Semantic Interoperability in the Norwegian Petroleum Industry

Jon Atle Gulla, Stein L. Tomassen and Darijus Strasunskas

Norwegian University of Science and Technology
NO-7491 Trondheim, Norway
jag@idi.ntnu.no

Abstract: The petroleum industry is a technically challenging business with high investments, complex projects and operational structures. There are numerous companies and public offices involved in the exploitation of a new oil field, and there is a high degree of specialization among them. Even though standardization has been considered important in this industry for many years, there is still very little integration across phases and disciplines. An industrially driven consortium launched the Integrated Information Platform project in 2004, in which semantic standards based on OWL and Semantic Web technologies were to be developed for the subsea petroleum industry. This paper presents the IIP project in more detail and discusses how industry-wide ontologies can support semantic interoperability and encourage more process integration.

1 Introduction

The petroleum industry in Norway is technically challenging with subsea installations and difficult climatic conditions. It is industrially still quite fragmented, in the sense that there is little collaboration between phases and disciplines in large petroleum projects. There are many specialized companies involved, though their databases and applications are not necessarily well integrated with each other. Research done by the Norwegian Oil Industry Association shows that there is a need for more collaboration and integration across phases, disciplines and companies [Ol05b]. The existing standards do not provide the necessary support for this, and the result is costly and risky projects and decisions based on wrong or outdated data.

This paper presents the visions and preliminary results of the Integrated Information Platform (IIP) project. The idea here is to extend and formalize an existing terminology standard for the petroleum industry, ISO 15926. Using Semantic Web technologies, we turn this standard into a real ontology that provides a consistent unambiguous terminology for subsea petroleum production systems. The results of the project so far are very promising, and the ontology under development is considered to be the first step towards real semantic interoperability in the Norwegian petroleum industry.

The paper is organized as follows. In Section 2 we go through the structures and challenges in the subsea petroleum industry, explaining the status of current standards and the vision of future integrated operations. Section 3 briefly presents the parts of the Semantic Web initiative relevant to this project. Whereas the IIP project is introduced in Section 4, we discuss the construction of our OWL ontology in Section 5. A preliminary assessment of our work is given in Section 6, and the conclusions are found in Section 7.

2 The Subsea Petroleum Industry

The Norwegian subsea petroleum industry is a technically challenging business. Sophisticated equipment and highly competent companies are needed, and the projects tend to be both large and expensive. Many disciplines and competences need to come together in these projects, and their success is highly affected by the way people and systems are able to collaborate and coordinate their work. On the Norwegian Continental Shelf (NCS) there are traditional oil companies like Statoil, Norsk Hydro and ElfTotalFina, but also specialized service companies like Schlumberger, Haliburton, Baker Hughes, Aker Kværner, FMC KongsbergSub, and smaller ICT service companies.

Both the projects and the subsequent production systems are information-intensive. When a well is put into operation, the production has to be monitored closely to detect any deviation or problems. The next generation subsea systems include numerous sensors that measure the status of the systems and send real-time production data back to certain operation center. For these centers to be effective, they need tools that allow them to understand this data, relate it to other relevant information, and help them deal with the situation at hand. There is a challenge in dealing with all this information, but also in interpreting information that is deeply rooted in very technical terminologies.

The Norwegian petroleum industry is now facing a number of challenges [Ol05a]:

- **Resource decline**. We produce more oil and gas than we add through the exploration and improved oil recovery additions. We add in the order of 100 million Sm3 oil equivalents (exploration and IOR) while we produce about 250 million Sm3 (standard cubic meters, 1 Sm3 = 6.29 barrels). This process of accumulated resource decline has been the case the last 5-10 years
- **Cost increase**. The future unit cost on all the bigger fields will increase significantly the coming years as the fields enter the decline phase.
- **Plethora of companies**. We now see a development on the NCS from the traditional bigger fields of 300-400 million Sm3 to fields of 3-5 million Sm3. In addition to traditional bigger service companies we see an increase of smaller highly specialized service companies entering the market.

