

Knowledge Transfer: From Maintenance to Engine Design

Sylvia C Wong, Richard M Crowder, Gary B Wills, and Nigel R Shadbolt

School of Electronics and Computer Science

University of Southampton, UK

sw2@ecs.soton.ac.uk, rmc@ecs.soton.ac.uk, gbw@ecs.soton.ac.uk, nrs@ecs.soton.ac.uk

July 20, 2007

Abstract

As manufacturers shift their focus from selling products to providing services, designers must therefore increasingly consider the life-cycle requirements in addition to conventional design parameters. To identify possible areas of concern, engineers must consider knowledge gained through the life cycle of a related product. However, because of the size and distributed nature of a company's operation, engineers often do not have access to front-line maintenance data. Additionally, the large number of documents generated during the design and operation of a product makes it impractical to manually review all documents thoroughly during a design task. This paper presents a prototype knowledge-based document repository for an aero-engine manufacturer. The developed system searches and analyzes distributed document resources, and provides engineers with a summary view of the underlying knowledge. The aim is to aid engineers in creating design requirements that incorporate maintenance issues. Unlike existing document repositories and digital libraries, our approach is knowledge-based, where users browse summary reports instead of following suggested links. To test the validity of our architecture, we have developed and deployed a prototype of our knowledge based document repository. The repository has been demonstrated to, and validated by the engine design community.

1 Introduction

The design and maintenance of large and complex engineering systems requires a significant amount of documentation, particularly if the system being considered is an aero-fan used on the current generation of aircraft. The aero-engine is amongst the most complex machine ever designed, incorporating a wide range technologies including high temperature materials, complex fluid dynamics and parts subjected to high mechanical stress due to high rotational speeds.

A fundamental shift is currently occurring in the aerospace industry away from selling products to providing services. Companies such as Rolls-Royce aim to make half its engine fleet subject to long-term maintenance service agreements by 2010 [13]. Essential to the success of this market shift is the significant cultural change from *offering a service to support a product to designing a service and the product to support it* [13]. In other words, new engines must be designed to provide lower and more predictable maintenance costs. To minimize maintenance costs through out an engine's life cycle, design engineers must obtain knowledge gained from maintenance histories of similar products during an engine's design phase. This will help engineers identify engine parts most likely to be problematic throughout the engine's life cycle. It should be noted that engine design is typically undertaken by a number of teams who are responsible for individual engine modules, e.g compressor, turbine, etc. Therefore it is impossible for any single member of a design team to access more than a fraction of the available documentation. As is widely recognized, information systems usually develop over time into a set of heterogeneous resources. Moreover these document resources usually have ill-defined metadata, which makes it difficult to find relevant information. As a result, it becomes difficult for engineers to follow a trail through

the resources [30]. The challenge for organizations is therefore to develop an information system that is both comprehensive and will satisfy the increasing demands from industry for up-to-date and easily accessible information.

In response to these challenges, we are implementing an intelligent, knowledge-based document repository to support engineers to design for the aftermarket. The intelligent document repository searches and analyses relevant maintenance records and design guidelines, and provides design engineers with easy access to such information. It is hoped that the summary reports provided by the intelligent document repository will help engineers in creating design documents that incorporate aftermarket issues into the design requirements.

This paper is organized as follows. Section 2 explains how knowledge can be re-used within a aeroengine manufacturing company. It also introduces a scenario explaining how maintenance records can be used to help improve the reliability of both existing and new products. Section 3 gives a brief overview of knowledge management. Section 4 discusses other works that aim to help user discover knowledge by integrating heterogeneous document sources. Section 5 describes our proposed architecture for an intelligent document repository. Section 6 discusses aspects of the system and the initial reaction of the user community and our conclusions are presented in Section 7.

2 Scenario

As is well recognized in engineering design, the use of past experiences and previously acquired knowledge, either from the designer's own experiences or from resources within their organization forms an important part of the design process. It has been estimated that 90% of industrial design activity is based on variant design [10], while during a redesign activity up to 70% of the information is taken from previous solutions [17]. A cursory consideration of these figures identified two immediate challenges – how to capture knowledge, and how to retrieve it. The purpose of the work reported in this paper is to develop an intelligent document repository that can be used within a manufacturing organization for the retrieval of knowledge from across the organization to support design activities.

Figure 1 shows the key information flow for the different stages in the life of an engine. Concept design is the first stage of an engine's life cycle. Given a set of broad requirements, such as thrust, range of the target aircraft and fuel burn, engineers determine the approximate dimensions, weight, power and other physical characteristics for the engine design. The engineers also make estimates of the manufacturing costs of the engine. In the design stage, engineers transform the preliminary abstract design into a set of concrete plans that can be used in production. In production, engines are built according to the design plans. Traditionally, after production, responsibility for the engine passed from the manufacturer, to the airlines, who own the engines. The airlines are responsible for maintaining the engines. This maintenance activity is supported by the manufacturer's technical support and operations team. To assist maintenance engineers to identify problems before a breakdown occurs, engines are commonly equipped with sensors for engine monitoring. This monitoring information can be analyzed for abnormal operating conditions, such as temperature or pressure. However until fairly recently, the monitoring data was only used to support maintenance activities, even though it is a rich source of information for the designers of future engines [15].

