



**E-Infrastructures
H2020-EINFRA-2015-1**

**EINFRA-5-2015: Centres of Excellence
for computing applications**

**EoCoE
Energy oriented Center of Excellence
for computing applications**

Grant Agreement Number: EINFRA-676629

**D3.2 - M24
Optimization of methods**

Project and Deliverable Information Sheet

EoCoE	Project Ref:	EINFRA-676629
	Project Title:	Energy oriented Centre of Excellence
	Project Web Site:	http://www.eocoe.eu
	Deliverable ID:	D3.2
	Lead Beneficiary:	CEA
	Contact:	Thierry Deutsch
	Contact's e-mail:	Thierry.Deutsch@cea.fr
	Deliverable Nature:	Report
	Dissemination Level:	PU*
	Contractual Date of Delivery:	30/09/2017
	Actual Date of Delivery:	18/12/2017
	EC Project Officer:	Carlos Morais-Pires

* - The dissemination level are indicated as follows: PU – Public, CO – Confidential, only for members of the consortium (including the Commission Services) CL – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	Title :	Optimization of methods
	ID :	D3.2
	Available at:	http://www.eocoe.eu
	Software tool:	L ^A T _E X
Authorship	Written by:	T. Deutsch
	Contributors:	M. Celino (ENEA), I. Duchemin (CEA), M. Salanne (MdlS), D. Borgis (MdlS), A. Walker (UBAH), U. Aeberhard (JUELICH), S. Giusepponi (ENEA), M. Gusso (ENEA)
	Reviewed by:	PEC members

Document Keywords:

DFT, Atomic-scale, iPV, Organic Photovoltaic (oPV), Batteries, Supercapacitors

Executive Summary:

Optimization of all numerical codes ported in the infrastructure and used for supercapacitors, PV and batteries.

Contents

1	Introduction	5
2	Optimization of the HPC environment for atomistic simulations for silicon based PV devices (from M12 to M24)	6
2.1	Porting on platforms	6
3	Optimizing large electronic structure methods to bridge the gap between DFT and force fields (BigDFT code)	9
4	Optimisation of the simulation of the electrolytes (Metalwalls code)	11
5	Optimisation of the ion dynamics in the perovskite cells	13
5.1	Drift-diffusion methods	13
5.2	Molecular dynamics results	14
5.3	BathKMC code	14
6	Optimisation of the GW and Bethe-Salpeter methods (Fiesta code)	15
7	Optimization of the NEGF method for photocarrier transport across hetero-interfaces in inorganic solar cells (PVnegf code)	16

1. Introduction

Computational materials modeling plays a crucial role in the design of devices for efficient low cost energy generation and storage. The accuracy of predicted macroscopic quantities depends on that of the atomic scale models describing the interatomic forces and how they are implemented on larger length and time scales. Despite its large demand on computer resources, materials modeling has a considerable impact in research and industry areas. Applications, for example inorganic and organic photovoltaics (PV), super-capacitors and batteries, can benefit from atomic scale design of materials to understand and improve charge transfer at molecular level.

The main objectives are:

- To provide a set of computational routines for morphology, electronic structure and transport properties of energy-related materials for PV, batteries and super-capacitors;
- To set up a screening methodology for designing materials for PV, rechargeable batteries and super-capacitors with optimal energy conversion and storage capabilities;
- To demonstrate how the computing infrastructure can address challenging problems in the field of energy by focusing on their atomic scale origin.

Deliverable D3.1 concerns the development of accurate numerical models for the simulation of materials of interest in the field of energy. This process leads to the design of materials at the atomic scale in view of the computation of macroscopic physical quantities in the fields of PV, batteries and supercapacitors. Thus it is very important to have atomic scale configurations that are able to reproduce the microscopic behavior that is at the origin of the macroscopic one. However this is not sufficient because robust and efficient algorithms are needed to compute both the structural and the electronic properties of the materials. The problem that scientists are facing now is that many algorithms implemented in the numerical codes are often not computationally efficient even very reliable to compute physical and chemical properties. Thus, to design an application that is able to scale up toward the exascale it is not sufficient to optimize the algorithms in terms of I/O, parallelism, libraries and so on. Indeed, we need to look for and implement more efficient numerical models that can be used to exploit the computational power of the new generation of high performance hardware.

In this respect, close to the design of models for the simulation of materials, new numerical approaches and algorithms are needed. Indeed the activities of WP3 are not only devoted to develop application lines (PV, batteries and supercapacitors) but also to new and improved methodologies to simulate materials and manage numerical codes. In this respect it is very important to benchmark a standard numerical code, as for example Quantum Espresso, on an application of interest for energy like PV to characterize its limits in terms of scalability and efficiency. At the same time, it is important to develop parametrization strategies, in which it is preserved the accuracy but it is improved the overall efficiency toward the exascale computers. In this respect, it is clear that it is crucial to develop new strategies to manage large numerical codes via modularity approaches and the coupling of different numerical codes to optimize their scalability.

The scope of deliverable D3.2 is to report the new advances in the field of materials

D3.2 Optimization of methods

for energy that comes from the search of new methodologies and models that could be more efficient on the new generation of computer hardware for exascale. In this respect, deliverable D3.2 is a transversal deliverable that report the new advances related to activities described in task T3.2 and those, described in T3.3 and T3.4, already started but not yet concluded.

2. Optimization of the HPC environment for atomistic simulations for silicon based PV devices (from M12 to M24)

As described in the first deliverable D3.1, an atomistic model has been developed for heterojunction silicon solar cells. The reliability of this model is confirmed by the good agreement with the experimental results, in terms of structural and electronic properties. The model is able to reproduce accurately the physical and chemical properties of an interface between crystalline silicon and amorphous hydrogenated silicon. The detailed knowledge of this interface is crucial for the performance of the PV cell because recombination processes are generally localized in this portion of the solar cell. However, to take into account all the possible microscopic processes that can happen at the interface a very large one should be considered. Enlarging the interface, however, on both sides of the interface, large systems should be designed to mimic the contribution of the infinite bulk materials and minimize the fictitious strains.

Moreover, to properly design the interface a quantum approach is needed to accurately reproduce the microscopic configurations and their electronic properties. In the following we report some benchmarks we have performed on this system by using the very popular Quantum Espresso package. This is useful to understand the performance of the code on this atomistic system and push the simulations to the limits. This step is crucial to test and trigger properly the numerical model by exploiting both the code and the available computational platforms. Given these computational results we will be able to scale up the model, eventually by testing different numerical codes (cp2k for example), to exploit larger sections of the HPC platforms.

2.1 Porting on platforms

ENEA and Jülich, performance evaluation session

In view of an intensive use of computational resources for an *ab initio* study of a crystalline Silicon (c-Si) and hydrogenated amorphous Silicon (a-Si:H) interface (see Fig. 1), a measure of the performances of three HPC infrastructures (CRESCO4, CRESCO5, JU-RECA) was carried out. This enables us to set the optimal computational resources and gives an estimation of the computational time. The benchmark is based on the time required to execute a self consistent calculation, for a single energy minimization in the computation of the ground-state energy of the c-Si/a-Si:H interface depicted in Fig. 1. The calculation is performed using the *ab initio* PWscf (Plane-Wave Self-Consistent Field) code of the Quantum ESPRESSO suite.

