Fairtrace
A Semantic-Web Oriented Traceability Solution Applied To The Textile Traceability

Bruno Alves[1], Michael Schumacher[1], Fabian Cretton[1], Anne Le Calvé[1], Gilles Cherix[2], David Werlen[2], Christian Gapany[2], Bertrand Baeryswil[3], Doris Gerber[4], Philippe Cloux[4]

[1]*IIG, University of Applied Sciences Western Switzerland, TechnoArk 3, Sierre, Switzerland*

[2]*Institut Icare, TechnoArk, Sierre, Switzerland*

[3]*Fairtrace, Technoark, Sierre, Swtizerland*

[4]*Importexa SA, 1095 Lutry, Switzerland*

{bruno.alves, michael.schumacher, fabian.cretton, anne.lecalve}@hevs.ch, {cherix, werlen, gapany}@icare.ch, bertrand@fairtrace.ch, dorisgerber@hotmail.com, philippe@importexa.com

Abstract: This paper presents solutions that leverage Semantic Web Technologies (SWT) to allow pragmatic traceability in supply-chains, especially for the textile industry. Objectives are the identification of the supply-chain, order management, tracking and problem reporting (such as dangerous substance detection). It is intended to be a generic platform supporting potentially any kind of industrial supply-chain, to be usable in harsh environments (mobile appliances) without any kind of communications possibility and to be fully usable to non-IT people, including for the modeling of the production processes. The developed solutions also allow the consumer to benefit from the traceability through information pages available by scanning the QR codes available on the finished products (clothes, clocks, etc.). This paper presents: i) the methodology applied to achieve those functionalities, ii) the design and implementation choices, and iii) the test results. The main value of this paper is the usage of the Semantic Web in real-world industrial traceability solutions, which were tested in real supply-chains in Switzerland and India. The commercialization of the developed solutions is in preparation

1 INTRODUCTION

In 2007, a new European regulation on chemicals called REACH[1] initiated the creation of a catalog of potentially dangerous substances actively used in everyday consumer goods. Many of these have been identified as potentially threats to human health and are therefore forbidden on the European territory.

REACH puts the responsibility on the industry for evaluating and managing the risks about chemicals they use or import. As a consequence, industrial actors henceforth have the duty of tracing all substances in use in the products they manufacture or import in Europe. Unfortunately, in practise this is barely the case. As an example, during the EURO 2012 soccer championship, Europe's Consumer Watchdog revealed unusually high concentrations of dangerous chemical substances in several team shirts that could potentially be harmful to fans' health[2]. Independent tests highlighted high concentrations of lead, nickel and organotin, a chemical that can irreversibly damage the human nervous system. That's just one example among many others, but it shows a major failure of brands and the clothing industry in their capacity of fully capturing or even understanding their own manufacturing processes. When a problem is finally discovered, it is often already too late and the cost of any corrective measure is usually too high.

The case shown above is unfortunately too common and stands as a motivation driver for the work described in this paper. Bringing supply-chain traceability to the end consumer and to the economic partners (resellers or brands) is a difficult task. Information related to the manufacturing process is usually not made available (obfuscated on purpose

---

[1]REACH - Registration, Evaluation, Authorization and Restriction of Chemical Substances - http://ec.europa.eu/environment/chemicals/reach/reach_intro.htm

[2]http://news.stepbystep.com/euro-2012-football-fans-warned-against-buying-toxic-shirts-313/

or by lack of sufficient means) or extremely opaque. However, in order to become REACH-compatible, industrial actors will have the difficult task of motivating each participant belonging to the process to commit to the gathering of the necessary information.

The work described in this paper provides the outline of a first answer to the following research questions: what is the best way to represent data related to a supply-chain, so that it is flexible for future adaptations and generic enough to be applied to several industrial domains? How to design a software solution that helps non-IT specialists in defining their supply-chains themselves and that assists their navigation through the acquired data? How to provide an easy tool-set for the acquisition of process information ? How to give a full transparent access to this information through the entire chain, so that we can really support traceability? To tackle those problems pragmatically, we initiated the Fairtrace project, a generic traceability framework. The prototype presented in this paper is applied to the textile industry.

The objectives of the project are primarily the identification of the supply-chain (activities), order management and monitoring, as well as problem reporting (such as dangerous substance detection). It is intended to be a highly generic platform supporting potentially any kind of industrial supply-chain (clock industry, cocoa, ...); to be usable in harsh environments (mobile appliances) without any kind of communication facilities and to be usable by non-IT people (including the modeling of their production processes). Our solution must also benefit the final consumer (buyer), allowing him/her to obtain information on the traceability with a simple scan of a QR code on finished products (clock, cloth, ...). To support our objectives, we decided to leverage Semantic Web Technologies (SWT) in the core of the Fairtrace software to assess potential advantages when applied to an industrial setting.

