Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Simulation of a Condition Monitoring Scheme for a Neutral Beam Injector Cryogenic Pump ⋆

N. Wright * **R. Dixon** * **C. Ward** *

* *School of Electrical, Electronic & Systems Engineering, Loughborough University, Leicestershire LE11 3TU, United Kingdom (Tel: 01509 227015/227018; e-mail: n.wright@lboro.ac.uk).*

Abstract: This paper presents the simulation testing of a model-based condition monitoring scheme for the Joint European Torus neutral beam cryogenic pumping system. The scheme was tested by examining its response to a range of realistic simulated faults. Ten faults were tested. The scheme detected each of the faults and uniquely identified six of them. The scheme was shown to have an improved performance compared to the current manual detection method.

1. INTRODUCTION

Commercial nuclear fusion power has, for several decades, been seen as a panacea to many of the problems associated with energy production. The prospect of harnessing this means of energy production is an enticing one, as it would potentially allow abundant, cheap energy to be produced with a relatively minor environmental impact (Bradshaw et al. [2011]). Although the science governing the fusion process is well understood, the engineering capability required to successfully harness it for useful power production is still under development.

The Culham Centre for Fusion Energy (CCFE), the United Kingdom's primary fusion research organisation, operate the Joint European Torus (JET) and MegaAmp Spherical Tokamak (MAST) experiments. Over the years, one of the main engineering challenges at CCFE has been to improve the reliability of the JET experiment. An important part of that process has been to improve the reliability of two essential plasma heating devices, known as neutral beam injectors (NBIs). Great progress has been made on this front, as evidenced by the much reduced failure rate of the devices since they started operating in the late 1980s (King et al. [2005]). Nonetheless, failures and unplanned maintenance activities remain difficult to deal with, even with their reduced frequency, owing to both the complexity of the devices and the practicality of carrying out repair work, as physical access to the device is often restricted.

The research presented in this paper concerns the simulation of a condition monitoring scheme for the cryogenic pumping systems that support the NBI devices, which was previously developed by Wright et al. [2013]. Currently, model-based condition monitoring has had very few fusion cryogenic applications, with only one other example appearing in the literature (Zhou et al. [2012]). The aim of this research was to demonstrate the utility of this condition monitoring scheme, and to ultimately support the case for the inclusion of model-based condition monitoring in current and future fusion experiments.

⋆ This work was sponsored by the Culham Centre for Fusion Energy, Oxfordshire UK, as part of a CASE studentship award.

2. PLANT DESCRIPTION

The cryogenic pumping systems, or cryopumping systems, are an important part of the NBIs. They create and maintain an ultra-high vacuum within which the NBIs operate ($\approx 10^{-9}$mbar). An ultra-high vacuum is an essential precondition for the successful operation of the NBIs, and consequently, the whole fusion device. The cryopumps are expected to operate for extended periods of time, typically several months, uninterrupted.

Fig. 1. The NBI cryopump panels prior to installation, courtesy CCFE

Physically, the pumps operate by condensing any trace gasses in the vacuum space onto two banks of extruded

aluminium panels located along the sides of the NBI. The panels are cooled to cryogenic temperatures by liquid helium ($\approx$ 4.23K @ 1bar). The liquid helium is transferred to and from the panels by means of a specially designed transmission line with several thermal insulation stages. Prior to delivery the helium is stored in a 10kl cryostat tank, that is periodically refilled.

