

High-performance computing for materials design to advance energy science

Mark T. Lusk and Ann E. Mattsson, Guest Editors

The development of new materials typically requires an iterative sequence of synthesis and characterization, but high-performance computing (HPC) adds another dimension to the process: materials can be synthesized and/or characterized virtually as well, and it is often an overlapping quilt of data from these four aspects of design that is used to develop a new material. This is made possible, in large measure, by the algorithms and hardware collectively referred to as HPC. Prominent within this developing approach to materials design is the increasingly important role that quantum mechanical analysis techniques have come to play. These techniques are reviewed with an emphasis on their application to materials design. This issue of *MRS Bulletin* highlights specific examples of how such HPC tools are used to advance energy science research in the areas of nuclear fission, electrochemical batteries, photovoltaic energy conversion, hydrocarbon catalysis, hydrogen storage, clathrate hydrates, and nuclear fusion.

Introduction

The future of modern society is tied to the availability of sustainable energy resources. While a global emphasis on conservation is certainly called for, there is a cultural imperative to create new ways to meet our energy needs (Please also see the April 2008 issue of *MRS Bulletin*). Positive responses to this mandate, in turn, will hinge on associated advances in materials design. This is where high-performance computing (HPC) plays an ever more valuable role. HPC in the energy sciences brings together materials researchers from such disparate fields as biology, fusion, inorganic chemistry, and photophysics to advance the computational interrogation of material structures (**Figure 1**). The common technical ground lies in the computational algorithms and hardware used to explore, design, and evaluate new materials and materials processes. The resulting discipline-bridging interactions have the potential to synergize new visions for how we might address the energy needs of our global society.

HPC relies on the development of algorithms that are parallelized (i.e., that can spread the computational workload out among a number of computer cores that are conducting calculations simultaneously) and that scale well. A code with perfect scalability, for instance, would run twice as fast on two cores than on one, and one thousand times faster on as many cores.

This is never the case in practice, but some types of algorithms, for instance those for which the overall task can be divided up into jobs that computer cores can work on with little communication, are much more scalable than others. Interestingly, most materials science papers involving HPC tend to dwell on the algorithms used rather than hardware when discussing their research. This does not imply that they have not spent significant time optimizing their code on a given computing platform. Rather, it reflects the fact that the codes are transferable between HPC systems, while the hardware may be very dissimilar. As a result, efforts made to optimize a code on a particular computing platform may not be of sufficiently general interest to the materials science community to be discussed in literature outside of computational science journals. This is certainly the case in this issue of *MRS Bulletin*, which showcases the HPC research of a number of leaders in the materials design community. Although mention is made of compute time requirements and the number of computer cores employed, the emphasis is on issues that are not hardware specific.

High-performance computing and first-principles algorithms

Perhaps correlating with research progress in the field, the role of algorithms associated with quantum mechanical calculations

Mark T. Lusk, Colorado School of Mines, Golden, Colorado 80401, USA; mlusk@mines.edu
Ann E. Mattsson, Sandia National Laboratories, Albuquerque, New Mexico 87185, USA; aematts@sandia.gov
DOI: 10.1557/mrs.2011.30