All these trends pose a challenge to the profitability of existing and future petroleum fields on NCS. While the costs of old large fields are increasing, the new ones are financially less attractive due to lack of scalability. The multitude of companies involved, with their own applications and databases, makes coordination and collaboration more important than in the past. For the industry as a whole, this severely hampers the integration of applications and organizations as well as the decision making processes in general:

- **Integration**. Even though there is some cooperation between companies in the petroleum sector, this cooperation tends to be set up on an ad-hoc basis for a particular purpose and supported by specifically designed mappings between applications and databases. There is little collaboration across disciplines and phases, as they usually have separate databases rooted in different goals, structures and terminologies. It is of course possible to map data from one database to another, but with the complexity of data and the multitude of companies and applications in the business this is not a viable approach for the industry as a whole.
- **Decision making**. A current problem is the lack of relevant high-quality information in decision making processes. Some data is available too late or not at all because of lack of integration of databases. In other cases relevant data is not found due to differences in terminology or format. And even when information is available, it is often difficult to interpret its real content and understand its limitations and premises. This is for example the case when companies report production figures to the government using slightly different terminologies and structures, making it very hard to compare figures from one company to another.

XML is already used extensively in the petroleum industry as a syntactic format for exchanging data. Over the last few years, there have been several initiatives for defining semantic standards to achieve semantic interoperability and information sharing in the business.

2.1 ISO 15926 Integration of Life-Cycle Data

ISO 15926 is a standard for integrating life-cycle data across phases (e.g. concept, design, construction, operation, decommissioning) and across disciplines (e.g. geology, reservoir, process, automation). It consists of 7 parts, of which parts 1, 2 and 4 are the most relevant to this work. Whereas part 1 gives a general introduction to the principles and purpose of the standard, part 2 specifies the modeling language for defining application-specific terminologies. Part 2 comes in the form of a data model and includes 201 entities that are related in a specialization hierarchy of types and sub-types. It is intended to provide the basic types necessary for defining any kind of industrial data. Being specified in EXPRESS [Is05], it has a formal definition based on set theory and first order logic.

Part 4 of ISO 15926 is comprised of application or discipline-specific terminologies, and is usually referred to as the Reference Data Library (RDL). These terminologies, described as RDL classes, are instances of the data types from part 2, are related to each other in a specialization hierarchy of classes and sub-classes as well as through memberships and relationships. Whereas part 2 defines the language for describing standardized terminologies, part 4 describes the semantics of these terminologies. There is ongoing work in the Norwegian offshore industry to provide a comprehensive standardized terminology for the petroleum industry in part 4. Part 4 today contains approximately 50.000 general concepts like motor, turbine, pump, pipes and valves.

ISO 15926 is still under development, and only Part 1 and 2 have so far become ISO standards. In addition to adding more RDL classes for new applications and disciplines in Part 4, there is also a discussion about standards for geometry and topology (Part 3), procedures for adding and maintaining reference data (Part 5 and 6), and methods for integrating distributed systems (Part 7). Neither ISO 15926 nor other standards have the scope and formality to enable proper integration of data across phases and disciplines in the petroleum industry.

2.2 The Vision of Integrated Operations

The Norwegian Oil Industry Association proposed the Integrated Operations program in 2004. The fundamental idea is to integrate processes and people onshore and offshore using new information and communication technologies. Facilities to improve onshore's abilities to support offshore operationally are considered vital in this program. Personnel onshore and offshore should have access to the same information in real-time and their work processes should be redefined to allow more collaboration and be less constrained by time and space. OLF has estimated that the implementation of integrated operations on the NCS can increase oil recovery by 3-4%, accelerate production by 5-10% and lower operational costs by 20-30% [Ol05b].

