

Capturing Industrial Information Models with Ontologies and Constraints

Evgeny Kharlamov¹(✉), Bernardo Cuenca Grau¹, Ernesto Jiménez-Ruiz¹,
Steffen Lamparter², Gulnar Mehdi², Martin Ringsquandl², Yavor Nenov¹,
Stephan Grimm², Mikhail Roshchin², and Ian Horrocks¹

¹ University of Oxford, Oxford, UK

evgeny.kharlamov@cs.ox.ac.uk

² Siemens AG, Corporate Technology, Munich, Germany

Abstract. This paper describes the outcomes of an ongoing collaboration between Siemens and the University of Oxford, with the goal of facilitating the design of ontologies and their deployment in applications. Ontologies are often used in industry to capture the conceptual information models underpinning applications. We start by describing the role that such models play in two use cases in the manufacturing and energy production sectors. Then, we discuss the formalisation of information models using ontologies, and the relevant reasoning services. Finally, we present SOMM—a tool that supports engineers with little background on semantic technologies in the creation of ontology-based models and in populating them with data. SOMM implements a fragment of OWL 2 RL extended with a form of integrity constraints for data validation, and it comes with support for schema and data reasoning, as well as for model integration. Our preliminary evaluation demonstrates the adequacy of SOMM’s functionality and performance.

1 Introduction

Software systems in the domain of industrial manufacturing have become increasingly important in recent years. Production machines, such as assembly line robots or industrial turbines, are equipped with and controlled by complex and costly pieces of software; according to a recent survey, over 40% of the total production cost of such machines is due to software development and the trend is for this number only to continue growing [35]. Additionally, many critical tasks within business, engineering, and production departments (e.g., control of production processes, resource allocation, reporting, business decision making) have also become increasingly dependent on complex software systems.

Recent global initiatives such as Industry 4.0 [9, 18, 34] aim at the development of *smart factories* based on fully computerised, software-driven, automation of production processes and enterprise-wide integration of software components. In smart factories, software systems monitor and control physical

This work was partially funded by the Royal Society under a University Research Fellowship, the EU project Optique [24] (FP7-ICT-318338), and the EPSRC projects MaSI3, DBOnto, ED3.

processes, effectively communicate and cooperate with each other as well as with humans, and are in charge of making decentralised decisions. The success of such ambitious initiatives relies on the seamless (re)development and integration of software components and services. This poses major challenges to an industry where software systems have historically been developed independently from each other.

There has been a great deal of research in recent years investigating key aspects of software development in industrial manufacturing domains, including life-cycle costs, dependability, compatibility, integration, and performance (e.g., see [41] for a survey). This research has highlighted the need for enterprise-wide *information models*—machine-readable conceptualisations describing the functionality of and information flow between different assets in a plant, such as equipment and production processes. The development information models based on ISA and IEC standards¹ has now become a common practice in modern companies [30] and Siemens is not an exception in this trend.

In practice, however, many types of models co-exist, and applications typically access data from different kinds of machines and processes designed according to different models. These information models have been independently developed in different (often incompatible) formats using different types of proprietary software; furthermore, they may not come with a well-defined semantics, and their specification can be ambiguous. As a result, model development, maintenance, and integration, as well as data exchange and sharing pose major challenges in practice.

Adoption of semantic technologies has been a recent development in many large companies such as IBM [11], the steel manufacturer Arcelor Mittal [2], the oil and gas company Statoil [21], and Siemens [1, 4, 19, 20, 22, 25, 32]. An important application of these technologies has been the formalisation of information models using OWL 2 ontologies and the use of RDF for storing application data. OWL 2 provides a rich and flexible modelling language that seems well-suited for describing industrial information models: it not only comes with an unambiguous, standardised, semantics, but also with a wide range of tools that can be used to develop, validate, integrate, and reason with such models. In turn, RDF data can not only be seamlessly accessed and exchanged, but also stored directly in highly scalable RDF triple stores and effectively queried in conjunction with the available ontologies. Moreover, legacy and other data that must remain in its original format and cannot be transformed into RDF can be virtualised as RDF using ontologies following the Ontology-Based Data Access (OBDA) approach [21, 23, 29].

In this paper, we describe the outcomes of an ongoing collaboration between Siemens Corporate Technology in Munich and the University of Oxford, with the goal of facilitating deployment of ontology-based industrial information models. We start by describing the key role that information models play in two use cases in the manufacturing and energy production sectors. Then, we present industrial information models that are used for describing manufacturing and

¹ International Society of Automation and International Electrotechnical Commission.

energy plants, and discuss how they can be captured using ontologies. In our discussion, we stress the modelling choices made when formalising these models as ontologies and identify the key OWL constructs required in this setting. Our analysis revealed the need for integrity constraints for data validation [27,37], which are not available in OWL 2. Hence, we discuss in detail what kinds of constraints are needed in industrial use cases and how to incorporate them. We then illustrate the use of reasoning services, such as concept satisfiability, data constraint validation, and query answering for addressing Siemens' application requirements.

Ontologies are currently being created and maintained in Siemens by qualified R&D personnel with expertise in ontology languages and ontology engineering. In order to widen the scope of application of semantic technologies in the company it is crucial to make ontology development accessible to other teams of engineers. To this end, we have developed the Siemens-Oxford Model Manager (SOMM)—a tool that has been designed to fulfil industrial requirements and which supports engineers with little background on semantic technologies in the creation and use of ontologies. SOMM provides a simple interface for ontology development and enables the introduction of instance data via automatically generated forms that are driven by the ontology and which help minimising errors in data entry. SOMM implements a fragment of the OWL 2 RL profile [26] extended with database integrity constraints for data validation; the supported language is sufficient to capture the main features of ISA and ICE based information models used by Siemens. SOMM is built on top of Web-Protégé [40], which provides built-in functionality for ontology versioning and collaborative development. It relies on HermiT [10] for ontology classification and LogMap [16] to support model alignment and merging. For query answering and constraint validation, SOMM requires a connection to a triple store or a rule inference system that supports Datalog reasoning and stratified negation-as-failure.

We showcase the practical benefits of our tool using two ontologies in the manufacturing and power generation domains. Both ontologies have been developed using SOMM by Siemens engineers to capture information models currently in use. Based on these ontologies, we conducted an empirical evaluation of SOMM's performance in supporting constraint validation and query answering over realistic manufacturing and gas turbine data. In our experiments, we coupled SOMM with the rule inference engine IRIS [3], which is available under the LGPL license.² Our evaluation demonstrates the adequacy of SOMM's functionality and performance for industrial applications.