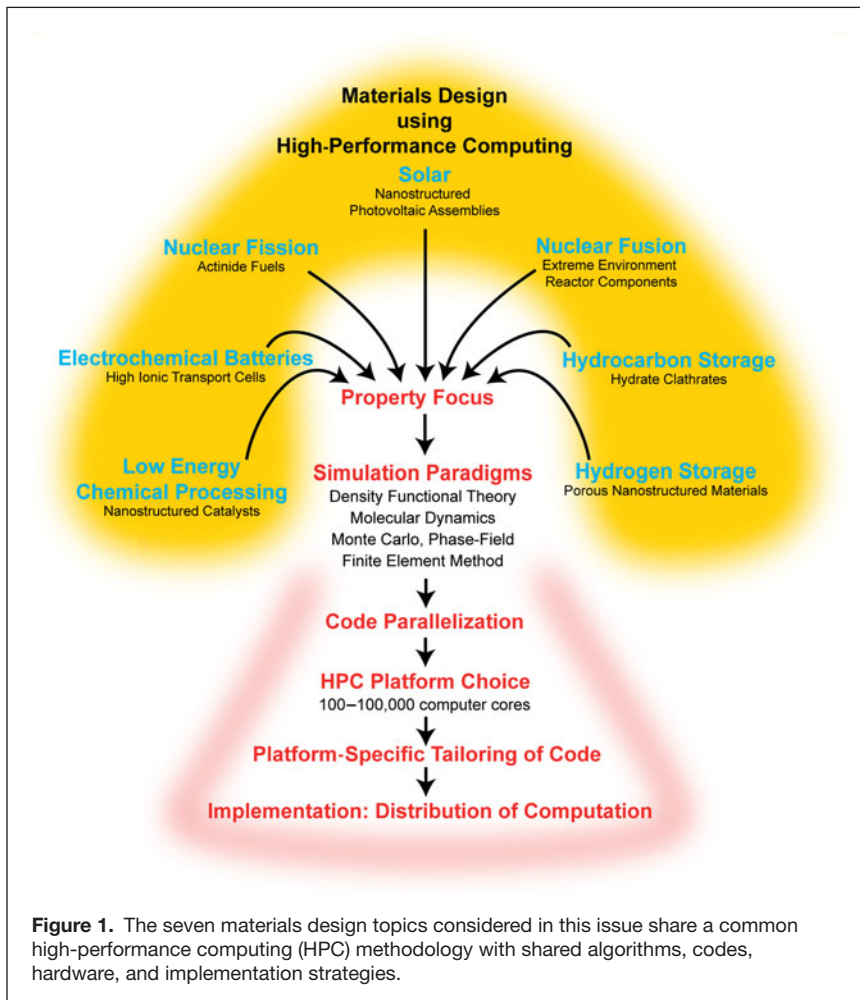


Figure 1. The seven materials design topics considered in this issue share a common high-performance computing (HPC) methodology with shared algorithms, codes, hardware, and implementation strategies.

continues to grow in the materials community. Although lofty sounding, it is a bit misleading to refer to these as *ab initio* approaches, since an exact solution of the underlying fundamental Dirac/Schrödinger equation (D/SE) is only possible for systems involving a small handful of electrons. Approximations are essential in solving these difficult many-body problems, and supplementary information is often brought to bear on the problem. On the other hand, the term first-principles is more reasonable, since our current understanding of the universe does not provide a more fundamental perspective of the physics involved. The common feature of all first-principles methods is that they seek to predict the spatial distribution and energies of the electrons that underlie the bonding between atoms.

The first-principles methods, which only require the types and approximate arrangement of atoms as input, are ideal for quantifying relationships between structure and properties (i.e., materials design). But there is a trade-off between accuracy/predictiveness and complexity/speed of the calculations. The direct solution of the SE is feasible only for systems involving very few electrons, and the most accurate systematically improvable quantum chemical methods, such as coupled cluster (CC) and configuration interaction (CI), can currently only treat small molecules.¹ Despite formally being a direct solution to the SE,

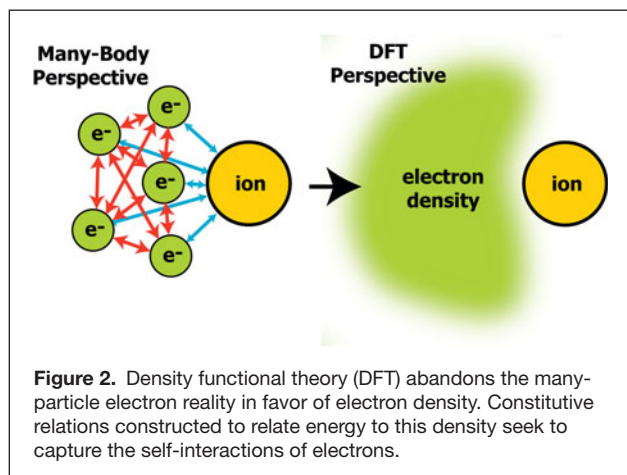
the numerical approximations needed in practical calculations currently render quantum Monte Carlo methods^{2,3} not accurate enough to compete with CC and CI, but future incarnations may remedy this, and the Monte Carlo paradigm can consider much larger systems than can either CC or CI. However, the currently preferred method of carrying out computational materials design is, without question, density functional theory.

