A Phased Mission Approach to Fault Propagation

by

Michael D. Lloyd, MEng

DOCTORAL THESIS

Submitted to

The University of Nottingham

for the degree of

PhilosophiæDoctor

July 2014

A Phased Mission Approach to Fault Propagation

Abstract: On complex systems with built-in health management systems, the faults diagnosed during a mission can number in the tens of thousands. When these faults are evaluated, many are found to be false. This work has, therefore, developed a technique by which diagnosed faults can be evaluated using known system data and a system modelling technique to automatically verify their legitimacy.

Petri nets (PNs) were selected as the modelling technique since they allow systems to be modelled in a componentistic and flexible way, that still provides a high level of accuracy. The PN technique was used to model the performance of an experimental facility, the BAE Systems fuel rig, which represents an aircraft fuel system. A wide range of faults were injected into the system and sensor outputs were recorded. By comparing the sensor outputs from the fuel rig to the PN predicted system behaviour, the faults were assessed as either genuine or false. The standard deviation technique is used as part of the comparison process as it provides a high level of detail with low computational requirements. A piece of software was written to automate the PN simulation and comparison of the output data.

The ability of the overall technique to verify diagnosed faults was demonstrated by a thorough consideration of failure modes in the fuel rig system. First and second order faults were evaluated and the results showed that the technique was very successful at identifying both genuine and false faults. Some issues were evident when hidden failures were considered and faults which were revealed for only short periods of time were injected.

The PN technique was also successfully used to model the behaviour of the fuel system of the Airbus A340 aircraft. This system contains a higher level of complexity in terms of both design and operation compared to the fuel rig. The behaviour of the system in normal operation was modelled to replicate that described in literature and a number of first and second order faults were modelled. The PN predicted behaviour of the fuel system in the presence of these faults matched well with that expected.

The PN technique can be used to obtain the output of sensors when failures occur, and such information can be used in the process of system design. An approach is presented by which a sensors value can be calculated and used to select sensors in a system. The technique considers the change in the value measured by a sensor as a result of faults for single sensors and their pairs.

Acknowledgements

Firstly I would like to thank my supervisors, Prof John Andrews, Dr Rasa Remenyte-Prescott and Dr Dovile Rama, whose longstanding guidance and support has been the cornerstone of this project. Their encouragement and efforts have been hugely influential and the quality of my thesis is a testament to their professionalism and expertise.

The project sponsor BAE Systems has kindly provided access to their facilities and the support of several staff members. In particular I'd like to recognise Dr John Pearson, whose technical advice, input and humour was invaluable to the success of the project. Pete Hubbard has provided countless hours of support with the fuel rig and has gone above and beyond to help wherever possible. Paul Thorley for his work co-ordinating the project, Tony Martin for his insight and poor quality golfing opposition and Ken Astley for his constantly positive outlook also deserve thanks for their efforts.

From an academic perspective I'd also like to thank The University of Nottingham for hosting and the EPSRC for funding my PhD studies. The guidance of Dr Chris Harvey in the field of everything, but in particular C++ programming, will not be forgotten.

For the past eight years I've enjoyed the experience of living and studying in Loughborough. To the lifelong friends I have made, all of whom have offered their unique forms of support and advice, I thank you. More people deserve to be mentioned than is possible but to former housemates, Loughborough Hockey teammates, the Cumberland Club, the BSG and everyone else, may you continue to walk on water.

The challenges of completing my work have been more enjoyable in the company of those from The University of Nottingham pursuing the same goal - Josh, Alan, Cinzia, Dermot and Matt. I will fondly remember our time spent together and that on a cricket pitch with Dr James Grenfell, Dr Adam Clare and Dr Matthew Byrne.

Finally I'd like to acknowledge my family. The successful completion of this work is due, in no small part, to their constant support over the last four years. To Nick, my ever dependable brother, and Steph, my greatest adversary and favourite sister, you both challenge and entertain me more than you could ever know. To my grandmother, your love and desire to find the best of people in life continues to influence and inspire me every day. Every candle you lit was worth it. To my parents, you have provided me with the advice and guidance to deal with every problem I have ever encountered. I cannot thank-you enough. For this I am eternally grateful and this thesis is dedicated to you all.

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

1.1 Health Management System Background

Complex systems can only fulfil their primary functional goals when they are operational; aircraft and power stations are two examples of such systems. The operation of these systems often involves an inherent amount of risk and therefore their safe operation must be closely monitored at all times. Furthermore the lifecycle costs associated with complex systems are often very high and it is of interest to the system operators to minimise these costs while also protecting their assets. To achieve these aims, health management systems are designed as part of complex systems to monitor a large number of sensors and built-in test (BIT) outputs in order to identify any unexpected system performance.

Sensors and BITs track the behaviour of variables within sub-systems and components, for example an aircraft's fuel tank level or a nuclear power plant's reactor temperature may be monitored. If a fault occurs on a complex system which causes an unexpected change in the system performance, health management systems attempt to diagnose the fault from the sensor and BIT outputs in order to generate an 'arising'. An arising is a record of the fault that has been diagnosed in the form of a fault code, as well as the time of the diagnosis. Every arising is then passed from the health management computer to a manual operator. This enables action to be taken by the operator, which in serious situations could prevent the fault propagating through the system and cause widespread damage. If faults are less severe and are captured prior to causing damage, preventative maintenance can also be planned in advance allowing it to be carried out when a system is scheduled to be out of service thereby minimising any loss of operation. In this sense, it is necessary for health management systems to provide both wide ranging and detailed

coverage of complex systems. This ensures that less severe component failures, which may cause an immediate but minor change to the system performance, are diagnosed promptly as these faults may become more severe or contribute to more critical faults occurring in the future.

While health management systems perform a key role in ensuring the safe operation of complex systems, they are known to generate a vast number of arisings even in normal system operation. On one type of military aircraft the number of arisings generated in a standard 6 hour flight is in the order of thousands. It is known from the flight data that the majority of these arisings are 'false positives'. Current techniques only allow a small number of these false arisings to be identified. Due to the high quantity of arisings, using a process of manual verification to identify genuine arisings is both overly time consuming and highly inefficient. The effect of high volume arisings is even greater on systems where there isn't an operator in the command loop. On manned systems, the occurrence of critical arisings can be manually verified by a human operator. However, on unmanned systems, such as increasingly popular unmanned aerial vehicles (UAVs), there is no human operator available to verify the arisings generated in-flight. As a result the operation of the aircraft would have to be terminated. An increase in aircraft downtime to verify potentially false arisings not only results in increased maintenance costs, but will also limit the capabilities available to the aircraft's operator. The operators of any form of unmanned complex system would experience similar cost increases and operational limitations.

Accurate sub-system and component monitoring is particularly important and yet challenging to provide in complex systems as many operate in several unique operational phases within a single mission. This presents a challenge as the behaviour of the system can vary between the different operational phases. On an aircraft for example the phases may include taxi to runway, take-off, climb, cruise, descend, land and taxi to gate. The behaviour of different sub-systems can vary between these phases and a suitable monitoring technique must be able to account for this variation in behaviour. Operational phases can also be differentiated by time or system configuration.

Phased mission monitoring is further complicated by the fact that a fault may occur and yet remain hidden in one phase but then be revealed in subsequent phases. Also, the performance of the system may alter when multiple faults occur, or become revealed due to a phase change for example. The behaviour of a system in every operational phase and

at any phase change must therefore be carefully monitored. This multifaceted behaviour of complex systems helps to illustrate one of the reasons why health management systems can generate so many false arisings.

1.2 Built-In Tests

Complex systems use BITs to monitor the behaviour of variables in sub-systems and components. BITs are capable of generating arisings based on the variable outputs that they monitor. Any faults diagnosed by a BIT are passed to the health management system to be merged with any other arisings.

An aircraft is an example of a complex system that makes use of BITs. On the BAE Systems Eurofighter Typhoon aircraft three types of BIT are used: powered built-in tests (PBITs), continuous built-in tests (CBITs) and interrupted built-in tests (IBITs). Each BIT performs a different function within the overall health monitoring process. PBITs are used when the aircraft is started-up to identify any components that have failed prior to power-up or exhibit a fault on power-up. PBITs are both thorough and wide-ranging to ensure that a detailed analysis of sub-systems and components is conducted prior to increased operational demands. CBITs are run throughout the operational mission of the aircraft and monitor variable behaviour in every operational phase. CBITs are active in the background of normal operational system activity and they do not offer the coverage or detail provided by PBITs. An IBIT is an on demand test that is initiated by a pilot or maintenance engineer usually at the end of an aircraft mission. IBITs provide a level of detail and coverage similar to that provided by PBITs. The combination of the three BITs described above provides an example of how detailed coverage of a complex system throughout its operation can be achieved.

BITs are modular stand-alone units that are often purchased through third party suppliers to supplement health management systems. The BIT is installed with the appropriate sub-system or component to form a line replaceable item (LRI). A BIT is specific to the LRI that it is installed on, but a single LRI can contain multiple types of BITs. A single LRI on the Typhoon aircraft could therefore contain a PBIT, a CBIT and an IBIT or any combination of these. A single LRI will use the same set of sensors as inputs to all of its BITs. The manner in which these inputs are dealt with and analysed is specific to each BIT.

On many systems if a BIT generates an arising, the entire LRI is replaced immediately. This procedure is not only expensive but in many cases it may be unnecessary if the fault identified by the BIT is shown to be false. However, with no process in place to verify these arisings large costs are incurred replacing LRIs on complex systems. It may also be possible that the BIT is faulty and as a result has generated an arising. Replacing the LRI as a result of this would again be unnecessary and costly but is unavoidable using current methods.

1.3 False Arisings

The number of false arisings generated by the health management system of a complex system has a significant impact on the maintenance costs and availability of the system. Identifying the causes of false arisings is therefore important in the overall effort to reduce the number generated.

1.3.1 False Arising Causes

1.3.1.1 Variable Power Supply at Start-Up

Complex systems are very large and may contain hundreds of sub-systems. When these complex systems start-up there may only be a finite amount of power available to bring all the sub-systems and components online. As a result, the power supply within the system will fluctuate. This means that the sub-systems and components within the system will start-up at different times, as and when the power supply becomes available to them. This inconsistent start-up process often creates false arisings.

Consider the flight control system on an aircraft. For redundancy purposes it may contain three identical flight control computers that perform the same tasks. If one computer comes online before the remaining units, the system sensor outputs will cause the health management system to generate an arising stating that the remaining computers have failed. However, these arisings would be false as the offline computers have not failed, they have only yet to finish starting up. Given that this problem can affect all of the electrically driven components and systems on an aircraft, the variable power supply available to an aircraft at start-up could cause a great number of false arisings.

1.3.1.2 Tolerance Levels

A tolerance level is factored into a health management system analysis of sensor outputs and in BITs in order to allow for noise within system variables. Noise will cause system variables, such as flow pressure or flow rate, to be reported by sensors as lower or higher than they actually are. For example, in an aircraft fuel system although a section of pipe may be operating at a constant flow rate, noise may cause the reported flow rate to vary at levels above or below this constant value. Health management systems and BITs are designed to apply tolerances to account for this. However if the tolerances are overly restrictive, noise could cause variables to exceed these tolerances causing an arising to be generated. If noise is the only cause of this arising occurring, the arising can be classified as false.

It would be possible to reduce the number of false arisings generated due to tolerance levels by relaxing the tolerances on the system variables. However, this would increase the likelihood of a genuine fault occurring and going undetected by the health management system or a BIT. Such a situation creates a 'true negative', which can be very dangerous. On complex systems, many of which are safety critical, narrow tolerances are applied to ensure that all faults are identified, even if it results in an increased number of false arisings.

The issue of false arisings generated due to tolerance levels is particularly prevalent on the start-up of complex systems. At this time many system variables exhibit significant changes in values, as related components come online and the system encounters a variable power supply. On complex systems this results in a large number of false arisings being generated, as tolerance levels are frequently exceeded. Once the system has completed its start-up processes and settled into its standard operating behaviour, such false arisings are less likely to be generated.

1.3.1.3 Sensor Failures

False arisings can be generated as a result of a fault with sensor equipment. If a sensor experiences a failure, the inputs to the health management system and BITs will fail to represent the true state of the component. This has the potential to cause multiple false arisings. The occurrence of such a sensor or instrument failure would effect every BIT type which relies on that sensor output.

The failure of a sensor also has more serious repercussions if a genuine fault of the related component occurs. The effect of the fault will not be represented in the output of the sensor and neither the health management system nor a BIT may be able to diagnose the presence of the fault.

1.3.1.4 BITs

There are a number of ways in which BITs can cause false arisings to be generated on complex systems. Any fault within a BIT unit can cause false arisings to be generated. In a similar fashion to sensor failures, hardware or software faults within a BIT can cause incorrect outputs to be produced. In the right conditions these can cause false arisings to be generated and passed to the health management system.

A second source of false arisings from BITs can be attributed to their tolerance levels - an issue which was identified previously.

The acquisition of BITs from third party manufactures creates a further source of false arisings in complex systems. BITs designed by third party companies are often not application specific and are subject to testing and analysis by the parent company only. As these BITs are purchased from suppliers, it is not possible for operators of complex systems to change or modify the units. When these BITs are installed on complex systems, their behaviour and fault testing performance can differ from the expected behaviour including the generation of false arisings. The occurrence of these false arisings can be reduced if the suppliers of BITs are made aware of falsely generated arisings. The supplier can then make appropriate adjustments to prevent the BIT from generating the unwanted arisings. However, this process is dependent on the supplier being made aware of and verifying the arising as being false using its own internal testing mechanism. If the manufacturer cannot verify a fault as being false in the equivalent operating conditions, the BIT will not be altered and the false arising will continue to be generated.

1.3.1.5 Variable Mission Types

Sub-system and component monitoring techniques are currently unable to account for the type of mission being undertaken by a complex system. As a result, a number of false arisings are generated. Consider a military aircraft undertaking a training mission. As no warfare-based phases of operation are conducted, a number of weapon and combat-based components may be omitted from the aircraft set-up process. As the health management

system and BITs cannot account for such information they will generate a number of false arisings stating that these missing components have failed.

1.4 Research Aim

The number of arisings generated by complex systems poses a problem to system operators who want to operate these systems safely but also effectively. Abandoning or delaying missions to investigate every arising is not practical. Missions are therefore undertaken in spite of unchecked arisings being present. While it is known that the majority of these arisings are false, there is no effective method available to identify which ones.

It has been shown that while false arisings can be generated throughout a system's operation/mission, the majority of false arisings are generated at start-up. On one military aircraft, of all the arisings generated during a mission almost 70% occurred during the start-up phase. Developing a technique to analyse arisings generated in the start-up phase of complex systems and identifying them as either true or false would therefore reduce the total number of arisings significantly.

The aim of the project is to develop a technique that can verify true arisings in the event of a component or system failure and filter those arisings that can be proven to be false. This technique is primarily concerned with those arisings that are generated in the start-up phases of a mission.

1.5 Research Objectives

In order to verify the accuracy of fault arisings it is necessary to replicate the behaviour of faults using a system model. There are a large number of techniques which have been devised, developed and utilised to model a full range of system types over many decades. Identifying the most suitable technique is the first stage of this research. Converting the modelling technique processes into a piece of software and testing it on a range of system types is also necessary in order to achieve the project aim. A number of objectives have been listed to achieve this project aim:

- Review and evaluate existing system modelling techniques. Demonstrate their capability through application to a theoretical system.

- Identify the most suitable system modelling technique and represent its behaviour in a software program. Use a further theoretical system to verify accuracy of software.
- Apply the chosen modelling technique to a physical system.
- Review existing fault verification techniques and through application to the physical system identify the most suitable technique(s) for project. Create a demonstrator that uses the chosen fault verification technique to verify fault arisings generated from the physical system using the physical system and system model outputs.
- Use the system modelling software and fault verification technique to demonstrate how fault arisings on the physical system can be verified as true or false. Consider a phased mission of the system.
- Demonstrate application of the technique on an industrial scale system
- Identify further areas where the modelling software and fault verification demonstrator can be used to improve the design/operation of systems.

CHAPTER 2

Literature Review

2.1 Introduction

The purpose of this literature review is to identify and evaluate existing system modelling techniques that may be suitable for fault propagation analysis. Through application to a simple, theoretical system the most appropriate modelling technique for this project will be identified.

The necessary characteristics required of the modelling technique for this project would include those listed below.

- Ability to model complex systems with accuracy
- Ability to deal with phased missions
- Offer a high level of flexibility

There are numerous techniques available to model systems [1] and consideration is given to three of the most developed and widely used; the decision table, digraph and the petri net (PN) techniques. The theoretical system that the modelling techniques will be applied to is a domestic hot water system. The system is considered in greater detail in Section 2.2.3.1. The most promising technique will also be used to model a tank level control system. This system is described in Section 2.4.4. Considering a further system will test the robustness of the modelling technique and help to confirm the observations made from its use when modelling the hot water system.

2.2 Decision Tables

2.2.1 Technique Details

Any process, be it chemical or manufacturing for example, is subject to a set of conditions that will determine its resulting action. A decision table is a tabular representation of these condition sets and the resulting actions. A complete decision table will list every relevant condition set, including failures, and the respective actions of a component. When modelling a physical system using the decision table modelling technique, every component is modelled by a decision table. A component's inputs and internal modes would be considered as a set of conditions while the output is the resulting action.

Consider a simple gas pipe with the conditions 'Gas Input' and 'Internal Mode'. The input condition can provide either supply (S) or no supply (NS) and the internal mode can either exhibit no blockage (NB) or a blockage (B). Depending on these conditions the resulting action 'Gas Output' will provide either a supply (S) or no supply (NS). The gas pipe has been modelled using a decision table in Table 2.1.

Table 2.1: Gas pipe decision table

Row Number	Gas Input	Internal Mode	Gas Output
1	NS	B	NS
2	NS	NB	NS
3	S	B	NS
4	S	NB	S

It can be seen from row four of Table 2.1 that a gas supply will only be output from the gas pipe when there is a supply at the input and there is no blockage in the pipe. All other combination of conditions will produce the same output - no supply. Although a decision table shows a range of condition sets and resulting actions, Table 2.1 demonstrates that only one row of a decision table can be true at any one time.

The gas pipe decision table shown above is one of the most basic that could be developed. A more complex component with multiple input and output variables would produce a much larger decision table. It is not uncommon for a detailed decision table to contain several columns and hundreds of rows. In order to consider only the minimum number of conditions to produce all the system outputs, a simplification rule is applied.

To simplify a decision table there must be at least two rows with matching outputs. These rows must also have identical values in all but one of the other columns. In this remaining column, if all the possible states of the condition are shown in the otherwise identical rows, then the rows can be simplified to one with the value in the non-identical column being set to a dash (-). This indicates the value of this variable does not matter, as it has no affect on the output.

In Table 2.1, when there is no gas supply at the input it is irrelevant what the internal mode of the pipe is, there will be no gas supply output. Rows one and two can therefore be combined. A similar outcome occurs when the pipe is blocked. The simplified gas pipe decision table is shown below.

Table 2.2: Simplified gas pipe decision table

Row Number	Gas Input	Internal Mode	Gas Output
1	NS	–	NS
2	–	B	NS
3	S	NB	S

When constructing the decision tables of an entire system it is necessary to consider how the components are linked together. Where two components are linked, all possible outputs from the first component must be considered as unique inputs to the second component. In cases where a component has multiple inputs each must be considered as a condition set. In the decision table all combination of input variable states must then be evaluated. Only in this way will the decision table technique provide the most accurate representation of the system behaviour.

2.2.2 Literature Review

The decision table modelling technique was first developed by Pollack [2]. The technique represented an expansion of Pollack's simpler yet more restrictive truth tables [3]. Truth tables limited the detail with which components could be modelled as any description of a system state had to be in a binary form. Contrastingly, decision tables do not place a limit how many discrete values can be used to describe a system variable [4].

Since their inception decision tables were designed to be able to model any system

type. Their application, therefore, is very wide ranging. Some of the systems modelled by decision tables include process plant systems [5], the residual heat removal system of a nuclear power plant [4] and a steel structure code-checking system [6]. These systems consider a range of components that are electrical, human and mechanical in nature.

One of the main advantages of using decision tables to model systems, as identified by Carpignano and Poucet, is the fact that they are context free [7]. This means that a decision table can be constructed knowing only the component type and its direct inputs. Information regarding its location in the overall system or its "underlying functional structure" is not required. Salem et al [4] also note that this characteristic means component decision tables can quickly account for changes in the system topology. This feature adds further value as it doesn't limit the decision table technique to analysing completed systems; it also allows it to be used as an effective system design tool once the decision tables have been constructed. For these reasons Kelly and Lees consider the decision table modelling technique to be very flexible [5].

As component decision tables are context free, they are not unique to a single system. If the same component is used in multiple systems, the same decision table can also be used. This means a library of component decision tables can be built over time as new components are considered. Salem et al also observe that by continuously building this library of decision tables, the time taken to model systems will continually decrease [4]. This is important because, as a number of authors note [5], [8], creating decision tables remains a manual task which requires significant resources. Kelly and Lees also state that the process of creating decision tables does not suit automation [5]. The reason for this is identified by Salem et al [4]; constructing component decision tables represents the only opportunity an analyst has to input their system knowledge into the model. Having a computer attempt to undertake this task risks key system knowledge being omitted and the system model being inaccurate.

In a review of the decision table technique, Carpignano and Poucet criticise "the simplified representation scheme of the system and component behaviour in which no functional information can be represented" [7]. The decision table technique fails to consider the functional details of a system as it uses a componentistic approach to modelling. The same authors state that techniques using this approach have difficulties modelling a system's global behaviour and control properties because they do not consider the functional role of components in a system. As a result the detail with which a system's true be-

haviour can be modelled using the decision table technique is limited. Other authors have made a similar observation about the poor consideration of control loops. Andrews and Henry, for example, state that an "inability to detect, classify and analyse control loops" is the decision table modelling technique's primary weakness [9]. It is the consideration of only individual system components and not all of the components collectively which leads to this poor consideration of control loops by the decision table technique. Andrews and Henry propose incorporating a positive/negative gain into the decision tables in order to correctly account for control loops.

A further shortcoming of the decision table modelling technique is its inability to account for the reverse propagation of faults. According to Carpignano and Poucet, this is due to the fact that decision tables are unidirectional [7]. As a result when a fault occurs in any one component, only those components downstream of it will exhibit the effects of the failure. Considering a system where flow moves from left to right, downstream components would be those to the right of a failed component. Components upstream, or to the left, of the failed component will not be able to account for the failure even though, in reality, a fault may propagate in both directions in a system. For this reason how components have been linked in the model will also significantly affect the propagation of any faults.

When a system is modelled for the first time, it is important to check that the model does not contain errors. Identifying modelling errors as early as possible is always preferable, however, it may not be until further analysis is undertaken that these errors become apparent. Having to correct errors at a later stage will take longer and require more resources, as the analysis process may have to be repeated. Given the above and the fact that decision tables consider components individually, Carpignano and Poucet state that system models produced by the tables are difficult to verify "for congruence and completeness" [7]. However many authors [10], [11], [12] take a different opinion and consider the technique easy to verify and validate. Vanthienen et al consider decision tables effective because their structured nature makes anomalies easy to identify [11]. Also unlike the PN technique, for example, decision tables do not have to be converted into their operational form in order to be verified. This means errors can be identified and corrected earlier in the decision table modelling process.

Majdara and Wakabayashi have recently attempted to overcome some limitations of the decision table technique by combining it with a number of functional modelling features

to create a single technique that doesn't ignore "any necessary details" [8]. The developed method is based on the decision table technique; components are modelled by both function and state transition tables. Function tables are the same as decision tables and can be simplified in the same manner, as can the state transition tables. The authors consider the state transition tables to be the tabular form of state diagrams. They are used for those components that have multiple operating states, i.e. a valve as it can be open or closed. The contribution from the functional modelling techniques comes in the form of considering the physical properties of flow through the system. Any physical property can be incorporated into the model, i.e. temperature, and these are defined at the start of modelling. Every component that can vary the defined physical properties has an extra column added to its function and/or state transition table that shows how the property is affected by the component. While the authors consider application of the developed technique a success, it is not clear how the physical properties are handled by the technique. The technique also continues to fail to account for reverse propagation and would require an even larger library of results to store the two types of tables.

2.2.3 Application to Hot Water System

2.2.3.1 System Description

The purpose of the hot water system is to provide a supply of hot water on demand from a user. To model this behaviour the system operation has been split into two phases. In the first phase the water is heated while the tap remains closed. In the second phase, the tap is opened and a supply of hot water is output from the system. A schematic diagram of the hot water system is shown in Figure 2.1

In phase one the water pipe is filled with water while the tap is closed. The water is heated by the pilot light which is constantly lit and receives a gas input from the supply. A feedback loop regulates the temperature in the water pipe using a temperature sensor, controller and a control valve in the gas pipe. The valve remains fully opened until the water temperature reaches its upper limit, at which point the valve is closed. The water in the pipe is then allowed to cool until its temperature falls to its lower limit. At this point the feedback loop opens the control valve and the heating process is repeated.

In phase two the tap is opened and hot water leaves the system. The flow of hot water leaving the system is continuously replaced by cold water entering the system. This water

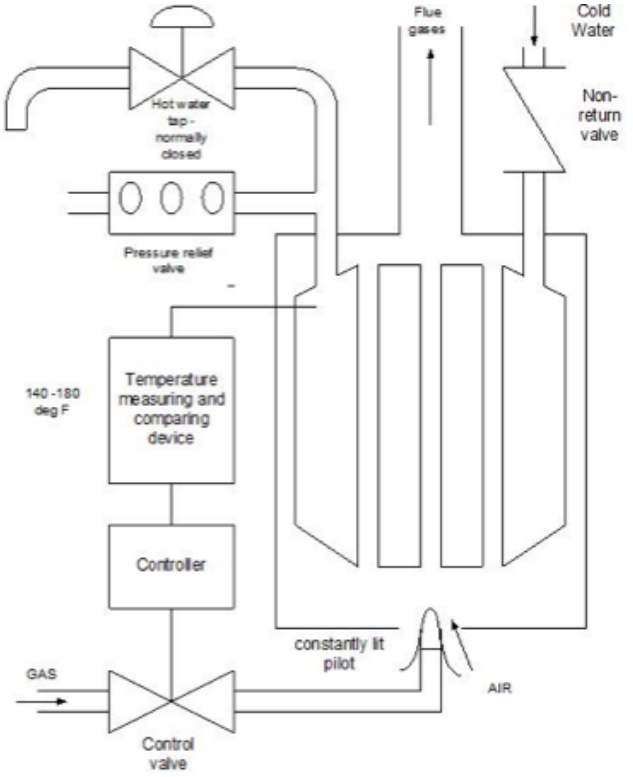

Figure 2.1: Hot water system

is heated by an unbroken gas supply to ensure a constant supply of hot water is provided to the tap. Before entering the water pipe, the pressurised water supply passes through a non-return valve (NRV). This valve is in place to prevent reverse flow should the hot water system become over-pressurised. If an excessive pressure does build up in the system it will be outlet through the pressure relief valve.

2.2.3.2 Phase 1 Component Decision Tables

The phase 1 decision tables of the control valve and water pipe components are shown below. These decision tables provide an example of how the hot water system has been modelled using the decision table technique. The decision tables of all the remaining system components are shown in Appendix A. The decision tables shown below represent the steady state behaviour of the hot water system components. It is necessary to consider the steady state behaviour, as decision tables cannot model the dynamic performance of components. The decision tables have been simplified where possible.

Table 2.3 presents the control valve decision table. The possible gas inputs and outputs

are no supply (NS) and supply (S). The controller input will be either signal to open (SO) or signal to close (SC). The potential internal modes are working (W), stuck open (StO) and stuck closed (StC).

Table 2.3: Control valve decision table

Row Number	Gas Input	Controller Signal Input	Internal Mode	Gas Output
1	NS	–	–	NS
2	–	SC	W	NS
3	–	–	StC	NS
4	S	SO	W	S
5	S	–	StO	S

From Table 2.3 it can be seen that a gas output supply requires a gas input supply and the control valve to be open. The control valve will be open if it is in a working state and receives the correct input signal or if the control valve is stuck open. All other condition states will produce no output supply.

Table 2.4 details the water pipe decision table. The heat and water input states will either be no supply (NS) or supply (S) while the internal mode will be one of secure (S), leaking (Lk) or ruptured (R). The output states describe the temperature, pressure and volume of the water in the pipe as either zero (0), low (L), normal (N) or high (H). A zero output can only occur when the pipe is ruptured and therefore represents atmospheric conditions.

It is assumed that the water supply to the system is pressurised. Row 4 of Table 2.4 shows that when there is a water input but no heat input the pipe pressure will be normal but the pipe temperature low. In the presence of a leak, it is assumed that there will be a loss of water and water pressure in the pipe however the system would be able to compensate for any heat lost. It is also assumed that the physical properties of the pipe prevent it from being filled above its maximum level, indicated in Table 2.4 as normal.

2.2.3.3 Phase 2 Decision Tables

In phase 2 the tap is opened to supply hot water. As hot water constantly leaves the system, fresh cold water must be added and heated to replace it. There is therefore a constant flow of water through the system in phase 2. In spite of this change in the

Table 2.4: Water pipe decision table - Phase 1

Row Number	Heat Input	Water Input	Internal Mode	Pipe Pres.	Pipe Temp.	Pipe Vol.
1	NS	NS	–	L	L	L
2	S	NS	–	L	H	L
3	S	S	S	H	H	N
4	NS	S	S	N	L	N
5	S	S	Lk	L	N	L
6	NS	S	Lk	L	L	L
7	–	–	R	0	0	0

system behaviour only two of the component decision tables change as a result; the water pipe and the tap. Table 2.5 shows the phase two decision tables for the water pipe. The codes used to represent the input states, internal modes and output states have not changed from those defined with the phase 1 decision table.

Table 2.5: Water pipe decision table - Phase 2

Row Number	Heat Input	Water Input	Internal Mode	Pipe Pres.	Pipe Temp.	Pipe Vol.
1	NS	NS	–	L	L	L
2	S	NS	–	L	H	L
3	S	S	S	N	N	N
4	NS	S	S	N	L	N
5	S	S	Lk	L	N	L
6	NS	S	Lk	L	L	L
7	–	–	R	0	0	0

The constant addition of cold water to the system in phase 2 means that only if there is no water supply will the pipe temperature reach a high level. At all other times the pipe temperature should not exceed normal.

2.2.3.4 Results

The ability of the decision table modelling technique to accurately model the behaviour of the hot water system and the effect of possible failure modes will now be assessed. By injecting individual failure modes into the system model the resultant symptoms can be determined. These will then be compared to the expected symptoms produced by the physical form of the system. A symptom will be considered as any unexpected and measurable component behaviour or state.

The expected symptoms have been determined theoretically through consideration of how a physical version of the hot water system would react to the occurrence of the failure modes. It is assumed that when a component failure occurs all of the remaining components are working normally.

All of the failure modes listed in the decision tables as internal mode states have been assessed and analysed. Several results of interest will now be considered in greater detail. Consider first the failure mode 'Control Valve Stuck Closed' when the system is in phase 1. This failure mode is injected into the system model in Table 2.3. Propagating its effect through the remaining phase 1 component decision tables a number of symptoms are produced. These are shown in Table 2.6 along with the expected system symptoms.

Table 2.6: Control valve stuck closed phase 1 results

Expected System Symptoms	Decision Table Symptoms
Control Valve Closed	Control Valve Closed
No Gas Output from Control Valve	No Gas Output from Control Valve
No Heat Supply from Pilot Light	No Heat Supply from Pilot Light
Low Water Pipe Temperature	Low Water Pipe Temperature

Table 2.6 shows that the decision table has accurately represented the behaviour of the hot water system in the presence of the control valve fault. All of the expected system symptoms have been found from the decision table model. In phases 1 and 2 the model predicted symptoms of several failure modes match exactly with those expected. However there are a number of failure modes where the expected and predicted symptoms are not the same.

Consider the results of the failure mode 'Pressure Relief Valve Stuck Open' in phase

2. This failure mode is injected into the decision table model from Table A.6 in Appendix A. When the pressure relief valve fails open in phase 2 there will be an above atmospheric pressure output through the valve. As a result the water pipe pressure will fall to zero and there will be insufficient pressure to enable water to flow out of the tap. With no flow out of the tap there will also not be any water flow into the system through the non-return valve. Any pressure increase that would be created through the continuous addition of heat to the system is lost through the pressure relief valve. The predicted and the expected symptoms are shown in Table 2.7.

Table 2.7: Pressure relief valve stuck open phase 2 results

Expected System Symptoms	Decision Table Symptoms
Above Atmospheric Pressure Output	Above Atmospheric Pressure Output
Water Pipe Pressure Zero	
No Output from Tap	
No Flow into NRV	
No Flow out of NRV	

Table 2.7 shows that there are five expected symptoms produced as a result of the pressure relief valve fault but only one has been identified by the decision table technique. The decision table model has correctly identified the above atmospheric pressure output from the pressure relief valve, however, the effect on the water pipe pressure, tap output and flow through the non-return valve has not been captured. This is due to the inability of the decision table technique to model the reverse propagation effect of faults. As there are no components that receive an input from the pressure relief valve the effect of the failure is limited to that component alone. This modelling limitation is seen in the results of several other failure modes.

Eighteen failure modes were considered in each phase of system operation, thirty-six in total. From these failure modes ninety-nine expected symptoms were identified. The decision table technique could only match seventy-four or 75% of these. It also produced twelve symptoms that were not expected. The decision table technique has therefore failed to provide a truly accurate representation of the hot water system in the presence of the failure modes under consideration.

2.3 Digraph

2.3.1 Technique Details

The digraph technique models a system by displaying its process variables on a directed graph. The system behaviour is captured in the model by defining all of the process variables that are present throughout the system as well as any relationships that exist between them. By including all possible relationships between variables, the behaviour of the system in both normal operation and with component failures is accounted for. A system model can be constructed by linking individual component digraphs. Linked components will share at least one process variable and by joining these together a complete system model can be created.

Process variables are represented on a digraph by circular nodes. Every node contains an alphanumeric label that indicates the variable type it represents and its location within the system structure. Typical variables considered by the digraph technique are pressure (P), temperature (T), mass flow (M) and signal (S). Relationships between two variables are represented on a digraph by directed edges, which connect the relevant nodes. The arrow on the directed edge points from the independent variable towards the dependent one. Coupled to every directed edge is a number that represents the strength of the relationship between the linked variables. This number is known as a gain and is the partial derivative of the dependant variable to the independent one. Relationships must be described in the digraph technique using one of five discrete variables; +10, +1, 0, -1, -10. Strong relationships are represented by ± 10 while moderate relationships are represented by ± 1. A zero gain represents the nullification of a relationship. This is commonly used with conditional edges, which require a certain condition to exist in order for the relationship to be active. Only one edge linking two variables can be active at a time.

To demonstrate the digraph features described above, consider a valve connecting two pieces of pipe, as shown in Figure 2.2. The valve can either be fully opened or fully closed.

In the digraph of Figure 2.2 the two nodes, M1 and M2, are connected by two directed edges, one normal and one conditional. As the arrow on the directed edges points away from M1, it is the independent variable. M2 is the dependent variable. The gain in Figure 2.2 can therefore be expressed as:

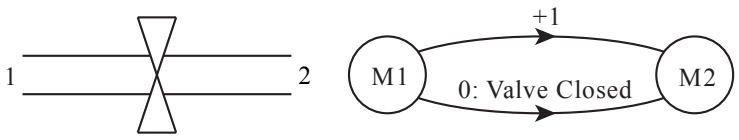

Figure 2.2: (a) Pipe schematic (b) Pipe digraph

$$Gain = \frac{\delta(M2)}{\delta(M1)} \tag{2.1}$$

In normal operation the valve is open. This allows mass flow through the pipe and is represented on the digraph by a gain of +1. If the valve is closed there is no mass flow through the pipe and therefore no relationship between the two nodes. This behaviour is modelled on the digraph by the conditional edge, which has a gain of zero.

