Published at the 2004 Int'l Conference on Knowledge-Based Intelligent Information & Engineering Systems (KES2004), New Zealand.

Decision Support System on The Grid

M. Ong, X. Ren, G. Allan, V. Kadirkamanathan, HA Thompson and PJ Fleming

Rolls-Royce Supported University Technology Centre in Control and Systems Engineering,
University of Sheffield,
Department of Automatic Control and Systems Engineering,
Mappin Street, Sheffield, S1 3JD, United Kingdom.

M.Ong@sheffield.ac.uk

Abstract. Aero engines are extremely reliable machines and operational failures are rare. However, currently great effort is being put into reducing the number of in-flight engine shutdowns, aborted take-offs and flight delays through the use of advanced engine health monitoring technology. This represents a benefit to society through reduced delays, reduced anxiety and reduced cost of ownership of aircraft. This is reflected in a change of emphasis within aero engine companies where, instead of selling engines to customers, there is a fundamental shift to adoption of power-by-the-hour contracts. In these contracts, airlines make fixed regular payments based on the hours flown and the engine manufacturer retains responsibility for maintaining the engine. To support this new approach, improvements in in-flight monitoring of engines are being introduced with the collection of much more detailed data on the operation of the engine. At the same time advances have been made in Internet technologies providing a worldwide network of computers that can be used to access and process that data. The explosion of available knowledge within those large datasets also presents its own problems and here it is necessary to work on advanced decision support systems to identify the useful information in the data and provide knowledge-based advice between individual aircraft, airline repair and overhaul bases, world-wide data warehouses and the engine manufacturer. This paper presents a practical framework in which to build such a system that is inherent in the emerging Grid computing paradigm that provides the necessary computing resources. A demonstrator system already developed and implemented in the UK E-Science Grid project, DAME, is introduced.

1 Introduction

Currently new aero engines [1] are being instrumented with engine monitoring units possessing significantly greater capability to record and analyse data. Each engine on a civil airliner could generate about 1Gbyte of data per flight. Rolls-Royce currently has over 50,000 engines in service with total operations of around 10M flying hours per month. As a result in future one can envisage many 100s of Gbytes of data being transmitted every day which will require analysis. In an analogy to the Internet the on-board processing units can be seen as single computers operating on top of the global network, which go on and off line as the aircraft lands and takes off. It is envisaged that such systems will become commonplace in the future with applications in other areas such as the automotive and marine industry.

The key objectives are to reduce delays and cost of ownership for the aircraft. The challenge is to provide the infrastructure to manage the large amounts of data, perform compute-intensive analysis and modelling to identify faults that have occurred but more importantly to perform knowledge-based maintenance advisory to identify the potential faults that require maintenance to prevent failures and aircraft downtime. It is this second feature of predictive maintenance that provides huge potential paybacks in terms of future systems giving much greater aircraft availability. The underlying research challenges are thus real-time knowledge-based decision support, intelligent feature extraction and intelligent data mining.

It is clear that in building a practical system, the expertise and software tools that support analysis need to be distributed around the world (Fig. 1). A user requires the ability to tune and choose the tools used to determine a diagnosis rather than being provided with a monolithic fixed system. This is particularly useful when there is any uncertainty regarding a diagnosis or prognosis. A key concern is data security and so the tools must also provide secure data encryption. These on-line analysis tools need to be able to look for and identify clusters of anomalies or novelties as they appear and support the appropriate diagnoses. The need for identification of clusters of faults arises from the desire to provide some insight into an underlying cause. For example, some operators may experience certain faults because of the way the engines are being operated, etc. Ideas regarding engine diagnoses can be tried out in a virtual environment using modelling techniques. This may be achieved through the use of engine performance models. Again these models must be provided as a Web Service. This may be particularly useful when a diagnosis or prognosis is unknown. For instance, if an abnormality is identified which has not been encountered before (novelty analysis), the system can consult previous history data stores for other engines for similar abnormalities. The knowledge gained from these previous instances stored along with this data may allow identification and rectification. In addition, the data stores (for all engines) can be analysed off line for fleet management.