A fully functional prototype has been designed and implemented. Field tests have been realized in India in February 2012. We have collected real-time information triggered by an order for an organic fair tee-shirt. A startup company founded by our commercial partner Importexa is currently preparing the commercialization of a product based on Fairtrace applied to the textile industry. Other markets will also follow soon. A part of the technology has already been patented (a dynamic formular creation system based on semantic data (Werlen et al., 2012)).

This paper is structured in five parts: the first part discusses some aspects of research in specific areas of supply-chain management. In the second part, we detail in descriptive terms the methodology of our work. Then, in the third, we present the results of the project and various discussion issues in the fourth part. We then conclude with some future directions and challenges still to overcome.

2 RELATED WORK

Fairtrace aims at achieving an agile traceability system that can help in the management and monitoring of supply-chains. Supply-chain management (SCM) is concerned with the coordination of activities for producing a product demanded by a customer (Mentzer et al., 2001)(Chandra and Kumar, 2001). Because thoughtful management of the supply-chain often results in substantial cost savings, the field has seen considerable industrial and research activity. It is no surprise then, the literature on the subject is so abundant that all aspects are well covered. (Lambert et al., 1998)(Lambert and Cooper, 2000) describe issues and opportunities of SCM. Other works propose formal SCM frameworks (Grubic and Fan, 2010), infrastructures (Fox et al., 2000)(Christopher, 2005) and models (Ye et al., 2008b). Logistics is a very active field where research focuses especially on optimization patterns (Bowersox et al., 2002). Supply-chain monitoring (SCMo) is a sub-branch of SCM that operates on information, inventory management and cash-flow (Lambert et al., 1998). The monitoring strives to rapidly identify problems in the supply-chain and solve them with the help of established procedures.

The focus of this paper is on supply-chain traceability (SCMt), which primarily addresses the problem of tracing goods from raw materials to finished products. An identifiable trend on SCMt is the use of external technologies such as RFID to enable traceability (Kärkkäinen, 2003)(Kelepouris et al., 2007), including in the textile industry (Kwok and Wu, 2009). Many publications also exist on SCMt processes and implications, such as (Opara, 2003). An important research track has been also developed on the composition of Web services in the supply-chain in order to ensure traceability among the partners of the chain (Kim and Jain, 2005). These works have been extended by enriching Web services with Semantic descriptions, in order to avoid ambiguity among the services and allowing an explicit meaning of the data interchange (Mocan et al., 2006). These Semantic Web Services (Paolucci et al., 2002) and their coordination (Schumacher

et al., 2008) allowed to improve mutual sharing of information for business-to-business integration.

In order to define a common understanding of supply-chain models, several research works proposed the use of supply-chain ontologies (Grubic and Fan, 2010). For instance, (Ye et al., 2008a) proposed quite a complex ontology that is not industry specific. Those work show that Semantic Technologies can have many advantages for the interoperability of the business partners in the chain. Research works also exist to use Semantics with RFID solutions (Virgilio et al., 2010). (Bechini et al., 2008) also proposes semantics description to ensure traceability. Other interesting research papers describe the usage of autonomous agents to deal with the information in the supply-chain (Fox et al., 2000), including its combination with Semantic Web descriptions (Datta et al., 2007).

In relation to the research above, the work presented in this paper can be specified as follows: Fairtrace is meant to trace all activities of the entire manufacturing process, including documents and products. It is able to identify the raw materials which were used, certificates, waste, etc. It can also be used to identify roles and responsibilities and to keep track of who did what. It does not currently include optimization techniques for the logistics. Fairtrace can track many variables, which can be used as performance indicators. A customizable chain of validators also allows to track problems arising along the supply-chain, such as the detection of dangerous chemical substances. Businesses can thus react very quickly to any unexpected condition. From a technological perspective, Fairtrace was conceived with a focus on Semantic Web Technologies (SWT). Every piece of data in the system is related to a concept in a semantic data repository. All models use semantic descriptions and traceability information can be displayed to the user hierarchically.

The next sections explain thoroughly how Fairtrace was conceived, and its main results.

3 METHODOLOGY

The scoping phase of the project consisted in capturing and understanding the requirements of traceability. An analysis of a typical textile manufacturing chain was made to get a rough idea on the process. A consultant specialized in the textile industry was commissioned to analyze Importexa's own supply-chain in India. Her mission was to visit factories, identify and document all activities throughout the manufacturing process. She captured

information about certificates, delivery challans, cotton lots, mixing lots and many other documents, but also took pictures of places and production machines. Our task was thereafter to analyze all those paper resources in order to sketch out a formal description of the longitudinal process. The very first model was a transcription of the complete, moderately detailed manufacturing process using the primitives of BPMN[3].