2 Industrial Information Models

Conceptual information models can be exploited in a wide range of manufacturing and energy production applications. In this Section, we discuss two concrete use cases and describe the underpinning models and their limitations.

² <http://www.iris-reasoner.org/>.

2.1 Applications in Manufacturing and Energy Production

In manufacturing and energy production plants it is essential that all processes and equipment run smoothly and without interruptions.

In a typical manufacturing plant, data is generated and stored whenever a piece of equipment consumes material or completes a task. This data is then accessed by plant operators using *manufacturing execution systems (MES)*—software programs that steer the production in a manufacturing plant. MESs are responsible for keeping track of the material inventory and tracing their consumption, thus ensuring that equipment and materials needed for each process are available at the relevant time [30]. Similarly, turbines in energy plants are equipped with sensors that are continuously generating data. This data is consumed by *remote monitoring systems (RMS)*, which analyse turbine data to prevent faults, report anomalies and ensure that the turbines operate without interruption. In both application scenarios, the use of information models is twofold.

1. Models are used to provide machine-readable specifications for the data generated by equipment and processes, and for the data flow across assets and processes in a plant.
2. Models provide a schema for constructing and executing complex queries. In particular, monitoring tasks in MESs are realised by means of queries issued to production machines and data hubs; similarly, anomaly detection in an RMS relies on queries spanning the structure of the turbines, the readings of their sensors, and the configuration of turbines within a plant.

2.2 Information Models Based on Industrial Standards

We next describe the information models in Siemens relevant to the aforementioned applications. These models have been developed in compliance with ISA, IEC, and ISO/TS international standards.

Manufacturing Models. For many manufacturing applications it is a common practice to rely on information models that are based on the international standard ISA-88/95.

The ISA-88/95 standard provides general guidelines for specifying the functionality of and interface between manufacturing software systems. The standard consists of UML-like diagrammatic descriptions accompanied with tables and unstructured text, which are used to extend the diagrams with additional information and examples. Figure 1 presents an excerpt of the ISA-88/95 standard modelling materials, equipment, personnel, and processes in a plant. For instance, one of these diagrams establishes that pieces of equipment can be composed by other pieces of equipment and are described by a number of specified ‘equipment properties’. The table complementing this diagram indicates that each piece of equipment must have a numeric ID and may have a textual

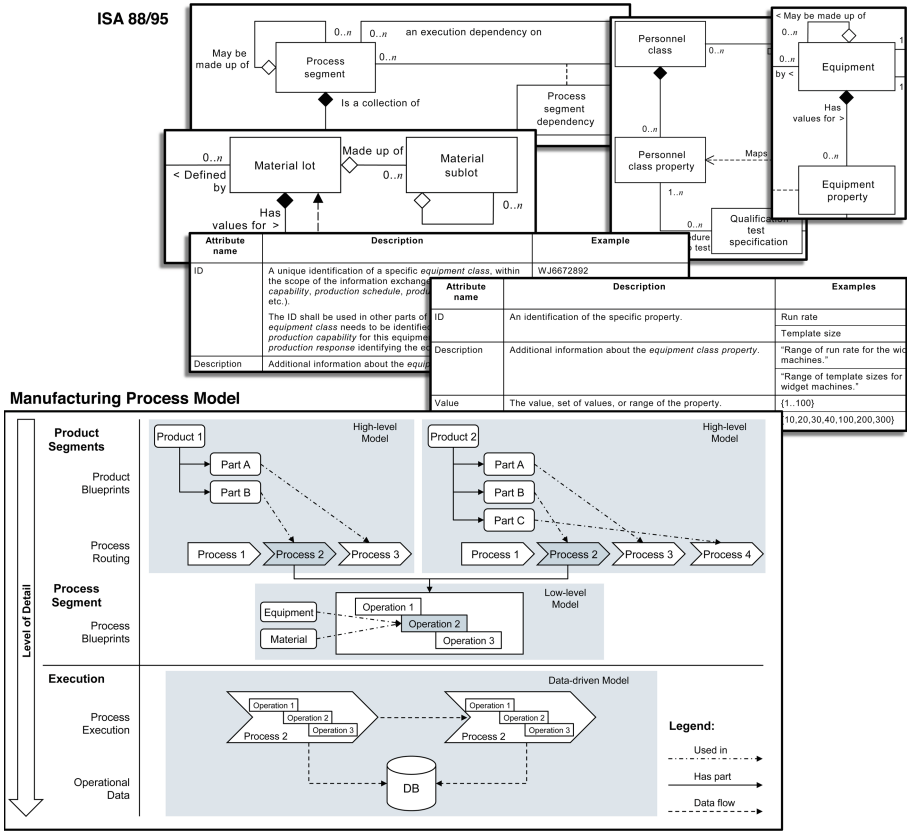


Fig. 1. Fragment of ISA 88/95 and an example model based on it.

description; additional properties of equipment can be introduced by providing an ID, a textual description of the property, and a value range.

Figure 1 provides a simplified version of an information model based on the standard ISA-88/95. The model is organised in three layers: *product*, *process*, and *execution*. On the product level, we can see the specification of two products and their relationship to production processes; for instance, *Product1* consists of *PartA* and *PartB*, which are manufactured by two consecutive processes. The process segment level provides more fine-grained specifications of the structure of each process; for instance, *Process2* consists of three operations, where the second one relies on specific kinds of materials and equipment. Finally, at the execution level, we can see how data is stored and accessed by individual processes.

Energy Plant Models. Information models for energy plants are often based on the *Reference Designation System for Power Plants (RDS-PP)* and *Kraftwerk-*