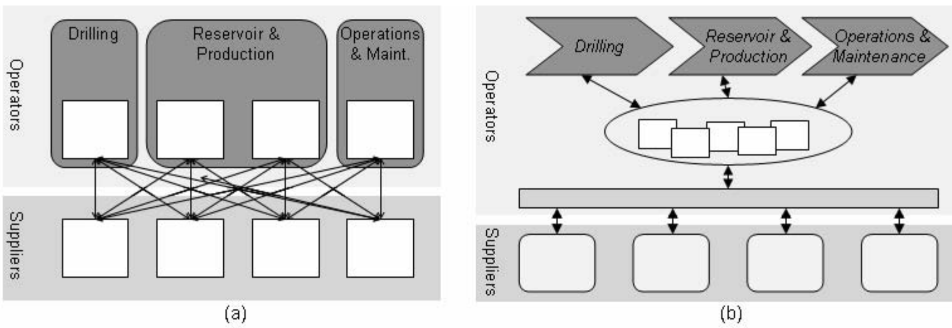

Figure 1(a) Current situation (b) The vision of integrated operations

Central in this program is the semantic and uniform manipulation of heterogeneous data.

Figure 1 illustrates the objectives of the integrated operations initiative. Whereas we in the current situation have numerous databases that need to be mapped to each other on an ad hoc basis, we envision a semantic standard in the future that supports integration and interoperability between data from all phases and disciplines. Suppliers' applications interact with the operators' data through standardized semantic interfaces, making sure that a unified terminology is used and data is consistent and unambiguous.

3 Semantic Web Technology and Interoperability

"The Semantic Web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation."

[BHL01]. The Semantic Web is a collaborative effort led by W3C with participation from a large number of researchers and industrial partners. The general idea in the Semantic Web is to annotate each piece of data with machine-processable semantic descriptions. These descriptions must be specified according to a certain grammar and with reference to a standardized domain vocabulary. The domain vocabulary is referred to as an ontology and is meant to represent a common conceptualization of some domain. The grammar is a semantic markup language, as for example the OWL web ontology language recommended by W3C. With these semantic annotations in place, intelligent applications can retrieve and combine documents and services at a semantic level, they can share, understand and reason about each other's data, and they can operate more independently and adapt to a changing environment by consulting a shared ontology [Sh02, ZLY02].

Interoperability can be defined as a state in which two application entities can accept and understand data from the other and perform a given task in a satisfactory manner without human intervention. We often distinguish between syntactic, structural and semantic interoperability [Ag05,Du04]:

- *Syntactic interoperability* denotes the ability of two or more systems to exchange and share information by marking up data in a similar fashion (e.g. using XML).
- *Structural interoperability* means that the systems share semantic schemas (data models) that enable them to exchange and structure information (e.g. using RDF).
- *Semantic interoperability* is the ability of systems to share and understand information at the level of formally defined and mutually accepted domain concepts, enabling machine-processable interpretation and reasoning.

For the Semantic Web technology to enable semantic interoperability in the petroleum industry, it needs to tackle the problem of *semantic conflicts,* also called *semantic heterogeneity*. Since the databases are developed by different companies and for different phases and/or disciplines, it is often difficult to relate information that is found in different applications. Even if they represent the same type of information, they may use formats or structures that prevent the computers from detecting the correspondence between data. For example, the tables *ORG_NAME* and *COMPNY* in two different applications may in fact contain the same information about organizations. Similarly, while a time period may be modeled with the variables *"StartTime"* and *"Endtime"* in one database, the same information may be represented with *"StartTime"* and *"Duration"* in another. Figure 2 sums up typical semantic conflicts among large enterprise systems.

The Semantic Web's approach to these problems is the construction of shared formal ontologies of all important domain concepts. These may be specified in OWL, which is a semantic markup language based on Description Logic. It has an XML syntax, is built on top of RDF(S)'s property statements and class hierarchies, and adds constraints for class membership, equivalence, consistency and classification [AFH05, W3C05].

4 The Integrated Information Platform Project

The Integrated Information Platform (IIP) project is a collaboration project between companies active on NCS and academic institutions, supported by the Norwegian Research Council [SM04]. Its objective is to ease the integration of data and processes across phases and disciplines by providing a comprehensive unambiguous and well accepted terminology standard that lends itself to machine-processable interpretation and reasoning. This should reduce risks and costs in petroleum projects and indirectly lead to faster, better and cheaper decisions.

