Modelling, Optimisation and Decision Support Using the Grid

A. Shenfield, M. Ong, G. M. Allan, X. Ren, V. Kadirkamanathan, H. A. Thompson and P. J. Fleming

Rolls-Royce University Technology Centre in Control and Systems Engineering
Department of Automatic Control and Systems Engineering
University of Sheffield
Sheffield, S1 3JD
UK

Email: {A.Shenfield, M.Ong, Jeff.Allan, Visakan, H.Thompson, P.Fleming}@sheffield.ac.uk

Abstract

Modelling, Optimisation and Decision Support tools are vital in most areas of engineering. As part of the Distributed Aircraft Maintenance Environment (DAME) e-Science project, a virtual “work-bench” has been developed that can aid the engineer in solving engineering design problems. While the problem that motivated the development of these tools was taken from the aerospace industry, this sort of approach has broader application in areas such as automotive and marine engineering, as well as in the medical industry.

1. Introduction

Aero-engines [1] are extremely reliable machines and operational failures are rare. However, great effort is currently being put in to reducing the number of in-flight engine shutdowns, aborted take-offs, and flight delays through the use of advanced engine health monitoring technology. The key benefits realised by these efforts are the reduction of delays and reductions in the cost of ownership of aircraft. This is reflected in a change in emphasis within aero-engine companies where, instead of selling the engines to customers, there is a fundamental shift to adoption of power-by-the-hour contracts. In these contracts airlines make regular fixed payments based on the hours flown by the engines, and the manufacturer retains the responsibility for maintenance of the engines. To support this new approach, improvements in in-flight monitoring of engines are being introduced, along with the collection of much more detailed data on the operation of the engine.

Each engine produces large amounts of detailed operating data (for example, a civil airliner produces approximately 1 Gbyte of data per engine per flight) which then has to be transmitted and analysed. Rolls-Royce currently has over 50,000 engines in service, with total operations of around 10M flying hours per month. In the future, one can envision many 100s of Gbytes of data being transmitted every day.

The reductions in delays and costs of ownership mentioned above can be achieved by providing an infrastructure capable of managing and analysing the large amounts of data produced by these fleets of engines. This infrastructure has to be capable of performing large compute-intensive optimisation, modelling and analysis to identify faults that have occurred, and, more importantly, to identify potential faults and provide knowledge-based maintenance advice to prevent failures and aircraft downtime.

It is clear that the software tools and expertise that comprise this "virtual work-bench" must be geographically distributed to cope with the demands of an international operation (see Figure 1). This work-bench must also be flexible so that the user can choose the appropriate tools to diagnose the potential problem, rather than being presented with a monolithic fixed system.

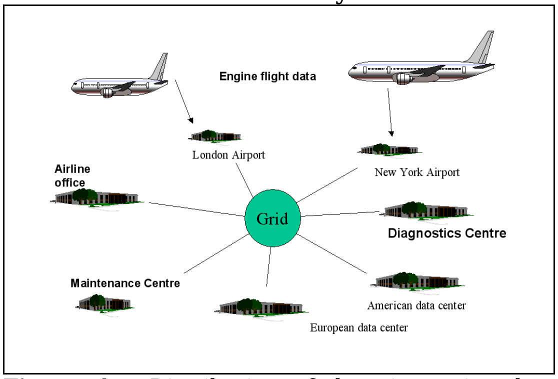

Figure 1 – Distribution of data in a virtual maintenance environment.

The tools comprising this virtual work-bench must be able to identify clusters of anomalies or novelties as they appear so as to give some insight into the underlying cause of a problem (for example, some operators may experience certain faults because of the way that the engines are being operated). The virtual work-bench must then allow the user to test ideas regarding engine diagnoses in a virtual environment using engine performance models. This latter stage is particularly useful when a diagnosis is unknown.

If an abnormality is identified which has not been encountered before (novelty analysis), then the work-bench provides the ability to consult previous engine histories to see if there are other engines that have shown similar abnormalities. The history of an engine with a similar abnormality may provide some knowledge as to the nature of the problem and its solution.

The work-bench also provides the ability to perform large-scale optimisation of both single-objective and multi-objective problems. This allows the user to perform robust optimisation on problems as diverse as engine controller design and the design of optimal maintenance strategies for fleets of aircraft. Using these tools, the user can investigate the impacts of changing operational schedules on cost and aircraft down-time.

The virtual work-bench described above lends itself to implementation using the paradigm of Grid Computing. Grid computing has the potential to mediate the task of diagnosis and prognosis within complex and dynamic operational business and engineering processes. The Grid is capable of providing high performance computing resources on demand, offering a resource for the computationally-intensive tasks of modelling, optimisation and analysis within the decision support process.

2. 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 [2]. Industrial partners in the DAME project are Rolls-Royce plc, who have provided the aero engine data for the diagnostic system (see Figure 2), Data Systems & Solutions, who 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.

