Current Opinion in Solid State and Materials Science 17 (2013) 298–304

Contents lists available at ScienceDirect

Current Opinion in Solid State and Materials Science

journal homepage: www.elsevier.com/locate/cossms

NSF cyberinfrastructures: A new paradigm for advancing materials simulation

Ellad B. Tadmor [a], Ryan S. Elliott [a], Simon R. Phillpot [b], Susan B. Sinnott [b,*]

[a] Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN 55455, United States
[b] Department of Materials Science and Engineering, University of Florida, Gainesville, FL 32611-6400, United States

ARTICLE INFO

Article history:
Available online 16 November 2013

Keywords:
Cyberinfrastructure
Materials modeling

ABSTRACT

This paper discusses the motivation for the creation of cyberinfrastructures to enhance specific technical areas of research. It then goes on to provide a review of two cyberinfrastructures supported by the National Science Foundation, OpenKIM and CAMS, which are geared towards enhancing materials modeling at the atomic scale. Their objectives, accomplishments, and future goals are discussed. Lastly, the future outlook for cyberinfrastructures such as these to impact materials modeling is discussed.

1. Introduction and motivation

In a recent National Science Foundation (NSF) Special Report on Cyberinfrastructure [1] "cyberinfrastructure" is defined as "a tightly integrated, planet-wide grid of computing, information, networking and sensor resources – a grid that we could one day tap as easily as we now use a power socket or a water faucet." Cyberinfrastructures are springing up in many areas of science and technology, and materials simulation is no exception.

Predictive materials modeling constitutes a key component of the U.S. Materials Genome Initiative (MGI) [2] introduced by the White House in 2011. The stated goal of MGI is to "discover, develop, manufacture, and deploy advanced materials at least twice as fast as possible today, at a fraction of the cost." Achieving this goal necessitates close integration between the traditional trial-and-error experimental approach to materials discovery and new computational techniques that have been developed in recent years. This is a challenging proposition given the inherent *multiscale* nature of materials [3]. In current-day engineering, the material is a "black box" whose response is characterized by a mathematical model fitted to experimental measurements. However, in reality a material is a complex dynamical system consisting of atoms (nuclei and electrons). There are about 10 sextillion (10^{22}) atoms in a gram of copper. The typical separation between atoms is of the order of 1 Å (10^{-10} m) and the characteristic time scale of atomic vibration is 1 fs (10^{-15} s). These atoms form patterns on many different scales between the atomistic and the macroscopic. For example a crystalline material such as copper is composed of a large number of single crystals (called grains) that are packed together to form the bulk material. A typical grain size in copper is 10 s of microns (100,000 times larger than the atomic separation). Grains contain large numbers of internal point, line and planar defects with different characteristic sizes and separations. This internal tapestry of structure, referred to collectively as *microstructure*, is what gives materials their properties. For example, as primitive humans discovered, mixing a bit of tin into copper leads to a new material called "bronze" which is much harder and tougher. This technological innovation was so important that this time in history is referred to as the "Bronze Age."

The internal complexity of materials is the reason that material design has been a largely experimental endeavor. With respect to material behavior, the basic laws of physics only provide information on the interaction of atoms through quantum mechanics and electromagnetic theory. The behaviors of materials at larger scales represent the coordinated actions of large numbers of atoms. In principle, one could predict the properties of a material by solving Schrödinger's equation, which describes the wavelike nature of particles in quantum mechanics, for all of the nuclei and electrons in the system. There are two main reasons why such a proposition is impractical. First, the solution of Schrödinger's equation is extremely computationally intensive. Depending on the level of desired accuracy, even with today's computational resources, the number of atoms that can be modeled ranges from dozens for quantum chemistry to a few thousands for the more approximate density functional theory methods. This is a tiny fraction of the number needed to model complex material microstructure. Second, the microstructure itself is typically unknown since it is a consequence of the manufacturing history of the material. For example, the microstructure in the bronze swords of the Bronze Age depends on the percentage of tin in the material and the repeated cycles of hammering, heating and cooling performed by the swordsmith.

