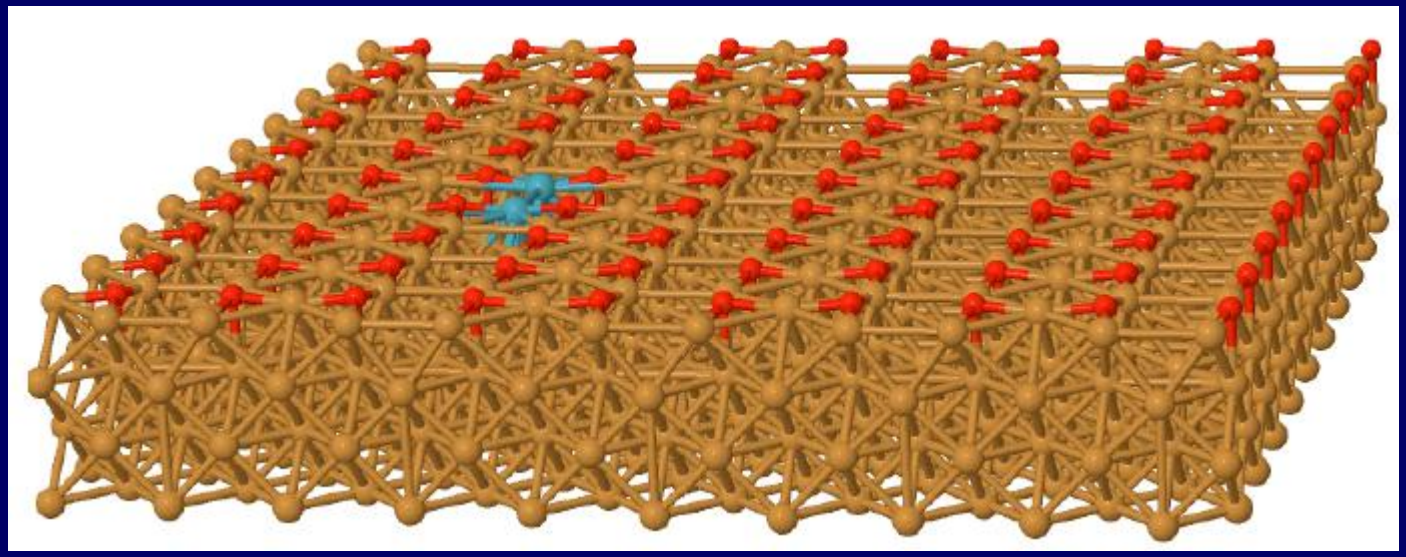

X 4

OpenMX scalability with hybrid MPI/OpenMP

Normalized Speed-up

Parallel Efficiency

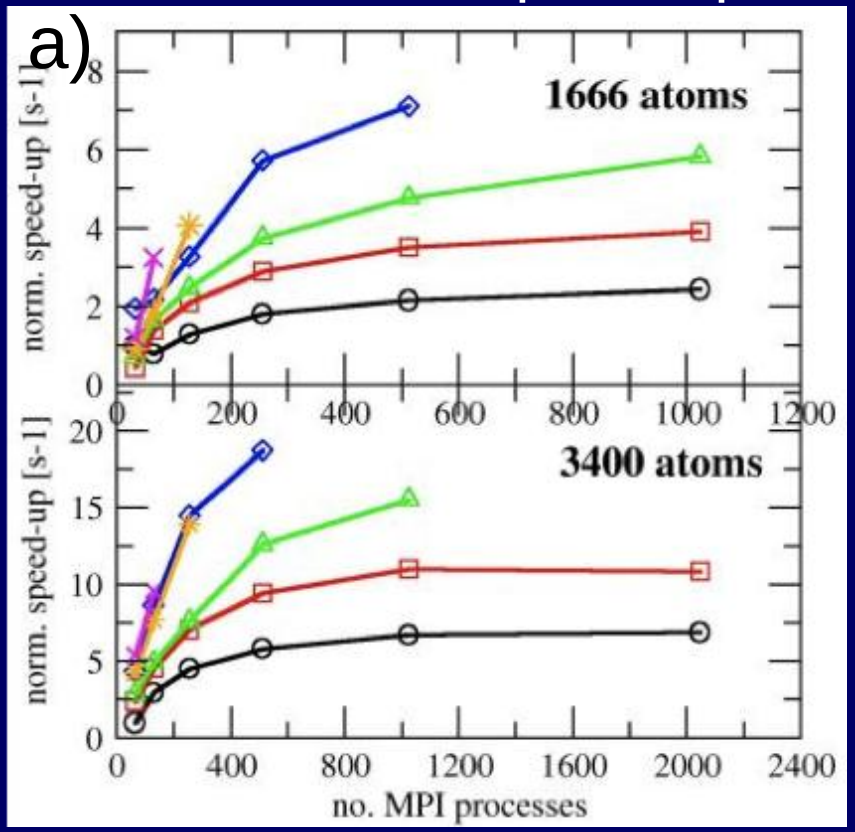

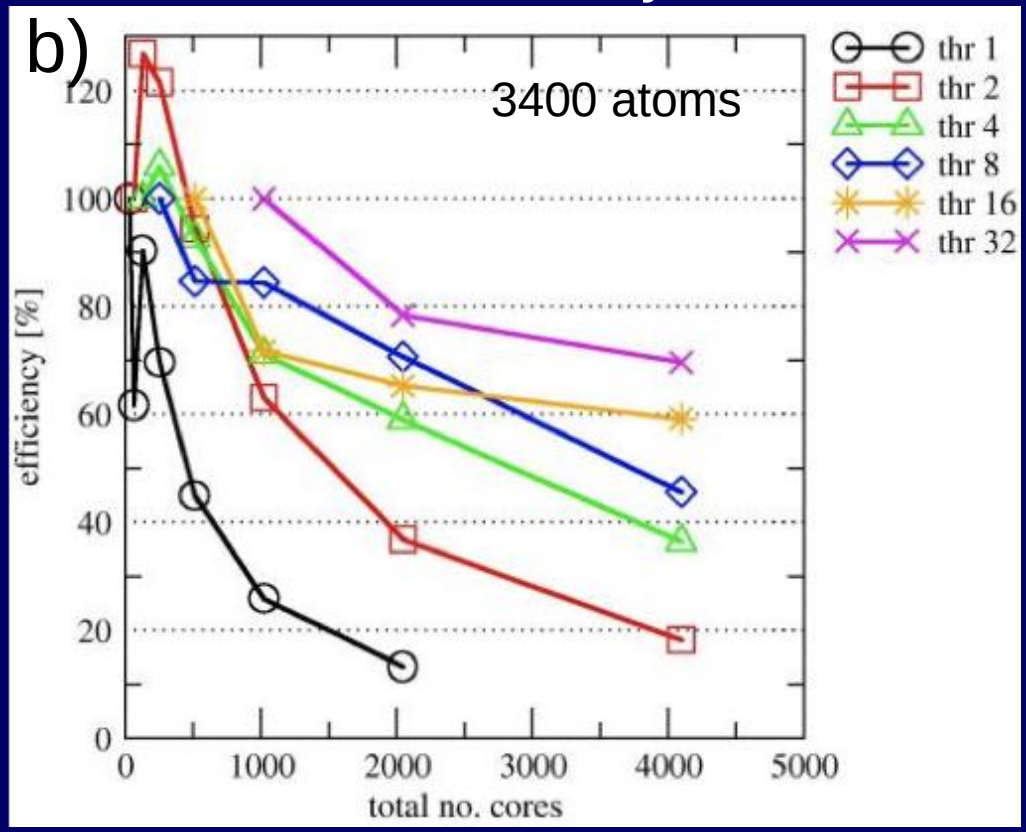

a) OpenMX speed-up performance with multithreading for large test calculation sizes (normalized to the first data point)

b) corresponding parallel efficiency computed for the 3400 atom test input.

