Informatics Infrastructure for the Materials Genome Initiative

ALDEN DIMA,^{1,3} SUNIL BHASKARLA,¹ CHANDLER BECKER,¹
MARY BRADY,¹ CARELYN CAMPBELL,¹ PHILIPPE DESSAUW,¹
ROBERT HANISCH,¹ URSULA KATTNER,¹ KENNETH KROENLEIN,²
MARCUS NEWROCK,¹ ADELE PESKIN,² RAYMOND PLANTE,¹
SHENG-YEN LI,¹ PIERRE-FRANÇOIS RIGODIAT,¹ GUILLAUME SOUSA
AMARAL,¹ ZACHARY TRAUTT,¹ XAVIER SCHMITT,¹ JAMES WARREN,¹
and SHARIEF YOUSSEF¹

1.—National Institute of Standards and Technology, Gaithersburg, MD, USA. 2.—National Institute of Standards and Technology, Boulder, CO, USA. 3.—e-mail: alden.dima@nist.gov

A materials data infrastructure that enables the sharing and transformation of a wide range of materials data is an essential part of achieving the goals of the Materials Genome Initiative. We describe two high-level requirements of such an infrastructure as well as an emerging open-source implementation consisting of the Materials Data Curation System and the National Institute of Standards and Technology Materials Resource Registry.

INTRODUCTION

New technologies are often limited by currently existing materials because the time to develop and deploy new materials generally exceeds the product design cycle. For example, it takes approximately 2 years to design a new jet engine using available materials, but it may take 10–15 years to design and certify the new materials needed for the engine.¹ Integrated computational materials engineering (ICME) approaches have proven successful at decreasing this gap between the materials development cycle and product development cycle,² but these approaches are not well developed for all classes and applications of materials, and there is a critical need for materials data and modeling tools that further enable these approaches.

To address the need to decrease the time and cost to develop and deploy new materials by 50%, President Obama announced the Materials Genome Initiative (MGI) in 2011.³ The MGI recognizes that advanced materials play a critical role in clean energy, human welfare, and national security. It is a multiagency initiative that focuses on the infrastructure needed to accelerate materials development, particularly in the following areas: (I) Computational Tools, (II) Experimental Tools, (III) Collaborative Networks, and (IV) Digital Data.

By facilitating the integration of data into developing ICME approaches and other computational approaches to materials discovery, design, development, and deployment, a materials data infrastructure that allows the wide range of materials data to be easily shared and transformed is essential to achieving the goals of the MGI.

As a part of this materials data infrastructure, the National Institute of Standards and Technology (NIST) is establishing essential data exchange protocols and the means to ensure the quality of materials data and models needed to foster widespread adoption of MGI approaches. This informatics infrastructure will play an important role, in particular, in the form of repositories that contain materials simulation and experimental data and metadata, models, and code. These repositories and other infrastructure will provide resources for use in the materials development process as researchers strive to create materials with targeted properties. NIST is particularly working to enable and enhance the exchange of materials resources across repositories, subdomains of the materials community, and industries. NIST is also working to assess and improve the quality of materials data, models, and infrastructure.

Users of these developing data resources come from diverse communities. Many informatics efforts are, by immediate necessity, ad hoc and organic as opposed to being top-down. Each community has its

own data, metadata, and tools that are often incompatible. NIST believes that there is a need for new methods to enable the rapid definition of data and metadata, as well as a need for tools to enable rapid discovery and integration of these diverse data.

HIGH-LEVEL REQUIREMENTS

We believe that, from an informatics perspective, the MGI goals of accelerating materials development and deployment will hinge on two high-level requirements:

- (1) Materials researchers require a platform for interoperable exchange of materials data and metadata, which supports an approach of modular community-developed data standards.
- (2) Materials researchers need a decentralized infrastructure to enable finding and sharing of materials resources.

To meet the first requirement, researchers must have a system of data templates that can be designed to form custom containers for their experimental and simulation data and its associated metadata. These custom data formats will, however, be made from combinations of standardized components including community-developed templates that describe particular experiments or simulations and low-level reusable data types that encode data values and metadata fields in a standard way. As a result, it is anticipated that many of the issues associated with the current diversity of materials data formats will disappear without requiring researchers to force fit their data into monolithic data formats ill-suited to their needs.