Given the above, it is clear that a fully quantum mechanics simulation of materials is not feasible. Instead a variety of different

* Corresponding author. Tel.: +1 352 846 3778; fax: +1 352 846 3355.
E-mail address: ssinn@mse.ufl.edu (S.B. Sinnott).

1359-0286/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cossms.2013.10.004

theoretical and computational tools have been developed to model the response of materials at different scales [3]. Molecular dynamics (MD) using empirical (fitted) interatomic potentials is applied routinely to simulate millions and even billions of atoms on today's parallel supercomputers. Concurrent multiscale codes [4] couple atomistic regions with a surrounding continuum to reach macroscopic scales. Discrete dislocation dynamics (DDD) codes can be used to simulate the response of collections of dislocations (crystal line defects) that are responsible for irreversible deformation (plasticity) in crystalline materials [5]. Crystal plasticity is a continuum-level theory of materials that incorporates slip (plasticity) due to the motion of dislocations. Finally, continuum mechanics [6] is a framework for describing the response of materials using mathematical models called "constitutive relations" for describing the macroscopic response of the material. These theoretical and computational approaches coexist with experimental characterization tools of ever increasing resolution and precision. For example, it is now possible to generate fully three-dimensional representations of material microstructure using experimental tools such as the "slice-and-view" focused ion beam (FIB) and scanning electron microscopy (SEM) active 3D reconstructive method [7].

A successful materials cyberinfrastructure will be one that connects between the various theoretical, computational and experimental approaches described above and has a well-defined and focused purpose. Broadly speaking, the goal is to provide materials scientists and engineers easy access to a curated repository of useful materials data (collected from experimental, theoretical, and computational studies), easy access to materials simulation resources, and a significant amount of educational resources and opportunities. Materials data should be provided in standardized formats via flexible search and scriptable access methods. Materials simulation resources should support multiple levels of user expertise and provide robust and reliable results. A materials cyberinfrastructure cannot build and maintain a critical mass of users without providing detailed and up-to-date documentation, tutorials, and accurate and active user forums. The above characteristics are essential because they will allow for the creation of an organized and systematic method for materials researchers with expertise in a certain area to contribute their research results to the collective community. In this way researchers will be able to avoid "reinventing the wheel." Instead they will spend their time more productively by concentrating on the development of new materials theories, experiments, and computations. Further, the availability of such materials cyberinfrastructures will allow non-experts to, in a very short time, obtain (through the cyberinfrastructures educational resources) a working knowledge of the available materials data or simulation resources. This will allow the materials scientist or engineer to quickly incorporate these data and simulation resources into their own work while maintaining a high level of confidence in the quality and reproducibility of their research or designs.

It is clear from the above description that cyberinfrastructures are highly complex online frameworks. Fortunately, the ongoing internet revolution leads to the continuous development of new tools for people to remotely share content, interact through social media, and work more productively. These tools can be used to create cyberinfrastructure frameworks. Some examples drawn from the OpenKIM project described below include: ClojureScript, Django, Jekyll, jQuery, Markdown, Mozilla Persona, nginx, Riemann, SPDY, Tornado, Twitter Bootstrap, and WebSockets for interactive web applications; JSON, YAML, Mustache, and Jinja for data representation; Dataomic, jclouds, MongoDB, and Solr for data storage and query; Beanstalkd and ZeroMQ for message queues and service coordination; Vagrant and VirtualBox for virtual machines; Canvas, JavaScript, D3.js, HTML5, ThreeJS, and WebGL for web-based visualization; Prismatic Schema for data validation and integrity; Sphinx for documentation; Git for revision control; and languages such as Clojure, Ruby, and Python to tie it all together.

The goals of the MGI, to halve the time required to discover, develop, and deploy new materials, will only be achieved if the materials community is able to develop a broad and comprehensive system (e.g. a network of materials cyberinfrastructures with the characteristics described above) within which its members can freely share and exchange information in order to build on the community's collective knowledge and accelerate its progress.