OpenMX scalability with hybrid MPI/OpenMP

Resources vs Threading

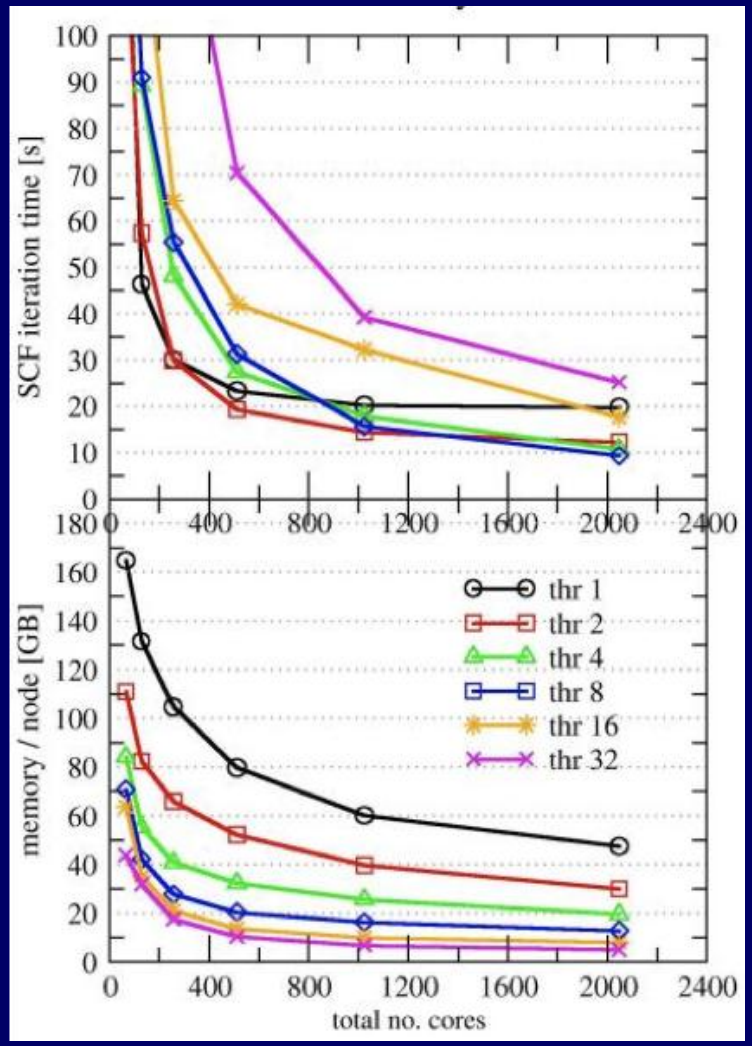

Variation in computational resources with multithreading for a constant number of cores (3400 atom calculation)

Linear scaling

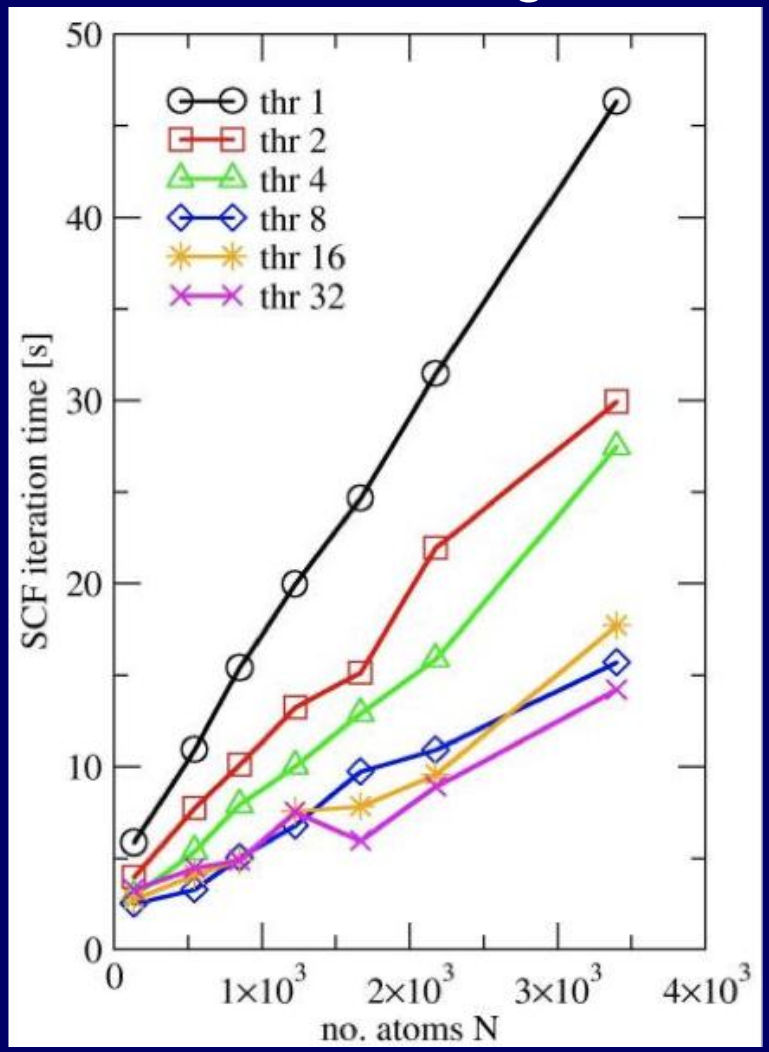

Linear-scaling of computational time with test system size (128 MPI procs with multithreading)

Problems in the testing stage

- Make sure that you test the same architecture you want to apply for: tests on the CURIE fat nodes (larger memory) BUT most of the PRACE time allocated on CURIE thin nodes)

  FIX: Repeat selected tests at the beginning of the project.

- Limited resources: < 4096 cores + small time allocation (while you are really intended to go for larger production jobs…)

  FIX: We got access to a larger machine (Cray XT5 @ JAIST, Japan)

+ (validation of scripts for scalability tests in Curie)

Applying for a PRACE Tier-0 project

- Technical review
  - Scalability
  - minimum request: 5 million hours
  - minimum job size > 512 cores (system dependent)
- Scientific review
  - Scientific excellence, novelty and impact
  - Focus on topics of major relevance for European research
  - Suggestion of three qualified reviewers.

