TIGER™: Continuous Diagnosis of Gas Turbines

Robert Milne [1] and **Charlie Nicol**[1]

Abstract. The TIGER system is in continuous use on 20 gas turbines across four continents. Continuously performing fault diagnosis of industrial gas turbines used for power generation. TIGER receives several hundred data inputs at a once per second interval and performs several layers of fault detection and diagnosis using expert system techniques in a temporal reasoning rulebase.

Development of the system has now extended for more than six years and more than 80 man-years. The diagnostics are supported by an extensive user interface, data archive, trending and display system. Across the 20 current locations, TIGER has helped to identify and diagnose over 700 incidents representing 200 different fault categories. Because of the high value of gas turbines and the high expenses in operating them, TIGER's benefit is considerable and has been documented at $150,000 or more for most of the locations.

1 INTRODUCTION

Industrial gas turbines are becoming the work horse of the power generation industry. New power plants can be built relatively quickly and cheaply, and when all goes well, they run continuously for long periods of time with very little attention required from operators. However, in the increasingly competitive electricity generation industry, profit margins are very tight, requiring good return on investment and small margin for the cost of unexpected maintenance problems or failures. In addition, most sites use the minimal level of manning that is practical and there is an increasing skill gap between the complexities of a gas turbine and the things that can go wrong, and the operators in charge of ensuring that the machine runs day-to-day.

In fact, a major trend is to run the gas turbines unattended with only remote support. Hence, in order to make a good profit for a power generation company, it is vital that the gas turbine has high availability, minimum unexpected maintenance costs, and at the same time, can be operated by someone who has relatively little skill and training in the intricacies of a gas turbine.

What is needed is a system that can help the plant increase its availability by avoiding shutdown time, either as an unplanned failure of the system or if a failure does occur, minimising the time that the system is down. To achieve this the goal is condition monitoring [1]. Assessing the condition of the gas turbine continuously and detecting at the earliest possible stage any developing faults or problems. When a problem has been detected it is also vital to diagnose it as completely as possible, and as quickly as possible.

Because the detailed knowledge of how a gas turbine operates and problems manifest themselves what is needed is a system which can make good diagnostic knowledge and information available to the people responsible to ensure the gas turbine is available when needed [2].

The aim is to detect problems at the earliest possible stage and be able to rectify them before they become expensive failures. If an unexpected failure does occur, then to rapidly understand what the problem was so that it can be fixed as efficiently as possible, and the machine returned back to power generation quickly. Another important factor is to get optimum efficiency out of the machine since machines burn $1.5 million worth of gas every month, 1-2% drop in efficiency can be very expensive.

The TIGER software was developed to address these requirements. It receives 600 measurements from the gas turbine controller at once per second intervals and performs an extensive fault detection and diagnosis every second. TIGER's task is to check the data inputs just as the best gas turbine engineer would, and ensure that everything is operating at the correct level. More importantly, as it makes dynamic changes such as an increase in power level, to confirm that the gas turbine has reacted properly in doing this.

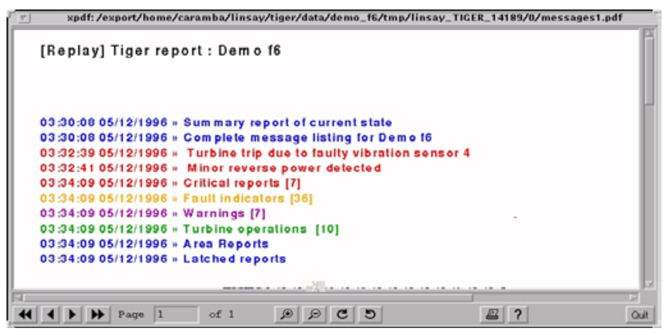

Figure 1. The TIGER fault manager top level

It is impossible to diagnose all faults which occur on a gas turbine, so TIGER's diagnostic system is structured in a layered way. At the lowest level is an extensive fault detection set of mechanisms which look for a wide variety of abnormalities. These are then combined using a patented temporal rulebase system. A process of hierarchical abstraction consolidates conclusions with high confidence until the system is able to reach a diagnostic conclusion for a particular event. Very often the cause of a problem is not known to TIGER from its data inputs or possibly not known at all to gas turbine engineers. In this case TIGER reports a partially consolidated conclusion to assist the engineer in focussing on the problem without actually diagnosing the exact cause. This strategy of always trying to provide a helpful focus for a problem greatly enhances its usefulness.