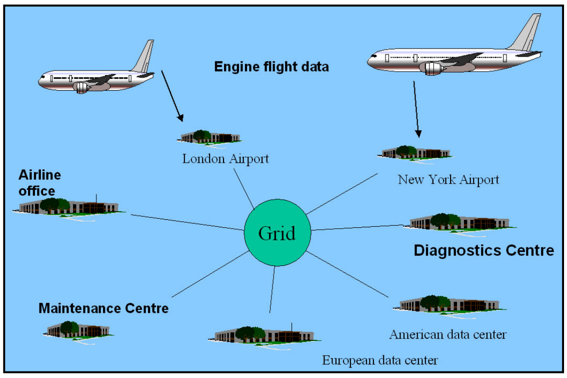

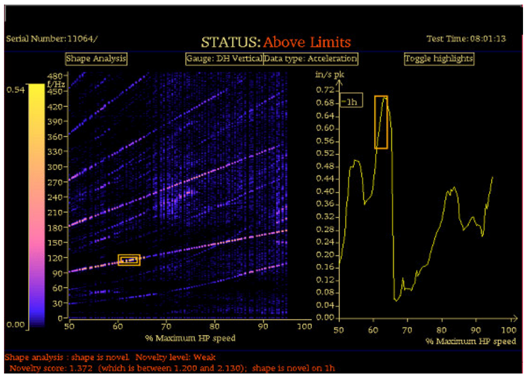

Fig. 1. (Left) Distribution of data in a virtual engine maintenance environment
Fig. 2. (Right) Representative sample of data downloadable from the engine on-wing monitoring system. Data includes engine performance parameters and vibration data throughout the flight as well as various other operating parameters.

In achieving the above-mentioned objectives, the concept of Grid computing is introduced. The Grid offers significant capability for the design and operation of complex decision support systems. Grid computing has the potential to mediate the task of diagnosis and prognosis, within complex and dynamic operational business and engineering processes. The Grid has the capability to provide high performance computing resource on demand. This is considered a given, and will offer resource for compute intensive tasks within the decision support process, such as data analysis and modelling.

2 Grid Computing

The Grid computing concept [2,3], first developed in the scientific community, was initially aimed to address the problems of sharing and working with large datasets. Grid computing is now moving towards a mainstream challenge of creating reliable, robust, scaleable and secure distributed systems. The Grid [4] is an aggregation of geographically dispersed computing, storage and network resources, co-ordinated to deliver improved performance, higher quality of service, better utilisation and easier access to data. It enables collaboration across “virtual organisations”, enabling the sharing of applications and data in an open, heterogeneous environment. Services previously considered to be host-centric can now be distributed throughout a network of powerful computers improving quality of service while also offering enhanced and improved capabilities. The emergence of Grid software such as the Globus Toolkit [5,6] provides the necessary middleware to implement a Grid system and includes services that tackle issues such as accessibility, security and resource management.

3 Distributed Aircraft Maintenance Environment (DAME)

The Distributed Aircraft Maintenance Environment (DAME) project is a pilot project supported under the United Kingdom e-Science research programme in Grid technologies [7]. Industrial partners in the DAME project are Rolls-Royce plc, who have provided the aero engine data for the diagnostic system (Fig. 2), Data Systems and Solutions, who are Rolls-Royce data systems providers and currently deliver commercial aero engine health monitoring services, and Cybula Ltd, who provide the high-speed pattern matching technology developed at York University. The university partners collaborating in the project are Sheffield, York, Leeds and Oxford.

DAME is particularly focussed on the notion of proof of concept, using the Globus tool kits and other emerging Grid service technologies to develop a demonstration system. This is known as the DAME Diagnostic/Prognostic Workbench. The demonstrator system tackles complex issues such as security and management of distributed and non-homogenous data repositories within a diagnostic analysis framework with distributed users and computing resources.

The Rolls-Royce supported University Technology Centre (UTC) in the Department of Automatic Control and Systems Engineering at the University of Sheffield is currently engaged within the DAME project contributing expertise in Decision Support, Modelling, Simulation and Workflow Optimisation. Grid computing expertise being gained on the DAME project is also influencing other work within the Sheffield UTC enabling existing in-house tools, models and services to be placed within a potentially new Grid-enabled framework.

