

Interactive Metadata Integration with Brick

Gabe Fierro¹, Anand Prakash², Cory Mosiman³,
Marco Pritoni², Paul Raferty¹, Michael Wetter² and
David E. Culler¹

¹University of California Berkeley, ²Lawrence Berkeley
National Laboratory, ³National Renewable Energy
Laboratory

Energy Technologies Area
November 2020

Disclaimer:

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

Demo Abstract: Interactive Metadata Integration with Brick

Gabe Fierro
UC Berkeley
gtfierro@cs.berkeley.edu

Marco Pritoni
Lawrence Berkeley National
Laboratory
mpritoni@lbl.gov

Anand Krishnan Prakash
Lawrence Berkeley National
Laboratory
akprakash@lbl.gov

Paul Raftery
UC Berkeley
p.raftery@berkeley.edu

David E. Culler
UC Berkeley
culler@cs.berkeley.edu

Cory Mosiman
National Renewable Energy
Laboratory
cory.mosiman@nrel.gov

Michael Wetter
Lawrence Berkeley National
Laboratory
mwetter@lbl.gov

ABSTRACT

Many different digital representations of a building are produced over the course of its lifecycle. While these representations individually contain metadata required to support different stages of the building's lifecycle, they are largely not interoperable due to differences in structure, syntax and semantics. This impedes the development and deployment of data-driven applications providing fault detection and diagnosis, virtual metering or optimal control. [4] introduces a new platform for the continuous curation of a unified, Brick [1]-based metadata model that can be maintained throughout the building lifecycle. In this demonstration, we present a live proof-of-concept implementation of the platform with support for several metadata representations in the context of a simulated building lifecycle.

CCS CONCEPTS

• **Information systems** → **Ontologies**; *Data encoding and canonicalization; Information extraction.*

KEYWORDS

Smart Buildings, Data Integration, Metadata, Ontologies, OWL, RDF, Brick

ACM Reference Format:

Gabe Fierro, Anand Krishnan Prakash, Cory Mosiman, Marco Pritoni, Paul Raftery, Michael Wetter, and David E. Culler. 2020. Demo Abstract: Interactive Metadata Integration with Brick. In *The 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '20)*, November 18–20, 2020, Virtual Event, Japan. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3408308.3431125>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

BuildSys '20, November 18–20, 2020, Virtual Event, Japan

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8061-4/20/11.

<https://doi.org/10.1145/3408308.3431125>




Figure 1: Illustration of the different sources of metadata produced over the course of a building's lifecycle

1 INTRODUCTION

Many different digital representations of a building are produced over the course of its lifecycle. These representations — sources of *metadata* — each support workflows within a given stage of a building's lifecycle (Figure 1) but largely lack the semantic content required to support many common data driven applications [2]. Recent building metadata efforts such as Brick [1] define models that can express the metadata required for such applications, but the challenge of how to bootstrap the creation of these models remains.

Current lines of research explore the use of human-in-the-loop machine learning techniques to derive Brick metadata from unstructured sources such as the labels found in building management systems [3, 5], or address one-off translations of existing sources of metadata to Brick [6]. In contrast, the work presented in [4] and demonstrated here explores the synthesis of a Brick metadata model from multiple existing sources of metadata.

While existing metadata representations like IFC, gbXML, Project Haystack, Modelica/CDL and BuildingSync are governed by a mix of standards, formats, and industry conventions, they are sufficiently different in their structure, syntax, and semantics that they are not interoperable. The lack of interoperability means that the different representations cannot be easily combined without significant manual effort. To complicate matters, because these representations are typically produced at different times and by different teams or individuals, the metadata they contain can often be incomplete, out of date or even change at a later date in response to retrofits or repairs.

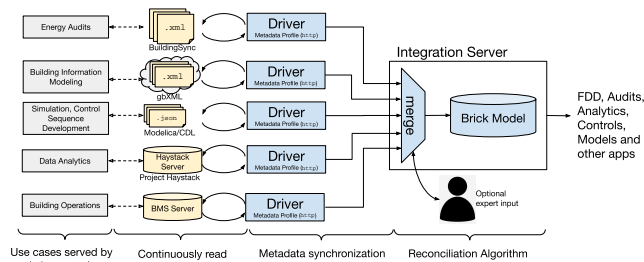


Figure 2: Architecture of the reconciliation platform

2 A METADATA INTEGRATION PLATFORM

The metadata integration platform introduced in [4] enables the creation and maintenance of a Brick metadata model for a building. In order to address the challenges outlined above, the platform divides its responsibilities into two components: *drivers* and the *integration server* (Figure 2).