[1] Intelligent Applications Ltd, 1 Michaelson Square, Livingston, West Lothian, EH54 7DP, Scotland, UK, email: rmilne@bcs.org.uk - email: cnicol@bcs.org.uk

The TIGER diagnostics are presented through a fault manager using hypertext. (See Figure 1) The most important problems are presented at the highest level and the hierarchical tree leading to the fault conclusions, can be viewed by exploration with a mouse. In fact, the output of the fault detection and diagnosis levels are presented three times: first in a complete time sequence; second graded by priority from critical, warning, detected incidents, or operational incidents; and finally, grouped by the sub-systems of the gas turbines.

To allow a gas turbine engineer to further troubleshoot problems, TIGER has an extensive set of displays, graphical and trending tools. The graphical system is used to display the relevant data inputs in a close time sequence to verify the conclusions of the diagnostic system. This has proven invaluable for permitting engineers to have confidence in a diagnosis.

The trend system automatically builds a 24 hour trend for all 600 inputs, each day. Weekly and monthly trends can be constructed easily. This provides not only a way to check for long term deterioration, but also to get a rapid overview of the data.

2 TIGER IN USE TODAY

TIGER is currently running at 20 locations on four continents. Although most of these locations are for power generation, one of them is in use at a process plant. A typical gas turbine, which TIGER monitors, is a General Electric Frame 6 gas turbine, generating 40 Mwatts of electricity. The output of the gas turbine is converted to steam through a heat recovery boiler, and this steam is used to generate further electricity. In many of the plants, the steam is also provided to an associated process such as a paper mill as an important product. The ownership of the plant varies from national power companies to small independent power producers.

TIGER is also in use on four oil platforms in the North Sea providing power generation on the oil platforms for major oil companies. In approximately half of the locations, TIGER can be accessed remotely so that a central monitoring office in Scotland is able to check each gas turbine on a daily basis and also provide remote diagnostics support. Companies using TIGER include: National Power, Powergen, Exxon Chemical, British Petroleum, Chevron, Oryx, and Kvaerner Energy.

There are two main groups of users for TIGER. At the plant level an operator can check TIGER every shift and see if any important warnings are being presented. It is also used locally to when a problem does occur. Generally, the operators don't have in depth knowledge of the gas turbine and are only concerned with faults requiring immediate action. The second group, and most common user type, is the engineer responsible for the operation of the gas turbine. The engineer needs to know about developing problems and will also troubleshoot more difficult situations. The range of problems detected by TIGER is most valuable to this type of engineer, hence giving him a high motivation to use the tool.

For locations where there are not engineers on site, the TIGER system is monitored and acted upon remotely by the TIGER monitoring team in Scotland. This remote technical backup has considerable benefit in saving travel costs and time, particularly for offshore oil platforms.

At every installation, TIGER has been invaluable in identifying a wide range of faults and problems which were previously undetected. It has now been used in a wide variety of situations to further troubleshoot and identify the cause of any problem so that rectification can be made.

3 TIGER DIAGNOSTICS

The TIGER system is a layered diagnostic system working at increasingly higher levels of abstraction. (See Figure 2) The lowest levels are a wide variety of specialist fault detection mechanisms. They are originally implemented in an expert system environment but have been coded directly into C for high efficiency. Many of the fault detection checkers are tuned to the performance of a particular gas turbine through an automatic tuning algorithm that examines one month's worth of data before setting the pattern of normal behaviour.

The fault detection mechanisms check that parameters are within normal operating limits for that gas turbine, but also look at rates of change and relationships to other variables. The output of the fault detection level is processed by a second level rulebase which uses time constraints between events to abstract events into actions on the turbine. Although it can be viewed as a classical knowledge base, it now has several thousand diagnostics based on the wide variety of incidents. These diagnostics are implemented in an expert system environment and then compiled to pure code for faster execution.