4 Decision Support in DAME

It is clear that in order to deal with the explosion in data available from complex engine health monitoring systems, it is necessary to design advanced decision support systems. These need to be able to identify faults based on knowledge of previous fault conditions and also perform analysis across fleets of engines. The Sheffield UTC has been actively exploring a variety of on-wing control system diagnostic techniques and also portable maintenance aid demonstrators. This work has been based on two fundamental underpinning technologies, both of which will be described in more detail.

- Case-Based Reasoning
- Integration of Model Based Fault Detection and Isolation Approaches

4.1 Case-Based Reasoning

Case-Based Reasoning (CBR) is a knowledge-based, problem-solving paradigm that resolves new problems by adapting the solutions used to solve problems of a similar nature in the past [8,9]. A further advantage of this approach is that it allows consolidation of rule knowledge and provides a reasoning engine that is capable of probabilistic-based matching. With CBR technology, development has taken place in an incremental fashion facilitating rapid prototyping of an initial system. The development of robust strategies for integration of multiple health information sources is achieved using reasoning algorithms of progressively increasing complexity.

In contrast to conventional search engines, CBR systems contain a knowledge model of the application domain in which it operates on. It is therefore not universal but specifically designed for the domain. Hence, it is possible to develop intelligent search abilities, which even show reasonable results when given fuzzy or incomplete requests. Moreover, the results are ranked and complemented by variants and alternatives, thus, not only matches are given but information is valued with "more suitable" or "less suitable".

Previous research work carried out within the Sheffield UTC include investigating the use of Case-Based Reasoning techniques for a portable PC-based Flightline Maintenance Advisor [10,11] to correlate and integrate fault indicators from the engine monitoring systems, BITE[*] reports, maintenance data and dialog with maintenance personnel to allow troubleshooting of faults (Fig. 3). The outcomes of the initiative included the implementation of a portable Flightline Maintenance Advisor that was trialed with Singapore Airlines (Fig. 4).

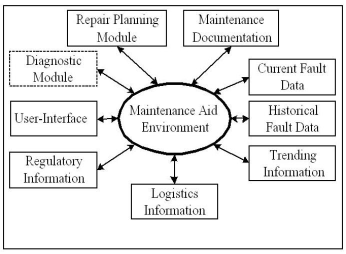

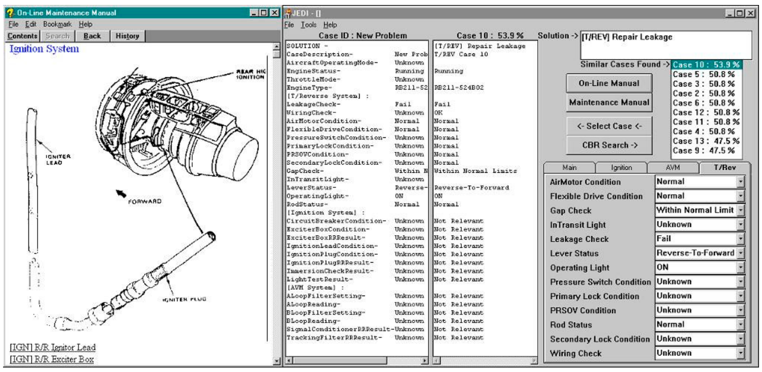

Fig. 3. (Left) Structure of Flightline Maintenance Advisor
Fig. 4. (Right) View of the Portable Maintenance Aid application in operation

[*] BITE – Built-In Test Equipment

Today, rather than using a portable computer which needs updating with new data as it becomes available, it is highly desirable for a CBR system to be accessed remotely by engineers over a computer network. The advantage of this is that it is easier to support and also allows search of an extensive casebase of historical maintenance incidents across an entire fleet of engines [12,13]. This allows identification of the most appropriate course of action to diagnose and rectify an engine problem with a prescribed set of fault symptoms.