The ENEA CRESCO computing facilities are based on the multi-core x86_64 architecture and is made up of various clusters: we tested the CRESCO4 and CRESCO5 clusters located in Portici ENEA centre CRESCO4 is composed by 304 compute nodes. Each node hosts a dual-socket 8 cores CPU Intel E5-2670 processor (Sandy Bridge architecture) which operates at a clock frequency of 2.6 GHz and has 20 MB of cache and AVX instructions set. The total number of cores is 4864. CRESCO5 is composed by 40

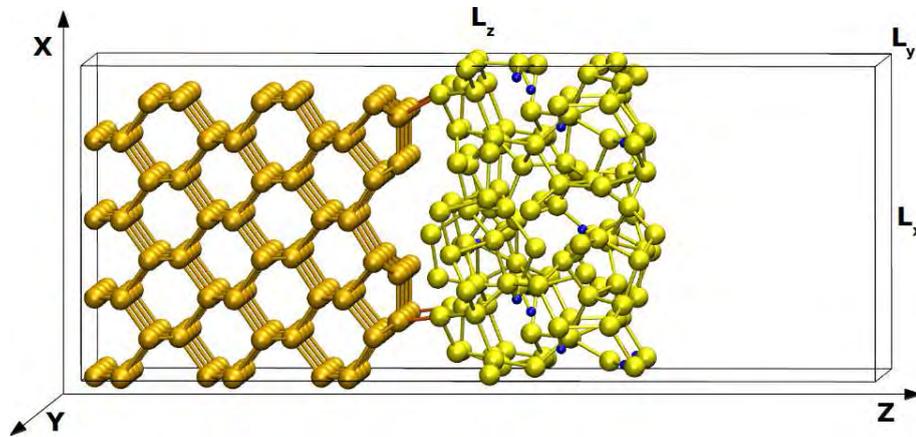


Figure 1: Snapshot of the initial a-Si:H/c-Si interface in the simulation box. The structure is infinitely extended in both x and y directions. A void region is considered to suppress the interaction, due to periodic boundary conditions, between the external surfaces. Free surfaces and a-Si:H/c-Si interface are perpendicular to the y axis. Hydrogen atoms are in blue, Silicon atoms are in orange in the c-Si side and are in yellow in the a-Si:H side.

compute nodes. Each node hosts a dual-socket 8 cores CPU Intel E5-2630 v3 processor (Haswell architecture) which operates at a clock frequency of 2.4 GHz and has 20 MB of cache and AVX 2.0 instructions set. The total number of cores is 640. Both clusters have 64 GB of RAM memory per node (4 GB per core) and are interconnected via an IB 4xQDR Qlogic/Intel 12800-180 switch. The Operating System for the clusters is the Red Hat Enterprise Linux ver. 2.6.32-220.7.1.el6.x86_64. Quantum Espresso ver. 5.1.2 is compiled with Intel compiler ver. 14.0 using the mathematical library MKL ver. 11.1 and OpenMPI ver. 1.4.3.

The third cluster is JURECA HPC hosted at Jülich Supercomputing Centre (JSC) JURECA is composed by 1872 compute nodes. Each node hosts a dual-socket 12 cores CPU Intel E5-2680 v3 processor (Haswell architecture) which operates at a clock frequency of 2.5 GHz and has 30 MB of cache and AVX 2.0 instructions set. The total number of cores is 44928, the cluster has 128 GB of RAM memory per node that are interconnected via a Mellanox EDR InfiniBand. The Operating System for the cluster is the CentOS 7 Linux ver. 3.10.0-327.10.1.el7.x86_64. Quantum Espresso ver. 5.1.1 is compiled with Intel compiler ver. 15.0 using the mathematical library MKL ver. 11.2 and ParaStationMPI ver. 5.0.

To compare the performance and scalability of the three clusters, we collected the time necessary to perform a self consistent calculation, for a single energy minimization in the computation of the ground-state energy of the c-Si/a-Si:H interface depicted in Fig. 1. Total wall time for the three clusters using 48, 96, 192 and 384 cores are illustrated in panel a of Fig. 2. The calculation is performed considering only the Γ point and using MPI parallelization. Moreover, in Fig. 2, there are shown the speed up (panel b) and the resulting efficiency (panel c) calculated for each cluster with respect to the time measured in the 48 cores execution. This type of computation has a good parallel scalability until 192 cores with an efficiency of about 75%. Then, this should be the optimal number of cores used for BOMD (Born-Oppenheimer Molecular Dynamics) simulation, in which at every timestep, the electronic ground-state of the system must be reached.