As shown in Figure 1, there are interactions and information flow between neighbouring stages in the production maintenance process. This is due to the iterative nature of engineering processes. Design knowledge is also passed to operations in the form of 'owner's manuals'. Sometimes, information also pass between unconnected stages, for example, between operations and engine concept. However, the flow is weak and may take the form of informal and personal networking between engineers.

While the process works very well, it does have significant disadvantages. In particular, design engineers are remote from the problems experienced in the field by operations. Due to the importance of increasing operational reliability and minimizing maintenance costs in the new market paradigm of product support, information gained in the operation of a fleet of engines needs to be

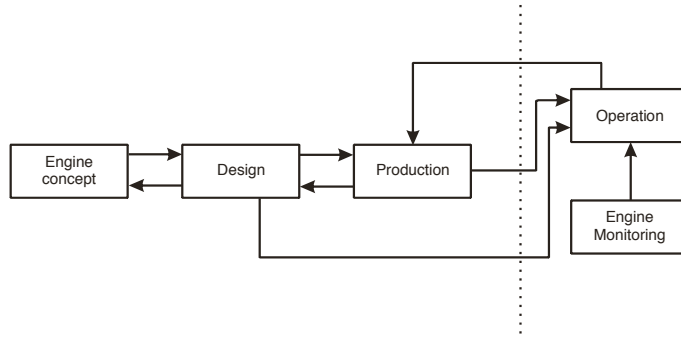


Figure 1: Information flow between the different stages in the life of an engine. The vertical line between production and operation represents the transfer of the engine from its manufacturer to its user, typically an airline. Operations is the generic term for maintenance and aftersales support.




Figure 2: We aim to facilitate the flow of information gained during the life-cycle of one engine variant to inform the design of the next variant. The dotted arrows indicate the flow of design rationale and related knowledge. The solid arrows represent all other information flows including design documentation and real time engine information.

fed back to the designers of subsequent engines. However, the current information infrastructure makes this difficult as concept design engineers do not have access to maintenance knowledge. Similarly, design engineers should consult existing maintenance documents to help design parts with more predictable maintenance costs.

As a result, we need to strengthen and help formalized the information flow between the company's aftermarket operations and the design teams. Figure 2 shows the information and knowledge flows our research aims to build. This would allow the knowledge gained during the design, production and operation of an engine to advice the design of the next variant.

The following scenario¹ illustrates the potential use and benefit from such a document repository. The scenario involves three separate and different groups of users that are involved in the life of a jet engine. Front-line maintenance engineers are involved in the day to day servicing of the engines, and are responsible for populating the document repository with maintenance reports and related documents.

During the regular pre-flight checks, a flight crew reported a problem with the leak from an engine's bleed air system. Inspection of the engine revealed that a duct had failed at a joint due to vibration. After repair, the aircraft was returned to service, and a full maintenance event report submitted to the document repository.

The document repository can then be used by technical support and operation engineers, who are responsible for improving the performance of existing engines. They can use information collected in the repository to monitor trends that develop over a fleet of engines. Modification can then be designed to mitigate any problems found:

Following a review of the maintenance events relating to a specific engine fleet, a trend was noticed in the higher than expected number of failure of a duct joint due vibration at a specific operating point. To maintain the reliability of the engine fleet, a modification was developed and implemented.

The same information in the repository will also be used by design engineers working on a new variant of the engine:

The design team for the next variant of this engine reviews the performance of the air bleed system across the fleet to learn from previous design rationale and operational history. The large number of failures were noted and detailed finite element analysis showed that a joint failure could occur due to vibration if certain operational conditions were met. It was therefore decided that the future variant of the engine would both eliminate the joint and reroute the duct work. The revised design costs 40% more than the original. However, the saving over the life of the engine will be substantial due to lower likelihood of in-service failure and consequential impact on operations.

The goal of our work can thus be summarized as follows: To feedback and harvest knowledge gained from the aftermarket operations documents to help (a) operations engineers in designing modifications to existing engines, and (b) design engineers in designing the next variant engine for the aftermarket.

3 Knowledge Engineering

Organizations have become increasingly concerned with knowledge management [27], amassing large amount of information into their corporate memory [14]. The aim is to use this repository of information to inform future discussions, decisions and activities. Figure 3 shows the four key activities in knowledge management – knowledge creation, knowledge mapping, knowledge retrieval and knowledge use.

¹The scenario is entirely fictitious and does not derive from any real event.




Figure 3: The four key activities in the knowledge management cycles.

The first step in creating a knowledge system is knowledge creation. In our scenario, this is the documentation written by engineers throughout the life time of an engine, from design to maintenance. In practice it can take many forms, including formal design reports, e-mails submitted by airlines and field service personnel, detailed inspection reports and the engine performance over individual flights recorded by engine monitoring systems.