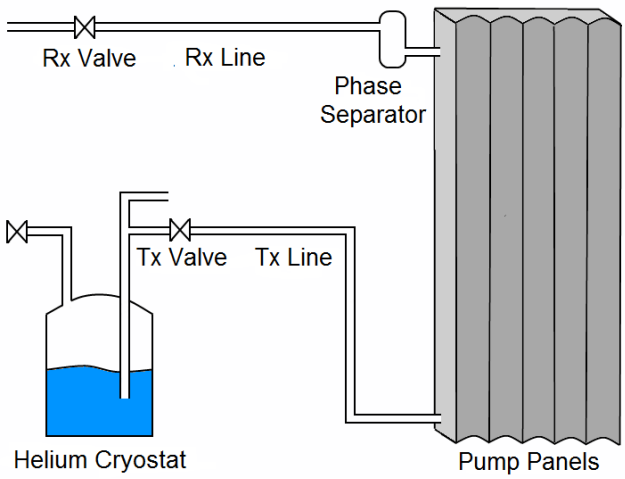

Fig. 2. An illustration of the NBI cryopump helium loop

As trace gasses are condensed by the cryopump, a thermal load proportional to the pumping rate is applied to the pump panels. In addition to this, constant heat loads are applied by both conduction from the supports holding the panels in place and by radiation from warmer bodies inside the NBI vessel. As a result, the liquid helium inside the pump panels boils off over time and must be replaced. A phase separator located on the top of each bank of panels allows the egress of helium gas and measures the amount of liquid helium remaining with a level gauge. Automatically controlled valves on both the supply (Tx) and return (Rx) portion of the transmission line mediate the flow of helium in either direction.

The pumps have several operational modes. The mode of interest for this research is the FULLDOWN mode, which is considered exclusively. In this mode the helium inside the pump is continuously topped up. This is generally active for the longest periods of time during a typical experimental campaign and at all times while the experiment is actively run.

2.1 Plant Model

The development of a dynamic simulation model of the plant was an essential step in this research. The busy operational schedule of the plant meant that it could not be used as a platform for physical experiments, owing to the disruption they would have caused. As such, a novel, accurate non-linear simulation model of the plant was developed. This model was presented by Wright et al. [2012]. This simulation model was validated using two different sets of recent process data to ensure it genuinely represented the plant during normal operation. Once validated, the simulation model was used as a surrogate for the plant during the research, allowing for the rapid design and execution of experiments to test the condition monitoring scheme.

The following plant process variables are predicted by the model:

Process Variable	Units
Helium Cryostat:	
Liquid Volume	m^3
Liquid Level	%
Supply Line:	
Liquid flow rate	$kg.s^{-1}$
Liquid pressure	Pa
Liquid temperature	K
Liquid density	$kg.m^{-3}$
Return Line:	
Gas flow rate	$kg.s^{-1}$
Gas pressure	Pa
Gas temperature	K
Gas density	$kg.m^{-3}$
Pump Panel Assembly:	
Liquid/gas void fraction	Unitless
Liquid/gas pressure	Pa
Liquid/gas temperature	K
Liquid level	%
Gas flow rate	$kg.s^{-1}$

Table 1. Simulated plant process variables

These variables fully describe the state of the plant. Many of these process variables are physically measured and recorded on the real plant, and were used to validate the model.

3. THE CONDITION MONITORING SCHEME

The illustration in Fig. 3 is a representation of a generic quantitative residual generation fault detection and isolation (FDI) scheme. This type of condition monitoring scheme was used for this research because a quick, deterministic response to faults was judged to be desirable, and these types of schemes typically have those characteristics (Venkatasubramanian [2003]).

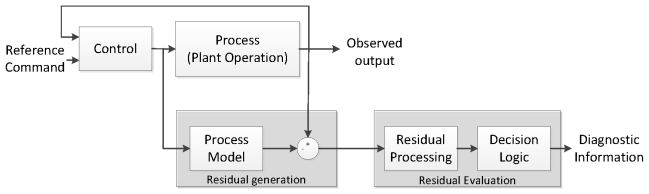

Fig. 3. A generic schematic for a residual generation FDI scheme

Fig. 4 is an illustration of the main components of the condition monitoring scheme presented by Wright et al. [2013]. It shows the flow of information and main processing stages, going from left to right. On the left the cryopumping system is represented (for this design a non-linear simulation model is used as a surrogate plant - see section 2.1). Data collected from the plant is passed to a bank of Kalman filters. This bank acts as a residual generator. The Kalman filter bank was tuned (i.e. weighted) according to the anticipated level of measurement and process noise. The Kalman filter bank estimates ten different plant process variables in a fault free state. The difference between the measurements and estimates (the residual) are passed to the middle block: the residual processing block.