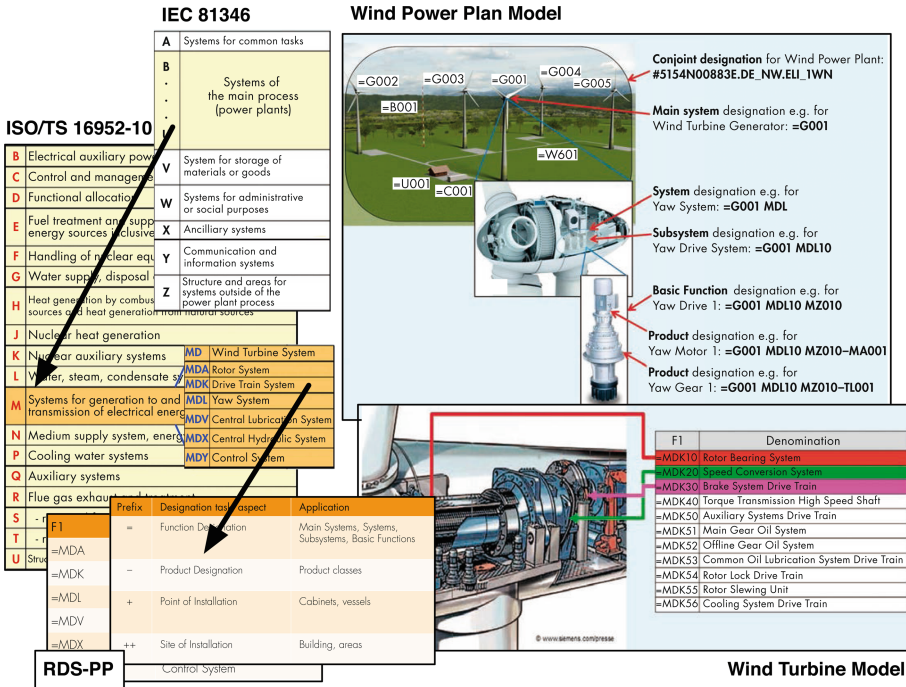


Fig. 2. Designation models IEC 81346, ISO/TS 16952-10, and RDS-PP and example information model for an energy plant [31].

Kennzeichensystem (KKS) standards, which are in turn extensions for the energy sector of the IEC 81346 and ISO/TS 16952-10 international standards.

IEC 81346 and ISO/TS 16952-10 provide a generic dictionary of codes for designating and classifying industrial equipment. Figure 2 provides an excerpt of these standards and their dependencies. For instance, in IEC-81346 letters ‘B’ to ‘U’ are used for generically designating systems in power plants. ISO/TS 16952-10 makes this specification more precise by indicating, for example, that letter ‘M’ refers to systems for generating and transmitting electricity, and that we can append ‘D’ to ‘M’ to refer to a wind turbine system. RDS PP and KKS provide a more extensive vocabulary of codes for equipment, their functionality and locations, as well as a system for combining such codes.

A typical energy plant model describes the structure of a plant by providing the functionality and location of each equipment component using RDS PP and KKS codes. Having this information in a machine-readable format is important for planning and construction, as well as for the software-driven operation and maintenance of the plant. Figure 2 shows how a specific plant is represented in a model; for instance, code =G001 MDL10 denotes that the yaw drive system number 10 of type MDL is located in the wind turbine generator number 001.

2.3 Technical Challenges

The development and use of information models in practice poses major challenges.

1. Model development is costly, as it requires specialised training and proprietary tools; as a result, model development often cannot keep up with the arrival of new equipment and introduction of new processes.
2. Models are difficult to integrate and share since they are often independently developed using different types of proprietary software and they are based on incompatible data formats.
3. Monitoring queries are difficult to compose and execute on top of information models: they must comply with the requirements of the models (e.g., refer to specific codes in the energy use case), and their execution requires access to heterogeneous data from different machines and processes.

In order to overcome these challenges Siemens has recently applied semantic technologies in a number of applications [13, 15, 19, 22, 32]. In particular, OWL 2 has been used for describing information models. The choice of OWL 2 is not surprising since it provides a rich and flexible modelling language that is well suited for addressing the aforementioned challenges: it comes with an unambiguous, standardised semantics, and a wide range of tools and infrastructure. Moreover, RDF provides a unified data exchange format, which can be used to seamlessly access and exchange data, and hence facilitate monitoring tasks based on complex queries.

3 From Information Models to Ontologies and Constraints

In this section we describe the ontologies that we have developed to capture manufacturing and energy production models presented in Sect. 2. The goal of our ontologies is to eventually replace their underpinning models in applications. Thus, their design has been driven towards fulfilling the same purposes as the models they originate from; that is, to act as schema-level templates for data generation and exchange, and to enable the formulation and execution of monitoring queries.

The representation of industrial information models and standards using ontologies has been widely acknowledged as a non-trivial task [5, 12, 14, 36]. In Sect. 3.1 we discuss the modelling choices underpinning the design of our ontologies and identify a fragment of OWL 2 RL that is sufficient to capture the basic aspects of the information models. Our analysis of the models, however, also revealed the need to incorporate database integrity constraints for data validation, which are not supported in OWL 2 [27, 37]. Thus, we also discuss the kinds of constraints that are relevant to our applications.

Finally, in Sect. 3.2 we discuss how the OWL 2 RL axioms and integrity constraints can be captured by means of rules with stratified negation for the

purpose of data validation and query answering. We assume basic familiarity with Datalog—the rule language underpinning OWL 2 RL and SWRL—as well as with stratified negation-as-failure (see [6] for an excellent survey on Logic Programming).

3.1 Modelling

From an ontological point of view, most building blocks of the the typical industrial information models are rather standard in conceptual design and naturally correspond to OWL 2 classes (e.g., *Turbine*, *Process*, *Product*), object properties (e.g., *hasPart*, *hasFunction*, *locatedIn*) and data properties (e.g., *ID*, *hasRotorSpeed*).

The main challenge that we encountered was to capture the constraints of the models using ontological axioms. We next describe how this was accomplished using a combination of OWL 2 RL axioms and integrity constraints.

Standard OWL 2 RL Axioms. The specification of the models suggests the arrangement of classes and properties according to subsumption hierarchies, which represent the skeleton of the model and establish the basic relationships between their components. For instance, in the energy plant model a *Turbine* is specified as a kind of *Equipment*, whereas *hasRotorSpeed* is seen as a more specific relation than *hasSpeed*. The models also suggest that certain properties must be declared as transitive, such as *hasPart* and *locatedIn*. Similarly, certain properties are naturally seen as inverse of each other (e.g., *hasPart* and *partOf*). These requirements are easily modelled in OWL 2 using the following axioms written in functional-style syntax:

$$\text{SubClassOf}(\textit{Turbine} \textit{Equipment}) \quad (1)$$

$$\text{SubDataPropertyOf}(\textit{hasRotorSpeed} \textit{hasSpeed}) \quad (2)$$

$$\text{TransitiveObjectProperty}(\textit{hasPart}) \quad (3)$$

$$\text{InverseObjectProperties}(\textit{hasPart} \textit{partOf}) \quad (4)$$

These axioms can be readily exploited by reasoners to support query answering; e.g., when asking for all equipment with a rotor, one would expect to see all turbines that contain a rotor as a part (either directly or indirectly).