The state of any system variable is expressed by the digraph technique using directed edges and the same set of values that describe relationship strengths. In its normal state a variable has a value of zero. If a variable experiences a moderate change in its value, its state will be ± 1. If a variable becomes much larger or smaller than its normal value, its state is expressed as ± 10. It should be noted that the digraph technique limits the value any variable can have to ± 10.

When a component fails or begins to operate in a way that does not represent its normal behaviour, it can create a disturbance within a system. Disturbances are modelled in the digraph technique using nodes with directed arrows. The node details the fault while the directed arrow shows which variable is initially affected by the disturbance. The magnitude of a disturbance is measured using the same set of values used to describe the variable states. A strong disturbance will therefore have a value of ± 10 while a moderate disturbance is given a value of ± 1. The sign indicates whether the disturbance will cause the process variable to take a value greater/lower than it would normally have. Multiple disturbances can affect a single node and disturbances can be conditional.

When a disturbance occurs, its effect may not be limited to the point where it first enters the system - the disturbance could propagate through the system. In a digraph, disturbances pass from independent variables to dependent ones along active directed edges. The effect of a disturbance on a dependent variable can be found by multiplying the magnitude of the disturbance in the independent variable by the respective gain. To show how disturbances propagate through a system consider again the valve section

in Figure 2.2. The digraph in Figure 2.3 includes consideration of the disturbance 'no supply'.

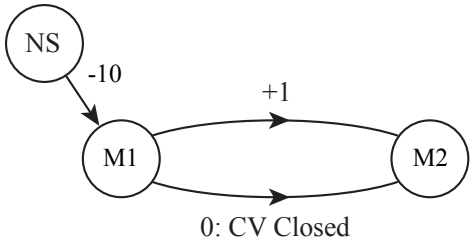

Figure 2.3: Pipe digraph with failure mode

When the pipe valve is open a loss of supply (NS) would require the system to create flow in order to maintain normal operating behaviour. As this is not possible a strong disturbance will enter the digraph at M1. The state of M1 is therefore expressed as M1(-10). As the valve is open and the normal relationship is active the disturbance will propagate to M2. Using the multiplication rule outlined above, disturbance multiplied by gain, the state of M2 becomes M2(-10). This represents a very low amount of flow at point two. Had the valve been closed, the zero gain on the conditional edge would have prevented the disturbance propagating to M2.

2.3.2 Literature Review

The digraph modelling technique was devised by Lapp and Powers as a tool to enable the computer aided construction of fault trees [13]. Its use was first published in April 1977. The digraph technique represented a novel approach to system modelling as it used the functional details of a system to model its behaviour. The advantage of using a functional modelling technique is the added detail with which a system can be represented.

As the digraph technique considers all of the relationships between process variables, it is able to model the reverse propagation effects of faults within in a system. For this reason functional modelling techniques, such as the digraph technique, are considered to be more effective at modelling the global behaviour of systems than componentistic approaches.

When Lapp and Powers created digraphs one of the motivating factors was the need for a systematic modelling technique. A systematic technique is desirable because it provides a methodical, ordered process which is not open to analyst subjectivity. The decision table modelling technique is an example of a non-systematic process as it allows variables

to be represented by any number of analyst defined states [4]. In this case an analysts' subjective input is required to model the system. Comparatively, in the digraph technique the strength of relationships and disturbances is limited to a set of five values. As a result this aspect of the technique is relatively objective. However, identifying system variable relationships and which conditions might have to be met in order for these relationships to become active still leaves the digraph technique open to analyst interpretation. This subjective input prevents it from being considered a truly systematic technique.

Although the set of five values, which are used to define relationships, disturbances and variable states in the digraph technique, provide an objective characteristic it also limits the accuracy with which systems can be modelled. There are only two values, for example, to describe the state of a variable that has exceeded its expected level; moderately high and very high. In an attempt to overcome this limitation Bartlett et al added two values to the set in order to model the fuel system of an aircraft with greater accuracy; ± 5 [14]. The authors consider the application of the expanded technique to be a success and state the extra values allowed partial component failures to be modelled. However using these additional values raises a number of issues. The use of ± 5 to model partial failures creates a conflict with the original definition of ± 1. It is no longer clear what ± 1 represents. Furthermore, as partial failures do not create uncontrollable faults it can be argued that they could have been modelled by ± 1 thereby making the use of ± 5 superfluous. Finally, the inclusion of the additional values makes the overall technique less objective as greater influence is placed on the analyst's interpretation of the system's behaviour.

A further limitation of the digraph technique, related to the limited number of values available to model component behaviour, is identified by Andrews and Brennan [15]. The digraph technique does not allow the expressed state of a process variable to be nil, i.e. 0L or 0L/min. The closest option available is very low (-10). In situations where a blocked pipe cuts off mass flow, using a very low level as opposed to zero immediately introduces a source of error into the model.

Andrews and Brennan have also identified that the digraph technique has issues modelling uncontrollable disturbances [15]. They cite an example where a system containing two feedback loops incorrectly propagates a large positive disturbance through the system. Despite the issues created when dealing with uncontrollable disturbances the digraph technique is still considered by many authors to be effective when dealing with control loops [7], [13], [15]. Many of the same authors also note that control loops are very effec-

tively displayed in digraphs, even when complex systems are considered. This makes them not only easy to identify but also help the analyst to ensure the behaviour of the system has been correctly modelled. In their description of the digraph technique, Andrews and Morgan [16] state that control loops are considered in a structured and effective manner. However, they noted that other researchers [17] had encountered problems when trying to model two-way flow within systems.

Two-way flow can also have a detrimental effect on control loops in a digraph. A control loop is built assuming one direction of flow (+ve). If a control loop sensor records reverse acting flow (-ve), a negative feedback loop will become overall positive. As a result the feedback loop would act to amplify a disturbance, not correct it. This could cause a controllable disturbance to become uncontrollable and fail the system.

Carpignano and Poucet have observed that component digraphs are context dependant; that is their structure is dependent on both the component type and how it is utilised in a system [7]. This means that individual component digraphs are unique to the system they are designed for and cannot be stored in a library of component models. By means of an example consider the non-return valve in the hot water system shown in Section 2.2.3.2. Its digraph model would not be correct if the same non-return valve was used in a cold water system. In the hot water system there is a need to consider the temperature variable within the valve's digraph but this would not be the case in a cold water system.

The digraph technique has previously been used to consider phased missions. However, there are conflicting opinions about its practicality. Andow [17] states that the consideration of a time base would require digraphs specific to each stage and therefore it would be a mismatch to use digraphs to model phased missions. However Allen believes that so long as the timing of every failure is carefully considered as well as "the repair of sleeping failures", digraphs can accurately represent phased missions [18].

Multiple authors have stressed that one of the key benefits of using the digraph technique is the knowledge that is gained by the analyst which can then be used to build in the system model [13], [15], [16]. This level of knowledge could only be gained through a manual technique. Andrews and Morgan state that in a fully automated technique a similarly full understanding of the system's behaviour could not be achieved [16]. This does, however, mean that the process of building a system digraph is a long, resource intensive task. The digraph technique is designed to allow any type of system to be modelled. Some of the systems modelled by the digraph technique include a nitric acid cooling system [13],

a pressure/volumetric regulation system [15] and the process stream of a butane vaporiser [16]. These system models have considered not only a range of mechanical and electrical components but also multiple complex control loop arrangements.

2.3.3 Application to Hot Water System

2.3.3.1 Digraph Topography

Figure 2.4 shows a breakdown of the hot water system and how the component input and output locations have been numbered for modelling by the digraph technique. The dummy tails are used when flow crosses over the system boundaries.

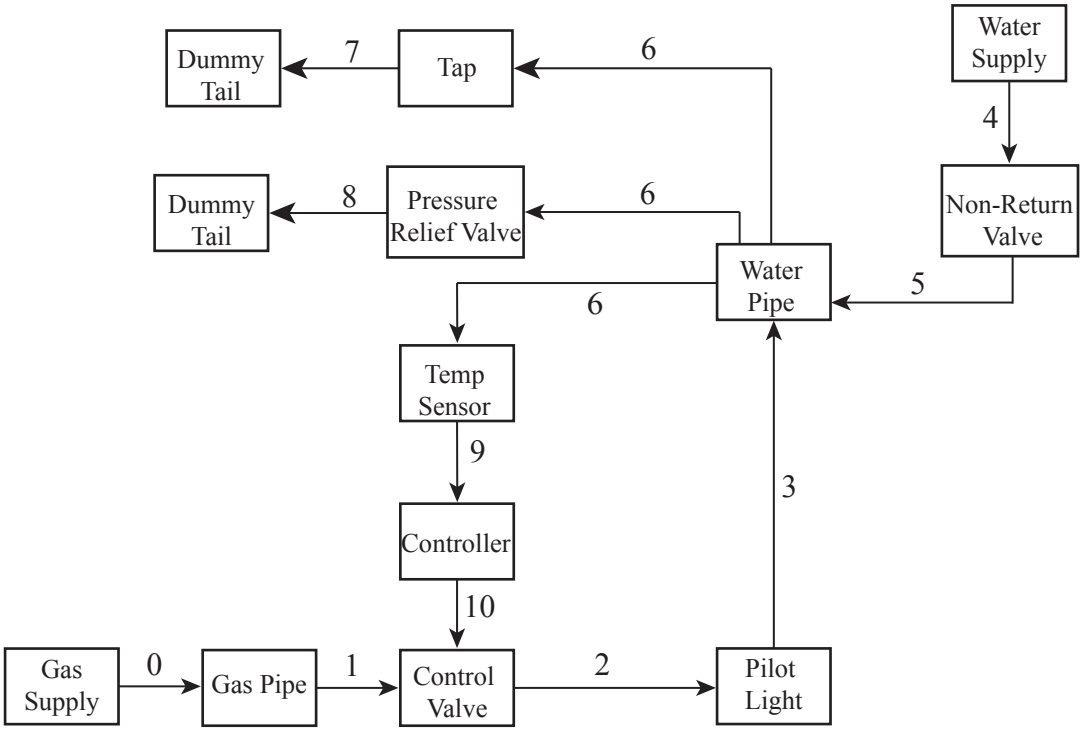

Figure 2.4: Hot water system digraph topography

The control loop in the hot water system can be clearly identified in Figure 2.4. The components on the control loop are the temperature sensor, controller, control valve, pilot light and water pipe.

Code names have been used to represent the failure modes on the component and system digraphs. Table 2.8 lists all of the component failure modes and their respective codes.

Table 2.8: Digraph codes - Component failure modes

Component	Failure Mode	Code
Gas Pipe	No Supply	NS
	Gas Pipe Blocked	P_{01}B
Control Valve	Stuck Open	CVSO
	Stuck Closed	CVSC
	Pipe to Pilot Blocked	P_{23}B
Controller	Failed High	CFH
	Failed Low	CFL
Pilot Light	Failed Off	PLFO
Non-Return Valve	Failed Open	NRVFO
	Failed Closed	NRVFC
	No Supply	NS
Water Pipe	Leaking	WPL
	Ruptured	WPR
Pressure Relief Valve	Stuck Open	PRVSO
	Stuck Closed	PRVSC
Tap	Stuck Open	TSO
	Stuck Closed	TSC

2.3.3.2 Phase 1 Digraph Models

In the construction of any system digraph, consideration is first given to the individual component digraphs. Having considered all the system components, the system digraphs for each phase will be presented. The control valve and water pipe component digraphs are shown below. The remaining hot water system component digraphs are presented in Appendix B.

Figure 2.5 shows the control valve component digraph. The control valve has two inputs; a mass flow from the gas pipe, M1, and a pressure input from the reverse-acting controller, P10. In normal system behaviour these inputs have direct relationships with the output variable M2. However, the conditional edges show that unconventional component behaviour, i.e. control valve stuck, causes these relationships to vary. If the control valve

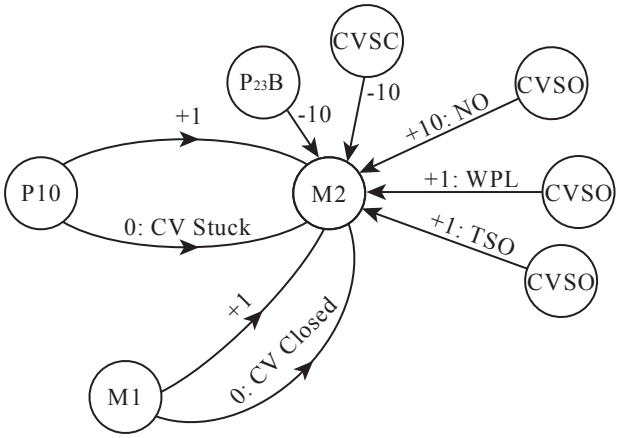

Figure 2.5: Control valve digraph - Phase 1

becomes stuck open or closed it will fail the control loop and introduce a large disturbance into the system, as the control loop cannot correct it. A large disturbance will also be created if the pipe between the control valve and the pilot light is blocked.

If the control valve fails open when there are no outputs from the system (NO), a large positive disturbance will be introduced. However should the control valve become stuck open when there is a leak in the water pipe (WPL) or when the tap is stuck open (TSO) the disturbance will only be moderate. This is because any output from the system will result in cold water being added to replace it. This will reduce the effect of the control valve becoming stuck open hence the smaller disturbance.

Figure 2.6 displays the digraph model of the water pipe component in phase 1. There are two inputs to the water pipe, water through the non-return valve and heat from the pilot light. In order to accurately model the pipe and its behaviour, three process variables are considered at its output - temperature, pressure and level.

A conditional relationship exists between the temperature input (T3) and the pipe pressure (P6). The temperature input will only have an effect on the pipe pressure when the pipe is full of water and when the water temperature in the pipe is very high. As water is incompressible, at all other times the temperature input will not have a measurable effect on the pipe pressure. There is also a conditional relationship between M5 and T6. Only when cold water enters the water pipe will the temperature in the pipe be affected.

If the water pipe ruptures, an uncontrollable disturbance will affect all of the water pipe variables. If the pipe is leaking there will be a loss of water and pressure in the pipe. This disturbance will be controllable because the water supply can replace any fluid and

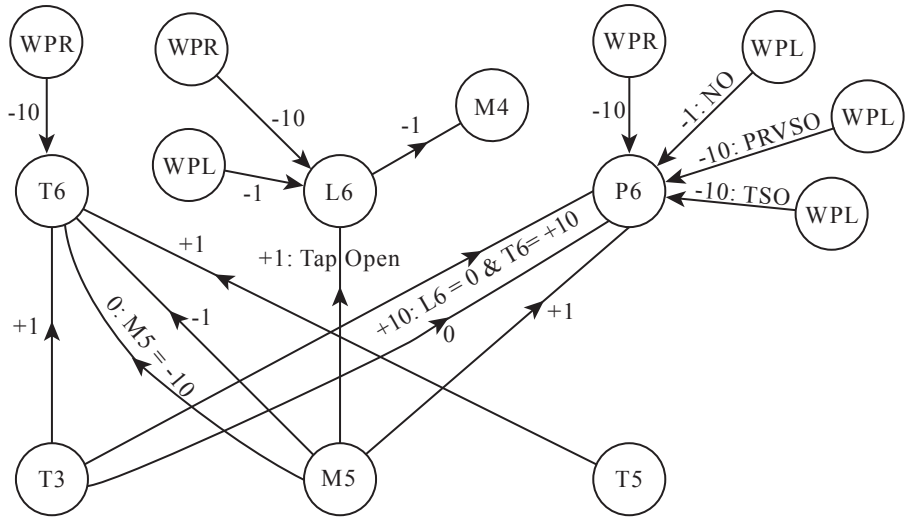

Figure 2.6: Water pipe digraph - Phase 1

pressure lost.

The complete phase 1 system digraph model of the hot water system is shown in Figure 2.7. The figure shows how the individual component digraphs are joined together to create a complete system digraph. The mass flow output from the control valve, M2, is an input to the pilot light, which outputs temperature to the water pipe through T3. The control loop can also be seen on the figure joining the M2, T3, T6, S9 and P10 nodes together.

2.3.3.3 Phase 2 Digraph Models

There are three components that produce identical digraphs in both the first and second phases of operation of the system; the gas pipe, the pilot light and the pressure relief valve. The remaining component digraphs are altered in some form to account for the fact that as the tap is open in phase 2 the control valve should always be open to provide gas to heat the constant addition of cold water to the system. The system digraph for phase 2 is shown in Figure 2.8.

2.3.3.4 Results

The digraph system model will be evaluated using the same process that was applied to the decision table system model. Failure modes will be injected into the model and any unexpected and measurable system symptoms will be recorded and compared to a list of expected symptoms.

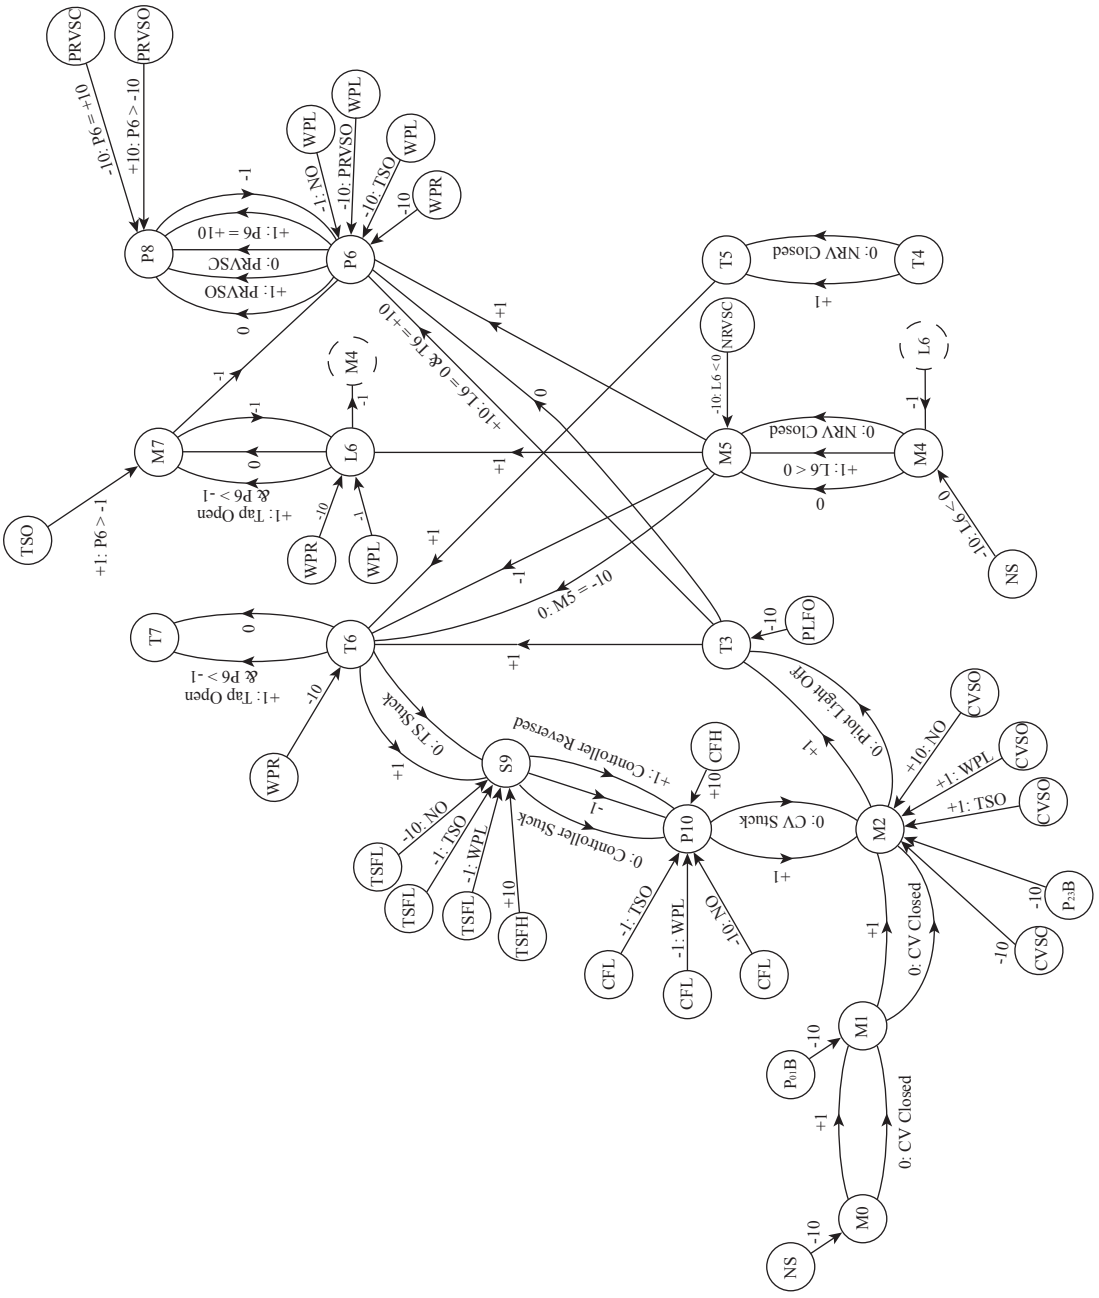

Figure 2.7: Hot water system digraph - Phase 1

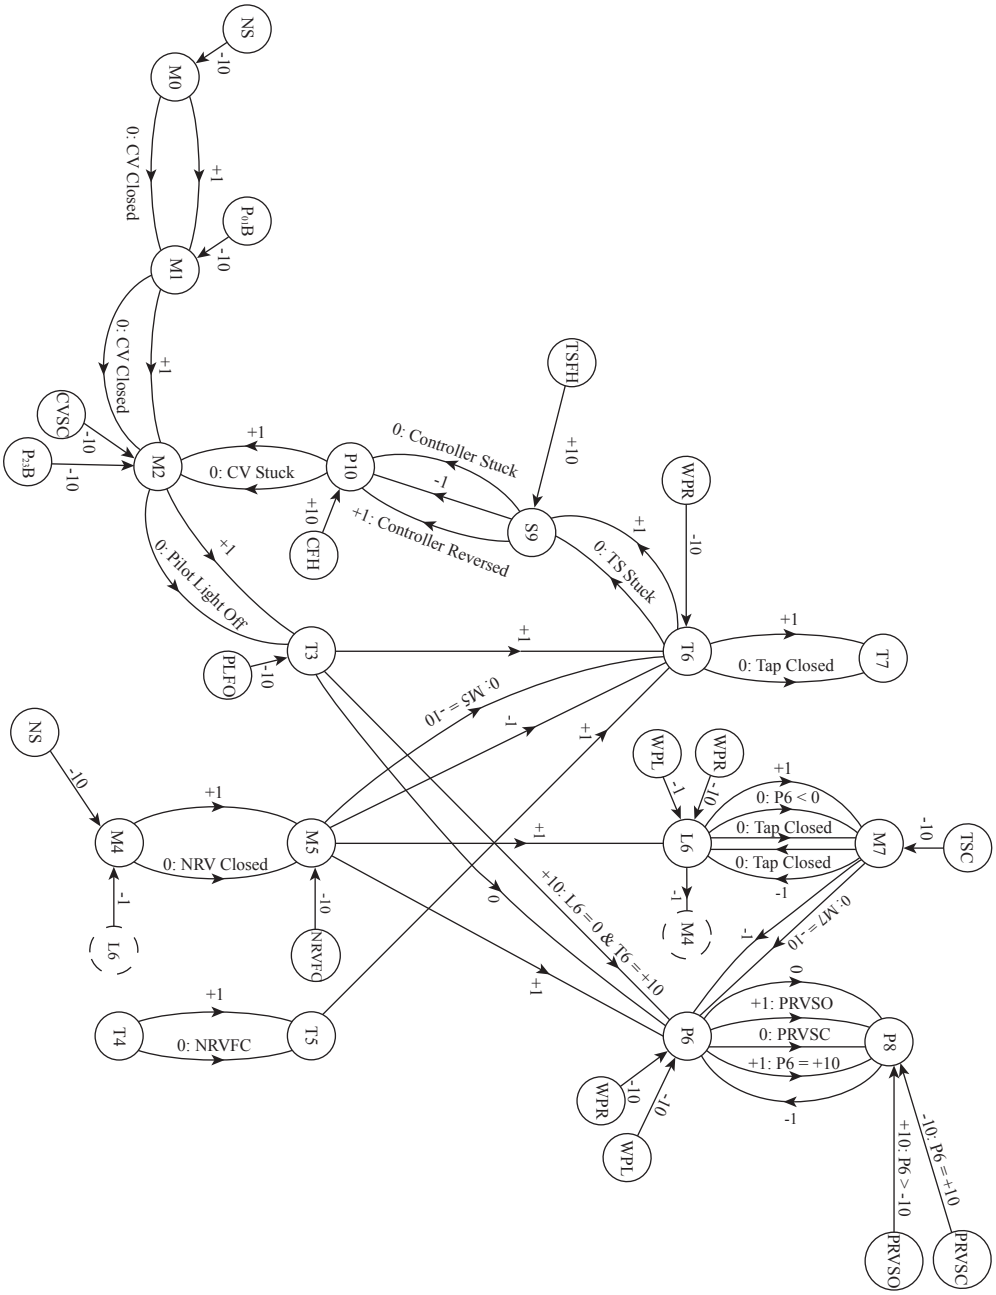

Figure 2.8: Hot water system digraph - Phase 2

The limited manner by which the digraph technique represents component states means that the symptoms identified from the model may not exactly match those expected. Consider the symptoms produced by the failure mode 'No Gas Supply to Gas Pipe' in phase 1 as shown in Table 2.9

Table 2.9: No gas supply to gas pipe phase 1 results

Expected System Symptoms	Digraph Symptoms
Control Valve Open	Control Valve Open
No Gas Out from Control Valve	V. Low Gas Out from Control Valve [M2(-10)]
Low Water Pipe Temperature	V. Low Water Pipe Temperature [T6(-10)]

Table 2.9 shows that while the digraph technique has produced three symptoms in the presence of the fault, two of them are not an exact match with that expected. The digraph technique cannot describe the state of a variable as zero and so the gas output from the control valve has to be displayed as very low. Furthermore the digraph has also predicted the water pipe temperature state to be very low in the presence of the gas pipe fault when it would only be expected to be low. Similar results can be seen when several other failure modes are considered. This demonstrates the lack of flexibility provided by the digraph technique when modelling systems, failure modes and the interaction between variables.

The digraph technique has also failed to identify symptoms where they would otherwise be expected. Consider the symptoms of the failure mode 'Pressure Relief Valve Stuck Open' in phase 2 as shown in Table 2.10.

Table 2.10: Pressure relief valve stuck open phase 2 results

Expected System Symptoms	Digraph Symptoms
Output from Pressure Relief Valve	Output from Pressure Relief Valve [P8(+10)]
Water Pipe Pressure Zero	Water Pipe Pressure Very Low [P6(-10)]
No Output from Tap	
No Flow into NRV	
No Flow out of NRV	

Table 2.10 shows that only two symptoms have been found from the digraph model.

This is due to the fact that when the failure mode is injected into the digraph its propagation is limited to the P6 and P8 nodes. This result can be attributed to the manner in which the model designer has interpreted the system. The relationship between the P6 and L6 variables has been accounted for in the directed edges between L6 and M7. While the failure mode does nullify the relationship between L6 and M7 it has not caused the variable states to change. This example demonstrates the issues that can be caused due to the analyst's interpretation of the system under consideration.

A number of other failure modes experience the issues identified by the results in Table 2.10. As a result the digraph model has failed to identify seven of the ninety-nine symptoms that were expected.

It should also be noted that two digraph models were required in order to model both the phases of the hot water system's operation. Constructing a system model of each phase of a systems operation will require greater amount of resources compared to a technique that could provide a single model capable of dealing with multiple phases.

2.4 Petri Nets

2.4.1 Technique Details

A PN is a bipartite, directed graph that can be used to model the structure and behaviour of a system. The technique uses four tools to generate a system model: a set of places (P), a set of transitions (T), an input function (I) and an output function (O). The combination of these tools forms the structure (C) of the PN, which is expressed mathematically as;

$$C = (P, T, I, O) \tag{2.2}$$

In a PN there are two node types; circles representing places and squares representing transitions. The input and output functions are modelled by directed edges connecting places to transitions and transitions to places respectively. PNs do not allow two nodes of the same type to be directly connected. The basic workings of a PN are shown in Figure 2.9. The modelled system requires two inputs in order to produce an output. The presence of an input or an output in the system is modelled using a small black dot known as a token.

The PNs shown in Figure 2.9 consist of three places and one transition. The input

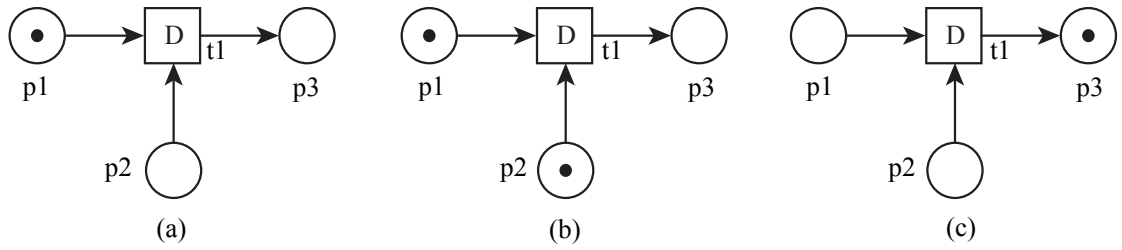

Figure 2.9: Petri net firing process

places, p1 and p2, are identified as circles with arrow headed edges pointing away from them. In (a) only one input is present, as denoted by the token in p1. In (b) both inputs are present. In this state, as the executing requirements of the system have been met, the transition in the centre of the PN, t1, becomes enabled. After a delay of time 'D', the transition fires and a token is taken from both the input places and a single token is added to the output place, p3, as seen in (c). The delay, D, in Figure 2.9 represents the time taken to transform the separate input entities into a single output. Originally the PN technique did not account for any transition delays, all transitions were immediate. However, in its current form transition delays can be set at a specific value, sampled from a distribution or set to zero. If a transition does not have a time delay, a '0' is used instead of a 'D'. If an edge connecting a place and a transition has arrows at both ends then when the transition fires a token will be taken from and returned to the same place. Figure 2.9 demonstrates that PNs do not have to contain a constant number of tokens. The process of firing a transition both destroys and creates tokens.

Figure 2.9 shows a single transition with multiple inputs. In a PN there is no limit on the number of inputs or outputs that a place or transition can have. Individual places are also not limited to containing one token; any number of tokens can be stored in a single place. The edges used to connect places and transitions can also be weighted to indicate the transfer of multiple tokens. This effectively allows a single edge to represent a multiple number of edges. If an edge is weighted a slash is placed through it and a numerical label is placed next to it to indicate its weighting. If there is no slash the edge is assumed to have a value of one. If a transition has any weighted inputs, it will not be enabled until the relevant input place(s) contains at least the same number of tokens as the value of the edge label. Figure 2.10 displays how many of the above features are modelled on a PN.

Due to the weighted edges in Figure 2.10 the transition, t1, requires two tokens in p1

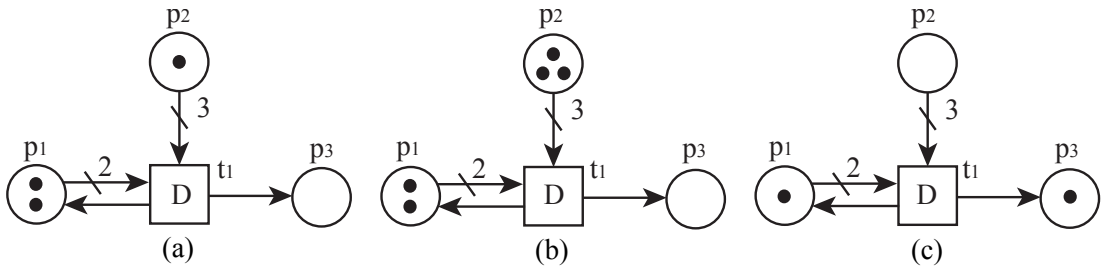

Figure 2.10: A weighted petri net graph

and three tokens in p2 before it will be enabled. In (a) p2 does not contain enough tokens therefore the transition is not enabled. In (b) enough tokens are present to enable the transition. Firing the transition after the time delay removes all the tokens from p1 and p2, places one token in p3 and one token back in p1, as is shown in (c). At this point the transition is no longer enabled.

A PN feature that does change the standard transition firing rules is the inhibit edge. These edges, when active, prevent the firing of transitions and therefore the flow of tokens through a system. Figure 2.11 demonstrates their application. When the inhibit edge, identified by a small circle as opposed to an arrowhead at its tip, is activated by placing a token in p2, the transition does not fire even though it is enabled and the time delay D has elapsed. The enabling of a transition therefore now requires that a suitable number of tokens be present in all its normal input places and no tokens be present in its inhibitor input places. This tool is useful when modelling component failures such as a blocked pipe.

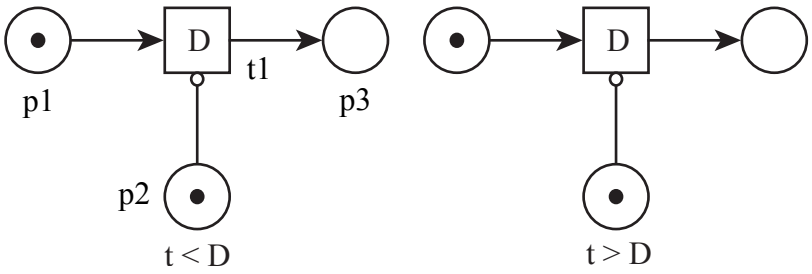

Figure 2.11: A petri net graph with inhibit edge

As can be seen in the above examples, the firing of transitions facilitates the flow of tokens around a PN model. This represents the flow of information, signal or mass flow around the system being modelled. The movement of these tokens not only allows PNs

to model dynamic systems but also makes the PN model itself dynamic. Finally, PNs can show the overall state of a system using unique 'system up' and 'system down' places. Using a simple example, where the firing of a single transition fails the system, Figure 2.12 shows how these places are used.

Figure 2.12: A petri net graph with system up and system down places

2.4.2 Literature Review

PNs were created by Carl Petri and first presented in his 1962 thesis [19]. Petri's initial work was mainly theoretical. The technique was developed by the Information System Theory Project of Applied Data Research and the Computation and Structures Group at M.I.T. This further development created the graphical form of the technique described above.

PNs were designed to be adept at modelling dynamic systems with concurrent processes; systems where multiple component activities occur at the same time [20]. The ability of the PN technique to effectively handle dynamic systems has been verified by multiple authors [21] [22]. Similarly, many authors state that the technique can be readily applied to systems containing concurrent and asynchronous processes [20] [21].

As a modelling tool, the PN technique is relatively simple. It contains only a small number of features that are used to model an entire system. The minimalistic nature of the technique is deceiving however. Russo and Sasso [21] state that "In spite of the apparent simplicity of Petri Nets, the systems that can be modelled are diverse, displaying significant breadth and complexity". It is the flexibility of the PN technique that allows this wide range of systems to be modelled. There is no limit to the number of inputs or outputs from places or transitions. Furthermore, there are no limitations on which places can be linked to which transitions. As a result, PN models can be structured to account for reverse propagation effects and two-way flow. The flexible nature of PNs not only allows complex systems to be accurately modelled but also smaller, simplified models to be produced where appropriate, i.e. time limited situations. These smaller models would

be appropriate in the initial stages of design, where topological system changes are likely to occur [20].

Evidence of the wide use of PNs can be seen in many pieces of literature. Beyond the computer and manufacturing systems already noted [21], PNs have been used to model economic systems [20], workflow processes and laboratory automation processes [21]. Murata [23] also lists a greater number of PN applications than can be considered here. In order to model the range of systems described above, the nodes on a PN have been interpreted in a number of different forms. Table 2.11 shows some of the typical place and transition interpretations.

Table 2.11: Potential petri net place and transition representations

Input Place	Transition	Output Place
Required Resources	Task	Released Resources
Input Data	Computation	Output Data
Input Signals	Signal Processor	Output Signals
Pre-conditions	Event	Post Conditions
Conditions	Logic Clause	Conclusion

Phased missions have also been successfully modelled using PNs. Schneeweiss [24] has demonstrated their application to smaller systems while Mura and Bondavalli [25] and Chew et al. [22] have modelled more complex systems. In the more complex cases multiple sub-nets were used within an overall PN to model the respective missions. Chew et al. used three sub-nets; a phase PN, a component PN and a master PN. This arrangement provided "a more structured modelling technique" that allowed more complex systems to be considered compared to the Mura and Bondavalli technique which only used two sub-nets.