The residual processing block represents the detection logic used to determine if the residuals generated by the

Kalman filter bank are indicative of a fault. Process noise and modelling errors will typically cause some degree of inaccuracy in a Kalman filter estimate, therefore a non-zero residual is to be expected, even in the absence of a fault. As such, the detection logic implemented for this design uses thresholds to determine if the residuals are sufficiently large as to be reasonably attributed to a fault. When a residual signal is large enough to cross a threshold, that threshold is flagged, and remains flagged until it is reset by an end user. Both positive and negative thresholds were assigned to each residual, so that both the magnitude and the direction of the residual could be checked.

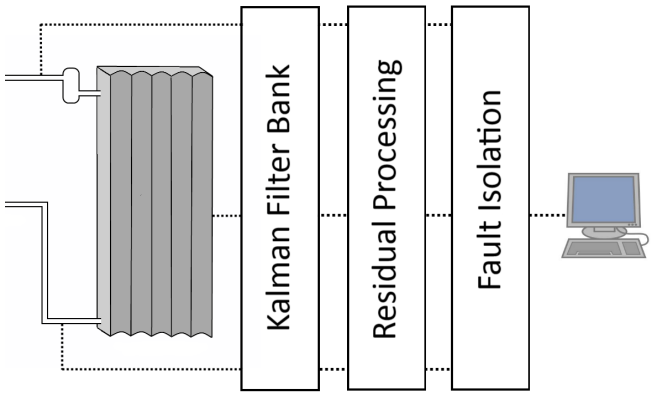

Fig. 4. The main components of the condition monitoring scheme

The third block (on the right of Fig. 4) represents the fault isolation logic. The purpose of the isolation block is to help identify what fault (or set of faults) has occurred, once it has been detected by the residual processing block. When a fault occurs, the residual processing block passes the status of all the threshold flags to this block. The combination of threshold flags can be used to isolate a set of faults (or, ideally, an individual fault), because different faults affect the plant in their own particular way. The faults therefore have a unique manifestation in the residual signals and the detection thresholds are flagged accordingly. This is known as a fault signature. For example, in a cold storage tank system, a leak might cause a drop in fluid level which would flag a threshold associated with a level gauge, but a thermal fault might flag a threshold associated with a temperature gauge. In this case, examination of the residual flags would allow the one fault to be differentiated from the other. For the design used in this research, a fault matrix which lists the combination of flags for each identified fault was created, and the isolation block compares each fault signature with this matrix to isolate faults.

The computer terminal on the right of the illustration represents the final part of the design: the user interface. A simple Graphical User Interface (GUI) was designed and implemented in Matlab (Mathworks Inc. [2013]), which presents the output of the fault isolation stage in an easy to interpret manner. When the isolation block matches the fault signature to a set of faults, the candidate faults are presented to the end user in order of their likelihood, together with a written description.

4. FAULT SIMULATIONS AND RESULTS

The effectiveness of the condition monitoring scheme was tested by simulation in Matlab/Simulink. The three performance criteria deemed to be important were: fault detection rate (percentage of faults detected), fault isolation rate (percentage of faults uniquely identified), and detection time (the time from the onset of a fault and its detection).

4.1 Fault Selection

The first step of the assessment procedure was to select a sample set of faults to simulate. As would be expected, a system of this complexity is potentially subject to a wide range of faults, of which only a small set have occurred historically. Therefore, in order to identify points of failure, it was necessary to use an analytical technique. Several tools are available, the most common of which are described by Rausand and Høyland [2004]. In this instance a Failure Mode Effect and Criticality Analysis (FMECA) was conducted. This method was used because it allows the generation of a comprehensive list of faults that can be ranked according to both their severity and likelihood.