Essential to the CBR system is the casebase that represents a knowledge repository containing detailed descriptions of engine faults and the best practice maintenance advice (solutions) gathered from engineers and experienced mechanics over the development and service life of the engine. For a new engine type, little information is known initially but the advantage of CBR techniques is that a casebase of independent records of fault and maintenance information can be developed in a piecemeal manner and updated as and when knowledge about the behaviour of the system is known. More importantly, the siting of the CBR system within a virtual maintenance facility also allows the integration of diagnostic knowledge from multiple health information sources which can improve the accuracy and coverage of the CBR system. Useful diagnostic information previously available from separate monitoring systems, when brought together into a single system, provides for a more powerful diagnostic tool.

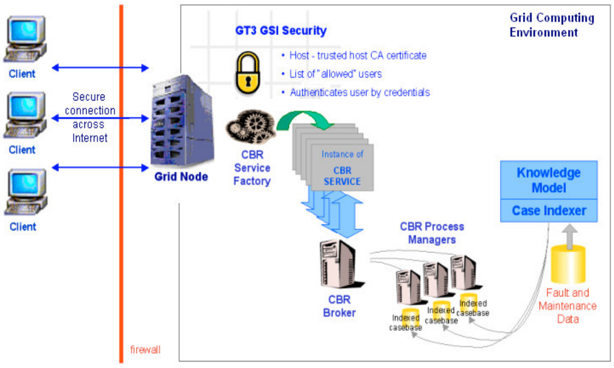

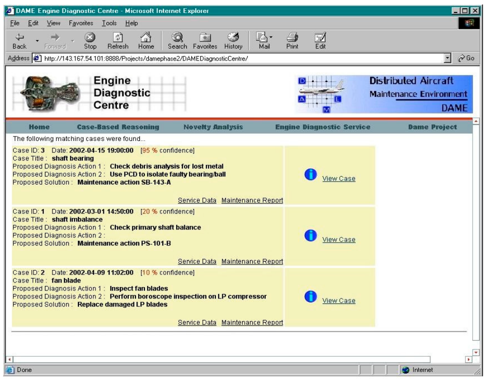

Fig. 5. (Left) The CBR service is available across the Internet, enabling maintenance personnel to access detailed fault information, maintenance advice and computing resources to support the problem-solving process. Grid service "factories" allow multiple instances of the CBR service to run in parallel on the Grid, each instance supporting an individual client request

Fig. 6. (Right) Web browser window displaying a generated list of cases that match a query for a particular engine fault. For each brief case listed, the detailed fault information and maintenance advice can be obtained by retrieving the full case details

In support of this, a CBR decision-support application has been developed and implemented at the Sheffield UTC as a web service within the Grid computing environment. Maintenance personnel can access this via a secure connection to the service using a web browser on any computer connected to the Internet (Fig. 5). Queries for matching cases can be submitted to the CBR service in two ways; directly via a Web browser window (Fig. 6) or automatically via an integrated client such as the one in the DAME automatic workflow system.

The CBR service provides maintenance personnel at various levels with access to stores of accumulated diagnostic knowledge and maintenance data as well as large computing resources to support the fault analysis and the decision-making process. This is a particularly important feature because it gives aero engine experts (considered as a high-value resource) the mobility to operate on large data and complex problems from a remote location. The system is implemented using available open standards such as Java, XML, and web service technologies. From the experience gained thus far, it is evident that the CBR decision support application displays much potential for integration of wireless-device access, thus enabling the use of mobile and wireless technologies in the near future.

In the future, the CBR system can be further escalated to accommodate a dynamic learning process. The CBR system may learn how an expert would typically troubleshoot a problem in a given situation. Using knowledge accumulated from various information sources, the system could in the future execute a set of diagnostic tests and present the data to the expert automatically thus saving valuable time and effort. At present, work is also in progress to investigate and develop a knowledge-based Workflow Advisory system within the DAME virtual maintenance environment.

4.2 Model-Based Fault Detection Isolation Approaches

To support the fleet management of engines, performance-analysis-based engine diagnostics is necessary. Here suites of modelling, estimation and analysis tools need to be integrated. There are several approaches to model-based fault detection and isolation (FDI). Requirements of modern fault diagnosis include promptness, accuracy and sensitivity to faults. It is commonly agreed that hybrid schemes would provide better solutions for future gas turbine diagnostic systems [14,15]. Additionally, it is important to consider how these approaches should be used in conjunction to provide the most accurate diagnosis in the decision-making process.