The results confirm that the more recent processors, E5-2630 v3 and E5-2680 v3

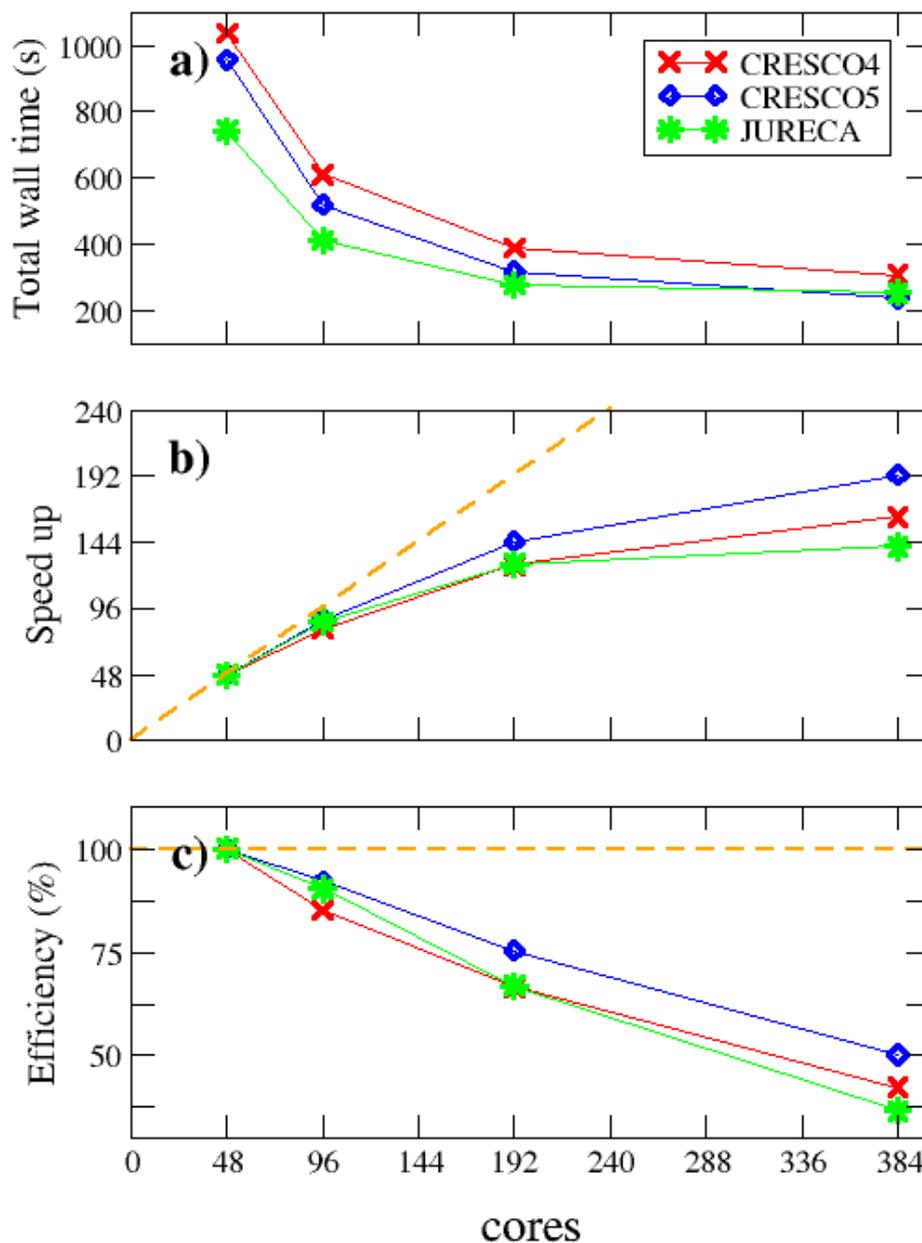


Figure 2: a) Total wall time (s) for a self consistent calculations for a single energy minimization. b) Speed up calculated with respect to the time in the 48 cores calculation, and c) the corresponding efficiency (%). Red lines(symbols) are for CRESCO4, blue lines(symbols) are for CRESCO5 and green lines(symbols) are for JURECA. Dashed orange lines represent ideal parallel scalability.

mounted on CRESCO5 and JURECA clusters, respectively, have the best performance with reductions of execution times, with respect to CRESCO4, in the range 30÷40% for JURECA and in the range 10÷20% for CRESCO5. Then, even if, they have the same architecture (Haswell) and the same instruction set (AVX 2.0), E5-2680 v3, being two and half times more expensive, legitimises its cost having better performance.

3. Optimizing large electronic structure methods to bridge the gap between DFT and force fields (BigDFT code)

The objective is to optimize the linear scaling version of BigDFT in order to have a large overlap between DFT and force fields in terms of system size. Then we can test the different approximations from Quantum Mechanics (QM) to Molecular Modelling (MM) through QM/MM flavors. We improved considerably the linear scaling version of BigDFT, mainly working on the parallelization. In September 2016, our record is 250,000 atoms using 10,000 cores without loss of accuracy comparing to the traditional cubic scaling version.

We wrote a review (see Fig. 3) showing the need of large DFT calculations¹. In this article, we have demonstrated the interest of large DFT calculations mainly to fill the gap between classical DFT and force fields. There already exist a large number of approximations to perform larger and larger calculations as fragment approaches and embedding schemes with implicit solvent to take into account an electrostatic environment given by a solvent. There is a need to test and control all these approximations which are dependent on the nature of the considered system. We shown than, thanks to the linear scaling method, it is now possible to control all these approximations, doing once one full quantum large calculation without approximations for a given system.

This is why we improved the fragment approaches in BigDFT and developed also a generalized Poisson and Poisson-Boltzman solver in order to test all these approximations.

In the article *“Soft-Sphere Continuum Solvation in Electronic-Structure Calculations”*², we present an implicit solvation approach (see Fig. 4) where the interface between the quantum-mechanical solute and the surrounding environment is described by a fully continuous permittivity built up with atomic-centered “soft” spheres. This approach combines many of the advantages of the self-consistent continuum solvation model in handling solutes and surfaces in contact with complex dielectric environments or electrolytes in electronic-structure calculations. In addition it is able to describe accurately both neutral and charged systems. The continuous function, describing the variation of the permittivity, allows to compute analytically the nonelectrostatic contributions to the solvation free energy that are described in terms of the quantum surface. The whole methodology is computationally stable, provides consistent energies and forces, and keeps the computational efforts and runtimes comparable to those of standard vacuum calculations. The capability to treat arbitrary molecular or slab-like geometries as well as charged molecules is key to tackle electrolytes within mixed explicit/implicit frameworks. We show that, with given, fixed atomic radii, two parameters are sufficient to give a mean absolute error of only 1.12 kcal/mol with respect to the experimental aqueous solvation energies for a set

¹ *Challenges in large scale quantum mechanical calculations*, Laura E. Ratcliff, Stephan Mohr, Georg Huhs, Thierry Deutsch, Michel Masella, and Luigi Genovese, WIREs Comput. Molecular Science (2016), doi: 10.1002/wcms.1290

² *Soft-Sphere Continuum Solvation in Electronic-Structure Calculations*, Giuseppe Fiscaro, Luigi Genovese, Oliviero Andreussi, Sagarmoy Mandal, Nisanth N. Nair, Nicola Marzari, and Stefan Goedecker, J. Chem. Theory Comput. (2017), 13 (8), pp 3829–3845