Knowledge mapping involves creating an *ontology*, which is a *specification of a conceptualization* [11]. Gruber explains that a common ontology defines the vocabulary with which queries and assertions are exchanged among agents (people or software). The ontology sets out all the entities (objects or concepts) that we are interested in and the relationships that connect these entities together. This is intended to be a *pragmatic* definition, i.e. it defines the vocabulary that is actually *in use*, and the concepts that are *useful* in problem-solving. It does not give the deep underlying philosophical vision of the fundamental entities in the field. Hence, in knowledge management, an ontology is a tool, whose quality is entirely dependent on its usefulness.

Retrieval is the step that transform mere *information* into *knowledge*. We believe that knowledge is relevant information delivered at the right time and context [22]. To deliver this knowledge, information needs to be semantically enriched so that it can be better reused. When a knowledge system has a shared ontology for its disparate information resources, software agents can handle the semantically enriched resources consistently. Thus semantics and ontology together help deliver the right information at the right time, hence generating knowledge.

In our scenario, knowledge use occurs when engineers apply the knowledge gained from the repository in their work. For example, a design engineer applies the knowledge gained to create a better engine for the aftermarket. Or an operations engineer discovers a recurring, but minor problem with an existing engine, which would indicate a larger problem than each individual incident suggests.

4 Related Work

The work described in this paper is an extension on our previous work with Rolls-Royce. In [8], we presented a future vision for the working practices of designers within a manufacturing organization. We have found that engineering design environments are highly distributed in nature and are characterized by a large number of information sources, which together with the designers forms a complex sociotechnical system. It is concluded that a range of knowledge management tools would be required to support this future vision of engineering design environments [29]. Therefore one of the objectives of our current work is to define a future engineering design environment, with particular emphasis on the social and technical systems that will support designers in their day-to-day activities.

In [30], we created a document repository from distributed and heterogeneous engineering document resources. When an engineer searches for documents within the repository, the system generates a list of documents ordered according to the engineer's role and its related concepts. Thus, the document retrieval process is *intelligent* and adapts to the user's role within the company. However, we found that engineers actually want the knowledge that is buried within the documents, instead of the actual documents themselves. This is due to the large volume of documents available within the repository which makes it very time consuming to peruse thoroughly. In other words, engineers prefers to see *summary reports* of documents archived. For example,

when searching for engine part failures, engineers want to see how many times the failure of a particular part leads to engine removal, but not the list of original maintenance documents. Another example of useful knowledge that can be extracted from engineering documents is an “expertise finder” [7]. The expertise knowledge can be obtained by integrating author information and the contents of documents. Thus, for our new engineering document repository, we aim to include the ability to provide analysis of information stored, in addition to simple document search.

Our work in [30] can be seen as a digital library, with the extension where information presented is adapted to the role of the user. Digital libraries concentrate on the problem of searching for documents distributed over multiple repositories. For example, Priebe and Pernul [24] developed a portal over multiple document repositories by using an integrated metadata store. As a result, users can search on both the content of the documents and their metadata. In contrast, document index functionality does not form part of our proposed infrastructure. However, global document indexes can be provided to our knowledge repository as *services* that implement document indexes and metadata indexes.

Another area digital libraries concentrate on is dynamic links generation to relevant document resources [2, 23, 19]. Dynamic links are injected into documents automatically during presentation time, and does not alter the original documents. These dynamic links can point to related documents, or even services such as searching annotation and peer reviews [23]. In comparison, our proposed system does not perform dynamic link injection on existing documents. However, it can provide a list of suggested documents as a *service*.

Carr et al extended the concept of digital libraries to include the dissemination process that leads to publications [3]. They have developed the Digital Review Journal (DRJ), which is a knowledge base where users from different roles can collaborate through the entire process on writing an academic paper. It includes support for experiment creation, execution to publication. In comparison, our work concentrates on extracting existing knowledge generated from outside the framework. Also, the DRJ does not handle integrating knowledge from heterogeneous document sources.

Another related area of work is design repositories. Szykman et. al. [28] defined a design repository as “an intelligent, knowledge-based artifact modeling system used to facilitate the representation, capture, sharing, and reuse of corporate design knowledge”. In [1], Bohm et. al. created a hierarchical representation to describe functional aspects of designs. This representation is stored in a SQL database for archival. A web application was written to provide entry to the knowledge base. In [18], Ma and Mili proposed a object-oriented representation of product designs with constraints. Semantic relationships among design knowledge is also covered. In comparison, the knowledge repository we proposed is not limited to design knowledge archival and reuse only. One of the explicit aim of the project is to transfer knowledge gained during the entire lifecycle of a product, including manufacturing and maintenance, into the design process. In our infrastructure, we do not restrict knowledge archived into a pre-defined model. Instead, knowledge is modeled via the use of ontologies. New ontologies can be easily added to the system when new concepts from design or maintenance need to be added to the repository. Sophisticated reasoning and queries are done via an open standard query language.

Finally, there are also projects working on creating new semantically marked up data for large knowledge repositories, such as [4] and [6]. Creating new documents is outside the scope of our project, as we concentrate on the problem of delivering knowledge from existing documents. However, employing techniques to generate semantic information automatically for new documents to be deposited into the knowledge repository will improve document analysis and searching inside our framework.

