The Semantic Web in the Internet of Production: A Strategic Approach with Use-Case Examples

Johannes Lipp
(i) Chair of Databases and Information Systems
RWTH Aachen University, Aachen, Germany
(ii) Fraunhofer Institute for Applied Information Technology
Sankt Augustin, Germany
orcid.org/0000-0002-2639-1949
email: lipp@dbis.rwth-aachen.de

Katrin Schilling
Laboratory for Machine Tools
and Production Engineering (WZL)
RWTH Aachen University, Aachen, Germany
email: k.schilling@wzl.rwth-aachen.de

***Abstract*—The semantic interoperability of data, models, systems, and knowledge in general is a core element of the Internet of Production, i.e., a cross-life cycle and interdisciplinary networking of all levels in manufacturing technology. Semantic Web technologies are a good choice for the implementation of such applications, but, despite numerous academic research projects, its true potential is still rarely used in practice. One reason is the lack of knowledge among practitioners about both the technology itself and possible application areas, as manufacturing engineers usually are no Semantic Web experts and vice versa. In this paper, we present five essential application areas for Semantic Web technologies in production engineering, and give five examples of how we use these in practice in the Internet of Production. Our two-folded presentation intends to clarify potentials within application areas, and at the same time support the ramp-up of practical applications based on our examples.**

Keywords*—*Semantic Web; Internet of Production; Use-Cases.

I. Introduction

The Semantic Web [1] and its community proposed multiple recommendations and standards to improve semantic interoperability in the interconnected World Wide Web. It addresses, among others, the tasks of knowledge sharing, validation, and reasoning. Users can tackle these tasks via combinations of a broad range of solutions, including Persistent IDentifiers (PIDs), ontologies, data shapes, and reasoning rules. Semantic Web solutions in general do not intend to replace other solutions like relational databases or machine learning, but aim to cooperate closely with them.

In the field of production technology, the idea of the Semantic Web got attention in various, mostly academic, research projects. Unfortunately, these technologies have not yet reached a broad acceptance or implementation in industry. This is mainly due to lacking ontology knowledge among employees, missing tool support, imprecise problem statements in industry use-cases, and unclear benefits like the return on invest for the extensive modeling effort. These issues in industry range among multiple levels and domains, and effect engineers, domain experts, ontology engineers, and C-level managers. For these reasons, even though some interesting concepts and the technical feasibility were analyzed in demo implementations, hardly any application was properly realized in a productive system or product. Most applications in production did never leave an experimental stage.

There indeed are strong reasons to continue the research on Semantic Web technologies for production engineering. Following the achievements in the vision of ”Industry 4.0” in the recent years, a proper infrastructure – a basic prerequisite for a networked production – has been created. Nowadays, the latest generation of products in automation technology are equipped with the necessary interfaces and communication protocols to enable distributed, data-driven applications. In particular, this means that ”data” is now available outside the devices and applications with low effort. Availability and accessibility of data alone however are not sufficient to match the vision of the Internet of Production [2], which requires the networking of all systems and data-based optimization along in the entire production process. For example, with a higher level of maturity for knowledge-based applications like artificial intelligence, lifting simple data to proper knowledge is a crucial factor. This need for semantic technologies is supported by the increased attention for protocols like OPC Unified Architecture (OPC UA) [3] as they add semantics to data interactions and also support interoperability. We argue that the development of these solutions did not fully take into account the previous achievements by the Semantic Web and thus tend to partially re-invent the wheel.

Both aspects, the better availability of data through an advanced infrastructure of production systems, and the increasing demand for semantically described data for new applications, show that a Semantic Web in action is required by the Internet of Production. In this paper, a state of the art overview is in Section II, before, in Section III, we demonstrate the benefits of the Semantic Web for both ontology experts and non-experts in order to convince all above-mentioned users in a handy way. We subsequently in Section IV present concrete use-cases that we observed in the research project *Internet of Production* and thus support industry and institutes in planning their use-cases, before we conclude our work in Section V.

