

Materials and Molecular Modelling at the Exascale

T. W. Keal

STFC Daresbury Laboratory

A. M. Elena

STFC Daresbury Laboratory

A. A. Sokol

University College London

K. Stoneham

University College London

M. I. J. Probert

University of York

C. S. Cucinotta

Imperial College London

D. J. Willock

Cardiff University

A. J. Logsdail

Cardiff University

A. Zen

University College London

P. J. Hasnip

University of York

I. J. Bush

STFC Rutherford Appleton Laboratory

M. Watkins

University of Lincoln

D. Alfè

UCL; Università di Napoli Federico II

C-K. Skylaris

University of Southampton

B. F. E. Curchod

Durham University

Q. Cai

University of Surrey

S. M. Woodley

University College London

The transition to exascale computing will make possible simulations of unprecedented accuracy and complexity. We focus on materials and molecular modelling (MMM) aspiring to high fidelity, *in silico* experiments on complex systems of technological interest. This progress will present unprecedented challenges to the software used, especially how to exploit the huge degree of parallelism and the associated problems of creating effective workflows and data management on such platforms. Within the UK's ExCALIBUR exascale computing initiative, our UK-led MMM Design and Development Working Group has worked with the broad MMM community to identify high priority applications, which will drive future exascale software developments. We present an overview of selected case studies that pose new

methodological challenges on exascale platforms and discuss the requirements, software challenges and impact of each application area.

Keywords: Materials Modelling; Exascale; Molecular Modeling; High Performance Computing; Research Software Engineering

■ **EXASCALE** supercomputing will enable new scientific endeavor in the materials and molecular modelling (MMM) domain. Members of the MMM community are investigating atomic and electronic structures with increasing realism to respond to major societal challenges – in areas including batteries, electrolytic cells, solar cells, computers, lighting and healthcare. We need to move beyond modelling phenomena with a single, well-defined length and time scale, and start to consider multiple processes, length- and timescales, including interactions between electrons, light, atomic vibrations and the environment. To enable these advances, there is a continual need to develop new methods and software to exploit increasing computational resources.

THE CHALLENGE OF MATERIALS AND MOLECULAR MODELLING

The chemical and physical properties of materials and molecules are determined fundamentally by their constituent electrons and nuclei. Modelling the detailed behavior of electrons requires quantum mechanics, with sub-nanometer (10^{-12} - 10^{-10} m) length-scales and attosecond (10^{-18} s) dynamics, but many important processes, such as corrosion, fracture and creep, occur on macroscopic length-scales visible to the naked eye (10^{-3} - 10^{-1} m), and over timescales of seconds or longer (10^0 - 10^5 s).

As can be seen in **Figure 1**, different modelling techniques are appropriate at different length and time-scales. At one end, electronic structure methods may be employed, with the most accurate techniques computing the many-electron wavefunction. Such calculations scale polynomially with simulation volume and *exponentially* with the number of electrons in the system, and so direct numerical solutions are impractical for most systems of interest; however, the exponential scaling can be avoided by employing sophisticated methods such as Coupled-Cluster or Quantum Monte Carlo (QMC). These methods represent the “gold standard” in terms of accuracy but are computationally expensive; current computers have only allowed limited simulations at these levels, for instance the adsorption of water on graphene [1]. Preparing the way for their routine use by

academic and industrial researchers in the UK is one of our aims.




Figure 1. Indicative length and time scales covered by different materials and molecular simulation techniques, including those key in electrochemical cell simulations.

An alternative quantum mechanical approach is density functional theory (DFT) which focuses instead on computing the electronic *density* – the number of electrons per unit volume at each point in space. DFT calculations are exact in principle and the approximations used in practice give good results for a wide range of systems. The low computational cost of DFT compared to other quantum mechanical models has made it one of the most popular and widespread MMM methods.

For processes that do not require a full description of electronic properties or reactivity, the behavior can be captured by parameterized interatomic potentials, removing the need for a quantum mechanical treatment altogether. These potentials are “atomistic”, i.e. functions of the positions of the atoms, which enable much more rapid calculations of atomic energies and forces. At even larger length- and time-scales, we might ignore interatomic interactions and employ coarse-grained (meso-scale) or continuum models (macro-scale); however, many important emerging applications

require both an accurate treatment of fast electronic and atomic processes *and* overall system dynamics on mesoscopic or macroscopic scales. The advent of exascale HPC can enable such studies for the first time.