Drivers interact directly with a source of metadata — e.g. gbXML files hosted on a file share or a Project Haystack model stored in a database — and produce a stream of *records* describing a set of *entities*. An entity is any virtual, logical or physical thing or data source in the building. Each record contains the Brick metadata for an entity that has been inferred or derived from the underlying metadata source, a pointer to the definition of the entity from the source and some supplementary metadata. A driver produces a set of records for each version of the original metadata source. As metadata sources change or become available, drivers transmit these metadata-containing records to the integration server.

The *integration server* catalogs the records received from the set of deployed drivers and merges the contained Brick metadata into a unified Brick model. The integration server *reconciles* the differences between the Brick metadata reported by each driver so that the unified Brick model is consistent with respect to the entities common to multiple metadata sources. The reconciliation process employed by the server consists of two phases. The first phase produces clusters of equivalent entities using a combination of properties for each entity, including their original labels, Brick relationships and other ontology-derived characteristics. Each cluster consists of one or more entities from one or more different metadata sources.

The second phase merges the Brick metadata for each entity in a cluster and validates the result using the formal axioms and rules defined by Brick. When validation fails or the clustering algorithm is unable to find equivalent entities, the integration server consults a human expert for additional metadata and reinitiates the reconciliation process. After validation succeeds, the integration server unifies the Brick metadata for all clusters into a model that can be made available to applications.

3 DEMONSTRATION OVERVIEW

In the demonstration, we plan to give participants the opportunity to interact with a live, proof-of-concept implementation of the metadata integration platform in the context of a simulated building lifecycle. A web interface (Figure 3) will illustrate a timeline

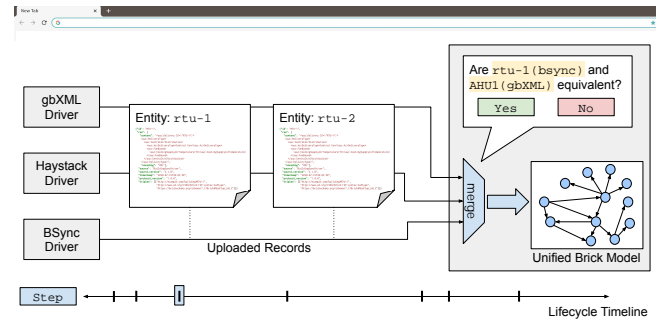


Figure 3: Mockup of interactive web interface showing timeline, transmitted records from drivers, integration server, and merged Brick model.

of a simulated building’s lifecycle and provide access to historical versions of a Project Haystack model, Modelica model and BuildingSync model for the simulated building. Users can advance through the timeline, inspect the produced records of inferred Brick metadata at each stage, and participate in the reconciliation algorithm resulting in a unified Brick model. The implementation of the metadata integration platform is open-source and available online at <https://github.com/gtfierro/shepherding-metadata/>.

4 ACKNOWLEDGEMENTS

This research is supported in part by Department of Energy grants DE-EE0008681, DE-AC36-08GO28308 and DE-AC02-05CH11231 and the CONIX research center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The opinions expressed belong solely to the authors, and not necessarily to the authors’ employers or funding agencies.

REFERENCES

- [1] Bharathan Balaji, Arka Bhattacharya, Gabriel Fierro, Jingkun Gao, Joshua Gluck, Dezhi Hong, Aslak Johansen, Jason Koh, Joern Ploennigs, Yuvraj Agarwal, et al. 2016. Brick: Towards a unified metadata schema for buildings. In *Proceedings of the ACM International Conference on Embedded Systems for Energy-Efficient Built Environments (BuildSys)*. ACM.
- [2] Arka Bhattacharya, Joern Ploennigs, and David Culler. 2015. Short Paper: Analyzing Metadata Schemas for Buildings: The Good, the Bad, and the Ugly. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*. ACM, 33–34.
- [3] Arka A Bhattacharya, Dezhi Hong, David Culler, Jorge Ortiz, Kamin Whitehouse, and Eugene Wu. 2015. Automated metadata construction to support portable building applications. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*. ACM, 3–12.
- [4] Gabe Fierro, Anand K. Prakash, Cory Mosiman, Marco Pritoni, Paul Raftery, Michael Wetter, and David E. Culler. 2020. Shepherding Metadata Through the Building Lifecycle. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*.
- [5] Jason Koh, Dezhi Hong, Rajesh Gupta, Kamin Whitehouse, Hongning Wang, and Yuvraj Agarwal. 2018. Plaster: An integration, benchmark, and development framework for metadata normalization methods. In *Proceedings of the 5th Conference on Systems for Built Environments*. 1–10.
- [6] Henrik Lange, Aslak Johansen, and Mikkel Baun Kjærgaard. 2018. Evaluation of the opportunities and limitations of using IFC models as source of building metadata. In *Proceedings of the 5th Conference on Systems for Built Environments*. 21–24.