Problems in the running stage (I)

- Beware of changing the parameters of the calculations (errors get amplified… and you are blacklisted !!!)
- Do not trust your internal functions to estimate the memory usage; use tools like Valgrind

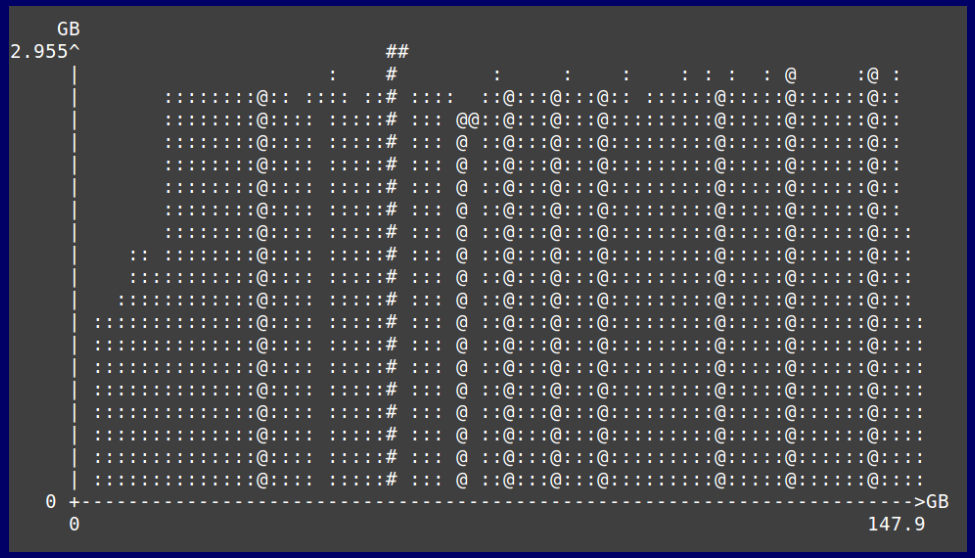

Memory consumption profile from Valgrind code for a test OpenMX calculation

(288 atoms, 13 orbitals, 2x2x1 grid for SCF & 10x10x1 DOS)

Mem. peak predicted by OpenMX: **1.542 Gb**

Mem. peak given by *ps ux* tracking: **2.932 Gb**

Mem. peak (real) given by Valgrind: **2.955 Gb**

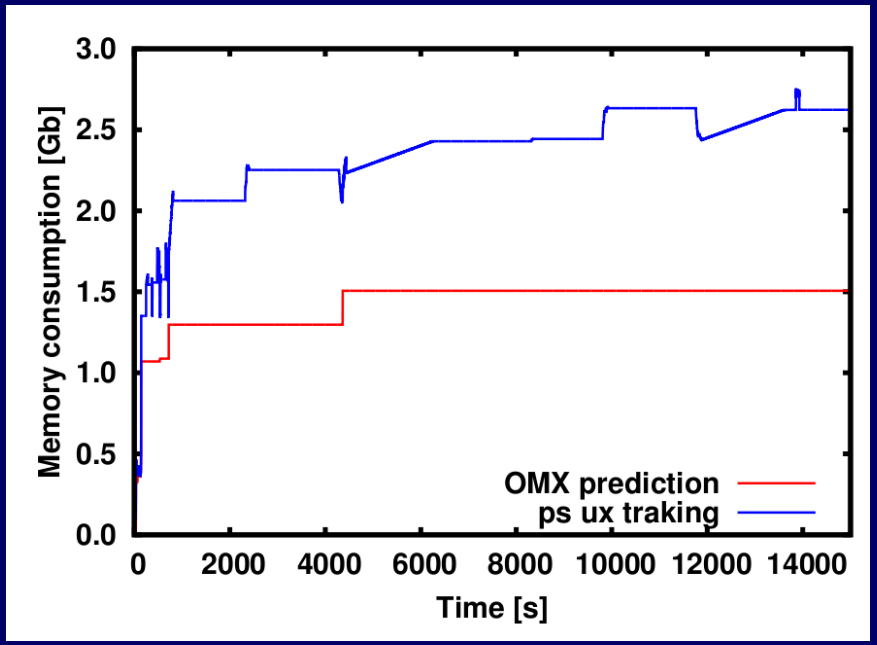

Problems in the running stage (II)

- Significant improvement in the performance of your code (3x for O(N), > 10x for diagonalization): you have asked for too much time…

v3.644 -O1 (lapack) vs v3.7 -O3 (elpa1 diagonalisation library)