5 Architecture

Two key technologies are used in integrating the distributed and heterogeneous data sources available from the document repository – Service-Oriented Architecture (SOA) and the Semantic Web. SOA is a software architectural concept that defines the use of services to support the

requirements of software users. In a SOA environment, functional components expose service behaviours accessible to other applications via loosely coupled standards-based interfaces. These components, termed Web Services, interoperate based on a formal definition independent of the underlying platform and programming language. Due to the nature of loose coupling in a SOA, applications can be developed and deployed incrementally. In addition, new features can be easily added after the system is deployed. This modularity and extensibility make SOA especially suitable as a platform for an integrated knowledge repository within large engineering organizations.

The Semantic Web is an application of the World Wide Web aimed at computational agents, so that *programs*, and not just humans, can interpret the meaning of documents on the Web (or an intranet). This allows the Web to be used for more than a human-browseable repository of information. The basis of this interpretation is an ontology, a structure which forms the backbone of the knowledge interpretation for an application.

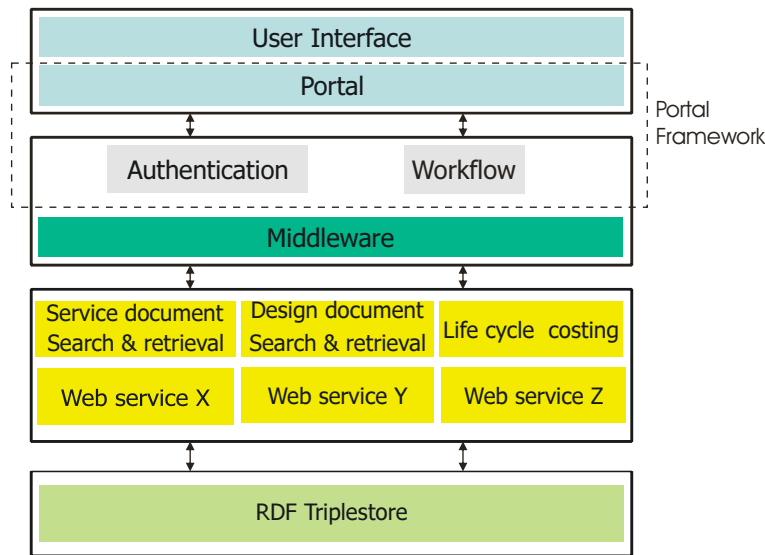


Figure 4: A conceptual view of the prototype’s architecture, showing three typical web services, two for search and retrieval, the third for interfacing to the life cycle cost modeler.

Our intelligent document repository is organized as a SOA, as shown in Figure 4. The user forms a query via a user interface. In the current system design, this user interface is in the form of web pages provided by a web portal. To answer the user’s query, the web portal finds and calls the web services that can provide it with relevant knowledge. These web services provide algorithms and functionalities that work on the underlying information set. There are no restrictions on the type of functionalities that the web services can provide, and it can range from an index of available documents, to finding the frequency of maintenance events.

The underlying information set is expressed as RDF triples [20]. RDF is the language for representing information about resources in the Semantic Web. RDF is used to represent both the data itself, and its associated metadata. An aerospace engineering ontology defined in OWL is used to describe the objects and concepts that appear within this information set. OWL [21] is the standard web ontology language for use between web agents in the Semantic Web. The development of this ontology was greatly helped by the closely controlled language used in the aerospace industry. For example, irrespective of a part’s description or name, it can be assigned a unique description based on its location within an engine using the ATA (Air Transport Association) descriptors [9].

Existing web standards are used wherever possible, to maximize tool reuse, compatibility and portability. A web browser access the web portal using the HTTP protocol. In turn, the web portal supplies the user interface to the web browser. The user interface is formed by a series of portlets.

Portlets are reusable components that display relevant information to portal users. In general, they appear to end users as rectangular sections of a web page offering a limited set of information. In our implementation, we use Liferay (<http://www.liferay.com>), an open source portal that conforms to the Java Portlet API standard, JSR 168². To create the information for display, the portlets access one or more web services. We envisaged that these web services are provided by different departments within the company, and can be distributed across multiple sites. These web services will perform more than document indexing that are already available in all document repositories. For example, it can provide an interface to existing manufacturing systems such as PDM (Product Data Management), or allow access to analysis tools such as Life-Cycle Cost Models and Finite Element Analysis. The web service interfaces are defined in WSDL [5], which is the standard interface language for defining web service interfaces. Specifically, we used the open source Apache Axis Java (<http://ws.apache.org/axis/java/index.html>) for implementing the web services. Documents in the repository will be in the form of RDF triples. Both the RDF triples and associated OWL ontologies are stored in 3store [12], an efficient RDF triplestore implementation that provides query and access abilities over HTTP. 3store accepts queries can be made in both RDQL [26] and SPARQL [25]. It should be noted that the architecture does not require all web services to access the underlying information sources via the triplestore. However, most, if not all, web services will operate on the semantically enabled data within the triplestore instead of the original heterogeneous documents.

