Short Paper: Analyzing Metadata Schemas for Buildings — The Good, The Bad, and The Ugly

Arka A. Bhattacharya
University of California, Berkeley
Berkeley, CA, USA
arka@eecs.berkeley.edu

Joern Ploennigs
IBM Research
Smarter Cities Technology Centre
Dublin, Ireland
joern.ploennigs@ie.ibm.com

David Culler
University of California, Berkeley
Berkeley, CA, USA
culler@berkeley.edu

ABSTRACT

Commercial buildings account for a large amount of delivered energy in the U. S., nearly 42 % of which is consumed in buildings with digital control systems [4]. These buildings are a ripe venue to deploy novel applications because of (a) access to sensors and actuators that are used in their digital control systems, (b) deployed wireless sensor networks, and (c) the advent of smart "internet-of-things" sensors. However, these novel applications face a fundamental scalability challenge because the sensor metadata across buildings do not follow any common schema. In this paper, we quantify the shortcomings of three metadata schemas which have gained traction in modeling the contextual, spatial and functional relationships between sensors in the built environment: (1) Project Haystack [2], (2) Industry Foundation Classes [12], (3) and Semantic Sensor Web [6] against three commercial buildings and an extensive list of smart-building applications.

Categories and Subject Descriptors

H.5.m [**Information Systems**]: Information Interfaces And Presentation—*Miscellaneous*

Keywords

Building Metadata Schemas; Portable Building Applications

1. INTRODUCTION

Most commercial buildings have an underlying wired sensor network that enables digital controls for managing and monitoring various aspects of the building. Vendors setting up these digital control systems often use customized metadata schemas to describe the sensors' context, function and semantic inter-relationships. Such schemas vary across buildings and across vendors. This non-uniform metadata problem is exacerbated by the custom metadata schemas of wireless sensor networks and novel sensors (such as BLE temperature and humidity sensors, etc.) deployed in buildings, making it extremely hard to deploy new applications like energy visualization, demand response, energy disaggregation, occupancy modeling, model predictive control, or anomaly detection in an unknown building, and porting the same application across buildings.

There have been many efforts to come up with a common metadata schema for all buildings. We evaluate three of these schemas: (1) Project Haystack [2], (2) Industry Foundation Classes [12], and (3) Semantic Sensor Networks [6]. We use sensor metadata from three large commercial buildings, each set up by a different Building Management System (BMS) vendor, and a list of 87 published smart-building applications[1] to evaluate the three schemas based on the following criteria:

- *Completeness:* Could all the distinct sensor metadata information (such as a sensor's location, type, etc.) contained in these buildings be represented by these schemas?
- *Ability to Capture Relationships:* Is it possible to express all the sensor relationships (required by the applications we study) through these schemas?
- *Flexibility:* How flexible are these schemas to capture uncertainty in the metadata (e. g. uncertainty over whether an air pressure sensor is located before or after a damper), or the emergence of new sensors (e. g. Apple iBeacon, Kinect, card swiping machines) and subsystems (e. g. a smart couch) ?

Our results show that none of the three schemas capture all the available sensor metadata, nor express all the relationships required by novel applications, nor capture any notion of uncertainty, nor allow easy extensibility to model novel sensors.

2. EXPERIMENTAL SETUP

We evaluate the effectiveness of the metadata schemas on three buildings set up by different BMS vendors [3]. One of the buildings is located in Dublin, Ireland and two in Berkeley, California, U.S.A. The metadata of the BMS sensors in these buildings were manually mapped against a set of Haystack tags to the extent possible. In cases where Haystack did not have the relevant tags, we developed our own consistent tag names.

Terminology: We define a *sensor label* as a collection of tags capturing a sensor's metadata. For example, in one of the buildings in our testbed, a sensor was encoded with the metadata `SODA1R465__ART`. This indicated that the sensor

BuildSys'15, November 4–5, 2015, Seoul, South Korea..
© 2015 ACM. ISBN 978-1-4503-3981-0/15/11 ...$15.00.
DOI: http://dx.doi.org/10.1145/2821650.2821669.