In the next section we describe two current efforts at creating materials cyberinfrastructures that address the two main limitations described above of a quantum mechanics-based approach to materials design, namely the computational expense and unknown microstructure. The computational intensity of quantum calculations has led to the creation of empirical interatomic potentials. These are fitted computational models that mimic the behavior of collections of atoms without solving Schrödinger's equation. They are therefore much faster than quantum approaches, but suffer from a limited *transferability* (i.e. an ability to describe behavior that they were not fitted to reproduce). The OpenKIM (Knowledgebase of Interatomic Models) [8,9] cyberinfrastructure currently being developed as part of NSF's Cyber-Enabled Discovery and Innovation (CDI) program is an effort to standardize the use of interatomic potentials and to address the issue of transferability in a rigorous fashion. The Cyberinfrastructure for Atomistic Materials Science (CAMS) project [10] is a portal for information and education on methods for modeling materials at the atomic-scale. It also serves as a repository for storing material microstructures obtained from both experimental and computational sources. As such, it addresses the second limitation of quantum mechanics-based approach to materials design, namely the lack of information on material microstructure.

2. Examples of cyberinfrastructures

2.1. Knowledgebase of interatomic models (OpenKIM)

The OpenKIM framework [8,9] (see Fig. 1) is an interactive, self-extending, cyberinfrastructure of interatomic models, self-contained simulation codes that test the predictions of these models, and reference data. The "open" in OpenKIM reflects the open source nature of this project. OpenKIM also involves establishment of standards for the field in collaboration with the National Institute of Standards and Technology (NIST) [11]. OpenKIM has the support of many key players at national labs, industry and other universities in the U.S. and abroad. The KIM project is supported through NSF's CDI program. The specific goals of the OpenKIM project are:

- Development of an online open resource for standardized testing and long-term warehousing of interatomic models (potentials and force fields) and data.
- Development of an application programming interface (API) standard for atomistic simulations, which will allow any interatomic model to work seamlessly with any atomistic simulation code.
- Development of a rigorous, extendable standard and data structure for describing arbitrary material properties.
- Development of a quantitative theory of transferability of interatomic models to provide guidance for selecting application-appropriate models based on rigorous criteria, and error bounds on results.
- Planning for the permanence of the OpenKIM project, including development of a sustainability plan, and establishment of a long-term home for storing and maintaining its contents.