The process was complemented by several additional documents more accurately describing the information to be captured on the supply-chain. A critical path — a minimal path from the starting activity (order) to the end (ginning)— was defined. The method consisted in searching through pairs of documents such as orders, challans or bills for matching identifiers to find a complete and continuous traceability path up to the origin of the cotton. The base model was augmented with additional attributes (GSM, weight, ...) extracted from the documents that we analyzed. This extra information helps tracking quantities and waste.

3.1 Requirements

Fairtrace objectives focus on usability, supply-chain data transparency, genericity and adaptability.

As a consequence of the first requirement, business partners in the supply-chain needed a simple way to enter data about their production directly into our system. Among all alternatives considered, a system of minimally constrained web formulars was retained. On one hand, it was important to give users the freedom to design themselves these formulars — possibly by drag-dropping GUI components — without having to resort to any particular IT knowledge or skills. On the other hand, formulars needed to the bound to the underlying data model. Web formulars coupled to a flexible data binding system offer the kind of usability that was required in this project. Web forms can be implemented in pure HTML and are supported natively by all kinds of devices. They can be adapted to any sort of display they are affected to, with relatively few lines of CSS code. They can be edited and pre-filled with data coming from the repository and customers can mark fields optional or required at wish.

The formular system was also designed to be adaptable. Each formular is dynamically linked to the underlying data model by a dynamic binding system that uses specific binding names to associate fields from the formulars to particular objects and properties

---

[3]BPMN - Business Process Modeling and Notation - http://www.bpmn.org

in the model. Coupled to an instantiation engine that creates model instances based on the data acquired, the whole system provides a very generic framework that can adapt to almost any model or domain, without requiring recompilation.

To provide transparency on the data, it is important that information can be accessed without any technological barriers. A dashboard system providing supply-chain monitoring support and capable of providing easy data exploration was designed as a consequence. Furthermore, we wanted the information to be presented to the customer both horizontally (information for a specific time frame) and vertically (hierarchy of data). Such a system would allow anyone entitled to do so, to descend into any part of the traceability data chain providing them with a full transparent access to its contents.

3.2 Domain modelling

The data model in Fairtrace was designed to capture many different aspects of the supply-chain, including details on the process itself and on the products deriving from it. We divided the modeling task into creating models for both.

To improve on usability, we wanted to empower the customers (or any one entitled to do so) to design graphically their supply-chains. We had to provide the necessary primitives (activity, flow) to allow it. For that reason, we did not hard-code the process into a model for each different customer, but decided instead to specify a language partly derived from BPMN describing business process concepts such as activities, flows, roles, users, partners, authorizations and collection points. Supply-chains would then be specified in terms of these primitives.

Currently, our model only provides support for a tiny subset of BPMN-like constructs. We added a few more custom primitives (collection points) that are used by the security granting mecanism. Before resigning ourserves to model everything from scratch, we tried on different approaches. We had a look to the PSL[4] specification and ontogy, M3PO(Haller et al., 2006) and BPMNO[5]. Unfortunately, the complexity of the semantics and the verbosity of those models seemed a bit "overkill" for our needs.

The product model is strongly bound to the industrial domain it describes; in this case, the domain was the textile supply-chain. We searched on-line for existing ontologies describing domain knowledge

---

[4]PSL - Process Specification Language - http://www.mel.nist.gov/psl/

[5]BPMO - BPMN Ontology - https://dkm.fbk.eu/index.php/BPMN_Ontology

on textile products, but unfortunately, we could not find any that more or less suited our needs. We decided thus to also model the product domain from scratch. To do so, we've literally analyzed dozens of documents brought back from India by the consultant.

3.3 Semantic technologies

The leveraging of Semantic Web Technologies was a prerequisite in this project. We didn't have an extensive knowledge on the subject when the project started. It was thus important for us to stay simple and practical at all times. In order to cope with this requirement at every step, we did a lot of testing on example datasets. We only started implementing the ontologies after we felt confident enough that there would be no uncontrolled side effects (i.e inferences that were not expected).

Our models didn't require a high level of expressiveness, since they were almost one-to-one mappings from an object-oriented hierarchy. We didn't use elaborated constructs like Restrictions, because we wanted to be sure of the decidability. We also needed the language to be sound, complete (all logical consequences are drawn) and monotonic (i.e all statements remain valid after inserting new knowledge). We translated the conceptual models (process, products) to RDFS instead of a more expressive OWL, not only because it was easier to work with, but also because semantics were easier to grasp. Our models are based on subsumption hierarchies that use sub-class and sub-property relations (`rdfs:subClassOf`, `rdfs:subPropertyOf`) extensively.

