

Materials and Molecular Modeling at the Exascale

Thomas W. Keal  and Alin-Marin Elena , STFC Daresbury Laboratory, WA4 4FX, Daresbury, U.K.

Alexey A. Sokol  and Karen Stoneham, University College London, WC1E 6BT, London, U.K.

Matt I. J. Probert , University of York, YO10 5DD, York, U.K.

Clotilde S. Cucinotta , Imperial College London, SW7 2BX, London, U.K.

David J. Willock  and Andrew J. Logsdail , Cardiff University, CF10 3AT, Cardiff, U.K.

Andrea Zen , University College London, WC1E 6BT, London, U.K.

Phil J. Hasnip , University of York, YO10 5DD, York, U.K.

Ian J. Bush , STFC Rutherford Appleton Laboratory, OX11 0QX, Oxfordshire, U.K.

Matthew Watkins , University of Lincoln, LN6 7TS, Lincoln, U.K.

Dario Alfe , University College London, WC1E 6BT, London, U.K.

Chris-Kriton Skylaris , University of Southampton, SO17 1BJ, Southampton, U.K.

Basile F. E. Curchod , Durham University, DH1 3LE, Durham, U.K.

Qiong Cai , University of Surrey, GU2 7XH, Guildford, U.K.

Scott M. Woodley , University College London, WC1E 6BT, London, U.K.

The transition to exascale computing will make possible simulations of unprecedented accuracy and complexity. We focus on materials and molecular modeling (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 U.K.'s ExCALIBUR computing initiative, our U.K.-led MMM Design and Development Working Group has worked with the broad MMM community to identify high-priority applications that 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.

Exascale supercomputing will enable new scientific endeavors in the materials and molecular modeling (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 modeling phenomena with a single, well-defined length- and time-scale, and start to consider multiple processes, length- and time-scales, 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.

1521-9615 © 2022 IEEE

Digital Object Identifier 10.1109/MCSE.2022.3141328

Date of publication 7 January 2022; date of current version 14 March 2022.

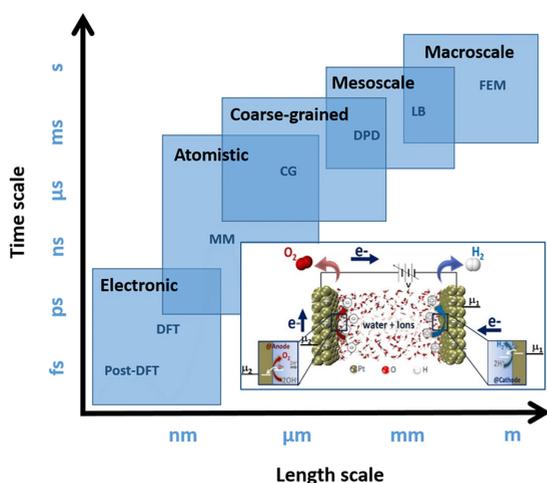


FIGURE 1. Indicative length- and time-scales covered by different materials and molecular simulation techniques, including those key in electrochemical cell simulations.

CHALLENGE OF MMM

The chemical and physical properties of materials and molecules are determined fundamentally by their constituent electrons and nuclei. Modeling the detailed behavior of electrons requires quantum mechanics, with subnanometer (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 time-scales of seconds or longer (10^0 – 10^5 s).

As can be seen in Figure 1, different modeling 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 U.K. is one of our aims.

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 (mesoscale) or continuum models (macroscale); 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 high-performance computing (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 modeling, statistical ensemble sampling, global optimization, and uncertainty quantification (UQ). Using petascale resources, materials toolkits and workflow software developed by the U.K. MMM community (e.g., ChemShell and KLMC²) 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.

MMM COMMUNITY

The MMM Software Design and Development Working Group (MMM DDWG) is one of ten software DDWGs funded by the U.K.'s ExCALIBUR research program,^a which aims to meet the challenge of preparing the U.K. 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 U.K. 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++^b for condensed matter simulations; CCP9,^c focused on the electronic properties of materials and molecules; CCP-NC^d for NMR Crystallography; and CCP-QC^e for quantum computing. The HECs in the DDWG are the U.K. Car-Parrinello Consortium (UKCP^f) and the Materials Chemistry Consortium (MCC^g), which foster collaboration between internationally recognized U.K. 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 U.K. MMM community has always been at the forefront of exploiting novel technologies and architectures, from early IBM mainframes to the current U.K. and international HPC facilities. For example, the MCC and UKCP accounted for over a third of the use of the previous U.K. national academic supercomputing resource (ARCHER) and published over 250 research articles per year. Our communities also have access to the MMM Hub,^h 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, CASTEP² was developed

in the U.K. by members of the UKCP and is used by over 800 academic and 500 industrial groups worldwide (including chemicals, pharmaceuticals, semiconductor manufacture, oil and gas, etc.). CRYSTAL,² developed by U.K. researchers in collaboration with Turin University and others, is licensed to over 700 groups worldwide, including ~50 outside academia. For classical molecular dynamics (MD) DL_POLY² was used by over 5000 licensees in 2019. The MMM Hub, MCC, and UKCP jointly created the MMM DDWG, with the aim of ensuring that MMM software is ready for exascale.