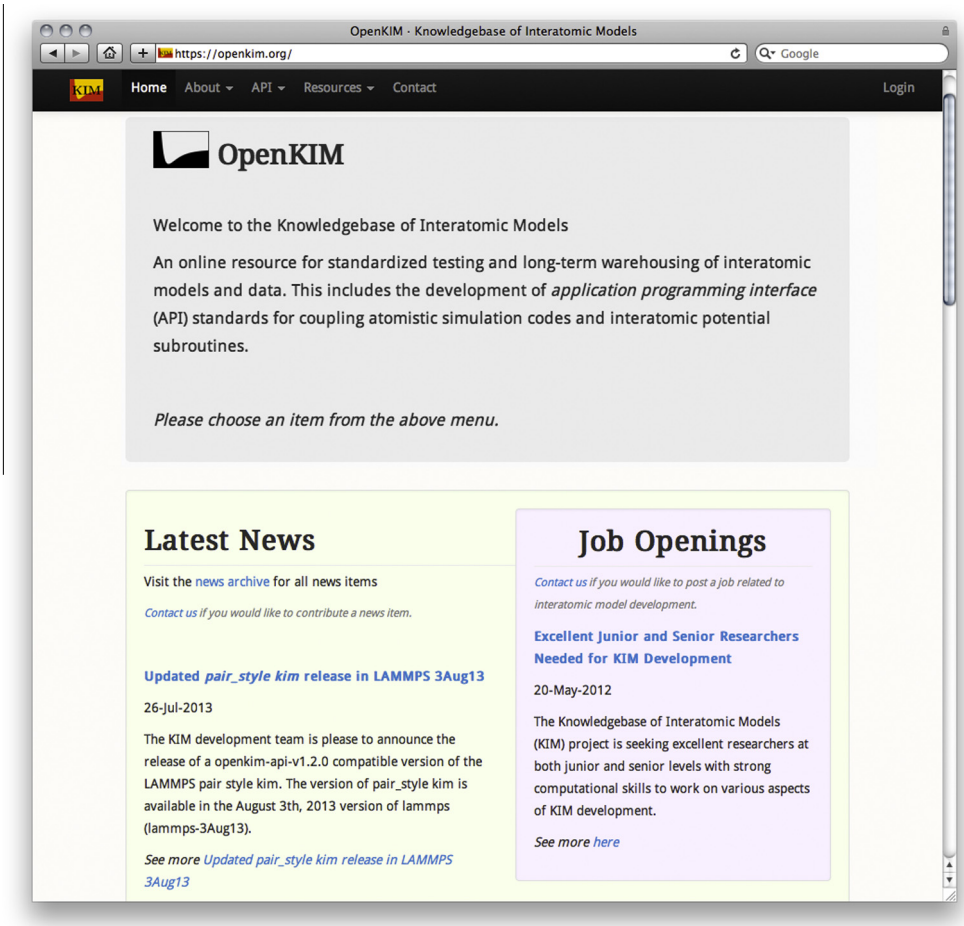

Fig. 1. Snapshot of the KIM homepage.

The OpenKIM cyberinfrastructure aims to answer the question: When and to what extent can we believe the results of atomistic simulations of materials? By "atomistic simulations" we mean simulations using approximate fitted interatomic potentials (or "Models" in OpenKIM terminology) as opposed to quantum simulations. The creation of the OpenKIM system will provide unprecedented, standardized, access to the state-of-the-art in atomistic modeling and simulation. This access will break down the barrier-to-entry to this field for traditionally underrepresented groups and institutions around the world and facilitate the efforts of industry in the U.S. and internationally to use interatomic potentials to advance their technological goals. OpenKIM will help to train the next generation of scientists and engineers by providing online tutorials and videos and conducting educational tutorials and workshops at popular materials conferences and other venues.

The OpenKIM project, from its beginning, has been focused on engaging the atomistic modeling community in order to develop a cyberinfrastructure that would be relevant, usable, and supported by the community. In this spirit, the first major accomplishment of the OpenKIM project was the development of the KIM "Requirements Document" (RD) [12], which was developed at the KIM Inaugural Meeting (held in San Diego, CA on Feb. 26,27, 2011). (This meeting was attended by invitation only by 63 participants from seven countries (Canada, Germany, Japan, South Korea, Sweden, UK and the US) and included many of the key interatomic potential and MD code developers. The objectives of the meeting were to clearly define the goals and structure of the OpenKIM project based on the needs of the atomistic simulation community, and to obtain the community's endorsement of OpenKIM. The event was highly successful on both counts. The KIM RD addresses the four main aspects of the project: (1) *governance* of OpenKIM including issues related to intellectual property rights; (2) the OpenKIM API; (3) the *web portal* for accessing OpenKIM content; and (4) *provenance and standards* issues needed to ensure the quality of OpenKIM content. The OpenKIM project currently has 378 registered members (July 2013).

A major accomplishment of the OpenKIM project has been the development and release of the "openkim-api" software package that implements the OpenKIM API for atomistic simulations. An API is an agreed-upon standard by which computer programs exchange information. An analogous example is the development of the DVD-Video format standard developed by the DVD Forum, an international consortium of hardware, software, media and content companies [13]. The existence of this standard is the reason that a DVD disc can be placed in DVD players made by different companies and work (modulo any zoning restrictions imposed for commercial reasons). In a similar manner, an interatomic potential written in any supported language (C, C++, FORTRAN 77, Fortran 90/95) and conforming to the OpenKIM API can be seamlessly run with any atomistic simulation code that supports this standard. The current package as this paper goes to press is "openkim-api-v1.2.1" (released July 2013). The latest version can be obtained from the KIM webpage [14] (see Fig. 2). The distribution includes the API code, examples of simulation codes for computing various properties (Tests), a sample collection of interatomic potentials (Models), and documentation. The OpenKIM API