The first stage of the FMECA procedure was to define the boundaries and operating state of the system. For the purpose of this procedure, the cryopumping system was treated as six discrete components corresponding to the physical assembly: The helium tank, transmission line, pump cold panels, pump shielding, valve box, and control and data acquisition.

For each of the plant components, the following analytical steps were completed and encoded in a table:

(1) The functions of each component were defined.
(2) The failure modes of each component were defined, in terms of non-fulfilment of the functions defined in the previous step.
(3) The potential causes of each failure mode were listed.
(4) The probable effects of each failure were listed, both for this component and for the other system components.
(5) The severity of the failure mode was estimated, according to the scale below.

The following definitions of fault severity were used:

(1) Catastrophic - Any failure that prevents the cryopump from pumping the NBHS gas space, or that results in damage to other systems within the NBHS.
(2) Critical - Any failure that impairs the cryopump performance to the extent that the NBHS vacuum pressure would likely increase.
(3) Major - Any failure that impairs the cryopump performance, but would not likely result in an increase in NBHS vacuum pressure before the next scheduled maintenance activity.
(4) Minor - Any failure that does not impair the cryopump performance.

In total, forty nine different failure modes were identified. A summary of the failure mode severity for each plant component is summarised below in Table 2.

	Catastrophic	Critical	Major	Minor
Helium Tank	3	2	1	0
Transmission Line	1	3	2	0
Pump Cold Panels	2	1	0	0
Pump Shielding	1	3	0	0
Valve Box	1	11	5	4
Controls	3	3	3	0
Totals	11	23	11	4

Table 2. A summary of the failure mode classes

A significant proportion of the failure modes categorised as catastrophic are associated with leak events (both liquid and gas) and events that prevent the transmission of cryogenic fluid. The cryopump relies on an uninterrupted supply of cryogenic fluid to maintain a low temperature, and as such, events resulting in their delivery being prevented have a significant effect on its operation. The valve box component has a relatively higher number of critical failure modes, because owing to its multiple functions and many moving parts, it can fail in more ways than a purely passive component, such as the helium cooled pumping surfaces.

In order to examine the response of the scheme to these faults, the non-linear simulation model presented by Wright et al. [2012] was adapted so that the faults could be injected into the Matlab simulation. The outputs of the fault simulations were passed to the condition monitoring scheme (also implemented in Matlab) to test its effectiveness at identifying and isolating the simulated faults.

A total of ten faults were simulated, each with a different failure mode. The most severe (and the most difficult to repair) and the most likely faults were chosen to be simulated, because these best demonstrate the scheme's utility. In addition to these, any relevant faults which have occurred recently (within the last five years) were also selected, as these were/are of particular interest to the engineers responsible for the operation of the plant.

The following faults were selected:

Fault	Description
Transmission Line:	
Leak Fault	Cryogenic fluid leaking from the transmission line to atmosphere
Ice Fault	Cryogenic fluid contamination or impurity resulting in ice formation in the transmission line
Insulation Fault	Compromised vacuum jacket resulting in deteriorated thermal insulation between the cryogenic fluid and atmosphere
Broken Valve Stem	Mechanical damage of the inlet valve stem resulting in loss of valve control and debris passing into the transmission line
Return Line:	
Worn Valve	Mechanical wear of return line valve resulting in loss of fine valve control
Leak Fault	Loss of helium gas to atmosphere or vacuum jacket from the return line
Insulation Fault	Compromised vacuum jacket resulting in deteriorated thermal insulation between the helium gas and atmosphere
Helium Panels:	
Heat Fault	Unplanned heat load on pumping surface owing to reduced thermal insulation or gas leak into NBI vacuum space
Leak Fault	Cryogenic fluid leaking from the heat exchanger to atmosphere or NBI vacuum space
Manifold Blockage	Ice formation on the return manifold owing to cryogenic vapour impurity

Table 3. Table of Selected Faults

4.2 Example Fault Simulation

The result of one fault simulations is shown in Fig. 5 to Fig. 8, as an example. The condition monitoring scheme produces ten residuals, but the first six are omitted here as they remain unaffected by this fault.