Model-based FDI can be used to track changes in engine parameters that may indicate impending faults. This predictive capability allows the fleet manager to schedule appropriate maintenance and minimise the downtime of an aircraft.

Advanced and accurate model-based FDI [16] may require intensive computing power for modelling and simulation. This processing need limits its application on large-scale complex systems. Thus there is a need for high-performance computing power to overcome these restrictions.

The Sheffield UTC has been actively exploring a variety of model-based techniques to identify engine faults and also performance degradation. Fundamental to model-based techniques is the provision of a detailed reference model that can be used for analysis. Considering the move from local diagnostics systems to remotely accessible systems, a major step has been taken for a future system at the Sheffield UTC through the development of a gas turbine engine performance model that can be run via a web service on the Grid (Fig. 7). With this service, a fleet maintenance engineer can perform engine performance simulations through a Web browser remotely without knowing details of the execution of the simulation. The simulation service itself is distributed among a set of high performance Grid computing nodes. In addition, the engine simulation web service is also made programmatically accessible through its public interface, enabling authorised users to further develop tools that may invoke this service within their own applications. Security of this interface is discussed in Section 5.

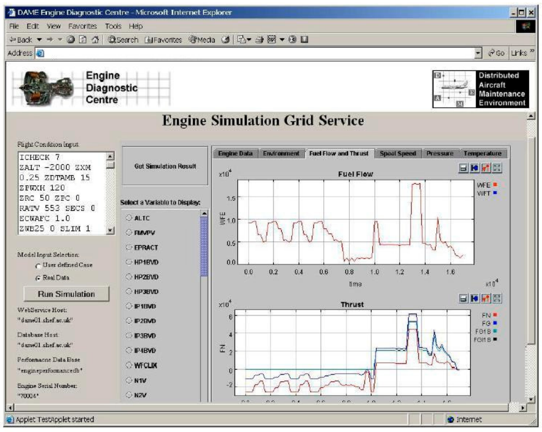

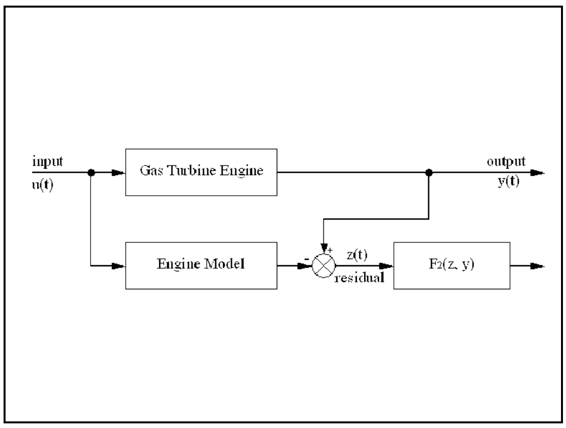

Fig. 7. (Left) Aero engine simulation service on the Grid is accessed via a Web browser connected to the Internet
Fig. 8. (Right) Simulation-based fault diagnosis

Figure 8 shows one basic usage of the engine simulation Grid service for fault diagnosis. When an accurate system performance simulation is available on the Grid, the experienced maintenance engineers can invoke this simulation against the real monitored process data. The system that is being analysed is compared against the simulation results and residuals are generated for the differences between the current state of the engine and the ideal model. These residuals then need to be intelligently analysed to form a decision about the current state of the engine. This can be used to track changes in engine parameters which may indicate impending faults.

A typical use case which encompasses both the engine simulation and CBR services in the fault analysis and maintenance process is described as follows. Data downloaded from an aircraft is first analysed for abnormalities. The existence of an abnormality and the possible cause can be checked against the engine simulation. If a novelty exists, then further information is extracted from the data and other available fault diagnostic services to form a query within the CBR services. The result returned to the maintenance personnel consists of previous similar fault cases, known solutions to the problem, as well as a confidence ranking for each case. The maintenance analyst and domain experts can further take advantage of the integrated fault diagnostic tools to confirm the fault diagnosis findings. For example, the domain experts can substantiate a proposed fault analysis by injecting the similar fault into an engine model and performing a simulation to check the degree of match.