Semantic conflicts	**Description**	**Example**
Data type	Different primitives or abstract types for same information	SSN as a VARCHAR vs. a NUM
Labeling	Synonyms/antonyms have different text labels	When ORG_NAME and COMPNY tables have data that mean the same thing
Aggregation	Different conceptions about the relationships among concepts in similar data sets.	Does a "motorcycle" have 1, 2, 3, 4 or more wheels, how are the constraints modeled in your schema?
Generalization	Different abstractions are used to model same domain	Are "cars" and "trucks" kinds of "vehicles" or are they top-level classes themselves?
Value representation	Different choices are made about what concepts are made explicit	"StartTime" plus "Duration" equals "EndTime"
Impedance mismatch	Fundamentally different data representations are used	Relational to Object mappings (key migrations, multiplicity, etc.)
Naming	Synonyms/antonyms exist in same/similar concept instance values	"Company" table has many entries: "DaimlerBenz", "Mercedes", etc. but they refer to the same thing
Scaling and unit	Different units of measures with incompatible scales	km vs. English mile
Confounding	Similar concepts with different definitions	"EarningsPerShare" object for a NASD application vs. a NYSE system
Domain	Fundamental incompatibilities in underlying domain	"MainAssembly" object in a Ford product system vs. a brake

		supplier system
Integrity	Disparity among the integrity constraints	Does an airline ticket have a primary key that uniquely IDs a passenger?

Figure 2. Semantic conflicts in enterprise systems (from [PH04])

The project is identifying an optimal set of real-time data from reservoirs, wells and subsea production facilities. The OWL web ontology language is chosen as the markup language for describing these terms semantically in an ontology. The entire standard is thus rooted in the formal properties of OWL, which has a model-theoretic interpretation and to some extent support formal reasoning. A major part of the project is to convert and formalize the terms already defined in ISO 15926 Part 2 (Data Model) and Part 4 (Reference Data Library), which we will come back to in the next section. Since the ISO standard addresses rather generic concepts, though, the ontology must also include more specialized terminologies for the oil and gas segment. Detailed terminologies for standard products and services are included from other dictionaries and initiatives (DISKOS,WITSML, ISO 13628/14224, SAS), and the project also opens for the inclusion of terms from particular processes and products at the bottom level. The IIP project started in June 2004 and will run until June 2007 with a total budget of 22 million NOK (about 2.8 million Euro). The participants include Det Norske Veritas, Statoil, Norsk Hydro, Cap Gemini, Poseidon, OLF, FMC Technologies, National Oilwell Varco, OilCamp, POSC and NTNU.

5 Ontology Engineering in IIP

The top structure of ISO 15926-2 Data Model is shown in EXPRESS notation in the upper part of Figure 4. The entity type *Thing* has two subtypes, *Abstract object* and *Possible individual*. The *Class* type has one super-type, *Abstract object*, and four sub-types (*Class of individual, Class of abstract object, Cardinality,* and *Role and domain*). The application-specific terminologies in ISO 15926-4 RDL are defined as instances of the relevant types in the Data Model. For example, the RDL class *Pump* is an instance of the *Class of individual* type from the Data Model. At the same time, pumps are specializations of the more general RDL class *IndustrialArtifact*. Similarly, the relationship between the classes Pump and Pipe is modeled as an instance of the Data Model type *Class of relationship*. Individual instances, like an actual pump in a particular subsea installation (shown as *#myPump* in Figure 3), are represented as instances of an RDL class (*Pump*) as well as an instance of a Data Model type (*Possible individual*). Note that this also applies to relationship between individual instances, like *#myConnection*.

The ISO 15926-2 Data Model is represented as a meta model in our ontology. For the representation of this Data Model in OWL, we use the concept RDFS:Class with subclass OWL:Class. Moreover, in the meta model OWL:DatatypeProperty and OWL::ObjectProperty are subclasses of RDFS:Property and instances of OWL:Class.