 ISBN: 978-1-61208-813-6

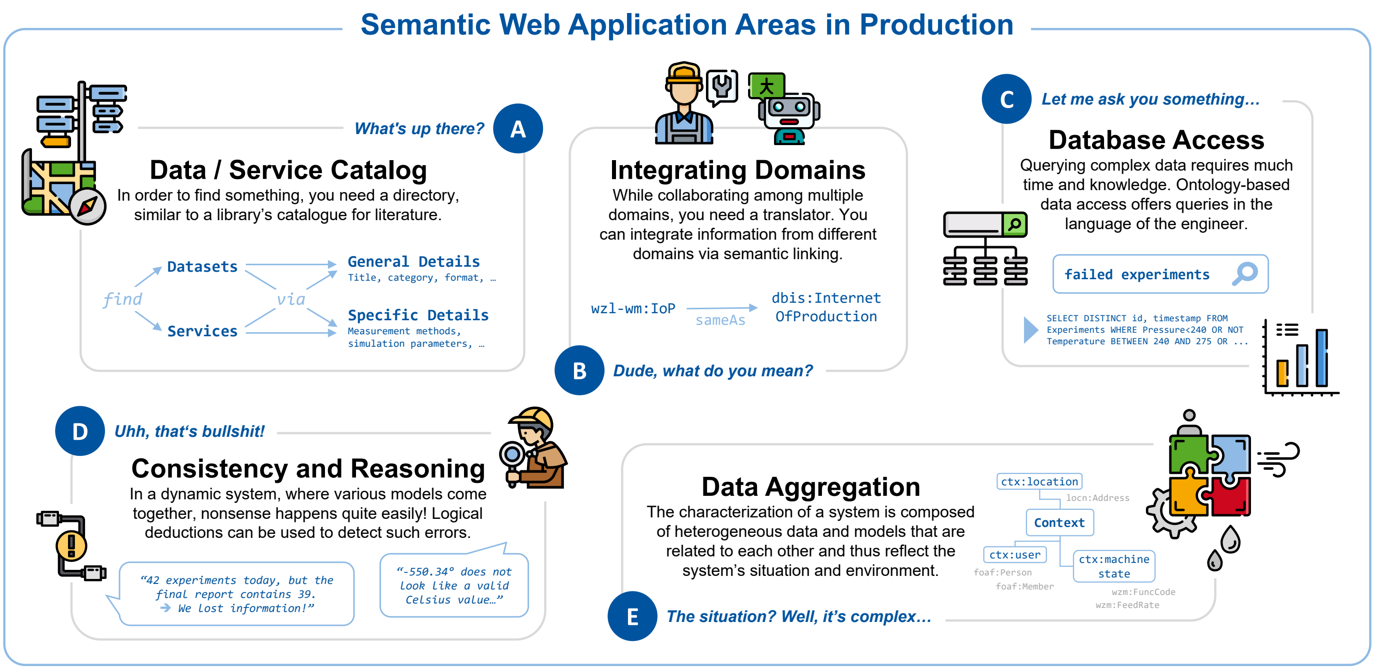

Fig. 1. Illustration of the main application areas for the Semantic Web in the context of production. This figure supports demonstrate the benefits to both experts and non-experts. These include, but are not limited to, the five areas data/service catalog, integrating domains, database access, consistency and reasoning, and data aggregation.

II. State of the Art

The potential of the Semantic Web idea in the context of production technology has been discussed in research projects. Upper ontologies for manufacturing, such as DOLCE [4], Cyc [5], SUMO [6] or MASON [7], as well as specific domain ontologies, were developed. An overview and comparison can be found in [8]. In particular, the challenge of breaking up silos and linking information across value chains is essential in the "Industry 4.0"; the concept of the administration shell is a concrete example [9]. It is still difficult to find these ontologies and reuse that work.

Semantic Web technologies have also been applied to solve a wide range of concrete research questions: From dynamic processor orchestration [10], over worker assistance [11], up to visualization via augmented reality [12], just to name a few examples. Furthermore, initiatives such as the Open Services for Lifecycle Collaboration (OSLC) try to establish the application in (engineering) tools through industry cooperations.