The following sections discuss each component within this architecture in greater details.

5.1 Documents and Ontology

For each maintenance event, maintenance engineers document information surrounding the event for later references. A maintenance event occurs whenever actions are performed on an engine. Usually, each maintenance event involves multiple actions. Information documented includes the circumstances of the event, actions taken, parts installed and removed and any other findings they observed. The engineers who created these documents are located in numerous sites around the world, in the manufacture's or third party repair bases or at airports. As a result, the maintenance documents are in a large variety of formats. Some of these reports are in the form of unstructured Word documents, with different sites using different templates. Other engineers record maintenance information using a centralized Service Data Manager (SDM).

To enable machines to interpret meanings stored within the documents, we created an ontology that captures all the terms and concepts used. Moreover, since the document repository is to be used by both design and service engineers, the ontology captures concepts from engineers working in both areas. The ontology is created by analyzing existing documents and conducting knowledge acquisition interviews with engineers [30]. The result of these interviews enabled us to identify, by specialism, the main concepts and the associated keyword for these concepts used by the particular type of engineer when searching for information.

The resulting ontology contains concepts ranging from engine deterioration mechanisms, engine models and parts to airport locations. Figure 5 shows a UML class diagram for the concepts associated with maintenance events in the application ontology. In the diagram, the square boxes are concepts, or *classes*, in the ontology. Underneath the name of the class is a list of attributes for the class. In the ontology, a UML attribute is modeled as an OWL property which has the specified class as its domain. The lines connecting two classes in the diagram are associations. Associations represent relationships between classes. The arrow on the line shows the direction of the association. A UML association is modeled with an OWL property which has the linked classes as its domain and range. In addition, the *Part* class contains a taxonomy of aero-engine parts, and the *Engine* class includes a list of exiting engine types. In total, the ontology has 896 classes and 133 properties. Most of the classes represent parts within the engine taxonomy.

Using this ontology, we have populated a triplestore with maintenance records extracted from the company's SDM database, together with a small subset of service reports provided as Word

²For further details see <http://jcp.org/jsr/detail/168.jsp>

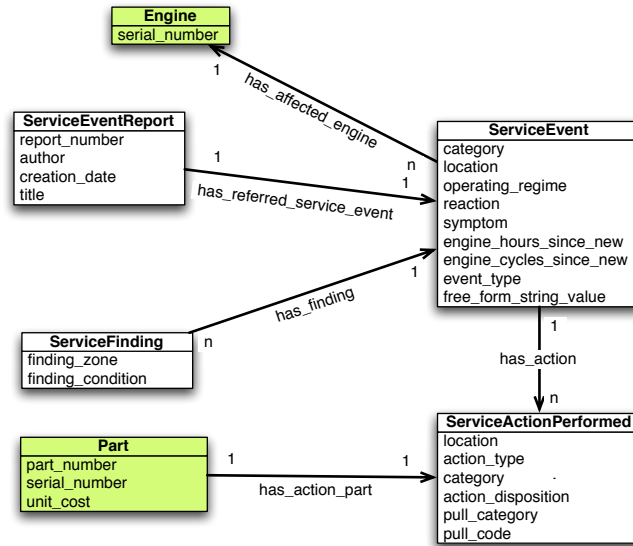


Figure 5: UML class diagram for the ontology used to describe maintenance events. This is a simplified view and does not show all properties and classes defined in the ontology.

documents. Population from the SDM database is done using a script that maps the SDM database schema to the ontology. Population from Word documents is currently undertaken manually, and is performed to test the validity of the ontology mapping. Automated extraction and population from Word documents is being investigated by one of our project partners that specializes in natural language processing. Currently, the triplestore contains approximately 3300 maintenance events, with around 31,000 actions. The populated triplestore contains 316 000 ground triples and 213 000 inferred triples. The current disk space usage is 12MB in data and 25MB in index.

5.2 Web Services

Web services provide processing functionalities that can be access from anywhere on the network. The web services are decoupled from the portlets. In other words, multiple portlets can execute the same web service, and a single portlet can execute multiple web services. The interface of a web service is defined by its WSDL document. This document is in XML format, and lists the operations a web service supports. The document also defines the syntactic types of inputs and outputs for each operation.

We have developed several web services that perform common requests on the maintenance documents. For example, obtaining a list of parts that are involved in the highest number of unscheduled maintenance events, tracing the maintenance record of an engine or a part, or retrieving details of a maintenance event. The web services construct SPARQL queries [26] according to their inputs. They then obtain results from the triplestore using the constructed query. Thus, the web services allow SPARQL queries to be reused across applications. Furthermore, they provide an abstraction over the data stored in the triplestore so developers who are unfamiliar with the ontology will still be able to perform queries. On the other hand, the use of an intermediary layer restricts the type of queries that can be performed. Therefore, for maximum flexibility, developers can access the triplestore directly by constructing their own SPARQL query.