4.3 Service-Oriented Architecture

Service-oriented architecture (SOA) is not a new concept. A SOA is essentially a collection of services, stressing on interoperability and location transparency. These services communicate with each other, and can be seen as unique tools performing different parts of a complex task. Communication can involve either simple data passing or it could involve two or more services co-ordinating some activity. Services and service-oriented architectures are about designing and building system using heterogenous network addressable software components. An important aspect of the service-oriented architecture is that it separates the service's implementation from its interface.

Web services, in the general meaning of the term, are services offered via the Web. In the DAME scenario, an application sends a request to a service at a given Internet address using the SOAP[1] protocol over HTTP[2]. The service receives the request, processes it, and returns a response. Based on the emerging standards such as XML[3], SOAP, UDDI[4], and WSDL[5], web services enable a distributed environment in which any number of applications, or application components, can inter-operate seamlessly among organisations in a platform-neutral, language-neutral fashion. Web service consumers view a service simply as an endpoint that supports a particular request format or contract. Web service consumers are not concerned with how the Web service goes about executing their requests; they expect only that it will. Consumers also expect that their interaction with the service will follow a contract, an agreed-upon interaction between two parties. The way the service executes tasks given to it by the service consumers is irrelevant, the only requirement is that the service sends the response back to the consumer in the agreed-upon format.

Grid services [17] on the other hand are based on the integration of Open Grid Services Architecture (OGSA) concepts and web service technologies. Specifically, Grid services benefit from both web service technologies as well as Grid functionality. As a prime example, Grid services can provide aero engine experts in any geographical location with remote access to powerful Grid computing resources, large knowledge repositories and datasets as well as diagnostic tools via a Web browser and Internet link.

5 Security

A Grid-enabled decision support system may contain potentially business-sensitive data and hence access to data and services should be restricted to authorised members within an organisation. For instance, both the engine faults knowledge base and engine models contain important information on its design characteristics and operating parameters. The use of the Grid Security Infrastructure (GSI) [18] enables secure authentication and communication over an open network. GSI consists of a number of security services including mutual authentication and single sign-on. This is based on public key encryption, X.509 certificates, and Secure Sockets Layer (SSL) communications. The implementation of GSI within the DAME decision support environment is composed of Globus Toolkit 3 (GT3) security elements conforming to the Generic Security Service API (GSS-API), which is a standard API for security systems promoted by the Internet Engineering Task Force (IETF).

At the core of the GT3 security infrastructure is client and host authorisation using X.509 identity certificates for both the service users and service hosts. Access to the decision support system and resources on the Grid require user authentication. Hence, all users and services need to have a certificate issued from a trusted Certificate Authority (CA). Because the CA is the heart of the security system, it is very important that Grid hosts and users only use their own trusted CA or utilise an established commercial CA. A CA's signing policy has to be placed in the Grid computing environment to allow that nodes to authenticate users holding valid certificates. On top of this, users must also have their user credentials listed on a Grid-Mapfile. A Grid-Mapfile is a local file used to store mappings between a user identity on the Grid to a local identity (an account name on the Grid computer being used). It is clear that users are only allowed to access the decision support services and

[1] SOAP - Simple Object Access Protocol
[2] HTTP - Hyper Text Transfer Protocol
[3] XML - eXtensible Markup Language
[4] UDDI - Universal Description, Discovery and Integration
[5] WSDL - Web Services Description Language

Grid resources on a Grid node if their verified credentials have been registered in the Grid environment by an administrator of that Grid.

6 Concluding Remarks

In this paper we have discussed the use of case-based reasoning techniques and model-based fault detection for the flightline maintenance and fault diagnosis of aero-engines. In particular, the paper has concentrated on particular aspects of this work highlighting how there is a move from local support for diagnostics to diagnostics from a centralised virtual environment operating across distributed resources. Although local diagnostics will always be performed on–engine, the adoption of a centralised service which maintenance personnel at various levels can access is highly advantageous as new technologies become available naturally with the explosion of the World Wide Web and Grid computing [19]. Aero-engine experts who are regarded as high-value resources can be mobilised to analyse complex problems globally from remote locations. The demonstration systems presented can be used to identify faults at flightline and also predict impending faults across fleets of engines through trending. The business benefits of this open, flexible, proactive approach to engine monitoring and maintenance are not only improved fault diagnosis performance, but also reusable service assemblies, better scalability, better maintainability, higher availability, reduction in unscheduled maintenance and resulting aircraft downtime.