Despite the success of Web-based search engines, they are in many ways not suited for searching for scientific resources. In this context, we use the term “resources” to include datasets and data collections or repositories, and information about organizations, application programming interfaces (APIs) and other information services, informational websites, and software. Simple text-based searches often return too many irrelevant results that require researchers to filter tediously through pages of output or to spend time devising clever search queries. Meeting the second requirement implies creating an informatics infrastructure that will enable materials researchers to search for materials data using metadata schemas with well-defined meanings. It will also enable them to make their data and other resources available to others using the same decentralized infrastructure.

The use of registries in informatics infrastructures is not new. In healthcare, registries support the task of identifying documents related to a patient in systems conforming to the Integrating the Healthcare Enterprise (IHE) Cross-Enterprise Document Sharing (XDS) integration profile.⁴ Metadata pertaining to a patient document are indexed in a registry that can be queried. In astronomy, the

Virtual Observatory⁵ provides astronomers with a distributed ecosystem for data-based research that includes community-established data protocols, formats, and tools. A key component of the discovery framework is federation of data resource registries that contain searchable metadata about archives, data collections, and services that are available.⁶

In addition, various other scientific registries and support tools are being developed.^{7–9} The Research Data Alliance (RDA) Data Type Registries Working Group has defined a data model for the collection of scientific data and has implemented a prototype data type registry¹⁰ to facilitate the understanding of scientific data collected by different research groups.

Also, a variety of materials science-based efforts exist to improve the exchange of materials-based data. The Materials Intelligence system from Granta Design* integrates materials data with a variety of software tools.

Boyce et al. worked to develop an integrated system by using HDF5 formats.¹¹ MatSeek¹² developed an ontology-focused system to federate search capabilities for materials data. The Materials Commons platform¹³ is a JavaScript Object Notation (JSON)-based modular system for data curation and provenance documentation. As far as we know, no materials informatics infrastructure currently exists that can easily and flexibly adapt with minimum development effort to the variety of needs described by our high-level requirements.

OVERALL ARCHITECTURE

After considering these previous and current efforts, we have chosen a Web-based approach that uses a Python-based Django framework, as illustrated in Fig. 1. User interaction can occur via a graphical user interface (GUI) or through scripts connected via a representational state transfer (REST) API. For data harvesting applications, we use the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) to query and retrieve data from known data providers such as repositories and other registries. Data, metadata, and binary large objects (BLOBs), such as images, are handled by a data management layer that ultimately stores the data and metadata contained in the Extensible Markup Language (XML) documents in a MongoDB NoSQL database. BLOBs are stored separately by default with MongoDB's GridFS, but other repositories such as DSpace can also be used. Our system can act as a data provider for harvesting via other OAI-PMH compliant systems.

*Certain commercial equipment or software is identified in this article to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the material or equipment identified is necessarily the best available for the purpose.

Fig. 1. Overall architecture of the informatics infrastructure under development.

Fig. 2. How the curation of data from literature fits into the overall architecture.

An important aspect of our architecture is the use of XML to structure data and metadata because this provides standardized methods for the encoding, interpretation, and transformation. We expect that

user communities will work together to generate shared data and metadata models expressed as XML Schema. Our infrastructure then dynamically renders a GUI based on the schema to allow users to input data

Materials Data Curation System
Welcome, admin. Thanks for logging in.

Logout | Dashboard | Help

Home | **Data Curation** | Data Exploration | Composer

Select Template | **Enter Data** | View Data

Data Curation

- Select Template
- Enter Data**
- View Data

Data Entry

Here you can fill in the Materials Data form. Once it is completed, you can go to 'View Data' to review what you have entered. You won't be able to reach the review page before the document is valid according to the selected template. From the review page, you will be able to curate the data. The 'Save Form' button allows you to save partial data that you may want to edit later. This will only save a temporary document and won't actually curate data. All grayed elements are optional. Thus, all elements written in black are required. The document may still be valid with empty elements. There are no validation on empty fields if no such constraint is defined in the template. Thus, an empty string of characters may not raise a validation error, but an empty number will.