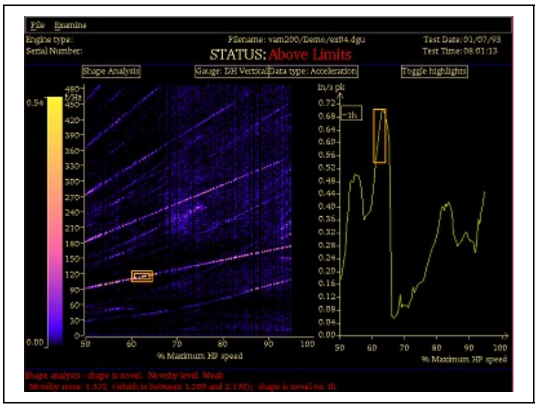

Figure 2 – Representative sample of data downloadable from the engine on-wing monitoring system.

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 Modelling, Optimisation and Decision Support. Grid computing expertise 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.

2.1. Diagnostic Scenario

A typical diagnostic scenario is as follows: The on-wing diagnostic system and its associated ground based system are used prior to the use of DAME. In addition to that initial on-wing diagnosis, DAME provides an automated diagnosis. This is desirable because DAME detects additional situations, for example:

- A recurring errant diagnosis
- A new condition that has not been yet been uploaded to the on-wing monitoring system
- A condition that can only be detected using tools that require extensive ground-based processing facilities

The resultant automatic diagnoses are then be assessed. In the vast majority of cases normal situations are indicated, however, if a condition is detected with a known cause then appropriate maintenance action can be planned. Additionally, in the rare case that a condition is detected without a clear cause then the situation is "escalated" to one of various remote experts who can look into the matter further. The Maintenance Analysts and Domain Experts have access to the data from the current engine flight, and can run searches on historical data, get workflow advice, and run signal processing and simulation tasks to gain an insight into any given event.

3. Modelling, Optimisation and Decision Support in DAME

It is clear that in order to deal with the explosion of 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 perform complex analysis – even across fleets of engines.

The Sheffield UTC has been actively exploring a variety of on-wing control system diagnostic methods and decision support techniques. The key technologies underpinning this work are:

- Case-Based Reasoning
- Integration of Model-Based Fault Detection and Isolation Approaches
- Many-Objective Optimisation using Genetic Algorithms

3.1. 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 [3, 4]. 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. Furthermore, a rank is given to the results, thus valuing the information as “more suitable” or “less suitable”.

Previous research work carried out within the Sheffield UTC includes investigating the use of Case-Based Reasoning techniques for a portable PC-based Flightline Maintenance Advisor [5, 6] 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 (see Figure 3). The outcomes of the initiative included the implementation of a portable Flightline Maintenance Advisor that was trialled with Singapore Airlines.

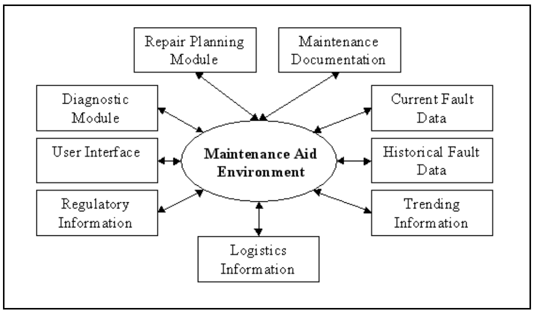

Figure 3 – Structure of the Flightline maintenance Advisor

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 case-base of historical maintenance incidents across an entire fleet of engines [7, 8]. 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 case-base 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 case-base 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.

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 (see Figure 4). Queries for matching cases can be submitted to the CBR service in two ways; directly via a Web browser window (see Figure 5) or automatically via an integrated client such as the one in the DAME automatic workflow system.

* BITE – Built-In Test Equipment

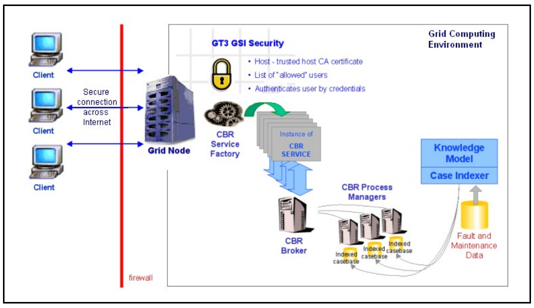

Figure 4 – The CBR service architecture.