In this example simulation a helium panel thermal fault was simulated. It had a magnitude of 20W and began at $t = 20$. A thermal fault like this could be caused by an atmospheric gas leak into the NBI vacuum space, by a fault in the panel supports,, by a fault in the panel radiation shield, or by a warm body coming into contact with the helium panels. In Fig. 5 to Fig. 8 the residuals signals are plotted in blue, and the upper and lower thresholds of detection are plotted in red. The residual signals correspond to the helium panel termination gas flow rate, the helium panel liquid level, the helium panel gas/liquid pressure, and the helium panel gas/liquid temperature, respectively.

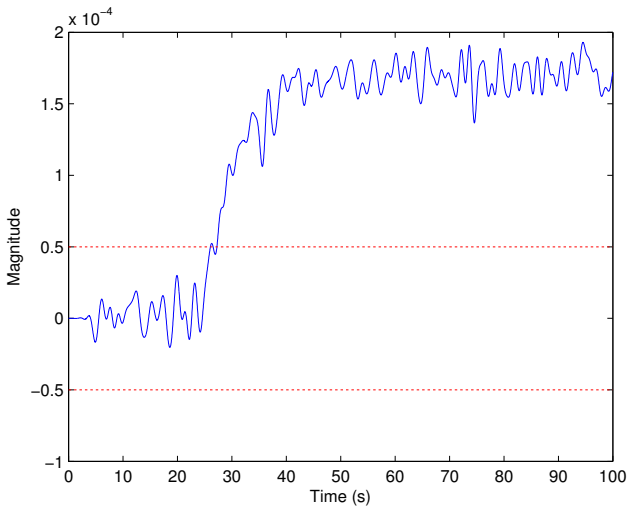

Fig. 5. Residual seven, R7 - Panel gas flow rate

Residual R7 crosses the upper detection threshold at $t \approx 26$, at which point it is flagged. Residual R9 similarly crosses its upper threshold, and is flagged at approximately the same time. At $t \approx 24$, residual R8 crosses the lower detection threshold and residual R10 crosses the upper threshold. Both R8 and R10 are flagged accordingly. As such, the thermal fault is detected at $t \approx 24$ and correctly isolated at $t \approx 26$. The final isolation flag status for this simulation is presented in Table 4.