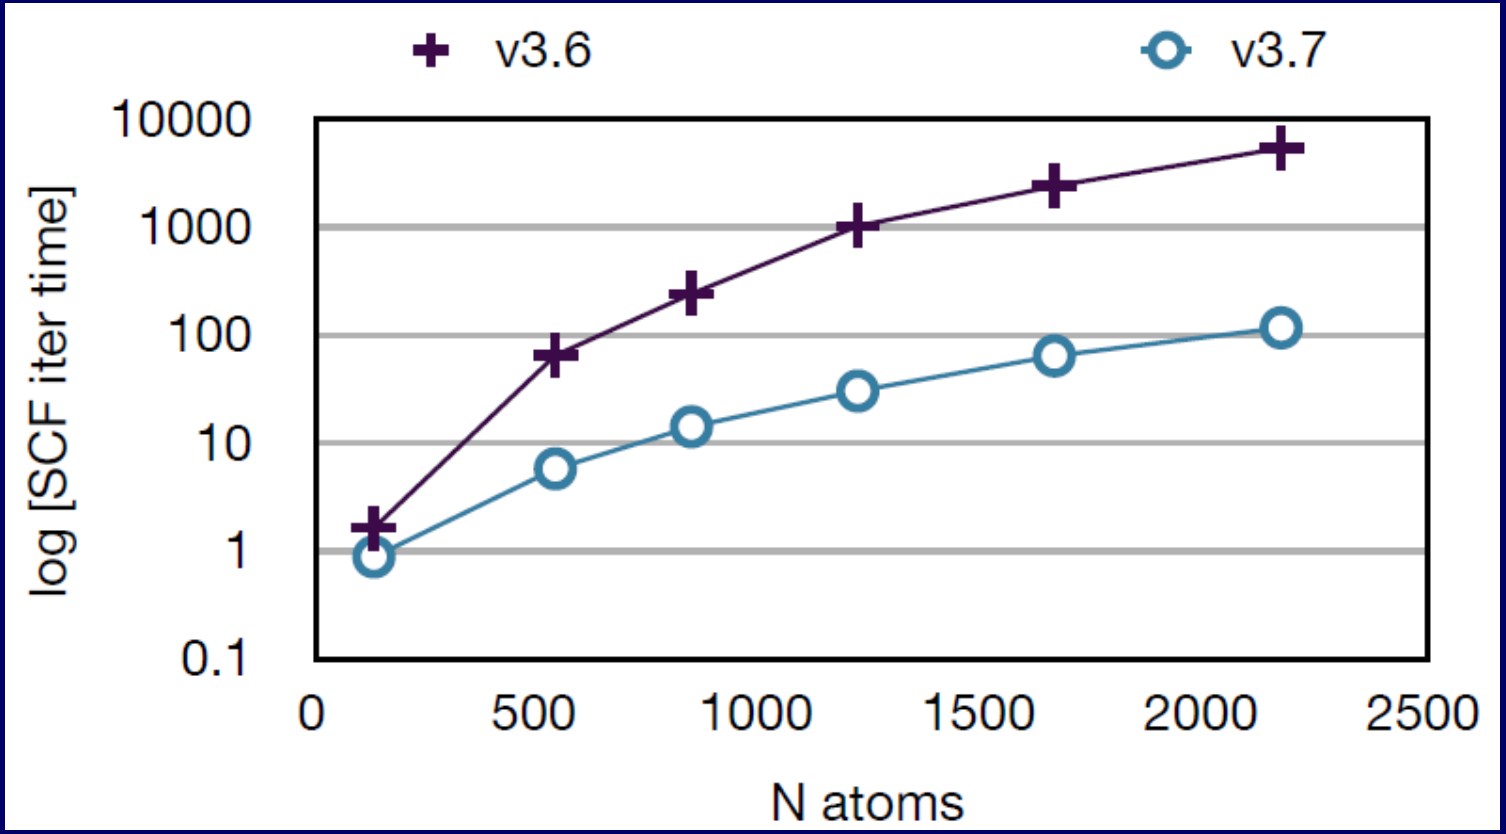

Outline

1. SPMTH Group: ab initio Forces and currents on nanostructures

2. PRACE Tier-0 Project Access
   - What are Tier-0 machines and PRACE?
   - Why to apply for Tier-0 resources?
   - How?: Preparatory Access: code scalability
   - Problems during the testing stage
   - Project Access / Problems during the running stage

3. What are we doing with these resources?:
   - O(N) vs diagonalization for total energy calculations
   - CuO: 3D STM/AFM Maps: Cuantification of forces near defects only seen on the STM channel
   - Vacancies on Graphene: electronic, magnetic and mechanical properties