Density functional theory

Density functional theory (DFT) offers a balance between accuracy and computational efficiency well suited to most materials science applications. The electron density only depends on the three spatial coordinates, while the many-body wave functions depend on all spatial coordinates of all electrons in the system. The electron density is thus a computationally simpler object, and this is the basis for the desirable computational properties of DFT (Figure 2). To completely reformulate the problem in terms of electron density, though, the Hamiltonian must be re-cast as a functional of this quantity. If such a divine,⁴ or exact, density-based exchange-correlation functional was known, the D/SE and the DFT equations would yield exactly the same (correct) answer for the ground state of a system. A roughly similar issue arises in classical fluid dynamics, where materials are characterized in terms of mass density instead of the position and momentum of every

atom, but where a direct tie between the energy description in the two paradigms is not easily ascertained. DFT amounts to treating the electrons as a fluid, and there are a host of approximations available that seek to give the best constitutive description of self-interactions of this fluid. These self-interactions are also known as exchange-correlation relations because they attempt to capture the dynamics of the correlated motion of electrons as well as the (exchange) forces that arise as a consequence of the indistinguishability of electrons. DFT currently allows quantitative data to be generated for systems on the order of one thousand atoms, often for materials in which each atom has electron counts on the order of one hundred. DFT even has a time-dependent extension that, with the correct accounting of electron-electron interactions, could precisely account for the excited state material properties.

As the internal energy is a directly accessible quantity in DFT, so are the forces between ions due to the electrons. This is increasingly used in molecular dynamics (MD) simulations. Here the usual methodology of calculating the forces between ions using interatomic potentials is replaced with a DFT calculation between each MD move. The increasing power of computers and the development of improved algorithms have facilitated a transition to these more accurate



quantum MD simulations, even though they still are limited to systems on the order of 500 ions. This methodology now plays a major role in providing improved equation of state tables for engineering finite element and finite difference codes⁵ by amending experimental data where such data are hard to obtain due to expensive and/or dangerous experiments and inaccessible experimental ranges. These DFT-based MD simulations are also crucial as benchmarks for classical inter-atomic potentials.⁶

DFT has two practical limitations. One is the need to approximate the exchange-correlation functional because the form of the divine functional is unknown. This is because the D/SE is formulated in terms of the many-body wave function, and there is no prescriptive method for deriving a re-formulation in terms of the electron density. The other is the limitation that properties calculated by DFT need to be formulated as functionals of the electron density. In addition to these fundamental limitations of the DFT method, various numerical approximations in codes can influence results. Some of these are discussed elsewhere.⁷

A significant effort has been made to develop new functionals that provide better and more universal accuracies. The two main areas where development is still needed are systems with heavy ions and those exhibiting physics governed by van der Waals interactions. The electron correlation in heavy electrons is notoriously complicated since *f*-electrons interact with both core and valence electrons, and this is exacerbated by the need to account for relativistic effects. However, overcoming these obstacles would allow DFT to be used in the design of nuclear fuels, for instance. The weak, long-range dispersion forces that comprise van der Waals interactions are associated with the time correlation of temporary dipoles setup within neighboring atoms. These are not accounted for in the standard DFT formulations,⁸ but their inclusion would enable the study of intermolecular interactions, catalytic processes, and the forces that neighboring nanostructures exert on one another.