The process model does not put any additional requirements on the predicates it declares. No particular semantics are specified, because this model is very similar to what would be done on an object-oriented programming language. The product model however, needs more expressiveness in order to model the existing links between objects of different classes, but also to compensate for potential gaps due to the lack of sufficient data. We had to model inverse properties (`owl:inverseOf`) and property chains axioms (only available in OWL 2). In order to avoid the extra cost of using OWL, we've decided to take an intermediate route and create extra entailment rules for these very few specialized cases. Here are some examples in *predicate logic*:

Simulating owl:inverseOf:
```
inOrder(X,Y), Type(X,Batch),
Type(Y,Order) => hasBatch(Y,X)
```

Simulating property chain axioms:
```
hasFabric(X,Y), hasYarn(Y,Z),
Type(X,Order), Type(Y,Fabric),
Type(Z,Yarn) => hasYarn(X,Z)
```

3.4 Infrastructure

The prototype infrastructure was designed to meet requirements for the three types of customers we wanted to target: business partners, end users and supply-chain partners.

The business partners category is mostly composed of brands and resellers that will license the future Fairtrace solution for monitoring their own supply-chain. Their needs are primarily basic-management controls, extensive information and problem reporting. Those requirements were implemented into the prototype as a supply-chain designer, a graphical formular edition interface and a dashboard to control all operations, assign users, roles and authorizations.

To end users we wanted to provid them with traceability information about a product on scanning a QR code on the finished product. We had to design a system to map these codes to specific views on the information collected from the data store. That information is then formatted to target the particular device the user browses on.

Finally, we had to provide a way for the supply-chain partners to send data to the system. By partners, we mean any entity working in the supply-chain to create and finalize the product. These can be factories, transporters, single users... The dynamic formular system we designed was meant to be simple and functional. It supports simple common controls such as date pickers, file selectors, picture boxes, text fields, lists, combo boxes and other types of selectors. It is thus possible to integrate certificates (such as for biological production) as PDF/A files as well as other kind of documents. Those binary resources are stored on the cloud and a link is kept inside the RDF data store. The formular design and the dynamic binding systems were thought to be used by designers not especially proficient with technology. We favored simplicity and wanted to hide the technical details (URIs, JSON object, ...).

Furthermore, we also considered both mobile and browser-enabled applications to display formulars. The mobile application was designed to manage and display the formulars in potentially non connected configurations. We had to take into account that wireless signal could be a problem in some places where the application was meant to be used (such as in cotton fields in India). We designed an offline mode, that would allow later synchronization with the Fairtrace server. The browser-enabled application was meant to be used directly in factories, where Internet is generally available. We considered this option too, because it was obviously easier and faster to input data from a keyboard rather than on a mobile.

3.5 Implementation

The actual implementation phase went surprisingly well. We opted for an agile development process with a relatively short feature development cycle (<1 month). We were able to finish both on time and on budget. We spent extra resources on polishing important features, so that the prototype would be ready for a demonstration to potential customers. The last step of the development was testing the prototype in a real situation. We modeled Importexa's supply-chain and requested them to initiate the process. Importexa started an order for an organic fair cotton tee-shirt. When the actual manufacturing began, we had the consultant in India to teach the functioning of the system and to check that everything was going as planned. We wanted to show off that the prototype was able to trace all the manufacturing activities and capture all the products of it.

4 RESULTS

This project yielded three main results, which are the functional prototype, two ontologies created from the process and product models and the system and the modeling of a textile manufacturing chain.

4.1 Prototype

The prototype shown on Figure 1 is a three-tier infrastructure composed of a data, a business and a presentation layers and was developed in Java as a Spring MVC application running on a Tomcat 6 servlet container. The data layer hosts the RDF store (semantic repository). The middleware is comprised of the business layer and acts as a message translator and rule enforcer between the application server and the RDF store. Finally, the presentation layer contains both the Business To Business (B2B) and front-end operators. A Business To Consumer (B2C) front-end solution maps QR codes that come on the final products with information views about the whole upstream chain to be presented to the final consumer.