Considering the potential of exascale computing, questions exist such as 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 U.K. MMM high priority use cases for Exascale HPC, which we discuss in the following. 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 U.K. 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.

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 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.

^a[Online]. Available: <https://excalibur.ac.uk>

^b[Online]. Available: ccp5.ac.uk

^c[Online]. Available: ccp9.ac.uk

^d[Online]. Available: ccpnc.ac.uk

^e[Online]. Available: ccp-qc.ac.uk

^f[Online]. Available: ukcp.ac.uk

^g[Online]. Available: mcc.hec.ac.uk

^h[Online]. Available: mmmhub.ac.uk

Long Time-Scale and Rare-Event Modeling

The challenges of reaching long time-scales 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 and 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.

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 time-scales of milliseconds (or longer), yet ultrafast electron dynamics take place on attosecond time-scales, 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.

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 superconductivity. Exascale computing resources will make post-DFT theories, such as QMC, more accessible.

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.

First Principles Multiscale Modeling (QM/MM and Other Hybrid Methods)

Modeling 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 are 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.

Coarse-Graining (CG) Methods

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, and 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 those for traditional atomistic models (see Figure 1), new challenges arise due to the drastic approximations made in generating the models. Some level of 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 the first two themes, the need for formal UQ stands it apart.

SELECTED APPLICATION CASE STUDIES

Modeling 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 modeling of *operando* electrochemical phenomena is still missing, even for simple interfaces.

Methodological developments will be required to enable realistic simulations of open systems undergoing electro-chemical 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 the following.

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 into 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 interfaces and their thermal, mechanical, and ionic transport properties. Current *ab initio* MD simulations are usually performed for picoseconds, and focus on a single inter-phase compound, yielding an incomplete and unrepresentative picture.⁶ Future modeling must encompass larger, more realistic systems, a wider range of chemical phases and boundaries, and longer simulated times (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.

Unraveling the thermodynamics and kinetics of nonequilibrium electrochemical transformations at the nanoscale requires accurate, multiscale simulations on large systems for long times, using a range of enhanced, multiple replica, and sampling techniques. These simulations will need to use huge computational resources, so they represent the ideal playground for exascale hardware. In this context, exascale solutions will be required for 1) coupling atomistic methods at one end with those for nonadiabatic electron dynamics, and at the other end with

continuum approaches to account for the longest range interactions; and 2) 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 plasma and generate the conditions necessary for fusion.

A key challenge in the Tokamak design is finding material(s) for the container walls that 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 one 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 nonlocal 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., ML atomistic potentials). This would have the advantage of reducing runtime 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 MD, 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 associations, 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 nonorthogonal 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.

QMC for Accurate Calculation and Calibration of Atomic Adsorption Energies

Molecular adsorption at surfaces is often governed by weak interactions that 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 $\sim 200\text{--}400$ 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 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 package² 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 the 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 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 is yet to be fulfilled.

Predictive Computational Catalysis

Exascale computing will enable the accurate modeling 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 modeling is required to develop next-generation materials to catalyze these reactions and optimize their operating conditions, through the 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 modeling of complex systems and their reactivity, providing both mechanistic insights into 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 modeling are required to enable the complex workflows that will become prevalent. Particular challenges in QM/MM 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 while 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:

- 1) discretization and load-balancing for multiple separable QM regions, allowing for investigation of dynamic reaction environments with separable components;
- 2) adoption of high-level quantum mechanical techniques in the QM core;
- 3) 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 toward 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 toward 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 article 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 rapidly transform 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 U.K.’s ExCALIBUR Project—we will have to prioritize:

- (i) Research software engineer training in the area of HPC, which is vital to ensure that there is a sufficient skill base to tackle the software challenges outlined previously, and is one key area the U.K. 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 the 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 that our software is “exascale ready” to maximize the science output per watt consumed. While we have emphasized the U.K. element of the research in this survey of our community, international collaboration will be essential to achieving these goals.

ACKNOWLEDGMENTS

The authors would like to thank all the participants in the materials and molecular modeling 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. The authors would also like to acknowledge the support and involvement of their 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).

This work was supported in part by the Engineering and Physical Sciences Research Council under Grant EP/V001507, Grant EP/R029431, and Grant EP/P022561; and in part by CoSeC, the Computational Science Centre for Research Communities, through MCC.