For a complete discussion of the mapping rules from EXPRESS to OWL, the reader is referred to [CVJ05].

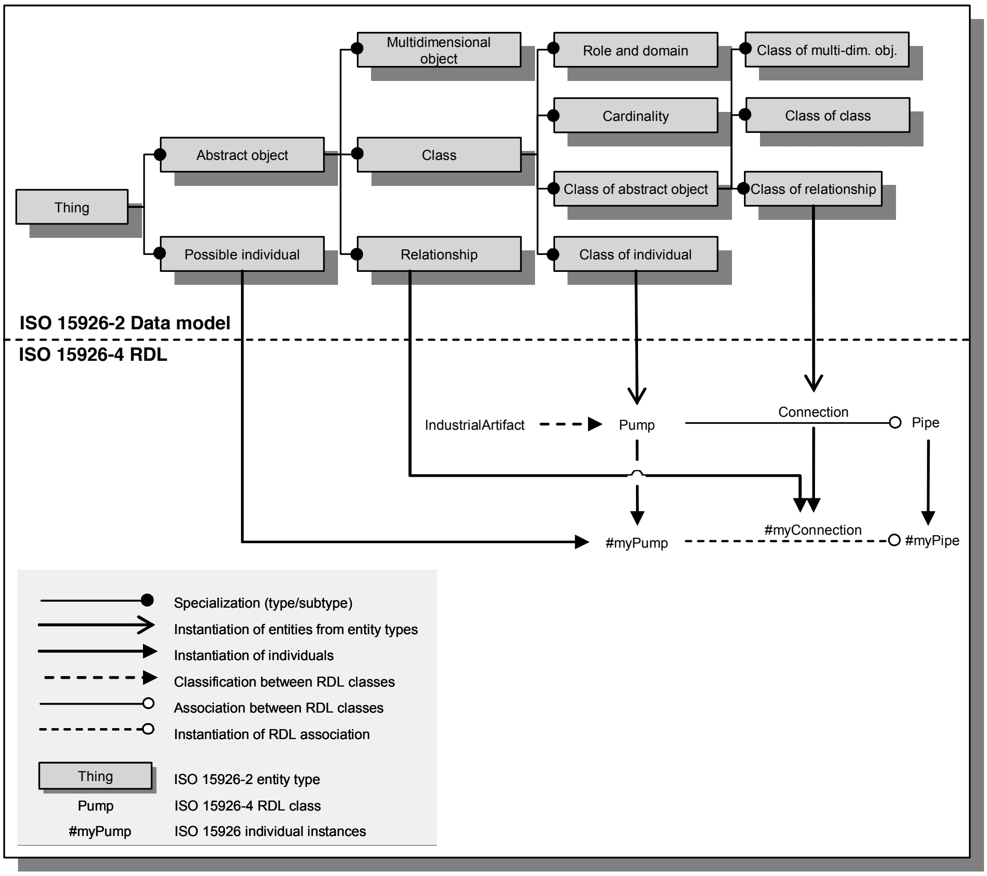

Figure 3. Structures of ISO 15926 part 2 (Data Model) and 4 (Reference Data Library)

The conversion of ISO 15926-2/4 from EXPRESS gives us an OWL hierarchy that forms the backbone of the new ontology. Additional terms are gradually and manually added to this hierarchy to reflect the larger scope of the new standard. In these initial stages it has been considered important to concentrate on hierarchical relationships between concepts. Relationships and constraints will be added gradually as the hierarchy matures and stabilizes. Let us now have a look at a particular term in this ontology. A *Christmas tree* is an assembly of parts that is connected to the top of a wellhead to control the flow out of the well. Its initial OWL definition (without relationships and constraints) is:

```
<owl:Class rdf:about="#CHRISTMAS_TREE">
        …
        <dc:description   rdf:datatype="http://www.w3.org/2001/XMLSchema#string">
                An artefact that is an assembly of pipes and piping parts, with valves
                and associated control equipment that is connected to the top of a
                wellhead and is intended for control of fluid from a well.
        </dc:description>
        <dc:title rdf:datatype="http://www.w3.org/2001/XMLSchema#string">
                CHRISTMAS TREE
        </dc:title>
        …
        <rdfs:subClassOf rdf:resource="#ARTEFACT"/>
</owl:Class>
```