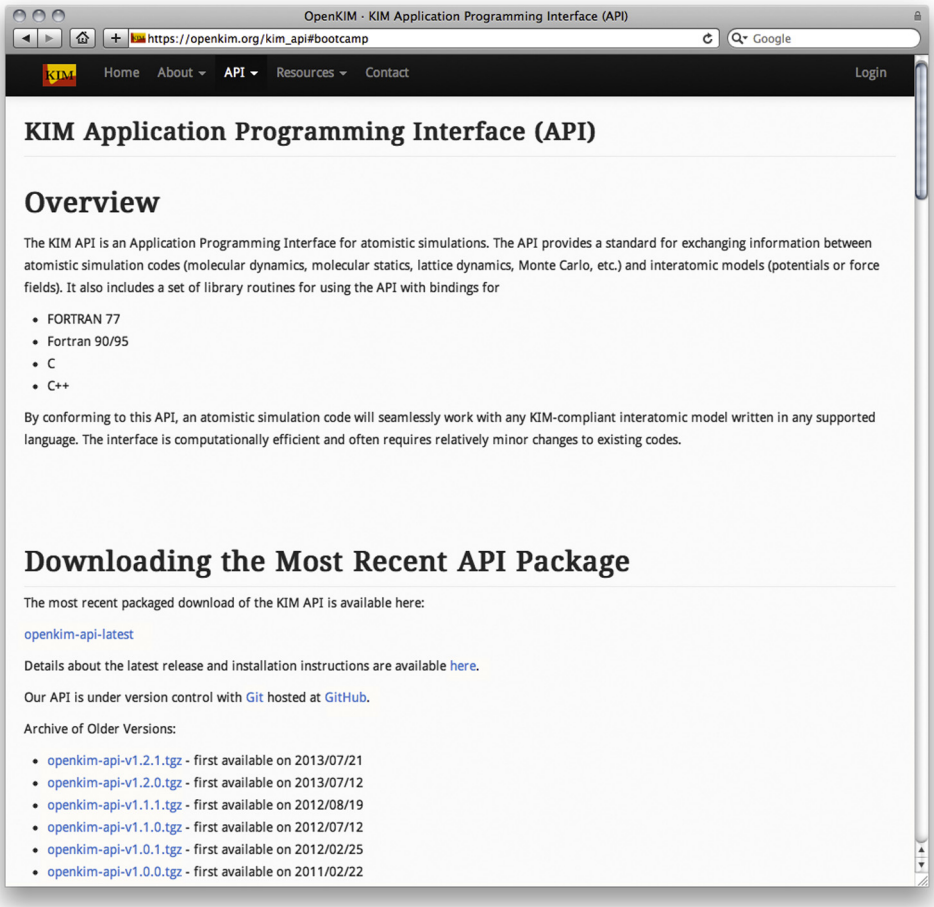

Fig. 2. Snapshot of the KIM API webpage.

is currently supported by the LAMMPS MD package [15,16], which is one of the key MD codes in use today. Other major atomistic codes are currently in the process of developing interfaces to support the OpenKIM API.

Another key development in the OpenKIM project is the definition of a robust extendible data structure for supporting arbitrary material properties. Recall that one of the objectives of OpenKIM is to extensively test interatomic potentials in order to develop quantitative measures for transferability (i.e. answer the question: Which potential is most suitable for a given calculation?). This is accomplished by including a host of simulation codes (referred to as "Tests") in the OpenKIM repository each of which computes a well-defined property of interest. Examples for crystalline materials include the equilibrium lattice constant, cohesive energy, elastic moduli, surface energies for various surfaces, phonon dispersion curves, and so on. A powerful aspect of this system is that users can upload new Tests to compute properties of interest that are not already in the OpenKIM repository. The calculations are carried out using a computational system called the KIM Processing Pipeline (KPP) that automatically matches compatible Tests and Models and executes the resulting simulation in order to generate a materials property prediction. An alpha version of this system, based on state-of-the art virtual machines and cloud computing technology, has been developed and is currently in testing. Storing model predictions in a way that is understandable to the OpenKIM system so that they can be compared with other calculations and reference data is highly challenging. A flexible extendable data structure for describing and storing material properties based on the idea of "material primitives" (distinct physical concepts associated with a property) has been developed and is currently undergoing testing.

