The Industrial Ontologies Foundry Proof-of-Concept Project

Boonserm (Serm) Kulvatunyou[1], Evan Wallace[1],
Dimitris Kiritsis[2], Barry Smith[3], Chris Will[4]

[1]Systems Integration Division, National Institute of Standards and Technologies
{serm, evan.wallace}@nist.gov
[2]EPFL, Switzerland
dimitris.kiritsis@epfl.ch
[3]State University of New York, Buffalo
phismith@buffalo.edu
[4]Dassault Systemes
chris.will@3ds.com

Abstract. The current industrial revolution is said to be driven by the digitization that exploits connected information across all aspects of manufacturing. Standards have been recognized as an important enabler. Ontology-based information standard may provide benefits not offered by current information standards. Although there have been ontologies developed in the industrial manufacturing domain, they have been fragmented and inconsistent, and little has received a standard status. With successes in developing coherent ontologies in the biological, biomedical, and financial domains, an effort called Industrial Ontologies Foundry (IOF) has been formed to pursue the same goal for the industrial manufacturing domain. However, developing a coherent ontology covering the entire industrial manufacturing domain has been known to be a mountainous challenge because of the multidisciplinary nature of manufacturing. To manage the scope and expectations, the IOF community kicked-off its effort with a proof-of-concept (POC) project. This paper describes the developments within the project. It also provides a brief update on the IOF organizational set up.

Keywords: Smart manufacturing, Industrie 4.0, Ontology, IOF, Industrial Ontologies Foundry

1 Introduction

The current industrial revolution is said to be driven by the digitization that exploits connected information across all aspects of manufacturing [1]. Standards have been recognized as an important enabler; meanwhile, ontology is considered as the next generation standard for connected information. Although there have been ontologies developed in the industrial manufacturing domain, they have been disparately developed with inconsistent principles and viewpoints. Hence, existing industrial ontologies are incoherent and not suitable for the connected information goal.

With successes in developing coherent ontologies in the biological, biomedical, and financial domains [2-4], an effort called Industrial Ontologies Foundry (IOF) has been formed to pursue the same goal for the industrial manufacturing domain [5]. Modern manufacturing, particularly with today's complex cyber-physical products and materials, however, requires diverse disciplines of engineering, information technology, and management. This nature makes the scoping and development of coherent ontology for the industrial manufacturing domain a mountainous challenge. To manage scope and expectations, the IOF community has devised a proof-of-concept (POC) project with the aim to prove the viability and values of the endeavor. This paper describes the developments within the POC project.

The paper first describes the IOF formation, organizational structure, and aims of its subgroups. It then describes the POC process, discusses current results, and finally concludes with future plans.

2 IOF Formation and Organization

The first IOF workshop was organized in December 2016 at the National Institute of Standards and Technology (NIST), Gaithersburg, USA [6]. Thereafter, the community has had weekly conference calls and yearly workshops [7, 8].

Following the workshop in 2017, the IOF charter has been drafted and became available on its web site as a community draft [9]. One of the most important messages from the charter that makes IOF unique from other standard organizations that have published engineering ontologies is the intention of the IOF for its ontologies to be freely open. Although the current charter stops short at indicating a particular intellectual property licensing agreement, the community recognizes the need for such an encompassing ontology to be freely reusable, so that the corresponding information can be truly connected. The intention of the community has always been gravitating toward one of the royalty free licenses (e.g., variations of the creative commons licenses [10], which only require some recognitions).

The charter also outlines goals that include not only publishing the freely available ontologies, but also providing principles, guidelines, and governance processes such that a suite of ontology modules can grow in an interoperable fashion. One of the components to enable this is the organizational structure of the IOF community.

The community has devised three kinds of committees: a governance board (GB), a technical oversight board (TOB), and working groups (WGs). There is one GB, one TOB, and as many WGs as deemed necessary by the community. To ensure interoperability, these boards have overlapping personals as shown in Fig. 1, where each circle represents the membership of each board.

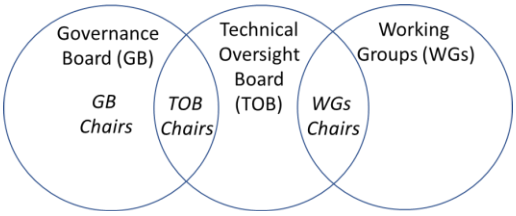

Fig. 1. IOF organizational structure

The primary role of the GB is to maintain the health and effective operation of the IOF organization. It sets the overall policy and manages legal aspects of the business. The other important role for the GB is to resolve conflicts unresolvable by the TOB.