These statements give us an informal definition of Christmas trees and reveal that they are subclasses of artefact. Looking at the excerpt of the class hierarchy in Figure 4, we see that there are at least three types of Christmas tree (subsea, vertical, and horizontal). It is a specialization of Artefact, which in turn is an Inanimate physical object that is made or given a shape by man. The Pipe class is also a specialization of Artefact, but it is also a specialization of two other classes. This is a quite natural, as the pipe both has a physical (artefact) and a functional dimension (pipeline or network connection).

The IIP project has now converted the ISO 15926 Part 2 (210 elements) and Part 4 (about 50.000) elements into OWL class hierarchies. In addition, we are in the process of incorporating additional terms from the following disciplines:

- Geometry and topology: 400 terms
- Drilling and logging: ca. 2.700 terms
- Production: ca. 2.000 terms
- Safety and automation: ca. 150 terms
- Subsea equipment: ca. 1.000 terms
- Reservoir characterization
- Reliability and maintenance

The Tyrihans oil field, operated by Statoil, is used as a case in the IIP project. This means that the initial terms included in the ontology are based on the Tyrihans specifications, though they have been generalized and verified against other specifications as well, like ISO 13628 “Petroleum and natural gas industries – Design and operation of subsea production systems”. The ontology will be the basis for developing new semantically interoperable applications, and IIP has already started experimenting with integrated visualization and information retrieval environments.

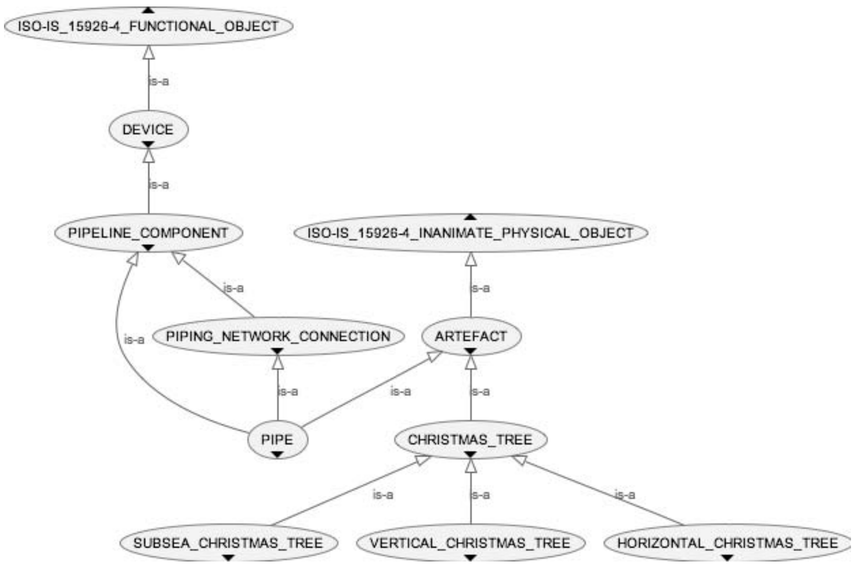

Figure 4. Christmas tree hierarchy

6 Discussion

In recent years a number of powerful new ontologies have been constructed and applied in selected domains. This is particularly true in medicine and biology, where Semantic Web technologies and web mining have been exploited in new intelligent applications [Ag05, Ge00, Pi04]. However, these disciplines are heavily influenced by government support and are not as commercially fragmented as the petroleum industry. Creating an industry-wide standard in a fragmented industry is a huge undertaking that should not be underestimated. In this particular case, we have been able to build on an existing standard, ISO 15926. This has ensured sufficient support from companies and public institutions. There is still an open question, though, what the coverage of such an ontology should be. There are other smaller standards out there, and many companies use their own internal terminologies for particular areas. The scope of this standard has been discussed throughout the project as the ontology grew and new companies signalled their interest. For any standard of this complexity, it is important also to decide where the ontology stops and to what extent hierarchical or complementing ontologies are to be encouraged. Techniques for handling ontology hierarchies and ontology alignment and enrichment must be considered in a broader perspective.