It has been shown that PNs provide a number of advantages as a modelling technique. However, there are also limitations to the technique. In poorly structured PNs there is a condition called 'conflict' that can affect a system model [21]. Figure 2.13 shows an example of a PN in conflict.

In Figure 2.13 both transitions are enabled but there are not enough tokens in p1 to allow both transitions to fire. Should this situation occur, it would be necessary for an analyst to instruct the model which transition should be fired. Alternatively a 'prioritisa-

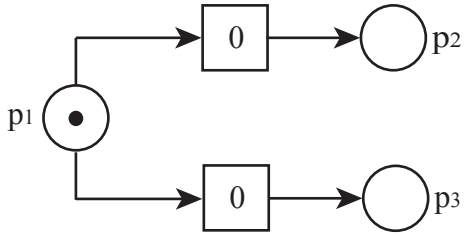

Figure 2.13: A petri net graph in conflict

tion scheme' could be built into the model to determine which transition should be fired [21].

PNs can also be affected by a situation known as 'deadlock'. This will occur when multiple processes, which are using shared resources, are all waiting for another process to finish using a resource. In this situation none of the processes can be completed and the system becomes stuck. A well-used example illustrating this problem is the dining philosophers' problem [20].

A practical limitation of the PN technique relates to the number of tokens within the model. Although there is no limit to the number of tokens allowed in a model, considering a large number of tokens will enable the number of reachable states within a model to increase significantly [26]. A reachable state is a PN marking that can occur through the firing of a number of transitions. In Figure 2.10 the PN in (c) is a reachable state of the PN in (b).

In a PN tokens do not contain any information or data. This means that every unique process in a system must be represented by its own set of places and transitions. Jensen [27] considers this inconvenience to only be "annoying for a small system" but states that when a system with a large number of processes has to be considered, the overall PN becomes unreadable due to its size. In an attempt to overcome this limitation, the PN technique has been developed to allow it to be applied "to more realistic situations" [21]. Russo and Sasso breakdown these developed techniques into two key groups; coloured PNs and hierarchical PNs.

Coloured PNs allow tokens to take on different colours. Each colour of token also contains specified information [27]. For example in a mechanical assembly process the tokens might be coloured and defined as follows; red = raw material one, blue = raw material two, green = machine available, yellow = finished product. With every resource

in the system represented by a different colour, multiple unique tokens can now occupy the same place. This removes the need for unique sub-nets within a PN and therefore reduces the size of the overall PN. The red, blue and green resource tokens, for example, could occupy a single place which inputs to a single transition representing machining taking place. When the transition fires a yellow and a green token would be output. The yellow token represents the finished product and the green token indicates that the machine is available to be used again. Although coloured tokens behave like normal ones in a PN, the rules governing the firing of transitions have changed. Transitions now require not only the correct number of tokens to be present in all input places but the colour of the tokens must also satisfy a condition known as a 'guard condition'. If all of the guard conditions on a transition's inputting edges are not satisfied then transition will not be enabled [21]. The work by Jensen provides much more detail and multiple examples on the topic of coloured PNs [27].

Hierarchical PNs are used for systems that would otherwise produce very large and/or complex PNs. They allow an individual place or transition to represent an entire sub-net that effectively exists on a lower level. The enabling of a hierarchical transition, for example, places a token in the start place of a sub-net. Once the token reaches the end place of the sub-net, the transition on the higher level is fired. The transition delay is therefore represented by the amount of time it takes the sub-net to execute. If a token arrives in a hierarchical place, it will place a token in the starting place of a sub-net and make the token on the higher level unavailable. When the sub-net finishes executing the token on the higher level becomes available again. This situation requires the sub-net to have a single start and a single end place. The key benefit of a hierarchical PN is the effective way in which it can deal with large PNs. This method can provide a simple, easy to interpret upper level PN with all the detail of a large PN effectively hidden in lower levels. It also makes large PNs easier to model as the system can be broken down into smaller sections [21].

In spite of the advantages and wide ranging uses of PNs, Schneeweiss [24] [28] states that "PN are not yet used to the extent that they can be used" and that people out with the reliability engineering community "are hardly aware of them and certainly don't use them". Chew et al propose that the reason for this is the numerous variations of the technique, such as those described above [22]. Schneeweiss believes that there remains great potential for the expanded uses of PNs in the wider world.

2.4.3 Application to Hot Water System

As PNs can model the dynamic behaviour of systems, the behaviour of the hot water system in both phases of operation can be modelled by a single PN. The water pipe and feedback loop sections of the PN model are shown below with the other PN models shown in Appendix C. A number of the same PN place nodes appear in multiple figures. In the system PN model however, these places would only appear once.

The water pipe section of the PN model is split over three figures; Figure 2.14, 2.15 and 2.16. The place descriptions are given in Tables 2.12, 2.13 and 2.14 respectively.

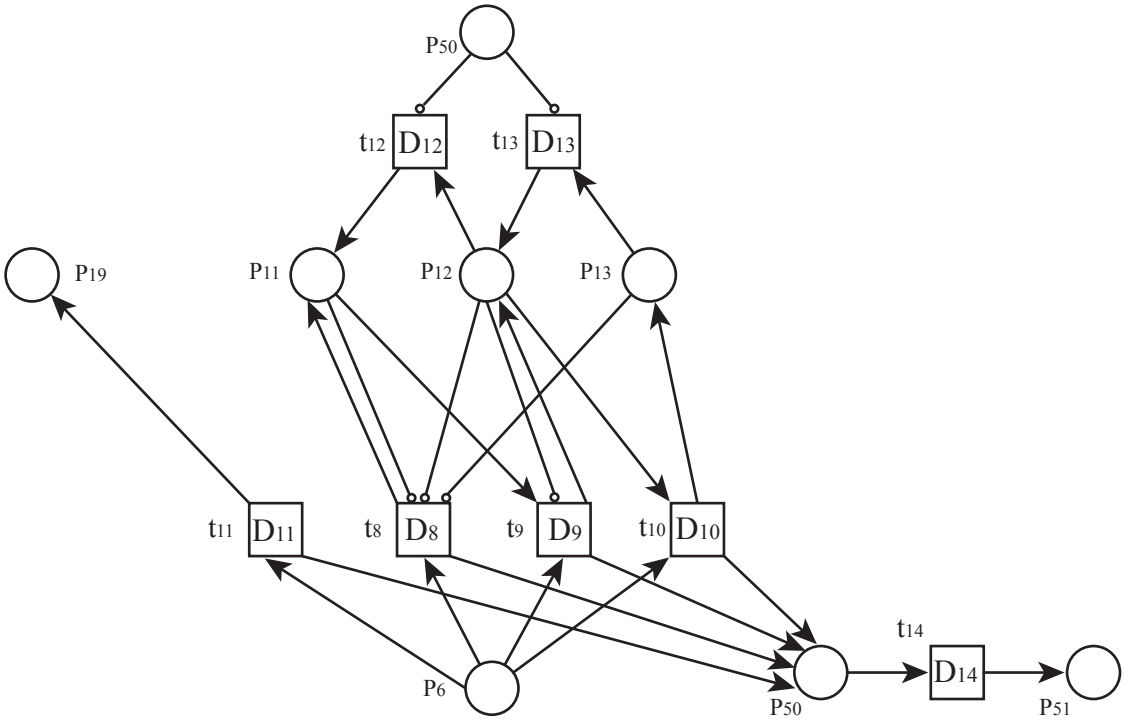

Figure 2.14: Water pipe petri net model [1/3]

Table 2.12: Water pipe petri net place descriptions [1/3]

Place No.	Description	Place No.	Description
6	Heat out of Pilot Light	19	Lost Heat
11	Low Water Pipe Temp	50	Block Cooling
12	Normal Water Pipe Temp	51	Dead Place
13	High Water Pipe Temp		

Figure 2.14 shows how the PN model represents the heating of the fluid in the water pipe. Heat out of the pilot light at p6 enables one of the transitions D8, D9 and D10, which in turn control the presence of tokens in the water pipe temperature places p11, p12 and p13. Should the water pipe temperature be high, the heat provided by the pilot light would be lost; p19. The firing of transitions D8, D9, D10 or D11 adds a token to place p50 that prevents the model decreasing the water pipe temperature. Tokens in p50 are removed to dead place p51 to ensure the water pipe temperature can decrease when there is no heat from the pilot light. The presence of conflict between t8, t9, t10 and t11 is avoided as D11 is greater than D8, D9 and D10.

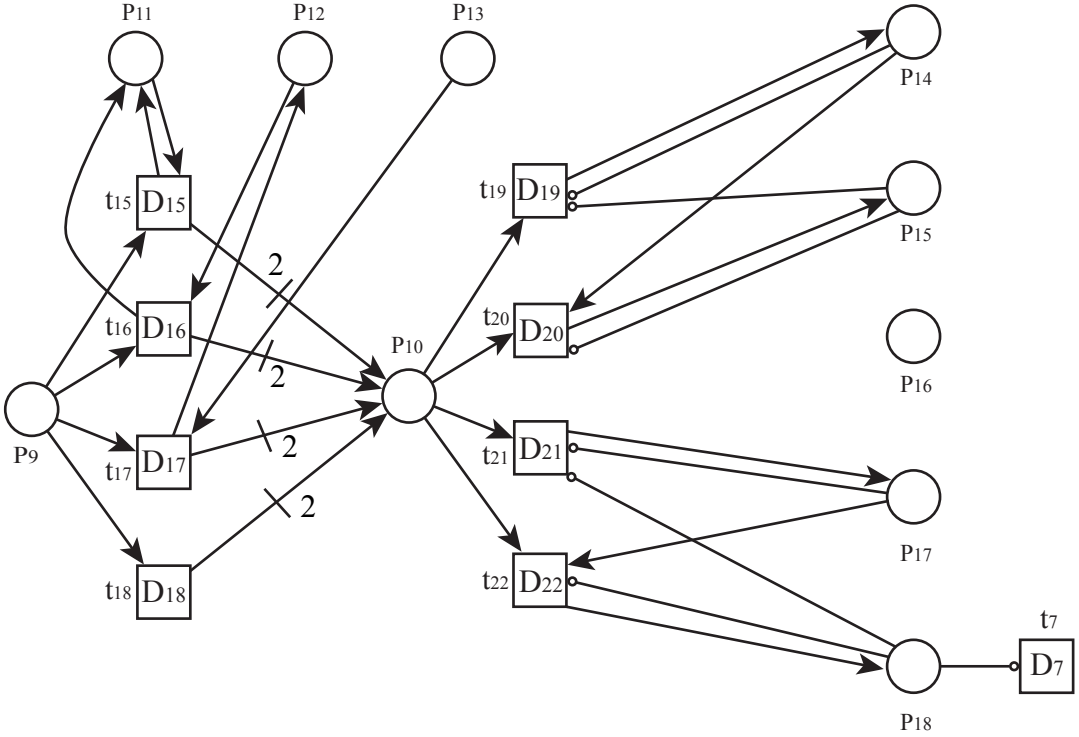

Figure 2.15: Water pipe petri net model [2/3]

Figure 2.15 shows how the addition of cold water from the non-return valve, p9, to the water pipe in p10, decreases the temperature of the water pipe contents as shown in places p11, p12 and p13. Transition D18 is necessary to model the flow into the water pipe when it is ruptured. Transitions D19 to D22 change the water pipe volume and pressure as water enters the piping. When the water pipe level reaches normal, p18, a token in this place will inhibit the further addition of water to the system.

Figure 2.16 shows the effect of a leak or rupture in the water pipe on the system. A

Table 2.13: Water pipe petri net place descriptions [2/3]

Place No.	Description	Place No.	Description
9	Water out of Non-Return Valve	14	Low Water Pipe Pres
10	Water into Water Pipe	15	Normal Water Pipe Pres
11	Low Water Pipe Temp	16	High Water Pipe Pres
12	Normal Water Pipe Temp	17	Low Water Pipe Volume
13	High Water Pipe Temp	18	Normal Water Pipe Volume

Table 2.14: Water pipe petri net place descriptions [3/3]

Place No.	Description	Place No.	Description
11	Low Water Pipe Temp	18	Normal Water Pipe Volume
12	Normal Water Pipe Temp	19	Lost Heat
13	High Water Pipe Temp	20	Lost Pressure
14	Low Water Pipe Pres	21	Lost Volume
15	Normal Water Pipe Pres	33	Water Pipe Ruptured
16	High Water Pipe Pres	34	Water Pipe Leaking
17	Low Water Pipe Volume		

rupture, indicated by a token in p33, removes any and all tokens from the temperature, pressure and volume places; p11 to p18. A leak, p34, reduces the volume and pressure in the water pipe. Places p19, p20 and p21 record the amount of heat, pressure and volume lost from the system.

Figure 2.17 shows the feedback loop section of the PN model. Table 2.15 lists the PN place descriptions. The figure shows how the water pipe temperature enables transitions that can change the control valve state. The failed states of the temperature sensor, controller and control valve also control the ability of these transitions to fire.

2.4.3.1 Results

Having constructed the hot water system PN model the effect of inputting the system failure modes into the model can be evaluated. The symptoms from the PN model have been determined using a software program. The program, written in C++, models the

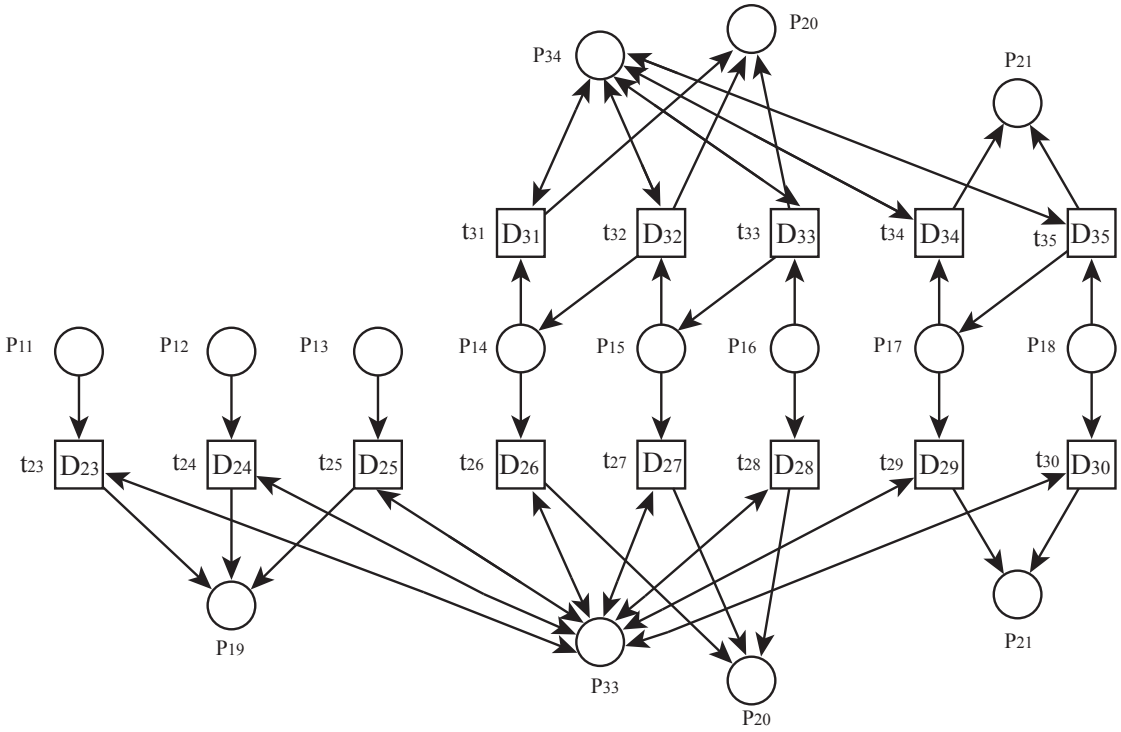

Figure 2.16: Water pipe petri net model [3/3]

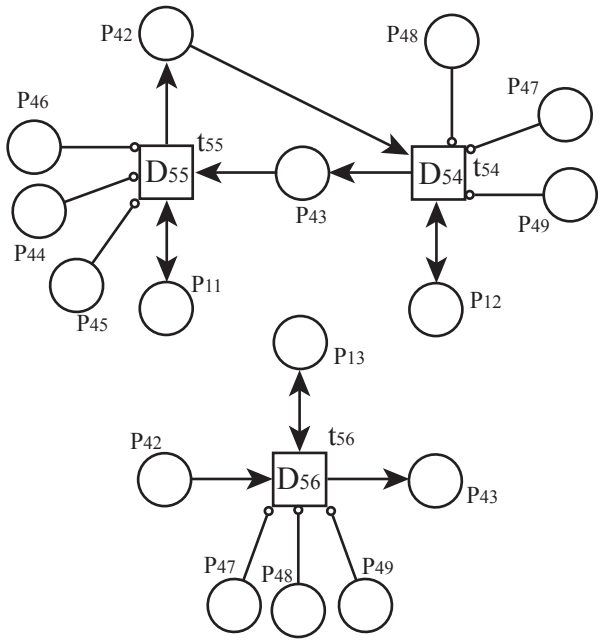

Figure 2.17: Feedback loop petri net model

behaviour of a PN using an input file which lists the places, transitions and initial marking of the model. The PN is given an initial marking that represents the normal behaviour

Table 2.15: Feedback loop petri net place descriptions

Place No.	Description	Place No.	Description
11	Low Water Pipe Temp	45	Temp Sensor Failed High
12	Normal Water Pipe Temp	46	Control Valve Stuck Closed
13	High Water Pipe Temp	47	Controller Failed Low
42	Control Valve Open	48	Temp Sensor Failed Low
43	Control Valve Closed	49	Control Valve Stuck Open
44	Controller Failed High		

of the system and a token is added in the relevant component failure place. The resultant behaviour of the system is then found and the symptoms present in the system are recorded. The process is then repeated with a different component failure until all of the possible failures have been considered.

The symptoms identified by the PN model are, in all failure mode cases, the same as those determined from the theoretical analysis. There are certain failure modes, such as 'Control Valve Stuck Open', where some of the tokens in the PN will always be moving. When this failure mode occurs the water pipe pressure will increase to high causing the excess pressure to escape through the pressure relief valve. In the PN model when there is flow out of the pressure relief valve, the water pipe pressure falls to normal. The continuous high temperature created by the valve failure however, increases the water pipe pressure to high and the process repeats itself.

The ability of the PN technique to correctly predict all of the system symptoms in a single system model demonstrates the flexibility of the technique and an ability to produce accurate models of phased mission systems. Given the success of the PN technique when utilised above, it will be applied to a further system to confirm these conclusions.

2.4.4 Application to Tank Level Control System

2.4.4.1 System Overview

The tank level control system is shown in Figure 2.18. Its principal aim is to maintain a suitable level of water in the tank and to provide a supply of water on demand when the outlet valve is open. If the water level in the tank is insufficient, the pump will be

activated and water will flow into the tank. When the water level reaches the required level the control sensor, L1, will induce a series of events that will turn the pump off. Should the water level ever exceed an acceptable level the trip sensor, L2, will deactivate the entire system.

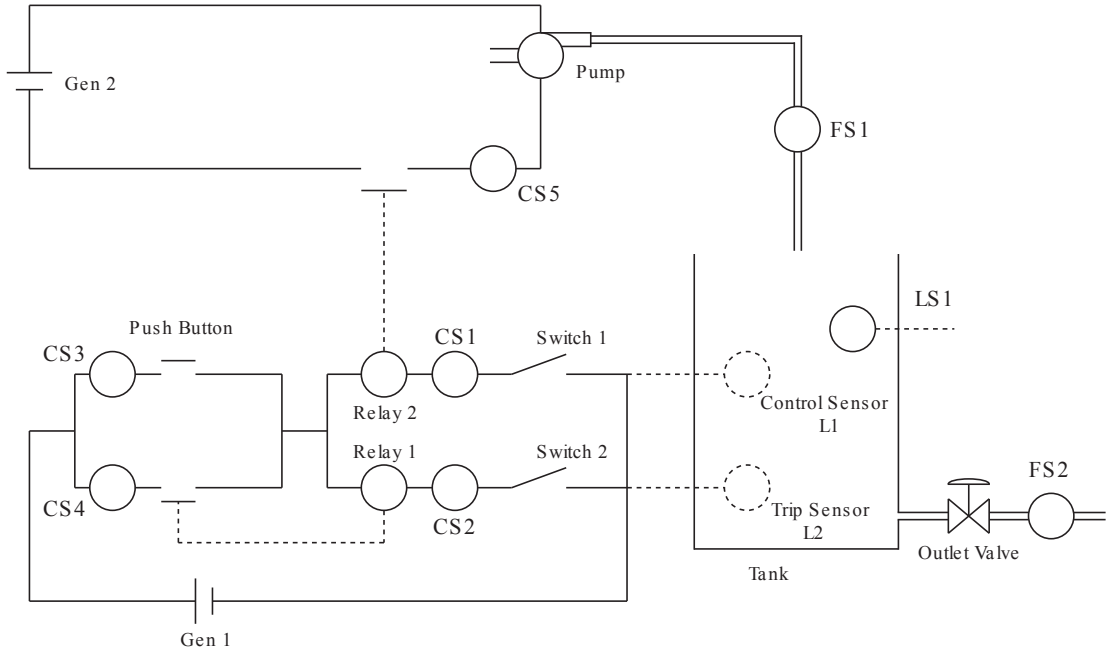

Figure 2.18: Tank level control system

2.4.4.2 System Description

Prior to start-up, the tank is assumed to be empty, the outlet valve and switches closed and the relays open. The start-up process is initiated by closing the push button. This completes the lower electric circuit and energises the relays. The circuit remains active when the push button is released, as relay 1 is self-latching. When relay 2 closes the pump circuit becomes active and the pump begins to fill the tank with water. The level of water in the tank is described by one of five discrete states; zero, low, normal, high or full. The system is set to maintain a normal level of water in the tank and the control loop enables this process. The control loop in the system is comprised of the control sensor L1, switch 1 and relay 2. Opening or closing switch 1, to de-energise or energise relay 2, controls the operation of the pump. If any component in the control loop fails in such a way as to keep the pump circuit active when the tank level is normal, the tank will continue to be filled. If the tank level becomes high or full the trip loop will shut down the entire system. The

trip sensor, L2, opens switch 2, which causes relays 1 and 2 to deactivate and open and the pump to turn off. If relay 1 is opened, the system is effectively shut-down and will require a push button input in order to be re-activated.

The rate at which water can be added to the tank from the pump or lost through the outlet valve is described as either high or low. In normal conditions a high flow rate will always be produced, however if the pump is experiencing a partial failure or the outlet valve pipe is partially blocked, a low flow rate will be experienced. A low flow rate is assumed to represent water flowing at 50% of a high flow rate. Two categories of leaks are also considered; small and large. A large leak is assumed to allow water out of the tank at a high flow rate while a small leak will allow water out at low rate. A leak at any height in the tank will allow the tank level to fall to that point but no further. Leak heights are categorised using the same discrete levels used to describe the water level in the tank.

The state of the tank level control system is monitored by eight sensors; five current sensors, CS1-5, two flow sensors, FS1-2, and a level sensor, LS1. The control sensor and trip sensor do not produce measureable outputs and as such cannot be used to monitor the state of the system. The location of all the system sensors can be seen in Figure 2.18. The current sensors indicate the presence, or absence, of a flow of electric current. The flow sensors indicate a water flow rate of either zero, low or high and the level sensor provides a measure of the water in the tank as either zero, low, normal, high or full. Unless otherwise specified it is assumed throughout that the sensors are working and accurate.

2.4.4.3 Mission Phases

The effect of failure modes on the tank level control system will be considered while it undertakes a five phase mission. The mission phases are described below.

Phase 1: Activation Phase one begins when a push button input is received and ends when the pump circuit has been activated.

Phase 2: Outlet Valve Closed and Pump On In phase two the pump fills the tank with water. When the tank level rises to normal the phase ends.

Phase 3: Outlet Valve Closed and Pump Off Throughout phase three the tank level is normal, the pump off and outlet valve closed. In this phase the system is in standby and the tank level should not change.

Phase 4: Outlet Valve Open and Pump Off The outlet valve is opened in phase four. In this state the tank level should fall as there is no water flow into the tank to replace that being lost.

Phase 5: Outlet Valve Open and Pump On The final operating state of the system has the outlet valve open and pump on. In this phase the rate of water flow out of the system should be equal to the flow rate into the system.

2.4.4.4 Failure Modes

In order to fully test the PN software an extensive number of failure modes in the tank level system have been considered. These are listed in Table 2.16. A number of second order failure modes have also been considered because they are known to override the control loop features of the system. An example of one of the second order failure modes considered is 'Level Sensor 1 Failed Low and Switch 2 Stuck Closed'.

2.4.4.5 Tank Level Control System Petri Net Model

This section presents two sub-nets from the tank level control system PN model. The remainder of the system model is shown in Appendix D. The fuel rig model has also been created as an input file to the PN software introduced previously and described in detail in Chapter 6.

Figure 2.19 shows the power-up, opening and closing processes for relay 1. It should be noted that as in the previous application example some place nodes appear multiple times to make the PN figures easier to interpret. The figures have been simplified further by combining transitions that have the same input and output places except for one place, which is unique to each transition. An example of this can be seen on the right side of Figure 2.19 with switch 2 working open (Sw2WO) and switch 2 failed open (Sw2FO) both listed next to a single place node. This arrangement is representative, on the actual PN model, of two transitions. The first transition has the inputs Re1WC, Re1E, Sw2WO and the outputs Sw2WO, Re1DE and Re1WO. The second transition has inputs Re1WC, Re1E, Sw2FO and outputs Sw2FO, Re1DE and Re1WO. Where there are multiple labels next to a single place it is only possible for one of these labels to be active or true at one time. In the example considered above switch 2 can be working open or failed open but never both at the same time. This means there will never be more than one token in a

Table 2.16: Tank level system first order failure modes

Code	Component Failure
Gen_iF (1,2)	Power supply Gen_i failed
PBSC	Push button stuck closed
SW_iSO/C (1,2)	Switch SW_i stuck open/closed
RE_iSO/C (1,2)	Relay RE_i stuck open/closed
PF	Pump failed
P-50	Pump working 50%
POPB	Pump output pipe blocked
L_1FH/L	Control Sensor L_1 fails high/low
L_2FH/L	Level Sensor L_2 fails high/low
OVFO/C	Output valve stuck open/closed
OVPPB	Output valve pipe partially blocked
OVPB	Output valve pipe blocked
Sml/Lrg-B-Lk	Small/large leak in tank base
Sml/Lrg-L-Lk	Small/large leak at low tank height
Sml/Lrg-N-Lk	Small/large leak at normal tank height
Sml/Lrg-H-Lk	Small/large leak at high tank height
Sml/Lrg-F-Lk	Small/large leak at full tank height

merged place at one time. Finally, any transitions without a specified delay are assumed to have a delay of zero.

Figure 2.19 shows that in order for relay 1 to become energised the generator must be working, switch 2 working closed or failed closed and there must be an input from the push button, or the push button must be failed closed. If all of these inputs are present t1 will fire. Once energised the relay will transition from working open to working closed, t2. If the generator fails or switch 2 is opened, relay 1 will become de-energised and open.

Figure 2.20 shows how the outlet valve state is changed and, as a result, how demand for water to flow out of the tank through the outlet valve is established. From a working closed state, transition t1 will fire when a demand on the outlet valve is present. The continued presence of this demand will inhibit the PN from returning the state of the outlet valve to working closed. In normal operation transition t2 would then fire to create

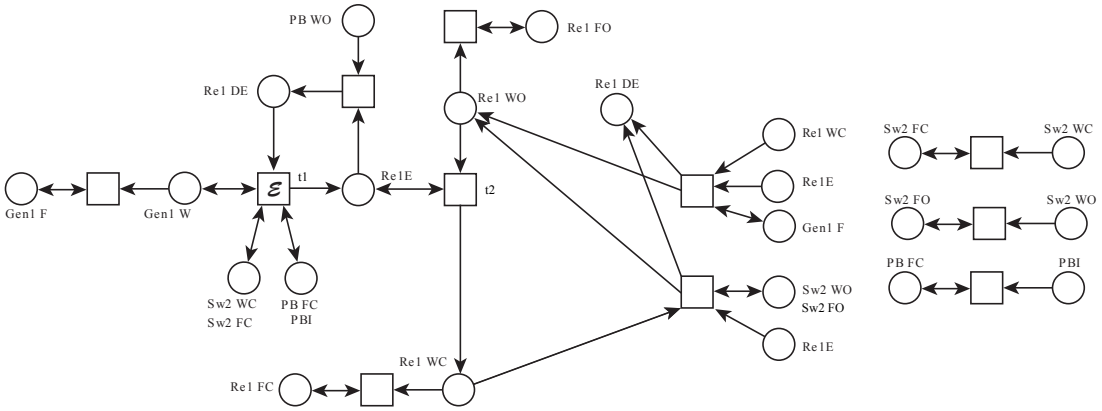

Figure 2.19: Relay 1 powering-up/down, opening and closing

two tokens in the 'Water out Tank' place. Tokens in this place represent the demand for water from the tank. Two tokens would also be added to this place if the outlet valve had failed open. If there is a partial blockage in the outlet valve pipe, only one token will be placed in the 'Water out Tank' place.

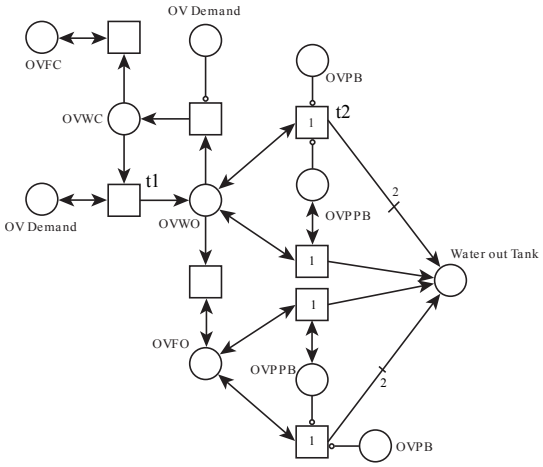

Figure 2.20: Output valve demand

2.4.4.6 Results

Before considering the effect of component failures on the tank level system, the behaviour of the system with no faults present is considered. Figures 2.21, 2.22 and 2.23 show the tank level, pump flow rate and outlet valve flow rate when no faults are present in the system. Two curves are shown on each figure. These represent the sensor outputs from the PN model and the sensor outputs that would be expected on a physical version of the

system.

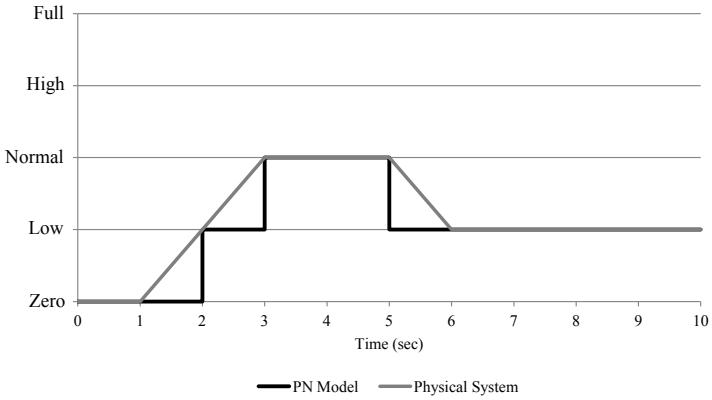

Figure 2.21: Tank level in fault free mission

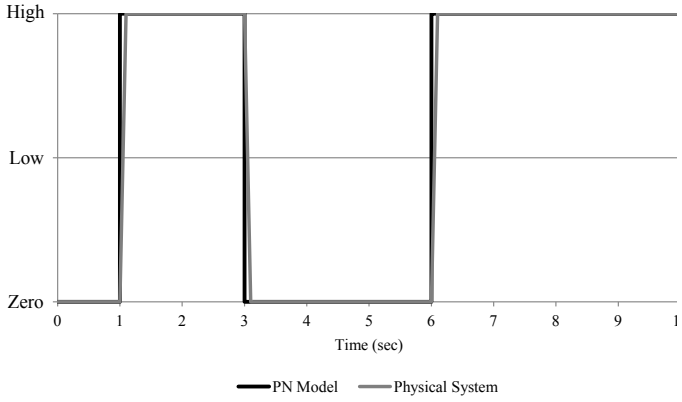

Figure 2.22: Pump flow rate in fault free mission

Figure 2.21 shows the tank level increasing in phase 2 at the same time as the pump flow rate is initially high in Figure 2.22. The decrease in tank level in phase 4 coincides with the increase in the outlet valve flow rate in Figure 2.23. The tank level remains constant in phases 1, 3 and 5 at the same points in time that the pump flow rate and outlet valve flow rate cancel each other out.

All three of the figures show that as the PN model considers only discrete variable states, it cannot model the linear change shown on the sensor output curves expected from the physical system. It would be possible to add more detail to the PN model in order to more accurately represent the behaviour of the tank level system. However given one second timesteps are being considered Figures 2.21, 2.22 and 2.23 show that the PN

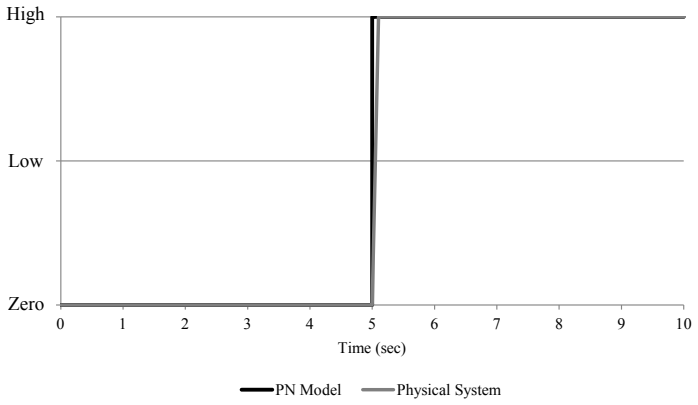

Figure 2.23: Outlet valve flow rate in fault free mission

model already provides a good level of accuracy compared to the expected system outputs.

Table 2.17 lists the current sensor readings that would be expected to be seen in each phase of system operation when no faults are present. The current sensors will either record a positive current flow (Y) or no flow (N). When no faults are modelled in the tank level PN the predicted current sensor outputs match those shown in Table 2.17.

Table 2.17: Tank level system expected current sensor readings

Operating Mode	CS1	CS2	CS3	CS4	CS5
Phase 1	Y	Y	Y	Y	Y
Phase 2	Y	Y	N	Y	Y
Phase 3	N	Y	N	Y	N
Phase 4	N	Y	N	Y	N
Phase 5	Y	Y	N	Y	Y

From Table 2.17 it can be seen that a current flow is only expected at CS3 in phase 1 when there is an input to the push button. As relay 1 is self-latching it can be seen that a current flow is recorded at CS2 and CS4 throughout the mission. Should the trip sensor L2 ever open switch 2, and therefore relay 1, there will be no current flow recorded by these sensors. Finally the output from sensors CS1 and CS5 changes as a result of the tank level. A current flow is recorded when the tank level is below normal, thereby activating the pump, while no current flow is recorded when the tank level is normal. At this time switch 1 is open which breaks the circuits that contain CS1 and CS5.

2.4.4.7 Overview

Thirty-two failure modes were considered when evaluating the tank level control system. All of the faults were modelled using the PN and the results compared to a theoretical prediction of the system's behaviour. In every case considered, the PN predicted system behaviour matched with the theoretical performance of the system in terms of flow rate, tank level and current sensor outputs. Section 2.4.4.8 provides a detailed example of one of the failure modes considered.

Ten failure modes did not cause the behaviour of the system to change from that where no faults were present. These hidden faults are; push button stuck closed, switch 2 stuck closed, relay 1 stuck closed, level sensor L2 fails low and any leak at the normal tank level height and above. All of the non-leak faults act to maintain an electrical current in the trip loop. In normal operation this loop is never broken and therefore the faults are hidden. The leak faults are also hidden as in normal operation the tank level doesn't exceed the normal tank level height.