We identify two potential simulation categories for the use of exascale facilities:

The Hero Run: applying a method at much larger length- or time-scales than is currently possible, using a significant fraction of available computing resources, for example to address large length-scale phenomena with high accuracy methods. The high computational requirement necessitates the use of scalable algorithms and the ability to make efficient use of evolving hardware; however, many of the applications used by the community rely on a complex interplay of algorithms (e.g. Fast Fourier Transforms, Matrix Diagonalization) that limit parallel scaling on current architectures; more efficient approaches are needed.

Complex Workflows: calculating a scientific quantity of interest through many individual runs that may vary considerably in resource requirements. Such simulations are essential for multiscale modelling, statistical ensemble sampling, global optimization, and uncertainty quantification. Using petascale resources, materials toolkits and workflow software developed by the UK MMM community (e.g. ChemShell and KLIMC [2]) are already employed in catalysis and structure prediction. Exascale computing will require increased automation, and job management must carefully balance available resources when running a mixed and varying workload, with efficient coupling of different applications and systematic error checking.

These two categories represent extremes of a spectrum of simulation types, but are useful archetypes when identifying fundamental issues that materials scientists will encounter on exascale machines.

MATERIALS AND MOLECULAR MODELLING COMMUNITY

The Materials and Molecular Modelling Software Design and Development Working Group (MMM DDWG) is one of 10 software DDWGs funded by the UK's ExCALIBUR research programme [<https://excalibur.ac.uk>], which aims to meet the challenge of preparing the UK for the arrival of exascale HPC. Each DDWG brings together a range of subject matter experts, including computational and mathematical scientists, and research software engineers (RSE).

The UK MMM community is composed of groups primarily based in Mathematics, Physics, Chemistry, Materials Science and Engineering departments, as well as Earth and Life Sciences. Within the DDWG several large consortia and networks are represented, including

Collaborative Computational Projects (CCP) and High End Computing (HEC) consortia, examples being CCP5++ (ccp5.ac.uk) for condensed matter simulations; CCP9 (ccp9.ac.uk), focused on the electronic properties of materials and molecules; CCP-NC (ccpnc.ac.uk) for NMR Crystallography; and CCP-QC (ccp-qc.ac.uk) for quantum computing. The HECs in the DDWG are the UK Car-Parrinello Consortium (UKCP; ukcp.ac.uk) and the Materials Chemistry Consortium (MCC; mcc.ac.uk), which foster collaboration between internationally recognized UK experts in MMM theory, applications and research software engineering. Both manage significant national HPC resources and have a strong track record in delivering world-leading research.

The UK MMM community has always been at the forefront of exploiting novel technologies and architectures, from early IBM mainframes to the current UK and international HPC facilities. For example, the MCC and UKCP accounted for over a third of the use of the previous UK national academic supercomputing resource (ARCHER) and publish over 250 research articles per year. Our communities also have access to the MMM Hub (mmmhub.ac.uk), which has hosted the “Thomas” and “Young” Tier-2 HPC facilities and runs a successful program of training and networking events. Researchers in the MMM domain have a strong track record of developing software, with numerous packages available and in use in both academia and industry worldwide. For instance, CASTEP [2] is developed in the UK by members of UKCP and used by over 800 academic and 500 industrial groups worldwide (including chemicals, pharmaceuticals, semiconductor manufacture, oil and gas, etc.). CRYSTAL [2], developed by UK researchers in collaboration with Turin University and others, is licensed to over 700 groups worldwide, including ~50 outside academia. For classical molecular dynamics DL_POLY [2] was used by over 5000 licensees in 2019. The MMM Hub, MCC and UKCP jointly created the MMM DDWG, with the aim of ensuring MMM software is ready for exascale.

Considering the potential of exascale computing, what scientific research can be done on the new machines, and how to prioritize areas with the greatest societal impact? From several open community meetings and surveys, the DDWG has identified recurrent themes and proposed UK MMM high priority use cases for Exascale HPC, which we discuss below. It is the need to address these areas that will drive the software development activities necessary for MMM research on exascale computing facilities to become a reality.

METHODOLOGICAL THEMES

In collaboration with the MMM UK community we have identified seven methodological themes, spanning all relevant length and time-scales and many different scientific areas and industrial sectors [3]. In this section, we begin by giving a summary of these themes, followed by further details of selected use cases.