TOB members are responsible for setting ontology principles and design guidelines used across the WGs. They have an important role to ensure that modules of the IOF ontologies developed by each WG are interoperable and consistent.

Each WG develops an ontology or a suite of ontologies of the IOF ontologies vetted by the TOB. Some WGs may be responsible for developing or adapting cross-cutting, domain independent ontologies such as for time or units of measurement. Fig. 2 shows types of ontologies anticipated within the IOF ontologies. A WGs may exist for the foundation ontology, each of bubbles, or a group of bubbles in the figure. The top two layers reflect specializations of the IOF ontologies for a particular use that may be private or licensed (developing such specializations is considered out of scope for the IOF).

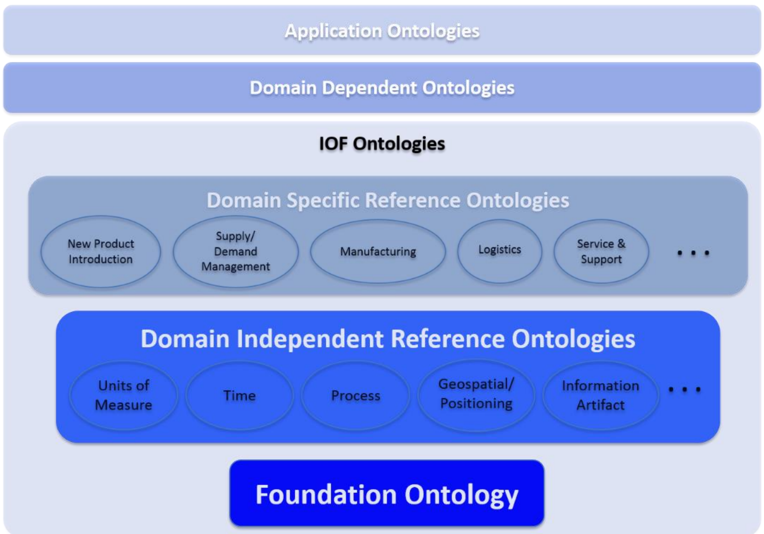

Fig. 2. Architecture of the IOF ontologies [9]

We can observe from Fig. 2 that it is a challenge to establish the scope of each WG, particularly at the Domain Specific Reference Ontologies level as information about a manufactured product overlaps across its life cycle phases and operational areas involved in development and manufacturing of the product. On the other hand, starting from bottom up looking at the choices of foundation ontology (e.g., BFO [11], DOLCE [**Error! Reference source not found.**]) and Domain Independent Reference Ontologies carry a risk of developing something don't meet the requirements of the domain. To manage such risk and challenge, the IOF community has agreed to start with a proof-of-concept (POC) project, which set the scope based on a consensus rather than the life cycle or operation areas of manufacturing and devised both bottom-up and top-down groups to test the feasibility and values of IOF ontologies. The next section describes this ongoing activity.

3 IOF POC

The IOF proof-of-concept (POC) project was intended to test the feasibility of IOF goals. Therefore, the objectives of the POC included not only producing a small, initial ontology, but also testing the organizational structure (described above) and producing and testing drafts principles and guidelines.

To set the scope, the POC started by asking for most interested manufacturing-related terms from the community. To set a very low bar, textual definitions were optional. Terms submitted should have a use case behind them so that a proof-of-value (POV) can be performed once the ontology is available. Each submission could include up to 50 terms in virtually any form of structured or unstructured file format. After collecting all the submissions, 20 terms were to be identified based on the frequencies of matches across the submissions. Each term has a (synonym) set of closely-matched terms; therefore, we call the output of this step the top-20 sets.

Getting the top-20 sets

At the time of this writing, 23 submissions were received, some of which went over the 50-term limit. Submission topic areas included product life cycle management, general manufacturing, manufacturing process, material (in the sense of material science), supply chain management, logistics, shop floor automation, manufacturing resources, production system engineering and analysis, and additive manufacturing.

To identify the top-20 sets, first submissions were transcribed into a web-based collaborative ontology editing tool [**Error! Reference source not found.**] using a Simple Knowledge Organization System (SKOS) representation [14]. In a few cases, submissions were provided in an Web Ontology Language (OWL) [15] format. The tool was able to import these, then SKOS concept assertions were manually added to their class URIs. In this way, all terms could be mapped using the SKOS *closeMatch* relationship.