Clear Fields | Save Form | Download

- experiment
 - experimentType
 - Choice:
 - tracerDiffusivity
 - material
 - materialName:
 - phase
 - name:
 - crystalStructure
 - spaceGroup:
 - wyckoffSequence:
 - Composition
 - quantityUnit:
 - constituents

Select Elements

 No element selected.
 - materialForm
 - Choice:

Fig. 3. A data entry form dynamically generated from an XML Schema file.

conforming to that schema. As MongoDB uses Binary JSON (BSON), a variant of JSON, to represent its data, we have created a translation layer that converts XML documents into the corresponding BSON and then back to XML as needed. The transformability of XML is also used to export retrieved XML documents to other formats. Currently, we allow for conversion to

other text-based formats such as comma separated values (CSVs), but in principle any format can be generated, including graphics.

Our architecture has been implemented for Windows, Mac OS X, and Linux and is currently the basis for four systems: the Materials Data Curation System (MDCS), the NIST Materials Resource

Materials Data Curation System

Welcome, admin. Thanks for logging in.

Logout | Dashboard | Help

Home | Data Curation | Data Exploration | Composer

Query by Example | Search by Keyword

Perform Search

Click on a field of the Query Builder to add an element to your query. The custom tree defined in Step 2 will appear and you will be able to select one of the highlighted elements. Click on plus/minus icons to add/remove criteria. You can save queries to build more complex queries and you will retrieve them on your next connection. When your query is done, please click on Submit Query to get XML documents that match the criterias. [Click here to get all documents of selected template.](#) or [Click here to get all documents from all versions of selected template.](#)

Query Builder

materialName is Aluminum

AND quantity < 5

Save Query Clear Submit Query

Saved Queries

Delete All

Add to Builder	Queries	Delete
↑	materialName is Aluminum	×
↑	element is Al	×
↑	quantity < 5	×
↑	(materialName is Aluminum AND quantity < 5)	×

Home | Data Curation | Data Exploration | Composer | Contact

Top ^

Fig. 4. Search interface generated from XML Schema.

Registry (NMRR), the MGI Code Catalog (MCC), and the National Metrology Institutes Resource Registry (NMIRR). The first two systems will be discussed in more detail here.

MATERIALS DATA CURATION SYSTEM

The MDCS was designed to address the first high-level informatics requirement of the MGI that materials researchers need modular data models that capture their data and metadata in community-developed templates using reusable data types. The MDCS source code and installation instructions are available from <https://github.com/usnistgov/MDCS>.

Scientific data exist in a multitude of formats, and similar data are often encoded in many ways. This diversity makes it difficult to combine data from multiple sources, understand and reuse existing data, find associated metadata, and transform data into new formats to support its reuse. Figure 2 shows how the MDCS fits in our overall architecture when data are curated from literature. By using a community-developed template expressed in XML Schema, a user can interact with a dynamically generated user interface to enter data and load images or other binary data into the MDCS. A similar user interface will allow the user to retrieve data already entered into the MDCS. Data are

Fig. 5. Interface of the template composer.

converted for storage and retrieval from MongoDB by a data management layer. Images and BLOBs are stored separately. An MDCS instance may act as a data provider to a registry; this functionality is available via the OAI-PMH data provider. The exporter functionality is also available to convert the data into other, possibly non-XML, formats.

Multiple instances of the MDCS can be connected to support federated searches. Figures 3 and 4 show the types of graphical user interfaces that can be dynamically generated from an XML Schema. The data entry form in Fig. 3 was generated from an XML Schema representing diffusion data, and Fig. 4 shows a search form also generated from

Fig. 6. Overriding the default rendering of an element in a schema using a UI module.

Fig. 7. UI modules support backend processing.

the schema. The ability to generate forms dynamically directly from XML Schema saves development effort and increases the flexibility of the MDCS.