Unfortunately, all this knowledge and experience is still not well known outside these participating disciplines. Especially classical engineering disciplines, which have had little contact with software engineering and information modeling, often face difficulties in transferring the often abstract paradigms to problems in their own domain. Our goal is not to replace any established or advanced technologies of these experts with Semantic Web technologies. Rather, the intention is to support the use of Semantic Web technologies as a "glue" to connect the specific technologies and expert domains by providing a descriptive set of application areas.

III. Application Areas in Production

This section presents possible application areas for the Semantic Web in production and it is intended to give a high-level overview for all relevant people. The following explanations refer to the graphical overview shown in Figure 1 and which we use as a one-page flyer to advertise this at partners.

A **Data / Service Catalog** (A) probably is the mostly used application and is well-known among most people. It is a directory of any data sources of interest, such as datasets, services, programs, people, projects, or sensors. Such a catalog enables people to find information based on given search details. Prominent examples are open data portals such as [13] or [14], where users typically can apply a wide range of parameters to their search, including keywords, usage rules, and both spatial and temporal ranges. Another frequently applied example is the dynamic management of semantically described functions for a service-oriented / skill-based management of production processes, described in [10], [15].

A catalog usually is deployed independent from the data itself, which means that it can be easily applied to any existing data management system. Note that, as depicted in Figure 1, it supports searching for general filters as well as specific details. The former represents domain-independent information that can be applied to most catalogs and thus can and should be shared among these. Concretely, this means that catalog developers should reuse existing (de-facto) standards such as the Data CATalog Vocabulary (DCAT) [16] to enable smooth interoperability on this level between different catalogs. The latter, namely specific details, stands for information that is

 ISBN: 978-1-61208-813-6

particular for a certain domain. Defining these requires much communication between both disciplines Semantic Web and the domain, as only the annotated pieces of information can be included in search requests afterwards.

The area **Integrating Domains** (B) supports human understanding as well as interoperability on machine level. Since it is not useful to re-invent the wheel for each small area one works on, people tie together knowledge from different domains in order to represent their particular use-case best. Combining pieces of different knowledge sources such as ontologies usually leads to intersections or overlaps, which often are not clear for humans and machines. With methods from the Semantic Web, we can solve these issues by introducing relations like `sameAs`, `broader`, or `narrower` between concepts from different domains.

Ontology Based Data Access (OBDA) provides the potential of simple **Database Access** (C) on a semantic level. By defining basic concepts such as ”failed experiments” for a specific scenario, the domain expert (without extensive database knowledge) is enabled to easily articulate even complicated queries. The goal is to separate a user-friendly wording of the queries from the concrete database structure.

The potential to check semantically described data for **Consistency** and to derive insights through **Reasoning** (D) is beneficial in the complex ecosystem of production technology: Errors can rapidly occur during the transition between different applications, systems from multiple manufacturers, and various standards along the product’s life cycle. But even logical conflicts within the data sets can be identified.

The **Aggregation of Data** and models (E) enables the mapping of complex situations and environmental conditions based on heterogeneous information sources. Semantic relationships of potentially very different aspects characterizing a situation allow an abstraction of concepts (such as location, states, persons) with their individual representation, even if they are represented in different structures.

IV. Application Examples in the Internet of Production

This section presents five use-case examples we identified in the research project *Internet of Production* and which we will fully implement in the near future. They all originate from open research questions in different domains, and aim to improve existent processes in terms of usability, stability, speed, or precision. These concrete examples are intended to be used by researchers as models for any application area in the future, and might even be translated into archetypes. Figure 2 clearly illustrates which application areas from Figure 1 are covered by the five example applications. All examples cover one or two areas, and all areas are covered by at least one example.