The next step in the evolution of the OpenKIM cyberinfrastructure is the development of the KIM Visualization System (KVS). The KVS is an interactive suite of visualization tools for exploring the large and ever increasing amount of materials data contained in the OpenKIM repository. The KVS will allow users to compare the prediction of a specific Model to those of other Models and to experimental and first-principles reference data. An alpha version of this system has been developed and is currently in testing. As in all aspects of the OpenKIM cyberinfrastructure, KVS is designed to allow for users to extend the system by creating their own visualizers. It is planned to release these exciting new tools to the materials research community in late 2013.

All OpenKIM content including Models, Tests and reference data are given unique OpenKIM identifiers (IDs) (along the lines of the DOI [17]). In the future, it is hoped that the OpenKIM ID for a Model will be included in journal publications that make use of that interatomic potential. This is critical for the reproducibility of published results, in particular given the fact that interatomic potentials continuously mutate as they are refit to new problems. Within the OpenKIM framework this is addressed by storing the refitted potential as a new Model. The existence of OpenKIM IDs for Tests constitutes a de facto standard for computing certain properties. For example, there may be multiple Tests within OpenKIM for computing the melting temperature of a material using different methods. By referring to the OpenKIM ID of a specific Test one

defines the particular standard approach for doing the calculation that was employed. These efforts towards increased standardization and reproducibility of results echo similar developments in the biological community, for example with the establishment of the protein data bank [18].

For OpenKIM (or any cyberinfrastructure) to succeed it is necessary to get a critical mass of members from the community to participate. In an effort to broaden the impact of the KIM API and develop an OpenKIM user base, the OpenKIM project has been conducting a series of KIM Content Carnivals (KCCs). These are intensive four-day workshops targeted at developers of interatomic potentials and atomistic simulation codes. So far, three such KCC events have taken place (Minneapolis, MN, USA (3/2012); Aachen, Germany (8/2012); and Singapore (10/2012)). In all, 55 researchers have attended KCC events where they were given instruction on how to use the OpenKIM API and then worked to implement OpenKIM compliant versions their previously developed interatomic potentials. Throughout these events the OpenKIM leadership and development team were available to answer questions, provide suggestions, and gather feedback from the KCC participants. Additional KCC events will be held in the near future where, in addition to Model developers, Test developers will be invited. These events will help to further broaden the OpenKIM community and to develop a critical mass of Models and Tests within the OpenKIM cyberinfrastructure.

The above efforts to broaden the user-base of OpenKIM are vital for the long-term health of the project and its sustainability. Once OpenKIM becomes an integral part of the workflow for researchers involved in atomistic simulations and development of interatomic potentials and MD codes, it is expected that resources will become available for the continued funding of the project. For this to occur, the OpenKIM repository must reach a critical mass where adding content to OpenKIM will be perceived to be in the interest of the researcher (due to the knowledge gained by doing so) as opposed to being seen as a public service. As of October 2013, the OpenKIM repository contains 127 Models and 1694 Tests and is beginning to approach critical mass.

2.2. Cyberinfrastructure for Atomistic Materials Science (CAMS)

The objective of this cyberinfrastructure, which is in the pilot-project stage, is to benefit the materials science community by serving as a hub for the dissemination of information about atomistic scale and related modeling methods. The particular focus of CAMS is the application of these methods to the modeling of material microstructures.

The CAMS cyberinfrastructure [10] houses a collection of "virtual samples" of material microstructures. This actively curated collection involves structures of atomic coordinates of various microstructural features, such as grain boundaries and dislocations, for crystal structures (see, for example, Fig. 3). Some of these structures are defined purely crystallographically and would thus be generic starting structures for the simulation of various materials of a given symmetry. For example the structure of a grain boundary in fluorite could be used as a starting point for simulations of, amongst others, UO_2, CeO_2, ZrO_2, HfO_2, CaF_2. Other structures have been relaxed at zero temperature for specific materials (e.g., UO_2) described by a specific classical interatomic potential. The structural database contains information regarding the source of the structure, such as the registered user, whether and how the structure was constructed and relaxed, and, where appropriate, ionic charges and other atomic-scale information. CAMS also contains links to other databases of microstructures to enable sharing and reduce duplication. This virtual sample collection will accelerate future atomistic simulations of material microstructures by allowing users to bypass the common, though often difficult, step associated with building these structures by hand. Such structures can thus be used for a range of studies to enable the systematic exploration of a range of properties of a number of well-defined microstructures. A citation structure similar to a DOI [16] is in place that allows the user who uploads microstructures to be cited by others when they make use of the structure in future studies. In the future, CAMS plans to additionally host or link to software for the generation or analysis of material microstructures.