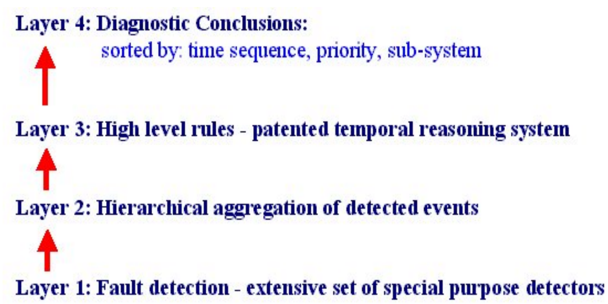

Figure 2. The layers of TIGER diagnostics

At the upper levels, a set of high level rules consolidate a combination of detected events and lower level rule conclusions into higher level conclusions. These might be a consolidation of faults or it might be a specific diagnosis such as a trip caused by faulty sensors. The approach to combine these is essentially a rulebased temporal reasoning system but using a patented mechanism.

The TIGER diagnostic system levels of hierarchy are summarised as follows.

Fault Detection

The purpose of this level is to detect any abnormalities or unusual events on the gas turbine. TIGER uses a variety of proprietary techniques and currently within the TIGER system, there are over 4000 fault detectors checked every second. These vary from simple context sensitive limit checkers to specialised mechanisms to ensure the dynamic response of the turbine is as it should be.

Diagnostic Rules

Using well established forward chaining techniques from knowledge based systems, TIGER uses diagnostic rules to specify the correct interaction of parameters and sub-systems within TIGER, and determines whether they are functioning properly or not. This provides a very sophisticated ability to detect whether the working of the turbine is correct. This mechanism uses temporal reasoning techniques to track sequences of events over time.

High Level Rules
The TIGER fault manager uses a patented [3] temporal reasoning to combine the results of the fault detection and diagnostic rule levels into higher level conclusions. In essence, this module is a forward chaining rule base. The temporal techniques allow it to collect related events in a time window. This resolves the problem that related events aren't always detected at the same time. For example, when there is a time window on a fault detector or a valve takes a few seconds to move. TIGER combines the thousands of diagnostic checks it performs in a hierarchical fashion so that the user sees a single fault conclusion for a major incident (See Figure 3).

Accounts For
Once a problem has been identified, the existence of this problem accounts for many other diagnostic messages. TIGER organises the wide range of phenomena that occurs when a problem exists through the 'accounts for' set, so that a single conclusion can account for many consequent results. This organises the many events, very simply for the user, and reduces the alarm explosion that can occur during a major incident.

Areas
When it is not possible to provide a concise diagnosis, the fault detection and low level diagnostic conclusions are grouped by areas, such as, all the outputs related to the fuel system, the inlet guide vanes or the combustion system. This allows an experienced engineer to rapidly focus on the core problem area and gives him support in determining what problem has just occurred.

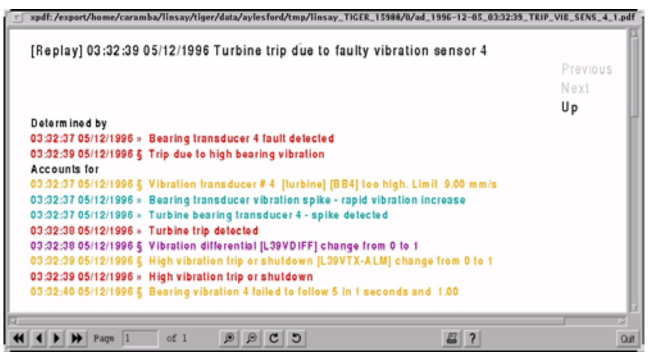

Figure 3. A TIGER rule display

Currently there are approximately 4000 rules and fault detection checkers within the TIGER system.

The diagnostics were built both 'bottom up' and 'top down'. For each sub-system of the gas turbine, a failure analysis was conducted identifying how the system should work, what might go wrong, and how that would be indicated by the TIGER variables. This led to a ground up development of a wide variety of rules. Currently the main driver of new diagnostics, is by examining incidents which actually occur and creating improvements or new diagnostic rules. A knowledge engineer carefully examines each incident which has been logged in the TIGER incident database, checks what diagnostic conclusions TIGER reach and then makes any necessary improvements or adjustments.