There are two primary types of self-interaction functionals, classified according to what density-related ingredients

they use: local density approximations (LDA) and generalized gradient approximations (GGA). LDA functionals utilize only the local value of the electron density to estimate the exchange-correlation energy, while GGA functionals incorporate the gradient of the electron density as well. There are also meta-GGA functionals, where either the kinetic energy density or the Laplacian of the electron density is an added ingredient. Functionals belonging to these classes are preferred, so long as they are accurate, because they provide fast results. For example, they can be used in the DFT-based molecular dynamics simulations just described. The trend, however, is to use more complicated approaches if these local and semi-local functionals do not provide sufficient accuracy. This development is also mixed with a connected effort to improve predictions for properties not easily obtained from density functionals, such as electronic bandgaps. While the DFT equations together with an exchange-correlation energy density functional can be seen as equivalent to the fundamental D/SE, another valid interpretation of these equations is as a starting point for approximate many-body interaction treatments. One approach used extensively in chemistry applications is the hybrid method,⁹ where in standard DFT functionals are mixed with Hartree-Fock (HF) functionals. The advantage here is that the exchange energy is exactly accounted for within HF theory.⁹ Due to the desire for chemical accuracy, the weighting of HF and DFT components of these hybrid functionals is usually set by empirical means, mainly by determining parameters in them to optimally reproduce properties of sets of atoms and molecules calculated with a more accurate method such as CC or CI. In this way, accuracy is enhanced for these classes of systems, but it also means that transferability to other types of systems, not included in the fitting set, is generally poor (i.e., the approach interpolates well but extrapolates poorly). Exact-exchange methods, distinct from the more *ad hoc* HF-DFT approach, offer a more fundamental and computationally expensive quasi-particle approach to account for the electron self-interaction.¹⁰ Other DFT strategies are based on fully embracing the mean field view of DFT and adding empirical corrections derived from model many-body systems. One such approach is DFT+*U*, where a term derived from the Hubbard model¹⁵ is added in order to correct for deficiencies in DFT functionals to properly account for electronic localization effects. Dynamic mean field theory is yet another method of this type under intensive development. The next step beyond DFT is many-body perturbation theory, in which screened green functions are used to derive a nonlocal, quasi-particle equation that is at least an order of magnitude more expensive to solve.^{11,12} This is particularly important in the estimation of electronic bandgaps. A still more sophisticated analysis, based on the Bethe-Salpeter equation,¹³ can be used to account for the electron-hole interactions associated with optical excitations and exciton dynamics. Even HPC environments can only afford to consider systems on the order of one-hundred atoms with these quasi-particle and excitonic corrections, but this is still an improvement over CC and CI.

Modeling across multiple length scales

First-principles calculations, such as DFT, are ground floor investigations that may be of direct use in materials design. Often, though, the first-principles data are used to fit the classical potentials associated with molecular dynamics simulations. Quantum mechanical data are also used to inform mesoscale design tools such as kinetic Monte Carlo methods and phase field models, in which the atomic picture is traded in for a continuum perspective of material. At the macroscopic level, information gleaned directly from quantum mechanical analyses is combined with data from atomistic and mesoscales, as well as experimental data, in order to construct equations used to predict bulk thermal, mechanical, and electromagnetic character and the rates at which chemical processes and microstructural features evolve. These simulators sometimes take on the task of information generators for still higher scale paradigms in which geometry and the role of boundary conditions play a crucial role. Significantly, HPC is used at each of these levels. More extensive overviews can be found in the literature (e.g., Reference 14) (see **Table I**).

Simulations at multiple length scales are increasingly being used together to carry out hierarchical interrogations for materials design. The methodology typically involves a range of accuracies and scales that are often embedded within sophisticated optimization routines. Such hybrid methodologies, which have showed significant dividends in the pharmaceuticals industry, are now being developed in order to efficiently assess and screen a wide range of prospective materials prior to carrying out costly synthesis and testing. As with the first-principles calculations, these hierarchical analyses are computationally intense and are only viable within an HPC environment.

HPC is an enabler of all of these computationally demanding tools for investigating material systems. While we have focused on the types of computational analyses that may be carried out, each must be formulated into algorithms that allow problems to be solved by the collective effort of many computer cores working simultaneously. Inter-core communication, with its inherent latency, is to be minimized, and there is often a trade-off between such cross-talk and the local memory demands required of each core. Because of this, hardware platforms tend to be designed with particular types of codes in mind.