Example 1: An example of an application for area (A) is the creation of a cross-disciplinary catalogue that provides a searchable overview of the various research activities and the data generated in the project. The catalogue makes it easier for researchers in computer science, mechanical engineering, economics and social sciences to find links between

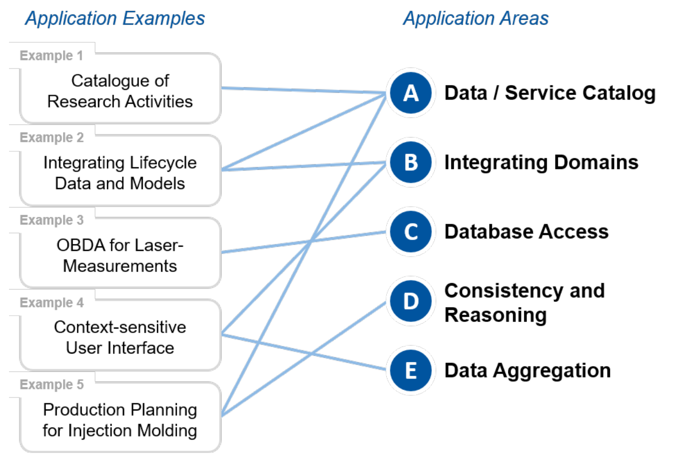

Fig. 2. Allocation of the use-case examples presented in Section IV to the application areas from Section III. Note that some examples cover multiple areas, and that all areas are covered.

(sub-) projects, solutions for similar problems or potential research partners. A concrete implementation plan includes both a distributed file system that stores the data, and Apache Jena Fuseki [17] metadata system that provides metadata management and a convenient query interface via the SPARQL Protocol and RDF Query Language (SPARQL) [18]. The catalogued information includes, but is not limited to, a vast amount of datasets consisting of sensor values collected from production machines or simulations. Required annotations and search filter in this example include responsible person, temporal characteristics, accrual periodicity, domain and file format. Note that providing data to others requires the data steward to add most of the above-mentioned annotations manually, as only some fields can be filled automatically. It is not a trivial task to motivate data providers to execute this step properly.

Example 2: The second application is the integration of different engineering models and to relate these with each other, which combines areas (A) and (B). In the product development process, a wide variety of models is created and their relations are mostly implicit knowledge only. Our partners asked for techniques to explicitly annotate important relations between models and query these afterwards. Please note that the models are very heterogeneous in terms of domain, file format, and level of detail. The file format, for instance, ranges from simulation scripts over 3D sketches to rich Computer-Aided Design (CAD) models. We tackle area (B) in this use-case by introducing a minimal ontology, which is depicted in Figure 3 and is aligned with DCAT. This ontology is used to properly relate models with each other, which includes to tell that (i) two models represent the same thing, (ii) an element in one model represents the same thing than one from another, or (iii) an element in a model or a complete model is more specific or general than another. The first realized concrete axioms state that a particular engine within a CAD model of a Audi A4 car represents the same as a blender 3D visualization’s part that depicts the engine of a Volkswagen Amarok. Since we did not only link models to

Copyright (c) IARIA, 2020. ISBN: 978-1-61208-813-6

each other via these properties, but also catalogize them in a model catalog, this use-case combines areas (A) and (B) and enables a holistic integration of individual (sub-)models along the life cycle.

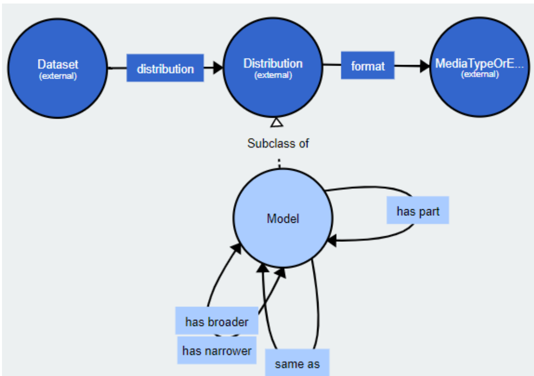

Fig. 3. The minimal ontology we created for the model catalog use-case. It is aligned with DCAT and represents any model as a distribution that has a file format (dark blue above). In this WebVOWL [19] screenshot, the light blue properties below show possible relations between models and elements within these, which are linked via the property *has part*.