2.4.4.8 Level Sensor 1 Failed Low

The results of the failure mode 'Level Sensor 1 Failed Low' will be considered in detail below. The failure of level sensor 1, the control sensor, should cause the pump to remain on even though the tank level has reached normal. As the control sensor has failed low it will never open switch 1 which would de-energise relay 2 and deactivate the pump circuit. As a result the tank level will continue to increase to high at which point the trip sensor will open switch 2 thereby de-energising relay 1 and the system as a whole. When the outlet valve is open the tank level should fall to zero. Figures 2.24 – 2.26 show how the system outputs as predicted by the PN model compare to the expected behaviour.

Figure 2.24 shows that the PN model has correctly modelled the increase in the tank level to high as a result of the control sensor fault. The pump flow rate figure also correctly illustrates a longer period of flow into the tank over this time. The system enters operational phase 3 after four seconds, where there is no flow into or out of the tank. Phase 4 begins after six seconds and continues till the end of the mission although there is only a change in the tank level between six and nine seconds. From nine seconds onwards the tank level is zero and therefore there is no flow through the outlet pipe although the outlet valve remains open. The system does not enter phase 5 in the presence of the level

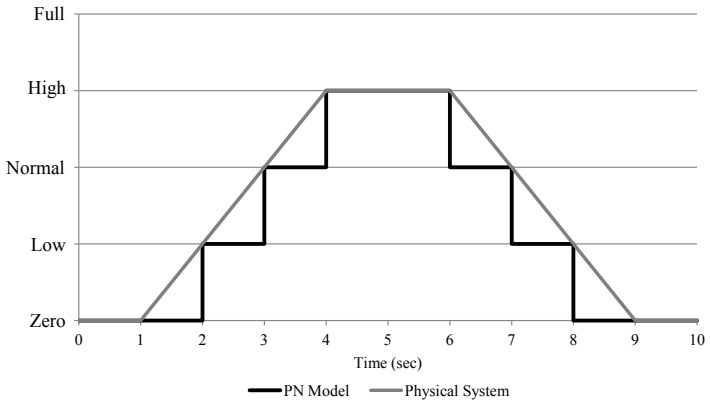

Figure 2.24: Level sensor 1 failed low tank level

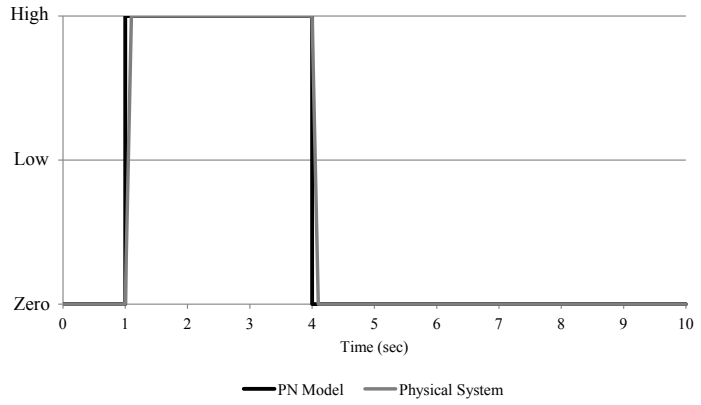

Figure 2.25: Level sensor 1 failed low pump flow rate

sensor fault, as the pump is not on when the outlet valve is open during the mission.

Table 2.18 shows the current sensor outputs from the tank level system in the presence of the 'Level Sensor 1 Failed Low' fault. The PN model outputs matched with those expected from a physical version of the system. The PN model has therefore correctly predicted the behaviour of the tank level control system in the presence of the fault.

2.4.4.9 Second Order Faults

Nine second order failure modes were also considered when evaluating the tank level control system. As mentioned previously, the majority of these were chosen as they were known to override the control loop features of the system. Four of the second order failure modes included the fault 'Control Sensor L1 Fails Low'. This fault was paired with one of the following; push button stuck closed, relay 1 stuck closed, switch 2 stuck closed, level sensor

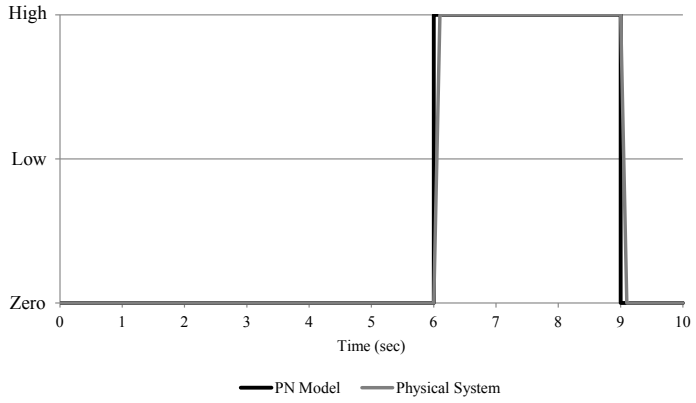

Figure 2.26: Level sensor 1 failed low outlet valve flow rate

Table 2.18: Tank level system expected current sensor readings

Operating Mode	CS1	CS2	CS3	CS4	CS5
Phase 1	Y	Y	Y	Y	Y
Phase 2	Y	Y	N	Y	Y
Phase 3	N	N	N	N	N
Phase 4	N	N	N	N	N

L2 fails low. Four further faults were also considered by using switch 1 stuck closed instead of the control sensor L1 fault. All of these second order failure modes caused the tank level in the system to become full as the pump was always on. The PN model accurately predicted the behaviour of the system when all of these second order faults were injected in the PN.

The final second order failure mode considered was a small leak in the tank base and the outlet valve pipe partially blocked. These faults caused a small deviation in several aspects of the system's behaviour from the second phase onwards. When these faults were included in the PN model, the outputs were very similar to those expected. The only variation appeared in phase 5 where the expected tank level is between the low and normal tank level heights modelled by the PN. As this tank level height was not modelled in the PN, a variation is inevitable. However, had the PN model considered the tank level variable in more detail, it would be expected that the model match the expected behaviour of the system.

The results of modelling the tank level control system and considering the range of faults described above has demonstrated the flexibility of the PN technique. The accuracy with which the tank level control system was modelled and the similarity of the PN outputs with those expected from a physical version of the system show that the technique can be used to provide high level of detail.

2.5 Modelling Technique Selection & Conclusion

All three of the modelling techniques considered by this work have been used to model the hot water system. The PN technique has also been used to model the tank level control system. A summary of the results of each technique is given below.

Decision Tables

Advantages

- Componentistic modelling approach would be easy to apply to large systems

Disadvantages

- Failed to identify 25/99 symptoms
- Identified symptoms that were not expected
- Issues with reverse propagation identified

Digraph

Advantages

- Captures global system behaviour

Disadvantages

- Requires multiple models for phased missions
- Failed to identify 7/99 symptoms
- Lacks modelling flexibility due to limited state and relationship values
- Model subject to analyst's interpretation of system behaviour

Petri Nets

Advantages

- Identified all expected symptoms from hot water and tank level control systems
- Offers greatest amount of modelling flexibility
- Can be applied modularly to systems

Disadvantages

- System models can become very large and difficult to interpret
- Excessive token count can create high number of reachable states

Of the three techniques evaluated above only the PN technique has been able to identify all of the expected symptoms for the failure modes under consideration. This technique has also been shown to offer a high level of modelling flexibility and an ability to model phased mission systems effectively. While the digraph technique also modelled the hot water system with a high level of accuracy it fails to provide the flexibility and phased mission capability offered by PNs. The decision table technique lacked modelling accuracy compared to the other techniques. Given these results the PN technique will be used as the modelling technique for the remainder of this work.

CHAPTER 3

Fuel Rig System and PN Model

3.1 Introduction

This chapter introduces the physical system that will be used as part of developing the fault verification technique and, later, its application. The aim of this chapter is to present the system and the accompanying PN model that has been constructed. The physical system under consideration is the Advanced Diagnostic Test Facility provided by BAE Systems. Figure 3.1 shows the system in the BAE Systems facility.

Figure 3.1: BAE Systems fuel system rig

The BAE Systems Advanced Diagnostic Test Facility, or fuel rig, is a mechanical system that is representative of a fuel system on an unmanned aerial vehicle (UAV). The fuel rig

is controlled by a manual input to a graphical user interface that has been developed by BAE Systems. It is possible to inject or artificially create faults in the fuel rig to replicate those that occur during the operation of a UAV fuel system. The fuel rig is also fitted with a range of sensors whose outputs are used to monitor the behaviour of the system. These sensor outputs can also be recorded for subsequent analysis. Using the recorded fuel rig sensor outputs and determining the PN model variable outputs, the actual and predicted behaviour of the system will be known. All of this data is then available to be used in the fault verification process.

3.2 System Description

3.2.1 System Operation

The fuel rig represents a four fuel tank UAV fuel system. The fuel rig contains three tanks with one split into two by means of a divider. Figure 3.2 shows the fuel rig system and how the two auxiliary tanks are created using a single, split tank. Within both of the wing tanks and the RH auxiliary fuel tanks is a level probe, or sensor, and a set of high and low level switches. The LH auxiliary tank only contains a level sensor. Each tank also contains a drain valve. The pumps on the system are peristaltic pumps. Water is used on the fuel rig to represent jet fuel.

Both wing tanks have direct connections to the left hand (LH) and right hand (RH) engines. Supply to each engine is controlled by a set of triple port L-valves (TPLVs) and the engine pumps. The possible valve settings are ON, CROSSFEED and OFF. The engine pumps operate with a rating between 0 and 100%. When the system is operating normally the valves will be in the ON setting, the engines will have a rating greater than 0% and fuel is supplied to the LH engine from the LH wing tank and to the RH engine from the RH wing tank. When a TPLV is in the CROSSFEED setting, fuel is supplied from the wing tank on the opposite side to the location of the engine. In the OFF setting the engine receives no fuel. If the engine ratings are 0% there will be no fuel supply to the engines irrelevant of the TPLV setting. The fuel supply from the auxiliary tanks to the wing tanks is controlled by the auxiliary engine pump ratings alone.

The fuel feed to each engine is monitored by two sensors; a fuel flow rate sensor and a flow pressure sensor. The engines are represented on the fuel rig by a single collector

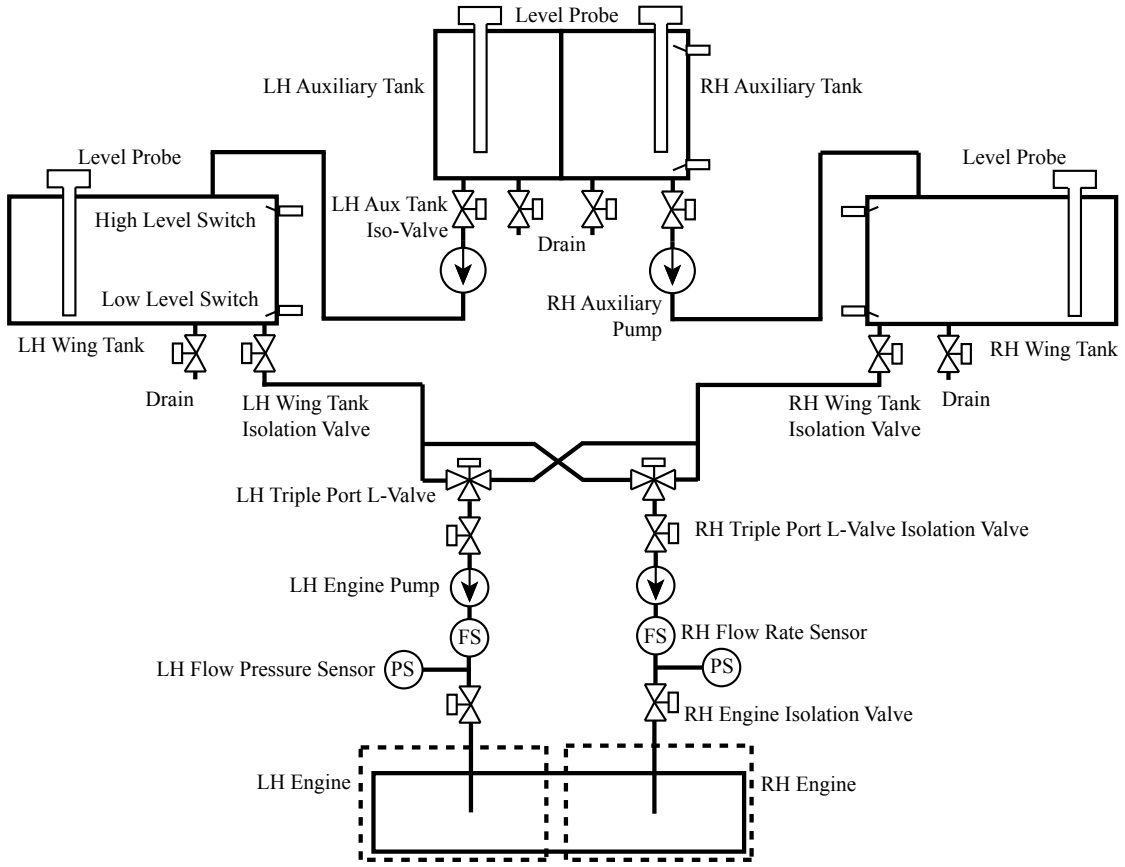

Figure 3.2: BAE Systems fuel rig system schematic

tank where water is accumulated to be re-used in the system. Water is pumped from the engine tank to the fuel tanks using a pump and piping system not shown on Figure 3.2.

3.2.2 Failure Modes

Faults can be injected into the fuel rig system in one of two ways; either manually or through the graphical user interface that controls the system. A valve blockage is an example of a fault that would be manually injected by closing one of the valves on the fuel rig while a sensor failure is an example of a fault that is injected using the graphical user interface. Table 3.1 lists all of the faults that are considered by this work. Further faults can be injected onto the system but were not considered due to logistical/availability issues. It is assumed throughout that if a component fails, it remains in that failed state for the remainder of the mission.

Table 3.1: Fuel rig system first order failure modes

Code	Component Failure
WT Lk	Wing Tank Leaking
AT Lk	Auxiliary Tank Leaking
Eng Pump FO	Engine Pump Failed Off
Eng Pump D	Engine Pump Degraded
Aux Pump FO	Auxiliary Pump Failed Off
Aux Pump D	Auxiliary Pump Degraded
Eng IV B/FC	Engine Isolation Valve Blocked/Failed Closed
TPLV IV B/FC	Triple Port L-Valve Isolation Valve Blocked/Failed Closed
WT IV B/FC	Wing Tank Isolation Valve Blocked/Failed Closed
AT IV B/FC	Auxiliary Tank Isolation Valve Blocked/Failed Closed
LS FH	Level Sensor Failed High
LS FL	Level Sensor Failed Low
LS FS	Level Sensor Failed Stuck
FS FH	Flow Rate Sensor Failed High
FS FO	Flow Rate Sensor Failed Off
FS FS	Flow Rate Sensor Failed Stuck
FP FH	Flow Pressure Sensor Failed High
FP FO	Flow Pressure Sensor Failed Off
FP FS	Flow Pressure Sensor Failed Stuck
H/LLSw FOn	High/Low Level Switch Failed On
H/LLSw FOff	High/Low Level Switch Failed Off
H/LLSw FS	High/Low Level Switch Failed Stuck

3.2.3 System Monitoring and Health Management

To allow for detailed operational and behavioural analysis to be undertaken all of the fuel rig sensor outputs are recorded when the system is in operation. The variable outputs are stored in a text (.txt) file known as the data log. Sensor outputs are written to the data

log at 0.5 second timesteps. The faults identified by the health management system are also recorded and are stored in a separate text document which is known as the health log.

3.3 Petri Net Model

The fuel rig system has been modelled using the PN technique. This PN model has been converted into an input file that can be read by the PN software described in Chapter 6. Section 3.3.1 describes several unique transition types that were used to model the fuel rig system. Section 3.3.2 presents the order in which PN modules are listed in the input file and Section 3.3.3 presents sub-net models of the fuel rig components and their interactions.

3.3.1 Specialist Transitions

The majority of transitions in the PN model are of the standard type as described in Section 2.4. There are however three specialist transition types that are utilised which do not follow the standard conventions.

3.3.1.1 Clear Transition

The 'Clear' transition has two input places and one output place. The purpose of the clear transition is to remove all of the tokens from one of the input places irrelevant of how many tokens there are. Upon firing token(s) are also added to the output place to satisfy the edge weighting. A capital 'C' is placed inside the transition to identify it as a clear transition as shown in Figure 3.3.

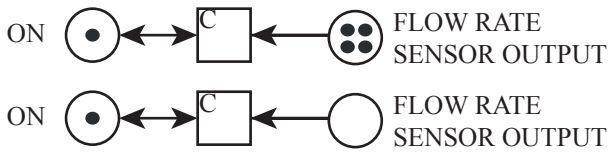

Figure 3.3: Clear transition firing process

3.3.1.2 If Transition

The 'If' transition is used in the PN model to determine if a leak is low enough to affect the water level in a tank. The transition is identified by the capital 'I' in the top left corner of the transition as seen in Figure 3.4.

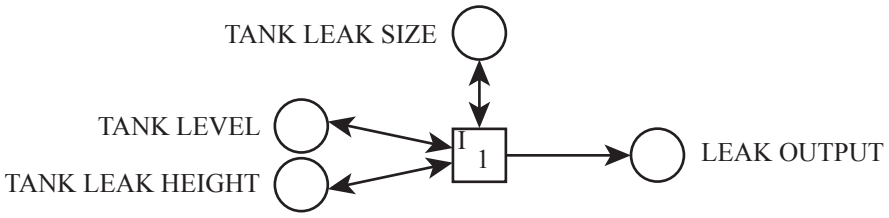

Figure 3.4: If transition

An 'if' transition must have three input places; leak height, leak size and tank level. When enabled, the transition compares the number of tokens in the tank level place to the number of tokens in the leak height place. If there are more tokens in the tank level place, which indicates the tank level is above that of the leak, the transition fires and a token is added to the 'Tank Leak Output' place. If the number of tokens in the tank level place is the same as or less than the tokens in the leak height place, the transition does not fire.

3.3.1.3 Single Transition

The 'Single' transition type is unique in that it can fire once and only once. It is used to model phase changes in the fuel rig PN model. Each 'single' transition also has a number associated with it. The single transitions can only fire in numerical order.

Consider the single transitions modelling phase changes in Figure 3.5. When single transition 'S1' fires after a delay of 10 seconds the engine pump ratings change from 0% to 50%. After a further delay of 30 seconds the engine pump ratings return to 0% as single transition 'S2' fires. Had normal transitions, as opposed to single transitions, been used transition S1 would then incorrectly fire again as it has the same inputs but a shorter delay than transition 'S3'. However as it has already fired, transition 'S3' fires and the mission is completed correctly.

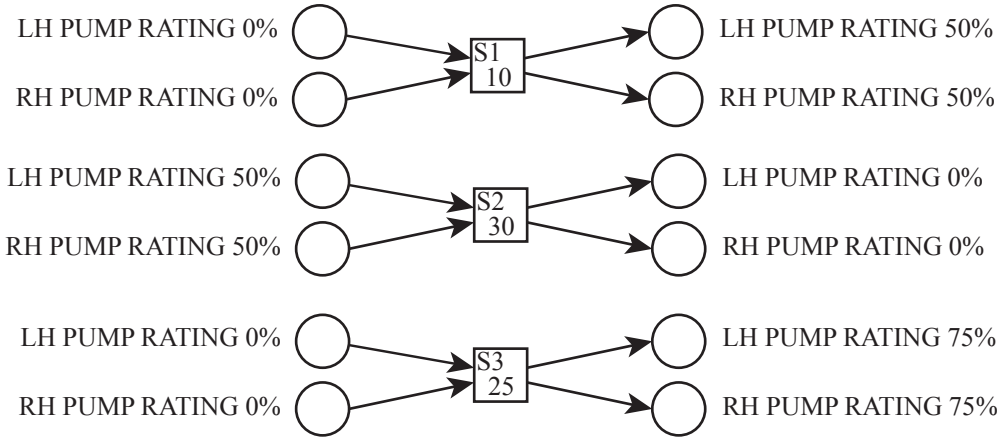

Figure 3.5: Set of single transitions

3.3.2 Model Structure Overview

The list below provides an overview of the fuel rig PN model and the order in which fuel rig components and component behaviour is considered. The order by which transitions are listed in the input file will also be the order by which they are evaluated by the PN software. The ordering of transitions in the input file is therefore important. Further detail regarding the operation of the PN software is presented in Chapter 6. Small sections or sub-nets of the system PN and further detail about each section is considered and presented beyond.

- Clear Sensor Output Places
- Level Sensor State Changes
- High and Low Level Switch State Changes
- Pipe State Changes
- Isolation Valve State Changes
- Pump State Changes
- Tank Demands
- Auxiliary Tank Outputs - Pump and Leak
- Auxiliary Tank Level Changes

- Auxiliary Tank High and Low Level Switch Changes
- Auxiliary Tank to Wing Tank Flow
- Wing Tank Level Changes
- Wing Tank Outputs - Pump and Leak
- Wing Tank Level Changes
- Wing Tank High and Low Level Switch Changes
- Wing Tank to Triple Port L-Valve Flow
- Fuel Flow Rate and Flow Pressure Outputs
- Fault Injections
- Phase Changes

3.3.3 Fuel Rig Component Sub-Net Models

The fuel rig system is structurally the same on the LH side as it is on the RH side. The PN model of the LH side is identical to that of the RH side. The PN sub-net models will therefore only show the RH side of the system.

3.3.3.1 Clear Sensor Output Places

The first transitions listed in the PN model input file are a set of clear transitions with their respective input and output places. Firing of these transitions remove all of the tokens from the flow rate and flow pressure sensor output places. Figure 3.6 shows these sub-nets for the RH sensors. The sub-nets for the LH sensors are identical but with the sensor output and sensor stuck places changed as appropriate.

The flow rate and flow pressure outputs are determined every timestep and therefore the output places must be cleared before these values are determined at the next timestep. It can be seen from Figure 3.6 that if the flow rate or flow pressure sensor is stuck then the clear transitions will be inhibited from firing.

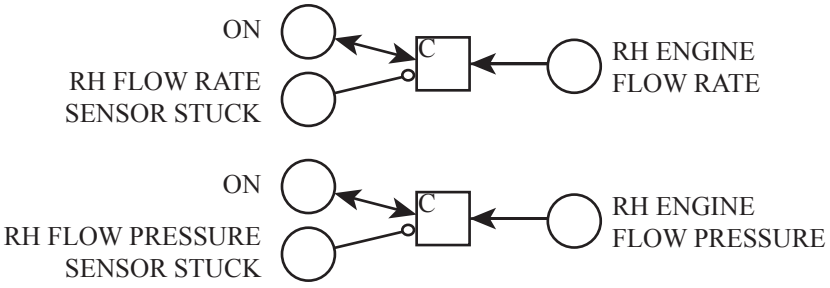

Figure 3.6: RH flow rate and flow pressure clear transitions

3.3.3.2 Level Sensor State Changes

The next group of transitions within the input file are those that control the state of the level sensors. Figure 3.7 shows the transitions that, on firing, change the state of the RH wing tank level sensor.

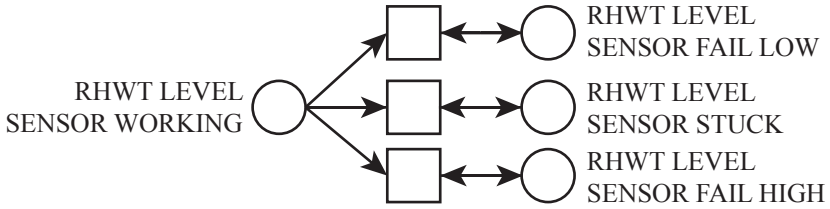

Figure 3.7: RH wing tank level sensor state

Figure 3.7 shows that the state of the level sensor can be changed from working to failed off, failed high or failed stuck. Equivalent places are also listed for the LH wing tank and both of the auxiliary tanks in the input file.

3.3.3.3 High and Low Level Switch State Changes

All of the low and high level switches on the fuel rig can fail both on and off. Figure 3.8 shows how this behaviour has been modelled for the RH wing tank high level switch.

Figure 3.8 shows that the high level switch can fail off from both working states. An electrical fault of the high level switch is an example of what could cause this switch to fail off from the working on state. This event could affect any of the switches on the system. If the switch fails stuck it will not be able to fail in another way. Equivalent transitions for all of the high and low level switches on the fuel rig have been constructed in the same manner as above.

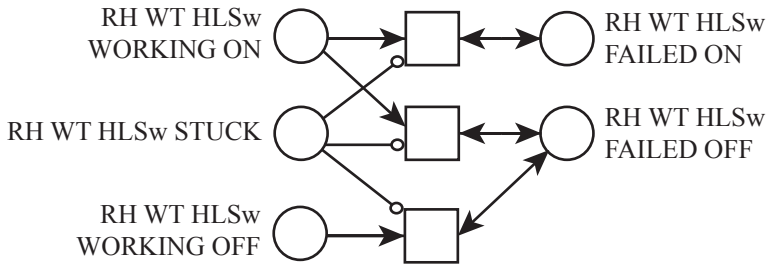

Figure 3.8: RH wing tank high level switch state

3.3.3.4 Pipe State Changes

The state of the piping in the fuel rig can change as a result of a blockage, which will cause a pipe to become obstructed. Pipe blockages have been modelled using two levels of severity; partially blocked and fully blocked. Figure 3.9 shows how these blockages have been modelled in the PN for the pipe between the RH wing tank and the RH TPLV.

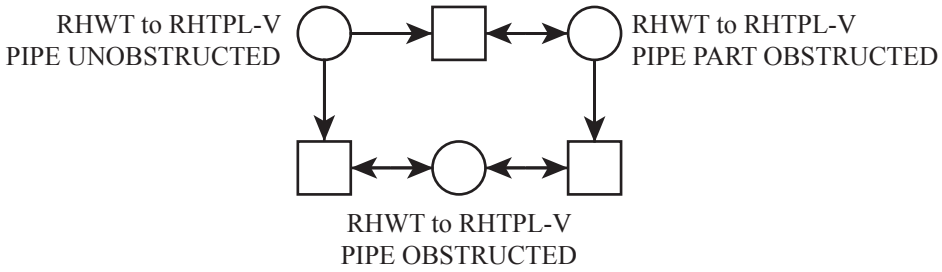

Figure 3.9: RH wing tank to triple port L-valve pipe state

Figure 3.9 shows that the pipe can change state from unobstructed to partially or fully obstructed as a result of a blockage. It can also change from partially obstructed to fully obstructed. The remaining pipe sections have equivalent PN transitions that control their state.

3.3.3.5 Isolation Valve State Changes

This section presents the PN sub-nets of the isolation valve operating states. There are four isolation valves on each side of the fuel rig system. The models of the operating states of the auxiliary tank isolation valve and the wing tank isolation valve are structurally the same. Figure 3.10 shows the PN sub-net of the RH wing tank isolation valve.

The left side of Figure 3.10 shows that the wing tank isolation valve can fail open from

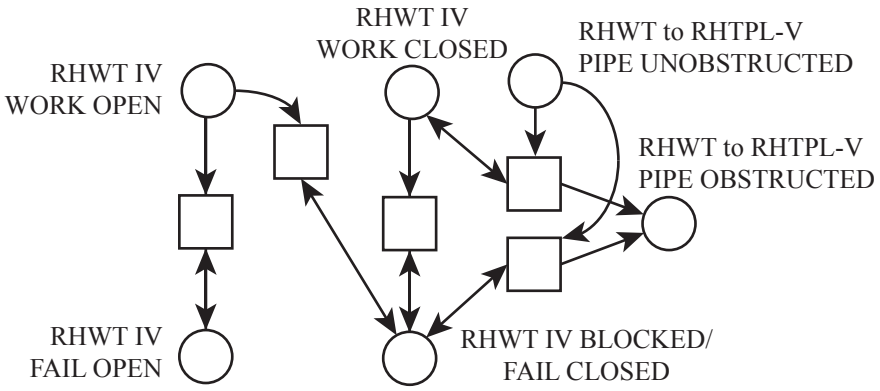

Figure 3.10: RH wing tank isolation valve state

an open state only. It can also experience a blockage when working open. The transitions on the right side of the figure show that if the isolation valve is working closed, blocked or failed closed, the piping section from the wing tank to the TPLV will change state from unobstructed to obstructed. The equivalent model for the auxiliary tank isolation valve changes the state of the piping that runs from the auxiliary tank to the wing tank.

Figure 3.11 shows how the state of the RH TPLV is modelled in the PN. The structure of the PN sub-net in Figure 3.11 is similar to that in Figure 3.10. The TPLV isolation valve can fail open or blocked from a working on state. It can also fail closed from a working closed state. The same PN structure is used for the engine isolation valves.

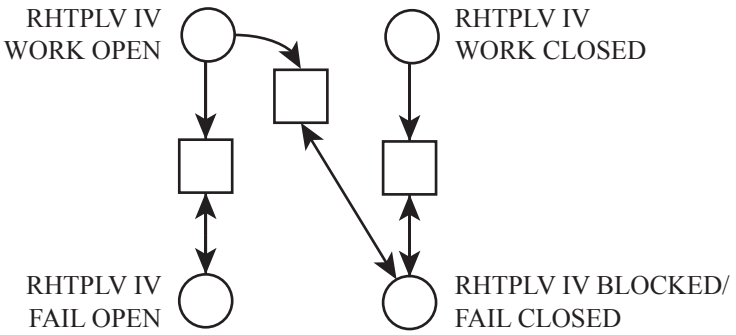

Figure 3.11: RH triple port L-valve isolation valve state

3.3.3.6 Pump State Changes

The pumps on the fuel rig have three possible states; operational, failed off and degraded-operational. Figure 3.12 shows how a RH engine pump failure affects the RH wing tank

demand and RH engine pump operational state. Equivalent sub-nets are also listed for

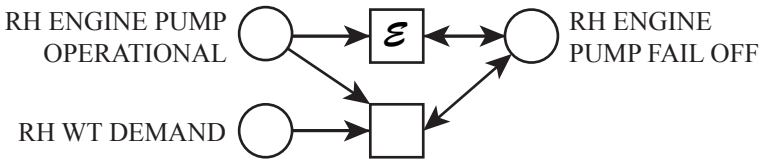

Figure 3.12: RH engine pump state

the LH engine and both auxiliary pumps.

Pumps are in a degraded state when their rating is limited to a maximum of 50%. If the pump is operating below this rating, the effect of the degradation fault will not be seen. However, if the engine rating is greater than 50% when the pump degradation fault occurs, the pump rating will fall to 50%. Figure 3.13 shows how this behaviour has been modelled in the PN for the RH engine pump.

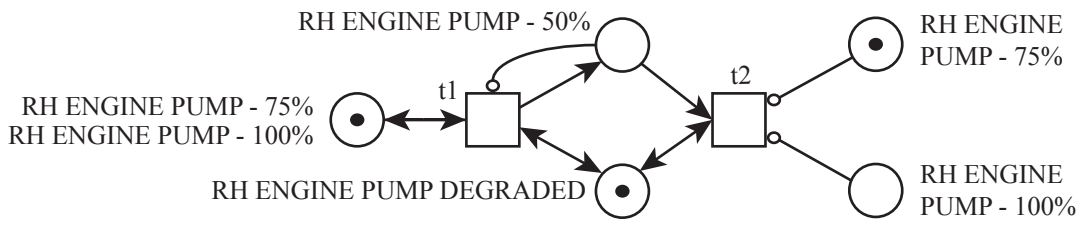

Figure 3.13: RH engine pump degraded state

Figure 3.13 shows how the PN would look immediately after the pump degraded fault is injected when the pump rating is 75%. The transition t1 would first fire adding a token to the 'RH Engine Pump - 50%' place. This ensures the PN simulates the fuel rig behaviour with a reduced engine rating.

Upon transition t1 firing a token will also remain in the 'RH Engine Pump - 75%' place. Should a future phase change reduce the engine pump rating to 50% or below, transition t2 would fire removing the token added to the 'RH Engine Pump - 50%' place by the firing of transition t1. The transitions in Figure 3.13 would continue to fire if the pump rating changed in subsequent phases. Issues that could be caused by having more than one engine rating place marked are resolved when the pump demand transitions are considered.

3.3.3.7 Tank Demands

A demand for water from a tank will be created if the pump is operational or degraded and the pump rating is greater than 0%. If the pump rating changes to 0% or the engine pump fails off, the demand for water will be lost. Figure 3.14 shows how demand from the RH auxiliary tank is modelled in the PN.

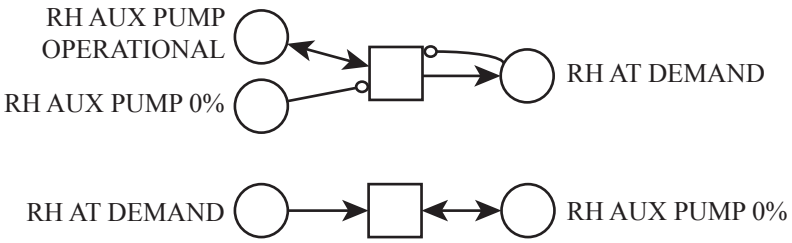

Figure 3.14: RH auxiliary tank demand

Figure 3.15 shows the sub-nets that model the state of demand for water from the RH wing tank. It can be seen that the ability to create demand is dependent on the engine pump states, the engine pump ratings and the TPLV states. Either the RH engine pump must be operational with the RH TPLV in the ON position or the LH engine must be operational with the LH TPLV in the CROSSFEED position. If at least one of these requirements are met, a demand for water from the RH wing tank will be established.

The lower part of Figure 3.15 shows how the RH wing tank demand is lost as a result of the TPLV state changing to OFF or the engine ratings changing to 0%. The lowest three sub-nets cater for the all the possible operating states of the fuel rig including; RH wing tank feeding RH engine, RH wing tank feeding LH engine and RH wing tank feeding both engines.