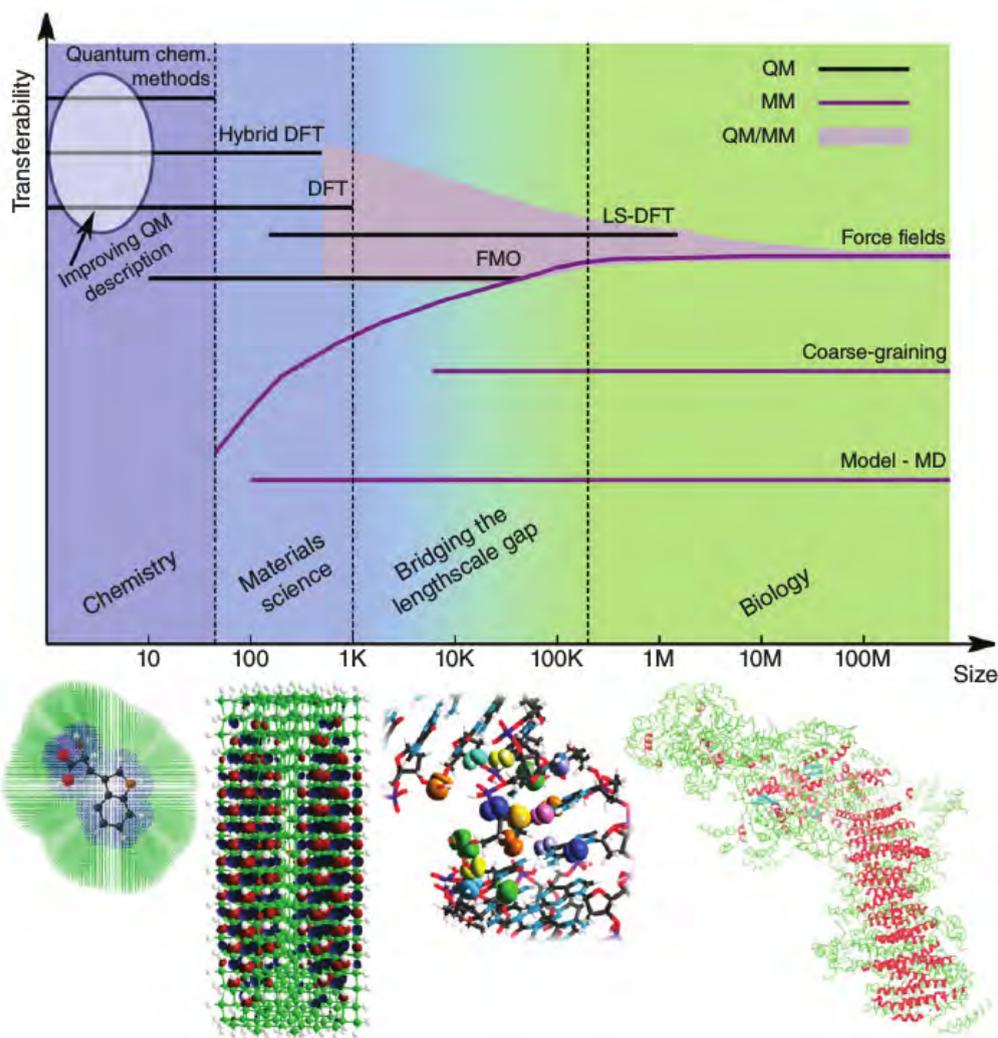


Figure 3: Overview of the popular methods used in simulations of systems with atomistic resolution, showing the typical length scales over which they are applied as well as the degree of transferability of each method, i.e., the extent to which they give accurate results across different systems without re-tuning. On the left hand side we have the Quantum Chemistry methods which are highly transferable but only applicable to a few tens of atoms; on the right hand side we see the less transferable (semi-)empirical methods, which can however express reliable results (as they are parametrized for) for systems containing millions of atoms; and in the middle we see the methods—in particular linear-scaling DFT which can bridge the gap between the two regimes. The vertical divisions and corresponding background colours give an indication of the fields in which the methods are typically applied, namely chemistry, materials science, biology, and an intermediary regime (‘bridging the length scale gap’) between materials science and biology. The line colours indicate whether a method is QM or MM, while the typical regime for QM/MM methods is indicated by the shaded region. Some representative systems for the different regimes are depicted along the bottom.

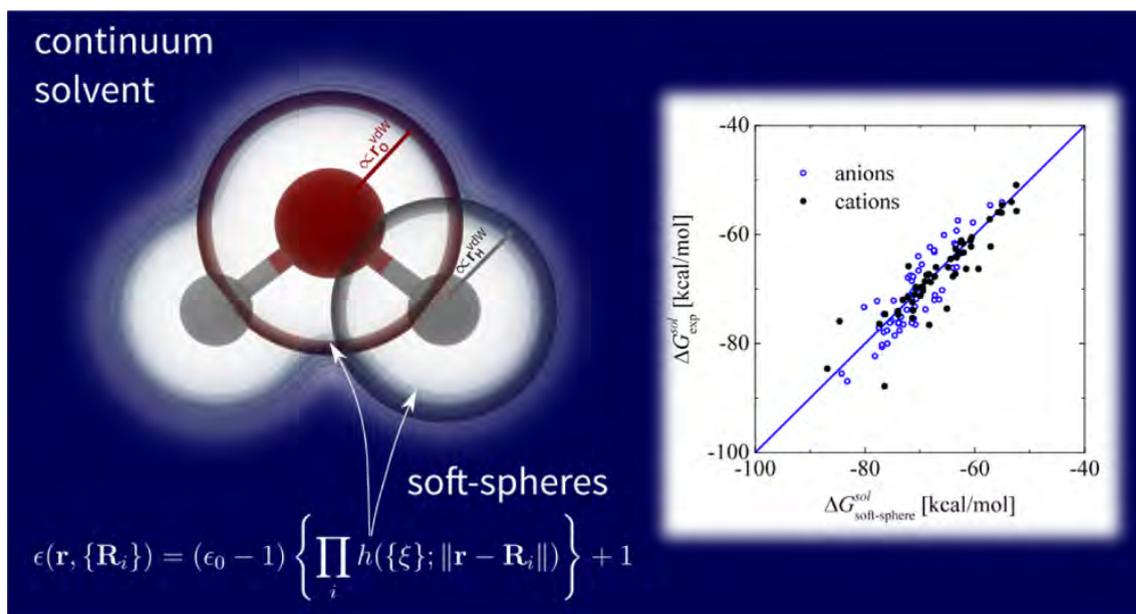


Figure 4: Illustration of our implicit solvation approach using a fully continuous permittivity built with atomic-centered "soft" spheres on the left. On the right panel, we show the experimental and *ab initio* aqueous solvation free energies for a set of 112 single-charged ions, of which 60 anions (empty blue circles) and 52 cations (full black circles), obtained with our method.

of 274 neutral solutes. For charged systems, the same set of parameters provides solvation energies for a set of 60 anions and 52 cations with an error of 2.96 and 2.13 kcal/mol, respectively, improving upon previous literature values. To tackle elements not present in most solvation databases, a new benchmark scheme on wettability and contact angles is proposed for solid-liquid interfaces and applied to the investigation of the stable terminations of a CdS (11 $\bar{2}$ 0) surface in an electrochemical medium.

Porting on platforms

Brice Videau, hired in our laboratory for this project is currently testing new architectures thanks to the BOAST technology developed in collaboration with LIG (Laboratoire d'Informatique de Grenoble, Prof. Jean-François Méhaut, WP1) as OpenPower and Xeon Phi. We have also improved the efficiency of our BigDFT code on massively parallel GPU architecture using GPUDirect especially the calculation of the exact exchange interaction. Thanks to BOAST, the goal is to develop a library of convolutions automatically tuned for a given architecture. This work is in progress.

4. Optimisation of the simulation of the electrolytes (Metalwalls code)

Metalwalls is a molecular dynamics code devoted to the study of electrochemical cells at the molecular scale. In our simulations, we model electrodes with various geometries and electrolytes of various natures. The originality of Metalwalls is that the simulations are performed at fixed **applied potential** between the electrodes, which implies a dynamic calculation of the charges on the electrode atoms at every step. Due to this particularity, the computational cost of the simulations is large compared to that of conventional Molecular Dynamics (MD) studies. In addition, the systems we aim to simulate, supercapacitors, require us to include a large number of carbon atoms in order to account for

D3.2 Optimization of methods

their complex topologies, i.e. porous carbons.

To apply a constant potential difference between two electrodes, our approach consists of imposing a constraint on the potential Ψ_i felt by each electrode atom i and to deduce the resulting set of electrode charges $\{q_j\}$ at each timestep^{3 4}. This results in a set of self-consistent equations:

$$\Psi_i(\{q_j\}) = V_i^0 \quad (1)$$

The target value V_i^0 is the preset electrode potential, it takes the same value for all ions in a particular electrode. The potential within the electrode is thus a constant, which is the electrostatic condition on the potential inside a conductor. Based on our experience with polarizable force fields, in the current version of Metalwalls the set of charges is not directly calculated by solving the equation 1, but rather by minimizing the total potential energy for our dynamical system,

$$U = U_c - \sum_i V_i^0 q_i, \quad (2)$$

where U_c is the Coulomb energy of the whole system (i.e. the electrode and the electrolyte charges), and the sum runs over the variable electrode charges. The second term is in fact added so that minimizing this function with respect to all the variable charges simultaneously gives the constant potential condition. A predictor step^{5 6} is made prior to the minimization, for which we use the conjugate gradients method. Typically, we need to perform ten minimization steps before reaching convergence for physical systems.