This led to the second step, creating a rough mapping. The objective of this mapping was not to do semantic alignment, but rather to set a scope for the POC. Hence, mappers were instructed to be quite liberal with their mapping and to not be concerned with the concept hierarchy. One relative rich submission was copied as a starting point for a canonical terminology (i.e. the IOFPOC skos:ConceptScheme) for mapping. The mapping task was divided among four individuals who performed the mapping consecutively, and the target for completion was a few weeks.

Even with the rough mapping in mind, it was not a trivial work for the last individual because the canonical grew as unmapped terms were added to it. The last individual reported that the search functionality in the tool became very handy. At the finish, the canonical had over 600 terms. Table 1 below shows the resulting top-20 sets. Due to space limitation, only the terms and counts are shown. Table 2 shows definitions provided for the term set with the most matches. Complete result will be available on the IOF web site [5].

Table 1. The top-20 sets of closely matched terms

	Term Set	Count
1	Product, Physical Product, Product Material, Manufactured Product	13
2	Material, Material Object, Engineered Material	12
3	Manufacturing Machine, Processor, Machinery, Machine Tool, Machine, Mechanism, Machine, Workstation	10
4	Tool, Tools, Tooling, Manufacturing Tool	9
5	Assembly, Part, Composition	8
6	Part, Physical Part, Sub Assembly, Product Component, Component	8
7	Process, Transformation	8
8	Supplier, Supplier Provider or Vendor, Material Supplier	8
9	Transportation Process, Move, Movement, Transfer, Transport, Act of Transportation, Transport	8
10	Quality, Indicator, General KPI, Quality, Engineering Quality, Physical Quality	7
11	Requirement, Requirement Specification, Control	7
12	Assembly Process, Assembly Operation, Technological Pair Positioning, Joining, Act of Assembly	6
13	Customer, Business Customer	6
14	Feature, Materials Property, CAD Model Feature	6
15	Process Plan, Work Instructions, Manufacturing Method, Operation, Process Plan, Manufacturing Process Plan	6
16	Resource	6
17	Task, Activity, Operation	6
18	Design, Design Process	5
19	Equipment, Machinery	5
20	Fixture, Work Holder	5

Table 2. Textual descriptions for the most frequent set

Term	Textual Description
Physical Product	Subclass of spatial region (derived from the class axiom)
Product	Product (for manufacturing industry) is a Material Object, manufactured to satisfy a need of the market (e.g. to be sold in order to provide profit and support customers by covering their needs).
Product	the output of a manufacturing process
Product	a material entity or service that is developed to be sold
Product	This is a tangible object manufactured to satisfy a need of the market. For the specific mould maker, common products are: moulds, dies, and high precision parts
Product	material and/or service sold to others. Note that in manufacturing enterprises 'product' often refers to a product type or class, which in supply chains may be differentiated by packaging, but 'product' may also refer to a product instance.
Product	A goods, idea, method, information, object or service created as a result of a process and serves a need or satisfies a want. It has a combination of tangible and intangible attributes (benefits, features, functions, uses) that a seller offers a buyer for purchase.
Product Material	A material entity produced by man or machine, including raw material, parts, semi-finished product, and finished product
Manufactured Product	a product that is created via a manufacturing process.
Product	A product is the subject of the activity.
Product	desired output or by-product of the processes of an enterprise. Note 1 to entry: A product can be an intermediate product, end product, or finished goods from a business perspective. [SOURCE: IEC 62264-1:2013-01, 3.1.27]
Product	No definition provided.

Discussion

It can be seen from Table 1 that judging only by the label, the terms in each set are quite semantically close to each other. However, Table 2 demonstrates that each submission in the set gave a variety of definitions; and a few distinct notions (or concepts) may be refactored from the set. All but three of the submissions provided textual definitions; though two of those three provided subsumption hierarchies. Due to the richness of the semantics provided in the submission, the mapping result frequently yielded closely related notions within each set. That is, we observed that notions in the same set were often either close to each other in a subsumption hierarchy, an action or process of another notion in the set, or a mereologically related notion to another notion in the set. There are however other complex cases where similar terms are used for overlapping notions and similar terms are used for subtly distinct notions. An example of the former case can be observed in Table 2 where 'product' was used exclusively for physical object, both physical object and service, only desired (or designed) output, and both desired and undesired (by-product) output. An example of the latter case is 'material'. In one notion, it is a chemical composition, while in the other, it is a part, assembled component, or raw substance (e.g., metal powder) supplied to the manufacturing activity.