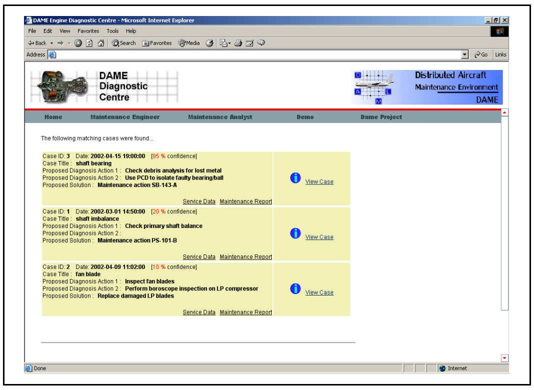

Figure 5 – The CBR service operating within a web-browser.

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.

Future development of the CBR system will incorporate a dynamic learning process whereby the CBR system will learn how a domain expert might troubleshoot a certain problem. The system will then automatically execute the set of diagnostic tests that the domain expert is most likely to run in that situation, and then present the data for further analysis, thus saving valuable time and effort.

3.2. Model-Based Fault Detection and 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 [9, 10]. 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 [11] 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 (see Figure 6). 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 4.

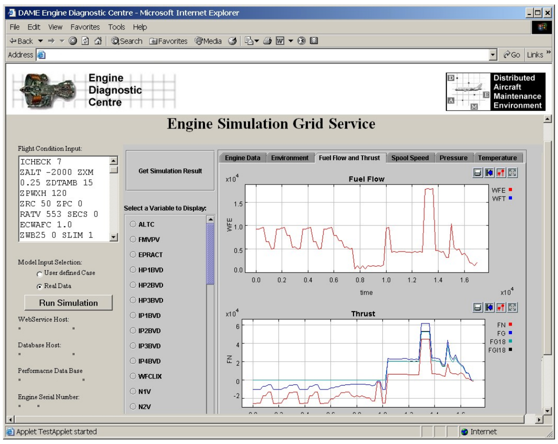

Figure 6 – The Engine Simulation Grid Service.

Figure 7 in the next column 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.

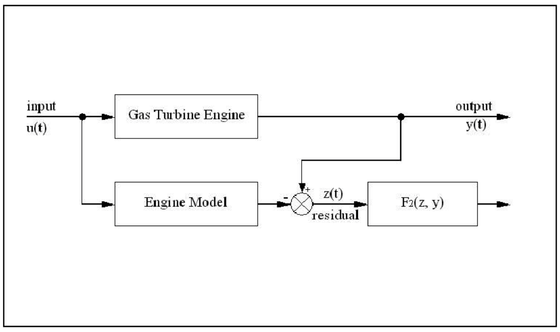

Figure 7 – Simulation-Based Fault Diagnosis.

3.3. Genetic Algorithms for Many-Objective Optimisation

Genetic Algorithms (GAs) are an optimisation technique utilising some of the mechanisms of natural selection [12]. GAs are an iterative method capable of both exploring the solution space of the problem and exploiting previous generations of candidate solutions. This combination of exploration and exploitation results in an intelligent search method that uses objective function pay-off information (gained from the candidate solutions in the previous generation) to guide the direction of the search towards potentially promising areas of the solution space, whilst preventing premature convergence to local optima by using exploratory operators such as mutation.

Many real-world optimisation problems involve the satisfaction of multiple conflicting objectives. In this case it is unlikely that a single ideal solution will be possible. Instead, the solution of a many-objective optimisation problem will lead to a family of Pareto optimal points, where any improvement in one objective will result in the degradation of one or more of the other objectives. Genetic Algorithms are particularly well suited to this kind of many-objective optimisation, because they search a population of candidate solutions. This enables the GA to find multiple solutions which form the Pareto optimal set.

GAs are often able to find high quality solutions to real-world problems, even those which conventional optimisation techniques (such as hill-climbing or constraint satisfaction) fail on or have difficulty with. This is due to the robust nature of GAs in dealing with the noisy or discontinuous solution spaces that real-world problems often have and the ability of GAs to handle 'black box' objective functions where nothing is known about the problem except the nature of the inputs and outputs.

Previous research in the Sheffield UTC has included using many-objective genetic algorithms to solve problems in tuning of aero-engine control systems [13], scheduling [14], and drug discovery [15]. For many real-world optimisation problems the evaluation of the objective function is performed by computer simulation of the system (for example, in [13] a computer model of the engine is used to evaluate the candidate solutions). However, to ensure that the results gained from the optimisation process are meaningful the simulation must be complex enough to capture all the relevant dynamics of the true system. This can result in an extremely computationally intensive simulation. This, coupled with the iterative and population-based nature of genetic algorithms (in a typical GA the simulation could be run 10,000 times), has motivated the development of a grid-enabled version of this many-objective genetic algorithm (see Figure 8) [16].