Acknowledgements. The authors gratefully acknowledge financial support of the Engineering and Physical Research Council in the UK under Grant Number GR/R67668/01 and input from engineers at Rolls-Royce and Data Systems & Solutions.

References

1. Rolls-Royce. The Jet Engine. Rolls Royce, Derby, England, 1986.
2. I. Foster, C. Kesselman, and S. Tuecke. The Anatomy of the Grid : Enabling Scalable Virtual Organizations. *International J. Supercomputer Applications*, 15(3), 2001.
3. I. Foster, C. Kesselman, J. Nick, and S. Tuecke. The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration, *Open Grid Service Infrastructure WG, Global Grid Forum*, June 22, 2002.
4. I. Foster, C. Kesselman. The Grid: Blueprint for a New Computing Infrastructure. *Morgan Kauffman*, 2003.
5. I. Foster, C. Kesselman. Globus: A Metacomputing Infrastructure Toolkit. In *Intl J. Supercomputer Applications*, 11(2): 115-128, 1997.
6. I. Foster, C. Kesselman, C. Lee, R. Lindell, K. Nahrstedt, and A. Roy. A Distributed Resource Management Architecture that Supports Advance Reservations and Co-Allocation. *Intl Workshop on Quality of Service*, 1999.
7. Distributed Aircraft Maintenance Environment (DAME) project; www.cs.york.ac.uk/dame
8. J. Kolodner. Case-Based Reasoning. *Morgan Kauffman*, 1993.
9. D. B. Leake. Case-Based Reasoning: Experiences, Lessons & Future Directions. *The MIT Press*, 1996.
10. S. M. Hargrave. Evaluation of Trent 800 Portable Maintenance Aid Demonstrator. *Rolls-Royce University Technology Centre, University of Sheffield,* Report No. RRUTC/Shef/R/98202, 1998.
11. S. M. Hargrave. Review of Performance-Based Diagnostic Tool, *Rolls-Royce University Technology Centre, University of Sheffield,* Report No. RRUTC/Shef/TN/98204, 1998.
12. R. Magaldi. CBR for Troubleshooting Aircraft on the Flightline, In *Proceedings of IEE Colloquium on Case Based Reasoning - Prospects for Applications*, Digest No 1994/057, 1994.
13. W. Alammari. Using Case-Base Reasoning for Improving Fault Diagnosis in Saudi Airline, *MSc Dissertation, Dept of Automatic Control & Systems Engineering, University of Sheffield*, 1996.
14. Y. G. Li. Performance-Analysis-Based Gas Turbine Diagnostics: A Review. In *Proc. Instn. Mech. Engrs.*, Vol 216, Part A: J Power and Energy, 363-377, 2002.
15. R. J. Patton, C. J. Lopez-Toribio, and F. J. Uppal. Intelligence Approaches to Fault Diagnosis. *Applied Mathematics and Computer Science*, Vol 9 No. 3, 471–518, 1999.
16. P. Li, and V. Kadirkamanathan. Particle Filtering Based Likelihood Ratio Approach to Fault Diagnosis in Nonlinear Stochastic Systems. In *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 31 No3, August 2001.
17. I. Foster, C. Kesselman, J. Nick, and S. Tuecke. Grid Services for Distributed System Integration, *Computer*, 35(6), 2002.
18. V. Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski, J. Gawor, C. Kesselman, S. Meder, L. Pearlman, and S. Tuecke. Security for Grid Services. *Twelfth International Symposium on High Performance Distributed Computing (HPDC-12)*, IEEE Press, 2003.
19. M. Ong, X. Ren, G. Allan, H. A. Thompson, and P.J. Fleming. Future Trends in Aircraft Engine Monitoring. In *IEE Professional Network on Aerospace seminar on "Aircraft Airborne Condition Monitoring"*, 2003.