3.3.3.8 Auxiliary Tank Outputs - Pump and Leak

Having established a demand for water from the auxiliary tank, the next set of sub-nets determine how much liquid will leave the tank. The determining factors in this process are the auxiliary pump rating, the auxiliary tank level and the state of the pipe between the auxiliary tank and the wing tank. For every pump, four pump ratings are considered; 25%, 50%, 75% and 100%. Figure 3.16 shows the transition that fires when the RH auxiliary pump rating is 75% and there is no blockage in the auxiliary to wing tank pipe. There are four input places to the transition. The first confirms the auxiliary tank outlet pipe

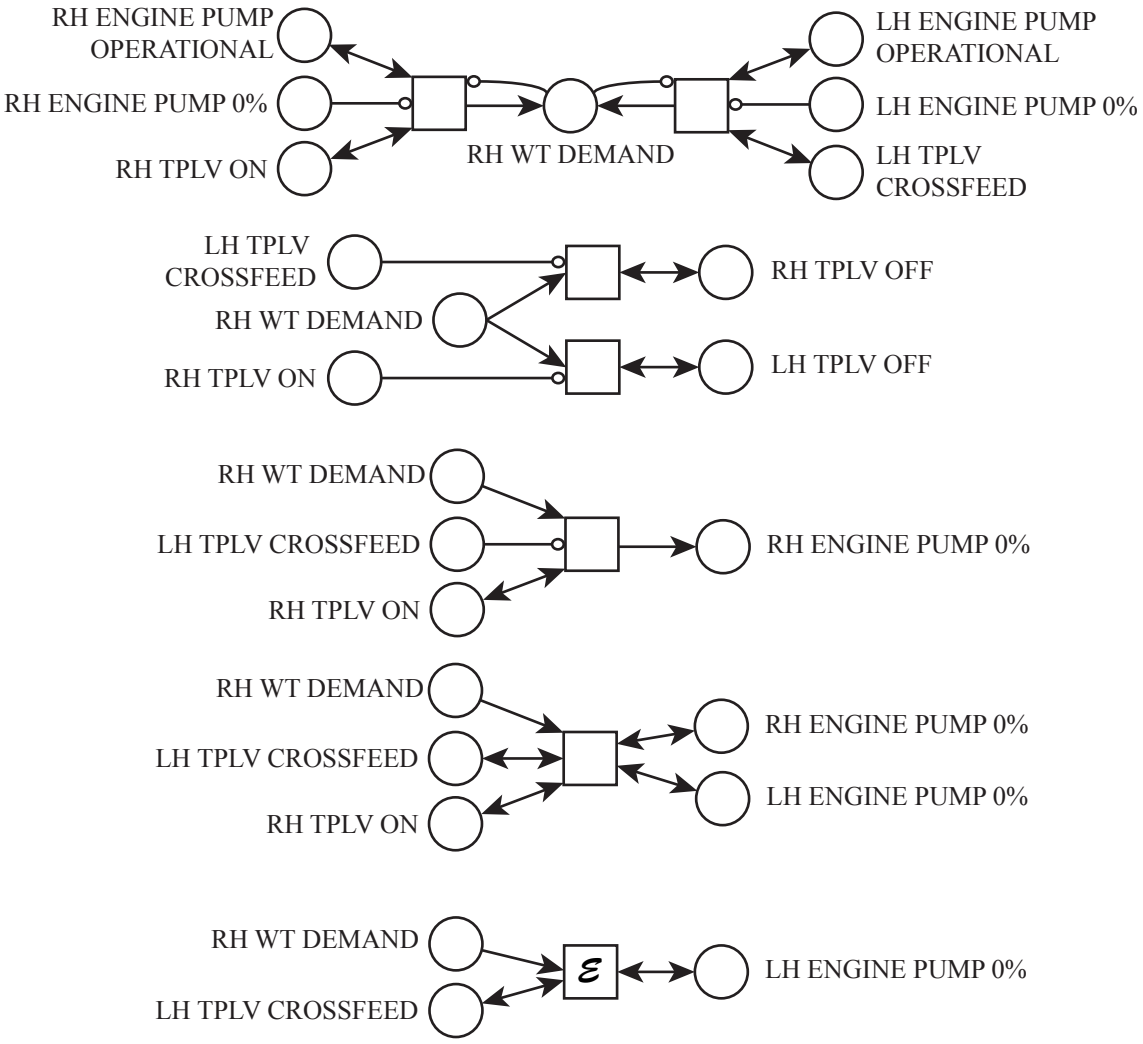

Figure 3.15: RH wing tank demand

is unobstructed. The second place provides the auxiliary pump engine rating. The third ensures there is a demand for water from the tank and the fourth input place confirms there is enough water in the tank to meet the demand. All of the input places are also output places. There are an additional three output places. The first of these indicates the flow out of the tank and into the pipe, tokens in this place will be used to increase the wing tank level. The tokens added to the second output only place will be used to reduce the auxiliary tank level. The tokens added to the auxiliary flow rate place will remain there to indicate the predicted flow rate out of the auxiliary tank. The inhibit edge prevents the transition firing if the pump is degraded.

There are equivalent sub-nets for all four of the possible engine ratings. The weightings of the output edges are adjusted as appropriate dependant on the engine rating. There

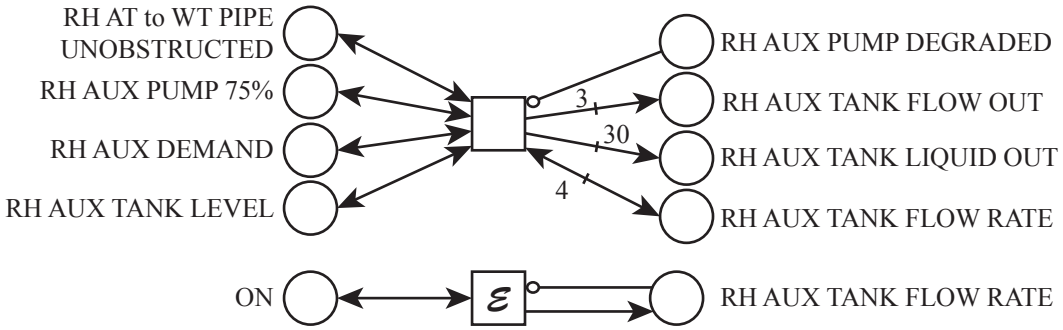

Figure 3.16: RH auxiliary tank output - Auxiliary pump 75%

is also a similar set of transitions to cater for scenarios where the auxiliary tank to wing tank pipe is partially blocked and where the auxiliary pump is degraded.

If for whatever reason there is no flow out of the auxiliary tank, the bottom transition in Figure 3.16 fires to add one token to the flow rate place. A single token in this place indicates a flow rate of 0L/min.

A leak is the other possible cause of an output from the auxiliary tank. A range of leak sizes have been considered in the model in addition to fifty possible leak heights - this reflects the maximum number of tokens used to model the tank level. The process used to determine leak sizes will be considered in more detail in Section 4.10.2. Figure 3.17 shows how tank leaks are modelled in the PN.

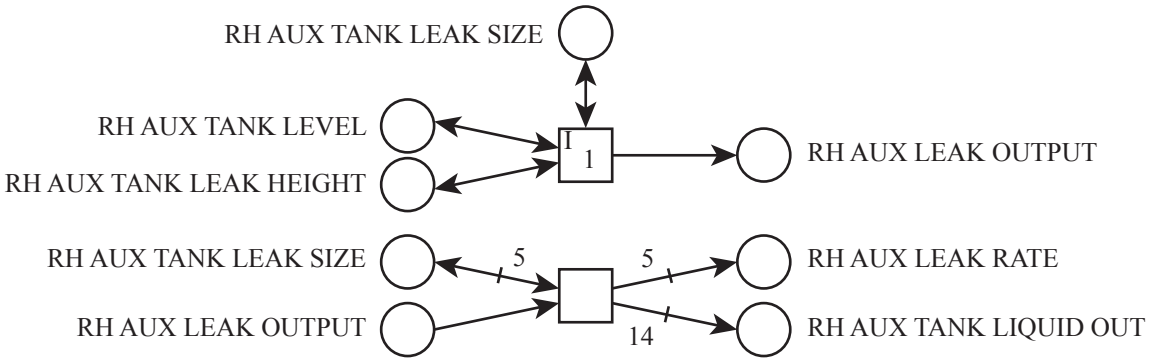

Figure 3.17: RH auxiliary leak output

The top transition in Figure 3.17 is an 'if' transition. As was discussed previously this transition checks to ensure the tank level is above the height of the leak. If this is determined as true a token is added to the leak output place.

The leak output place and the leak size place are inputs to the lower transition. The

tokens added to the leak rate output place remain there to indicate the effect of the leak at each timestep. The tokens added the liquid out place join those added due to the pump outputs, as shown in Figure 3.16, and are used to reduce the auxiliary tank level. There are equivalent transitions for every leak size considered.

3.3.3.9 Auxiliary Tank Level Changes

The previous sub-nets determined how much water will leave the auxiliary tanks due to the pump operation or a leak. The following sub-nets will adjust the auxiliary tank level to account for these losses.

The level in each auxiliary tank is measured by two PN places. For the RH auxiliary tank these places are 'RH Aux Tank Level' and 'RH Aux TL-ve'. The first of these places can contain up to fifty tokens, each of which represents 1.2cm of water to give a maximum tank level of 60cm. Using this place alone however did not provide enough accuracy with which to measure the tank level. The 'RH Aux TL-ve' place was therefore added. This place can contain up to six hundred tokens which, combined, are equivalent to one token in the tank level place. Figure 3.18 shows how the RH auxiliary tank level is changed.

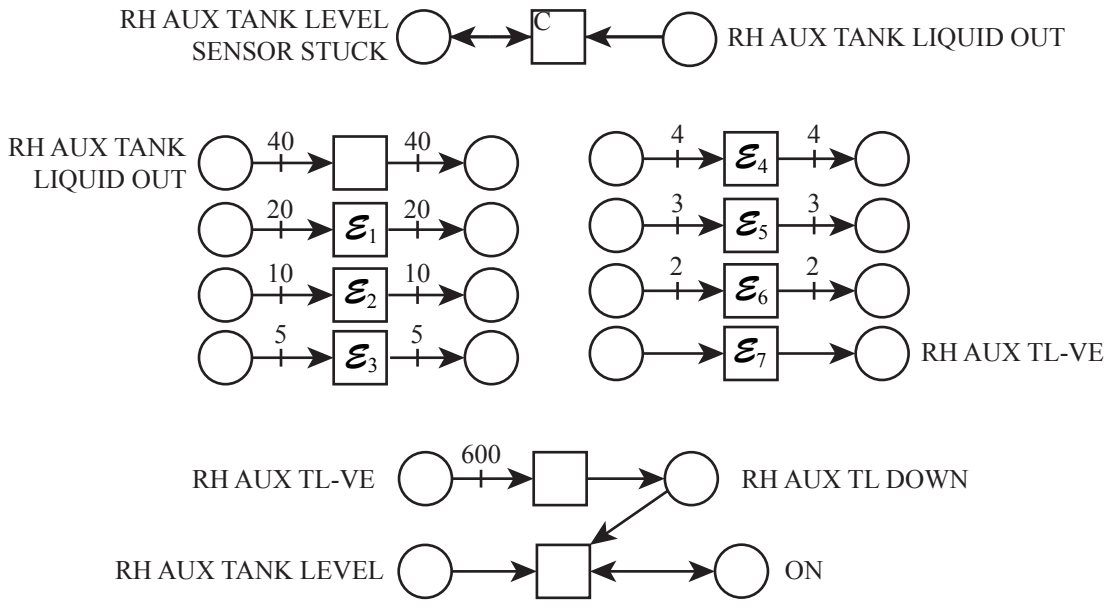

Figure 3.18: RH auxiliary tank level change

The top transition in Figure 3.18 will only fire if the RH auxiliary tank level sensor has

failed stuck. It is a clear transition that removes all of the tokens from the 'RH Aux Tank Liquid Out' place. Removing the tokens from this place will not change the RH auxiliary tank level.

The central block of transitions in Figure 3.18 moves tokens from the 'RH Aux Tank Liquid Out' place to the 'RH Aux TL-ve' place. Using the set of transitions shown will ensure that no matter how many tokens are in the liquid out place, they will all be moved to the TL-ve place. The different engine ratings and leak sizes will all produce a different number of tokens in the liquid out place and using the transitions shown is the easiest way to account for them all and all possible combinations. The delay lengths of these transitions are such that $\varepsilon_1 < \varepsilon_2 < \varepsilon_3 < \varepsilon_4 < \varepsilon_5 < \varepsilon_6 < \varepsilon_7$.

The sub-nets shown at the bottom of Figure 3.18 illustrate how tokens are removed from the 'RH Aux Tank Level' place. Once six hundred tokens are in the TL-ve place, the upper transition fires adding a token to the 'RH Aux TL Down' place. This will then enable the final transition, which removes a token from the tank level place. The PN software will use the token count in both the tank level and TL-ve places when calculating the PN determined auxiliary tank level values.

3.3.3.10 Auxiliary Tank High and Low Level Switch Changes

Having potentially reduced the tank level in Figure 3.18, it is necessary to consider the state of the auxiliary tank low level switches. Figure 3.19 shows how the RH auxiliary tank low level switch state is changed.

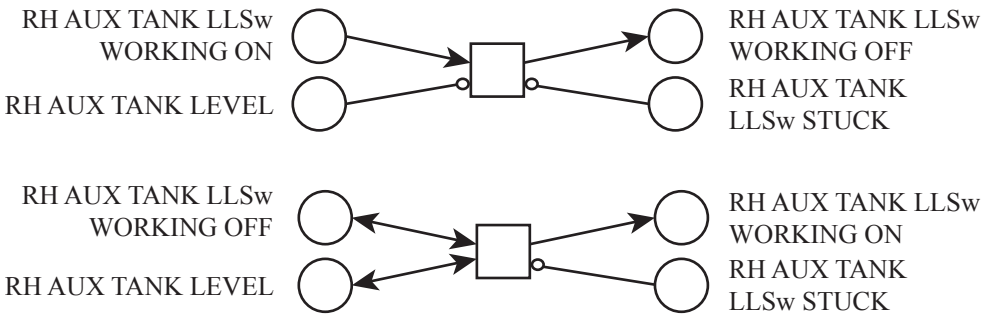

Figure 3.19: RH auxiliary tank low level switch change

The top transition in Figure 3.19 will only fire and change the low level switch state from on to off if there are no tokens left in the 'RH Aux Tank Level' place. The transition is also inhibited from changing the state of the switch if it has failed stuck. The lower

sub-net shows how the low level switch state would change from off to on. The sub-net structure in Figure 3.19 can be applied to all of the low level switches on the fuel rig system by adjusting the input and output places as appropriate.

The high level switch changes are shown for the RH auxiliary tank in Figure 3.20.

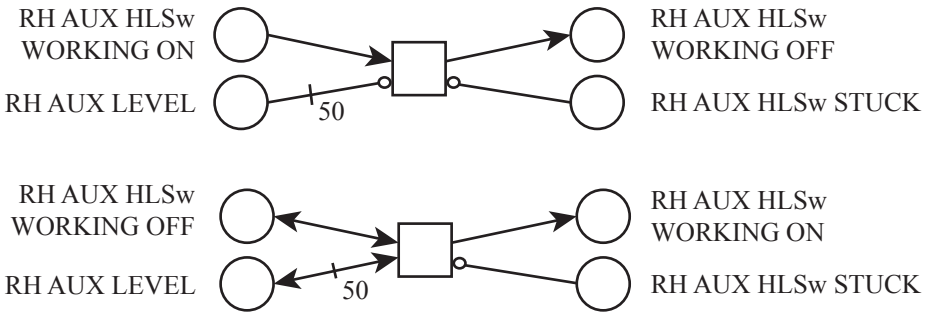

Figure 3.20: RH wing tank high level switch change

Figure 3.20 shows that the high level switch will only be working on if the auxiliary tank level place contains fifty tokens as a result of the tank being full. Otherwise the high level switch state will be 'working off'. These sub-nets can be applied to all of the high level switches on the fuel rig by changing the output and input places as appropriate.

3.3.3.11 Auxiliary Tank to Wing Tank Flow

The sub-nets described to this point have determined the amount of flow to leave the auxiliary tanks and the resultant effect on the auxiliary tank levels. The water flow from the auxiliary tank to the wing tank and the resultant increase in the wing tank level will now be considered.

The piping between the auxiliary tank and the wing tank can fail as a result of a blockage or a leak. A fault with the auxiliary tank isolation valve can also affect this section of the system. The effect of a pipe or isolation valve blockage on the flow out of the auxiliary tank was accounted for in Sections 3.3.3.4, 3.3.3.5 and 3.3.3.8. Figure 3.21 shows how the effect of pipe leaks are included in the PN model when the auxiliary pump rating is 100%.

Figure 3.21 shows that three pipe leak sizes have been considered; small, medium and large. A large leak will cause all of the flow out of the RH auxiliary tank to be lost. Smaller leaks will lose proportionally smaller amounts of flow. The bottom sub-net in Figure 3.21 shows that when there are no leaks in the pipe all of the flow out of the auxiliary tank

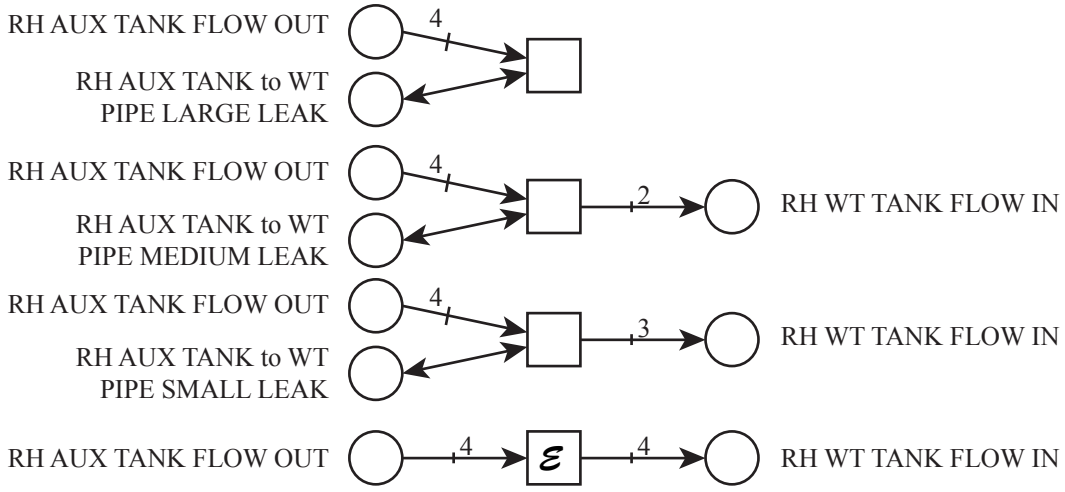

Figure 3.21: RH auxiliary tank to wing tank flow

goes to the wing tank.

Whereas the auxiliary tank level was represented by two PN places, the wing tanks are represented by three PN places. The three places for the RH wing tank are 'RH WT Level', 'RH WT TL-ve' and 'RH WT TL+ve'. The additional place, TL+ve, is necessary to account for the increases in the wing tank level as a result of the auxiliary tank input. Both the TL-ve and TL+ve place can hold a maximum of six hundred tokens and this represents one token in the 'RH WT Level' place. The PN software accounts for tokens in all three places when determining the wing tank level values.

Having determined what proportion of the flow out of the auxiliary tank will reach the wing tank, the next consideration is to convert this into tokens that fill the 'RH WT Liquid In' place. Figure 3.22 shows, given the different input flow rates, how many tokens are added to the wing tank liquid in place. In terms of delays, $\varepsilon_1 < \varepsilon_2 < \varepsilon_3$.

Figure 3.22 shows that each of the possible flow rates into the wing tank will cause a different number of tokens to be added to the 'RH WT Liquid In' place. If the wing tank level sensor is stuck all of the tokens added to the 'RH WT Liquid In' place are removed to prevent the wing tank level from being changed.

The number of tokens in the 'RH WT Liquid In' place represents the increase in the wing tank level. The tokens in this place are moved to the 'RH WT TL+ve' place using equivalent transitions to those shown in the centre of Figure 3.18. If the respective input conditions are met the transitions in Figure 3.23 will then fire to increase the RH WT Level token count by one.

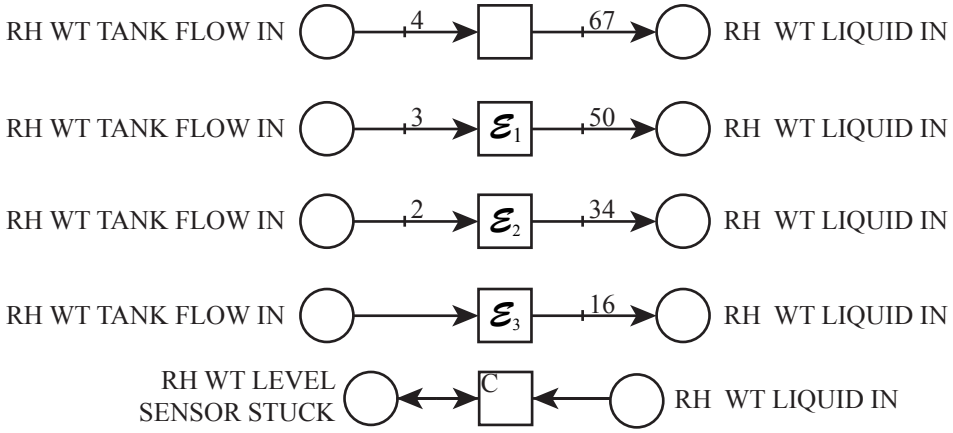

Figure 3.22: RH wing tank level increase

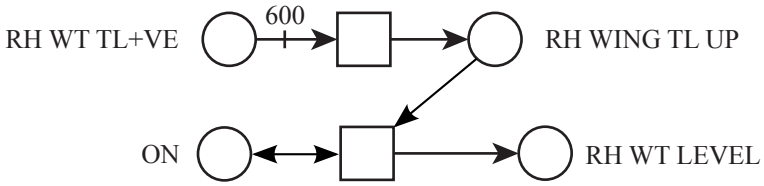

Figure 3.23: RH wing tank level change

3.3.3.12 Wing Tank Outputs - Pump and Leak

Figure 3.15 showed how demand for water from the wing tanks was established. The transitions which then fire to satisfy this demand are very similar to those in Figure 3.16 which showed the RH auxiliary tank pump output. Figure 3.24 shows the transition that fires when the RH wing tank experiences a demand from the RH engine which has a rating of 50%.

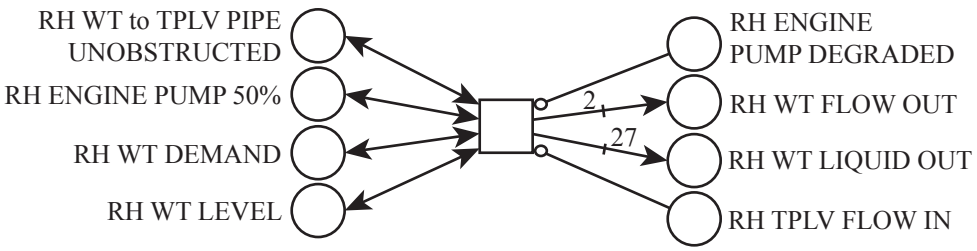

Figure 3.24: RH wing tank output - RH engine pump 50%

Figure 3.24 shows that the inputs to the transitions are equivalent to those in Figure

3.16. There are two output only places however, flow out and liquid out. The tokens in the flow out place will be used to determine the flow rate and flow pressure sensor outputs downstream. Tokens in the liquid out place will be used to reduce the wing tank level. The inhibit edge originating from the 'RH TPLV Flow In' place stops the transition from firing if there is already water at the input to the TPLV. The cause of such a backlog will be seen in subsequent transitions.

Equivalent transitions are in place for all of the remaining engine pump ratings and when the engine pump is degraded. Dependent upon the input conditions, the number of tokens added to the flow out and liquid out places varies as appropriate. Leak outputs from the wing tanks are modelled in the same way as leak outputs from the auxiliary tanks.

3.3.3.13 Wing Tank Level Changes

The PN sub-nets that control the decrease in the wing tank levels as a result of pump and leak outputs have the same structure as the sub-nets which fulfil the function for the auxiliary tanks. The transitions shown in Figure 3.18 are duplicated with the relevant wing tank input and output places used instead of the auxiliary tank places.

3.3.3.14 Wing Tank High and Low Level Switch Changes

The sub-nets that model changes in the operating state of the low level switches on the fuel rig were shown in Figure 3.19. Equivalent sub-nets are used to model the RH wing tank low level switch states.

The sub-nets that model the state of the RH wing tank high level switch are equivalent to those shown in Figure 3.20.

3.3.3.15 Wing Tank to Triple Port L-Valve Flow

Flow leaving the wing tanks will travel through a pipe to either the LH or RH TPLV. Figure 3.21 showed different scenarios of flow between the auxiliary and wing tanks in the presence of leaks and no leaks. Equivalent sub-nets are used to model flow between the wing tanks and the TPLVs. Figure 3.25 shows the sub-nets modelling the flow from the RH wing tank to the LH and RH TPLVs when there are no leaks in the piping.

The first two transitions in Figure 3.25 show that the flow out of the RH wing tank will reach the input of the TPLVs subject to the operating state of the respective valve

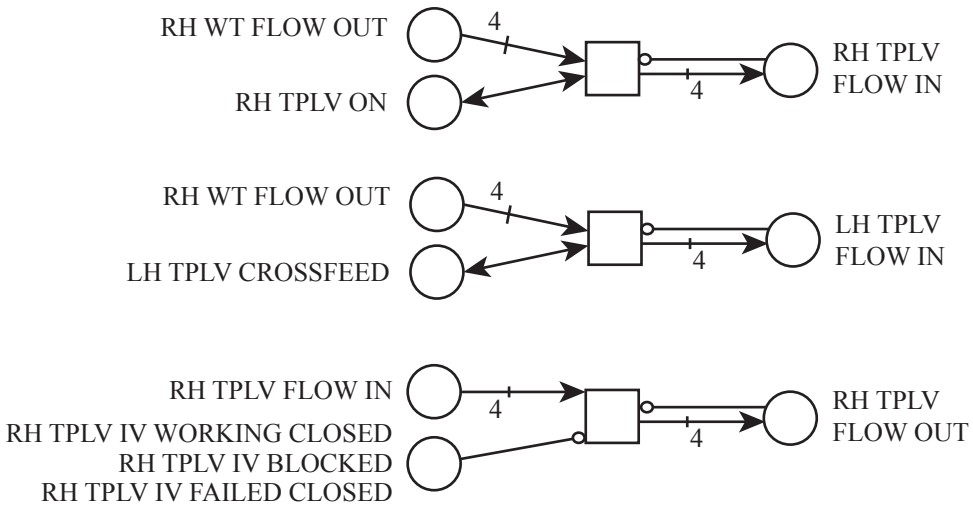

Figure 3.25: RH wing tank to TPLVs flow

and the input not containing water already. The TPLV input may contain a backlog of water if, for example, an isolation valve downstream has become blocked or failed closed. The bottom sub-net shows that water will flow from the TPLV input to the TPLV output if the TPLV isolation valve is not closed or blocked. For the transition to fire it is also required that there is not a backlog of water at the valve exit. A backlog at this point could be caused by a fault with the engine isolation valve. Equivalent sub-nets to those shown above are also included in the input file to account for all of the possible number of tokens in the 'RH WT Flow Out' place.

3.3.3.16 Fuel Flow Rate and Flow Pressure Outputs

Once water reaches the exit of the TPLV isolation valves, it will pass through the engine pumps and engine isolation valves before reaching the 'engines'. The sub-nets at this point in the input file therefore determine the outputs from the flow rate and flow pressure sensors while also accounting for the operating state of the engine isolation valve and the engine pump rating. Figure 3.26 shows one of the transitions that could fire when the RH engine isolation valve is working open and the RH engine pump rating is 100%.

The transition in Figure 3.26 shows that when the engine isolation valve is open and the sensors have not failed stuck, tokens will be added to the flow rate and flow pressure places. It is not necessary to have the engine rating places as inputs because the fact that there are tokens in the 'RH TPLV Flow Out' place indicates the engine rating must be

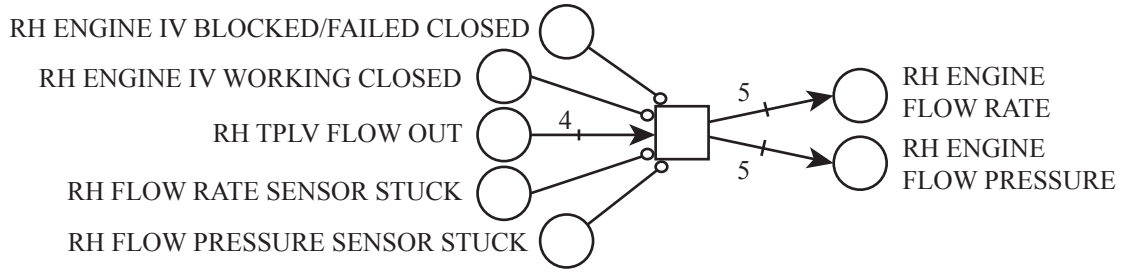

Figure 3.26: RH sensor outputs - Sensors working, engine IV open, engine rating 100%

greater than 0%. Equivalent sub-nets account for all possible number of tokens that can be present in the 'RH TPLV Flow Out' place. The number of tokens added to the sensor output places is proportional to the tokens in the flow out of the TPLV place.

If there are no tokens in the 'RH TPLV Flow Out' place but the engine rating is greater than 0% the transition in Figure 3.27 would fire. This scenario may occur, among other reasons, as a result of a blockage or fault with the wing tank isolation valve or TPLV isolation valve.

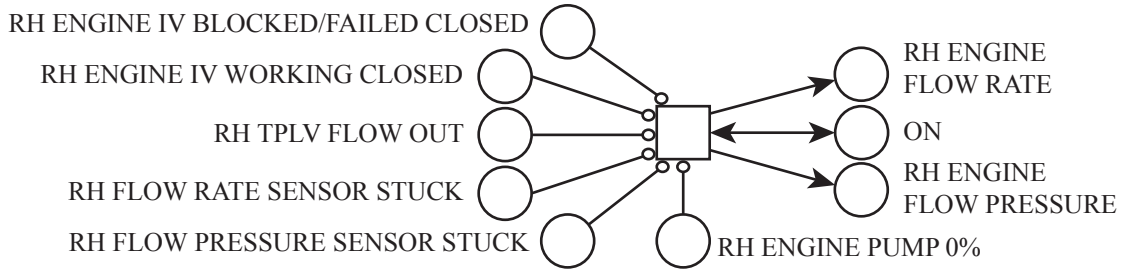

Figure 3.27: RH sensor outputs - Sensors working, engine IV open, engine rating > 0%

In Figure 3.27 a token is added to the flow rate and flow pressure places even though there is no water flow. A token has been added to these places to enable the software to differentiate between when there is no flow in the system and when the sensors have failed off. If no tokens are present in the sensor output places the software will interpret this as the sensors having failed off. Placing a single token in each place when there is no flow is therefore critical for the PN outputs to be accurate. Figure 3.28 shows that a single token is also placed in each of the sensor output places when the engine rating is 0%.

If the engine isolation valve becomes blocked or fails closed when the engine pump is on a small amount of water will be trapped between the pump and the engine isolation valve. As the engine pump remains on the trapped fluid will become highly pressurised even

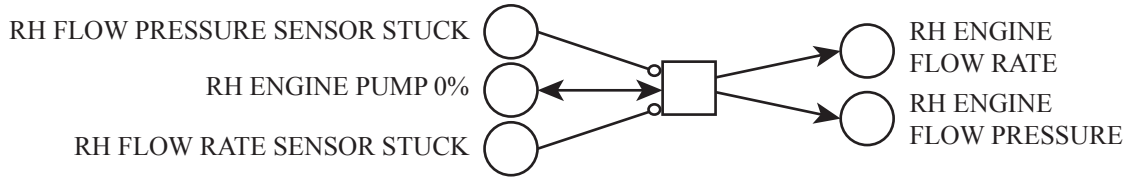

Figure 3.28: RH sensor outputs - Sensors working, engine IV open, engine rating 0%

though the flow remains stationary. This is true irrelevant of the engine rating. Figure 3.29 shows how this behaviour has been modelled in the PN.

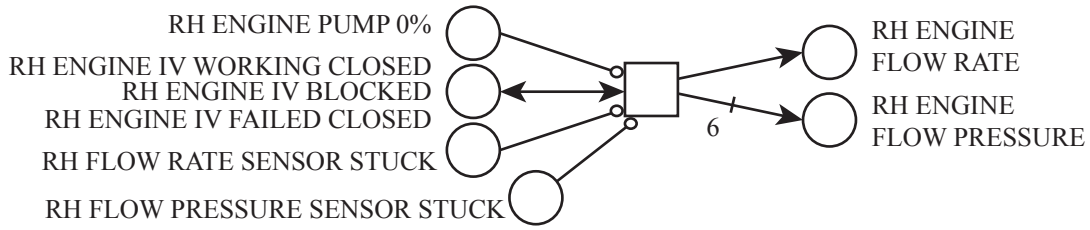

Figure 3.29: RH sensor output - Sensors working, eng IV blocked/closed, eng rating $> 0\%$

Figure 3.29 shows that, so long as the sensors have not failed stuck, a single token will be added to the flow rate place to indicate a zero flow rate while six tokens are added to the flow pressure place to indicate very high pressure. This is the only scenario in which six tokens will be present in the flow pressure place.

The subsequent sets of sub-nets consider the cases where either the flow rate or flow pressure sensors have failed stuck. In these cases tokens are not added to the respective sensor output places. Figure 3.6 showed that when these sensors are stuck the clear transitions are inhibited from changing thereby maintaining the number of tokens in the sensor outlet place.

Figure 3.30 shows transitions where the flow pressure sensor is stuck. In the first sub-net the engine pump rating is 100% and the engine isolation valve is open. In the second sub-net the engine pump rating is greater than 0% and the engine isolation valve is closed. In the third sub-net the engine pump rating is 0%.

The first sub-net in Figure 3.30 is similar to that shown in Figure 3.26 expect there is only an output to the flow rate place. A similar comparison can be made between the second sub-net and Figure 3.29 and the third sub-net and Figure 3.28. In all of these cases the behaviour of the system remains the same except that tokens are not added to the

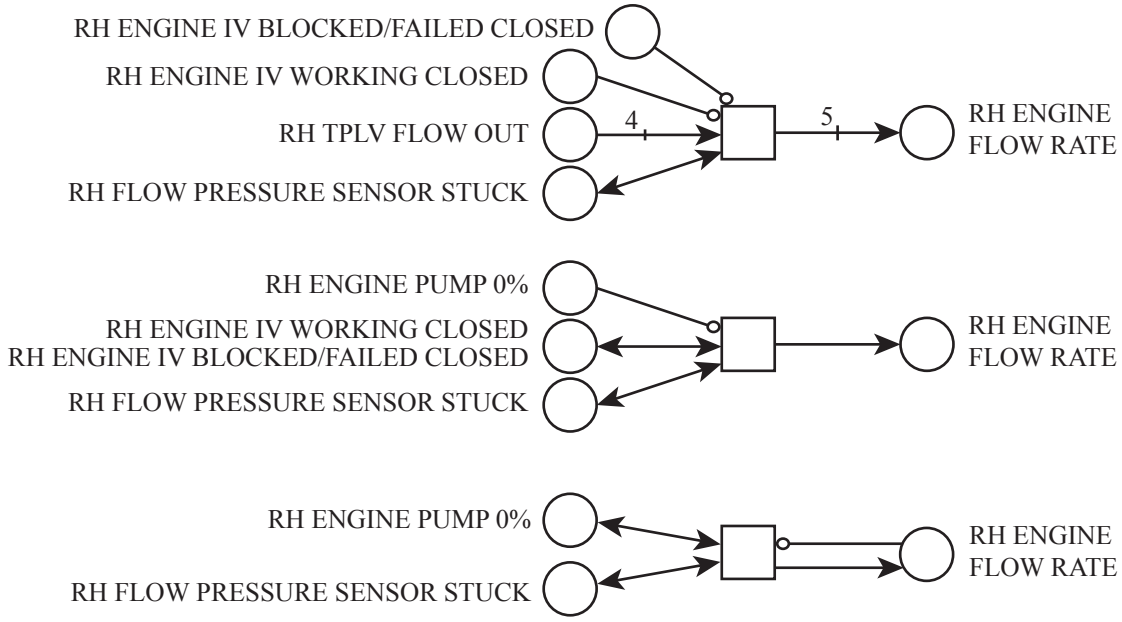

Figure 3.30: RH sensor outputs - Flow pressure sensor stuck

flow pressure sensor output place. In this way, the behaviour of the system in the presence of the fault is accurately modelled.

A similar, but opposite, effect can be seen when the flow rate sensor is failed stuck as can be seen in Figure 3.31.