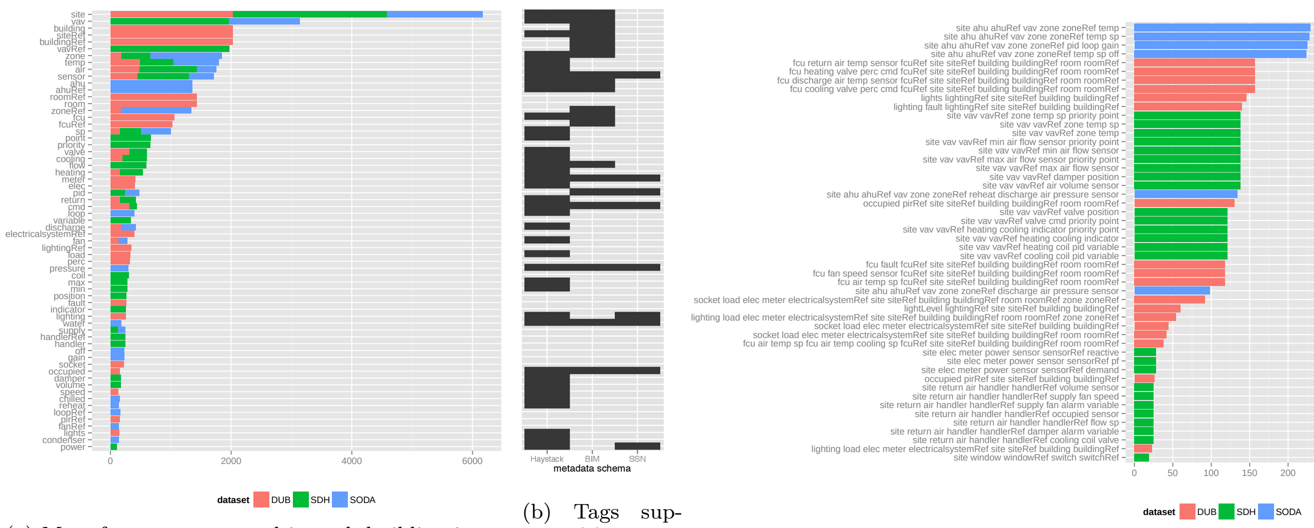

(a) Most frequent tags used in each building in our testbed. The list of tags apply also to Figure 1b.

(b) Tags supported by metadata schema.

(c) Most frequent sensor labels in each building in our testbed.

Figure 1: Analysis of tags and sensor labels in the three buildings in our testbed – DUB, SODA and SDH

was a *zone air temp sensor* (`ART`) in *zone* 465 (*zoneRef* : `465`) which is served by *ahu* 1 (*ahuRef* : `1`) in the *site* SOD (`SOD`). Hence its sensor label is *'site ahu ahuRef zone zoneRef zone return temp sensor'* , each of the terms in the label being a 'tag'. Sensor labels indicate which tags the building vendors put together while framing the custom metadata of a sensor. The three buildings have 2028, 2551, and 1586 sensors respectively, comprising of 510, 148, and 281 distinct sensor labels.

The histogram of the tags is shown in Fig. 1a. Many tags, such as *site, zone, sensor* are common accross the buildings. This illustrates that a taxonomy of common tags exists between various buildings. However, Fig. 1b shows that none of the three studied schemas has a taxonomy that includes all the tags used in these buildings. Hence, none of the three schemas are complete. The most common sensor labels in our dataset are shown in Fig. 1c. Although tags are common across buildings, sensor labels are not, showing that different BMS vendors use different combination of tags while describing a sensor which might perform the same function. Also note that the frequency of occurrence of individual tags are Pareto distributed, implying that getting a common taxonomy of a few tags can result in normalizing a large amount of building metadata.

3. EVALUATION OF SCHEMAS

We quantify the extent to which the tags identified in our testbed are expressible in the three studied metadata schemas. We also analyze a set of 87 applications from the building application literature[1] to quantify whether the schemas can capture the required sensor relationships.

The main classes of applications and their metadata requirements are listed in Table 1. The first column lists the needed relationships between different dimensions of metadata. For example, 'sensor ↔ function' states that a link between the sensor entity and the functional semantic of the sensor is required. Hierarchical relationships are possible such as 'location ↔ location' refers to relating the sublocations to a location. The second column provides an intuitive query example for each relationship. the applications listed in the columns 3 to 10. An 'X' indicates that a specific relationship is required to automate the application. The last three columns show which relationships are supported by the studied metadata schema. The last three rows in the table then compute the *applicability* of each schemata as the fraction of the required relationships the is captured by a particular schema.