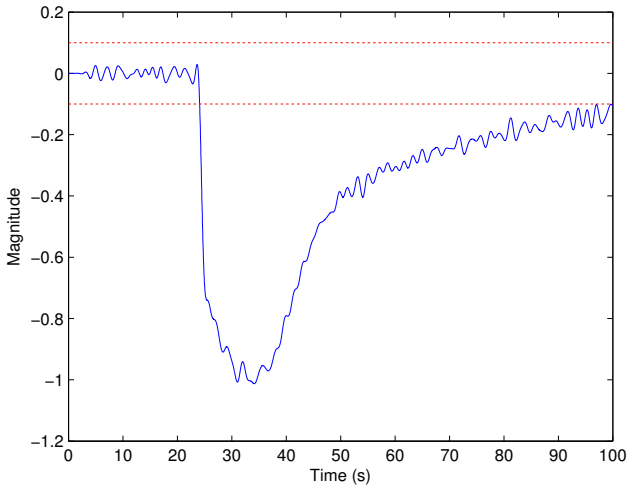

Fig. 6. Residual eight, R8 - Panel liquid level

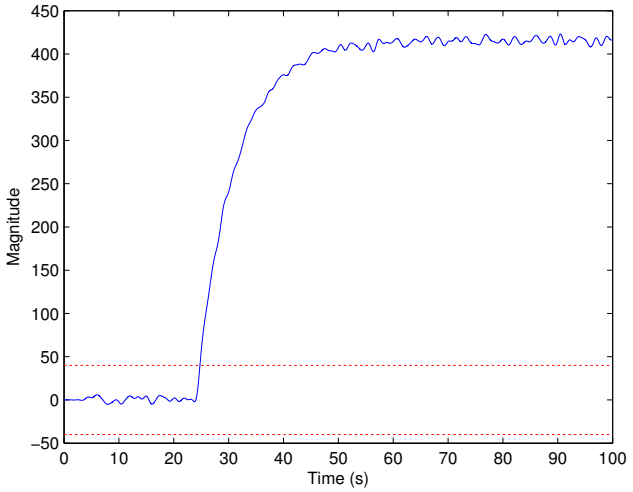

Fig. 7. Residual nine, R9 - Panel gas pressure

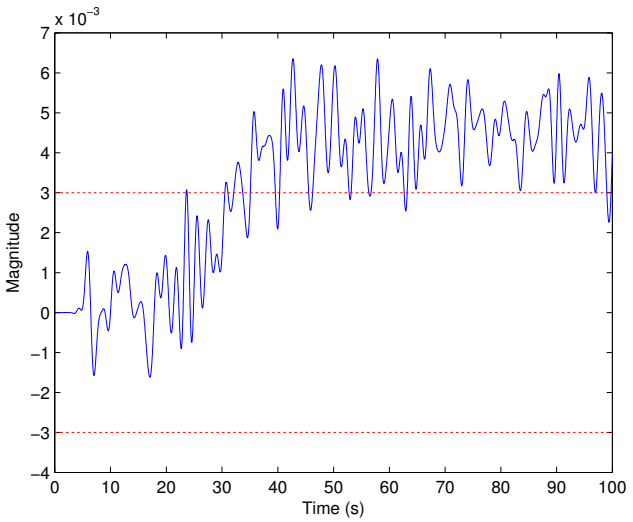

Fig. 8. Residual ten, R10 - Panel temperature

	R1 to R6	R7	R8	R9	R10
Upper Threshold		✓		✓	✓
Lower Threshold			✓		

Table 4. Isolation matrix - Thermal fault

4.3 Results

All ten faults were simulated in the same manner as the example above. Using this scheme, all ten faults caused at least one residual to deviate across a detection threshold and were thus detected by the condition monitoring scheme. Six of the ten faults produced a unique combination of threshold crossings (a signature), allowing them to be isolated. Four of the ten faults shared a signature.

A full list of all the thresholds flagged during the fault simulations is presented below in Table 5. In the right column, the thresholds are referred to by their index number and trailed by a plus or minus symbol (or both) to denote which of the two thresholds were flagged.

Fault name	Flagged Thresholds
Supply Line Leak Fault	R1+-
Supply Line Ice Fault	R1+-
Supply Line Thermal Fault	R1+-
Supply Line Valve Fault	R1+-
Return Line Valve Fault	R4+- R5+ R6+
Return Line Leak Fault	R4- R5- R6-
Return Line Thermal Fault	R5+ R6+
Helium Panel Thermal Fault	R7+ R8- R9+ R10+
Helium Panel Thermal Fault	R8+
Helium Panel Thermal Fault	R7- R8+- R9+ R10+

Table 5. The thresholds flagged after each fault simulation

In Table 6, below, the final detection and isolation times for each fault are listed. The detection time is defined as the length of time between the injection of the fault and at least one threshold being flagged. The isolation time is defined as the length of time between the injection of the fault and all of the associated thresholds being crossed and flagged.