1. Large scale simulations and sampling, without direct electronic structure calculations

Empirically fitted atomistic force fields can be employed to model systems at a scale unreachable by quantum mechanical methods. Avoiding calculation of the electronic structure enables efficient optimizations over large numbers of structures to find low-lying energy minima [4], or to run large scale molecular dynamics (MD) simulations with hundreds of millions of particles. Exascale resources will allow even larger scale simulations to be performed and enable increasing use of novel machine-learned (ML) force fields that require significant resources both in their training and application. Large scale classical simulations are used in structure prediction, phase diagram searches, phase transitions, radiation damage, materials failure at micrometer scales, crack propagation and protein-protein interactions, etc.

2. Long timescale and rare event modelling

The challenges of reaching long timescales are even greater than the challenges of increasing the length-scale, with 22 orders of magnitude between a typical simulation timestep (~1 fs) and slow mechanical and chemical processes such as cracking or corrosion (years). Moreover, parallelizing causality is inherently more difficult than parallelizing over system size. In some cases, a simple parallelization strategy is possible, as in replica MD, where parallel walkers can be used. Some methods, like kinetic Monte Carlo and temperature-accelerated MD, simulate rare events using transition-state theory, allowing an adaptive time sampling which focuses on the events themselves. This requires the calculation of the dynamic rates of many possible events. Methods of evaluating free energy surfaces (e.g. metadynamics, umbrella sampling) also are commonly used. Exascale challenges in this area will involve managing large numbers of calculations and processing the results, in applications such as solution chemistry, advanced composites and membrane neuroreceptors.

3. Systems beyond equilibrium: *ab initio* MD, ground and excited states

This theme covers dynamical simulations of complex systems at the nanoscale, requiring *ab initio* accuracy. Such studies face a wide range of challenges. For example, *ab initio* simulations require atomic length-scales to represent electronic states, yet many important dynamical processes, such as ion diffusion, take place over length-scales many orders of magnitude larger. Moreover, materials processes such as fracture dynamics take place on timescales of milliseconds (or longer), yet ultrafast electron dynamics take place on attosecond timescales, and even accurate atomic dynamics often require time-steps of 1 fs or less. Materials' behavior is commonly stochastic in nature, and many trajectories must be analyzed to generate a representative sample. These challenges often arise in combination, making the problems particularly suited to exascale resources. Applications include *ab initio* electrochemistry, green technologies, chemical fuel production, water-surface interactions including water pollution, design of biocompatible materials, excited-state dynamics of metal complexes, photocatalysis, and emitters in OLEDs.

4. Post-DFT methodologies

While DFT-based methods often represent a sweet spot between computational accuracy and expense for many MMM studies, there are scientific challenges for which the current DFT functionals are not sufficiently accurate. Post-DFT methods can be necessary for molecular adsorption at surfaces, which is often governed by weak interactions that DFT struggles to predict, and strongly correlated materials where higher-level *ab initio* theories are required to study problems such as unconventional super-conductivity. Exascale computing resources will make post-DFT theories, such as quantum Monte Carlo, more accessible.

5. Scaling up DFT

In principle, increased computational power allows larger physical systems to be studied. However, the ability to do this will be limited by the scalability of current codes; common algorithms such as the Fast Fourier Transform or Matrix Diagonalization have fundamental limitations in their parallel scalability, and scaling of memory usage of current codes is often a bottleneck. Another direction is linear-scaling DFT, such as in ONETEP and CP2K [2], where the computational effort increases linearly with the number of atoms, N , instead of N^3 as in conventional DFT. To be able to perform meaningful *ab initio* MD simulations we need to reduce the

time to solution from several hours (on tens of thousands of cores) to a few seconds, so we will need to re-engineer the codes to enable strong scaling to several million cores. Many areas of research would benefit from scaling up DFT, including extended defects, accurate simulation of biomolecular association events in proteins, enzyme catalysis; next generation energy materials, chemistry at interfaces in batteries, drug design, catalysis and electrochemistry.

6. First principles multiscale modelling (QM/MM and other hybrid methods)

Modelling approaches that combine high-accuracy localized calculations embedded in a more affordable representation of the long-range environment are a well-established means to describe complex systems with efficiency [5]. The high-accuracy region herein is treated through a range of QM methods, nowadays commonly DFT, with the embedding environment based on classical molecular mechanics (MM) force fields, or interatomic potentials. Increased computational resources offer the opportunity of more expensive embedding methods (e.g. high-level quantum chemistry embedded in DFT for greater accuracy) and to a wider range of systems via QM embedding with machine learning potentials. Models of interest can be solid, liquid or gas-phase, or their interfaces, and used to model properties, processes or reactivity. Areas of interest include catalytic reactions, energy storage, and drug discovery; structures of mesoporous materials; defects in semiconductors of commercial relevance; and materials failure through fracture or deformation.