Although gas turbines are relatively well understood, it is incredible how much they get up to in unexpected ways. Many strange behaviours and actions occur which engineers are surprised with and do not understand immediately, hence, being data driven has been very productive to the improvement of the TIGER diagnostics. This data driven approach has been driven by over 700 incidents logged in the TIGER incident database, and approximately 25 fired years of gas turbine operation.

From the Artificial Intelligence viewpoint, these are relatively strange forward techniques and technical approaches. However, the complexity of the system requires considerable experience to develop and evolve.

4 TIGER BENEFITS

The following is a summary of locations where TIGER has produced a saving of at least $150,000. This represents the majority of the installations which have been in place for at least one year.

Aylesford Newsprint
TIGER was used in cataloguing and identifying problems with the Nitrogen Oxide (NOx) steam injection system. It has been used to greatly reduce the troubleshooting time for a number of trips by helping to rapidly identify the cause of a trip. In many cases, it was able to identify that the cause of the trip was not part of the gas turbine, and not damaging to the gas turbine itself. TIGER is the key tool being used to identify a problem with hydraulic pressure which is affecting the fuel valves and behaviour of the gas turbine.

BASF
TIGER helped to identify and correct a number of initial commissioning and installation problems. One of these was a problem with the stability of the steam supply for NOx injection. Based on the documented behaviour of the gas turbine through TIGER, a refurbishment of the steam injection supply system was conducted which eliminated many of the problems.

Alba Platform
TIGER was instrumental in identifying how liquid fuel start problems were because of instability of the liquid fuel system at low load. This was causing problems with starting the gas turbines on liquid fuel and directly affecting their availability.

Dubal
TIGER identified over 30 problems that were affecting the turbines behaviour and operation on a regular basis, but which were unknown prior to its installation. In addition, there were two failures of the atomising air compressor, which were rapidly diagnosed using TIGER and helped to considerably reduce the downtime.

Exxon Chemical
TIGER has become the primary tool for the engineers to monitor and analyse the behaviour of the gas turbine. This then provided critical data in the de-bottlenecking process for the plant. A major bottleneck was the reaction of the second stage nozzles at high power. TIGER was used to documentation and find a resolution to this problem.

At higher power levels, the problem again reoccurred and TIGER was used to understand and resolve the issues. As Exxon extracted higher power out of the machine, the window into the gas turbine, provided by TIGER, was critical in supporting this process and enabling them to get more power out of the turbine and a higher output out of the plant.

Fellside
Prior to the installation of TIGER it was noted that first stage nozzles were having to be changed more frequently than expected. Almost immediately after TIGER was installed, the problem was

identified as major oscillations caused by the fuel control system. TIGER was instrumental in the troubleshooting and fault isolation process. Using the knowledge gained through the TIGER analysis, the problem was rectified with considerable savings.

Hjorring

TIGER was able to identify a number of intermittent problems with the valve movements and flame scanners, as well as the DLN1 NOx control system. The experienced site engineers used TIGER to track down a range of minor problems and faults affecting the behaviour of the gas turbine and improve its operation.

Kemsley Mill

Due to a developing problem in the combustion system, the turbine was operating with a very high exhaust spread, however, it was not convenient to shut the turbine down to rectify it. Using TIGER as a window into the behaviour of the gas turbine, the site was able to nurse the gas turbine operation through to a weekend when they were able to have the appropriate parts and maintenance engineers ready to correct the problem as quickly as possible.

Through TIGER they were able to get the maximum output of the turbine without tripping it through the period until it was appropriate to shutdown. By delaying the shutdown until they had all the equipment and personnel ready, plus at a lower tariff time, it provided considerable return to the site.

ETAPS Turbine 1 & 2

On this oil platform, neither gas turbine was able to start on liquid fuel as gas fuel was not available. An engineer used TIGER to rapidly identify the precise failure within the liquid fuel system and direct the repair work as efficiently as possible. After expensive downtime, by using TIGER, the site was able to get the gas turbines up and running quickly, providing considerable return on the oil platform.

TIGER was also used to catalogue all the problems and deficiencies from the initial installation of the gas turbine. This enabled the customer and the current holder of the maintenance contract to make a warranty claim against the initial installation. TIGER documented all the problems affecting the behaviour, efficiency and operation of the gas turbine at a very early stage so that they could be rectified before they caused further long-term problems.