Additionally, the models describe *optional relationships* between entities. In the manufacturing model certain materials are optional to certain processes, i.e., they are compatible with the process but they are not always required. Similarly, certain processes can optionally be followed by other processes (e.g., conveying may be followed by packaging). Universal (i.e., *AllValuesFrom*) restrictions are well-suited for attaching an optional property to a class. For instance, the axiom

$$\text{SubClassOf}(\textit{Conveying} \text{ObjectAllValuesFrom}(\textit{followedBy} \textit{Packaging})) \quad (5)$$

states that only packaging processes can follow conveying processes; that is, a conveying process can be either terminal (i.e., not followed by any other process)

or it is followed by a packaging process. As a result, when introducing a new conveying process we are not forced to provide a follow-up process, but if we do so it must be an instance of *Packaging*.

All the aforementioned types of axioms are included in the OWL 2 RL profile. This has many practical advantages for reasoning since OWL 2 RL is amenable to efficient implementation using rule-based technologies.

Constraint Axioms. In addition to optional relationships, the information models from Sect. 2 also describe relationships that are inherently *mandatory*, e.g., when introducing a new turbine, the energy model requires that we also provide its rotors.

This behaviour is naturally captured by an integrity constraint: whenever a turbine is added and its rotors are not provided, the application should flag an error. Integrity constraints are not supported in OWL 2; for instance, the axiom

$$\text{SubClassOf}(\textit{Turbine} \text{ ObjectSomeValuesFrom}(\textit{hasPart} \textit{Rotor})) \quad (6)$$

states that every turbine must contain a rotor as a part; such rotor, however, can be possibly unknown or unspecified.

The information models also impose cardinality restrictions on relationships. For instance, each double rotor turbine in the energy plant model is specified as having exactly two rotors. This can be modelled in OWL 2 using the axioms

$$\text{SubClassOf}(\textit{TwoRotorTurbine} \text{ ObjectMinCardinality}(2 \textit{ hasPart} \textit{Rotor})) \quad (7)$$

$$\text{SubClassOf}(\textit{TwoRotorTurbine} \text{ ObjectMaxCardinality}(2 \textit{ hasPart} \textit{Rotor})) \quad (8)$$

Such cardinality restrictions are interpreted as integrity constraints in many applications: when introducing a specific double rotor turbine, the model requires that we also provide its two rotors. The semantics of axioms (7) and (8) is not well-suited for this purpose: on the one hand, (7) does not enforce a double rotor turbine to explicitly contain any rotors at all; on the other hand, if more than two rotors are provided, then (8) non-deterministically enforces at least two of them to be equal.

There have been several proposals to extend OWL 2 with integrity constraints [27, 37]. In these approaches, the ontology developer explicitly designates a subset of the OWL 2 axioms as constraints. Similarly to constraints in databases, these axioms are used as checks over the given data and do not participate in query answering once the data has been validated. The specifics of how this is accomplished semantically differ amongst each of the proposals; however, all approaches largely coincide if the standard axioms are in OWL 2 RL.

3.2 Data Validation and Query Answering

Our approach to data validation and query answering follows the standard approaches in the literature [27, 37]: given a query Q , dataset \mathcal{D} , and OWL 2 ontology \mathcal{O} consisting of a set \mathcal{S} of standard OWL 2 RL axioms and a set \mathcal{C} of axioms marked as constraints, we proceed according to Steps 1–4 given next.

Table 1. OWL 2 RL axioms as rules. All entities mentioned in the axioms are named. By abuse of notation, we use `SubPropertyOf` and `AllValuesFrom` to refer to both their Object and Data versions in functional syntax.

OWL 2 Axiom	Datalog Rules
<code>SubClassOf(A B)</code>	$B(?x) \leftarrow A(?x)$
<code>SubPropertyOf($P_1 P_2$)</code>	$P_2(?x, ?y) \leftarrow P_1(?x, ?y)$
<code>TransitiveObjectProperty(P)</code>	$P(?x, ?z) \leftarrow P(?x, ?y) \wedge P(?y, ?z)$
<code>InverseObjectProperties(P_1, P_2)</code>	$P_2(?y, ?x) \leftarrow P_1(?x, ?y)$ and $P_1(?y, ?x) \leftarrow P_2(?x, ?y)$
<code>SubClassOf($A \text{ AllValuesFrom}(P B)$)</code>	$B(?y) \leftarrow P(?x, ?y) \wedge A(?x)$

1. Translate the standard axioms \mathcal{S} into a Datalog program $\Pi_{\mathcal{S}}$ using the well-known correspondence between OWL 2 RL and Datalog.
2. Translate the integrity constraints \mathcal{C} into a Datalog program $\Pi_{\mathcal{C}}$ with stratified negation-as-failure containing a distinguished binary predicate *Violation* for recording the individuals and axioms involved in a constraint violation.
3. Retrieve and flag all integrity constraint violations. This can be done by computing the extension of the *Violation* predicate.
4. If no constraints are violated, answer the user's query Q using the query answering facilities provided by the reasoner.

Steps 3 and 4 can be implemented on top of RDF triple stores with support for OWL 2 RL and stratified negation (e.g., [28]), as well as on top of generic rule inference systems (e.g., [3]). In the remainder of this Section we illustrate Steps 1 and 2, where standard axioms and constraints are translated into rules.

Standard Axioms. Table 1 provides the standard OWL 2 RL axioms needed to capture the information models of Sect. 2 and their translation into negation-free rules. In particular, the axioms (1)–(5) are equivalent to the following rules:

$$\textit{Equipment}(?x) \leftarrow \textit{Turbine}(?x) \quad (9)$$

$$\textit{hasSpeed}(?x, ?y) \leftarrow \textit{hasRotorSpeed}(?x, ?y) \quad (10)$$

$$\textit{hasPart}(?x, ?z) \leftarrow \textit{hasPart}(?x, ?y) \wedge \textit{hasPart}(?y, ?z) \quad (11)$$

$$\textit{Packaging}(?y) \leftarrow \textit{Conveying}(?x) \wedge \textit{followedBy}(?x, ?y) \quad (12)$$

Constraint Axioms. Table 2 provides the constraint axioms required to capture the models of Sect. 2 together with their translation into rules with negation. Our translation assigns a unique id to each individual axiom marked as an integrity constraint in the ontology, and it introduces predicates not occurring in the ontology in the heads of all rules. Constraint violations are recorded using the fresh predicate *Violation* relating individuals to constraint axiom ids.