Example 3: Another promising use-case is an implementation of the OBDA approach (C) for Ultrashort Pulse Laser Processing (USP). In this process, we record time series of a laser's three-dimensional position as well as temperature data of four locations, and store it in a relational database. Analyzing the data requires the USP domain experts to design complex Structured Query Language (SQL) queries, which however is not part of their expertise. We avoid the time-consuming and error-prone individual process via OBDA mappings and a minimal ontology, which are both designed cooperatively by the USP domain experts and ourselves. The current demonstrator can be queried locally via the Ontop plugin for Protégé and allows the engineers in particular query failed experiments, crucial temperature developments, stable runs etc. in their own wording via SPARQL. A full implementation of this use-case includes to identify and understand all existing SQL queries, create new ones where required, and specify proper OBDA mappings that are easy to understand for the end users.

Example 4: An example for the data aggregation (E) as well as the integration of domains (B) is a context-sensitive user interface that adapts the user's position to show relevant information regarding the nearest, dynamic environment. For this purpose, a predefined information object is labeled with contextual tags (e.g., a location, device category, user role, machine state). Depending on the user's devices (e.g., tablet or glasses) the localization can be determined in different ways: An indoor tracking system such as Bluetooth Low Energy beacons refers to the referencing anchors; image recognition enables tracking based on visual significant features in the environment; augmented reality frameworks (e.g., Google ARCore) combine multiple technologies and define virtual anchors. A semantic description of the spatial references links them to the concrete information object via the concept of localization. This is applied in the same way to other tags such as machine state or the user role.

Example 5: The last use-case we present models production planning, logistics, and control for injection molding in plastics processing. It combines areas (D) and (A), as we construct and manage both reasoning rules and instances, respectively. In this example, we together with the experts from plastics processing fully model the required complexity of production planning in this domain. That are in particular dependencies and consequences between possible choices, and support to infer new knowledge from given the input in form of annotated instances. Possible outcomes of this use-case include the ability to produce optimal production plans from given inputs, as well as to derive new knowledge in that area, which can be shared among humans and machines. In order to complete this, all necessary information on the machines' availability, incoming orders, and matching rules need to be extracted in a semi-automatic way from an Enterprise Resource Planning (ERP) system.

This section presented five exemplary use-cases that we observed, and which are intended to support future researchers in their tasks to leverage the Semantic Web in their projects. As shown in Figure 2, these examples cover one or two application areas from Figure 1 each, and all areas are covered. The presented examples tackle practical problems occurring in different domains, ranging from data access and management to analysis and reasoning. In the concrete implementation, domain experts work together with Semantic Web experts to build target-oriented solutions for practical use.

V. Conclusion and Future Work

In this paper we argued that, especially in the recent years, the push of the Semantic Web matches well with the pull from the ever growing amount of networked information sources in the Internet of Production. This leads to an increased need for an actual application of Semantic Web technologies within various domains including production. We grouped the major strengths of the Semantic Web in the production domain into five areas that are intended to support motivating these to different people in research and industry.

With five exemplary use-cases that we observed in the project *Internet of Production*, we demonstrate possible solutions and their effectiveness to future researchers. These use-cases show that a strong collaboration of experts from both the Semantic Web and the application domain is essential indeed. Our paper is a good step towards bridging these domains, as we showed important matches between possibilities on the one side and requirements in use-cases on the other side.

Future work includes to further design, implement, and document these five use-cases. Further leveraging the strengths of the Semantic Web and its community in production will enable a semantically interconnected Internet of Production. The importance of collaboration between experts from both fields remains, and is crucial to drive both domains semantics and production.

 ISBN: 978-1-61208-813-6

ACKNOWLEDGEMENT

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production – 390621612.

REFERENCES

[1] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, vol. 284, no. 5, pp. 34–43, 2001.

[2] J. Pennekamp et al., "Towards an Infrastructure Enabling the Internet of Production," in *2019 IEEE International Conference on Industrial Cyber Physical Systems (ICPS)*, May 2019, pp. 31–37.