5.3 Web Portal

Users access our knowledge framework via the web portal. The portal uses username/password authentication, and role-based access control. Using role-based access control, we can customize the content of the portal according to the engineer's specialization. The roles defined in the system can reflect job functions of the engineers. Thus, engineers with different specialization can be served a different set of portal pages. For example, when studying deterioration mechanisms, support and operations engineers can be presented with information from individual engine parts, while design engineers will be presented with information across the entire fleet. These different requirements arises from the different tasks the engineers have to perform. Operations engineers are interested in locating and fixing problems of existing engines in service. On the other hand, design engineers are interested in overall performance to aid the design of future variants. Another possible customization is the use of navigation/drill down over the engine taxonomy. Since an engine contains a very large number of parts, a simplified taxonomy for navigation will greatly speed up information browsing. The navigation taxonomy can be a subset of the entire engine taxonomy, based on the engineer's specialization.

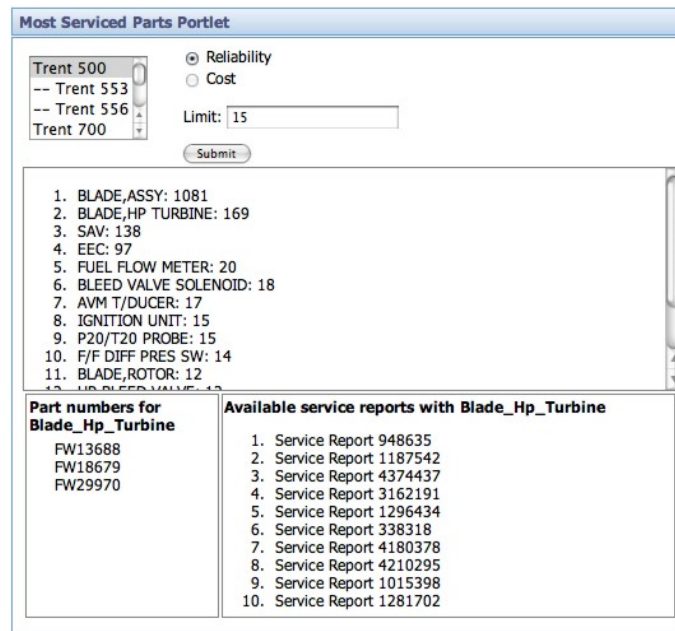


Figure 6: Screenshot from the current portal implementation.

A questionnaire based study was carried out to better understand what kind of knowledge engineers would like to discover from the document repository [16]. In the questionnaire, engineers were presented with a list of questions relating to maintenance experience with a product. They are asked how often they might ask them when designing a new product. They are also asked what other questions they might want to ask. The result of this questionnaire tells us what are the most important and most common life cycle information design engineers seek from maintenance documents. Typical questions that engineers ranked highly included:

1. What are the common deterioration mechanisms associated with this part?
2. What are the typical sequences of events which lead to this part's deterioration?
3. When this part fails, what other parts are typically affected as a result?
4. Which parts dominate the cost and reliability drivers in this engine?

```
SELECT ?a
WHERE (?a <rdf:type> <i:ServiceActionPerformed>)
      (?a <i:hasActionPart> ?c)
      (?c <i:hasPartNumber> "ET10935")
```

Figure 7: RDQL query for obtaining all maintenance actions that involve the part ET10935.

```
SELECT ?a
WHERE (?a <rdf:type> <i:ServiceActionPerformed>)
      (?a <i:hasActionPart> ?c)
      (?c <rdf:type> <i:AirDuctJoint>)
```

Figure 8: RDQL query for obtaining all maintenance actions that involve air duct joints.

5. What is the time between inspections/overhauls/repairs for this part?

Based on the results of the user study, we developed portal applications that help engineers answer these questions using information from the document repository. These portal applications are called portlets. A portal page usually consists of a series of portlets. Each portlet within the user interface performs a single presentation task. Figure 6 shows a screenshot of one of the portlets we developed. This portlet shows engineers the list of parts that are involved in the highest number of service actions in the selected engine. Users can obtain further information about any of the parts in the list. For example, they can see what part numbers make up a functional part. They can also view service reports that mention a part. To generate the list of most serviced parts, the web service has to consult two different information sources. The engine taxonomy and part names of a particular engine model are obtained from design documentation. Maintenance actions carried out through the engine's life are obtained from maintenance documentation. Within the maintenance documents, parts are only referred to using their part numbers, for example ET10935. Also, parts with different part numbers can be identical in functionality. Therefore, to generate results meaningful to the engineers, the web service must combine information from both design and maintenance. By correctly instantiating the parts involved in the *Part* hierarchy in the ontology, we can group functionally identical parts together within an RDQL query. As a result, by using the ontology, the web service can return a list of most serviced parts based on functionality groupings, and not simply distinct part numbers. This is illustrated in the query statements shown in Figure 7 and Figure 8.

6 Discussion

In this paper, we presented an intelligent document repository for an aero-engine manufacturer. However, the system architecture is generic and can be applied to organizations outside of aerospace engineering. Most large organizations generate a substantial amount of documents everyday. This is regardless of the sector the organization operates in, be it healthcare or finance. Our proposed architecture are applicable to document repositories in any industry, where the number of documents are too large to peruse, and that users are interested in knowledge extracted from documents deposited.