7. Coarse-graining methods

Coarse-graining (CG) methods can be used to study problems in self-assembly and phase phenomena in soft/condensed matter systems, including polymers and nanomaterials. These complex states of matter occur in many real-life fields, including combustion, food processing, fundamental biological and physiological processes such as protein dynamics, and blood flow. Due to the length-scales in coarse grained models, a few orders of magnitude bigger than for traditional atomistic models (see Figure 1), new challenges arise due to the drastic approximations made in generating the models. Some level of uncertainty quantification (UQ) is required to understand and keep under control the errors introduced by the crude models. As a result, while CG models share a lot of the exascale challenges with the other themes, especially those in themes 1 and 2, the need for formal UQ stands it apart.

SELECTED APPLICATION CASE STUDIES

Modelling of *operando* electrochemical interfaces

Electrochemical transformation at the electrified interface (EI) between electrodes and electrolyte solution underpins many technologically and industrially critical processes and devices, from corrosion to neurotransmission, to fuel cells, supercapacitors, batteries, electrochemical sensors and memristors. *In silico* design and control with atomic and electronic resolution in *operando* conditions will enable massive energy and cost savings but has remained out of reach for almost all experimental techniques. Comprehensive computational modelling of *operando* electrochemical phenomena is still missing, even for simple interfaces.

Methodological developments will be required to enable realistic simulations of open systems undergoing electrochemical transformations on exascale HPC, to model fundamental electrocatalytic processes at EIs, including the response to changes in the potential applied to the cell, and understanding time resolved electron dynamics, including the effect of currents and applied potential on atomic dynamics, energy levels, forces, vibrations and redox transformation at the nanoscale.

Examples of electrochemical applications of huge technological relevance include:

Chemical fuel production: Important examples of electrochemical cells are electrolyzers (see Figure 1) that produce pure hydrogen out of water. In turn, environmentally clean and efficient hydrogen fuel cells convert the chemical energy of hydrogen and oxidants to electricity at the point of need. Achieving nanoscale insight into electrified solid/liquid interfaces and exploring factors that control reactivity or limit catalytic performance are essential to develop more active and longer lasting devices. The complexity of interfacial charge localization, dynamical redistribution and energy level alignment at EI, under operating conditions, require *ab initio* accuracy beyond conventional DFT in combination with large space- and time-scale *ab initio* MD, performed under bias and explicit open-boundary conditions.

Energy storage: The electrode-electrolyte interface layers play a key role in battery performance. Battery electrode materials must form stable interfaces with the electrolyte and allow easy diffusion of ions across them. Computational studies are essential for understanding the formation mechanisms and structures of the interphases and their thermal, mechanical, and ionic transport properties. Current *ab initio* MD simulations are usually performed for picoseconds, and focused on a single interphase compound,

yielding an incomplete and unrepresentative picture [6]. Future modelling must encompass larger, more realistic systems, a wider range of chemical phases and boundaries, and longer simulated time (e.g. nanoseconds), in order to fully capture reaction processes and pathways, and ultimately design improved electrode and electrolyte materials with enhanced battery performance and cycling stability.

Electrochemical sensing and memory applications: memory, sensing and logic devices applications, e.g. electrochemical metallization memories and gating devices, operate in the electrochemical environment under the combined effects of applied potential, current flow and chemical transformation. The relevant electrochemical processes underlying the functionality of these devices occur at the nanoscale, across molecular junctions or at the tiny EIS between ionic and electronic conductors. Achieving atomistic insight into these processes is key to their atom-based optimization.

Unravelling the thermodynamics and kinetics of non-equilibrium electrochemical transformation at the nanoscale requires accurate, multiscale simulations on large systems for long times, using a range of enhanced, multiple replica, sampling techniques. These simulations will need to use huge computational resources, so represent the ideal playground for exascale hardware. In this context, exascale solutions will be required for (i) coupling atomistic methods at one end with those for non-adiabatic electron dynamics, and at the other end with continuum approaches to account for the longest-range interactions; and (ii) the complexity of data, memory and tasks management, including problematic use of FFT at this scale, and the need for transferability of developed codes to multiple architectures. Suitable software for the developments includes CP2K and ONETEP and their seamless coupling to MD and computational fluid dynamics methods.