Another way in which CAMS disseminates information is through education. For example, in May of 2013 CAMS hosted a summer school on "Simulation of Complex Microstructure in Materials" at the University of Florida. It was attended by over 50 participants from 20 different U.S. institutions and one German university. The participants included a mix of graduate students, undergraduate students, postdocs, and junior faculty. Experts in diverse aspects of the modeling and characterization of material microstructure from two U.S. National Laboratories, one German university, and multiple U.S. academic institutions were instructors for the school; copies of their presentations, along with video of their lectures, are accessible to registered users of CAMS. Travel assistance was provided by CAMS to both the participants and the instructors to ensure broad participation.

The summer school included opportunities for the participants to discuss the lectures in small groups and develop in-depth questions. Participants had the opportunity to present research results from their home institution in a poster session. In addition, time was devoted to professional development activities, such as a questions and answer session on what it is like to be a junior staff person at a national laboratory or a junior faculty member at an academic institution, and a presentation on entrepreneurship. The participants also worked together in teams to develop research proposals based on the topics presented at the school to provide them with experience in translating scientific concepts to proposals; awards were given for best posters and best proposals. A second summer school is planned for May of 2014 at the University of Wisconsin, Madison, which will be hosted cooperatively with other NSF-funded projects, including a Materials Research Science and Engineering Center (MRSEC).

Future activities will depend on the ability of CAMS to evolve from the pilot stage to a more expansive cyberinfrastructure. Its long-term sustainability will thus depend on the ability of CAMS to meet the needs of the materials community, the involvement of this community in its progression, the support of funding agencies, and the availability of resources to sustain its development.

3. Summary and future outlook

As the examples above illustrate, cyberinfrastructures are poised to make substantial impacts on the area of materials modeling and have already made great strides in this direction. While the cyberinfrastructures featured here are distinctly different from one another, they share several common features, including initiatives to catalog, curate, and disseminate user-developed content in a manner that gives the developer credit, removes ambiguities from the literature, and makes it easier for the non-expert to navigate the field. Both cyberinfrastructures have also taken an active role in the dissemination of knowledge to the next generation of scientists.

Other cyberinfrastructures devoted to materials modeling have a complementary yet distinctly different focus. For example, the objective of the Materials Project [19] at MIT is to apply first-principles, density functional theory calculations in combination with other computational tools to determine the properties of materials for energy storage applications, such as batteries and solar cells. Users are able to search through property databases or compute

Fig. 3. Snapshot of the CAMS microstructure library and an illustrative sample.

properties such as solid-state reaction enthalpies, phase diagrams, crystal structures under specified conditions, and aqueous stability. Another example is the AFLOWLIB.ORG cyberinfrastructure, a consortium of material data and software [20,21]. Databases are available with a range of material properties, including structural, electronic, thermoelectric, and scintillator, along with a database of structural properties for metal alloys predicted from first-principles calculations. An additional attractive feature allows users to run density-functional theory calculations through a user interface located on the infrastructure.

The importance of materials modeling cyberinfrastructures is only expected to increase as the use of first-principles calculations and atomic-scale simulations becomes more widespread. Contributing to their increased importance is improvement in internet connectivity speed, the ubiquitous nature of electronic devices that facilitate greater access to cyberinfrastructures, and rapid increases in computing power coupled with decreases in their costs.

In the future it is envisioned that the impacts of current cyberinfrastructures will be magnified by linking them together. In some cases, this may involve a central portal to direct users to the appropriate site in a manner similar to the Materials Innovation @TMS website. In other cases, the linkages may seamlessly exploit the capabilities of each site. For instance, a user may download a material microstructure from CAMS and begin examining it in a simulation using a potential from the OpenKIM site without having to visit each cyberinfrastructure separately.