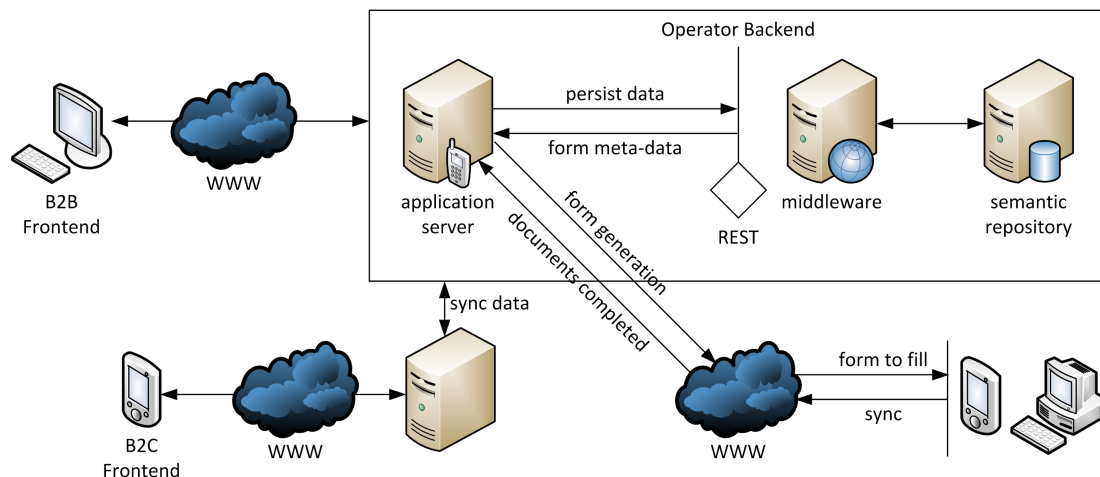The semantic functionality (i.e. semantic store) is provided by the BigOWLim 3.5 engine from

Figure 1: Overview of the architecture.

OntoText [6] running as an extension of the OpenRDF[7] Sesame 2.4.0 runtime. A set of Data Transfer Object (DTO) classes provide seamless data access and querying functionality. The prototype exclusively uses SPARQL 1.0 queries, since SPARQL 1.1 UPDATE was not supported by the versions of OWLim and Sesame used in the prototype. Data had to be programmatically inserted, updated or removed. Queries for traceability data are made against the implicit graph containing the inferred triples. Queries on the process itself are made against the explicit graph and returned to the API caller via the DTOs.

Listing 1: Product object encoded in JSON

```
property_set: {
  properties: [
    {
      property: {
        name: "order_id",
        label: {
          "en": "original order number"
        },
        value: "CA16576"
      }
    },
    {
      property: {
        name: "reference_id",
        label: {
          "en": "purchase order number"
        },
        value: "2219"
      }
    }
  ],
  object_id: "order001"
}
```

The business layer is comprised of a set of classes that receive requests through a series of REST Web service endpoints. The requests and responses are encoded as JSON objects. Payloads are translated, verified, business rules are enforced, security applied and then forwarded as a data access requests. Data extraction is done with SPARQL `SELECT` and `CONSTRUCT` queries. Repository modifications are done in transactions and are rolled back in case of problem. Service endpoints that manage process resources (objects) use the CRUD paradigm (Create, Retrieve, Update, Delete) and produce JSON-encoded messages. Data requested for products is returned as a group of properties along with the identifier of the individual (Listing 1).

The application server was built on the latest version of Ruby on Rails and coordinates all requests between the presentation layer components and the back-end storage. It also enforces a certain number of business rules, but its role is more to coordinate B2C and B2B actions, unlike the back-end, which role is to coordinate data requests. The B2B front-end is the Web-based control panel of the system. It is intended to be customizable to businesses willing to license the Fairtrace ecosystem. It allows to dynamically and graphically model the supply-chains, creating the formulars and binding them to any step of the associated process. It can monitor order progress by showing various indicators and track any upstream issue such dangerous chemical detection. Users can then navigate vertically through traced data (Figure 2) or horizontally (Figure 3) by querying the date inside a specific time span.

The Operator front-end is aimed at supply-chain partners working in coordination. It presents them formulars that were previously created in the B2B front-end for them. It allows them to send data

▷ 📁 fabric	fabric001
▷ 📁 fabric	processed_fabric001
▷ 📁 yarn	yarn001
▷ 📁 yarn	yarn013
▷ 📁 color	color001
📄 order id	CA16576
▷ 📁 batch	batch001
▷ 📁 transaction	transaction001
📄 fairtrade id	18753
📄 reference id	2219
📄 buyer code	E 150
▷ 📁 processed fabric	processed_fabric001

Figure 2: Vertical navigation through data hierarchy.

Figure 3: Horizontal navigation through data.

related to their activities into the system. On the mobile appliances, the frontend is an HTML5 container application that allows to download pending formulars created on the B2B front-end, fill them, update them and synchronize them back with the application server. It provides support for an offine mode, where data can be synchronized back later.