Thermal conductivity of liquid metals

There is a growing demand for environmentally-friendly energy production, and nuclear fusion is a promising candidate technology. Unlike nuclear fission, fusion does not require rare elements or produce long-lived radioactive waste, but the design and development of nuclear fusion reactors is an ongoing challenge. One leading design is the Tokamak, which uses magnetic fields to compress and contain a plasma and generate the conditions necessary for fusion.

A key challenge in the Tokamak design is finding material(s) for the container walls which can withstand the neutron bombardment. One popular suggestion is tungsten with a layer of liquid lithium behind it to act as a heat

exchanger. To design this, it is crucial to thoroughly understand the thermal properties of liquid lithium under typical Tokamak operating conditions. This need for detailed knowledge of liquid metals under extreme conditions goes far beyond what is practical in experimental studies, but it is amenable to computer simulation.

Quantum simulations are a natural route to getting reliable data, and at these high temperatures, an MD approach is the best way to handle the anharmonic effects. There are established algorithms for the simulation of thermal conductivity by MD [7], but the challenge is that the length and time scales required make quantum simulation very expensive. Exascale computing will be required to handle large quantum simulations with long time-scale MD.

The necessary MD simulations can be performed using CASTEP as it can do MD of metals already and has the thermal conductivity simulation algorithms. The trajectories will need to be analyzed to extract the ionic contributions to the conductivity, and snapshots of the electronic density of states can be used to calculate the electronic contributions.

To perform an accurate simulation of the thermal conductivity of a liquid metal requires many atoms to accurately describe the structural disorder, for a long simulation time to get the correct distributions and statistics, and with an accurate treatment of the interactions at high temperatures. Based upon preliminary studies of solid lithium with CASTEP, it appears to need at least 1000 and potentially 10,000 atoms in the simulation. The MD required will need to be at least 10 ps (i.e. 10,000 MD steps with a typical time step of 1 fs), and this will need to be repeated for each temperature of interest.

Such calculations using a conventional GGA-type approach to DFT will require a large resource (at least 50,000 cores) for a significant amount of runtime (at least 1 week per simulation) and be a challenge to scale. It remains to be seen whether a more accurate approach to the interactions is appropriate at these temperatures, such as time dependent-DFT or non-local functionals within DFT. If so, this would significantly increase the cost of each MD step.

One big gain would be to add parallel-in-time capabilities, such as the parareal algorithm [8], which requires coupling of an accurate but slow method (e.g. DFT) with a fast and less accurate method (e.g. machine-learned atomistic potentials). This would have the advantage of reducing run-time by increasing the core-count and would be a natural way to exploit an exascale resource to handle these large simulations efficiently. CASTEP has recently been coupled to the QUIP library for on-the-fly machine learning within molecular

dynamics, which should greatly increase the speed of these calculations, and if coupled with parareal, could significantly improve the scaling as well.

The insights gained from this project into the thermal properties of liquid metals will be a key input to future fusion reactor designs, including the first prototype “net power generation” nuclear fusion reactor which will be the successor to the ITER reactor currently being built in France.

Drug design from large-scale DFT calculations

Computational drug optimization is a key tool in the drug development process and focuses on accurate predictions of biomolecular association such as protein-ligand, protein-protein and other interactions with chemical accuracy. Rigorous methods rooted in statistical mechanics, such as thermodynamic integration, have been developed to compute the free energies of binding which determine the activity of a potential drug. These methods require extensive MD simulations to compute the “alchemical” transformation between ligands. While very successful, these methods fall in the accuracy of the underlying force fields which are often not able to describe the interactions of ligands with unsupported chemical functional groups. On the other hand, DFT is a promising alternative as it can describe these interactions with much higher accuracy, but only if the calculations can be made computationally feasible. Preliminary work on simple systems (hydration free energies of ligands in water) has shown that we can use alchemical transformations to augment the thermodynamic cycles of free energy calculations from the force field description to the DFT description, with the long-term aim of calculating free energies of binding entirely within DFT.

The ONETEP linear-scaling DFT program is ideal for these simulations as it has been developed to achieve routinely plane wave accuracy, it can treat entire proteins and other biomolecules with thousands of atoms, has a wide range of exchange-correlation functionals and MD capabilities and advanced solvation models, hence it combines all the physical models required. Examples of other suitable codes for such applications are CP2K and FHI-aims [2].