Cu(100)-0: An oxide ML on top of a metal

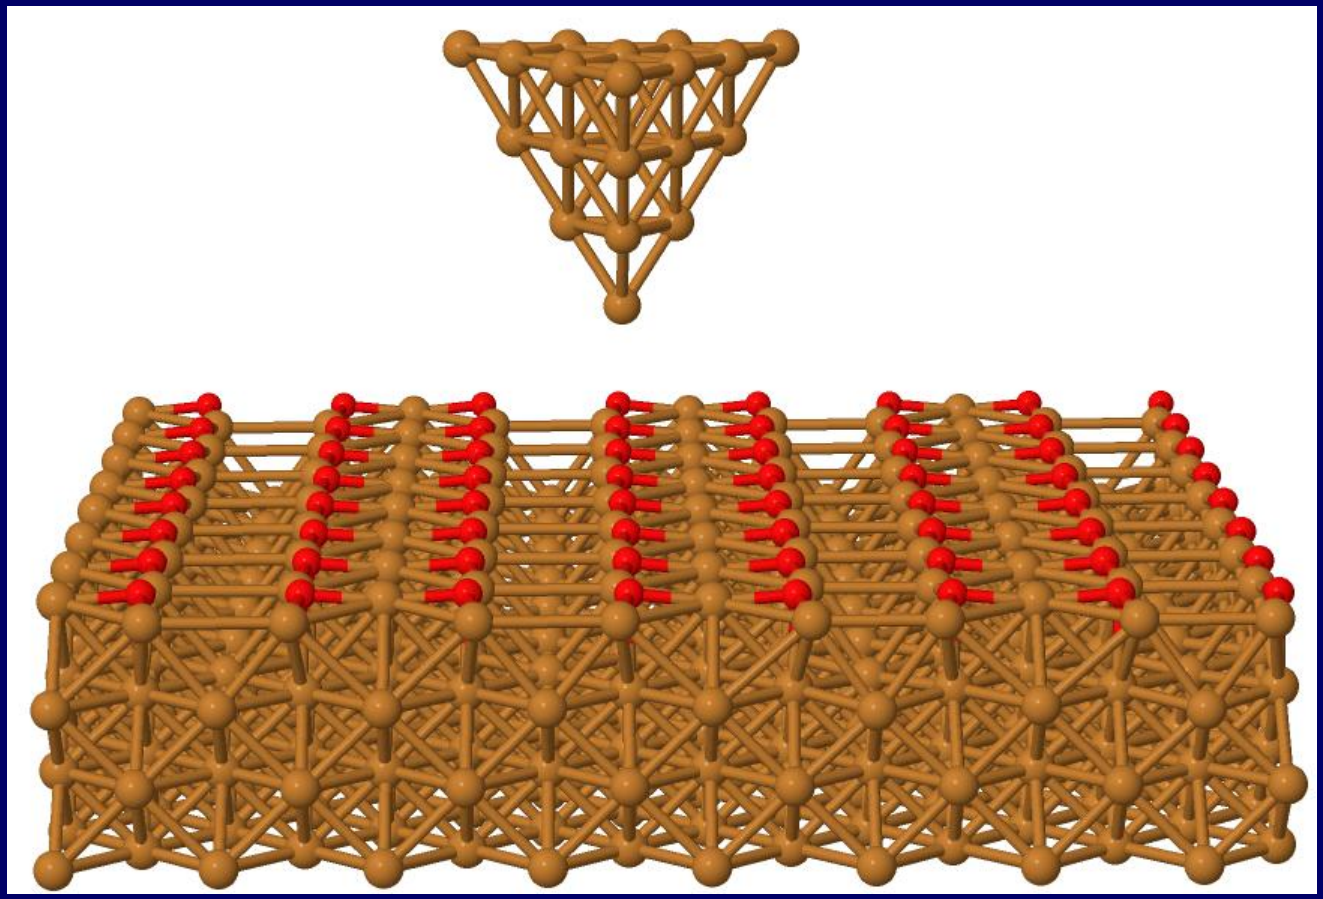

Cu tip on Cu(100)-0: O(N) vs diagonalization

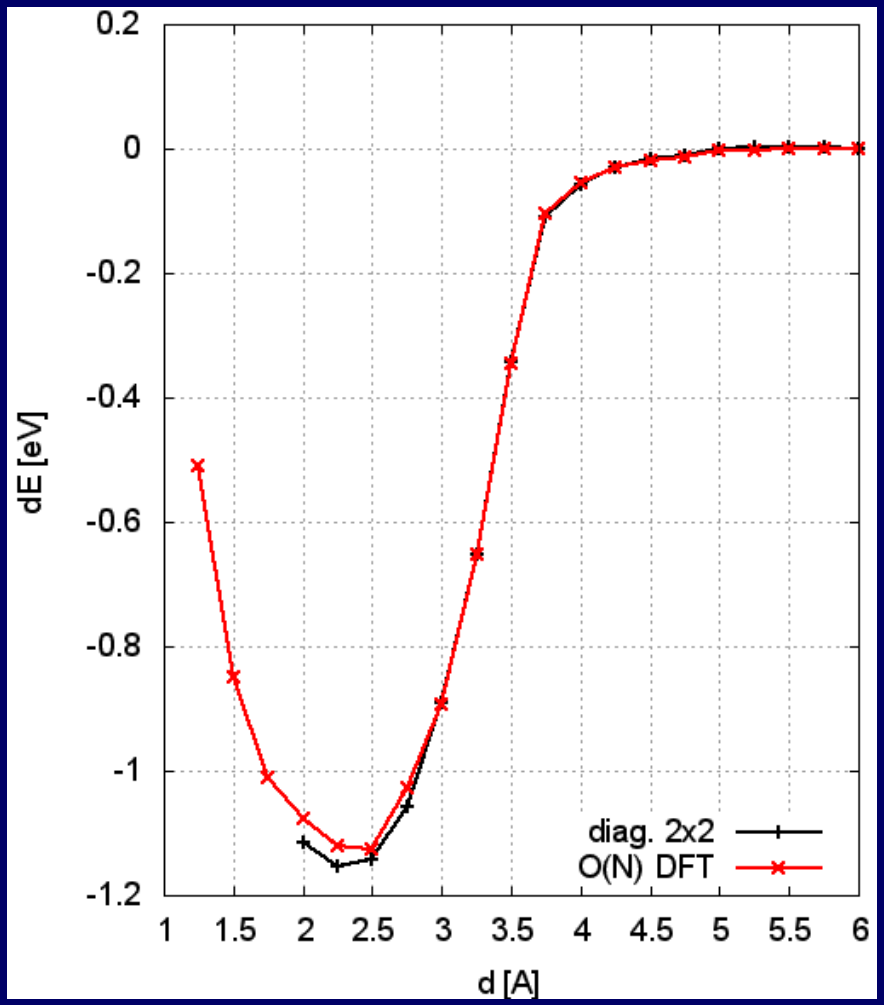

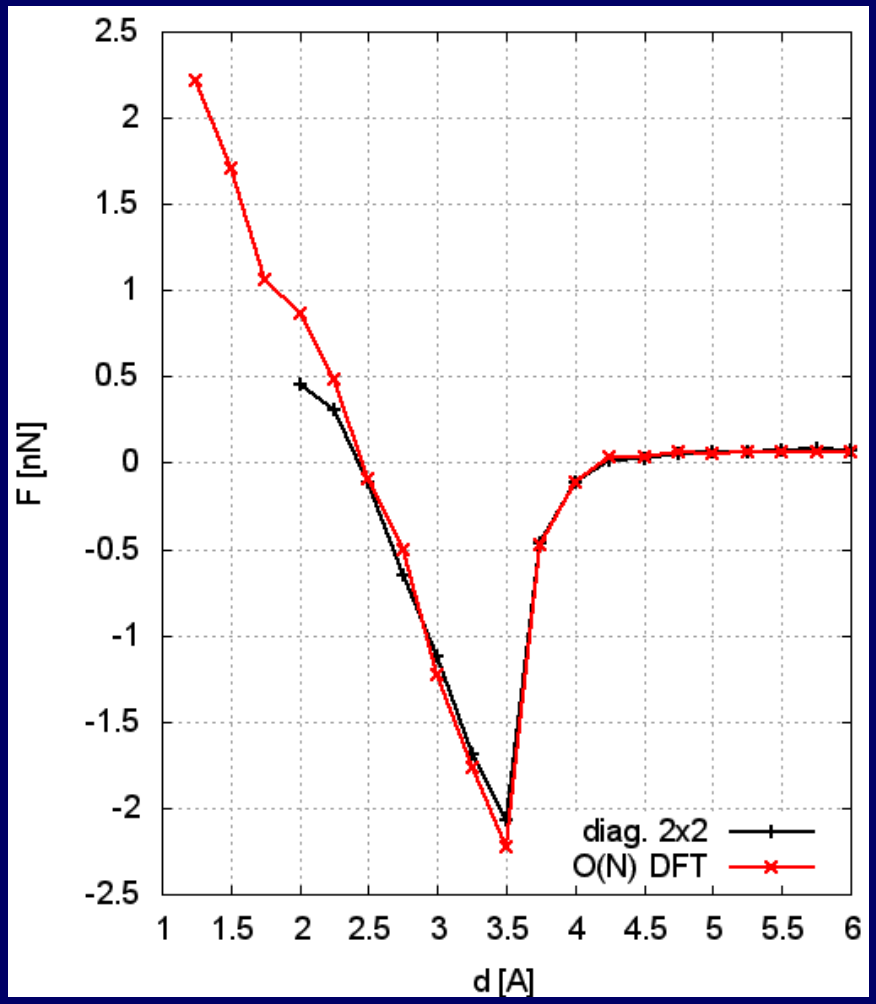

Large Diagonalization calcs possible with the elpa1 library !!!