In the prototype, the detection of dangerous substances was not codified in the model ontologies. We decided to develop that functionality with a chain a validators. Each validator is created as a plug-in, which can analyze a subset of the dataset and generate messages on certain conditions. There are three types of messages: exception, warning and informational. Validation can be trigged manually from the dashboard or can be set to be triggered automatically upon any modification of the data. The validator for dangerous chemicals validates each individual of class `ft:Substance` by matching it against a database of known harmful chemicals. One goal set for the next iteration of Fairtrace is to use external sources directly integrated with our models in order to be able to infer the `fto:DangerousSubstance` class and take advantage

of the inferencing mechanisms.

4.2 Ontologies

Implementing the models was a bit delicate in that a clear separation of concerns was necessary. Having a multi-disciplinary team, we needed a modularized design in order to allow parallel and effective development. In order to create the ontologies based on our models a RDF repository to store ontology data was required. The choice of BigOWLim as the semantic repository and reasoner is the result of an objective test, where different semantic repositories were compared based on a few metrics (size, simplicity, ..). The BigOWLim 3.5 engine works as an extension of a SAIL repository in Sesame. It is easy to setup, provides really fast inferences on vast amounts of data, supports geo-spatial and full-text search capabilities.

BigOWLim Triple Reasoning and Rule Entailment Engine (BigTRREE) uses a total materialization reasoning strategy that computes inferences after insert/update/delete cycles. That kind of strategy has the disadvantage of making inferences a bit slower (especially on very large datasets), because it computes the complete closure (optimizations aside) on each transaction. Even so, it does allow extremely fast queries, almost on par with traditional DBMS systems. The Fairtrace prototype heavily relies on queries, so BigOWLim was certainly a righteous choice. An additional feature that also helped in our choice was that it was actually very easy to deploy and write entailment rules for. We made an agreement with OntoText to use their repository technology free of charge in our research projects.

Almost every object used by the web and mobile platforms is stored in the repository A-BOX as an individual from either the process ontology or the product ontologies. Each ontology is the implementation of a conceptual model in RDF Schema[8].

The process ontology defines the building blocks for modeling business processes. It is a straight mapping from an object model to RDF. It does not define any particular semantics, except the ones explicitly given in the RDF Schema specification (`?x rdfs:type rdf:Resource` for instance). It builds around concepts like steps (activities), flows (links), users, roles, authorizations, partners, collection points and formulars. Formulars are assigned to a particular collection point linked to a specific step. Authorizations are defined for collection points

---

[8]`http://www.w3.org/TR/rdf-schema/`

to state who and on what role can enter data and download the formulars. The intent of these primitives is not to fully map the BPMN specification, but to present a synthetic view of the supply-chain. Conditions, loops and other advanced primitives are not yet supported. We did not reuse existing ontologies, because we needed a simple and efficient ontology.

The product ontology models the domain objects. In contains definitions for textile industry objects, such as fabrics, parts, dyes, yarns, mixing lots. Concepts are linked together – in a critical path – so a customer can eventually navigate through individuals from the initial order down to the geographical place where the cotton was produced. Every step and every product of the supply-chain can thus be identified. Documents and special objects, such as certificates are also modeled and can be linked to binary data (pictures). They can be used to prove that cotton is really organic and fair and that all authorizations have been obtained by a partner for instance. Importexa can enforce that business partners insert them in the system. In terms of semantics, the model uses mostly subsumption hierarchies. We assert for example:

```
fto:Certificate rdfs:subClassOf
fto:Document .
fto:ProcessedFabric rdfs:subClassOf
fto:Fabric .
fto:hasCertificateName rdfs:subPropertyOf
rdfs:label .
```

Furthermore, to compensate for the lack of property-chains (only available in OWL 2), we defined custom rule-sets (see Section 3.3). These are necessary in order to bridge the gaps in the model that can occur due to the lack of sufficient data, causing thus an impossibility to descend in the hierarchy.

4.3 Supply-chain modelling

We also modeled a simple textile supply-chain obtained from the analysis of the information taken in the supply-chains in India. According to our consultant, our process model, even simple matches almost 90% of existing real chains. It models all upstream steps such as Mixing, Spinning, Dyeing, Processing, Transforming, Weaving, Sewing and Packaging. Ginning was purposely left aside, since information about the provenance of cotton was readily available at the Mixing step.