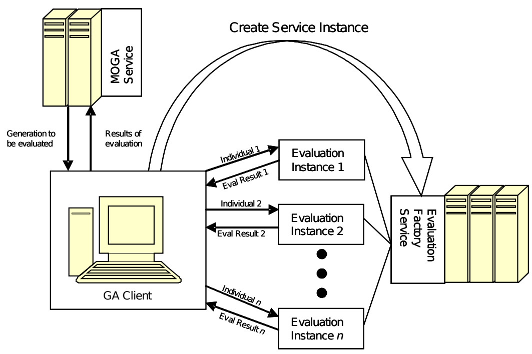

Figure 8 – A Grid-Enabled Many-Objective Genetic Algorithm.

One potential use for this grid-enabled genetic algorithm tool is the evolution of optimal maintenance strategies. Rolls-Royce and Data Systems and Solutions have developed a simulation capable of modelling the lifecycle of engine components across fleets of engines [17]. This Modular Engine Arisings Repair and Overhaul Simulation (MEAROS) is a complex stochastic model that uses engine module failure distributions to predict the required maintenance schedules for fleets of engines.

MEAROS is a computationally intensive simulation, and because of its stochastic nature, each candidate solution has to be run multiple times. The high performance, on-demand nature of the Grid offers an ideal platform for the optimisation of this problem. By grid-enabling this simulation it can be run in a distributed manner to speed up the evaluation of candidate solutions. The grid also provides a secure environment for this evaluation process, allowing the protection of commercially sensitive evaluation function such as those including cost models.

3.4. Use of the Integrated Decision Support System

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.

In the future, the domain experts could incorporate the knowledge gained using these tools (such as the impact of changing operational schedules on engine component lifecycles) into the optimisation of maintenance strategies. This could help to reduce both the operating costs and aircraft downtime.

4. Security

A Grid-enabled modelling, optimisation and 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 knowledge base of engine faults and the engine models contain important information on the design characteristics and operating parameters of the engine. 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 Grid resources on a Grid node if their verified credentials have been registered in the Grid environment by an administrator of that Grid.

5. Concluding Remarks

In this paper we have discussed the use of model-based fault detection and isolation, case-based reasoning and optimisation in advanced decision support systems for maintenance and fault diagnosis of aero-engines.

In particular, this 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 and is made possible by 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 tools presented can be used to identify faults when they happen or predict impending faults on engines through trending. They can also be applied across fleets of engines to perform fleet management or optimisation of engine maintenance strategies.

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 reduction in 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. Distributed Aircraft Maintenance Environment (DAME) project; www.cs.york.ac.uk/dame
3. J. Kolodner. *Case-Based Reasoning*. Morgan Kauffman, 1993.
4. D. B. Leake. *Case-Based Reasoning: Experiences, Lessons & Future Directions.* The MIT Press, 1996.
5. S. M. Hargrave. *Evaluation of Trent 800 Portable Maintenance Aid Demonstrator.* Rolls-Royce University Technology Centre, University of Sheffield, Technical Report RRUTC/Shef/R/98202, 1998.
6. S. M. Hargrave. *Review of Performance-Based Diagnostic Tool*, Rolls-Royce University Technology Centre, University of Sheffield, Technical Report RRUTC/Shef/TN/98204, 1998.
7. 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.
8. 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.
9. Y. G. Li. *Performance-Analysis-Based Gas Turbine Diagnostics: A Review.* In Proceedings of I Mech. E, Part A: J Power and Energy, **216(5)**, pp. 363-377, 2002.
10. R. J. Patton, C. J. Lopez-Toribio, and F. J. Uppal. *Intelligence Approaches to Fault Diagnosis.* In Applied Mathematics and Computer Science, **9(3)**, pp. 471–518, 1999.
11. 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, **31(3)**, August 2001.
12. D.E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning,* Addison-Wesley, 1999.
13. P.J. Fleming, R.C. Purshouse, A.J. Chipperfield, I.A. Griffin, and H.A. Thompson, *Control Systems Design with Multiple Objectives: An Evolutionary Computing Approach*, Workshop in the 15th IFAC World Congress, Barcalona, 2002.
14. K.J. Shaw and P.J. Fleming, *Initial Study of Practical Multi-Objective Genetic Algorithms for Scheduling the Production of Chilled Ready Meals,* In Proceedings of Mendel '96, the 2nd International Mendel Conference on Genetic Algorithms, 1996.
15. O. Nicolotti, V.J. Gillet, P.J. Fleming, and D.V.S Green, *Multiobjective Optimisation in Quantitative Structure-Activity Relationships: Deriving Accurate and Interpretable QSARs,* In Journal of Medicinal Chemistry, **45**, pp. 5069-5080, 2002.
16. A. Shenfield and P.J. Fleming, *A Service Oriented Architecture for Decision Making in Engineering Design*, In Advances in Grid Computing - EGC 2005, LNCS 3470, eds. P.M.A. Sloot, et. al., pp. 334-343, 2005.
17. J. Crocker, A Methodology for the Prediction of Maintenance and Support of Fleets of Repairable Systems, PhD Thesis, University of Exeter, 2001.
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.* In 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.