Heterogeneous document sources are integrated by a shared vocabulary – an ontology. To answer some of the questions from design engineers listed in Section 5.3, we need to combine information from multiple sources. For example, in Question 4, the engineer wants to know which parts dominate the cost drivers in the engine. To answer this, we combined reliability data from maintenance event reports with the engine taxonomy and costing information. These data are contained in different databases within the company. Integration is made possible with the use of the ontology, which allows software agents to reason over the different resources.

Users access the document repository via a standard web browser. System components are hosted on distributed servers. As a result, the software can be deployed and updated centrally, without changing the configurations of thousands of desktop computers. Also, users can access the document repository without special software installed on their computers.

The prototype has been demonstrated to Rolls-Royce design engineers and while it has limited functionality the approach taken has received positive reviews. The design community recognizes the need for this approach as we are able to summarize large document collections (both in size and type). The customisation of the portal is highly relevant to the both the user and the company, as it bring clear benefits. The prototype is in the process of being extended to incorporate further documents collects, allowing more summarization and questions to be addressed.

7 Conclusions

In this paper, we discuss the development of a knowledge-based document repository. The system allows users from across a company to access knowledge that they need to undertake their activities. In particular, we have developed a prototype for a large aero-engine manufacturer. A fundamental shift is occurring in the aerospace industry away from selling products to providing services. This shift in market focus means that new engines must be designed to provide lower and more predictable life cycle costs. To achieve this, engineers must obtain knowledge gained from the entire life of an engine.

However, design engineers often do not have access to front-line maintenance data due to the size and distributed nature of aerospace manufacturers. Moreover, an engine model remains in service for a long period of time and each engine in service generates a large number of maintenance documents during its life. As a result, it is impossible for design engineers to examine all information available thoroughly. Therefore, unlike existing document repository or digital libraries, our system is knowledge, and not information, based. It is based on services, and not document indexes or document linking. Services can search and analyze relevant maintenance records and design guidelines, and provides engineers with a ‘summary report’ to such information. It is hoped that this summary view provided will help engineers in creating design documents that incorporate aftermarket issues into the design requirements.

Two key technologies are used in integrating the distributed and heterogeneous data sources – the Semantic Web and Service-Oriented Architecture. The Semantic Web provides the framework allowing computer programs to interpret and reason over the heterogeneous document sources. The documents are integrated using an ontology, which captures the terms and concepts used in aerospace engineering. Service-Oriented Architecture allows knowledge extraction and analysis functionalities to be added to the system as modules called web services.

To test our proposed architecture, we have developed a working prototype. The prototype consists of a web portal, several web services and a RDF triplestore. Users access the document repository via a web portal, which is customized according to users’ roles. The web services perform processing functionalities for queries we found engineers wanted to seek most from maintenance documents. The RDF triplestore contains real maintenance documents and the OWL aerospace ontology.

8 Acknowledgments

This research was undertaken as part of the IPAS project (DTI Grant TP/2/IC/6/I/10292). The authors would also like to thank the project partners for providing us with data and ontologies. Specifically, we would like to thank Derek Sleeman and David Fowler from Aberdeen AKT for the ontology, and Alymer Johnson and Santosh Jagtap from Cambridge EDC for the user requirement analysis.

References

- [1] M. R. Bohm, R. B. Stone, and S. Szykman. Enhancing virtual product representations for advanced design repository systems. *Journal of Computing and Information Science Engineering*, 5(4):360–372, 2005.
- [2] L. Carr, W. Hall, S. Bechhofer, and C. Goble. Conceptual linking: Ontology-based open hypermedia. In *Proceedings of 10th International World Wide Web Conference (WWW)*, pages 334–342, Hong Kong, 2001.
- [3] L. Carr, T. Miles-Board, G. Wills, and S. Grange. Extending the role of digital library: Computer support for creating articles. In *Proceedings of The 15th ACM Conference on Hypertext and Hypermedia*, pages 12–21, Santa Cruz, USA, 2004.
- [4] L. Carr, T. Miles-Board, A. Woukeu, G. Wills, and W. Hall. The case for explicit knowledge in documents. In *Proceedings of ACM Symposium on Document Engineering*, pages 90–98, Milwaukee, Wisconsin, 2004.
- [5] E. Christensen, F. Curbera, G. Meredith, and S. Weerawarana. Web services description language (WSDL) 1.1. Technical report, W3C Note, <http://www.w3.org/TR/wsdl>, Mar. 2001.
- [6] F. Ciravegna and Y. Wilks. Designing adaptive information extraction for the semantic web in amilcare. In S. Handschuh and S. Staab, editors, *Annotation for the Semantic Web*. IOS Press, 2003.
- [7] A. L. Cohen, P. P. Maglio, and R. Baret. The expertise browser: How to leverage distributed organizational knowledge. In *Workshop on collaborative and cooperative information seeking in digital information environment at the International Conference on Computer Supported Cooperative Working*, Seattle, WA, 1998.
- [8] R. Crowder, R. Bracewell, G. Hughes, M. Kerr, D. Knott, M. Moss, C. Clegg, W. Hall, K. Wallace, and P. Waterson. A future vision for the engineering design environment: A future sociotechnical scenario. In A. Folkesson, K. Gralen, M. Norell, and U. Sellgren, editors, *Proceedings of 14th International Conference on Engineering Design*, pages 249–250, Stockholm, 2003.
- [9] D. Fischer. *A theory of presentation and its implications for the design of online technical documentation*. PhD thesis, Coventry University, 1997.
- [10] Y. Gao, I. Zeid, and T. Bardasz. Characteristics of an effective design plan system to support reuse in case-based mechanical design. *Knowledge-Based Systems Knowledge-Based Systems Knowledge-Based Systems*, 10(6):337–350, Apr. 1998.
- [11] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
- [12] S. Harris and N. Gibbins. 3store: Efficient bulk RDF storage. In *Proceedings of 1st International Workshop on Practical and Scalable Semantic Systems (PSSS'03)*, pages 1–15, Sanibel Island, Florida, 2003.
- [13] A. Harrison. Design for service – harmonising product design with a services strategy. In *Proceedings of GT2006, ASME Turbo Expo 2006: Power for Land, Sea and Air*, Barcelona, Spain, May 2006.
- [14] I. Heath, G. Wills, R. Crowder, , W. Hall, and J. Ballantyne. Towards a new authoring methodology for large-scale hypermedia applications. *Multimedia Tools and Applications*, 12(2-3):129–144, Nov. 2000.