In this issue

This issue of *MRS Bulletin* offers a sampling of topics intended to give readers deeper insight into the ways in which advances

in computational materials science is helping to address the energy needs of our society. We consider the following seven topics:

- Nuclear fission fuels for harnessing nuclear binding energy
- Batteries for storing energy electrochemically
- Nanostructured photovoltaics for improving the efficiency with which solar energy is converted to electricity
- Porous, nanostructured materials for high capacity storage of molecular hydrogen to facilitate the use of hydrogen as a reversible, dispatchable means of storing energy in chemical bonds
- Hydrate clathrates, which need to be prevented from forming in oil recovery operations but also may be hiding huge reserves of energy in trapped hydrocarbons
- Nanostructured catalysts that have the potential to significantly reduce the energy cost of chemical processing, and
- Reactor materials that have yet to be designed but will be required for the successful exploitation of energy released in nuclear fusion.

The first article is an overview of first-principles design of nuclear fission fuels, given the lead-off position because the setting requires the most fundamental approach of any considered. In particular, the actinide elements comprising these fuels have massive nuclei, which cause electrons to move at speeds that are a substantial fraction of the speed of light. The relativistic behavior, in turn, can have a substantial influence on material properties. Together with the problem of treating the large number of electrons surrounding these nuclei, it gives rise to the so-called *f-electron* problem.¹⁴ The wide-ranging zone of influence of *f*-shell electrons precludes the straightforward adoption of a number of the computationally helpful approximations that are acceptable in materials composed of lighter elements. In addition, energy functionals that describe the self-interactions of electrons should be generalized to properly account for their dynamics within a relativistic framework. However, useful results can still be obtained for many properties of heavier elements using conventional DFT, if careful analysis is done where insights from calculations are paired with insights from other sources, such as experiments. Yun and Oppeneer describe the state of the art in predicting material properties, such as binding energies, lattice constants, electrical conductivity, magnetic character, and activation energies via first-principles analyses. A long-term goal is to use this approach to design next-generation nuclear fuels that are more reliable, tamper-resistant, and more easily recycled.

Attention is then turned to the design of batteries that use transition metal oxides to efficiently store and move lithium ions. During discharging, these ions detach from the lithium electrode, travel through a layer of electrolyte, and are intercalated into the crystal structure of the transition metal oxide cathode. This is where electrons, originating on the anode but traveling through the external circuit,

Table I. Current high-performance computing size scale limits for common computational methods.

Method	Interaction Level	System Size Scales
Density functional theory	Electron–electron, electron–ion	10 ³ atoms, 5 nm ³
Molecular dynamics	Atom–atom	10 ¹¹ atoms, 1 μm ³
Monte Carlo/phase field	Mesoscale continuum	1 cm ³
Finite element method	Macroscale continuum	1 m ³

reduce the oxide to form transition metal ions that are attracted to their lithium counterparts. The design objectives here are to create an electrolyte that can support the largest voltage, since this correlates with the energy capacity of the battery and a system that can be charged and discharged as quickly as possible. Ceder et al. overview this design approach and show how DFT can be used to craft LiFePO_4 cathodes with improved ionic transport properties. The rather complicated ion diffusion process was analyzed with a combination of density functional and Monte Carlo methods and was aided by a first-principles generation of phase diagrams. It is suggested that such an approach has the potential to identify materials that would allow battery charge and discharge times of only a few seconds. The authors also discuss the problems associated with the application of standard DFT methods to predict the electrochemical potential developed across electrolytes. This is akin to the tendency of the method to underestimate semiconductor bandgaps and stems from the inability of the method to account for the correlation between the ground-state electrons and new electrons added to the system.

Solar energy production is considered next, where the role of nanostructuring is explored as a means of improving the efficiency of photovoltaic energy conversion. The approach offers the prospect of exploiting the unusual behavior that results when material dimensions are so small that they constrain the extent to which quantum mechanical wave functions can spread out. A first-principles approach is considered here as well, but it is not a density functional technique. This is because, in addition to the bandgap problem identified in the previous article on batteries and its related shortcomings in predicting effective mass, another limitation of the commonly used ground-state DFT method becomes relevant: the method does not account for the interaction of photo-excited electrons with the holes that they leave behind. Since it is the separation of electron-hole-pairs (excitons) that generates electrical current, this is a serious deficiency, and the time-dependent extension of DFT that, in principle, would be able to deal with those excitons is not yet well developed. Although many-body perturbation techniques exist for such problems, they are computationally intractable for the many thousands of atoms in the nanostructures that Franceschetti is designing. He shows us how an alternative to DFT, using empirically fitted parameters, can be used to predict and design according to the behavior of such nanostructured materials.