Recently, another implementation was proposed by Wang *et al.*⁷, in which the mathematical problem is solved by rewriting it as follows

$$\mathbf{A}\mathbf{Q} = \mathbf{V} - \mathbf{B} \quad (3)$$

where \mathbf{A} is a matrix containing terms that depend on the position of the electrode atoms $\{R_i\}$ only, \mathbf{Q} is a vector containing all the fluctuating electrode charges, \mathbf{V} a vector containing the target potential of each electrode atom and \mathbf{B} is a vector in which the terms depend on the position of the electrolyte atoms $\{r_i\}$. Since the electrode atoms are held fixed during the simulations, the matrix \mathbf{A} only has to be calculated (and inverted) once, at the first step of the simulation. On the contrary, \mathbf{B} is updated at each timestep, its evaluation and subsequent multiplication for the precalculated inverse of \mathbf{A} , thus becoming the costly part of the calculation.

In order to test this approach, we have implemented it in Metalwalls and compared its performance to that of the conjugate gradient approach for a generic test case, consisting

³ *Influence of surface topology and electrostatic potential on water/electrode systems*, J. I. Siepmann and M. Sprik, *J. Chem. Phys.*, 102:511--524, 1995.

⁴ *Electrochemical interface between an ionic liquid and a model metallic electrode*, S. K. Reed, O. J. Lanning and P. A. Madden, *J. Chem. Phys.*, 126:084704, 2007.

⁵ *Time-reversible always stable predictor--corrector method for molecular dynamics of polarizable molecules*, J. Kolafa, *J. Comput. Chem.*, 25:335--342, 2004.

⁶ *Efficient and accurate Car-Parrinello-like approach to Born-Oppenheimer molecular dynamics*, T. D. Kühne, M. Krack, F. R. Mohamed and M. Parrinello, *Phys. Rev. Lett.*, 98:066401, 2007.

⁷ *Evaluation of the constant potential method in simulating electric double-layer capacitors*, Z. Wang, Y. Yang, D. L. Olmsted, M. Asta and B. B. Laird, *J. Chem. Phys.*, 141:184102, 2014

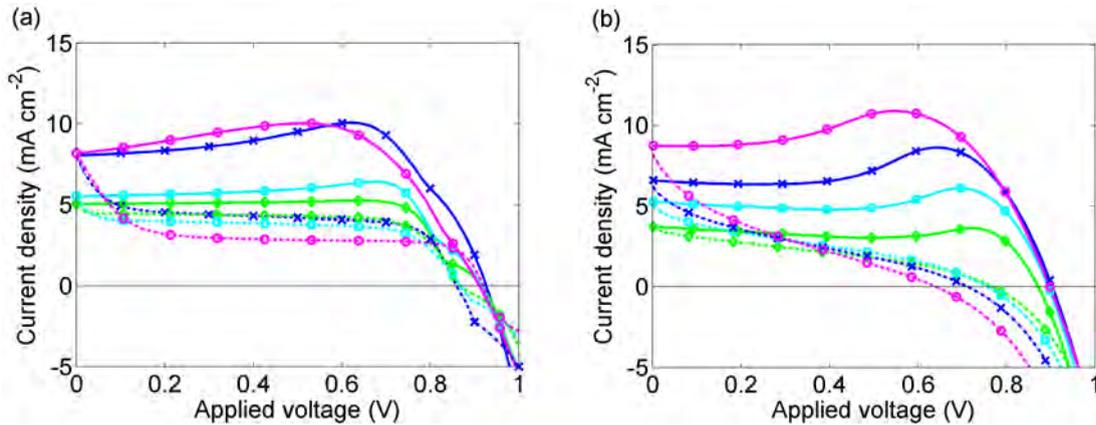


Figure 5: a) Calculated current-voltage curves without allowing for preconditioning; (b) measured current-voltage curves for cell 1 after 2 previous scan cycles and preconditioning for 5s at FB=1.2 V. Solid lines show the FB-SC scan; broken lines show the SC-FB scan. Scan rates are 1 V/s (magenta, circles), 500 mV/s (blue, crosses), 250 mV/s (cyan, filled squares), 100 mV/s (green, diamonds)

of an ionic liquid electrolyte put in contact with two graphite electrodes. The results are the following:

- Each of the conjugate gradient step takes 0.188 seconds. The total contribution of the minimization procedure to the MD step is therefore of ≈ 1.88 seconds (i.e. 90 % of the total cost).
- The calculation of the matrix \mathbf{A} takes 146 seconds, while its inversion takes 4 seconds. The calculation of the vector \mathbf{B} takes 0.36 seconds.

The first MD step with the matrix-based approach therefore costs 75 times more than that with our current implementation of conjugate gradients, but then every MD step will cost ≈ 4 to 5 times less. Since in a typical MD simulation at least one million steps are performed, the matrix-based approach will be much more efficient for systems where the electrodes are held fixed. So far this test was performed on a single-core. We now need to implement an MPI version in order to compare the relative performance on supercomputers, and to test other systems with more complex geometries.

5. Optimisation of the ion dynamics in the perovskite cells

5.1 Drift-diffusion methods

Perovskite cells suffer from hysteresis which greatly affects their performance. To follow on from results reported for D3.1, the following paper⁸ has been submitted to Physical Review B acknowledging EoCoE. In this paper we describe an asymptotic series approach that replaces the ion dynamics by interfacial (nonlinear) capacitances on the edges of the perovskite layer. Favourable comparisons result between the results of the numerical and asymptotic models over a wide range of cell operating conditions and with experimental observations. This method is widely applicable to other double layer problems occurring in electrochemical applications such as the evolution of transmembrane potentials in living cells.

This method has been used to optimise drift diffusion simulations of perovskite cells

⁸ *Systematic derivation of a surface polarization model for planar perovskite solar cells.* N. E. Courtier, J. M. Foster, S. E. J. O’Kane, A. B. Walker, G. Richardson

D3.2 Optimization of methods

that include the charge transport layers as well as the perovskite layers. The following paper⁹ is in preparation. This result was obtained by a continuum (drift diffusion) model. Task 3.1.1 is developing a multiscale model that links microscopic calculations of ion vacancies to the macroscopic drift diffusion model that produced Figure 5.

5.2 Molecular dynamics results

Existing molecular dynamics codes have difficulty in isolating the motion of defects in the perovskite layer from that of the non-mobile ions.

Our work on the use of the CP2K code to understand ion motion in perovskites at a microscopic level has resulted in a submitted paper¹⁰. In this paper we explore why perovskite compositions with a mixture of cation species on the A-site show the best performance and possess higher stability than those with a single cation species. We show that substitution of low concentrations of smaller cations on the A-site in formamidinium lead iodide ($\text{CH}(\text{NH}_2)_2\text{PbI}_3$) results in a global 'locking' of the PbI_6 octahedra tilting. In the locked structure the octahedra tilt at a larger angle but undergo a much reduced amplitude of rocking motion. A key impact of this feature is that the rotational or tumbling motion of the $\text{CH}(\text{NH}_2)_2^+$ molecular ion in a locked cage is severely restricted. We discuss the impact of locking on the photovoltaic performance and stability.