3.1 Project Haystack

Project Haystack is one of the main open-source initiatives trying to facilitate interoperability of applications across buildings. It defines tagging models, data formats and data structures to exchange data over HTTP using REST APIs. Project Haystack specifies entities, such as air handling units, variable air volume units, etc. and uses tags (or name/-value) pairs to add attributes to those entities. Entities can be linked together using references (similar to foreign keys in databases). The haystack taxonomy defines a finite set of tags and entities.

Completeness: Figure 1b shows the tags supported by Haystack in our three testbed buildings. Project Haystack supports 54 % of the unique tags used in the three datasets. Weighted by the occurrence frequency of the tags in the dataset, Haystack supports 63 % of the tags. While Haystack has well-defined tags to capture very commonly used sensor types, e. g. zone temperature sensors, its tags are incapable of representing (a) building-specific sensors such as a fault-detecting sensor that is monitoring the status of a pump in the condensor water loop, (b) common sensors outside the HVAC system, such as whether a light array is monochromatic or has hue controls, or whether an entire light panel is controlled by a single sensor or by multiple sensors.

Ability to Capture Relationships: Table 1 compares the relationships required by common applications toward the support in Haystack. Although Project Haystack is capable of capturing relationships between HVAC subsystems, it cannot model relationships between spatial elements, such as the list of rooms in a floor. Also, Haystack is unable

Table 1: Metadata relationships required by building applications and their expressibility in the three metadata schemas

Needed Relationships		Common Applications								Schemata		
Relationships	Example	Occupancy Modelling[1]	Energy Apportionment[2]	Web Displays[3]	Model-Predictive Control[4]	Participatory Feedback[5]	Fault Detection and Diagnosis[6]	Non-Intrusive Load Monitoring[7]	Demand-Response[8]	Haystack	IFC	Semantic Web
sensor ↔ function	sensorsOf(type)	X								X	X	X
sensor ↔ location	sensorsIn(room)	X	X		X	X	X		X	X	X	X
location ↔ location	roomsIn(building)	X	X								X	X
asset ↔ location	ahuOf(room)	X			X	X	X		X	X	X	
sensor ↔ asset	sensorsOf(AHU)	X	X		X	X	X		X	X	X	
asset ↔ asset	ahuSupplying(vav)	X			X	X	X		X	X	X	
location ↔ persons	occupantOf(room)	X										
location ↔ organisation	ownerOf(room)	X	X								X	
gadget ↔ persons	macAddrOfPhone(user)	X										
gadget ↔ location	computerIn(room)	X										
meter ↔ location	meterOf(room)		X					X		X	X	X
meter ↔ gadget	meterOf(computer)	X										
meter ↔ asset	meterOf(ahu)	X										
meter ↔ sensor	sensorUnder(meter)		X					X				
Applicability	Haystack	42%	50%	75%	100%	100%	100%	50%	100%	77 %		
	IFC	58%	83%	100%	100%	100%	100%	50%	100%		86 %	
	Semantic	25%	50%	100%	25%	25%	25%	50%	25%			41 %

References: [1][16, 15]; [2][9, 13]; [3][5]; [4][18, 20]; [5][11, 14]; [6][19, 21]; [7][17]; [8][22]
A more complete set of citations is available in [1].
Applicability denotes what fraction of the required relationships is captured by a particular schema.

to capture relationships between different views (spatial, HVAC, power) of a building, e. g. which set of rooms/floors an HVAC zone comprises of.

Flexibility: The schema has no way to capture any uncertainty, does not capture any novel sensors/electronic gadgets (such as computers, iPhones, etc.), and requires consensus among the community to include new sensor tags into the vocabulary.

3.2 Building Information Models (IFC)

Building Information Models (BIM) is widely discussed as the solution to the problem of information management during a building's entire life-cycle. Common exchange formats include Industry Foundation Classes (IFC), COBie, and gbXML. IFC is the most comprehensive format of the three. It was first standardized in 2000 as ISO 16739 [12] and supports explicit modeling of sensors, actuators and controllers since version 4 (published in 2013). COBie is a subset (view) of IFC that focusses on simple export formats such as Excel [8]. The Green Building XML schema gbXML [10] is another format that concentrates on energy performance analysis tools and only has rudimentary elements to model sensors, and no dedicated taxonomy of semantic types for sensors. In this study, we only evaluate IFC.