The constraint (6) from Sect. 3.1 is captured by the following rules:

$$\textit{hasPart_Rotor}(?x) \leftarrow \textit{hasPart}(?x, ?y) \wedge \textit{Rotor}(?y) \quad (13)$$

$$\textit{Violation}(?x, \alpha) \leftarrow \textit{Turbine}(?x) \wedge \mathbf{not} \textit{hasPart_Rotor}(?x) \quad (14)$$

Table 2. Constraints axioms as rules. All entities are named, $n \geq 1$, and α is the unique id for the given constraint. *SomeValuesFrom*, *HasValue*, *FunctionalProperty*, *MaxCardinality* and *MinCardinality* denote both their Object and Data versions.

OWL Axiom	Datalog rules
SubClassOf(A SomeValuesFrom(R B))	$R_B(?x) \leftarrow R(?x, ?y) \wedge B(?y)$ and $Violation(?x, \alpha) \leftarrow A(?x) \wedge \mathbf{not} R_B(?x)$
SubClassOf(A HasValue(R b))	$Violation(?x, \alpha) \leftarrow A(?x) \wedge \mathbf{not} R(?x, b)$
FunctionalProperty(R)	$R_2(?x) \leftarrow R(?x, ?y_1) \wedge R(?x, ?y_2) \wedge$ $\mathbf{not} owl:sameAs(?y_1, ?y_2)$ and $Violation(?x, \alpha) \leftarrow R_2(?x)$
SubClassOf(A MaxCardinality(n R B))	$R_-(n+1)_B(?x) \leftarrow \bigwedge_{1 \leq i \leq n+1} (R(?x, ?y_i) \wedge B(?y_i))$ $\bigwedge_{1 \leq i < j \leq n+1} (\mathbf{not} owl:sameAs(?y_i, ?y_j))$ and $Violation(?x, \alpha) \leftarrow A(?x) \wedge R_-(n+1)_B(?x)$
SubClassOf(A MinCardinality(n R B))	$R_n_B(?x) \leftarrow \bigwedge_{1 \leq i \leq n} (R(?x, ?y_i) \wedge B(?y_i))$ $\bigwedge_{1 \leq i < j \leq n} (\mathbf{not} owl:sameAs(?y_i, ?y_j))$ and $Violation(?x, \alpha) \leftarrow A(?x) \wedge \mathbf{not} R_n_B(?x)$

Rule (13) identifies all individuals with a rotor as a part, and stores them as instances of the auxiliary predicate *hasPart_Rotor*. In turn, Rule (14) identifies all turbines that are not known to be instances of *hasPart_Rotor* (i.e., those with no known rotor as a part) and links them to the constraint α they violate.

Integrity constraints based on cardinalities require the use of the OWL 2 equality predicate *owl:sameAs*. For instance, the constraint axiom (7) from Sect. 3.1, to which we assign the id β_1 , is translated into the following rules:

$$\begin{aligned}
 hasPart_2_Rotor(?x) &\leftarrow \bigwedge_{1 \leq i \leq 2} (hasPart(?x, ?y_i) \wedge Rotor(?y_i)) \wedge \\
 &\quad \wedge (\mathbf{not} owl:sameAs(?y_1, ?y_2)) \\
 Violation(?x, \beta_1) &\leftarrow TwoRotorTurbine(?x) \wedge \mathbf{not} hasPart_2_Rotor(?x)
 \end{aligned}$$

The first rule infers an instance of the auxiliary predicate *hasPart_2_Rotor* if it is connected to two instances of *Rotor* that are not known to be equal; in turn, the second rule infers that all instances of *TwoRotorTurbine* that are not known to be instances of the auxiliary predicate violate the constraint (7). Similarly, axiom (8), to which we assign the id β_2 , is translated as follows:

$$\begin{aligned}
 hasPart_3_Rotor(?x) &\leftarrow \bigwedge_{1 \leq i \leq 3} (hasPart(?x, ?y_i) \wedge Rotor(?y_i)) \wedge \\
 &\quad \wedge \bigwedge_{1 \leq i < j \leq 3} (\mathbf{not} owl:sameAs(?y_i, ?y_j)) \\
 Violation(?x, \beta_2) &\leftarrow TwoRotorTurbine(?x) \wedge hasPart_3_Rotor(?x)
 \end{aligned}$$

Analogously to the previous case, the first rule infers that an individual is an instance of *hasPart_3_Rotor* if it is connected to three instances of *Rotor* that are

not known to be equal; in turn, the second rule infers that every such individual that is also an instance of *TwoRotorTurbine* violates the constraint axiom (8).

To conclude this section, we note that our translation in Table 2 yields a stratified program for any set \mathcal{C} of constraints. We can always define a stratification where the lowest stratum consists of the predicates in \mathcal{C} and *owl:sameAs*, the intermediate stratum contains all predicates of the form *R_B*, *R_n_B*, and *R_n*, and the uppermost stratum contains the special *Violation* predicate.

4 SOMM: An Industrial Ontology Management System

We have developed the *Siemens-Oxford Ontology Management (SOMM)* tool³ to support engineers in building ontologies and inserting data based on their information models. The interface of SOMM is restricted to support only the kinds of standard OWL 2 RL axioms and constraints discussed in Sect. 3.

SOMM is built on top of the Web-Protégé platform [40] by extending its front-end with new visual components and its back-end to access a Datalog-based triple store or a generic rule inference system for query answering and constraint validation, the OWL 2 reasoner HermiT [33] for ontology classification, and LogMap [16] to support ontology alignment and merging. Our choice of WebProtégé was based on Siemens' requirements for the platform underpinning SOMM, namely that it (i) can be used as a Web application; (ii) is under active development; (iii) is open-source and modular; (iv) includes built-in functionality for ontology versioning and collaborative development; (v) provides a form-based and end-user oriented interface; and (vi) enables the automatic generation of forms to insert instance data. Although we considered other alternatives such as Protégé-desktop [39], NeON toolkit [8], OBO-Edit [7], and TopBraid Composer [38], we found that only WebProtégé satisfied all the aforementioned requirements.

In the remainder of this section, we describe the main features of SOMM.