Fault name	Detection (s)	Isolation (s)
Supply Line Leak Fault	4	4
Supply Line Ice Fault	3	4
Supply Line Insulation Fault	5	10
Supply Line Valve Fault	10	13
Return Line Valve Fault	10	35
Return Line Leak Fault	10	20
Return Line Insulation Fault	105	180
Helium Panel Thermal Fault	4	6
Helium Panel Thermal Fault	5	5
Helium Panel Thermal Fault	6	70

Table 6. Approximate detection & isolation times

The detection and isolation times for all of the faults are significantly faster in these simulations than they currently are in practice, where they are diagnosed by manual inspection of the process variables. The supply line faults all share the same signature, however, indicating that while the scheme is effective at detecting supply line

faults, it is less effective at discriminating between them. The return line insulation fault is a relatively slow acting fault, and this is reflected in its detection and isolation time.

The three performance criteria for this scheme were as follows: detection rate, detection time, and isolation time. In all three areas the scheme met a satisfactory standard, and for the last two criteria, the simulation showed an improvement over current practice.

5. CONCLUSION

The objective of this research was to demonstrate the usefulness of the condition monitoring scheme as a tool for diagnosing faults in fusion engineering setting. By simulating the operation of the scheme in Matlab/Simulink, and testing its response to a selection of ten faults, it has been demonstrated that the scheme can be used to rapidly detect and diagnose faults on the neutral beam injector cryopumping system. Given that even small faults on this device can disrupt the fusion experiment as a whole and have historically been difficult to diagnose and repair, these results imply that this scheme could be useful for the engineering staff responsible for its operation.

It is important to note, however, that the total range of faults that could affect the cryopump extend beyond those examined here. Future work should begin with testing the scheme's effectiveness at detecting a wider range of faults, including those that affect the pump boundary conditions.

Future fusion experiments, in particular the ITER experiment (Tesini and Palmer [2008]), will be increasingly reliant on remote and robotic maintenance activities. The challenge of maintaining fusion devices will increase accordingly. Given the typical high degree of instrumentation required for fusion experiments and the amount of real-time data that is available, there is significant scope for the application of model-based condition monitoring. It is hoped this research project will support the case for the further uptake of these techniques for nuclear fusion applications.

REFERENCES

A M Bradshaw, T Hamacher, and U Fischer. Is nuclear fusion a sustainable energy form? *Fusion Engineering and Design*, 86(911):2770–2773, 2011.

R King, C Challis, and D Ciric. A review of JET neutral beam system performance 19942003. *Fusion Engineering and Design*, 74(1-4):455–459, November 2005.

Mathworks Inc. MATLAB Documentation (Online). http://www.mathworks.co.uk/help/matlab/, 2013. Accessed: 2013-10-1.

M Rausand and A Høyland. *System Reliability Theory: Models, Statistical Methods, and Applications*. John Wiley and Sons, 2004. ISBN 047147133X.

A Tesini and J Palmer. The {ITER} remote maintenance system. *Fusion Engineering and Design*, 83(79):810–816, 2008.

V Venkatasubramanian. A review of process fault detection and diagnosis Part I: Quantitative model-based methods. *Computers and Chemical Engineering*, 27(3): 293–311, March 2003.

N Wright, R Dixon, and R Verhoeven. Model of a Fusion Cryopumping System for Condition Monitoring. *Proceedings of UKACC International Conference on CONTROL, Cardiff, UK, 3-5 September 2012*, pages TueA1–3, 2012.

N. Wright, R. Dixon, and R. Verhoeven. Condition monitoring for a neutral beam injector cryopumping system. *Fusion Engineering and Design*, 88(68):1236 – 1239, 2013. Proceedings of the 27th Symposium On Fusion Technology (SOFT-27); Lige, Belgium, September 24-28, 2012.

Z Zhou, M Zhuang, X Lu, L Hu, and G Xia. Design of a real-time fault diagnosis expert system for the EAST cryoplant. *Fusion Engineering and Design*, 87(12):2002–2006, 2012.