As far as the construction of the ontology is concerned, there is a need for both domain experts and ontology engineers. Since both the syntax and the semantics of OWL are non-trivial, it cannot be assumed that domain experts do the modeling themselves. To handle the complexity, the IIP project decided to model only the hierarchical relations in the first round, delaying relationships and constraints until the hierarchies are stable. For later update and quality assessment, it may be useful to use text mining techniques for automatic term extraction [GBK04, Ma02, Ve05].

The quality of ontologies is a delicate topic. It is important to choose an appropriate level of granularity. In this project we have been fortunate to have an existing standard to start with. What was considered satisfactory in ISO 15926 may however not be optimal for the ontology-driven applications that will make use of the future ontology. Ultimately, we need to consider how the ontology will be used in these applications and the nature of the source data to be annotated with ontological descriptions. For a search tool developed in IIP, we are adding a mechanism for enriching the ontology with weighted words to reflect the user needs and the nature of the document collection [TGS06].

Since the Semantic Web is still a rather immature technology, there are still open issues that need to be addressed in the future. One problem in the IIP project is that we need the full expressive power of OWL (OWL Full) to represent the structures of ISO 15926-2/4. Reasoning with OWL specifications is then incomplete. The lack of industrial SW applications is another issue worth taking into consideration. There may be performance and maintenance complexities that are still unclear with such an untested technology. However, there is now a large community promoting SW technologies and developing innovative applications, and the first commercial products have also emerged. Additionally, the tool development in IIP indicates that the technology can form the semantic foundation for a new generation of intelligent, interoperable information services.

The success of the new ontology, and standardization work in general, depends on the users' willingness to commit to the standard and devote the necessary resources. If people do not find it worthwhile to take the effort to follow the new terminology, it will be difficult to build up the necessary support. This means that it is important to provide environments and tools that simplify the use and maintenance of the ontology. Intelligent ontology-driven applications must demonstrate the benefits of the new technology and convince the users that the additional sophistication pays off.

7 Conclusions

The Integrated Information Platform project is one of the first attempts at applying state-of-the-art Semantic Web technologies in an industrial setting. Existing standards are now being converted and extended into a comprehensive OWL ontology for reservoir and subsea production systems. The intention is that this ontology will later be approved as an ISO standard and form a basis for developing interoperable applications in the industry.

With the new ontology at hand, the industry will have taken the first step towards integrated operations on the Norwegian Continental Shelf. Data can then be related across phases and disciplines, helping people collaborate and reducing costs and risks.

However, there are costs associated with building and maintaining such an ambitious ontology. It remains to be seen if the industry is able to take full advantage of the additional expressive power and formality of the new ontology. The work in IIP indicates that both information retrieval systems and sensor monitoring systems can benefit from having access to an underlying ontology for analyzing data and interpreting user needs.

As the class hierarchies in the ontology are completed, the emphasis of the IIP project will be put on adding relationships and constraints to the ontology. This also includes specifying rules that will be used to analyze anomalies in the real-time data from the subsea sensors. At that point we can start exploiting the logical properties of OWL and start experimenting with the next generation rule-based notification systems. Also, we can build more complete semantic descriptions of documents and add more reasoning capabilities to our information retrieval tools. We will then see if a strong semantic foundation makes it easier for us to handle and interpret the vast amount of data that are so typical to the petroleum industry.

Acknowledgements

This research is funded by the Integrated Information Platform for reservoir and subsea production systems (IIP) project, which is supported by the Norwegian Research Council (NFR project number 163457/S30).

Bibliography

[AFH05] Antoniou, G.; Franconi, E.; van Harmelen, F. Introduction to Semantic Web Ontology Languages. In Eisinger, N. and Maluszynski, J. (Ed.), Reasoning Web, First International Summer School 2005, Chap. 1, Malta, July 2005. Springer.