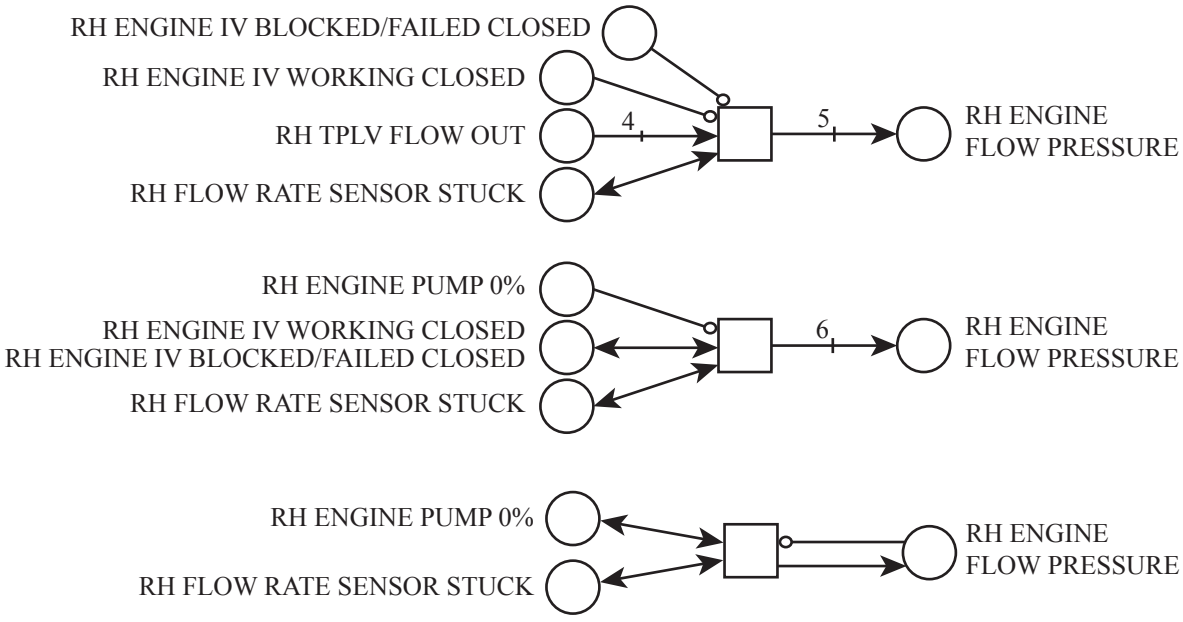

Figure 3.31: RH sensor outputs - Flow rate sensor stuck

Again, the sub-nets in Figure 3.31 are similar to those where no sensor faults are present. The only difference being that no tokens are added to the flow rate sensor output place. Tokens are added to the flow pressure sensor output place in the same number as they were when no sensor faults were present.

If the flow rate sensors have failed high or failed off, the above transitions will still be able to fire. This ensures the correct number of tokens are still added to the flow pressure sensor output place. The relevant transition in Figure 3.32 will then fire to correct the number of tokens in the flow rate sensor output place.

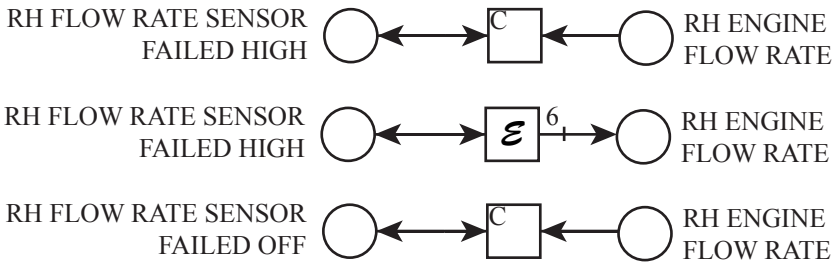

Figure 3.32: RH sensor outputs - Flow rate sensor failed high/off

If the flow rate sensor has failed high, the top transition in Figure 3.32, a 'clear' transition, will fire to remove tokens from the flow rate place. The middle transition will then fire to add six tokens to the flow rate sensor output place. Using this arrangement minimises the number of transitions required to place the correct number of tokens in the output place.

If the flow rate sensor fails off the bottom transition in Figure 3.32 will fire to remove the tokens from the flow rate output place. As mentioned previously, the software will identify a lack of tokens as a sign that the flow rate sensor has failed off.

Equivalent sub-nets to those shown in Figure 3.32 are used to model the system behaviour when the flow pressure sensor fails high or off.

3.3.3.17 Fault Injection

The penultimate sub-nets listed in the input file are those that inject faults into the PN model. Figure 3.33 shows what a typical set of nodes would look like when a 'RH Flow Rate Sensor Failed Off' fault is to be injected into the model, sixty seconds into the mission.

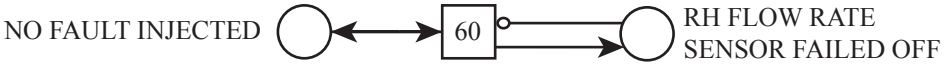

Figure 3.33: Failure mode inject transition

3.3.3.18 Phase Changes

The final set of sub-nets listed in the input file are the phase change nodes. These sub-nets model the change in the operating state of the system either through increasing/decreasing the pump ratings or altering the TPLV settings. Phase change sub-nets have transitions of the specialist 'single transition' type as described in Section 3.3.1.3.

3.3.3.19 Faults Modelled in Software

Figure 3.7 showed how the state of the level sensor is changed in the PN model. The effect of the level sensors failing stuck on the system was shown in Figures 3.18 and 3.22. The effect of the level sensor failing high or low, however, is not shown. The effects of these failure modes are accounted for in the PN software and not the PN model. The PN software evaluates the PN marking when computing the tank levels to determine if any level sensor has failed high or low. If a level sensor has failed high, the PN software will record the tank level as 60cm. If a level sensor has failed low, the the level sensor will record the tank level as 0cm. These actions are described further in Chapter 6.

To model high and low level sensor faults in the PN model would require extra place nodes to represent the actual tank level and the tank level output by the level sensors. However, by accounting for these failure modes in the PN code, the size and complexity of the PN model can be minimised.

3.4 Conclusion

This chapter aimed to introduce the physical system that will be used to develop and test the fault verification technique. The BAE Systems fuel rig system has been described in detail and its modes of operation have also been discussed. Potential component failure modes that can be injected into the fuel rig have been presented. These faults can either be injected into the fuel rig manually or through a computer interface.

In order to model the fuel rig system with a suitable level of accuracy, three specialist

transitions were introduced; clear, if and single. Each of these were introduced and their area of application discussed. The fuel rig PN was then presented in accordance with the order in which the transitions are listed in the PN software input file. Sub-nets were presented that showed how the fuel rig normal operating behaviour has been modelled as well as the fuel rig behaviour in the presence of faults.

The fuel rig PN model can now be used to accurately predict the behaviour of the fuel rig system in a range of operating phases and in the presence of numerous failure modes. Knowing the outputs from the fuel rig sensors, a comparison technique can be applied to the data sets to potentially identify the validity of arisings.

CHAPTER 4

Fault Verification Techniques

4.1 Introduction

The process of verifying arisings from the fuel rig system will be undertaken by comparing the variable outputs from the fuel rig with the predicted variable outputs from the PN model. This chapter will investigate how to compare these variable outputs by assessing a range of different techniques. Techniques will be applied to two different scenarios to determine their ability to compare the recorded and predicted behaviour of the fuel rig with and without faults present. The PN software will be used to conduct all of the analysis.

A number of techniques for comparing the variable outputs have been identified from literature and others designed independently. The standard deviation and dynamic time warping techniques have both been sourced from literature. The point-by-point, delta, binary and time techniques were developed separately. The point-by-point and delta techniques were developed with a focus on the fuel rig tank level variables as, if a fault were to occur in a fuel system, it is likely that the effect of that fault will be seen in the behaviour of the system's fuel tanks and their respective levels.

A thorough description of each technique is given below and the result of their application to the fuel rig with and without faults is also presented. The techniques have been identified and developed as being suitable for fault conditions generated by PBITs. They are therefore applied retrospectively and the system behaviour throughout the fuel rig's operation is known.

4.2 Application of Comparison Techniques to Fuel Rig Details

4.2.1 Application to Fuel Rig

All of the comparison techniques will be tested by evaluating the RH wing tank level data from a mission of the fuel rig in two different scenarios. These scenarios will test the ability of the technique to operate successfully in the presence of a fault and without a fault present. This is representative of actual scenarios that would be encountered by the system. The first scenario will involve the fuel rig undertaking a phased mission with no faults present. In the second scenario the same phased mission is carried out and a leak is induced in the base of the RH wing tank after 90 seconds. Each scenario has been conducted four times to test the robustness of the techniques and the fuel rig data from each mission was recorded.

4.2.2 Mission Description

The phased mission in both scenarios is the same. It is a three phase mission lasting 200 seconds. The mission progresses from phase 1 to phase 2 after 20 seconds and from phase 2 to phase 3 after 180 seconds. In phases 1 and 3 all of the engine ratings on the fuel rig are 0% and there are no flow paths. In phase 2 the RH TPLV is set to on and the RH engine rating is set to 50%. Figure 4.1 shows a set of the recorded and PN predicted RH wing tank level values from scenario 1. Figure 4.2 shows a set of the recorded and PN predicted tank level values from scenario 2. Also shown on Figure 4.2 is a set of PN predicted values, where the leak fault has not been included in the PN model. This illustrates the variation in the tank level behaviour as a result of the leak fault.

4.2.3 Initial Tank Level

In order to run a PN simulation of the fuel rig, the initial tank level values have to be assigned to the system model. While it would be easier to use the first tank level value for each tank

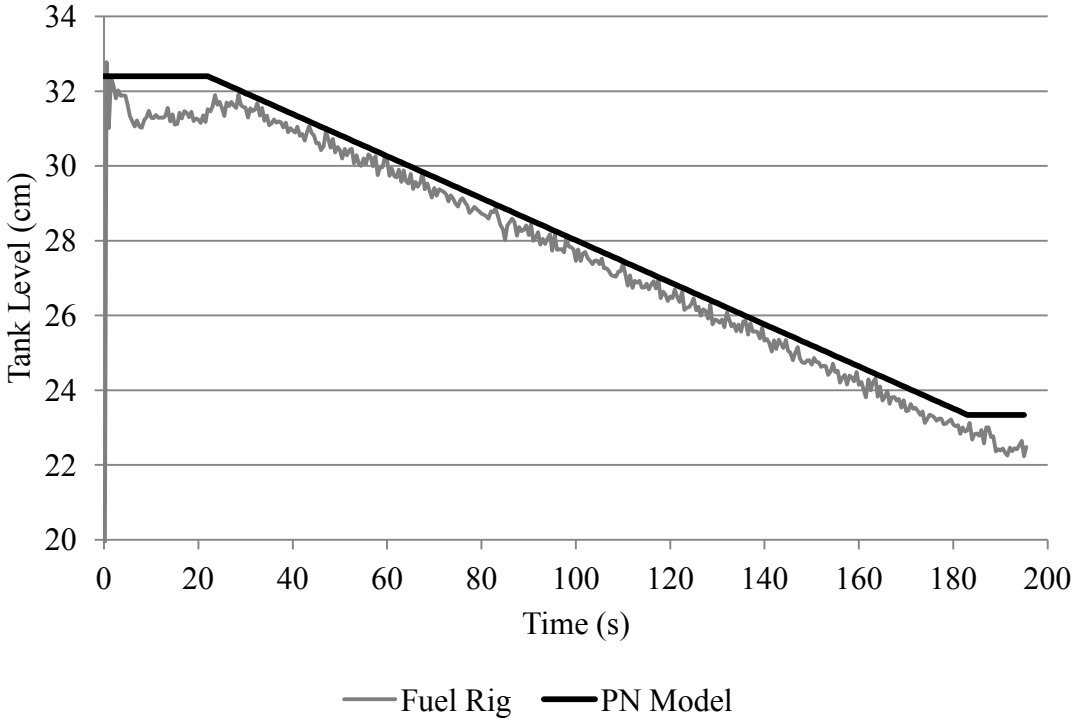

Figure 4.1: Scenario 1 RH wing tank level

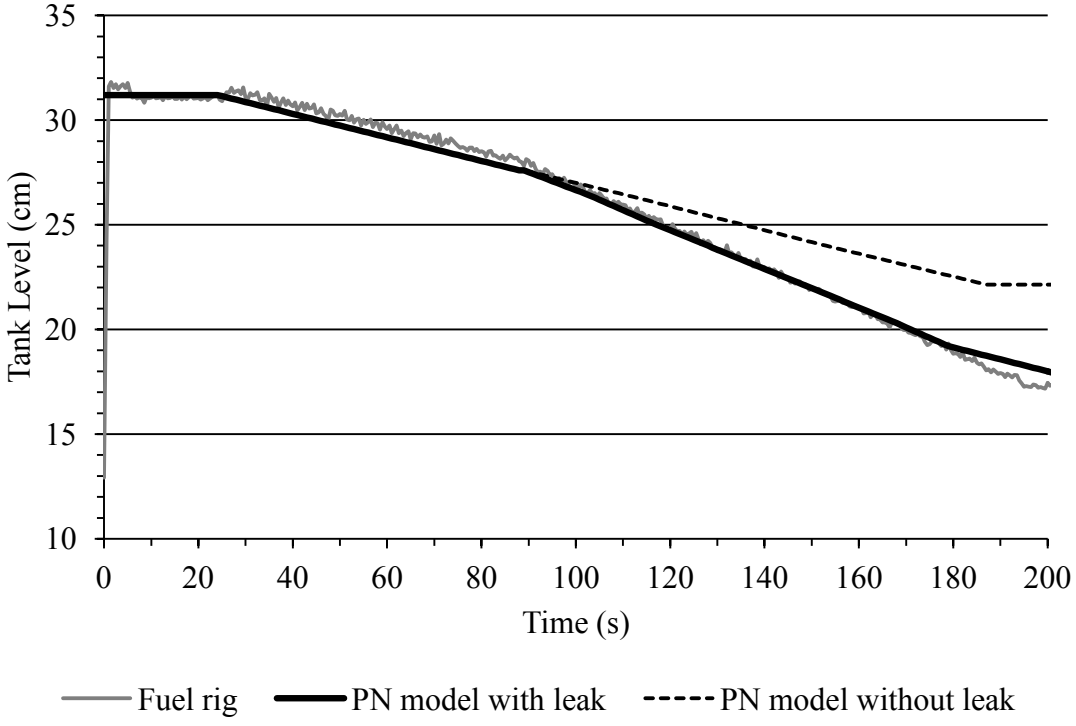

Figure 4.2: Scenario 2 RH wing tank levels with and without leak in PN model

from the data log file, this is not always sensible on the fuel rig. Figure 4.3 shows the RH wing tank level values recorded during four test runs in scenario 1 which illustrates the issue. In all test runs the RH pump was activated after approximately 20 seconds.

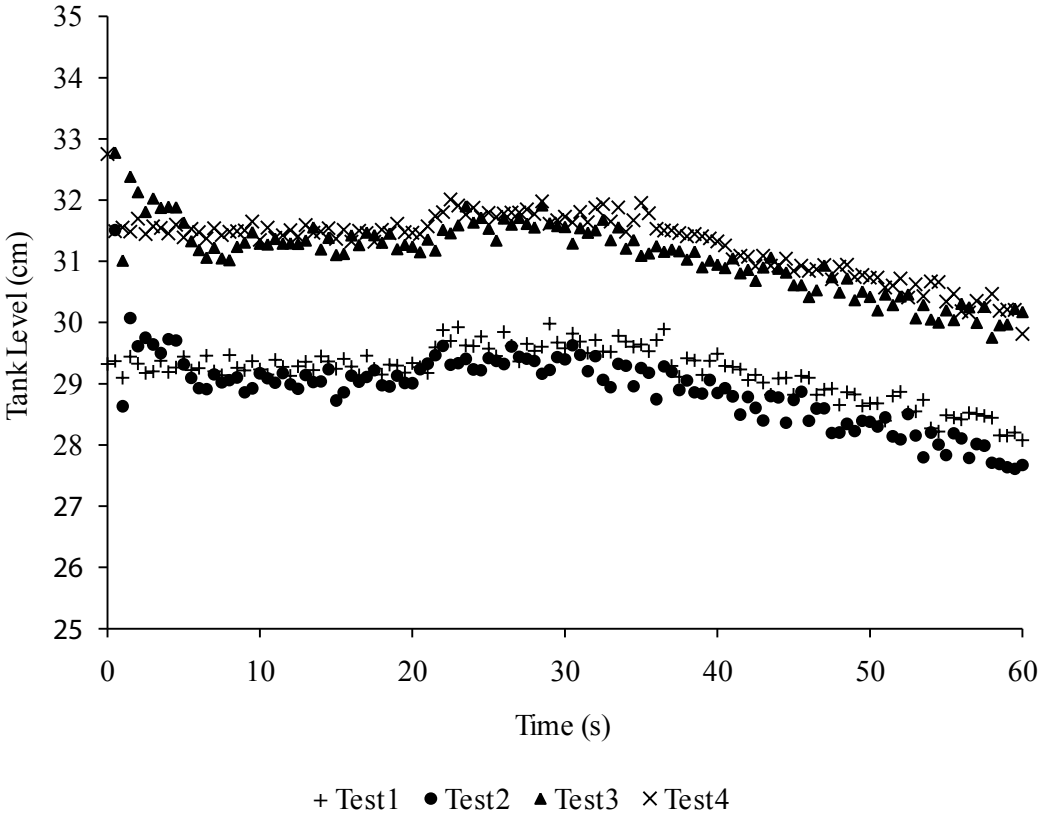

Figure 4.3: RH wing tank level values at pump switch on

It can be seen in Figure 4.3 that the tank level variable has an unstable nature at the start of the mission. This is a result of the time taken for the tank level sensor to fully engage and settle on the actual level of water in the tank. Using the first tank level value listed in the data log as the initial tank level value may, therefore, be inaccurate. Figure 4.3 shows that it takes approximately 5 seconds for the tank level values in all the tests to begin to settle. In order to account for this activity and still use an accurate initial tank level, the PN software averages the recorded tank levels between 5 and 8 seconds. This process will be suitable even if the pump is on or a leak is present from the start of the mission, as the time period being averaged is small enough that minor level changes will not have a significant effect on the initial tank level value assigned to the model.

4.2.4 Phase Change Effects

While evaluating the behaviour of the fuel rig, it was also found that at phase changes the system variables were liable to unexpectedly fluctuate for a short time. This came as a result of the pumps on the system turning on and physically vibrating the system, or turning off and again disturbing the state of the system. This effect can be seen in Figure 4.3 when the engine pumps are turned on after 20 seconds and the tank level values increase before settling at a higher value. A similar effect was seen when many faults were injected into the rig, for example valve blockages, which also cause the system to vibrate. The effect of these short term variable fluctuations is not limited to the tank level variable, but can also be seen in other system variables. As a result of these fluctuations, tolerances were often exceeded unexpectedly.

In order to account for this short term fluctuation, the comparison techniques described in this chapter ignore the first 5 seconds of data immediately after a phase change or arising. In doing so, the system variables on the fuel rig are able to fluctuate and then settle again without requiring every comparison technique to attempt to quantify the variable fluctuations.

4.3 Point-by-Point Technique

4.3.1 Description

The point-by-point comparison technique has been designed to provide a linear comparison of tank levels. Comparisons will be made every second. This enables the most thorough analysis possible, as the PN software utilises larger 1 second timesteps whereas the fuel rig records values every 0.5 seconds.

In order to carry out the point-by-point analysis, vectors of the tank levels recorded from the fuel rig and predicted by the PN model must be constructed in the PN software. The fuel rig tank level data is read directly from the data log file with only data recorded at whole second timestamps saved to a first vector. The PN tank levels are determined from the token counts of the relevant places at each timestep and added to a second vector. Having constructed two equally sized vectors of tank level values, the residual values in the tank level at each time step can be found by subtracting the fuel rig tank level from the value predicted by the PN model at the same timestep. Table 4.1 lists a group of

predicted and recorded RH wing tank level values, as well as the resultant residuals from a fuel rig mission where no faults are present. The data in the table covers a period of five timesteps where there are no phase changes.

Table 4.1: Calculating residual values in point-by-point technique

Time (s)	98	99	100	101	102
Predicted Tank Level (cm)	25.6909	25.6364	25.5818	25.5273	25.4727
Fuel Rig Tank Level (cm)	26.2250	25.9873	25.9302	26.0405	25.5086
Residual (cm)	-0.5341	-0.3509	-0.3484	-0.5132	-0.0359

In order to determine if all of the tank level residuals are small enough, it is necessary to apply a tolerance test. The successful application of tolerances to any variable on a complex system requires detailed testing and analysis, that balances the need for accuracy with the desire to prevent false alarms by allowing for noise in the system. Assuming an initial tolerance of ±1.2cm, the value of one PN tank level token, all of the results in Table 4.1 would pass the tolerance test. However, if the tolerance were reduced to ±0.5cm, the residual at time $t_i = 98$ and $t_i = 101$ seconds would fall outside the tolerance and therefore fail the test. It could therefore be argued that a tolerance of ±0.5cm is overly restrictive as the results, with no faults present, should be similar. However, further testing and analysis would be required to justify this.

All of the residuals determined by the point-by-point technique are subject to a tolerance test and the number of residuals that pass and fail is recorded. To determine if the PN predicted wing tank levels are similar enough to those recorded from the fuel rig, a minimum percentage of passed tolerance tests can be defined, i.e. 97%. Using a value below 100% is reasonable, as it prevents a nominal number of spurious rig results, which may have been subject to higher levels of noise, from influencing the final result. Using a single cut-off percentage will provide a definitive pass or fail result. Alternatively, it would also be possible to use a number of success bands, i.e. 93-95%, 95-97%, 97%+, to describe a level of confidence regarding the similarity of the two sets of data.

4.3.2 Application to Fuel Rig

4.3.2.1 Scenario 1

Having determined the initial tank levels, the relevant simulations can be conducted and the application of the point-by-point technique can be carried out. Table 4.2 shows the percentage of tank level residuals that fall within the 1.2cm tolerance limit.

Table 4.2: Percentage of residuals within tolerance levels - Scenario 1

Test 1	Test 2	Test 3	Test 4
100%	99.5%	99.5%	100%

The results in Table 4.2 show that nearly every residual value was within the tolerance limit for the RH wing tank. Two of the tests produced a 100% success rate, while 99.5% of the residuals in the remaining tests were within the tolerance limit. These results are in line with those that could be expected given the similarity of the curves in Figure 4.1.

The results in Table 4.2 demonstrate that the model provides an accurate representation of the fuel rig system and also increases confidence in its application in future tests. The results also demonstrate that the techniques employed to determine the initial tank levels and deal with the phase change effects, as described in Section 4.2.1, have been successful.

4.3.2.2 Scenario 2

The second scenario considers when a leak is injected into the base of the RH wing tank of the fuel rig. The top line of Table 4.3 shows how many residuals are within the 1.2cm tolerance when the leak is injected in the fuel rig and modelled in the PN. The second row in the table shows the results produced by the point-by-point technique when the leak is injected into the fuel rig but is not included in the PN model. This arrangement, where the fault is only injected in the fuel rig, will show the results of using the point-by-point technique when comparing different data sets.

It can be seen in the top row of Table 4.3 that when the PN model includes the leak failure mode the number of residuals within the tolerance limit is at least 99.0%. This shows that the tank level values in the data sets are similar. When the leak is not injected into the model the number of residuals falls to between 55 and 61%. As the fault is injected

Table 4.3: Percentage of residuals within tolerance levels - Scenario 2

	Test 1	Test 2	Test 3	Test 4
PN Model with Leak	99.5%	99.5%	99.0%	99.0%
PN Model without Leak	55.8%	60.0%	57.6%	60.3%

into the fuel rig after 90 seconds, the outputs from the fuel rig and PN should be similar until this time. The number of residuals within the tolerance limit reflects this. Beyond 90 seconds the two outputs should be significantly different and the results show this to be the case. Figure 4.2 also illustrates this point.

Importantly the above results show that the PN model is accurately representing the behaviour of the fuel rig in both normal operation and with a fault present. The results from the second row of Table 4.3 also show that the technique can identify deviations between outputs. These results, along with those in Table 4.2, suggest that the tolerances being used in this technique are suitable.

One potential disadvantage of the point-by-point technique would be revealed when variables containing high levels of noise are considered. In these cases, where the fuel rig output shows a large amount of variation in a normal operating state, the tolerances applied by this technique would have to be wider to allow for the increased noise levels. As mentioned previously, this increases the risk of genuine faults not being verified. The application of the point-by-point technique could, therefore, be limited to variables that exhibit lower levels of noise.

4.4 Delta Technique

4.4.1 Description

The delta comparison technique uses the change in tank level over time, or tank level gradient, to compare the PN predicted tank levels and those recorded from the fuel rig. The tank level gradient in each phase of the fuel rig's operation is compared using this technique in order to determine the overall similarity of the two sets of data.

The technique uses two tank level vectors, one listing the tank level values recorded from the fuel rig and a second listing those predicted by the PN model. Two further vectors

are then constructed that list the phase start and end times, along with the arising time, if one is present. One of these vectors lists the phase start and arising times with a 5 second delay to account for the fuel rig phase change effects. Table 4.4 shows, using a set of actual phase times, what times are used when calculating the tank level gradients from the PN model and fuel rig tank level data. In the case of Table 4.4 a fault is injected into the fuel rig after 90secs.

Table 4.4: Phase times used to calculate phase gradients

	Actual	PN Model	Fuel Rig
Phase 1	0 – 21	0 – 21	5 – 21
Phase 2	21 – 90	21 – 90	26 – 90
Phase 3 (Arising)	90 – 183	90 – 183	95 – 183
Phase 4	183 – 190	183 – 190	188 – 190

Table 4.4 shows that the gradients are calculated from the PN data in line with the actual phase times. The fuel rig tank data is evaluated after a 5 second delay from the start of the phase or from the arising time. Once the times at which the tank levels will be evaluated are known, the tank levels are identified from the respective vectors. In order to reduce the effect of noise and erroneous values from the fuel rig data, every fuel rig tank level value used is averaged from three data log values. In calculating this value at the phase start time, the tank level at the time searched for and the two subsequent values are used. The averaged value at the end of the phase is found from the tank level at the time searched for and the two previous values. This ensures no values recorded at a phase change are used. The PN model tank level values are not averaged, as they do not contain noise. Having found all of the tank levels at the start and end of each phase of operation and knowing the time at which these phases start and end, the respective phase gradients can be found. Equation 4.1 is used to determine the gradient of each operational phase, where y_2 represents the tank level at time x_2 and y_1 is the tank level at time x_1.

$$m = \frac{y_2 - y_1}{x_2 - x_1} \tag{4.1}$$

Gradient residuals in each phase are then found by taking the fuel rig gradient from the model predicted gradient. Each gradient residual will then be subject to a tolerance

test to ensure that the actual and predicted values are within a suitable range. If all of the gradient residuals are within the tolerance limit, the actual and predicted tank level behaviour is considered to be suitably similar. Determining a suitable tolerance level will come from data analysis of the fuel rig.

4.4.2 Application to Fuel Rig

The accuracy and ability of the delta comparison technique to identify the presence of a fault will now be tested using the two fuel rig scenarios described in Section 4.2.1.

4.4.2.1 Scenario 1

Table 4.5 lists the predicted and actual tank levels at the start and end of every phase in test one. The fuel rig tank levels are the averaged tank level values. The tank level gradients in each mission phase have also been calculated and are displayed in the table. In this scenario, as no faults are present, there are three phases of operation to be considered.

Table 4.5: Tank levels (cm) and phase gradients (cm/sec) - Scenario 1

	Ph. Time	Model TLs	Model Grad	Rig TLs	Rig Grad
Phase 1	0 – 21	30.000 – 30.000	0.000	29.300 – 29.262	-0.0024
Phase 2	21 – 182	30.000 – 21.000	-0.0559	29.625 – 21.086	-0.0530
Phase 3	182 – 191	21.000 – 21.000	0.000	20.908 – 20.596	-0.0782

Table 4.5 shows that the RH wing tank level gradients determined from the PN model and fuel rig data are very similar in phases 1 and 2. There is only a small amount of variation in the gradient values and the tank level values are also similar. There is a greater amount of variation between the gradients in phase 3. This could be due to that fact that phase 3 is relatively short, at 9 seconds long, and therefore any noise in the tank level values will have a greater affect on the gradient value. Figure 4.4 also shows how small tank level changes in a short phase can have a significant effect on the gradient values.

Figure 4.4 shows two curves plotted from fuel rig data. The figure shows data from phase 1 where the tank level should remain constant. In test 1 the tank level decreases by 0.04cm, whereas in test 2 the tank level decrease is 0.45cm. Compared to the tank level of

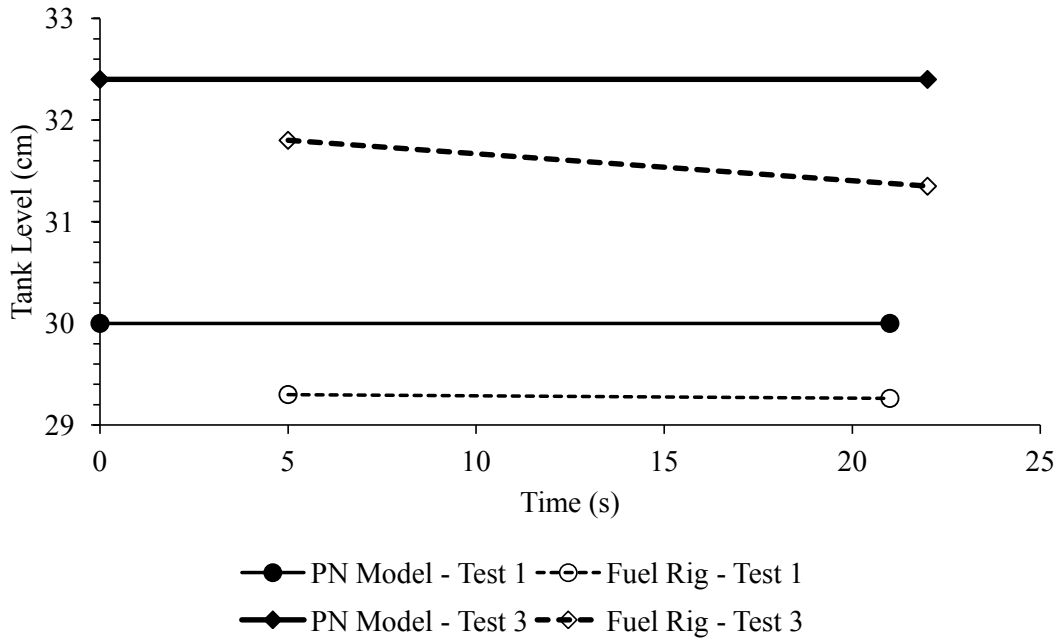

Figure 4.4: Phase 1 tank level phase gradients

approximately 30cm, these changes are very small, however, the short phase length means that the gradient of each curve is significantly different. The test 1 curve has a gradient of -0.002cm/sec, while the test 2 gradient is -0.027cm/sec. These results show that noise can have a significant impact on the results from the delta technique. Table 4.6 shows the gradient residuals in each phase of the four tests conducted in scenario 1.

Table 4.6: Gradient residuals (cm/sec) - Scenario 1

	Test 1	Test 2	Test 3	Test 4
Phase 1	0.0024	0.0060	0.0266	-0.0037
Phase 2	0.0029	0.0033	-0.0038	-0.0030
Phase 3	0.0782	0.0254	0.0100	-0.0222

The phase 2 results in Table 4.6 show that the PN and fuel rig tank level gradients match well, as the residual values are all relatively small. A similar result can be seen in three of the phase 1 results (Test 1, test 2 and test 4). The cause of the anomalous result from test 3 has already been discussed. The results from phase 3 show a consistently larger residual value than in the previous phases. As the engines are off in this phase,

there will be no change in the PN model tank level. The cause of the high residual values can, therefore, be attributed to the fuel rig data and the short phase length. These results give further proof that the delta technique can produce large residual values as a result of noise and in particular as a result of noise and a short phase length.

From these results it can be hypothesised that identifying a suitable tolerance level for the delta technique may be problematic. Considering only the results from phase 2, a tolerance value of ±0.005cm/sec would appear to be sensible. However, if all of the results are considered, and all tests need to pass, a much larger tolerance closer to ±0.03cm/sec would be required.

4.4.2.2 Scenario 2

In scenario 2 a leak is injected into the base of the fuel tank after 90 seconds. As a result all of the tests in scenario 2 have four phases, as the injection of a fault is treated as a phase change. Phase 3 now represents the point from which the leak is present in the system. All other operational phases remain the same. The phase gradient residuals for each of the four tests in scenario 2 with the leak modelled in the PN are shown in Table 4.7.

Table 4.7: Gradient residuals (cm/sec) - Scenario 2

	Test 1	Test 2	Test 3	Test 4
Phase 1	0.0006	0.0058	0.0102	0.0029
Phase 2	0.0103	0.0080	0.0065	0.0140
Phase 3	0.0105	0.0116	0.0138	0.0088
Phase 4	-0.0513	-0.0372	-0.0339	-0.0581

The phase 1 and 2 results in Table 4.7 are in line with those that were produced in scenario 1. This result would be expected given that the fuel rig is in the same operational state as in scenario 1 and there are no faults present. The phase 3 gradient residuals are, on the whole, larger than those from phases 1 and 2 but not significantly larger. There are a number of potential causes for this; the PN may not be accurately modelling the effect of the leak on the tank level or noise could be impacting the fuel rig results. Looking at Figure 4.2 it is likely that noise is causing the increase in the residual values. Noise

and a short phase length again appear to have caused high residual values in phase 4. To determine if the phase 4 residual values could be used as a tolerance guide, the results of not including the leak fault in the PN model are considered. Table 4.8 presents the residuals for the four tests when the leak is injected into the fuel rig but not modelled in the PN.

Table 4.8: Gradient residuals without fault in model - Scenario 2

	Test 1	Test 2	Test 3	Test 4
Phase 1	0.0006	0.0058	0.0102	0.0029
Phase 2	0.0121	0.0085	0.0085	0.0147
Phase 3	0.0449	0.0461	0.0470	0.0436
Phase 4	-0.0353	-0.0346	-0.0285	-0.0420

Table 4.8 shows that the residuals in phase 3, the first phase after the fault is injected, are significantly larger than those in Table 4.7. This shows that the effect of the leak can be seen in the gradient residual results. However, it is also clear that the results in phase 4 have not changed. Despite the differing tank level behaviours being modelled in scenario 2, the residual values in phase 4 remain consistently high. If the tolerance limit was set at the highest residual value from Table 4.7, ± 0.0581, all of the residual results in Table 4.8 would have passed when they should all have failed. Therefore, the results from scenarios 1 and 2 suggest that there are issues with the delta comparison technique in its current form. It has been shown that comparing the PN and fuel rig tank levels over a short phase length does not allow an accurate comparison to be made. Furthermore, the noise in the tank level data has a significant impact on the delta results. Identifying a suitable tolerance level would prove to be challenging and may require different tolerance levels dependent upon a number of factors, i.e. phase length, arising type, etc.

4.5 Standard Deviation Technique

4.5.1 Description

The standard deviation (SD) comparison technique determines the variation that is present in a single data set. In order to compare the results recorded from the fuel rig and predicted

by the PN model, a residual data set will be evaluated to determine the SD value.

Márquez *et al* demonstrated the use of the SD technique as a means of comparing two sets of data from a railway points system [29]. A known 'good' set of data was compared to the actual data set from the railway points to try and identify the occurrence of faults. If the SD of the residual data set exceeded a specified tolerance, a fault was considered to have occurred. The technique was successfully applied by Márquez *et al* to 760 data sets with all faults successfully detected and no false alarms registered.

In order to determine the residual values, the SD comparison technique must first create vectors of the tank level values recorded from the fuel rig and predicted by the PN model. Only tank level values recorded from the fuel rig at whole second timesteps are read from the data file and saved to the vector. The PN tank level values are found in the normal way. Two vectors of equal size will therefore have been created. From these vectors the SD of the residuals can be computed.

Equation 4.2 shows how the SD is found.

$$SD = \sqrt{\frac{\sum_{i=1}^{n}(x_i - \bar{x})^2}{(n-1)}} \tag{4.2}$$

The final aspect of the technique is to apply a tolerance test to the SD value. A tolerance limit was also applied by Márquez *et al* in their work with the railway points system. Identifying a suitable tolerance level, as has been mentioned, will require consideration of the variable behaviour, noise levels and system safety considerations amongst others.

4.5.2 Application to the Fuel Rig

The SD technique will now be used to compare the RH wing tank level curves produced by the PN model and the fuel rig in two scenarios.

4.5.2.1 Scenario 1

The result of applying the SD technique to the four sets of test data from scenario 1 is shown in Table 4.9.

The results of Table 4.9 show that the SD values calculated are very consistent. Given the similarity of the two curves in test 1, as shown in Figure 4.1, these SD results can be assumed to show a good level of similarity between the two curves. The results also

Table 4.9: Standard deviation values - Scenario 1

	Test 1	Test 2	Test 3	Test 4
SD (cm)	0.2829	0.2232	0.1959	0.2647

indicate that a tolerance level of at least 0.3cm would be necessary for the RH wing tank level variable in order to prevent arisings being incorrectly verified.