[3] T. Hannelius, M. Salmenpera, and S. Kuikka, "Roadmap to Adopting OPC UA," in *2008 6th IEEE International Conference on Industrial Informatics*. IEEE, 2008, pp. 756–761.

[4] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider, "Sweetening Ontologies with DOLCE," in *International Conference on Knowledge Engineering and Knowledge Management*. Springer, 2002, pp. 166–181.

[5] D. B. Lenat and R. V. Guha, *Building Large Knowledge-based Systems; Representation and Inference in the Cyc Project*. Addison-Wesley Longman Publishing Co., Inc., 1989.

[6] I. Niles and A. Pease, "Towards a Standard Upper Ontology," in *Proceedings of the international conference on Formal Ontology in Information Systems-Volume 2001*, 2001, pp. 2–9.

[7] S. Lemaignan, A. Siadat, J. Dantan, and A. Semenenko, "Mason: A proposal for an ontology of manufacturing domain," in *IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications (DIS'06)*, 2006, pp. 195–200.

[8] N. Khilwani, J. A. Harding, and A. K. Choudhary, "Semantic Web in Manufacturing," *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, vol. 223, no. 7, pp. 905–924, 2009.

[9] I. Grangel-González, L. Halilaj, S. Auer, S. Lohmann, C. Lange, and D. Collarana, "An RDF-based Approach for Implementing Industry 4.0 Components with Administration Shells," in *2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2016, pp. 1–8.

[10] M. Loskyll, J. Schlick, S. Hodek, L. Ollinger, T. Gerber, and B. Pîrvu, "Semantic Service Discovery and Orchestration for Manufacturing Processes," in *ETFA2011*, 2011, pp. 1–8.

[11] D. Gorecky, M. Loskyll, and C. Stahl, "Semantic Digital Factory – Using Engineering Knowledge to Create Ontologies for Virtual Training," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 7825 – 7830, 2014, 19th IFAC World Congress. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1474667016428454

[12] I. Mizutani, M. Kritzler, K. Garcia, and F. Michahelles, "Intuitive Interaction with Semantics Using Augmented Reality: A Case Study of Workforce Management in an Industrial Setting," in *Proceedings of the Seventh International Conference on the Internet of Things*, ser. IoT '17. New York, NY, USA: Association for Computing Machinery, 2017. [Online]. Available: https://doi.org/10.1145/3131542.3131550

[13] "GOVData," https://govdata.de, accessed: 2020-10-21.

[14] "EU Open Data Portal," https://data.europa.eu/euodp, accessed: 2020-10-21.

[15] C. Brecher, E. Kusmenko, A. Lindt, B. Rumpe, S. Storms, S. Wein, M. von Wenckstern, and A. Wortmann, "Multi-Level Modeling Framework for Machine as a Service Applications Based on Product Process Resource Models," in *Proceedings of the 2nd International Symposium on Computer Science and Intelligent Control*, ser. ISCSIC '18. New York, NY, USA: Association for Computing Machinery, 2018. [Online]. Available: https://doi.org/10.1145/3284557.3284714

[16] R. Albertoni, D. Browning, S. Cox, A. G. Beltran, A. Perego, and P. Winstanley, "Data Catalog Vocabulary (DCAT)-Version 2," *W3C Recommendation*, vol. 3, 2019, [Accessed 2020.09.15].

[17] The Apache Software Foundation, "Apache Jena Fuseki," https://jena.apache.org/documentation/fuseki2/, 2011, [Accessed 2020.09.15].

[18] E. Prud'hommeaux and A. Seaborne, "SPARQL 1.1 language to query and manipulate RDF graph content, W3C Recommendation, 21 March 2013," https://www.w3.org/TR/sparql11-overview/, 2013, [Accessed 2020.09.15].

[19] S. Lohmann, V. Link, E. Marbach, and S. Negru, "WebVOWL: Web-based Visualization of Ontologies," in *International Conference on Knowledge Engineering and Knowledge Management*. Springer, 2014, pp. 154–158.