[Ag05] Aguilar, A. Semantic Interoperability in the Context of e-Health. 2005. CDH Seminar. http://m3pe.org/seminar/aguilar.pdf.

[BHL01] Berners-Lee, T.; Hendler, J.; Lassila, O. The Semantic Web." Scientific American, Vol. 284, No. 5, 2001, pp. 34-43.

[CJV05] Christiansen, T.;Jensen, M; Valen-Sendstad, M. Defining ISO 15926-4 Reference Data Library classes in OWL. The Norwegian Oil Industry Association. http://www.olf.no/io/kunnskapsind/?28140.pdf. Accessed 28 December 2005.

[Du04] Dublin Core. Dublin Core Metadata Glossary. http://library.csun.edu/mwoodley/dublincoreglossary.html. 2004.

[Ge00] Gene Ontology Consortium. Gene Ontology: tool for the unification of biology. Nature Genet. Vol. 25, pp 25-29. 2000.

[GBK04] Gulla, J. A., Brasethvik, T., Kaada, H. A Flexible Workbench for Document Analysis and Text Mining. In Proceedings of the 9th International Conference on Applications of Natural Language to Information Systems (NLDB'04), pp. 336-347, Manchester, June 2004. Springer.

[Is05] International Standards Association. Industrial automation systems and integration – Product data representation and exchange. Part 11: Description methods: The EXPRESS language reference manual. http://www.iso.org/iso/en/CatalogueDetailPage.CatalogoueDetail?CSNUMBER=18348.

[Ma02] Maedche, A. Ontology Learning for the Semantic Web. Kluwer Academic Publ. 2002.

[Ol05a] OLF. Digital Infrastructure Offshore – Common Network Operation Management for Digital Infrastructure Offshore on The Norwegian Continental Shelf. The Norwegian Oil Industry Association. http://www.olf.no/io/digitalinfrastruktur/?29220.pdf. Accessed 27 December 2005.

[Ol05b] OLF. Integrated Work Processes: Future work processes on the Norwegian Continental Shelf. The Norwegian Oil Industry Association. Accessed 28 December 2005.

[PH04] Pollock, J. T. and R. Hodgson. Adaptive Information: Improving Business Through Semantic Interoperability, Grid Computing, and Enterprise Integration. Wiley Publishers, 2004.

[Pi04] Pisanelli, D. M. (Ed.). Ontologies in Medicine. Volume 102 Studies in Health Technology and Informatics. IOS Press. 2004.

[Sh02] Sheth, A.; Bertram, C.; Avant, D.; Hammond, B.; Kochut, K.; Warke, Y. Managing Semantic Content for the Web. IEEE Internet Computing, July/August 2002, pp. 80-87.

[SM04] Sandsmark, N.; Mehta, S. Integrated Information Platform for Reservoir and Subsea Production Systems. In Proceedings of the 13th Product Data Technology Europe Symposium (PDT 2004), October, Stockholm 2004.

[TGS06] Tomassen, S. L.; Gulla, J. A.; Strasunskas, D. Query Enrichment with Domain Adapted Ontology. Accepted for 11th International Conference on Applications of Natural Language to Information Systems (NLDB'2006), Klagenfurt, June 2006.

[Ve05] P. Velardi, P.; Navigli, R,; Cucchiarelli, A.; Neri, F.. Evaluation of OntoLearn, a Methodology for Automatic Learning of Domain Ontologies, Ontology Learning from Text: Methods, Evaluation and Applications (P. Buitelaar, P. Cimiano, and B. Magnini, Eds.), Series information for Frontiers in Artificial Intelligence and Applications, IOS Press, 2005.

[W3C04] W3C. OWL Web Ontology Language Overview. http://www.w3.org/TR/owl-features/ (Accessed 4 January 2006).

[ZLY02] Zhong, N.; Liu, J.; Yao, Y. In Search of the Wisdom Web. Computer, November 2002, pp 27-31.