The vision and support of the National Science Foundation has proved to be critical to launch and establish cyberinfrastructures to multiple fields, including the area of materials modeling. Ensuring that these entities live up to their potential will require the support of the computational community in addition to the efforts of the dedicated individuals who establish and maintain them. It is clear that there is tremendous interest on the part of the community in taking advantage of these resources to minimize duplication of effort and the barriers to atomic-scale modeling by non-experts. Thus, the future outlook is bright for the use of the tools of cyberinfrastructure with materials modeling.

Acknowledgements

EBT and RSE acknowledge the support of the National Science Foundation (PHY-0941493). SRP and SBS acknowledge the support of the National Science Foundation (DMR-1246173). The OpenKIM project is an effort of many people. EBT and RSE acknowledge Prof. James P. Sethna (co-PI on the OpenKIM project) and the OpenKIM team members, Alex Alemi, Matt Bierbaum, Yangiun Chen, Daniel Karls, Adam Ludvik and Trevor Wennblom, for contributing to all aspects of the OpenKIM project.

References

[1] Hart D. NSF special report on infrastructure. National Science Foundation; June 28, 2006. <http://www.nsf.gov/news/special_reports/cyber>.
[2] Materials genome initiative for global competitiveness. U.S. office of science and technology policy (OSTP); Washington, USA: 2011. <http://www.whitehouse.gov/mgi>.
[3] Tadmor EB, Miller RE. Modeling materials: continuum, atomistic and multiscale techniques. Cambridge University Press; 2011.
[4] Miller RE, Tadmor EB. A unified framework and performance benchmark of fourteen multiscale atomistic/continuum coupling methods. Model Simul Mater Sci Eng 2009;17:053001.
[5] Bulatov VV, Cai W. Computer simulations of dislocations. Oxford University Press; 2006.
[6] Tadmor EB, Miller RE, Elliott RS. Continuum mechanics and thermodynamics: from fundamental concepts to governing equations. Cambridge University Press; 2012.
[7] Cao S, Tirry W, van den Broek W, Schryvers D. Optimization of a FIB/SEM slice-and-view study of the 3D distribution of Ni_4Ti_3 precipitates in Ni–Ti. J Microsc 2009;233(1):61–8.
[8] Tadmor EB, Elliott RS, Sethna JP, Miller RE, Becker CA. The potential of atomistic simulations and the knowledgebase of interatomic models. JOM 2011;63:17.
[9] Knowledgebase of Interatomic Models (KIM). <http://OpenKIM.org>.

[10] Cyberinfrastructure for Atomistic Materials Science (CAMS). <http://cams.mse.ufl.edu>.
[11] Becker CA. NIST interatomic potentials repository project. <http://www.ctcms.nist.gov/~cbecker/>.
[12] KIM requirements document <http://OpenKIM.org/requirements>.
[13] DVD forum. <http://www.dvdforum.org>.
[14] OpenKIM-api software package. <http://OpenKIM.org/kim_api#download>.
[15] Plimpton S. Fast parallel algorithms for short-range molecular dynamics. J Comp Phys 1995;117:1–19.
[16] Large-scale atomic/molecular massively parallel simulator (LAMMPS). <http://lammps.sandia.gov>.
[17] The DOI system. <http://www.doi.org>.
[18] Bernstein FC, Koetzle TF, Williams GJB, Meyer Jr EF, Brice MD, Rodgers JR, et al. The protein data bank: a computer-based archival file for macromolecular structures. J Mol Biol 1977;112:535–42.
[19] Jain A, Ong P, Hautier G, Chen W, Richards WD, Dacek S, et al. The materials project: a materials genome approach to accelerating materials innovation. Appl Phys Lett Mater 2013;1:011002.
[20] Setyawan W, Curtarolo S. High-throughput electronic structure calculations: challenges and tools. Comp Mat Sci 2010;49:299–312.
[21] Curtarolo S, Setyawan W, Wang S, Xue J, Yang K, Taylor RH, et al. AFLOWLIB.ORG: a distributed materials properties repository from high-throughput ab initio calculations. Comp Mat Sci 2012;58:227–35.