Insertion of axioms and constraints. We have implemented a form-based interface for editing standard axioms and constraints. Figure 3 shows a screenshot of the SOMM class editor representing the following axioms about *SteamTurbine* (abbreviated below as *ST*), where all but the last axiom represent constraints.

```

SubClassOf(ST ObjectSomeValuesFrom(hasState State))
SubClassOf(ST DataSomeValuesFrom(hasId xsd:string))
SubClassOf(ST ObjectMinCardinality(1 hasConfig STConfig))
SubClassOf(ST ObjectMaxCardinality(3 hasConfig STConfig))
SubClassOf(ST ObjectAllValuesFrom(hasProductLine ProductLine))

```

The interface shows that the class *SteamTurbine* has three mandatory properties (*hasState*, *hasID* and *hasConfig*) marked as 'Required' and interpreted as constraints, and an optional property (*hasProductLine*) interpreted as a standard

³ <http://www.cs.ox.ac.uk/isg/tools/SOMM/>.

axiom. Object and data properties are indicated by blue and green rectangles, respectively. For each property we can specify their filler using a WebProtégé autocompletion field. Finally, the fields ‘Min’ and ‘Max’ are used to represent cardinality constraints on mandatory properties.

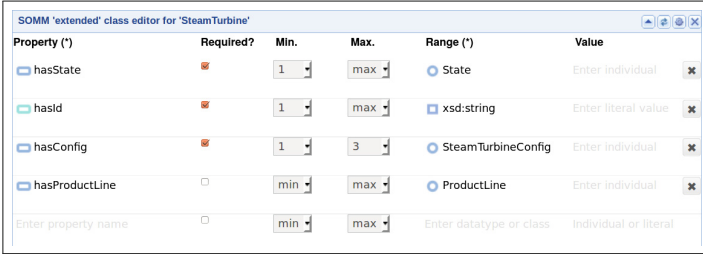


Fig. 3. SOMM editor to attach properties to classes.

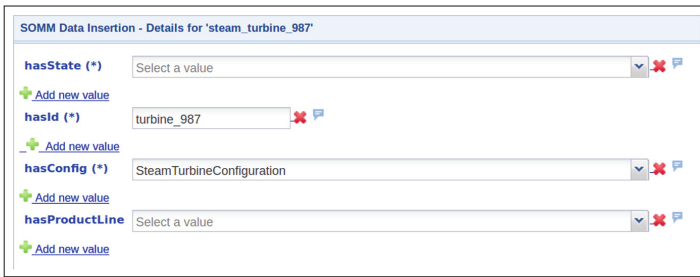


Fig. 4. Data insertion in SOMM.

Automatically generated data forms. SOMM exploits the capabilities of the ‘knowledge acquisition forms’ in Web-Protégé to guide engineers during data entry. The main use of data forms that we envision is ontology validation during the time of ontology development. The forms are automatically generated for each class based on its relevant mandatory and optional properties. For this, SOMM considers (i) the explicitly provided properties; (ii) the inherited properties; and (iii) the properties explicitly attached to its descendant classes. The latter were deemed useful by Siemens engineers, e.g., although *Turbine* does not have directly attached properties, the SOMM interface would suggest adding data for the properties attached to its subclass *SteamTurbine*. Figure 4 shows an example of the property fields for an instance of the class *SteamTurbine*, where required fields (i.e., those for which a value must be provided) are marked with (*).

Extended hierarchies. In addition to subsumption hierarchies, SOMM allows also for hierarchies based on arbitrary properties. These can be seen as a generalisation of partonomy hierarchies, and assume that the dependencies between

classes or individuals based on the relevant property are ‘tree-shaped’. Figures 5a and b show the hierarchy for the *follows* property, which determines which kinds of processes can follow other processes; for instance, *Conveying* follows *Loading* and is followed by *Testing*.

Alignment. SOMM integrates the system LogMap [16] to support model alignment and merging. Users can select and merge two Web-Protégé projects, or import and merge an ontology into the active Web-Protégé project. Although LogMap supports interactive alignment [17], it is currently used in SOMM in an automatic mode; we are planning to extend SOMM’s interface to support user interaction in the alignment process.

Reasoning. SOMM relies on Hermit [10] to support standard reasoning services such as class satisfiability and ontology classification. Data validation and query answering support is currently provided on top of the IRIS reasoner [3], as described in Sect. 3.2. Figures 5c and d illustrates the supported reasoning services. The left-hand-side of the figure shows that the class *GasTurbineModes* is satisfiable and *Process* is an inferred superclass. On the right-hand-side we can see that *steam_turbine_987* violates one of the integrity constraints; indeed, as shown in Fig. 4, *steam_turbine_987* is missing data for the property *hasState*, which is mandatory for all steam turbines (see Fig. 3).

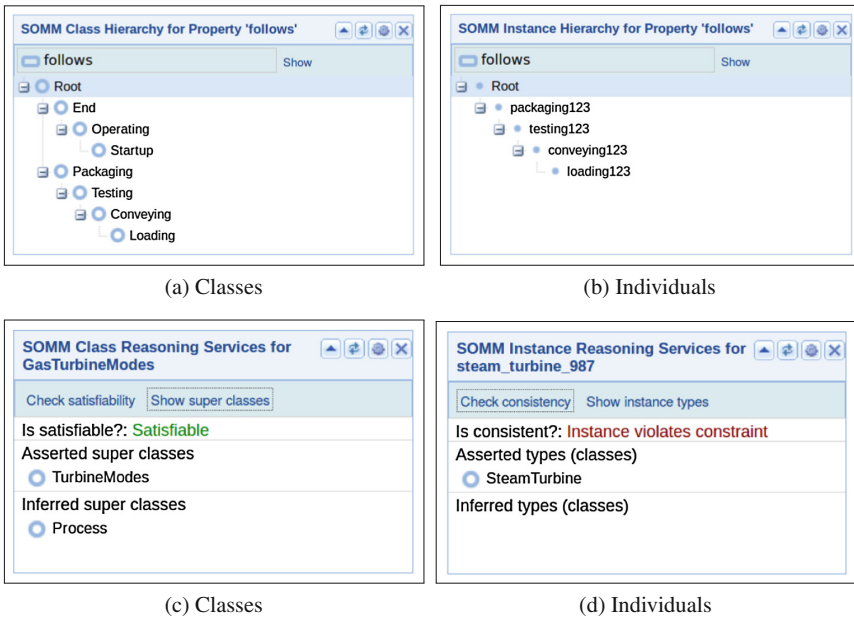


Fig. 5. Above: tree-like navigation of the ontology classes and individuals in SOMM. Below: reasoning services for ontology classes and individuals in SOMM

5 Evaluation

We have evaluated the practical feasibility of the data validation and query answering services provided by SOMM. For this, we have conducted two sets of experiments for the manufacturing and energy turbine scenarios, respectively. In the first experiment, we simulated the operation of a manufacturing plant using a synthetic generator that produces realistic product manufacturing data of varying size; in the second experiment, we used real anonymised turbine data.⁴ All our experiments were conducted on a laptop with an Intel Core i7-4600U CPU at 2.10 GHz and 16 GB of RAM running Ubuntu 14.04 (64 bits). We allocated 15 GB to Java 8 and set up IRIS with its default configuration.