4.4 Testing

The infrastructure was deployed and tested in India. The objective was purely to assess its flexibility. We did not acquire quantitative indicators such as number of queries for each order, etc., because it was materially (i.e. financially) not possible to launch enough real orders to gather enough data for the results to turn out meaningful. We concentrated thus on a feasibility study following the project plan methodology:

- We selected the supply-chain partners in India that usually work with the textile company Importexa in Switzerland ;
- A consultant from Importexa visited the partners on site in order to thoroughly document the supply-chains ;
- The consultant described formally the supply-chains using a Business Process Modeling notation ;
- On the basis of the consultant documentation (data description, business process model, etc.), we designed in our Web-based B2B Frontend the processes and formulars acquiring information inside those formulars ;
- Those formulars were instantaneously deployed for data acquisition, both as HTLM5-enabled mobile and Web formulars ;
- Importexa launched a real order for an organic fair cotton tee-shirt, whose production data should be acquired with our solutions ;
- A consultant was then sent to India in order to teach working partners how to input data on the forms and send it for treatment ;
- A final report described the results of the tests .

The tests ran for many days and a couple of formulars had to be re-adjusted to better match the partner's requirements. We were reactive and could almost immediately update the formulars and they were able to use it straight away. As the tests were made on a real Importexa order, we could monitor data incoming from all steps including all chemicals and certificates (GOTS, Fairtrade, GMO free, and so on) in realtime.

5 DISCUSSION

Fairtrace is regarded as a success on management, collaboration and technology support levels. The project encountered high commercial concrete

Figure 4: Test data acquisition in India.

interest when it was presented and demonstrated in a technical session to potential customers. One customer is ready to deploy the solution for his supply-chains. The project plans to be integrated in a start-up with the help of the CTI Start-up[9], an important Swiss coaching institution. Although we had no quantitative metrics for the tests, we could assess on its success in a real scenario. The testing in India has faced many unexpected real-world conditions that could not be foreseen. We were able to respond to problems and modification requests quickly (in a matter of minutes) and could attest of the flexibility of the prototype.

Importexa is now heading to the creation of a start-up that will commercialize the finished product. Fairtrace is going be improved and adapted to market segments other than the textile industry. A survey on potential markets has already been done and first discussions have already taken place with interested brands and Swiss resellers, in order to adapt the tool to brand new markets (clock industry, controlled cheese production, cocoa, ...). Supply-chain traceability can be a strong selling argument, because it implies that companies master the entire production chain, can react quickly to potential problems (such as to follow the REACH regulation) and can prove their customers that they exactly know what they produce.

In this project the use of Semantic Web Technologies was a prerequisite, because it was a good opportunity to explore its possibilities for future projects. We were already a bit familiar with the tools and techniques, but not sufficiently aware of its benefits or complexities to apply it in future projects. We had to work out many issues and carefully think about our solutions. Still, Fairtrace proved that SWT are safe for use in commercial projects, because they are now sufficiently mature. We kept things very

simple at the moment, because we had a real need for pragmatism, so we only had a glimpse at the full range of possibilities that Semantic Web has to offer. Our model was not thought in terms of inferences it could entail, but rather in terms of object orientation. We needed to translate RDF data into Java classes and thereby lost part of the expressiveness power of these technologies. To compensate for this, the prototype relied heavily on queries. The lack of SPARQL 1.1 UPDATE support has also somewhat hindered our efforts. Programmatically inserting/updating/deleting triples in the store was difficult, verbose and error prone.

Although we did only use a small subset of it, Semantic Web Technologies still provided a real competitive advantage, because they allowed a greater modeling flexibility in comparison to other database technologies. Translating the logical models to RDF Schema was simple and quick. The expressiveness of the language coupled to the rule-based OWLIM reasoning engine was good enough to allow sufficiently complex constructs. There is still a lot we can do to improve the current state of things: we can link geographical places to GeoNames[10] features for instance, redirect specific concepts to DBpedia[11] definitions or even automatically categorize items as dangerous by making use of bridges to external datasets. In next iteration of Fairtrace, we'll definitely investigate on integrating OWL features like (inverse) functional, symmetric and invese properties, as well OWL2 property chain axioms vs. creating our own rules, also on integrating geo-spatial extensions (for marking places) to reason about distances and, of course, the new SPARQL 1.1 constructs.

6 CONCLUSION

This paper presented a pragmatic semantic-based research solution to the traceability problem in supply-chains. The first prototype is finished and has been successfully tested in a Swiss and Indian supply-chain in the textile industry. However, Fairtrace is still an ongoing work and the roadmap contains many improvements. SWT will continue playing a central role in the future infrastructure. We are planning to use the next-generation OWLIM semantic repositories and take advantage on the new constructs offered by OWL2 and SPARQL 1.1 UPDATE. We are also

planning important improvements on the creation of formulars, on the customization of the system and on the primitives for modeling supply-chains. Finally, we consider intelligent agent techniques for automatically queryingthe semantic repositories to help the decision support of users.