The consideration of nanostructures using first-principles techniques is also central to the porous, nanostructured materials considered by Jhi and Ihm for high capacity storage of molecular hydrogen. Advances in this area are central to the success of an envisioned hydrogen economy in which molecular hydrogen is rapidly, reversibly, and safely stored in containers that can be easily and economically transported. In the research of Jhi and Ihm, carbon and silicon are nanostructured for high surface area and serve as a scaffold, or backbone, which is typically decorated with transition metal atoms. While carbon nanotubes have received a great deal of attention, Jhi and Ihm draw

attention to the potential of using fully hydrogenated graphene or graphane, the corrugated form of graphene that results when oxygen is allowed to bond to its surface. Non-transition metal functionalization is also a possibility. Computationally based design of hydrogen storage materials must be able to predict the weight percent of hydrogen adsorbed. This requires a calculation of the adsorption energy, and sometimes the dissociation energy, of hydrogen molecules as a function of loading, the rate of diffusion of hydrogen atoms that have become dissociated, the tendency of metal atoms to cluster and/or oxidize, and the spin state of the transition metal atoms since this influences the binding of hydrogen. Other articles in this issue of *MRS Bulletin* identify weaknesses with the density functional methodology, but here a new problem is identified: Current DFT functionals are not capable of correctly capturing the electron correlation that is responsible for van der Waals attraction, a key force in the design of hydrogen storage materials. In spite of this, a number of authors report what must be serendipitous success in accounting for such dispersion forces with functionals not intended to capture such physics.

Naturally occurring nanostructures can also be used to store both hydrogen and hydrocarbons. Cages created by hydrogen-bonded water molecules, for instance, can be used to trap such dispatchable fuels, and these are known as clathrate hydrates. The structure amounts to a rigid, ordered arrangement of nanometer-scale shells of water molecules that can each hold, in principle, a guest molecule of sufficiently small size. Clathrate hydrates can be a nuisance in oil and gas pipelines but are also a rich energy resource in extensive deposits beneath the deep ocean and in the permafrost—frozen soil found at extreme altitudes or latitudes such as in the arctic tundra. Sum et al. use molecular dynamics to elucidate the conditions under which methane hydrates nucleate and the rate of subsequent growth. This allows them to identify mechanisms that can either help or hinder the formation of the encapsulating hydrate cages. High-performance computing was used to examine a wide range of pressure, temperature, and cage/guest molecule stoichiometry conditions in a search that ultimately led to the direct observation of nucleation and growth events. Significantly, a number of other clathrate structures can also be synthesized, including those composed of silicon, germanium, and possibly even carbon. The methodology developed by Sum et al. shows us how to explore these nanostructures as well.

Hydrogen and hydrocarbons are also considered in the article by Mpourmpakis and Vlachos. Their focus, though, is on the computational design of catalysts that can improve energy efficiency in the industrial processing of chemicals, especially petrochemicals. They use first-principles analyses to consider ways of reducing the activation barrier for chemical reactions, but then go one step further by also designing for improved selectivity (i.e., designing catalysts that result in a larger proportion of a desired product being formed and not just a higher yield of all potential products). This conserves raw materials and reduces the demands on downstream separation

operations. Reminiscent of the computational design of pharmaceuticals, the idea is to construct a library of activation energies, assumed to be interpolative, for both desired and undesired reactions as a function of make up for bi-metallic catalysts. Mpourmpakis and Vlachos take one step further, though, by adding a dimension of molecular architecture to the library so that inhomogeneous noncatalyst designs can also be considered.