Completeness: IFC supports 11 predefined semantic types, such as CO2, heat, temperature and sound sensors, with additional properties specifying units, values, type, etc. 22 generic measurement types such as count, electric current, length and time are available. IFC provides concepts for 29 % of the unique tags in our dataset (Figure 1b), and 60% when weighted by tag occurrence frequency.

Ability to Capture Relationships: IFC has its foundation in 3D geometrical modelling and provide comprehensive ways to model spatial and asset relationships (Table 1). The main shortcoming with IFC is that metadata is primarily related via 3D objects, making it hard to query even simple things, such as if two devices are located within the same room. It is also not possible to assign multiple sensor types to a device, e. g. modeling multi-sensors that sample semantically different values. Concepts for humans and novel sensors do not exist.

Flexibility: Although, IFC and gbXML formats are built to be extensible, an extension to the standardized core model requires consensus from a community primarily focused on building design. Also it is unable to capture uncertainty of 3D objects such as room size, wall thickness, etc.

3.3 Semantic Sensor Web

The goal of the Sensor Web is to link sensors together using semantic relationships, and make their measurements available over the web. The Semantic Sensor Network (SSN) ontology is a set of domain independent concepts to model sensors [6], defined by the W3C consortium. The study on smart appliance interoperability in [7] provides a good summary on the availability of semantic models in the smart building domain. The study analyzes 47 semantic assets and identifies 20 recurring concepts within them that are modeled in a new Smart Appliances REFerence (SAREF) ontology.

Completeness: SAREF classifies device functionality roughly into Sensing, Actuating, Metering. For each function a specific sensor type can be defined as a literal. By default 5 types (Temperature, Occupancy, Humidity, Motion, Smoke, Pressure) are provided. Meter types such as Water, Gas, Pressure, Energy, and Power are provided. This small set of default tags results in a coverage of only 11% of the tags in our dataset, (8% when weighted by frequency).

Ability to Capture Relationships: The SAREF ontology allows modeling hierarchies of spatial elements, but does not specify modeling of assets (such as AHUs, VAVs, etc.). This strongly reduces the applicability of the SAREF ontology to the use cases defined in Table 1. Units are specified using the external ontology – Units of Measure (OM). SAREF links 20 ontologies that define concepts for other areas. This demonstrates an inherent strength of the semantic web, as additional ontologies can be linked and reused to model aspects such as organizational structure, novel devices, etc. However, a strict guideline of which ontology should be used is not provided by SAREF.

Flexibility: Ontologies differentiate between classes (kind of things) with attributes (properties of things) and individuals (instances of classes) with relationships (links between instances). While such ontologies are helpful in defining a clear and verifiable meta-model, it only allows capturing uncertainty on an individual level, such as multi-lingual names. Also, addition of novel sensors and taxonomies requires a consensus by the standards body.

4. CONCLUSION

In this study we studied three possible metadata schemas — Project Haystack, BIM, and Semantic Sensor Web — that could be used to normalize the metadata schemas across buildings, and quantified their shortcomings in three areas: (a) completeness, (b) ability to capture sensor relationships, (c) ability to easily express novel sensors and capture metadata uncertainty. Our analysis has three main takeaways:

- None of the metadata schemas investigated is complete or expressive enough to capture all the tags and semantic information in buildings. As long as this problem is not solved, BMS vendors will come up with their own metadata schemas for either (a) the sensors that do not have a standardized representation, or (b) all the sensors in a building in order to maintain a constant naming schema. This would, at best, lead to fragmentation of the schemas used for naming building sensors.
- No existing metadata schema is flexible to capture novel sensors e. g. a kindle, an iBeacon, etc., or capture any notion of uncertainty, e.g whether a temperature sensor in a large multi-thermal zone hall is capturing heat-exchange processes in adjoining zones, or uncertainty in whether a supply air pressure sensor is placed before or after the VAV damper.
- Semantic sensor web ontologies are too generic and fragmented to be of practical relevance. They need to have (a) a well-defined taxonomy of common building functions, (b) concepts for modelling building assets and persons, (c) have tools which can make them easy to use and validate by domain experts.

5. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation grants CPS-1239552, CNS-0931843 and Intel. We would also like to thank Gabe Fierro for his helpful feedback.

6. REFERENCES

[1] List of building applications. `https://ibm.biz/UCB-IBM-Apps`.

[2] Project haystack. `http://project-haystack.org/`.

[3] Used building metadata. `https://ibm.biz/UCB-IBM-Data`.

[4] E. I. ADMINISTRATION. Commercial buildings energy consumption survey,. Technical report, 1999.

[5] B. Balaji, H. Teraoka, R. Gupta, and Y. Agarwal. Zonepac: Zonal power estimation and control via hvac metering and occupant feedback. In *BuildSys*. ACM, 2013.

[6] M. Compton, P. Barnaghi, L. Bermudezc, and et al. The SSN ontology of the W3C semantic sensor network incubator group. *Web Semantics: Science, Services and Agents on the World Wide Web*, 2012.

[7] L. Daniele, F. den Hartog, and J. Roes. Study on semantic assets for smart appliances interoperability, Mar. 2015.

[8] E. W. East. Construction operations building information exchange (Cobie). Technical report, DTIC Document, 2007.

[9] D. Garnier-Moiroux, F. Silveira, and A. Sheth. Towards user identification in the home from appliance usage patterns. In *2013 ACM Conf. on Pervasive and ubiquitous computing adjunct publication*, pages 861–868. ACM, 2013.

[10] GBXML. Green building xml. `http://www.gbxml.org/`.

[11] L. A. Hang-yat and D. Wang. Carrying my environment with me: A participatory-sensing approach to enhance thermal comfort. In *BuildSys*. ACM, 2013.

[12] ISO 16739:2005 - Industry Foundation Classes, Release 2x, Platform Specification (IFC2x Platform), 2005.

[13] M. Jahn, T. Schwartz, J. Simon, and M. Jentsch. Energypulse: tracking sustainable behavior in office environments. In *2nd Int. Conf. on Energy-Efficient Computing and Networking*, pages 87–96. ACM, 2011.

[14] F. Jazizadeh and B. Becerik-Gerber. Toward adaptive comfort management in office buildings using participatory sensing for end user driven control. In *BuildSys*, 2012.

[15] D. Jung, V. B. Krishna, N. Q. M. Khiem, H. H. Nguyen, and D. K. Yau. Energytrack: Sensor-driven energy use analysis system. In *BuildSys*. ACM, 2013.

[16] W. Kleiminger, S. Santini, and F. Mattern. Smart heating control with occupancy prediction: How much can one save? In *ACM Int. Joint Conf. on Pervasive and Ubiquitous Computing*, pages 947–954. ACM, 2014.

[17] T. Liu, Y. Che, Y. Liu, Z. Xu, Y. Duan, and S. Chen. A user demand and preference profiling method for residential energy management. In *ACM Int. Joint Conf. on Pervasive and Ubiquitous Computing*, pages 911–918. ACM, 2014.

[18] A. Parisio, D. Varagnolo, D. Risberg, G. Pattarello, M. Molinari, and K. H. Johansson. Randomized model predictive control for HVAC systems. In *BuildSys*, 2013.

[19] J. Ploennigs, B. Chen, A. Schumann, and N. Brady. Exploiting generalized additive models for diagnosing abnormal energy use in buildings. In *BuildSys*. ACM, 2013.

[20] D. Sturzenegger, D. Gyalistras, M. Morari, and R. S. Smith. Semi-automated modular modeling of buildings for model predictive control. In *BuildSys*, pages 99–106. ACM, 2012.

[21] J. Weimer, S. A. Ahmadi, J. Araujo, F. M. Mele, D. Papale, I. Shames, H. Sandberg, and K. H. Johansson. Active actuator fault detection and diagnostics in HVAC systems. In *BuildSys*, pages 107–114. ACM, 2012.

[22] T. Weng, B. Balaji, S. Dutta, R. Gupta, and Y. Agarwal. Managing plug-loads for demand response within buildings. In *BuildSys*, pages 13–18. ACM, 2011.