We investigated azetidinium as a potential third cation for synthesizing organic-inorganic perovskites and have submitted a paper¹¹ to J. Materials Chem. A. Azetidinium lead iodide is a stable, bright orange material which does not appear to form a 3-dimensional or 2-dimensional perovskite. It was successfully used as the absorber layer in solar cells. We show with our electronic structure calculations that it is possible to make mixed cation devices by adding the azetidinium cation to methylammonium lead iodide. Mixed azetidinium-methylammonium cells show improved performance and reduced hysteresis compared to methylammonium lead iodide cells.

In another paper that will be submitted early Sept 2017¹², UBAH calculated thermodynamic stability and electronic properties of germanium and tin based binary metal lead-free hybrid perovskite solid solutions for all compositions that agree well with our experimental measurements. Our findings demonstrate an attractive family of lead-free (and hence less toxic) perovskite semiconductors with a favourable bandgap range for efficient single junction solar cells. These calculations were performed with the VASP electronic structure code. The novelty here lies in the new materials being investigated.

5.3 BathKMC code

A Kinetic Monte Carlo, KMC, simulation code for perovskite cells is being developed as a mesoscopic simulation of charge transport. In this way, we can bridge the length

⁹ *How device architecture influences planar perovskite solar cell performance*, N.E. Courtier, J. Cave, A.B. Walker, J.M. Foster, & G. Richardson

¹⁰ *Good Vibrations: Locking of Octahedral Tilting in Mixed-Cation Iodide Perovskites for Solar Cells*, D Ghosh, P Walsh Atkins, M S Islam, A B Walker, C Eames

¹¹ *Azetidinium Lead Iodide for Perovskite Solar Cells*, S Pering, W Deng, J R Troughton, P Kubiak, D Ghosh, R Niemann, F Brivio, F E Jeffrey, A B Walker, M S Islam, T M Watson, P Raithby, A Johnson, S Lewis, P J Cameron, submitted to J. Materials Chem. A

¹² *Lead Free Germanium-Tin Solid Solutions Based Perovskite Semiconductors: Structural and Optoelectronic Properties*, S Nagane, D Ghosh, R L Z Hoye, B Zhao, M S Islam, A B Walker, S Ogale and A Sadhanala

D3.2 Optimization of methods

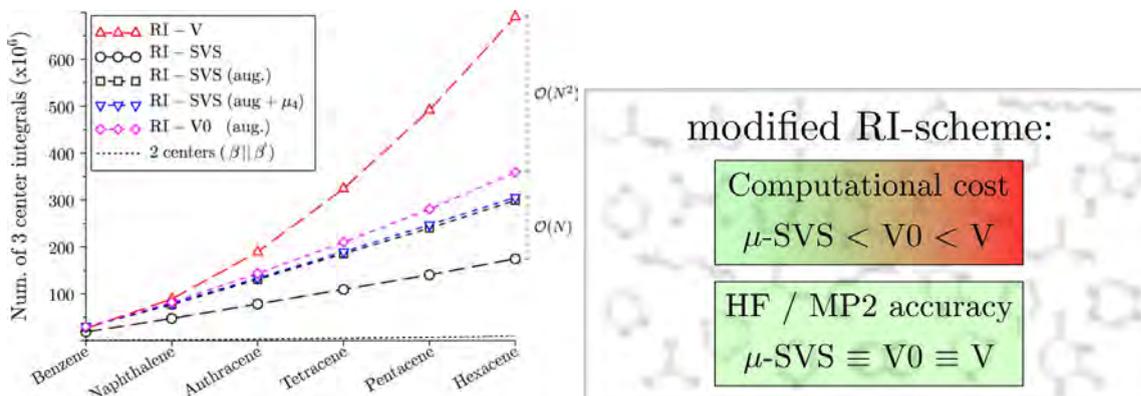


Figure 6: Number of three-center integrals involved through the different fitting schemes over the acene series, for a nonzero integral 10^{-9} cutoff value set to 10 Hartree. We compare here the Coulomb (RI-V) and density (RI-SVS) fitting schemes using the Cartesian cc-pVTZ-RI auxiliary basis, with the standard and constrained (up to fourth order momenta μ_4) RI-SVS scheme and the RI-V0 scheme using the augmented Cartesian aug-cc-pVTZ-RI auxiliary basis. The corresponding number of two-center Coulomb integrals between elements of the Cartesian cc-pVTZ-RI auxiliary basis is also reported.

scales between the atomistic molecular dynamics simulations and our drift diffusion code. It is now clear that the charge carriers in perovskites are delocalised. We are therefore developing a device Monte Carlo simulation code using the methods developed for inorganic semiconductor devices described in the book of R W Hockney and J W Eastwood¹³. These device simulations require a fast Poisson solver that allows for the mobile ion charges and that will be developed with R Scheichl, UBAH, WP1.

6. Optimisation of the GW and Bethe-Salpeter methods (Fiesta code)

The code FIESTA, developed both by CNRS (X. Blase) and CEA (I. Duchemin), is based on the Many-Body Perturbation theory and more specifically on the GW method and Bethe-Salpeter equation (BSE). This code is used to calculate the optical spectrum of molecules (GW) and the energies of the excitons (BSE) It uses Gaussian basis set to calculate the spectrum of quasiparticles and first representing the orbitals. In order to represent the density and calculate the Coulomb integral, an auxiliary basis set is used instead of the natural product basis reducing considerably the cost of calculations. The introduction of auxiliary bases to approximate molecular orbital products has paved the way to significant savings in the evaluation of four-center two-electron Coulomb integrals.

In an accepted article¹⁴, I. Duchemin et X. Blase present a generalized dual space strategy that sheds a new light on variants over the standard density and Coulomb-fitting schemes, including the possibility of introducing minimization constraints. They improve in particular the charge- or multipole-preserving strategies introduced respectively by Baerends and Van Alsenoy that they compare to a simple scheme where the Coulomb metric is used for lowest angular momentum auxiliary orbitals only. they explore the merits of these approaches on the basis of extensive Hartree-Fock and MP2 calculations over a standard set of medium size molecules.