As XML Schema plays a central role in the MDCS, a concern is that reliance on this technology might prove an obstacle to widespread use of the MDCS by users who are not versed in schema development and use. In recognizing this, we have created a template composer as part of the MDCS that allows users to either start with an existing XML Schema and modify it or use an existing collection of lower level templates to create an entirely new template. Figure 5 shows a screenshot of the MDCS template composer. We plan to leverage public registries in the future to enable an ecosystem centered on the creation and sharing of MDCS templates.

The dynamically generated user interfaces are limited to generating default user interface widgets for a given schema. Certain schema elements, such as the one representing the elements of the Periodic Table, would by default be rendered as a long pull-down list. This is unnecessarily tedious, and the

MDCS provides facilities to override the default user interface elements with custom widgets. Figure 6 shows how the default Periodic Table element pull-down list is replaced by a custom Periodic Table. The custom widget was developed by a programmer and is associated with the XML Schema elements in the Admin dashboard by the MDCS system administrator. Subsequent uses of the template now render the Periodic Table in a more familiar format.

The user interface (UI) module system can do more than just override the default rendering of XML tags in the input form. It can also be used to create entire mini applications that can do backend processing to support the overall use of the MDCS for curating particular types of data. Figure 7 shows that the UI module architecture is capable of interacting with the server, remote data sources, and external programs to support data processing and validation. Figure 8 shows the administrative user interface that allows a module to be associated with elements from a schema. This will effectively

Fig. 8. The module manager associates UI modules with elements in documents generated from an XML Schema.

allow the default rendering behavior associated with an element to be replaced by another specified behavior in the module.

The MDCS allows for automated curation of data via user scripts written in languages such as Python that interact via the MDCS REST API (Fig. 9). The REST API enables the full functionality of the MDCS to be accessed by using a wide variety of programming languages without using the graphical user interface. Scientific equipment will often generate output data in a text format. It is a relatively simple manner to write code that will convert the text data into an XML document and then submit it to the MDCS for storage. We have used Swagger to expose and document the MDCS REST API to users via a Web browser. This should greatly facilitate its use.

One of the great strengths of XML is its ability to be transformed into other formats by using standard tools such as Extensible Stylesheet Language

Transformations (XSLT), a programming language that uses XML syntax. The MDCS Exporter allows for the XML documents stored in the MDCS to be transformed into other formats such as CSV by using an XSLT stylesheet associated with the schema. This enables data stored in XML to be converted into tool-specific formats for use as part of scientific workflows.

NIST MATERIALS RESOURCE REGISTRY

The NIST Materials Resource Registry (NMRR) was developed to address the second high-level MGI informatics requirement that materials researchers need to be able to find and share materials resources in a decentralized way. The source code for the NMRR is available from <https://github.com/usnistgov/MaterialsResourceRegistry>.

Figure 10 shows how the NMRR fits within our overall architecture. NMRR users can publish metadata describing their resources using

swagger		http://127.0.0.1:8000/docs/apiapi-docs/	api_key	Explore
		Django REST Swagger		
blob		Show/Hide	List Operations	Expand Operations
POST	/rest/blob	GET http://localhost/rest/blob?id=id		
GET	/rest/blob	GET http://localhost/rest/blob?id=id		
curate		Show/Hide	List Operations	Expand Operations
POST	/rest/curate	POST http://localhost/rest/curate		
explore		Show/Hide	List Operations	Expand Operations
GET	/rest/explore/select/all	GET http://localhost/rest/explore/select/all		
GET	/rest/explore/data/download	GET http://localhost/rest/explore/data/download		
GET	/rest/explore/select	GET http://localhost/rest/explore/select		
DELETE	/rest/explore/delete	GET http://localhost/rest/explore/delete		
POST	/rest/explore/query-by-example	POST http://localhost/rest/explore/query-by-example		
exporter		Show/Hide	List Operations	Expand Operations
GET	/rest/exporter/select/all	GET http://localhost/rest/exporter/select/all		
GET	/rest/exporter/select	GET http://localhost/rest/exporter/select?param1=value1¶m2=value2		
POST	/rest/exporter/xslt/add	POST http://localhost/rest/exporter/xslt/add		
DELETE	/rest/exporter/xslt/delete	GET http://localhost/rest/exporter/xslt/delete?id=id		
POST	/rest/exporter/export	POST http://localhost/rest/exporter/export		
repositories		Show/Hide	List Operations	Expand Operations
GET	/rest/repositories/select/all	GET http://localhost/rest/repositories/select/all		
GET	/rest/repositories/select	GET http://localhost/rest/repositories/select?param1=value1¶m2=value2		
POST	/rest/repositories/add	POST http://localhost/rest/repositories/add		
DELETE	/rest/repositories/delete	GET http://localhost/rest/repositories/delete?id=IDtodelete		