Cu(100)-0: Force contrast around the defect

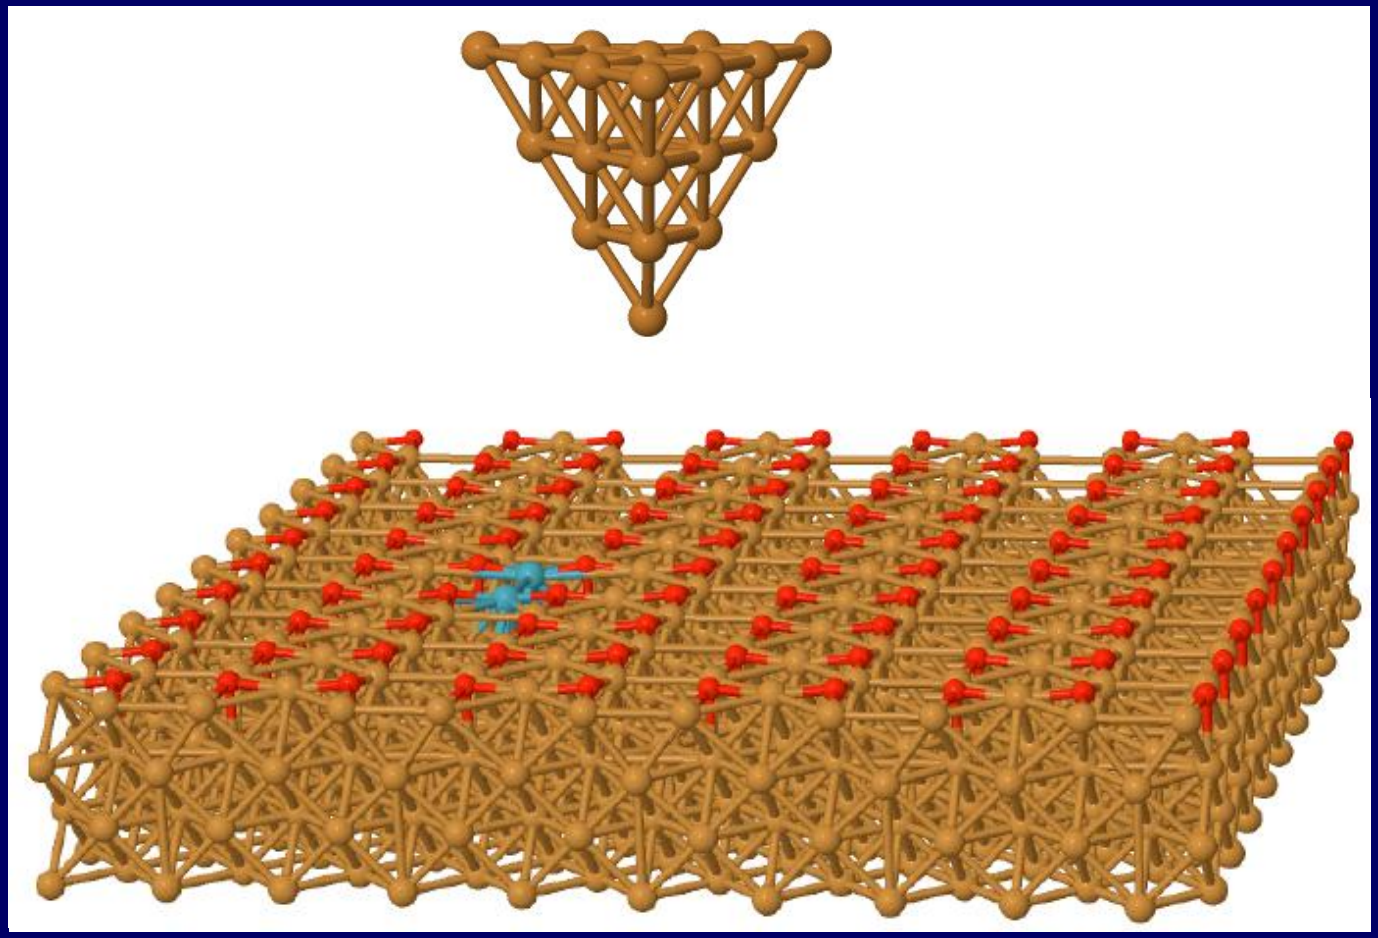

Vacancies in Graphene: Magnetism??

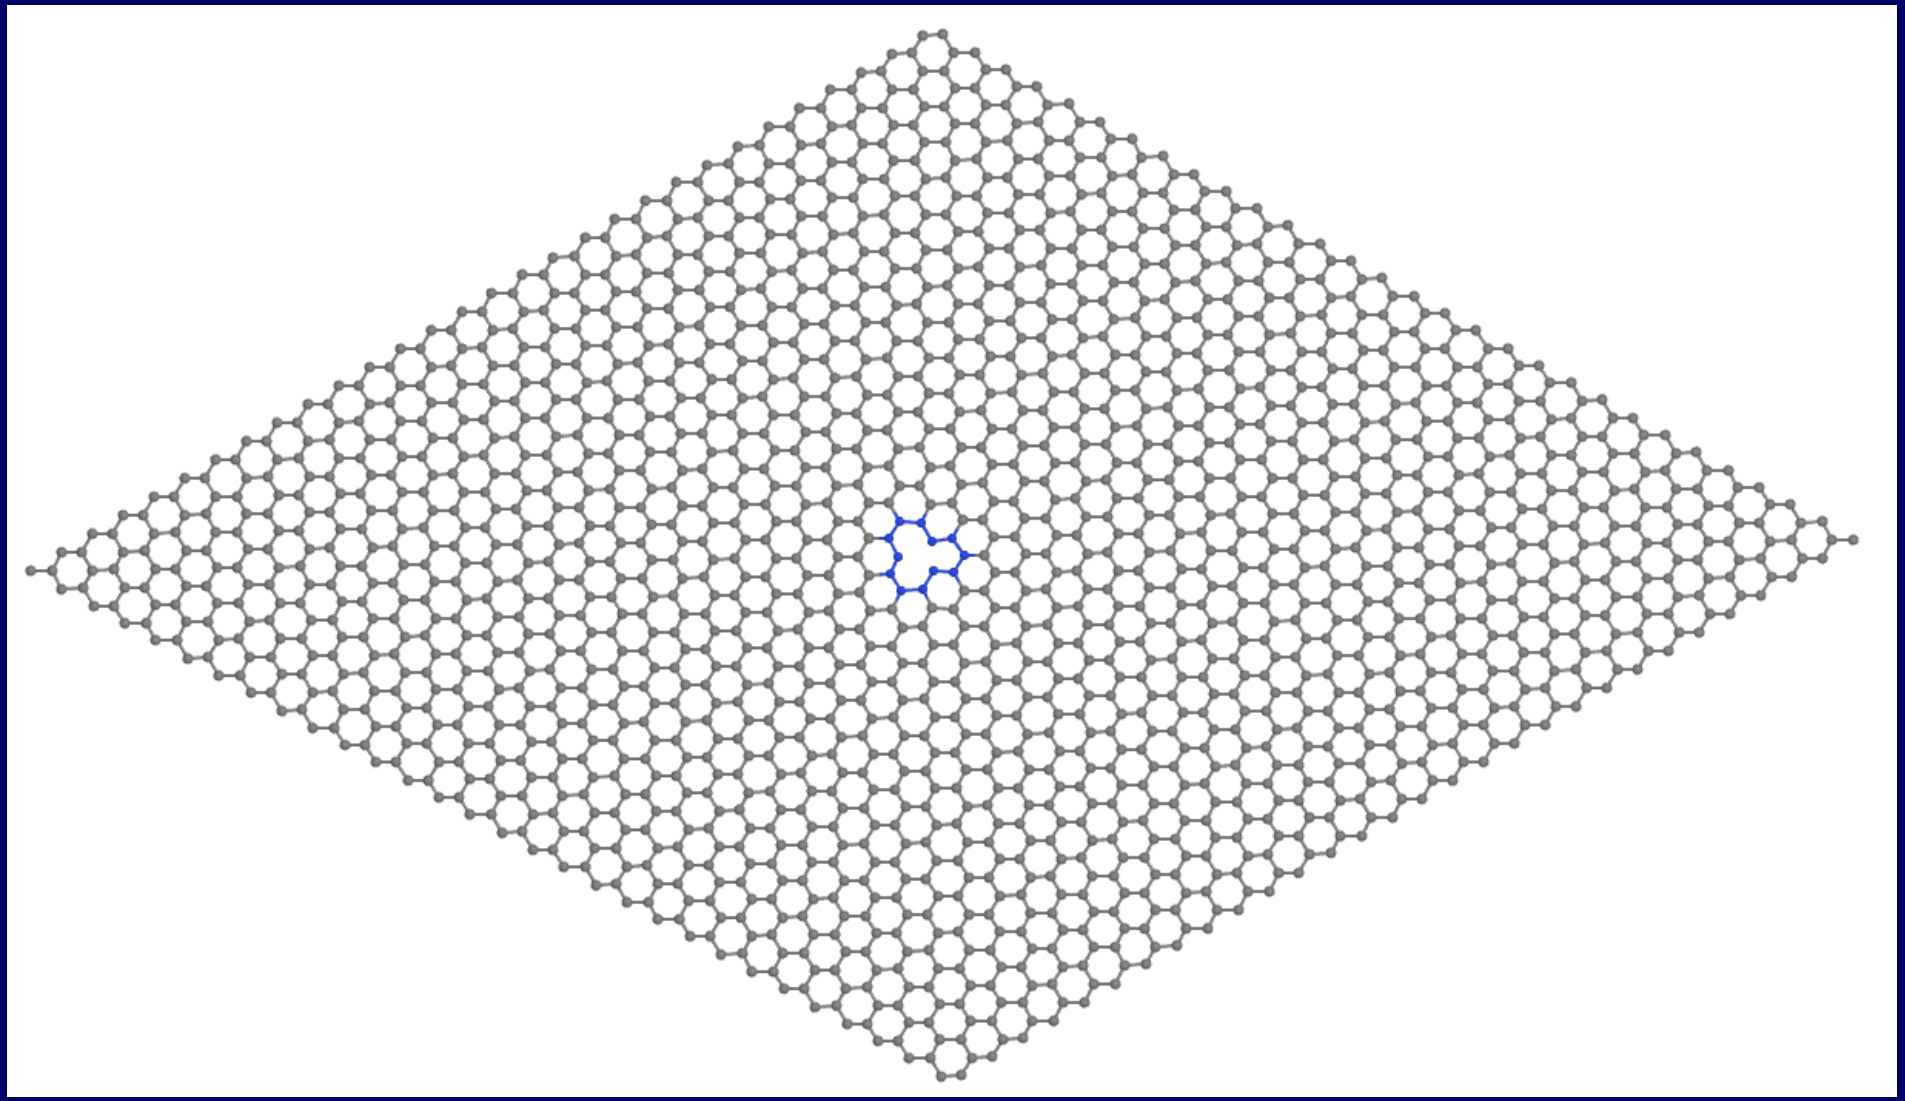

Why do we need these large unit cells and fine k-sampling?