¹³ *Computer Simulation Using Particles*, R. W. Hockney and J. W. Eastwood, Taylor & Francis, Group, Abingdon 1988

¹⁴ *Hybrid and Constrained Resolution-of-Identity Techniques for Coulomb integrals*, I. Duchemin, J. Li, X. Blase *J. Chem. Theory Comput.* (2017), 13, 1199, DOI: 10.1021/acs.jctc.6b01215

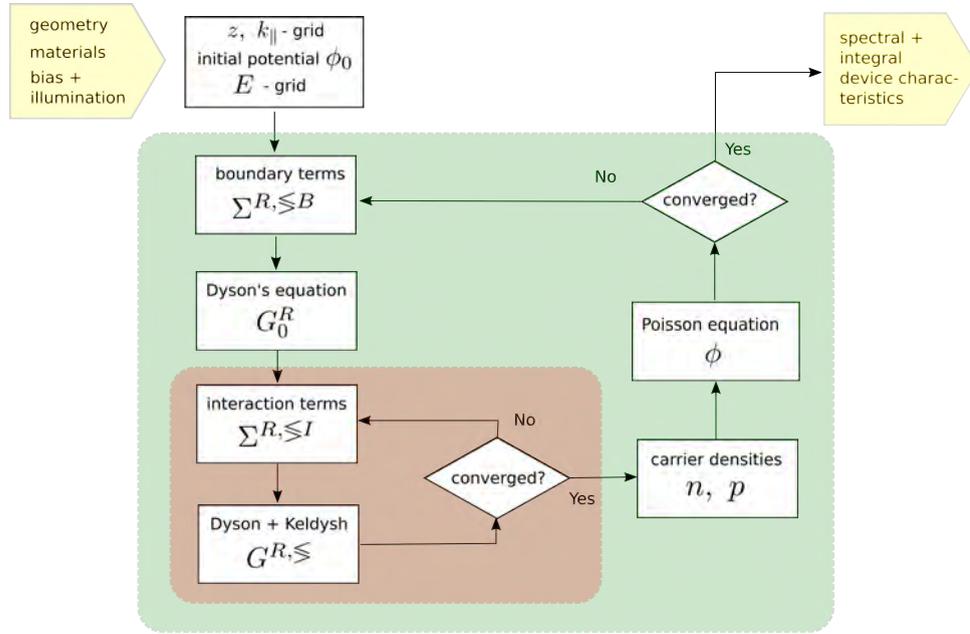


Figure 7: Computational procedure of the NEGF approach to photovoltaic device simulation. The computational meshes in real space (z), transverse momentum (\mathbf{k}_{\parallel}) and electronic energy (E) are determined by the device geometry and corresponding electronic and vibrational modes. The boundary self-energies Σ^B - encoding contact properties - are used in a Dyson equation to provide the retarded Green's function G_0^R of the noninteracting system. In an inner loop, interaction self-energies for inelastic scattering of electrons with phonons and photons are computed self-consistently with the retarded GF component G^R (related to the density of states) and the correlation functions $G^< >$ (related to the density). An outer self-consistency loop couples the evaluation of the carrier density to the computation of the electrostatic potential via solution of Poisson's equation.

Porting on platforms

All these developments were ported and tested on massively parallel architecture based on Xeon Intel processors and Infiniband fast network. We took care about the load-balancing between processors.

7. Optimization of the NEGF method for photocarrier transport across hetero-interfaces in inorganic solar cells (PVnegf code)

The *PVnegf* code solves the Dyson and Keldysh equations of the Non-equilibrium Green's functions formalism for steady-state operation of nanostructure-based thin film solar cell devices with planar geometry¹⁵. The approach provides macroscopic device characteristics as a function of bias voltage and illumination spectrum under consideration of the local electronic, vibrational and photonic structure and based on rigorous non-equilibrium quantum statistical mechanics. The current implementation is restricted to an effective mass approximation or simple tight-binding model for a two band Hamiltonian, but includes coupling to both phonons and classical light (for absorption) as well as quantized photon fields (for spontaneous emission). In the scope of WP3 of the EoCoE project, the code is used to investigate inelastic photocarrier transport at heterointerfaces

¹⁵ *Theory and simulation of quantum photovoltaic devices based on the non-equilibrium Green's function formalism*, U. Aeberhard, J. Comput. Electron. 10, 394 (2011).

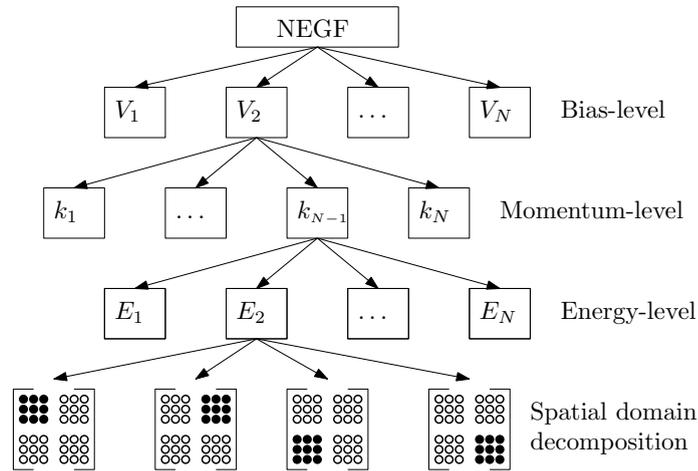


Figure 8: Levels of parallelism in the parameter space of the NEGF formalism for planar solar cell devices. While the level of the bias voltage is trivially parallelized due to complete independence of individual bias points, energy and momentum level parallelization suffers from the coupling due to inelastic scattering. In the current MPI implementation, these two levels are distributed together, while spatial domain decomposition is not considered at this point.

in silicon heterojunction solar cells under consideration of non-classical effects arising from tunneling and confinement in the proximity of potential barriers.

The general computational algorithm is sketched in Figure 7. It consists of two nested self-consistency loops. In the inner loop, the Green's function and self-energy components are iterated, which arises from the use of the self-consistent Born approximation (SCBA) for the self-energy of electron-phonon and electron-photon interaction. The evaluation of the self-energy and the solution of the linearized Dyson equation for the retarded Green's function represent the most compute-intensive parts of the algorithm. Convergence of this inner loop depends strongly on the specific device structure and operating conditions, with strong interactions and localized states as prominent hampering factors¹⁶. The outer loop encodes the self-consistent evaluation of the electrostatic potential via solution of Poisson's equation under consideration of the charge density derived from the Green's functions. In practice, the two loops can often be combined by solving Poisson equation at each inner iteration step, with marginal impact on the convergence of the inner loop.

The challenges for an efficient numerical implementation of the NEGF formalism are manifold. First of all, the Green's functions span a vast and complex parameter space in energy, transverse momentum and real space. Resolution in energy and transverse momentum is dictated by the resonance structure of electronic and vibrational modes of the device, while spatial resolution and mesh size depends on the device geometry and extension. For full band – full spectrum simulation of photocurrent generation and transport in a nanostructured absorber, a huge number of large linear systems results, necessitating a diversified parallelization strategy for the different levels corresponding to the specific parameters (Fig. 8). Another challenge that influences directly the choice of viable parallelization approaches is the extensive coupling in energy and momentum space induced by

¹⁶ *The non-equilibrium Green's function picture of inelastic processes in nanostructure photovoltaics*, U. Aeberhard, Journal of Computational Electronics 15, 1219 (2016)