REFERENCES

1. J. G. Brandenburg *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, doi: [10.1021/acs.jpcllett.8b03679](https://doi.org/10.1021/acs.jpcllett.8b03679).
2. A summary of MMM software packages, Accessed: Feb. 04, 2022. [Online]. Available: <https://mmmhub.ac.uk/mmm-software/>
3. The full set of application case studies, Accessed: Feb. 04, 2022. [Online]. Available: <https://mmmhub.ac.uk/mmm-themes/>
4. S. M. Woodley, G. M. Day, and R. Catlow. “Structure prediction of crystals, surfaces and nanoparticles,” *Philos. Trans. R. Soc. A*, vol. 378, no. 2186, 2020, Art. no. 20190600, doi: [10.1098/rsta.2019.0600](https://doi.org/10.1098/rsta.2019.0600).
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. Accounts*, vol. 117, 2007, Art. no. 185, doi: [10.1007/s00214-006-0143-z](https://doi.org/10.1007/s00214-006-0143-z).
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, 2018, Art. no. 15, doi: [10.1038/s41524-018-0064-0](https://doi.org/10.1038/s41524-018-0064-0).
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, doi: [10.1063/1.473271](https://doi.org/10.1063/1.473271).
8. Y. Maday and O. Mula, “An adaptive parareal algorithm,” *J. Comput. Appl. Math.*, vol. 377, 2020, Art. no. 112915, doi: [10.1016/j.cam.2020.112915](https://doi.org/10.1016/j.cam.2020.112915).
9. M. Parsons *et al.*, “EXCALIBUR research software engineer knowledge integration landscape review,” 2021, doi: [10.5281/zenodo.4986062](https://doi.org/10.5281/zenodo.4986062).

THOMAS W KEAL is currently a Principal Scientist with STFC Daresbury Laboratory, Warrington, U.K. His research focuses on the development and application of multiscale QM/MM methods. Contact him at thomas.keal@stfc.ac.uk.

ALIN-MARIN ELENA is currently a senior Computational Scientist with University College London, London, U.K. His research interests include computational statistical physics, MD, and MC methods. Contact him at alin-marin.elena@stfc.ac.uk.

ALEXEY A SOKOL is currently a principal researcher with University College London, London, U.K. His research focuses on methods for modeling materials. Contact him at a.sokol@ucl.ac.uk.

KAREN STONEHAM is currently a Manager of the MMM Hub and a coordinator of the Thomas Young Centre, University College London, London, U.K. Contact her at k.stoneham@ucl.ac.uk.

MATT I J PROBERT is currently a Chair of Computational Physics with the University of York, Heslington, U.K. He is a principal author of CASTEP. Contact him at matt.probert@york.ac.uk.

CLOTILDE S CUCINOTTA is currently with Imperial College London, London, U.K. Her research focuses on electrochemical transformation phenomena. She is an EPSRC Fellow (EP/P033555). Contact her at c.cucinotta@imperial.ac.uk.

DAVID J WILLOCK is currently a Reader of Chemistry with Cardiff University, Cardiff, U.K. His research interests include materials chemistry and catalysis. Contact him at willockdj@cardiff.ac.uk.

ANDREW J LOGSDAIL is currently a Lecturer in computational and catalytic chemistry with Cardiff University, Cardiff, U.K. He is also a UKRI future leaders Fellow (MR/T018372/1). Contact him at logsdaila@cardiff.ac.uk.

ANDREA ZEN is currently an Assistant Professor with Università di Napoli Federico II, Naples, Italy, and Honorary Research Fellow with University College London, London, U.K. His research focuses on developing and applying QMC to modeling materials. Contact him at andrea.zen@unina.it.

PHIL J HASNIP is currently an EPSRC Research Software Engineer Fellow with the University of York, Heslington, York, U.K. He is a principal author of CASTEP. Contact him at phil.hasnip@york.ac.uk.

IAN J BUSH is currently a Principal Research Software Engineer with STFC Rutherford Appleton Laboratory, Oxfordshire, U.K. He is an author of CRYSTAL and a contributor to DL_POLY. Contact him at ian.bush@stfc.ac.uk.

MATTHEW WATKINS is currently an Associate Professor with the University of Lincoln, Lincoln, U.K. He is a contributor to CP2K. Contact him at mwatkins@lincoln.ac.uk.

DARIO ALFÈ is currently a Professor of Physics with University College London, London, U.K., and Università di Napoli Federico II, Napoli, Italy. His interest focuses on developing and applying quantum MC to modeling materials. Contact him at d.alfè@ucl.ac.uk.

CHRIS-KRITON SKYLARIS is currently a Professor of Computational Chemistry with the University of Southampton, Southampton, U.K. He is a principal author of ONETEP. Contact him at c.skylaris@soton.ac.uk.

BASILE F E CURCHOD is currently an Assistant Professor with Durham University, Durham, U.K. His research focuses on developing and applying methods for nonadiabatic MD. Contact him at basile.f.curchod@durham.ac.uk.

QIONG CAI is currently a senior Lecturer with the University of Surrey, Guildford, U.K. His research interests include multi-scale materials design for electrochemical energy conversion and storage applications. Contact him at q.cai@surrey.ac.uk.

SCOTT M WOODLEY is currently a Professor of Computational Chemistry and Physics with University College London, London, U.K. He is also a lead of MMM Hub, MCC, and the MMM DDWG (EP/V001507). Contact him at scott.woodley@ucl.ac.uk.



www.computer.org/cga

IEEE Computer Graphics and Applications bridges the theory and practice of computer graphics. Subscribe to *CG&A* and

- stay current on the latest tools and applications and gain invaluable practical and research knowledge,
- discover cutting-edge applications and learn more about the latest techniques, and
- benefit from *CG&A*'s active and connected editorial board.

 **IEEE COMPUTER SOCIETY**

 **IEEE**