Next step

In the next step, the IOF community is creating a WG to formalize these top-20 sets. Two kinds of WGs are in the plan, top-down and bottom-up. At the time of this writing, a top-down WG has been formed. Its approach is to classify the notions in the top-20 set using the Basic Formal Ontology (BFO) [11] and conversely to see if BFO needs changes.

The formations of the bottom-up WGs are driven by use cases. At this time, the TOB is soliciting use cases from the community. With use case information, the TOB will cluster use cases together to form bottom-up WGs.

Both types of WGs will allow the submitters an opportunity to complete their submissions with detail textual definitions before the WGs proceed. It is expected that each bottom-up WGs would consider notions beyond the top-20 set. They will take the focal point at ensuring that the harmonized definitions or ontology in each bottom-up domain board satisfy their functional requirements and viewpoints. It is unclear at this point whether each bottom-up WG will produce a more formal definition (or model) such as class diagram or even an ontology. They will however perform gap analysis between their notions and the output from the top-down group by classifying their notions with respect to the top-down view. It is yet to be determined whether the gap analysis will be performed concurrently across all the bottom-up groups or sequentially. The community seems excited about what is coming before them.

4 Conclusion and Future Plans

The IOF is a growing community. After a round of inactivity elimination, over 60 participants have registered and are actively involved with representative across all populous continents except Africa. Both the contingent governance (5 seats) and technical oversight (12 seats) boards consist of diverse representatives from private companies, research institutes, academia, and standard development organizations. Both boards have a lot of deliverables lying ahead. The governance board should set up membership policies, development infrastructure, and intellectual property policy. The technical oversight board should draft ontology design rules (e.g., naming convention, minimal ontological commitment, URIs, versioning) and design principles (e.g., modularity, interaction with existing standards). It is anticipated that the proof-of-concept

(POC) project, which is still ongoing, will be the platform for developing these documents. At the present time there is no fee to participate in the IOF; however, the governance board is tasked with developing a business model to sustain the community given a successful POC. Interested individual are invited to submit a request for participation on the IOF web site.

Acknowledgement and a Disclaimer

The authors wish to thank all the Industrial Ontologies Foundry (IOF) members who have contributed their expertise and content to the founding of IOF and the POC project. Any mention of commercial products is for information only; it does not imply recommendation or endorsement by NIST

Reference

1. The Economist. April 21 issue on the Third Industrial Revolution (2012).
2. Open Biological and Biomedical Ontology Foundry Web Site – http://obofoundry.org, last access 2018/07/11.
3. Financial Industry Business Ontology press release - http://www.edmcouncil.org/downloads/20171026_FIBO_Release.pdf, last access 2018/07/11.
4. FIBO spec page - https://www.edmcouncil.org/financialbusiness, last access 2018/07/11.
5. Industrial Ontologies Foundry (IOF) Website - https://sites.google.com/view/industrialontologies/home, last access 2018/07/11.
6. Kulvatunyou, B. and Morris, K.: Working Towards an Industrial Ontology Foundry to Facilitate Interoperability - http://blog.mesa.org/2017/03/working-towards-industrial-ontology.html, last access 2018/07/11.
7. Ivezic, N., et al.: 2017 NIST/OAGi Workshop: Enabling Composable Service-Oriented Manufacturing Systems. NIST Advanced Manufacturing Series 100-15 (2018).
8. 2018 NIST/OAGi Workshop Enabling Composable Service-Oriented Manufacturing Systems - https://www.nist.gov/news-events/events/2018/04/2018-nistoagi-workshop-enabling-composable-service-oriented-manufacturing, last access 2018/07/11.
9. IOF Charter - https://sites.google.com/view/industrialontologies/about/charter, last accessed 2018/07/11
10. Creative Common Licenses - https://creativecommons.org/licenses/, last accessed 2018/07/11.
11. Smith, B., et al. Basic Formal Ontology Home Page - http://ontology.buffalo.edu/bfo/, last accessed 2018/07/11.
12. Gangemi A, et al.: Sweetening Ontologies with DOLCE. LNCS, vol 2473. Springer, Berlin, Heidelberg (2002).
13. Mobi - a decentralized, federated and distributed graph data platform for teams and communities - https://github.com/matonto/matontohttps://creativecommons.org/licenses/, last accessed 2018/07/11.
14. W3C: Simple Knowledge Organization System Reference (2009)
15. OWL Document Overview - https://www.w3.org/TR/owl2-overview/, last accessed 2018/07/11.