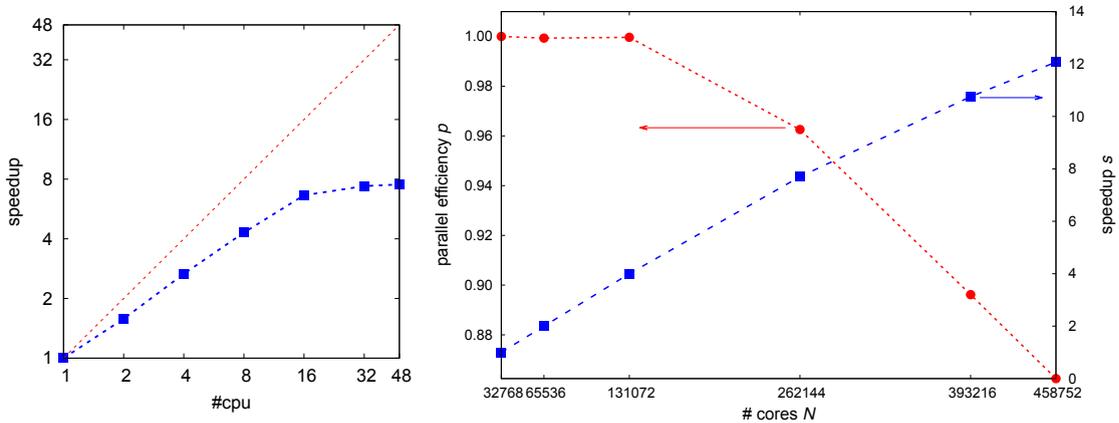


Figure 9: Left side: Speedup achieved on a single node of JURECA for the OpenMP version of the full *PVnegf* code. Right side: Speedup and parallel efficiency on the entire supercomputer JUQUEEN for the MPI implementation of the nested energy and momentum loops of the inelastic self-energy evaluation (\rightarrow computational bottleneck).

the inelastic scattering of electrons with photons and phonons. Finally, the nested loop structure of the algorithm itself increases substantially the need for communication, with the associated computational overhead. As a consequence of the difficulty of implementation, only a few massively parallel NEGF transport codes exist, and none designed for photovoltaic device simulation.

The original (serial) code is written in Fortran 90 and makes use of optimized BLAS and LAPACK routines for the solution of the discretized Dyson and Keldysh equations. The first step in the optimization of the code in the scope of the EoCoE project was to allow for shared memory parallelism by adding an OpenMP layer, distributing the evaluation of the transverse momentum loops on the different threads of a multicore machine. The left side of Fig. 9 shows the corresponding speed-up obtained on the Jülich supercomputer JURECA in the course of an EoCoE - supported performance optimization workshop at the Jülich supercomputing center (JSC).

In a second step, the core of the code including the inner self-consistency for inelastic scattering as described by a SCBA self-energy was parallelized in the frame of a Msc thesis at RWTH and the Simlab *ab initio* (now: Quantum Materials) at JSC¹⁷. Both pure MPI and hybrid parallelization was implemented and tested on the JURECA cluster as well as on the Blue Gene/Q machine JUQUEEN. Excellent scaling and parallel efficiency could be achieved for both architectures. The results for JUQUEEN are shown on the right side in Fig. 9. A paper containing these results is currently in preparation.

¹⁷ *Distributed parallel non-equilibrium Green's function approach to inelastic charge transport across Schottky junctions*, A. S. Achilles, Msc thesis, RWTH Aachen University, 2017

List of Figures

1 Snapshot of the initial a-Si:H/c-Si interface in the simulation box. The structure is infinitely extended in both x and y directions. A void region is considered to suppress the interaction, due to periodic boundary conditions, between the external surfaces. Free surfaces and a-Si:H/c-Si interface are perpendicular to the y axis. Hydrogen atoms are in blue, Silicon atoms are in orange in the c-Si side and are in yellow in the a-Si:H side. 7

2 a) Total wall time (s) for a self consistent calculations for a single energy minimization. b) Speed up calculated with respect to the time in the 48 cores calculation, and c) the corresponding efficiency (%). Red lines(symbols) are for CRESCO4, blue lines(symbols) are for CRESCO5 and green lines(symbols) are for JURECA. Dashed orange lines represent ideal parallel scalability. 8

3 Overview of the popular methods used in simulations of systems with atomistic resolution, showing the typical length scales over which they are applied as well as the degree of transferability of each method, i.e., the extent to which they give accurate results across different systems without re-tuning. On the left hand side we have the Quantum Chemistry methods which are highly transferable but only applicable to a few tens of atoms; on the right hand side we see the less transferable (semi-)empirical methods, which can however express reliable results (as they are parametrized for) for systems containing millions of atoms; and in the middle we see the methods—in particular linear-scaling DFT which can bridge the gap between the two regimes. The vertical divisions and corresponding background colours give an indication of the fields in which the methods are typically applied, namely chemistry, materials science, biology, and an intermediary regime (‘bridging the length scale gap’) between materials science and biology. The line colours indicate whether a method is QM or MM, while the typical regime for QM/MM methods is indicated by the shaded region. Some representative systems for the different regimes are depicted along the bottom. 10

4 Illustration of our implicit solvation approach using a fully continuous permittivity built with atomic-centered ‘soft’ spheres on the left. On the right panel, we show the experimental and *ab initio* aqueous solvation free energies for a set of 112 single-charged ions, of which 60 anions (empty blue circles) and 52 cations (full black circles), obtained with our method. 11

5 a) Calculated current-voltage curves without allowing for preconditioning; (b) measured current-voltage curves for cell 1 after 2 previous scan cycles and preconditioning for 5s at FB=1.2 V. Solid lines show the FB-SC scan; broken lines show the SC-FB scan. Scan rates are 1 V/s (magenta, circles), 500 mV/s (blue, crosses), 250 mV/s (cyan, filled squares), 100 mV/s (green, diamonds) 13

6	Number of three-center integrals involved through the different fitting schemes over the acene series, for a nonzero integral 10^{-9} cutoff value set to 10 Hartree. We compare here the Coulomb (RI-V) and density (RI-SVS) fitting schemes using the Cartesian cc-pVTZ-RI auxiliary basis, with the standard and constrained (up to fourth order momenta μ_4) RI-SVS scheme and the RI-V0 scheme using the augmented Cartesian aug-cc-pVTZ-RI auxiliary basis. The corresponding number of two-center Coulomb integrals between elements of the Cartesian cc-pVTZ-RI auxiliary basis is also reported. . . .	15
7	Computational procedure of the NEGF approach to photovoltaic device simulation. The computational meshes in real space (z), transverse momentum (\mathbf{k}_{\parallel}) and electronic energy (E) are determined by the device geometry and corresponding electronic and vibrational modes. The boundary self-energies Σ^B - encoding contact properties - are used in a Dyson equation to provide the retarded Green's function G_0^R of the noninteracting system. In an inner loop, interaction self-energies for inelastic scattering of electrons with phonons and photons are computed self-consistently with the retarded GF component G^R (related to the density of states) and the correlation functions G^{\lessgtr} (related to the density). An outer self-consistency loop couples the evaluation of the carrier density to the computation of the electrostatic potential via solution of Poisson's equation.	16
8	Levels of parallelism in the parameter space of the NEGF formalism for planar solar cell devices. While the level of the bias voltage is trivially parallelized due to complete independence of individual bias points, energy and momentum level parallelization suffers from the coupling due to inelastic scattering. In the current MPI implementation, these two levels are distributed together, while spatial domain decomposition is not considered at this point.	17
9	Left side: Speedup achieved on a single node of JURECA for the OpenMP version of the full <i>PVnegf</i> code. Right side: Speedup and parallel efficiency on the entire supercomputer JUQUEEN for the MPI implementation of the nested energy an momentum loops of the inelastic self-energy evaluation (\rightarrow computational bottleneck).	18