TIGER has consistently identified a number of faults and problems on gas turbines including many of which were not known, but their rectification will lead to a more efficient and smoother long-term operation.

5 TIGER DEVELOPMENT

The first prototype of the TIGER system was built with Exxon Chemical in Scotland, approximately seven years ago. This then led to a three year ESPRIT supported project to develop the initial TIGER prototype [4], [5], [6]. That project involved a collaboration of Intelligent Applications Ltd, Exxon Chemical Olefins Inc, John Brown Engineering, Dassault Aviation, CNRS-LAAS, and University Polytechnic of Catalonia (See Figure 4).

After the ESPRIT project, John Brown Engineering and Intelligent Applications worked together to commercialise the system and have continued to further develop the software and diagnostics.

A key aspect of the development team is that it combines artificial intelligence developers, gas turbine experts and software engineers. Each provides expertise in a complementary area.

At this point, the investment in TIGER has continued over a six year period and totals more than 80 man years and $5 million of development. Development is still continuing at a rapid pace. Future developments include continued expansion of the diagnostics, widening the range of gas turbines for which TIGER is available, enhancing the performance monitoring functionality and the addition of model based prediction and isolation techniques.

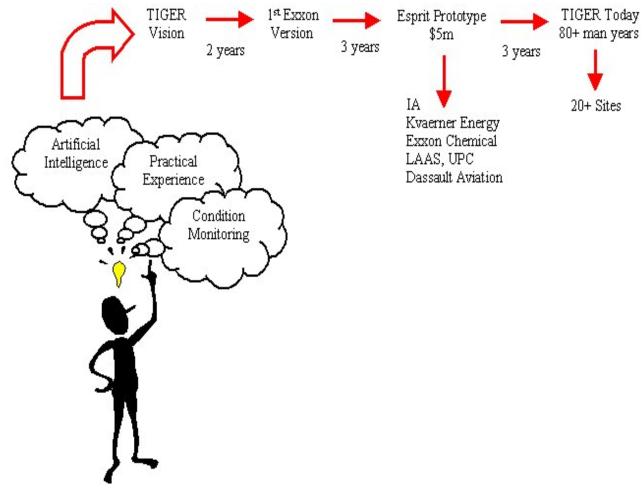

Figure 4. The development of TIGER

Each TIGER installation is 'almost' standard. TIGER has to be configured to the exact sensors which are installed on each gas turbine. Although these are very similar, the detailed differences have to be identified and the configuration adjusted. TIGER will then auto tune its fault detection systems to the way a particular gas turbine operates. The diagnostics automatically activate (or deactivate) depending on which inputs are available. The configuration for a TIGER for a Frame 6 gas turbine takes about two days. Initial user training takes about one day, reflecting the ease of use of the system.

6 MODEL BASED FAULT DETECTION AND DIAGNOSIS

Although the core of TIGER is a temporal rulebased system, there is an important class of faults to detect which require a model based approach, that is, the predicted value for a measurement depends on a complex combination of several inputs. One area where this is particularly important is predicting the expected power output of the gas turbine, giving a combination of atmospheric air conditions, fuel setting, and steam injected for pollution control.

Recently TIGER has been extended to include a neural net mechanism to cross predict these values. The neural network is able to collect training data automatically and on-line and re-train itself after each major change to the turbine. Testing to date has demonstrated that it makes very accurate and reliable predictions of the expected power output as well as four other key indicators.

The focus of current development is to use a qualitative interval predictor based on the Ca-En software [7], [8]. This combines model based fault detection ability by predicting what expected outputs should be with a model based fault isolation mechanism providing this additional leverage over rulebased implementations [9].

The qualitative predicted predictions system is being used on critical sub-systems of the gas turbine such as the gas and liquid fuel systems. Figure 5 shows the display of the fault isolation for the gas fuel systems as use by the model-based module. Ca~En fault isolation generates diagnostic messages which are used by the rest of the TIGER diagnostic system. One of these messages is visible in the middle of the diagnostics window at the lower left of Figure 5.