Fig. 9. The MDCS REST API as exposed by Swagger.

community-developed metadata templates rendered by a graphical user interface, and they can also search and discover existing resources. Additionally, resource metadata can be published in an automated fashion by using the REST API or it can be harvested from registered data providers (such as repositories and other registries) using OAI-PMH. An NMRR registry can also serve as a data provider for other OAI-PMH compliant registries. In this fashion, multiple NMRR installations can be interconnected to create a decentralized federation of registries. Figure 11 shows the interface presented to users searching for resources. The resource search and retrieval process begins when

a user submits a query to the NMRR search interface. The NMRR then responds with a list of available resources that match the query. The user then selects the link to the appropriate resource and the user's browser is redirected to that resource.

The NMRR and the MDCS are complementary systems where the MDCS can be used to make materials data accessible and the NMRR can be used to make materials data discoverable. From the perspective of the data consumer, a search on the NMRR returns candidate instances of the MDCS and other repositories. The user can then search an individual repository for candidate datasets.

Fig. 10. How the NMRR fits into the overall architecture.

Fig. 11. NMRR search interface.

DISCUSSION

The goal of the MDCS is to facilitate the collection, use, and reuse of materials data and to provide the needed informatics infrastructure to facilitate

the implementation of ICME approaches. Several collaborators are using the MDCS for their own work. Northwestern University's NanoMine, an online platform for the prediction of polymer

Fig. 12. (a) MRR harvests metadata from a variety of data sources including both MDCS and data repositories around the world. (b) User sends data request to MRR. (c) MRR sends the user a variety of potential data sets and enables the use of a REST-API to push data to user.

nanocomposites, uses the MDCS to curate nanocomposite processing, structure, and property data reported in literature and then to link it to a variety of modeling tools.¹⁴ Raymundo Arroyave's group at Texas A&M University is using the MDCS to collect data from computational materials science simulations and measurements of differential scanning calorimetry. At NIST, work is being done to curate both literature and experimental thermodynamic data with the MDCS. The NIST Thermodynamic Research Center is expanding ThermoML^{15,16} to include data on metals and plans to integrate their efforts with the MDCS.

The MDCS is also being integrated with the Interatomic Potentials Repository (IPR) Project.¹⁷ A recent article summarized the expanded scope of the IPR Project as a response to the MGI.¹⁸ Prior to the creation of the MDCS, metadata for interatomic potentials were manually curated in semistructured text files. As the project is working to enable selection of interatomic potentials based on material properties and other metadata, the MDCS is being used to curate all supporting data and metadata. Furthermore, rapid property calculation tools are being developed and directly integrated with the MDCS via its API. This combined toolset could also be used to develop new potentials, where local instances of IPR tools and the MDCS address data

management issues associated with developing many different iterations or variants of interatomic potentials, as part of the typical development process.

A 2014 whitepaper indicated that high-throughput experiments (HTEs) are uniquely suited to meet many needs within the MGI by generating large volumes of high-quality experimental data suitable for model validation or model input.¹⁹ Efforts at NIST are focused on capturing data as it is generated on the synthesis or measurement apparatus and automatically transforming applicable data and metadata into XML formats, which are compliant with the MDCS. This effort is part of a broader effort to exchange samples and data across institutions to advance HTE metrology.

As the use of the MDCS expands, users of this software will be able to register datasets to share using the Materials Resource Registry. Registering a dataset will allow the metadata to be harvested, enabling potential users to find it. Figure 12 illustrates how the potential user might use the Materials Resource Registry to locate data stored in Materials Data Curators and a variety of other data repositories.