The main issue for scaling to the exascale regime is the time to solution (strong parallel scaling). Currently ONETEP can scale to thousands of cores for calculations with thousands of atoms which may take 10 minutes for a single point energy. The core operations of the code need to be rewritten (calculations of localized non-orthogonal Wannier function integrals, local potential and density, gradients of energy with respect to density matrix and Wannier function coefficients) so that it can scale to millions of cores and reduce the time to solution to fractions of a second to make these calculations

viable. This is a step change from current performance as even the “fastest” (in terms of parallel strong scaling) DFT codes take about 10 seconds per MD step on supercomputers.

These developments in DFT are led by the MMM community, but link with other communities and consortia in biosimulations. It will have a transformative impact on the research and market of drug development and will introduce an extra level of refinement in the “funnel” hierarchy of methods of computational drug development, saving billions of pounds from chemical syntheses and *in vivo* assay experiments in the lab. The developments that are needed to make this happen will also result in a code capable of simulating whole catalytic cycles, thus opening another area of large industrial impact, and electrochemical simulations of energy storage devices, linking with the *ab initio* MD and multiscale QM/MM themes.

Quantum Monte Carlo for accurate calculation and calibration of atomic adsorption energies

Molecular adsorption at surfaces is often governed by weak interactions, which determine fundamental material properties such as lubrication, filtering, gas sensing, or storage for energy purposes. Conventional computational materials science methods, including DFT, are not accurate enough to make reliable predictions for these types of systems, in general. For example, a material with good hydrogen storage properties should be able to bind hydrogen with an energy in the range ~ 200-400 meV/molecule, but the spread of binding energies obtained with various DFT functionals is often much larger than these binding energies, making the method unreliable. The quantum Monte Carlo (QMC) method has been demonstrated to be able to deliver useful accuracy (less than ~ 40 meV, and often much better than that), and so it is a prime candidate to be employed for these systems. Despite its much larger computational cost (at least three orders of magnitude over standard DFT), QMC is well suited to run on massively parallel computers, and some QMC codes are already able to exploit the power of accelerators such as GPUs. Future developments in the CASINO software package [2] will target exascale computing requirements. Exascale resources will allow the method to be deployed at scale, not just for single total energy calculations, but also to sample the phase space to gather thermodynamic properties. This will allow, for example, to obtain quantities such as the vapor pressure of a hydrogen gas in solid state allowing the study of confinement as a function of pressure and temperature, a crucial quantity to separate promising materials from nonviable ones. Parameter space sampling is an easily scalable (“embarrassing”) form of parallelization, and the difference between a petascale and an exascale machine is the

Department Head

difference between being restricted to zero temperature and gaining access to high temperature properties.

A method that is demonstrated to be predictive can become a tool complementary to experiments, with the advantage of being able to explore chemical space much more effectively. It will be possible to design new materials *in silico*, a promise that has yet to be fulfilled.

Predictive computational catalysis

Exascale computing will enable the accurate modelling of catalytic processes that serve vital societal and economic needs. Two key areas that must be addressed are CO₂ activation, where developments in catalytic technology are essential for achieving the net zero target, and nitrogen fixation, which underpins global nutritional security and currently relies on energy intensive technology. Computational modelling is required to develop next generation materials to catalyze these reactions and optimize their operating conditions, through development and application of *operando* models with high-level accuracy.

Computational methods already have a major impact on catalytic science, and multiscale methods are used to model problems such as CO₂ activation. Exascale computing will offer the opportunity for predictive and chemically accurate modelling of complex systems and their reactivity, providing both mechanistic insights in and kinetics of reaction processes. With such developments, computation will enable optimization of current reaction processes and design of new catalysts, inverting the current protocols of laboratory-led chemistry.

To make full use of future exascale resources, methodological improvements in multiscale modelling are required to enable the complex workflows that will become prevalent. Particular challenges include the need to run calculations with multiple QM regions; integration of adaptive (varying) QM regions in a synergistic manner, without compromising the discretization of the system; efficiently parallelized electronic structure code, both DFT and post-DFT methods; and “break-out” capability, i.e. for sampling-based approaches the ability to sample dynamically components of the system whilst core (QM) calculations are continuing.

The ChemShell multiscale QM/MM environment [2], which has a strong focus on performance on HPC systems, offers a suitable platform for the required developments. These need to deliver on three key aspects: (i) discretization and load-balancing for multiple separable QM regions, allowing for investigation of dynamic reaction environments with separable components; (ii) adoption of high-level

quantum mechanical techniques in the QM core; (iii) coupling QM/MM to higher length-scales, e.g. by incorporation of continuum methods into the models.