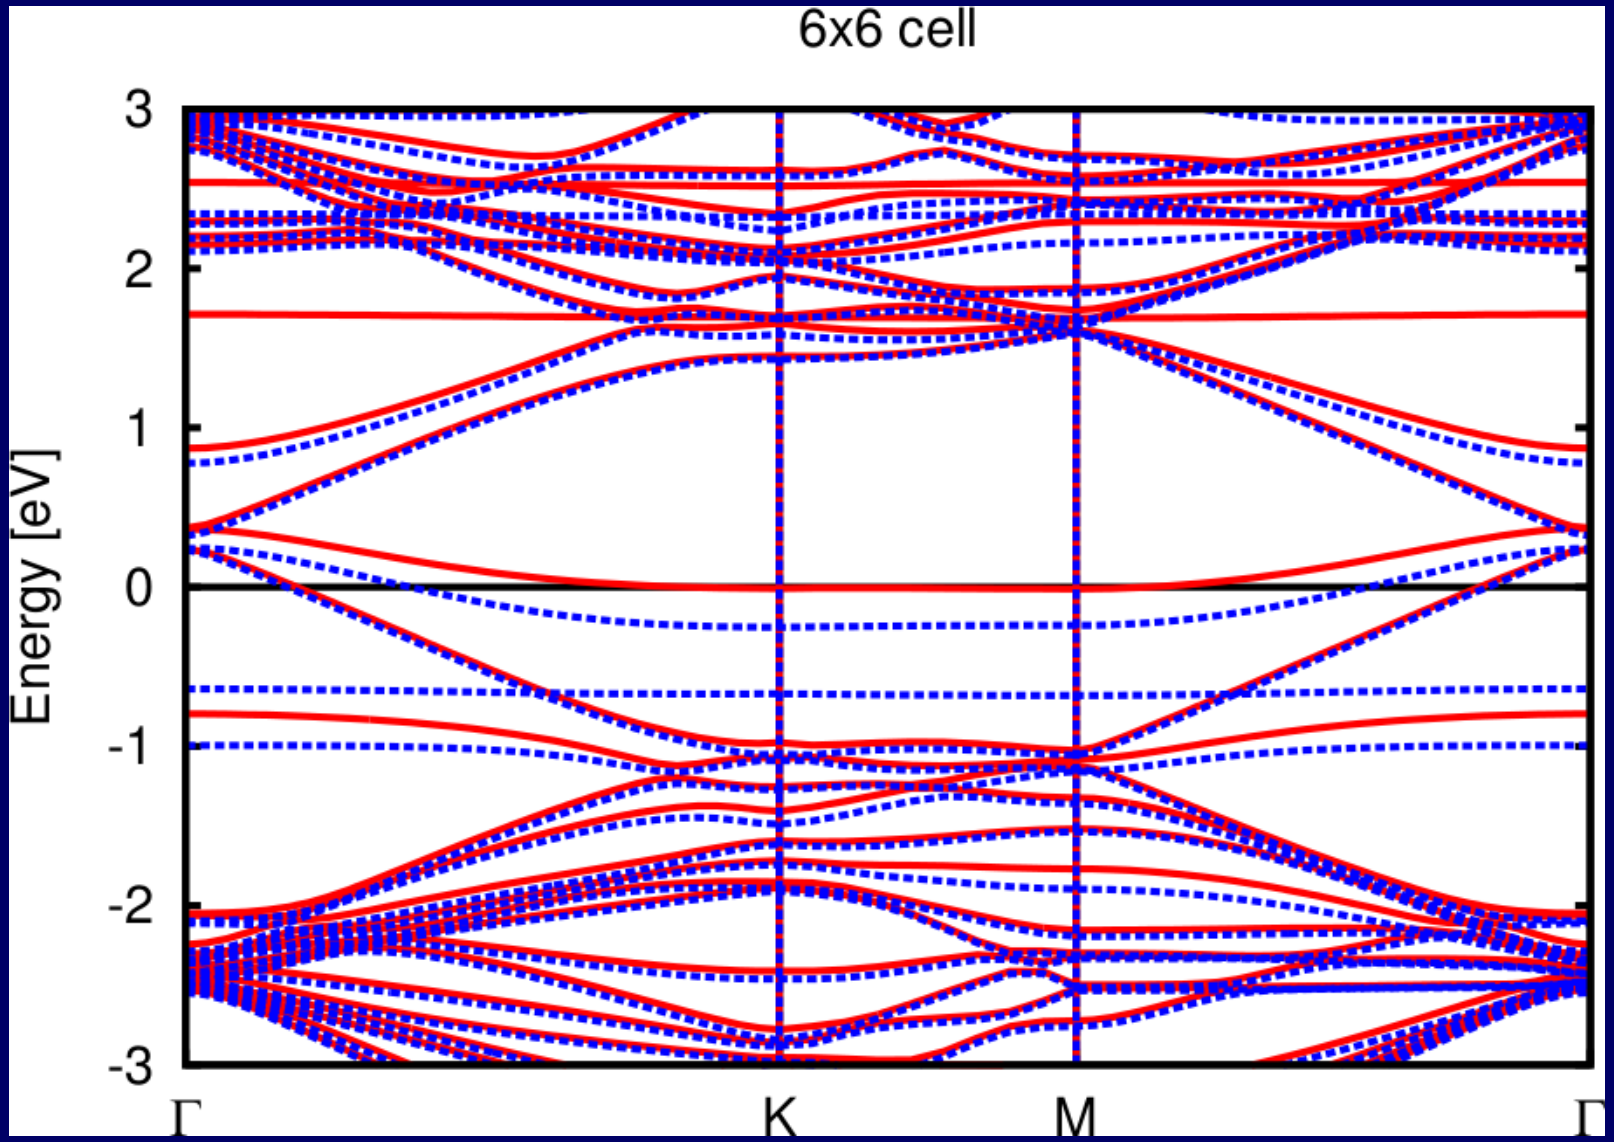

Why do we need these large unit cells and fine k-sampling?