Manufacturing Experiments. In our experiments for the manufacturing use case we used the ontology, data and queries given next.

- The ontology capturing the manufacturing model illustrated in Fig. 1 from Sect. 2.1. The ontology contains 79 standard axioms and 20 constraints.
- A data generator used by Siemens engineers to simulate manufacturing of products of two types based on the aforementioned model. We used two configurations of the generator: configuration (*C1*) simulates a situation where products were manufactured in violation of the model specifications (e.g., they used too much material of some kind); in (*C2*), each product is manufactured according to specifications.
- A sample of three monitoring queries commonly used in practice. The first query asks for all products that use material from a given lot; the second asks for all material lots used in a given product; finally, the third one asks for the total quantity of material in lots of a specific kind.

We generated data for 6 different sizes, ranging from 50 triples to 1 million triples. For each size, we generated one dataset for each configuration of the generator. We set up configuration *C1* so that 35% of the manufactured products violate specification. Our experiments follow Steps 1–4 in Sect. 3.2. We checked validity of each dataset against the ontology using Steps 1–3; then, for each dataset created using *C2* we also answered all test queries (Step 4). We repeated the experiment 5 times for each dataset and configuration (i.e., 10 times for each dataset size).

Our results are summarised in Fig. 6. Times for each data size are *wall clock* time averages (in *ms*). Constraint validation time (grey bar) correspond to Step 3 in Sect. 3.2. Query answering times (blue bar) measure the time for answering the use case queries (Step 4); here, only datasets satisfying the constraints (i.e., generated using *C2*) are considered. The figure also provides the average number of constraint violations in data generated according to *C1*, and the number of triples after constraint validation.

Our results demonstrate the feasibility of our ontology-based approach to model validation and query answering in realistic manufacturing scenarios. In

⁴ We are in the process of sorting out the licenses for the ontologies and data used in our experiments; they cannot be made publicly available at this point.

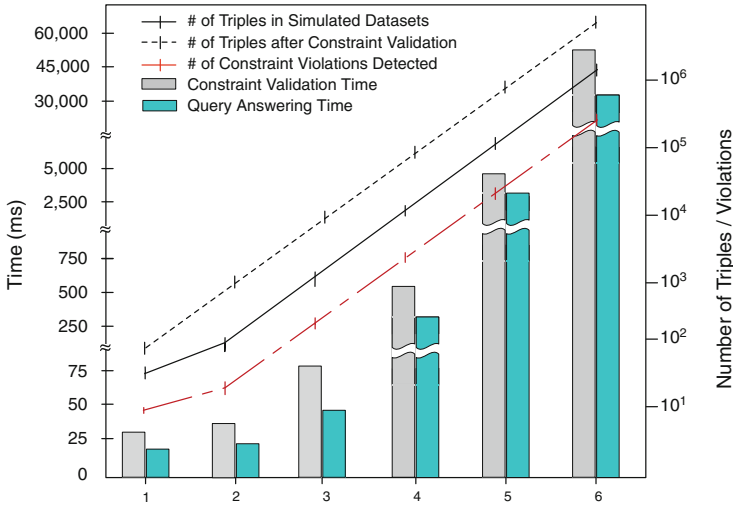


Fig. 6. Experimental results

particular, constraint validation and query answering were feasible within 87s on stock hardware over datasets containing over 1 million triples.

Gas Turbine Experiment. In this experiment we used the following data:

- The ontology capturing the energy plant model illustrated in Fig. 2 from Sect. 2. The ontology contains 121 standard axioms and 25 constraints.
- An anonymised dataset describing the structure of 800 real gas turbines, their sensor readings (temperature, pressure, rotor speed and position), and associated processes (e.g., expansion, compression, start up, shut down). The dataset was converted from a relational DB into RDF, and contains 25, 090 triples involving 4, 076 individuals.
- Three commonly used test queries. The first query asks for the core parts, equipment and current state of all turbines of a given type; the second asks for all components involved in a compression process; the last query asks for the temperature readings of turbines of a given type.

We followed the same steps as in the previous experiments, with very positive results. Constraint checking was completed in 2s and generated 27, 007 additional triples; we found 1, 582 constraint violations, which is especially interesting given that the data is real. Query answering over the valid subset took 1s on average.

6 Lessons Learned and Future Work

We have studied the use of ontologies to capture industrial information models in manufacturing and energy production applications.

Our study of the requirements of information models revealed that many key aspects of information models naturally correspond to integrity constraints and hence cannot be captured by standard OWL 2 ontologies. This demonstrates intrinsic limitations of OWL 2 for industrial modelling and gives a clear evidence of why constraints are essential for such modelling.

We also learned that even a rather simple form-based interface such as the one of SOMM is sufficient to capture most of the manufacturing and energy information models based on ISA and ICE standards. This was an important insight for us since at the beginning of this research project it was unclear whether designing such a simple tool to write ontologies of practical interest to our use cases would be feasible.

Finally, we have received a very positive feedback from Siemens engineers about the usability of SOMM at informal workshops organised as part of the project. This was encouraging since the development of a tool that is accessible to users without background in semantic technologies was one of the main motivations of our work.

In the future, we plan to conduct a formal user study where—with the help of SOMM—Siemens engineers will design elaborate information models and perform various tasks on these models, including validation and merging. We also plan to conduct more extensive scalability experiments. SOMM is a research prototype and, depending on the outcome these studies, we would like to deploy it in production departments.