These requirements are based on the necessity to harness exascale computing to tackle complex, dynamic environments as encountered under catalytic operating conditions. Such models may contain catalyst, solvent, reagent and products, interacting dynamically, and therefore one needs high-level simulation management capable of handling an adapting QM region, especially given the inclusion of explicit solvent; and with capability to handle multiple QM regions. Such approaches need careful attention to the separability of the QM regions, and the overall embedding summations, but with exascale resources these new functionalities will present opportunities for broader and more accurate reaction sampling. They will be combined with other strategies to improve sampling, e.g., by averaging QM/MM calculations on much larger numbers of “snapshots” from MD simulations, to achieve reliable statistics.

Complementary to the objective of more dynamic multiscale models, at the core of such calculations is a desire for high-level accuracy. Immediate efforts will be towards coupling with quantum mechanical software specifically designed to maximize the potential of exascale computing (including those developed by other exascale initiatives such as the US Exascale Computing Project), as well as working with development teams towards software that is suitable for next generation hybrid architectures (e.g. GPU-enabled). The workflow management features of ChemShell will be further developed to ensure that the most efficient use of heterogeneous computing environments is made when calling external QM and MM software.

CONCLUSION

This paper outlines the myriad challenges and possibilities faced by the MMM community by the transition to exascale computing.

Investment in exascale computing is vital given the potential impact on wider society in contributing to fusion energy, combating pandemics, new pharmaceuticals, feeding the global population, and realizing net zero carbon emissions. To succeed we will need to transform rapidly our MMM software to accommodate the changing parallel programming software ecosystem and hardware landscape. As our methods rapidly evolve - following the four pillars set out by the UK’s ExCALIBUR project - we will have to prioritize:

(i) *Research Software Engineer training* in the area of high performance computing which is vital to ensure that there is a

sufficient skill base to tackle the software challenges outlined above, and is one key area the UK is actively pursuing [9].

(ii) *Separation of concerns* – our methods work with well documented APIs allowing multiscale problems to be tackled, as far as possible, in separate subprojects, but allowing inventive practitioners to splice them together to create new algorithms and solutions to MMM challenges.

(iii) *Co-design* – looking beyond our community as there are common bottlenecks, e.g. many communities currently perform Fast Fourier Transforms and/or diagonalize a matrix, which are not scalable to exascale HPC. Future parallel strategies are likely to include asynchronous task-based methods for compositional parallelization, as well as new ways to exploit heterogeneous HPC architectures.

(iv) *Data science* – there are a wealth of opportunities for using the data generated by the new facilities. It is essential that the data are reproducible and stored with appropriate metadata to allow data science domain experts to harness opportunities to use data for rapid materials design and optimization.

The driver for these activities is the science that will be achieved and the impact on society of the discoveries made. The impact already achieved by research that was enabled by the latest HPC has led to the desire for ever-increasing computing resources. There is a significant cost of running such facilities, so we have an additional responsibility of ensuring our software is “exascale ready” to maximize the science output per Watt consumed. Whilst we have emphasized the UK element of the research in this survey of our community, international collaboration will be essential to achieving these goals.

ACKNOWLEDGMENTS

We would like to thank all the participants in the materials and molecular modelling application case study workshops for their contributions to the full set of 19 MMM exascale case studies [3], including: Dmitry Nerukh, Gábor Csányi, James Elliot, Stewart Clark, Steven Hepplestone, Marco Molinari, James Ewen, Arash Mostofi, Laura Ratcliff, Carla Molteni, Dimitar Pashov, Mark von Schilfgaarde, Jamieson Christie, Paola Carbone, Sergey Karabasov, Kostya Trachenko, John Harding, Andrey Brukhno, John Purton, Vlad Sokhan, Ilian Todorov, Mozhdeh Mohammadpour, Jochen Blumberger, David Bowler, Richard Catlow, Peter Coveney, James Kermode, Nicholas Hine, Aires Ferreira and Keith McKenna. This work was supported in part by EPSRC grants (EP/V001507, EP/R029431, EP/P022561); and by CoSeC, the Computational Science Centre for Research Communities, through MCC. We acknowledge the support and involvement of our industrial and international project partners and collaborators on EP/V001507, namely Mike Dewar (NAG), Volker Blum (Duke University), Anibal Ramirez-Cuest and Ada

Sedova (Oak Ridge National Lab), and Andrea Ferretti (MAX centre of excellence).