4.5.2.2 Scenario 2

The SD technique will now be applied to four test data sets from the fuel rig with a leak injected in the base of the RH fuel tank. The PN simulations will be carried out twice. In the first set of simulations the leak fault will be included in the model. In the second set the leak fault will be omitted. This will show how the SD technique deals with differing system behaviour. Table 4.10 shows the SD values produced from the results of each set of simulations.

Table 4.10: Standard deviation values - Scenario 2

	Test 1	Test 2	Test 3	Test 4
SD (cm) - PN Model with Leak	0.4081	0.3818	0.4956	0.4052
SD (cm) - PN Model without Leak	1.7210	1.6781	1.7239	1.7223

The results of Table 4.10 show a clear difference in the SD results when the leak fault is included in the PN model and when it is not. Looking first at the results where the leak is included in the PN model, it can be seen that all of the SD values are higher than those listed in Table 4.9. This shows that there is a greater variation between the recorded and predicted tank level values when a fault occurs in the mission than when there is no fault. Nonetheless, the difference between the SD values is modest and considering the similarity of the respective curves in Figure 4.2, it can be concluded that the SD values in top line of Table 4.10 have been found from a data set of residual values that contain a small amount of variation.

The SD values produced when the leak fault is not included in the PN model are several times higher than when the fault is included. The variation between the respective curves in Figure 4.2 illustrates why these larger SD values have been produced. These results

suggest that the SD technique is sensitive to the difference in data sets being compared. As a result, it may be possible to use wider tolerances with the SD technique, which would reduce the likelihood of false positive arisings being verified. From the results of Table 4.9 and 4.10 a SD tolerance of 1.0cm could be applied without the risk of false positives arisings being generated.

The SD technique method is versatile enough that it could be applied to any variable on the fuel rig system. As long as a residual vector of values can be created, the SD technique can be applied. However, as with the other fault verification techniques considered, unexpected levels of noise in a data set could cause higher than expected SD values to be produced. Despite this, it has been shown that as the technique is sensitive to variations between data sets, wider tolerances may be applied which could reduce/negate this issue.

4.6 Dynamic Time Warping Technique

4.6.1 Description

Dynamic Time Warping (DTW) is a means of comparing two curves, or sets of data, that is considered by some to be more intuitive than a direct comparison, such as that used in the point-by-point technique described in Section 4.3. The DTW technique was first devised in 1983 and is unique in that it allows the time axis to be distorted to account for sources of variation and error between two curves [30]. The original use of DTW was in the field of handwriting recognition [30] [31] however, more recently, Atamuradov *et al* have demonstrated its application in fault diagnosis with a railway turnout/points system [32].

The DTW technique attempts to find the optimal match of two curves by determining the smallest distance between the points on the first curve and the points on the second curve. In order to use DTW it is unnecessary for the two curves to contain the same number of points. Considering the fuel rig system and the tank level variable specifically, the DTW technique will be used to find which data point recorded from the fuel rig is the shortest distance from each data point on the PN tank level curve, a process which will be known as matching.

The first stage in the technique is to write all of the recorded fuel rig tank level values and PN predicted values to individual vectors. The next stage of the technique is to

determine the shortest distance between every data point on the PN tank level curve and a data point on the fuel rig tank level curve. Initially, every data point from the fuel rig curve can be considered when finding the shortest distance. This is known as 'complete matching'. However, DTW can limit the potential number of data point matches in order, for example, to reduce computational demands. Three constraints can be applied, as described below:

1. Continuity Condition

The continuity condition directly controls the number of points on the second curve which can be considered as potential matches for each point on the first curve. Vuori *et al* [33] formalised this in the form of an equation and gave more information on the continuity condition. Figure 4.5 shows how the variable c from the continuity condition equation affects the points that can be matched on two curves of eight and six points respectively. The boxes that are crossed out indicate that a match is not permitted by the continuity condition. The figure shows how increasing the value of c relaxes the constraints on which points can be matched. One noteworthy concern is the higher the value of c, the more calculations will have to be carried out to find the smallest distance between all the possible point combinations, which will increase the computational requirements.

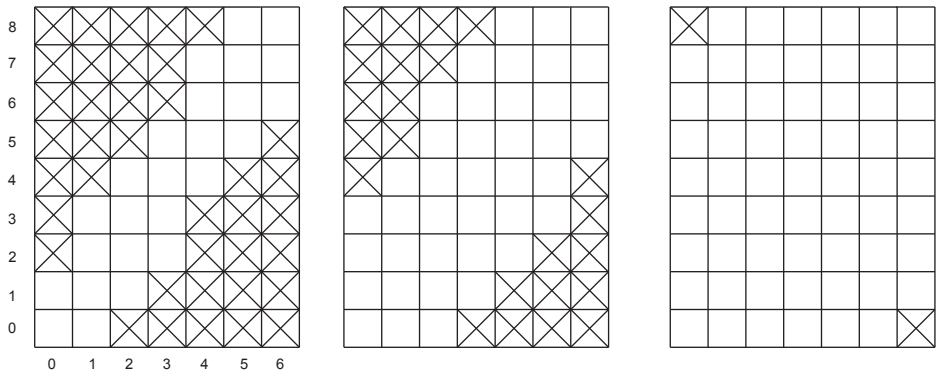

(a) $c = 0.2$ (b) $c = 0.4$ (c) $c = 0.8$

Figure 4.5: Effect of varying c on point matches

2. Boundary Condition

The boundary condition is an optional constraint that ensures that the first and last points of each curve are automatically matched. While this condition may be appropriate in some systems or variables, there is an obvious issue with the fuel rig tank level variable. It has been shown that the first tank level values recorded by the fuel rig are often

erroneous.

3. Monotonicity Condition

The monotonicity condition is another condition that limits possible point matches. Niels [31] describes this in more detail. The monotonicity condition appears to be useful in cases where curves have a circular nature to them, such as in alphabetical letters. However, given such curves are not likely to be found when considering the fuel rig tank level variable, it is unlikely that enforcing the monotonicity condition will have a significant effect on the overall result of the DTW technique.

Having applied the chosen constraints, the distance between each of the possible point matches can be calculated. The distance between the points is measured in the DTW technique by Euclidean distance. Euclidean geometry states that the distance between two points, P = (x_1, y_1) and Q = (x_2, y_2), is found using Equation 4.3.

$$|PQ| = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2} \tag{4.3}$$

The shortest distance between each point on the first curve and a point on the second curve is then summed together and divided by the number of matched points to give the average distance between the curves, also called the DTW value. This value gives an indication of how similar the curves are, while the lower the value the smaller the distance between the curves. To determine whether the DTW value is small enough to consider the curves a good match, the result will be subject to a tolerance test.

4.6.2 Application to Fuel Rig

The DTW technique will now be applied to the two fuel rig scenarios used to assess the comparison techniques. Throughout the application of the DTW technique the continuity condition parameter c will be set to 0.4, as determined from a short period of testing, and the boundary and monotonicity conditions will not be enforced.

4.6.2.1 Scenario 1

Once the PN simulation has been completed, the DTW technique is applied. Table 4.11 shows, for each of the four tests in scenario 1, which points from the fuel rig curve provided the closest match to a range of points at the start and middle of the predicted tank level curve. The values in the table indicate the time, or x-axis value, of the points on the

respective curves. The DTW value for each test is also given.

Table 4.11: DTW point matches - Scenario 1

Model Time	Test 1	Test 2	Test 3	Test 4
0	0	1.5	0.5	0
1	1.5	1.5	1.5	1
2	2	2	2	2
⋮	⋮	⋮	⋮	⋮
100	100	100	100	100
101	101	101	101	101
102	102	102	102	102
⋮	⋮	⋮	⋮	⋮
DTW Value (cm)	0.3300	0.2078	0.4167	0.2446

From Table 4.11 it can be seen that in the lower group, 100 – 102 seconds, all of the PN data points were matched with a fuel rig data point recorded at the same time. In the upper group, 0 – 2 seconds, the matching is not as linear. Given the noise that is present in the fuel rig data at this time, it is not unexpected to find that some of the data point times do not match. Nonetheless, the variation between the two curves can be seen to be a maximum of 1.5 seconds. These results illustrate that when noise is present in the fuel rig data the DTW technique provides some flexibility to deal with this issue. Evaluation of the final DTW values will be considered with the results of scenario 2.

4.6.2.2 Scenario 2

Table 4.12 shows, for each of the four tests in scenario 2, which points from the fuel rig curve are the closest match to a range of points at the start and middle of the predicted tank level curve.

The results in Table 4.12 are similar in both groups to those results listed in Table 4.11. The noise present in the fuel rig data at the start of the mission has resulted in several data points being matched that were not recorded at the same time. However, the variation between the data point times is less than 2 seconds. It could also be concluded from these results, and those in Table 4.11, that the PN model is accurately representing

Table 4.12: DTW point matches - Scenario 2

Model Time	Test 1	Test 2	Test 3	Test 4
0	1	1	1.5	1.5
1	1	1	1.5	1.5
2	2	2	1.5	1.5
⋮	⋮	⋮	⋮	⋮
100	100	100	100	100
101	101	101	101	101
102	102	102	102	102
⋮	⋮	⋮	⋮	⋮
DTW Value (cm)	0.3413	0.3215	0.3971	0.3764

the behaviour of the fuel rig when in normal operation and with a leak in the base of the RH wing tank.

To investigate the sensitivity of the DTW technique, a comparison between a fuel rig tank level curve and a PN tank level curve where the leak is not included in the PN model but is injected in the fuel rig is carried out. All four sets of the fuel rig data with a leak present are analysed and Table 4.13 shows the respective data points that produce the shortest Euclidean distance for three small sections of the mission, and the overall DTW values.

Table 4.13 shows that the matching points in each of the sections considered are close together and never more than 1.5 seconds apart. This is true even after the leak has been injected into the fuel rig after 90 seconds and the curves diverge, as shown in Figure 4.2. In the third section, long after the leak has been injected into the fuel rig, the largest variation in the x-axis values of the matching points is 0.5 seconds. Given the difference between the curves at these times, it might have been expected to see greater variation in the matched points x-axis values. The cause of these results can be attributed to the fact that the Euclidean distance equation gives equal weighting to the distance between points on the x-axis, as it does to the y-axis. This means that, unless there is a significant difference between multiple points on the y-axis at the same time, due to persistent noise or many erroneous values in the data log for example, point matches will always be made

Table 4.13: DTW point matches without fault - Scenario 2

Model Time	Test 1	Test 2	Test 3	Test 4
0	1	1	1.5	1.5
1	1	1	1.5	1.5
2	2	2	1.5	1.5
⋮	⋮	⋮	⋮	⋮
100	100	100	100	100
101	100.5	101	101	101
102	102	102	102	102
⋮	⋮	⋮	⋮	⋮
174	173.5	173.5	174.5	174
175	175.5	175.5	174.5	175
176	175.5	176.5	175	175
⋮	⋮	⋮	⋮	⋮
DTW Value (cm)	1.4799	1.4015	1.4466	1.4445

with very similar x-axis values.

A further conclusion that can be made from Table 4.13 relates to the DTW values. The result of every test shows that the DTW values are all significantly higher than those in Table 4.11 and Table 4.12. All of the DTW values in the previous tables were less than 0.5cm. By comparison all of the DTW values in Table 4.13 are greater than 1.4cm. Given these results, a tolerance of 0.7cm could be safely proposed for the RH wing tank level variable. The results from Table 4.12 and 4.13 indicate that using this tolerance value would prevent false positives from occurring.

The DTW results have shown that there is some flexibility in the technique and that it can be used to overcome some issues caused by the presence of noise in variables. It was possible to use the results to confirm the presence of the leak in the wing tank. The results were also shown to change significantly, when the leak fault was present in the fuel rig but not in the PN model. One potential issue with the DTW technique is the number of calculations that are required, although this is dependant on the constraint conditions applied. Compared to the other techniques considered, significantly many more

calculations are necessary in order to generate a result. When considering all of the fuel rig variables, this requirement could increase computational needs significantly.

4.7 Binary Technique

4.7.1 Description

The comparison techniques considered up to this point have dealt with the tank level variable alone. The binary technique, however, was designed to allow the output from the low and high level switches to be compared.

The binary technique compares the state of the low/high level switch only at the end of the mission. The switch state from the fuel rig is read from the data log file, as either on or off. In the PN model, however, the switch states can be working on/off or failed on/off. In order to account for the greater number of model states, a positive, or 'on', output from the fuel rig is assumed to be verified if the PN contains a token in the switch 'working on' or 'failed on' places. It is necessary to include the 'failed on' place, because without it, if the 'high/low level switch failed on' fault were to occur, the binary technique would fail to correctly verify it. Similarly, a negative, or 'off', output is verified by a token in the 'working off' or 'failed off' PN places. If the predicted and actual state of the low/high level switches match, the binary test is considered to have been passed.

The greatest advantage of the binary technique is that it has very low computational requirements. It may, therefore, be of benefit to use a variation of the binary technique to compare other system variables at the end of or throughout a period of fuel rig activity. Possibilities include; at the start and end of every phase, or every x seconds. Such a comparison, although very limited in detail, could provide a quick check to ensure that the respective data points are within a sensible range before a more computationally demanding and exhaustive technique is applied.

Although the binary technique allows variables beyond the tank level to be compared, its application with the low and high level switches could be limited. For safety reasons the fuel tanks on many systems are not permitted to 'run dry' or become completely empty. As a result, the low level switches should always be on. Also, when in operation, the high level switch is likely to be off for the majority, if not all, of the mission. As a result, comparing the low and high level switch states will allow only a small number of faults

to be verified. The suggested application of an adjusted binary technique as an initial comparison technique, however, could prove useful if the computational times associated with other techniques become too great.

4.8 Time Comparison Technique

4.8.1 Description

The time comparison technique is an extension of the binary technique, that considers the state of low and high level tank switches. The distinctive feature of the time technique is that it compares the actual and predicted time at which the state of a switch changes.

Before comparing the time at which a switch changes state, the time comparison technique assesses the PN predicted results and fuel rig data log to ensure that a change of state has occurred in both results. Once it has been identified that a state change occurs, the time at which the state change occurred in the fuel rig and PN model is determined. If the difference between these times is within a set tolerance, this stage of the test is passed. The state of the switch at the end of the mission must also be consistent in both the recorded results and predicted data in order for the result to stand.

The time comparison technique assumes that the tank switches can only change state once over the time period under consideration. In the case of the high and low level switches on the fuel rig this is reasonable, as once either switch turns off it would require fuel to be created for them to change state again.

The time comparison technique is similar to the binary technique, since it has very low computational requirements. However, it is also similar in the sense that it has only limited application potential in terms of fault verification. There is only a small number of faults that can be verified by the time comparison technique alone and the techniques considered in Section 4.3 to Section 4.6 may be able to verify these faults with greater accuracy.

4.9 Fault Verification Technique Selection

The aim of this chapter is to identify the most suitable technique for the comparison of fuel rig and PN variable outputs in order to verify arisings. While all of the techniques considered above displayed a positive comparison feature, that in some cases is highly

specific, several exhibit poor performance characteristics. The binary and time comparison techniques have very low computational requirements, however do not provide a similar level of detail, compared to the other techniques. The delta technique was shown to have issues when dealing with even small amounts of noise, which are constantly present on the fuel rig. Another technique that was identified to have potential issues with noise is the point-by-point technique. Of the two remaining techniques, the DTW technique was identified as the one with far greater computational demands. Additional time would also have to be spent determining which constraint conditions to apply to which variables and, if the continuity condition were to be applied, identifying the optimum value of c. As a result, the SD technique was identified as the most robust and suitable for application to the fuel rig system. The technique was successfully applied to the tank level variable and it is clear how this can be extended to the other variables on the fuel rig system. It has also been used in the process of variable comparison in the past and will be used in this work going forward. Tokens representing different variables in the fuel rig PN will have unique values associated with them. A token in a tank level place for example, will have a different value from a token in a flow rate place. Establishing the true value of each variable in the PN will therefore require a unique piece of code for each variable. Once these values have been established however, they are all compared in the same manner using the SD technique presented in Section 4.5.

4.10 Fuel Rig Specific Features

In order to make the SD fault verification technique as effective as possible, a number of fuel rig specific features have been developed for in the PN software. Implementing these will improve the accuracy and capability of the technique.

4.10.1 Auxiliary Tank Vibration Effect

The engine and auxiliary pumps are represented on the fuel rig using peristaltic pumps. When operational, the rotational motor within these pumps causes the fuel rig to vibrate. The effect is greater as the pump rating is increased and the motor turns at a greater speed. Testing identified that the effect of this vibration was greater on the auxiliary tanks than on the wing tanks. Figure 4.6 shows the effect of the auxiliary pump vibrations on the

RH auxiliary fuel tank level. The figure shows the tank levels recorded from a phased mission, where the auxiliary tank isolation valve is closed to prevent any fluid leaving the tank. The auxiliary pumps are only active in phase 3.

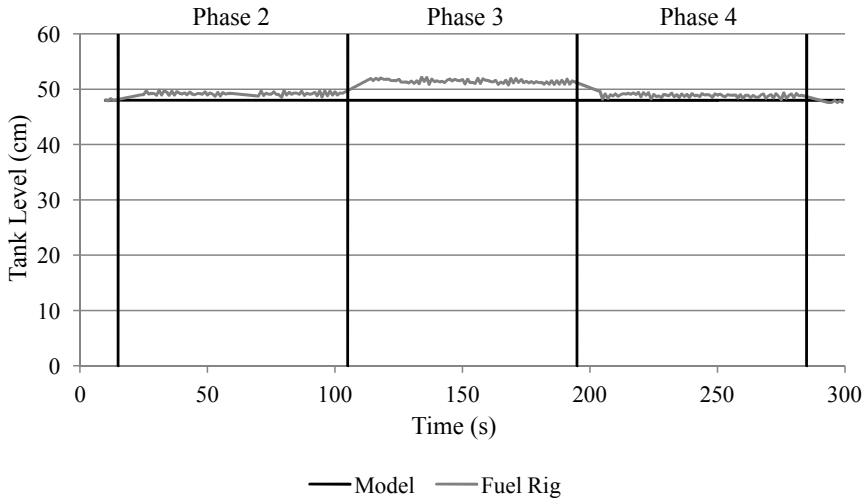

Figure 4.6: RH auxiliary tank level

It can be seen that at the start of phase 3, as the auxiliary pump rating is established, there is a significant increase in the tank level. A similar, but smaller, effect can be seen at the start of phase 2 when the engine pump demand is established. Both of these increases are related to the pump induced rig vibrations that cause the fluid in the fuel tanks to shake. As would be expected once the pump demands are removed, and the rig stops shaking, at the start of phases 4 and 5 the tank level falls. Figure 4.7 shows the RH wing tank level over the course of the same phased mission. There will be no flow into the wing tank.

Figure 4.7 shows an increase in the wing tank level at the start of both phases 2 and 3. A decrease in the wing tank level can be identified at the start of phases 4 and 5. These results indicate that the pump induced system vibrations are affecting both the auxiliary and wing tanks on the fuel rig. While the engine pump alone is having an effect on the tank levels, there is a much larger effect when both the engine and auxiliary pumps are on. It is therefore sensible to attempt to quantify the effect of the vibration when both pumps are on. This will enable accurate comparisons to be made between the fuel rig and PN predicted tank level behaviour.

The vibration effects will be quantified by comparing the recorded tank levels, when only the engine pumps are on and when the engine and auxiliary pumps are on. During the

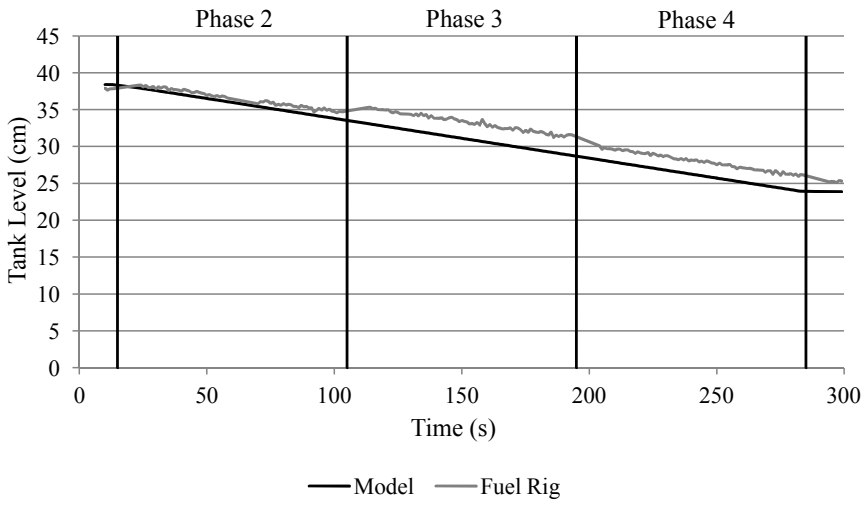

Figure 4.7: RH wing tank level

tests the auxiliary and engine tank isolation valves were closed to prevent any flow leaving the tanks. The tests were performed at a number of tank level heights to investigate whether the vibration effect varied with tank level. Figure 4.8 and Figure 4.9 show the tank level values for the situation with the auxiliary engine off vs auxiliary engine on, for the RH auxiliary and wing tank respectively. Both figures show a clear increase in the tank level when the auxiliary tanks are on, and that the effect is greater at higher tank levels. Using the gradient of a linear trend line through the points plotted when the auxiliary engine is on, the vibration effect on each of the four tanks can be quantified.

Using a process of trial and error, it was identified that the effect of the vibration on both wing tanks could be accounted for by a single equation. Therefore during the process of fault verification, all wing tank level values predicted by the PN when both the wing and auxiliary pumps are on, are multiplied by 1.033. The auxiliary tanks have to be considered individually, with LH auxiliary tank levels multiplied by 1.05 and the RH auxiliary levels subject to the formula shown in Equation 4.4. In the equation, L' represents the adjusted RH auxiliary tank level and L represents the initially predicted tank level.

$$L' = (L * 1.08) + 0.3 \tag{4.4}$$

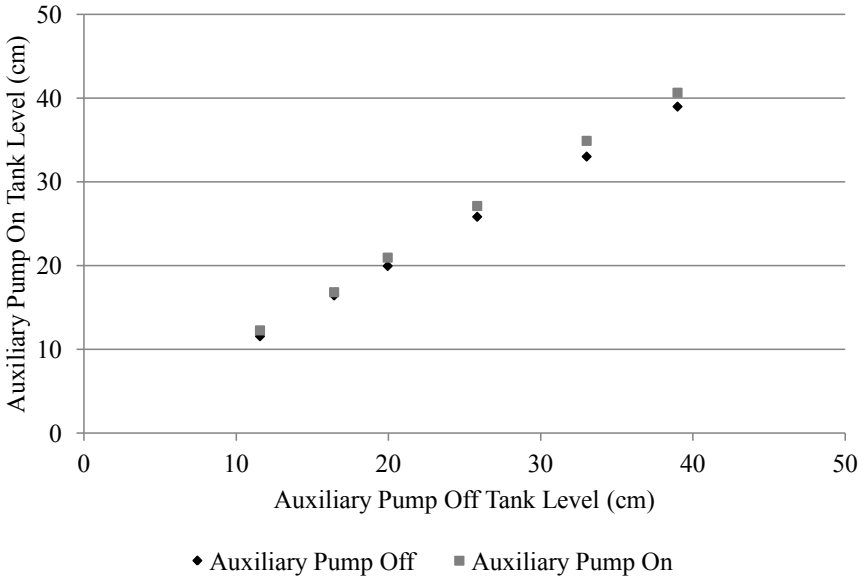

Figure 4.8: Vibration test RH auxiliary tank level

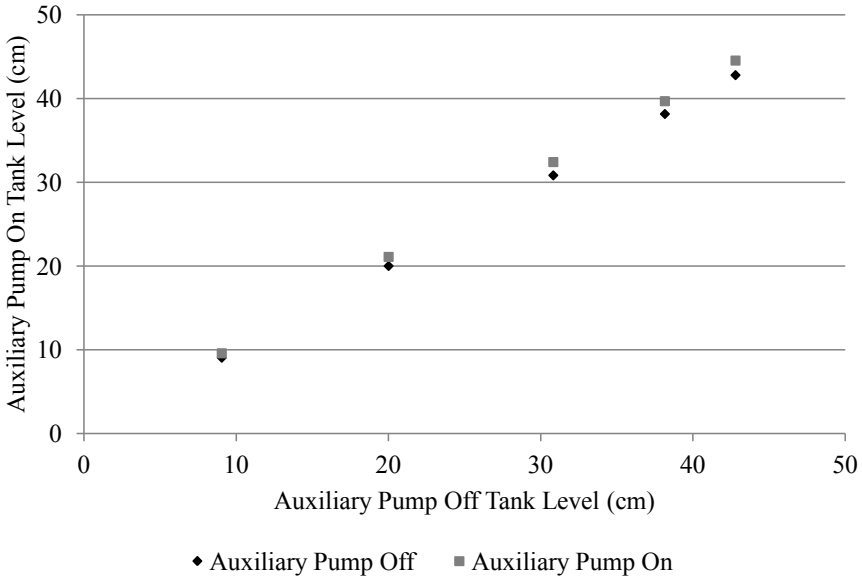

Figure 4.9: Vibration test RH wing tank level

4.10.2 Leak Faults

The only fault type that will not be evaluated using the SD technique, described in Section 4.5, is a tank leak. By considering the leak faults separately it is possible to determine the size and location of the leak when verifying the fault. Determining these additional pieces of information using the SD technique would require every possible size and location of

leak to be evaluated. Leak faults are evaluated using only data recorded from the fuel rig.

The process created to verify a leak arising considers both the tank level and flow rate variables of the tank under consideration. In order to compare the two variables directly requires flow rate monitors at the input to and exit from the tank. Initially the flow rate data must be converted into tank level values. Equation 4.5 is used to convert the flow rate data (FR_i) at every timestep (Δt) in the mission into the volume of liquid that leaves/enters the tank (V). Both flow rates out of and into the tank under consideration must be subject to Equation 4.5. Equation 4.6 then converts the volume into a change in the tank level (L') using the tank cross sectional area (CSA). Given the initial tank level, found using the method presented in Section 4.2.3, this value can be used to help calculate the tank level from the flow rate outputs throughout the mission.

$$V = \frac{FR_i + FR_{i+1}}{2} * \Delta t * 1000 \tag{4.5}$$

$$L' = \frac{V}{CSA} \tag{4.6}$$

In Equation 4.5 the flow rate data is measured in L/sec. The volume is expressed in cm^3. The cross sectional area of the tank is measured in cm^2 and the change in tank level is expressed in cm.

To reduce the noise effects seen in the tank level output, a 10-point moving average has been applied to the level sensor data. This filter determines a tank level by averaging the nine previous data points with the 10th point under consideration. This significantly reduces the noise in the level sensor output and enables a more accurate comparison with the flow rate data to be carried out.

To verify the presence of a leak arising, the tank level gradients prior to and after the arising time are assessed. Gradients, m, are calculated from the data recorded by both the level sensor and flow rate data and Equation 4.1. The tank level gradient prior to the arising is found using data points from the start of the phase in which the arising occurs and at the arising time. The first data point is taken as that 20 seconds after the phase start time. The second data point is at the time of the arising. The tank level gradient after the arising is found using data points 20 and 30 seconds after the arising. The 20 second delay from the start of the phase/arising is necessary to allow phase change effects and the moving average results to settle.

The gradient residual before and after the arising is then found. The gradient residual is found by subtracting the flow rate tank level gradient from the level sensor value. If a leak is present, the gradient residual after the arising will be lower than that prior, as only the level sensor output after the arising will show any leak effects. To allow for the presence of noise in the tank level variable, the gradient residual must fall by 0.019cm/sec in order for a leak to be verified. This value was identified from testing of a variety of leak sizes and locations on the fuel rig. Equation 4.7 expresses how a leak is verified, where $R_{Grad-Pre}$ is the gradient residual prior to the arising, and $R_{Grad-Post}$ is the gradient residual after the arising.

$$R_{Grad-Post} < R_{Grad-Pre} - 0.0190 \tag{4.7}$$

The final step, having verified the presence of a leak, is to identify the location or height of the leak. As was shown above, when a leak occurs the gradient residual value decreases. It follows then that, if the tank level were to fall below the height of the leak, the gradient residual value would increase to a value approaching that found prior to the leak appearing. It was shown previously, that in order for a leak to be verified the gradient residual value after the arising had to be at least 0.019cm/sec lower than the value prior to the report. Therefore, in order to confirm the tank level has fallen below the leak height, the residual gradient must be greater than the residual gradient prior to the arising less 0.019cm/sec. The leak height will have been identified if Equation 4.8 is satisfied. In the equation $R_{Grad-Interval}$ is the gradient residual at an interval some time after the arising.

$$R_{Grad-Interval} > R_{Grad-Pre} - 0.0190 \tag{4.8}$$

To find the leak height as accurately as possible, gradient residual values are found at 15 second intervals starting from the last time considered to find the post fault report gradients, i.e. 30 seconds after the arising. If Equation 4.8 is satisfied by any of these, then the leak height is determined from the level sensor outputs at that time. If a leak is verified but the gradient residuals never exceed the minimum gradient residual, it is possible that the leak could be present anywhere between the base of the tank and the tank level at the end of the mission.

The leak arising technique will now by demonstrated by means of an example. A five phased mission was undertaken by the fuel rig and a leak was injected into the side of the

LH auxiliary tank of the fuel rig after 60 seconds. Only in phase 3 the LH auxiliary pump rating is set to 75%. In all other phases the pump rating is 0%. For this test a flow rate meter was placed at the outlet from the LH auxiliary tank, thereby ensuring that the flow rates out of the auxiliary tank were measured. Figure 4.10 shows the tank levels for the LH auxiliary tank over the course of the mission, as determined from the moving averaged level sensor data and the flow rate data.

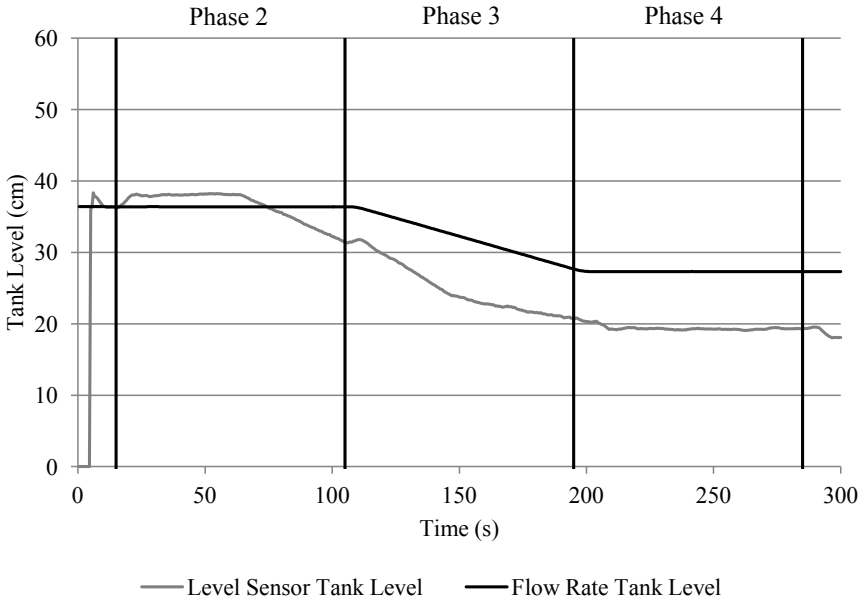

Figure 4.10: LH auxiliary tank levels

Figure 4.10 shows that the effects of the leak are only visible in the level sensor data. As none of the leak flow passes through the flow rate meter, its effects are not captured in the flow rate data. The tank level calculated from the flow rate data only falls in phase 3, as expected, when the auxiliary pump is on. At the start of phase 3 the level sensor curve also becomes steeper indicating an increased flow out of the tank. This effect, however, only lasts until the middle of the phase, when the gradient becomes more gradual. This change is a result of the tank level falling below the height of the leak and no longer having an effect. Beyond this point the gradients of the level sensor and flow rate tank level curves are similar.

In the phased mission considered above the arising occurs in the second phase of the mission, which began after 15 seconds. The leak arising was recorded at 60 seconds. The pre arising tank level gradients are, therefore, found using the tank level and flow rate sensor data points at 35 and 60 seconds. The post arising gradients are found using the

data points at 80 and 90 seconds. Table 4.14 lists the tank level gradient values found using Equation 4.1 and the gradient residual values for the phased mission being considered.

Table 4.14: Level sensor and flow rate determined tank level gradients

	Pre-Arising (cm/sec)	Post-Arising (cm/sec)
Level Sensor	-0.0015	-0.1752
Flow Rate Sensor	-0.0002	-0.0001
Gradient Residual	-0.0013	-0.1751

It can be seen that while the gradients determined from the flow rate data show only a small amount of change, the gradients determined from the level sensor data show a much larger amount of change. The gradient residual value has decreased by 0.1738cm/sec due to the leak. As Equation 4.9 has been satisfied, a leak can be verified. The size of the leak is equivalent to the change of the gradient residual values, 0.1738cm/sec. To determine the height of the leak the $R_{Grad-Interval}$ term in Equation 4.10 has to be found. Inserting the relevant values it can be found that, if the gradient residual at any interval is greater than -0.0203cm/sec, then the tank level will have dropped below the leak height. Table 4.15 lists the interval gradient residuals for the phased mission. Intervals including data that falls within the first ten seconds of a phase are ignored due to phase transition effects.

The results show that the gradient residuals calculated at the first three intervals are all lower than the value required to identify the leak height. The interval from 150 – 165 seconds represents the first time that the gradient residual is greater than -0.0203cm/sec. It can also be seen that all subsequent values are greater than this value. It is clear therefore that at 150 seconds the tank level has fallen below the height of the leak. The leak height can therefore be found from the level sensor data at 150 seconds. A leak height of 23.6cm is found by averaging the tank level values over the first 3 seconds of the interval. This result also matches well with the observations that were made of the tank level curves in Figure 4.10.

4.10.3 Fuel Rig Variable Tolerance Limits

In order to fully apply the SD verification technique to the fuel rig system, it is necessary to determine the tolerance limits for each of the system variables. In Section 4.5, a

Table 4.15: Residual values after arising

Interval (sec)	Interval Residual (cm/sec)
90 – 105	-0.1592
120 – 135	-0.1127
135 – 150	-0.0865
150 – 165	0.0190
165 – 180	0.0415
180 – 195	0.0488
210 – 225	0.0038
225 – 240	-0.0063
240 – 255	-0.0004
255 – 270	-0.0015
270 – 285	0.0052

tolerance limit for the RH wing tank level variable was proposed based on the results shown. However, the behaviour of each variable must be considered in the presence of many possible failure modes on the system and in a range of operating modes in order to determine the most suitable limits. Table 4.16 lists all of the SD tolerance limits that have been chosen for the fuel rig variables. These tolerance limits have been identified to provide the best differentiation between faulty and non-faulty behaviour in the fuel rig system. The application of these tolerances with the overall fault verification technique should allow arisings generated on the fuel rig to be correctly identified as either genuine or false.

4.11 Conclusion

This chapter has considered a number of variable comparison techniques that could be used to verify the occurrence of faults on the fuel rig system. Two of the techniques, SD and DTW, were found to offer a superior comparison of the predicted and actual behaviour of the tank level variable. A number of other techniques showed greater evidence of issues with noise, accuracy and failed to provide as detailed an analysis. All of the techniques were applied to two scenarios involving the fuel rig system undertaking a phased mission.