- [15] T. Jackson, J. Austin, M. Jessop, B. Linag, A. Pasley, M. Ong, G. Allan, V. Kadiramanathan, H. Thompson, and P. Fleming. Distributed health monitoring for aero-engines on the grid: DAME. In *Proceedings of IEEE Aerospace*, Montana, March 2005.
- [16] S. Jagtap, A. Johnson, M. Aurisicchio, and K. Wallace. Pilot empirical study: Interviews with product designers and service engineers. Technical Report 140 CUED/C-EDC/TR140-March 2006, Engineering Design Centre, University of Cambridge, 2006.
- [17] D. V. Khadilkar and L. A. Stauffer. An experimental evaluation of design information reuse during conceptual design. *Journal of Engineering Design*, 7(4):331–339, 1996.
- [18] Z. M. Ma and F. Mili. Knowledge comparison in design repositories. *Engineering Applications of Artificial Intelligence*, 16:203–211, 2003.
- [19] K. Malcolm, S. Poltrock, and D. Schuler. Industrial strength hypermedia: Requirements for a large engineering enterprise. In *Proceedings Hypertext '91*, pages 13–24, Seattle, 1991.
- [20] F. Manola and E. Miller. RDF primer. Technical report, W3C Recommendation, <http://www.w3.org/TR/rdf-primer>, Feb. 2004.
- [21] D. L. McGuinness and F. van Harmelen. OWL web ontology language overview. Technical report, W3C Recommendation, <http://www.w3.org/TR/owl-features>, Feb. 2004.
- [22] D. E. Millard, F. Tao, K. Doody, A. Woukeu, and H. C. Davis. The knowledge life cycle for e-learning. *International Journal of Continuing Engineering Education and Lifelong Learning*, 16:110–121, 2006.
- [23] N. Nnadi and M. Bieber. Lightweight integration of documents and services. In *DocEng '04: Proceedings of the 2004 ACM symposium on Document engineering*, pages 51–53, New York, NY, USA, 2004. ACM Press.
- [24] T. Priebe and G. Pernul. Towards integrative enterprise knowledge portals. In *CIKM '03: Proceedings of the twelfth international conference on Information and knowledge management*, pages 216–223, New York, NY, USA, 2003. ACM Press.
- [25] E. Prud'hommeaux and A. Seaborne. SPARQL query language for RDF. Technical report, W3C Working Draft, <http://www.w3.org/TR/rdf-sparql-query>, Oct. 2006.
- [26] A. Seaborne. RDQL - a query language for RDF. Technical report, W3C Member Submission, <http://www.w3.org/Submission/RDQL>, Jan. 2004.
- [27] N. Shadbolt and N. Milton. From knowledge engineering to knowledge management. *British Journal of Management*, 10(4):309–322, Dec. 1999.
- [28] S. Szykman, R. D. Sriram, C. Bochenek, J. W. Racz, and S. Jocelyn. Design repositories: Engineering design's new knowledge base. *IEEE Intelligent Systems*, 15(3):48–55, 2000.
- [29] K. M. Wallace, C. Clegg, and A. Keane. Visions for engineering design: a multi-disciplinary perspective. In *Proceedings of 13th International Conference on Engineering Design*, pages 107–114, Glasgow, Scotland, 2001.
- [30] G. Wills, D. Fowler, D. Sleeman, R. Crowder, S. Kampa, L. Carr, and D. Knott. Issues in moving to a semantic web for a large corporation. In *Proceedings of 5th International Conference on Practical Aspects of Knowledge Management (PAKM)*, volume 3336 of *Lecture Notes in Artificial Intelligence*, pages 378–388. Springer, 2004.