ACKNOWLEDGEMENTS

REFERENCES

Bechini, A., Cimino, M. G. C. A., Marcelloni, F., and Tomasi, A. (2008). Patterns and technologies for enabling supply chain traceability through collaborative e-business. *Inf. Softw. Technol.*, 50(4):342–359.

Bowersox, D., Closs, D., and Cooper, M. (2002). *Supply chain logistics management*, volume 2. McGraw-Hill New York.

Chandra, C. and Kumar, S. (2001). Enterprise architectural framework for supply-chain integration. *Industrial Management and Data Systems*, 101(6):290–304.

Christopher, M. (2005). *Logistics and supply chain management: creating value-added networks*. Pearson education.

Datta, S., Lyu, J., and Ping-Shun, C. (2007). Decision support and systems interoperability in global business management. *International Journal of Electronic Business Management*.

Fox, M., Barbuceanu, M., and Teigen, R. (2000). Agent-oriented supply-chain management. *International Journal of Flexible Manufacturing Systems*, 12(2):165–188.

Grubic, T. and Fan, I.-S. (2010). Supply chain ontology: Review, analysis and synthesis. *Comput. Ind.*, 61(8):776–786.

Haller, A., Oren, E., and Kotinurmi, P. (2006). m3po: An ontology to relate choreographies to workflow models. In *Services Computing, 2006. SCC'06. IEEE International Conference on*, pages 19–27. IEEE.

Kärkkäinen, M. (2003). Increasing efficiency in the supply chain for short shelf life goods using rfid tagging. *International Journal of Retail & Distribution Management*, 31(10):529–536.

Kelepouris, T., Pramatari, K., and Doukidis, G. (2007). Rfid-enabled traceability in the food supply chain. *Industrial Management & Data Systems*, 107(2):183–200.

Kim, J. W. and Jain, R. (2005). Web services composition with traceability centered on dependency. In *Proceedings of the Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05) - Track 3 - Volume 03*, HICSS '05, pages 89–, Washington, DC, USA. IEEE Computer Society.

Kwok, S. K. and Wu, K. K. W. (2009). Rfid-based intra-supply chain in textile industry. *Industrial Management and Data Systems*, 109(9):1166–1178.

Lambert, D. and Cooper, M. (2000). Issues in supply chain management. *Industrial marketing management*, 29(1):65–83.

Lambert, D., Cooper, M., and Pagh, J. (1998). Supply chain management: implementation issues and research opportunities. *International Journal of Logistics Management, The*, 9(2):1–20.

Mentzer, J., DeWitt, W., Keebler, J., Min, S., Nix, N., Smith, C., and Zacharia, Z. (2001). Defining supply chain management. *Journal of Business logistics*, 22(2):1–25.

Mocan, A., Moran, M., Cimpian, E., and Zaremba, M. (2006). Filling the gap - extending service oriented architectures with semantics. In *Proceedings of the IEEE International Conference on e-Business Engineering*, ICEBE '06, pages 594–601, Washington, DC, USA. IEEE Computer Society.

Opara, L. (2003). Traceability in agriculture and food supply chain: a review of basic concepts, technological implications, and future prospects. *JOURNAL OF FOOD AGRICULTURE AND ENVIRONMENT*, 1:101–106.

Paolucci, M., Kawamura, T., Payne, T., and Sycara, K. (2002). Semantic matching of web services capabilities. *The Semantic WebISWC 2002*, pages 333–347.

Schumacher, M., Helin, H., and Schuldt, H. (2008). *CASCOM: Intelligent Service Coordination in the Semantic Web*. Birkhäuser Basel, 1st edition.

Virgilio, R., Sciascio, E., Ruta, M., Scioscia, F., and Torlone, R. (2010). Semantic-based rfid data management. In Ranasinghe, D. C. C., Sheng, Q. Z. Z., and Zeadally, S., editors, *Unique Radio Innovation for the 21st Century*, pages 111–141. Springer Berlin Heidelberg. 10.1007/978-3-642-03462-6˙6.

Werlen, D., Alves, B., Cherix, G., Gapany, C., and Schumacher, M. (2012). Procédé d'édition de formulaires pour la saisie de données en différents points de collecte. Swiss patent application CH00104/12 filed on Jan. 25, 2012.

Ye, Y., Yang, D., Jiang, Z., and Tong, L. (2008a). An ontology-based architecture for implementing semantic integration of supply chain management. *Int. J. Comput. Integr. Manuf.*, 21(1):1–18.

Ye, Y., Yang, D., Jiang, Z., and Tong, L. (2008b). Ontology-based semantic models for supply chain management. *The International Journal of Advanced Manufacturing Technology*, 37(11):1250–1260.

---

[12] http://www.kti-cti.ch