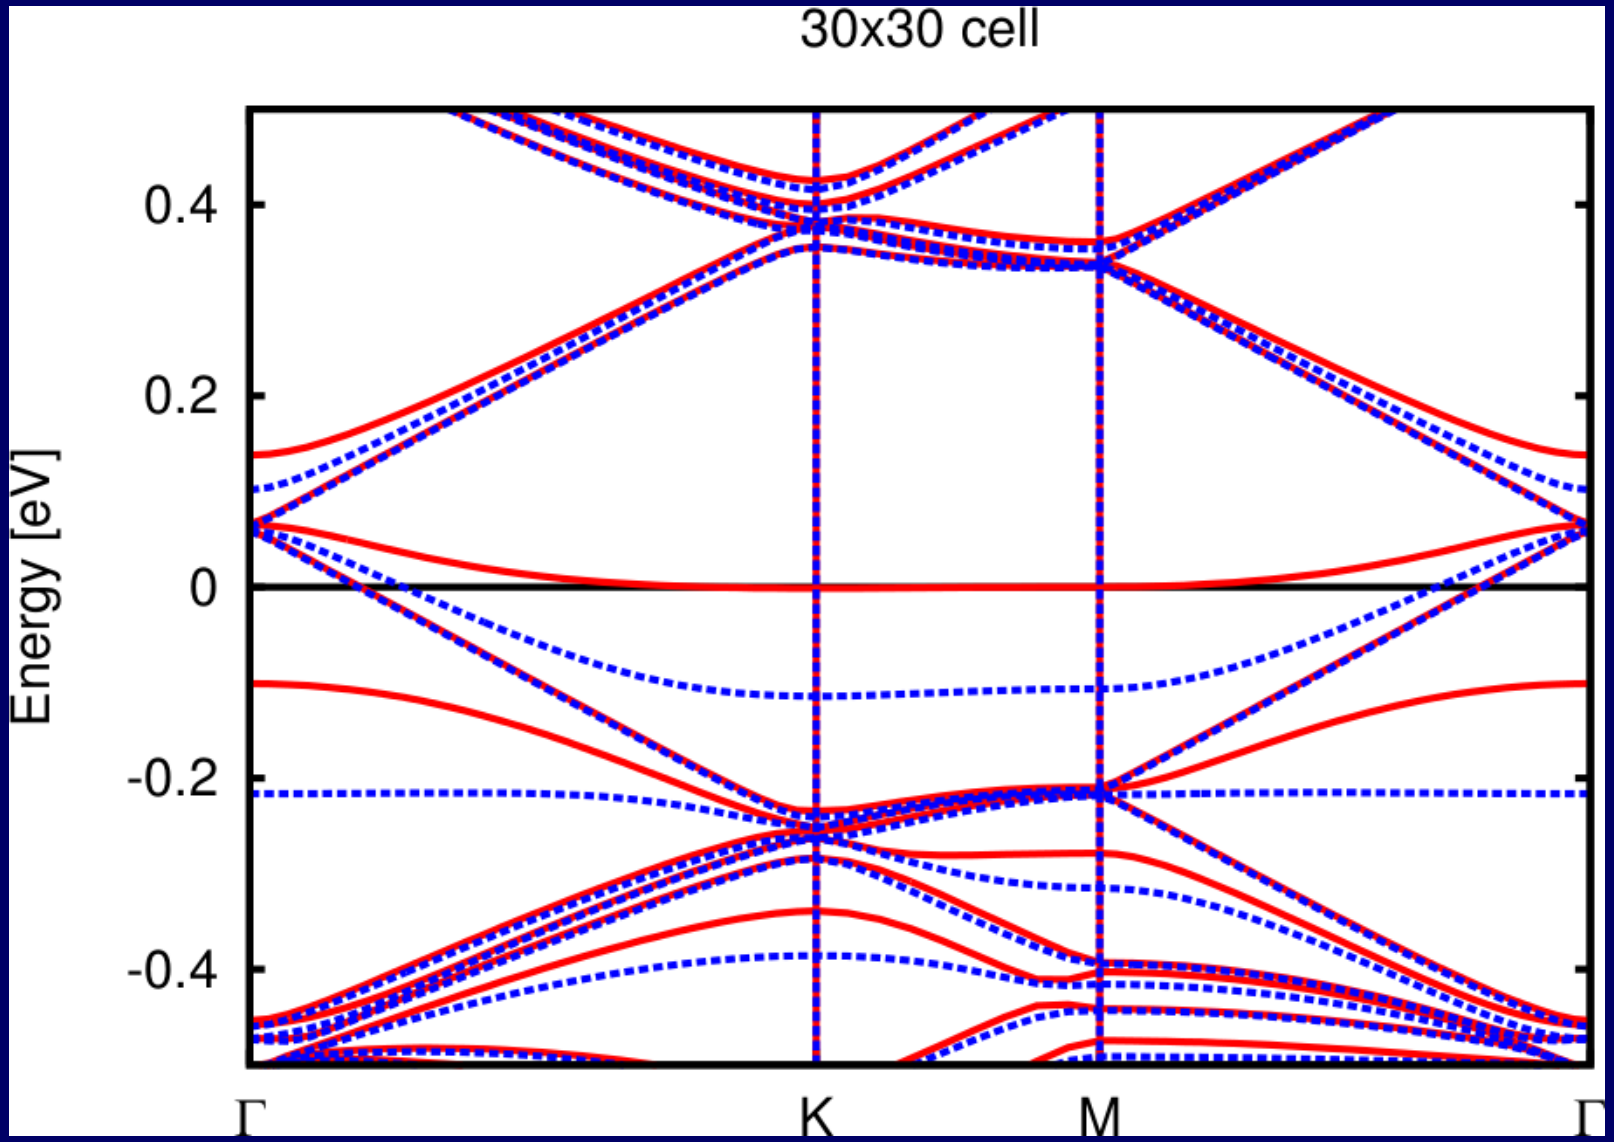

Conclusions

1. Ab initio O(N) energies and forces on metallic nanostructures

CuO: large scale defect structures; domain nucleation and growth, link between forces and chemical reactivity

Vacancies on Graphene: magnetic and mechanical properties

2. PRACE Tier-0 Project Access: Key points for success

Preparatory Access: demonstrate code scalability

Project Access: World-class science project

Apply in the next PRACE call !!

SPM Theory & Nanomechanics Group

SPM-TH

- www.uam.es/spmth

- **Group Leader:**

  Prof. Rubén Pérez

- ***Postdoctoral Researchers*:**

○ *Pablo Pou (RyC)*

Nanotubes, graphene

○ *Delia Fernández*

Oxides

○ *Milica Todorovic*

○ *Guilherme Vilhena*

AFM in biology

- ***PhD Students*:**

○ *Lucía Rodrigo*

Molecular self-assembly

○ *Carlos Romero*

KPFM Simulations

○ *Diego Rodriguez*

Oxides

○ *Michael Ellner*

Acknowledgements & Funding

MINECO (Spain): Projects MAT2008-02929-NAN,

MAT2008-02939-E (FANAS EUROCORES, ESF),

MAT2011-23627

PLE2009-0061 (Japan-Spain collaboration)

CSD2010-00204 (CONSOLIDER Force-for-Future)

Spanish Supercomputing Network (RES):

Mare Nostrum (BSC, Barcelona)

PRACE (Curie, France)

ruben.perez@uam.es

http://www.uam.es/spmth