REFERENCES

1. J. G. Brandenburg, A. Zen, M. Fitzner et al, “Physisorption of Water on Graphene: Subchemical Accuracy from Many-Body Electronic Structure Methods”, *J. Phys. Chem. Lett.*, vol. 10, no. 3, pp. 358-368, 2019. (journal)
2. A summary of MMM software packages can be found at <https://mmmhub.ac.uk/mmm-software/>
3. The full set of application case studies can be viewed online at: <https://mmmhub.ac.uk/mmm-themes/>
4. S. M. Woodley, G. M. Day, and R. Catlow. “Structure prediction of crystals, surfaces and nanoparticles”, *Phil. Trans. R. Soc. A*, vol. 378, no. 2186, pp. 20190600-1-23, 2020. (journal)
5. H. Lin, and D. G. Truhlar, “QM/MM: what have we learned, where are we, and where do we go from here?”, *Theor. Chem. Acc.*, vol. 117, no. 185, pp. 185-199, 2007. (journal)
6. A. Wang, S. Kadam, H. Li, A. Shi, and Y. Qi, “Review on modeling of the anode solid electrolyte interphase (SEI) for lithium-ion batteries”, *npj Comput. Mater.*, vol. 4, no. 15, pp. 1-26, 2018. (journal)
7. F. Müller-Plathe, “A simple nonequilibrium molecular dynamics method for calculating the thermal conductivity”, *J. Chem. Phys.*, vol. 106, no. 14, pp. 6082-6085, 1997. (journal)
8. Y. Maday, and O. Mula, “An adaptive parareal algorithm”, *J. Comput. Appl. Math.*, vol. 377, no. 112915, pp. 1-18, 2020. (journal)
9. M. Parsons, A. Basden, R. Bower et al, “ExCALIBUR Research Software Engineer Knowledge Integration Landscape Review”, <http://doi.org/10.5281/zenodo.4986062>, 2021.

Thomas Keal is a Principal Scientist specializing in the development and application of multiscale QM/MM methods. (thomas.keal@stfc.ac.uk)

Alin Elena is a Senior Computational Scientist specializing in computational statistical physics, MD and MC methods. (alinmarin.elena@stfc.ac.uk)

Alexey Sokol is a Principal Researcher specializing in methods for modeling materials. (a.sokol@ucl.ac.uk)

Karen Stoneham is the manager of the MMM Hub and Coordinator of the Thomas Young Centre. (k.stoneham@ucl.ac.uk)

Matthew Probert is a Chair of Computational Physics and a principal author of CASTEP. (matt.probert@york.ac.uk)

Clotilde Cucinotta is an EPSRC fellow (EP/P033555) specializing in electrochemical transformation phenomena. (c.cucinotta@imperial.ac.uk)

David Willock is a Reader of Chemistry specializing in materials chemistry and catalysis. (willockdj@cardiff.ac.uk)

Department Head

Andrew Logsdail is a Lecturer in Computational and Catalytic Chemistry, and UKRI Future Leaders Fellow (MR/T018372/1). (logsdail@cardiff.ac.uk)

Andrea Zen is an Assistant Professor specializing in developing and applying QMC to modeling materials. (andrea.zen@unina.it)

Phil Hasnip is an EPSRC Research Software Engineer Fellow and a principal author of CASTEP (phil.hasnip@york.ac.uk)

Ian Bush is a Principal Research Software Engineer and author of CRYSTAL and contributor to DL_POLY. (ian.bush@stfc.ac.uk)

Matt Watkins is an Associate Professor and contributor to CP2K. (mwatkins@lincoln.ac.uk)

Dario Alfe is Full Professor and interested in developing and applying quantum MC to modelling materials. (d.alf@ucl.ac.uk; dario.alf@unina.it)

Chris-Kriton Skylaris is Professor of Computational Chemistry and a principal author of ONETEP. (c.skylaris@soton.ac.uk)

Basile Curchod is Assistant Professor and specializing in developing and applying methods for nonadiabatic MD. (basile.f.curchod@durham.ac.uk)

Qiong Cai is a Senior Lecturer specializing in multiscale materials design for electrochemical energy conversion and storage applications. (q.cai@surrey.ac.uk)

Scott Woodley is a Professor of Computational Chemistry and Physics and lead of MMM Hub, MCC, and the MMM DDWG (EP/V001507). (scott.woodley@ucl.ac.uk)