The open source software infrastructure presented in this work supports both data curation using modular data schema models for data exchange and decentralized data search platform. The MDCS

will enable the materials science community to build and share community-based data models for the curation of specific data types. The Materials Resource Registry will improve the ability to find and share data with the metadata harvestable by other registries. Both the MDCS and the NMRR are designed to work with other data curation and sharing tools to further the aims of the MGI.

ACKNOWLEDGEMENTS

The authors would like to thank Raymundo Arroyave, Cate Brinson, Yannick Congo, Lucas Hale, Ya-Shian Li-Baboud, Greta Lindwall, Chris Muzny, Pierre Savonitto, Daniel Saucedo, and Richard Zhao for their support and contributions.

REFERENCES

1. C. Rae, *Mater. Sci. Techn.* 25, 479 (2009).
2. W. Xiong and G.B. Olson, *MRS Bull.* 40, 1035 (2015).
3. National Science and Technology Council, *Materials Genome Initiative for Global Competitiveness* (Washington: Office of Science and Technology Policy, 2011).
4. Integrating the Healthcare Enterprise (IHE International 2015), <http://www.ihe.net/>. Accessed 24 March 2016.
5. R.J. Hanisch, G.B. Berriman, T.J.W. Lazio, S. EmeryBunn, J. Evans, T.A. McGlynn, and R. Plante, *Astron. Comput.* 11, 190 (2015).
6. M. Demleitner, G. Greene, P. Le Sidaner, and R.L. Plante, *Astron. Comput.* 7–8, 101 (2014).
7. Corda (Corporation of National Research Initiatives, January 2016), <https://www.corda.org/>. Accessed 24 March 2016.
8. 2nd Generation of Open Access Infrastructure for Research in Europe, OpenAIRE (Openaire Consortium, February 2016), <https://www.openaire.eu/>, Accessed 22 March 2016.
9. Research Data Switchboard (Research Data Alliance, March 2015), <http://www.rd-switchboard.org/>. Accessed 21 March 2016.
10. Data Type Registry (Corporation of National Research Initiatives, August 2014) <http://typeregistry.org/registrar/>. Accessed 13 July 2015.
11. D.E. Boyce, P.R. Dawson, and M.P. Miller, *Metall. Mater. Trans. A* 40A, 2301 (2009).
12. K. Cheung, J. Hunter, and J. Drennan, *Intell. Syst.* 24, 47 (2009).
13. B. Puchala, G. Tarcea, E.A. Marquis, M. Hedstrom, H.V. Jagadish, and J.E. Allison, *JOM* (2016). doi:10.1007/s11837-016-1998-7.
14. C. Brinson, H.R. Zhao, NanoMine, <http://brinson.mech.northwestern.edu/research/Nanomine.html>. Accessed Feb 2016.
15. M. Frenkel, R.D. Chirico, V.V. Diky, Q. Dong, S. Frenkel, P.R. Franchois, D.L. Embry, T.L. Teague, K.N. Marsh, and R.C. Wilhoit, *J. Chem. Eng. Data* 48, 2 (2003).
16. R.D. Chirico, M. Frenkel, V.V. Diky, K.N. Marsh, and R.C. Wilhoit, *J. Chem. Eng. Data* 48, 1344 (2003).
17. C.A. Becker, F. Tavazza, Z.T. Trautt, and R.A. Buarque de Macedo, *Curr. Opin. Solid State Mater. Sci.* 17, 277 (2013).
18. Z.T. Trautt, F. Tavazza, and C. Becker, *Model. Simul. Mater. Sci. Eng.* 23, 074009 (2015).
19. M.L. Green, J.R. Hattrick-Simpers, C.L. Choi, I. Takeuchi, A.M. Joshi, S.C. Barron, T. Chiang, A. Davydov, S. Empedocles, J. Gregoire, and A. Mehta, *Fulfilling the Promise of the Materials Genome Initiative with High-Throughput Experimentation* (MR Society, 2014), <http://www.mrs.org/mgi-workshop-full-report/>. Accessed 24 March 2016.