References

1. Abele, L., Legat, C., Grimm, S., Muller, A.W.: Ontology-based of plant models. In: INDIN (2013)
2. Arancón, J., Polo, L., Berrueta, D., Lesaffre, F., Abajo, N., Campos, A.M.: Ontology- based knowledge management in the steel industry. In: *The Semantic Web: Real-World Applications from Industry* (2007)
3. Bishop, B., Ficsher, F.: IRIS - integrated rule inference system. In: *Workshop on Advancing Reasoning on the Web* (2008)
4. Calvanese, D., et al.: Optique: OBDA solution for big data. In: *ESWC (Satellite Events), Revised Selected Papers* (2013)
5. Classification and Product Description. <http://www.eclass.eu/>
6. Dantsin, E., Eiter, T., Gottlob, G., Voronkov, A.: Complexity, expressive power of logic programming. *ACM Comput. Surv.* **33**(3), 374–425 (2001)
7. Day-Richter, J., Harris, M.A., Haendel, M., Lewis, S.: OBO-Edit - an ontology editor for biologists. *Bioinformatics* **23**(16), 2198–2200 (2007)
8. Erdmann, M., Waterfeld, W.: Overview of the NeOn toolkit. In: *Ontology Engineering in a Networked World* (2012)
9. Forschungsunion. Fokus: Das Zukunftsprojekt Industrie 4.0, Handlungsempfehlungen zur Umsetzung. In: *Bericht der Promotorengruppe KOMMUNIKATION* (2012)
10. Glimm, B., Horrocks, I., Motik, B., Stoilos, G., Wang, Z.: HermiT: An OWL 2 reasoner. *J. Autom. Reasoning* **53**(3), 245–269 (2014)
11. Gliozzo, A., Biran, O., Patwardhan, S., McKeown, K.: Semantic technologies in IBM watson. In: *Teaching NLP and CL Workshop (TNLP) at ACL* (2013)

12. Grangel-González, I., Halilaj, L., Coskun, G., Auer, S., Collarana, D., Hoffmeister, M.: Towards a semantic administrative shell for industry 4.0 components. In: ICSC (2016)
13. Grimm, S., Watzke, M., Hubauer, T., Cescolini, F., Embedded \mathcal{EL} + reasoning on programmable logic controllers. In: ISWC (2012)
14. Hepp, M., de Bruijn, J.: GenTax: a generic methodology for deriving OWL and RDF-S ontologies from hierarchical classifications, thesauri, and inconsistent taxonomies. In: ESWC (2007)
15. Hubauer, T., Lamparter, S., Pirker, M.: Automata-based abduction for tractable diagnosis. In: DL (2010)
16. Jiménez-Ruiz, E., Cuenca Grau, B.: LogMap: logic-based and scalable ontology matching. In: ISWC (2011)
17. Jiménez-Ruiz, E., Cuenca Grau, B., Zhou, Y., Horrocks, I.: Large-scale interactive ontology matching: algorithms and implementation. In: ECAI (2012)
18. Kagermann, H., Lukas, W.-D.: Industrie 4.0: Mit dem internet der Dinge auf dem Weg zur 4. industriellen Revolution. In: VDI Nachrichten (2011)
19. Kharlamov, E., et al.: Enabling semantic access to static, streaming distributed data with optique: demo. In: ACM DEBS (2016)
20. Kharlamov, E.: How semantic technologies can enhance data access at siemens energy. In: ISWC (2014)
21. Kharlamov, E., et al.: Ontology based access to exploration data at statoil. In: Arenas, M., et al. (eds.) ISWC 2015. LNCS, vol. 9367, pp. 93–112. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-25010-6_6](https://doi.org/10.1007/978-3-319-25010-6_6)
22. Kharlamov, E.: Ontology-based integration of streaming and static relational data with optique. In: ACM SIGMOD (2016)
23. Kharlamov, E., et al.: Optique: ontology-based data access platform. In: ISWC (P&D) (2015)
24. Kharlamov, E., et al.: Optique: towards OBDA systems for industry. In: Cimiano, P., Fernández, M., Lopez, V., Schlobach, S., Völker, J. (eds.) ESWC 2013. LNCS, vol. 7955, pp. 125–140. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-41242-4_11](https://doi.org/10.1007/978-3-642-41242-4_11)
25. Kharlamov, E., et al.: Semantic access to siemens streaming data: the optique way. In: ISWC (P&D) (2015)
26. Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C.: OWL 2 web ontology language profiles (Second Edition). W3C Recommendation (2012)
27. Motik, B., Horrocks, I., Sattler, U.: Bridging the gap between OWL and relational databases. *J. Web Sem.* **7**(2), 41–60 (2009)
28. Nenov, Y., Piro, R., Motik, B., Horrocks, I., Wu, Z., Banerjee, J.: RDFox: a highly-scalable RDF store. In: ISWC (2015)
29. Poggi, A., Lembo, D., Calvanese, D., De Giacomo, G., Lenzerini, M., Rosati, R.: Linking data to ontologies. *J. Data Semant.* **10**, 237–271 (2008)
30. Qiu, R.G., Zhou, M.: Mighty MESs; state-of-the-art, future manufacturing execution systems. *IEEE Robot. Automat. Mag.* **11**(1) (2004)
31. Richnow, J., Rossi, C., Wank, H.: Designation of wind power plants with the reference designation system for power plants - RDS-pp. *VGB PowerTech.* **94**, 2 (2014)
32. Ringsquandl, M., Lamparter, S., Brandt, S., Hubauer, T., Lepratti, R.: Semantic-guided feature selection for industrial automation systems. In: Arenas, M., et al. (eds.) ISWC 2015. LNCS, vol. 9367, pp. 225–240. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-25010-6_13](https://doi.org/10.1007/978-3-319-25010-6_13)
33. Shearer, R., Motik, B., Horrocks, I.: HermiT: a highly-efficient OWL reasoner. In: OWLED (2008)

34. Siemens. Modeling new perspectives: digitalization - the key to increased productivity, efficiency and flexibility (White Paper). In: DER SPIEGEL, 6 2015
35. Stetter, R.: Software im maschinenbau-laestiges anhangsel oder chance marktfuehrerschaft? In: VDMA, ITQ (2011). <http://www.software-kompetenz.de/en/>
36. Stolz, A., Rodriguez-Castro, B., Radinger, A., Hepp, M.: PCS2OWL: a generic approach for deriving web ontologies from product classification systems. In: Presutti, V., d'Amato, C., Gandon, F., d'Aquin, M., Staab, S., Tordai, A. (eds.) ESWC 2014. LNCS, vol. 8465, pp. 644–658. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-07443-6_43](https://doi.org/10.1007/978-3-319-07443-6_43)
37. Tao, J., Sirin, E., Bao, J., McGuinness, D.L.: Integrity constraints in OWL. In: AAAI (2010)
38. Top Quadrant. TopBraid Composer. <http://www.topquadrant.com/>
39. Tudorache, T., Noy, N.F., Tu, S.W., Musen, M.A.: Supporting collaborative ontologydevelopment in protégé. In: ISWC (2008)
40. Tudorache, T., Nyulas, C., Noy, N.F., Musen, M.A.: WebProtégé: a collaborative ontology editor and knowledge acquisition tool for the web. In: Semantic Web 4.1 (2013)
41. Vyatkin, V., Engineering, S.: Software engineering in industrial automation: state-of-the-art review. IEEE Trans. Ind. Inf. **9**(3), 2351–2362 (2013)