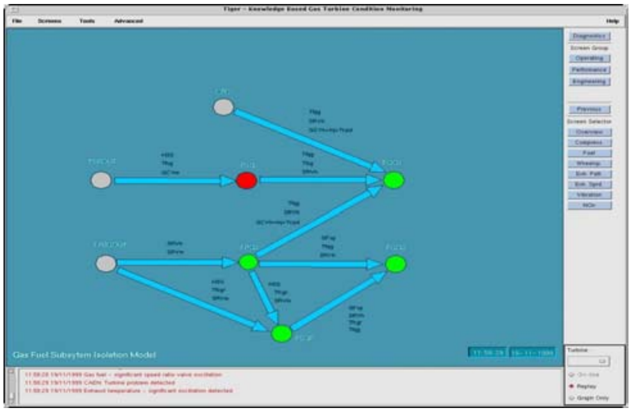

Figure 5. Gas Fuel Sub-System Isolation Model

The model based system's power is restricted primarily by the availability of key measurements to allow it to isolate faults with high resolution. Unfortunately the measurements of a gas turbine are generally provided to support monitoring that it is working safely rather than diagnosing and isolating faults. The Ca-En fault prediction and diagnostics are now implemented and undergoing extensive testing on site.

The TIGER system is being steadily installed on more gas turbines and the number of users continues to increase.

7 SUMMARY

This paper has presented the TIGER knowledge based condition monitoring system for gas turbines. It is in use at over 20 locations providing continuous fault detection and diagnosis for gas turbines used in power generation. TIGER has consistently found a wide variety of faults and problems for gas turbines and has produced benefits in excess of $150,000 at all the locations for which it has been installed for more than one year.

The system uses a temporal reasoning knowledge based system and was developed with classic knowledge acquisition techniques primarily driven by real examples from gas turbines. It is interesting that the AI techniques are not complex and innovative. Instead, well-established forward chaining rule based techniques have been demonstrated to be adequate. The power is the knowledge, not the technical approach. The system continuous to be developed and is being installed on increasing numbers of gas turbines.

8 REFERENCES

[1] Handbook of Condition Monitoring - Techniques and Methodology. Chapter 22, 521-540. Editor: A. Davies. Publishers: Chapman & Hall (1998).

[2] R. Milne, Knowledge Based Systems for Condition Monitoring, *Handbook of Condition Monitoring - Techniques and Methodology*. Chapter 22, 521-540. Editor: A. Davies. Publishers: Chapman & Hall (1998).

[3] United Kingdom Patent Application No: GB2323197A. Filed in 1997 by applicants Intelligent Applications Ltd.

[4] R. Milne, C. Nicol, M. Ghallab, L. Trave-Massuyes, K. Bousson, C. Dousson, J. Quevedo, J. Aguilar and A. Guasch, Real-Time Situation Assessment of Dynamic Systems, *Intelligent Systems Engineering Journal*. Autumn, 103-124 (1994).

[5] R. Milne, C. Nicol, L. Trave-Massuyes, J. Quevedo, TIGER: Knowledge Based Gas Turbine Condition Monitoring, *AI Communications Journal*. Vol 9, N° 3, 1-17. Publishers: IOS Press (1996).

[6] L. Trave-Massuyes and R. Milne, Gas Turbine Condition Monitoring Using Qualitative Model Based Diagnosis, *IEEE Expert magazine, "Intelligent Systems & Their Applications*". Vol 12, N° 3, 21-31. Publishers: IEEE Computer Society (1997).

[7] K. Bousson and L. Trave-Massuyes, Fuzzy Causal Simulation in Process Engineering, *Proceedings of IJCAI-93* (1993).

[8] K. Bousson and L. Trave-Massuyes, Putting More Numbers in the Qualitative Simulator Ca~En, *Proceedings Second International Conference Intelligent Systems Engineering*, IEE Publications, Stevenage, UK, 62-69 (1994).

[9] L. Trave-Massuyes and R. Milne, Diagnosis of Dynamic Systems Based on Explicit and Implicit Behavioural Models: An Application to Gas Turbines in ESPRIT Project TIGER, *Applied Artificial Intelligence*. Vol 10, N° 3, 257-277. Publishers: Taylor & Francis (1996).