Our last article in this issue brings us back to nuclear power, but here the focus is on fusion. As with fission systems, computational materials design offers a means of exploring performance under conditions that are expensive and difficult to reproduce experimentally. Unlike fission reactors, though, there are operating conditions that cannot currently be produced in practice, and some of the demands of fusion reactors are beyond the capability of known materials. Wirth et al. emphasize the need to develop multiscale suites of computational design tools and focus on two materials design challenges associated with fusion reactors. They first consider the prompt chemical sputtering that occurs during the exposure of graphite to hydrogenic plasmas. This results in the low-energy erosion of carbon. The second example of fusion materials challenge taken up using HPC actually combines experimental investigation and computational multiscale modeling to investigate microstructural evolution in irradiated iron. At issue is the high rate of transmutations in the fusion energy environment that results in a significant helium concentration that interacts with point defects, which can lead to material failure.

Summary

In each of the articles in this issue, an effort is made to give a perspective on how high-performance computing is used to help design and characterize new materials. The works tend to emphasize the ways in which first-principles algorithms, abstracted onto highly parallelized computer codes and implemented on computers containing many computing cores, have changed the business of materials design in their respective disciplines. While each article stands on its own technical merit, we stress that the true worth of a compilation of such summaries is that the reader has the opportunity to make connections among material, performance goals, and the computational approach for materials design across a broad swath of energy science applications.

Acknowledgments

M.T.L., director of the Golden Energy Computing Organization at the Colorado School of Mines, would like to thank the Renewable Energy Materials Research Science and Engineering Center (REMERSEC) for its support and encouragement of this theme issue. Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin company, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

1. A. Szabo, N.S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (McGraw-Hill, New York, 1989).
2. B.L. Hammond, W.A. Lester, Jr., P.J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry* (World Scientific, Hackensack, NJ, 1994).
3. W.M.C. Foulkes, L. Mitas, R.J. Needs, G. Rajagopal, *Reviews of Modern Physics* **73**, 33 (2001).
4. A.E. Mattsson, *Science* **298**, 759 (2002).
5. S. Root, R.J. Magyar, J.H. Carpenter, D.L. Hanson, T.R. Mattsson, *Phys. Rev. Lett.* **105**, 085501 (2010).
6. T.R. Mattsson, J.M.D. Lane, K.R. Cochrane, M.P. Desjarlais, A.P. Thompson, F. Pierce, G.S. Grest, *Phys. Rev. B* **81**, 054103 (2010).
7. A.E. Mattsson, P.A. Schultz, M.P. Desjarlais, T.R. Mattsson, K. Leung, *Modelling Simul. Mater. Sci. Eng.* **13**, R1–R31 (2005).
8. P.L. Silvestrelli, *Phys. Rev. Lett.* **100**, 053002 (2008).
9. A.D. Becke, *J. Chem. Phys.* **98**, 1372 (1993).
10. W.G. Aulbur, M. Städele, A. Görling, *Phys. Rev. B*, **62**, 7121 (2000).
11. M. Shishkin, G. Kresse, *Phys. Rev. B* **75**, 235102 (2007).
12. D.J. Appelhans, Z. Lin, M.T. Lusk, *Phys. Rev. B*, **82**, 073410 (2010).
13. N. Nakanishi, *Progress of Theoretical Physics Supplement* **43**, 1 (1969).
14. R. Devanathan, L. Van Brutzel, A. Chartier, C. Guéneau, A.E. Mattsson, V. Tikare, T. Bartel, T. Besmann, M. Stan, P. Van Uffelen, *Energy Environ. Sci.* **3**, 1406 (2010).
15. J. Hubbard, *Proc. Roy. Soc. A* **276**, 238 (1963). □



Join the Conversation in the
Materials for Energy Blog

www.materialsforenergy.org

JANIS

MRS BOOTH 309

Cryogenic Systems



Does your research require low temperatures?
Contact Janis today. Our engineers will assist you
in choosing the best system for your application.

- 10 mK to 800 K
- Magnet Systems
- Cryocoolers
- Dilution Refrigerator Systems
- LHe/LN₂ Cryostats
- Micro-manipulated Probe Stations

Janis Research Company

2 Jewel Drive Wilmington, MA 01887 USA
TEL +1 978 657-8750 FAX +1 978 658-0349 sales@janis.com
Visit our website at www.janis.com