Table 4.16: Fuel rig SD tolerance limits

Fuel Rig Variable	Tolerance
LH Auxiliary Tank Level	1.500cm
RH Auxiliary Tank Level	1.500cm
LH Wing Tank Level	1.500cm
RH Wing Tank Level	1.500cm
LH Flow Rate	0.30L/min
RH Flow Rate	0.30L/min
LH Fuel Flow Pressure	9,000Pa
RH Fuel Flow Pressure	9,000Pa
LH Wing Tank High Level Switch	0.1
RH Wing Tank High Level Switch	0.1
LH Wing Tank Low Level Switch	0.1
RH Wing Tank Low Level Switch	0.1
RH Auxiliary Tank High Level Switch	0.1
RH Auxiliary Tank Low Level Switch	0.1

In the first scenario no faults were considered. In the second scenario a fault was injected in the fuel rig system and modelled in the PN. The SD and DTW techniques demonstrated an ability to accurately compare the fuel rig and PN tank level behaviour, when no faults were present and in the presence of faults. Owing to the ease with which the SD technique could be applied to the other variables on the fuel rig system, it was chosen as the most suitable fault comparison technique. The SD technique also created smaller computational demands compared to the DTW technique.

The penultimate section of this chapter considered features of the fault verification technique that are specific to the fuel rig system. It was found that the peristaltic pumps installed on the system cause the fuel rig to vibrate and the water in the tanks to shake. This effect was amplified when both the auxiliary and engine pumps were active. A number of formula were, therefore, derived to allow the PN predicted tank levels to account for this phenomena, when all the pumps were active. A specific leak fault verification technique was also introduced. This technique not only verifies the presence of a leak, but can also estimate the leak size and, where possible, the leak height in the tank. Finally, the fuel

rig variable SD tolerances were also presented.

The complete fault verification technique can now be used to assess the legitimacy of arisings produced by the fuel rig system.

CHAPTER 5

Fuel Rig System Results

5.1 Introduction

Previous chapters have identified a modelling technique and variable comparison technique with which to verify arisings from complex systems. The operation of the BAE Systems fuel rig was described and modelled in Chapter 3. All of these resources will now be combined to verify arisings from fuel rig.

This chapter will consider the results of applying the fault verification process to arisings generated during a phased mission of the fuel rig system. All of the first order faults listed in Table 3.1 have been induced in the fuel rig and its behaviour has been recorded in the form of the output variables. Using the PN model of the fuel rig and the SD comparison technique, the ability of the fault verification process to correctly verify these faults is demonstrated. A scenario where the fault verification process is used to identify a single genuine fault from a list of arisings is also investigated. A number of second order faults are considered. The results will also consider the ability of the fault verification technique to identify false arisings.

5.2 Phased Mission Description

All of the failure modes considered in this chapter were individually injected into the fuel rig while it was progressing through the same phased mission. The mission undertaken has five phases and a duration of 300 seconds. The system pump demands and TPLV states throughout the mission are displayed in Table 5.1. In phase 2 the TPLVs are set to ON and a 50% demand is applied to both engine pumps. This creates flow paths from

the LH and RH wing tanks to the LH and RH engines respectively. Phase 2 lasts for 90 seconds.

All of the first order failure modes were injected after 60 seconds, during phase 2. Injecting the faults at this time allowed the effect of the fault to fully propagate through the model. It also allows the fault to be present while the system operates in several different phases. This will reduce the number of hidden failure modes considered.

Table 5.1: Fuel rig pump and valve states in phased mission

	Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
LH Auxiliary Pump	0%	0%	75%	0%	0%
RH Auxiliary Pump	0%	0%	75%	0%	0%
LH Engine Pump	0%	50%	50%	50%	0%
RH Engine Pump	0%	50%	50%	50%	0%
LH TPL-Valve	OFF	ON	ON	ON	OFF
RH TPL-Valve	OFF	ON	ON	ON	OFF
Phase Length	15s	90s	90s	90s	15s
Cumulative Mission Length	15s	105s	195s	285s	300s

5.3 Normal Operating Behaviour

The ability of the fault verification technique to assess the legitimacy of arisings is based on the fuel rig variable outputs and the PN model prediction of those outputs. The figures below, from Figure 5.1 to Figure 5.13, show these outputs, as recorded from the fuel rig and as predicted by the PN model, when no faults are present in the system. These 'clean' outputs will be used as a baseline performance indicator that future outputs can be compared to in order to identify the effects of faults on the system.

Table 5.2 lists the SD values calculated from the recorded and predicted fuel rig system variable outputs in the 'clean' arrangement. None of the SD values exceed the respective tolerances. These results can now be used to see the effect of different failure modes on the performance of the system in the phased mission.

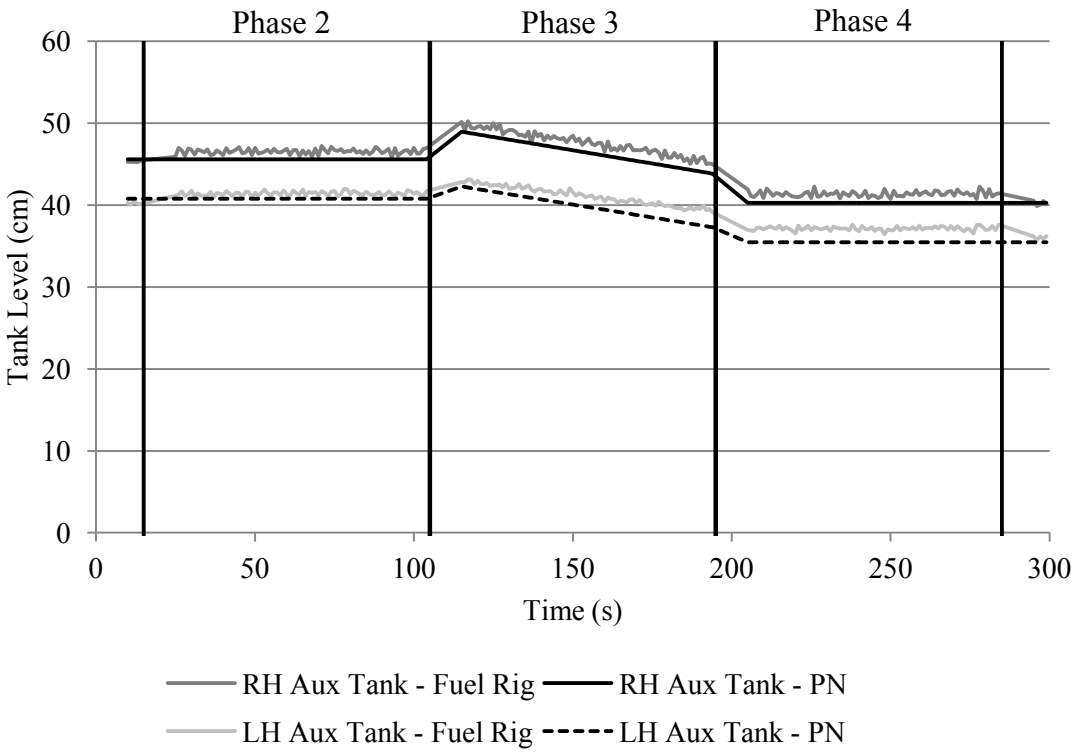

Figure 5.1: 'Clean' fuel rig arrangement - Auxiliary tank levels

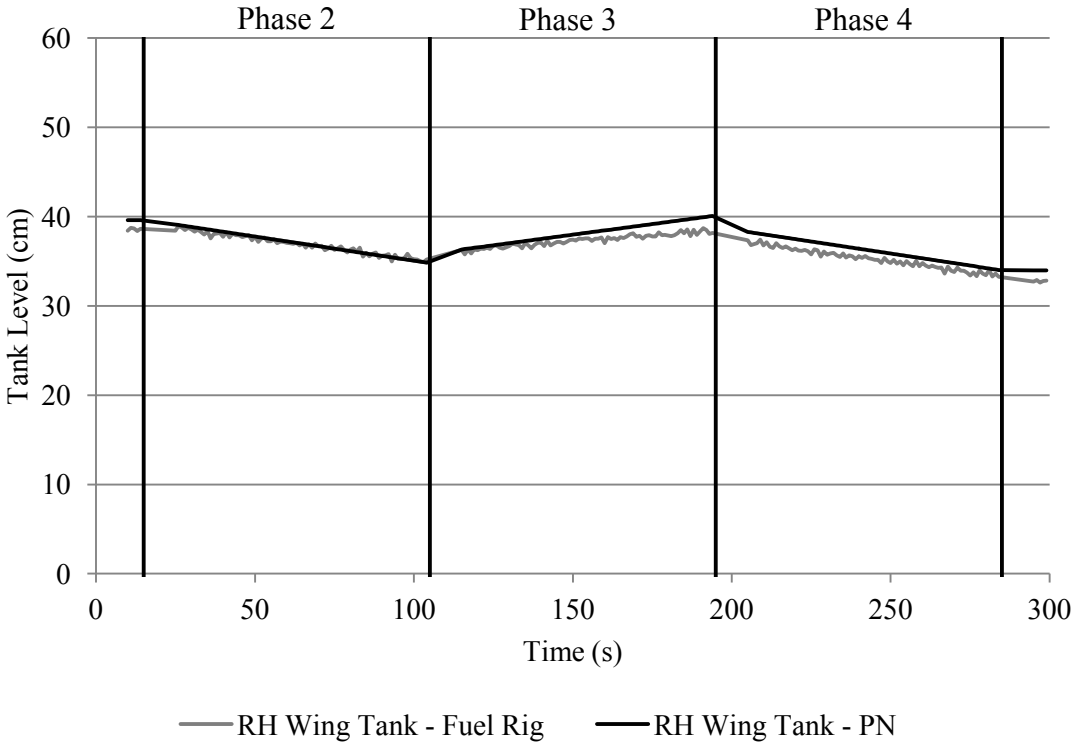

Figure 5.2: 'Clean' fuel rig arrangement - RH wing tank level

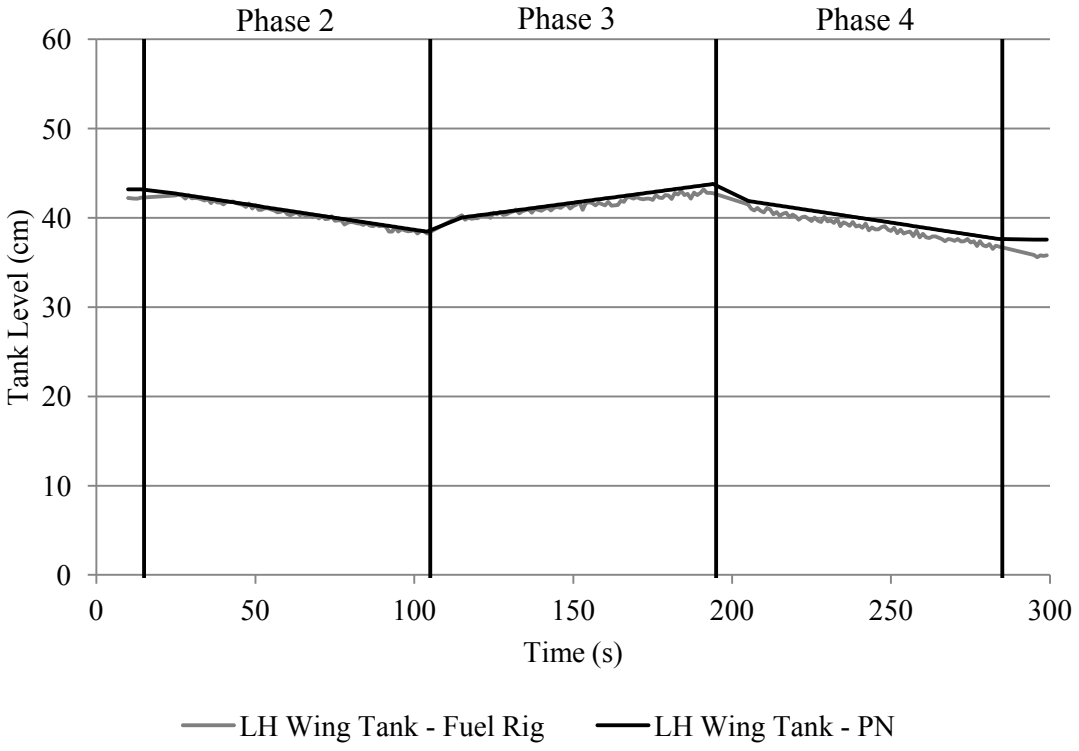

Figure 5.3: 'Clean' fuel rig arrangement - LH wing tank level

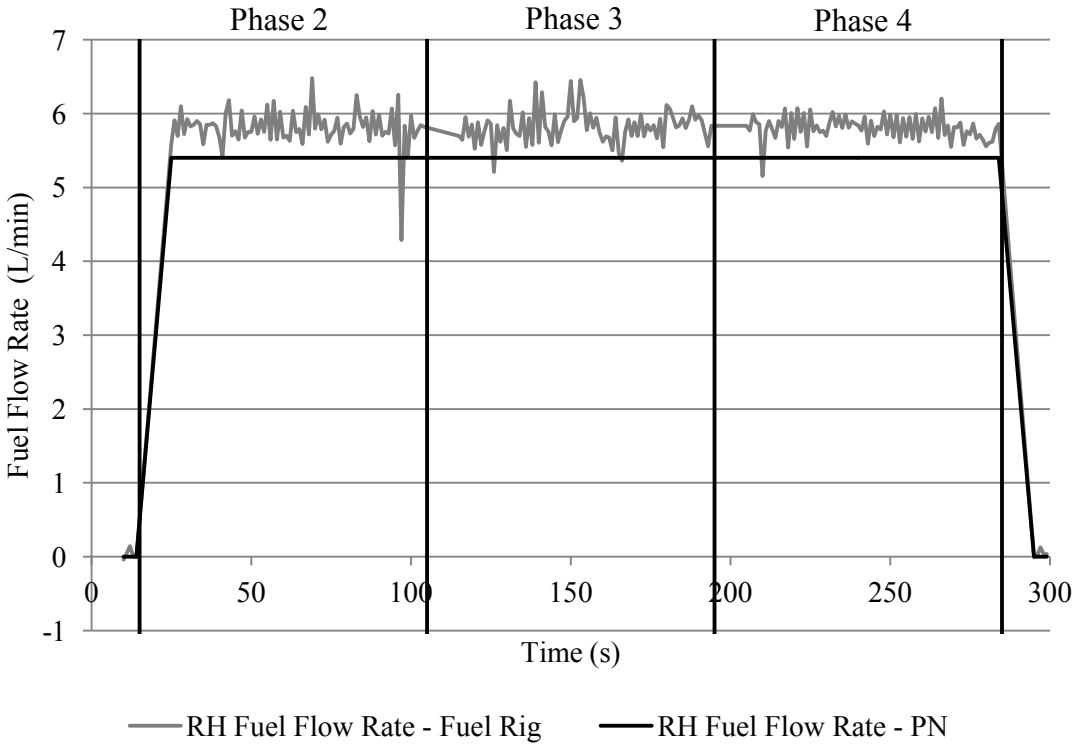

Figure 5.4: 'Clean' fuel rig arrangement - RH fuel flow rate

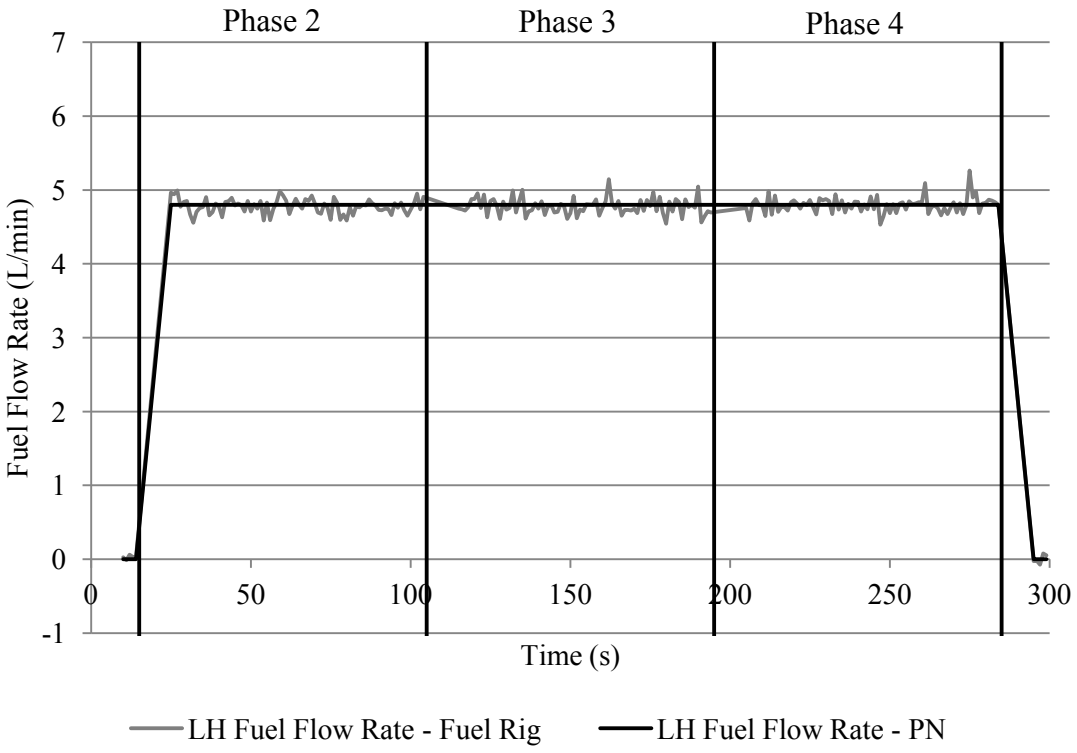

Figure 5.5: 'Clean' fuel rig arrangement - LH fuel flow rate

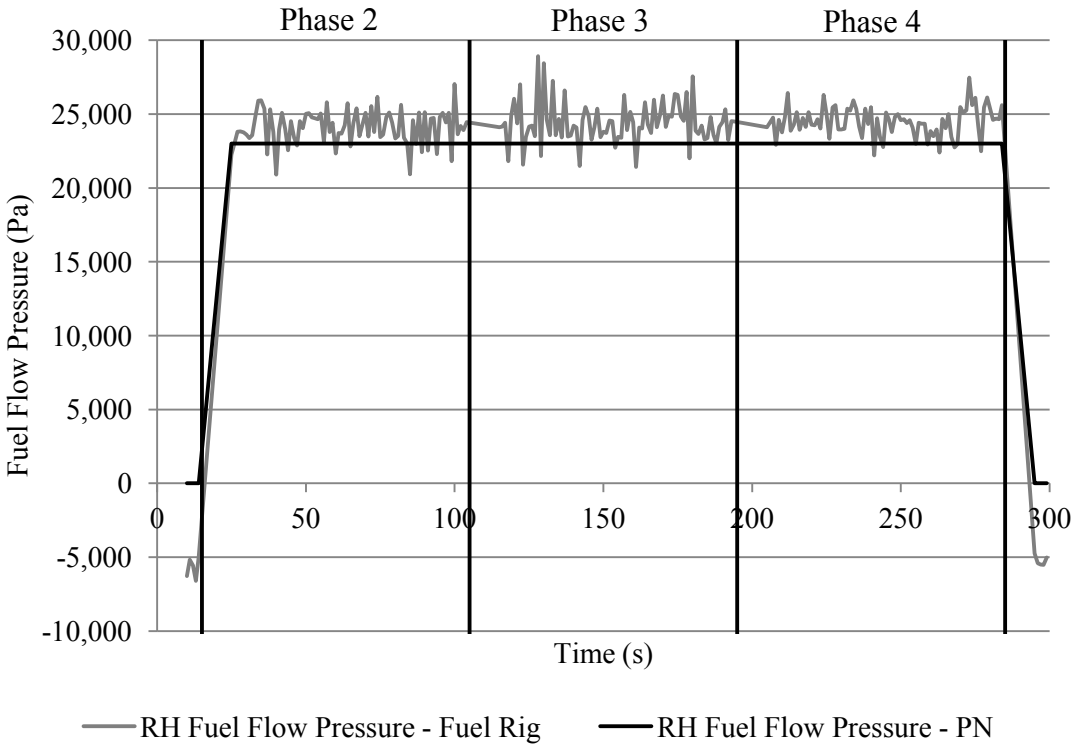

Figure 5.6: 'Clean' fuel rig arrangement - RH fuel flow pressure

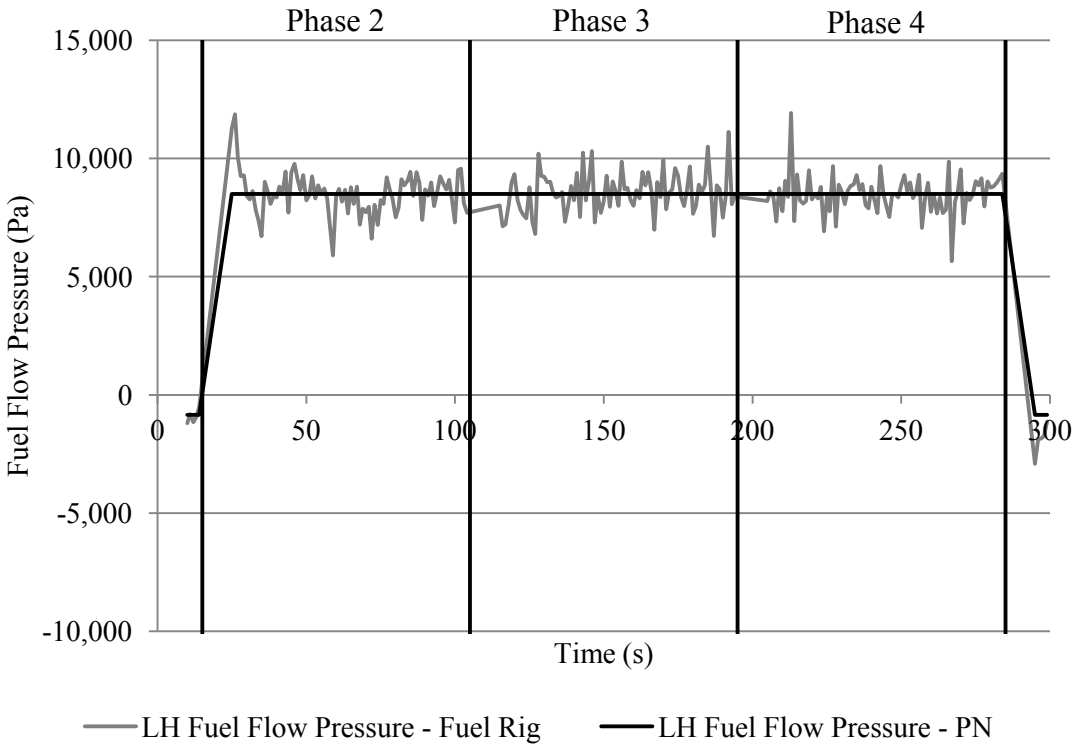

Figure 5.7: 'Clean' fuel rig arrangement - LH fuel flow pressure

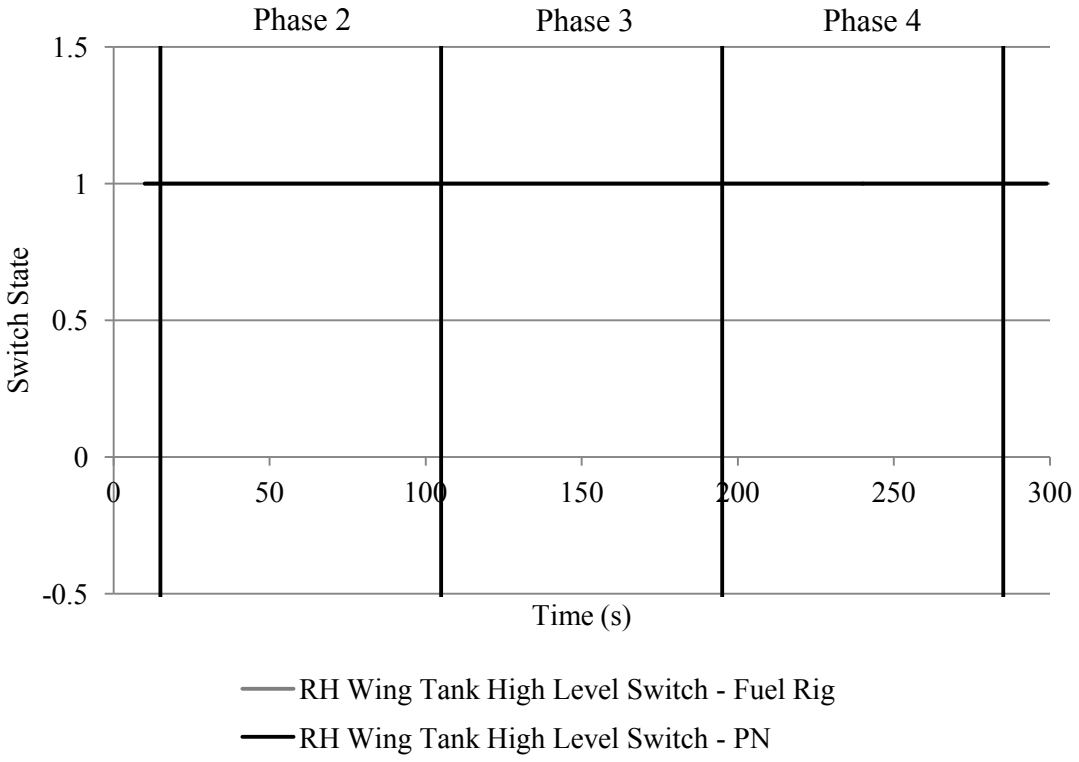

Figure 5.8: 'Clean' fuel rig arrangement - RH wing tank high level switch state

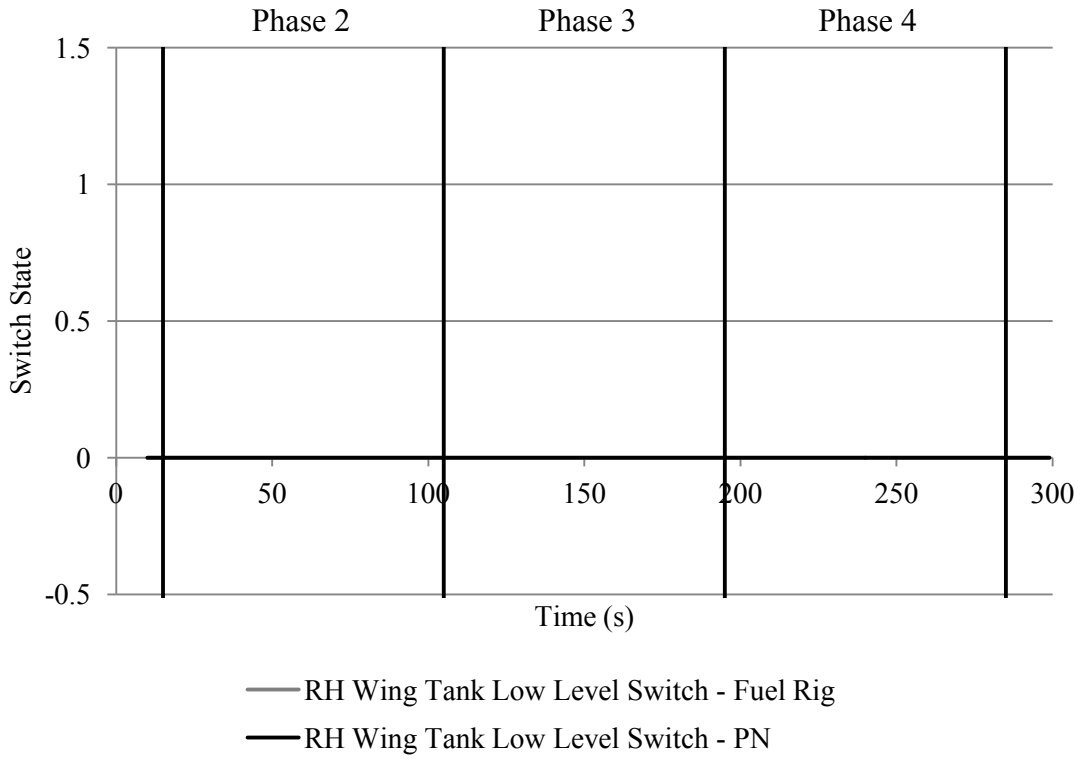

Figure 5.9: 'Clean' fuel rig arrangement - RH wing tank low level switch state

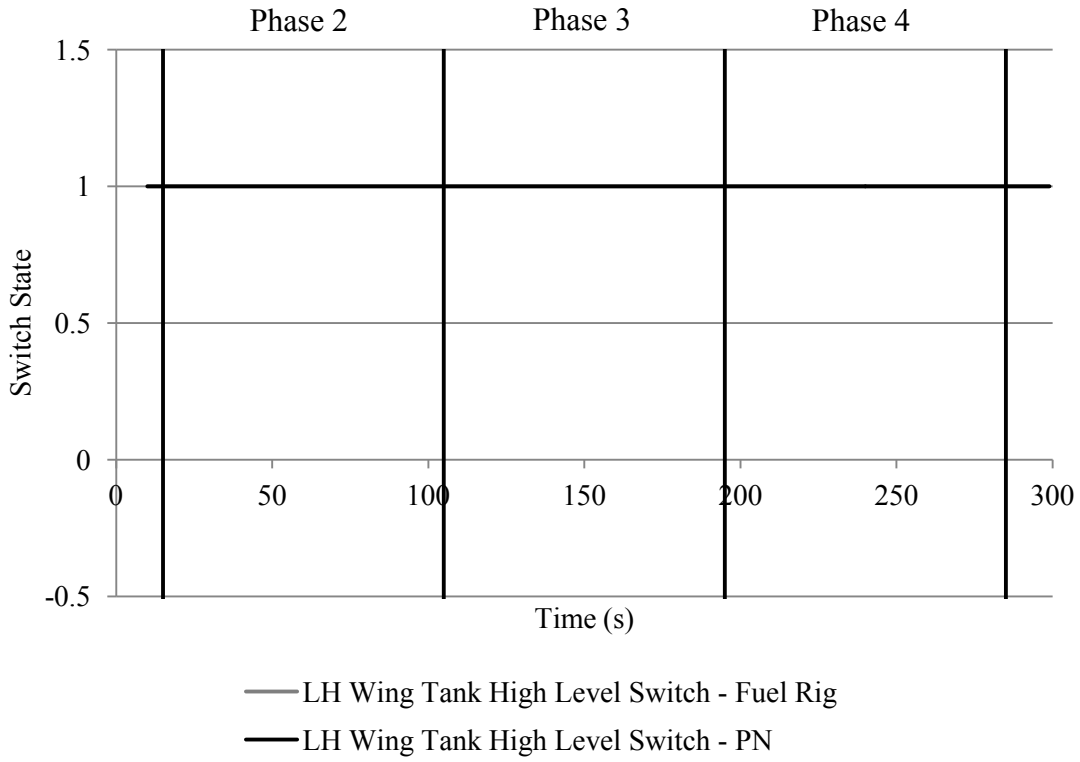

Figure 5.10: 'Clean' fuel rig arrangement - LH wing tank high level switch state

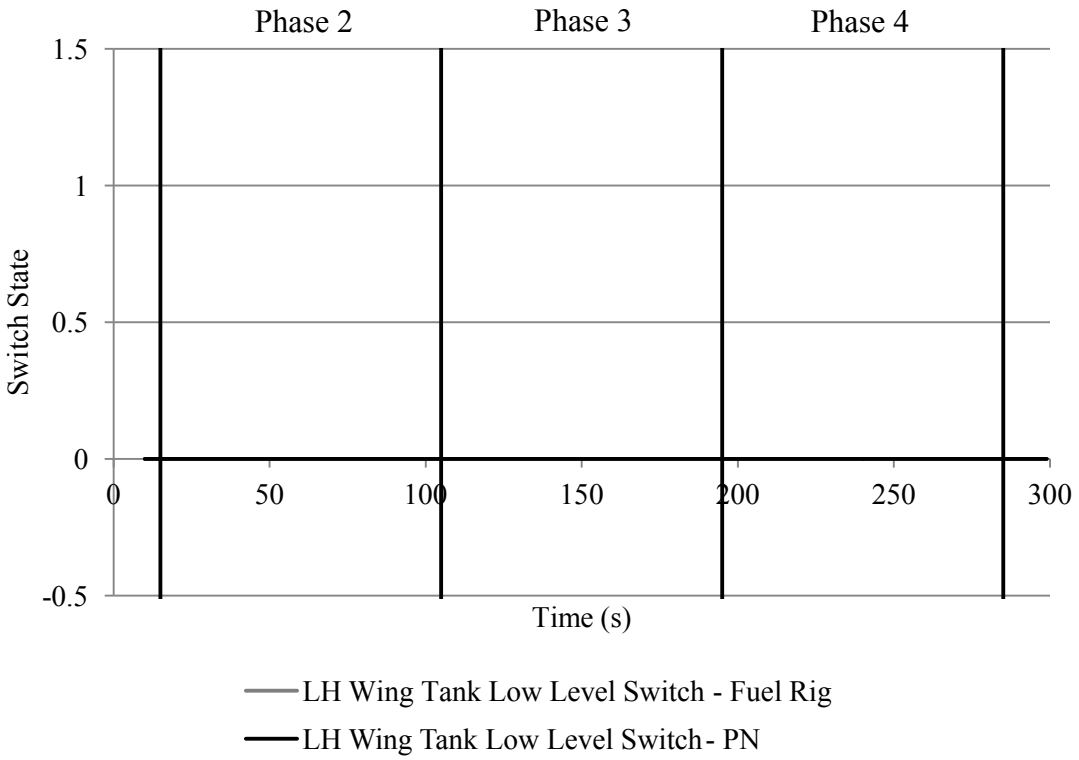

Figure 5.11: 'Clean' fuel rig arrangement - LH wing tank low level switch state

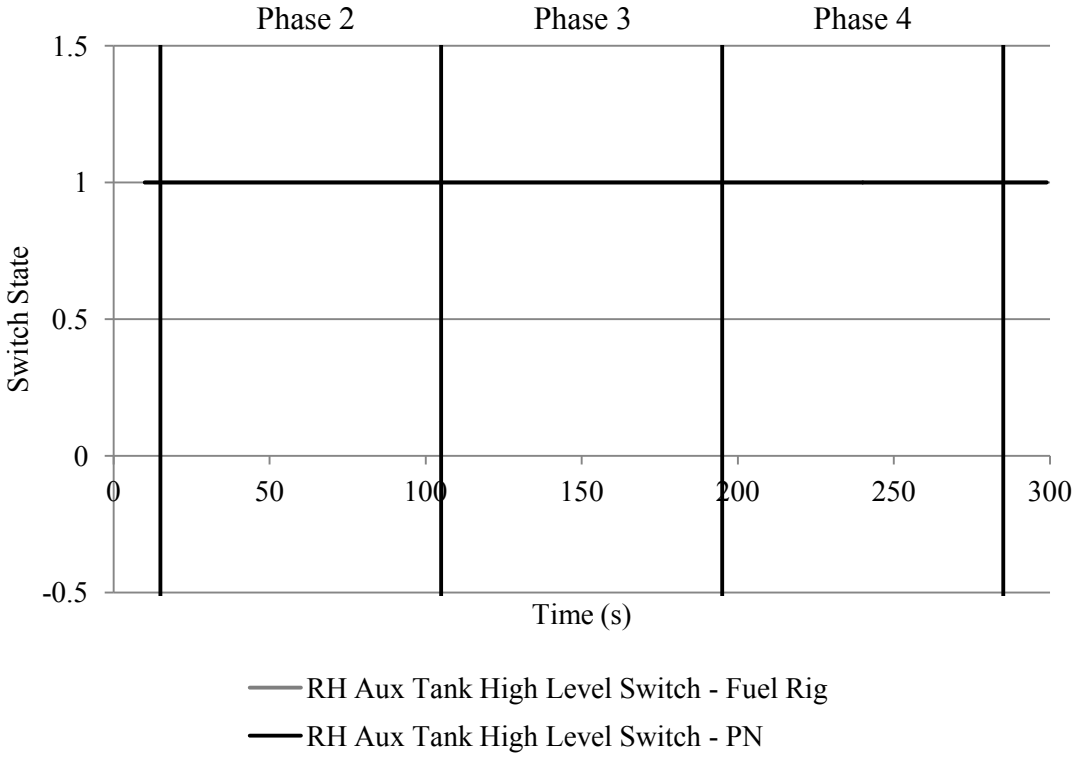

Figure 5.12: 'Clean' fuel rig arrangement - RH auxiliary tank high level switch state

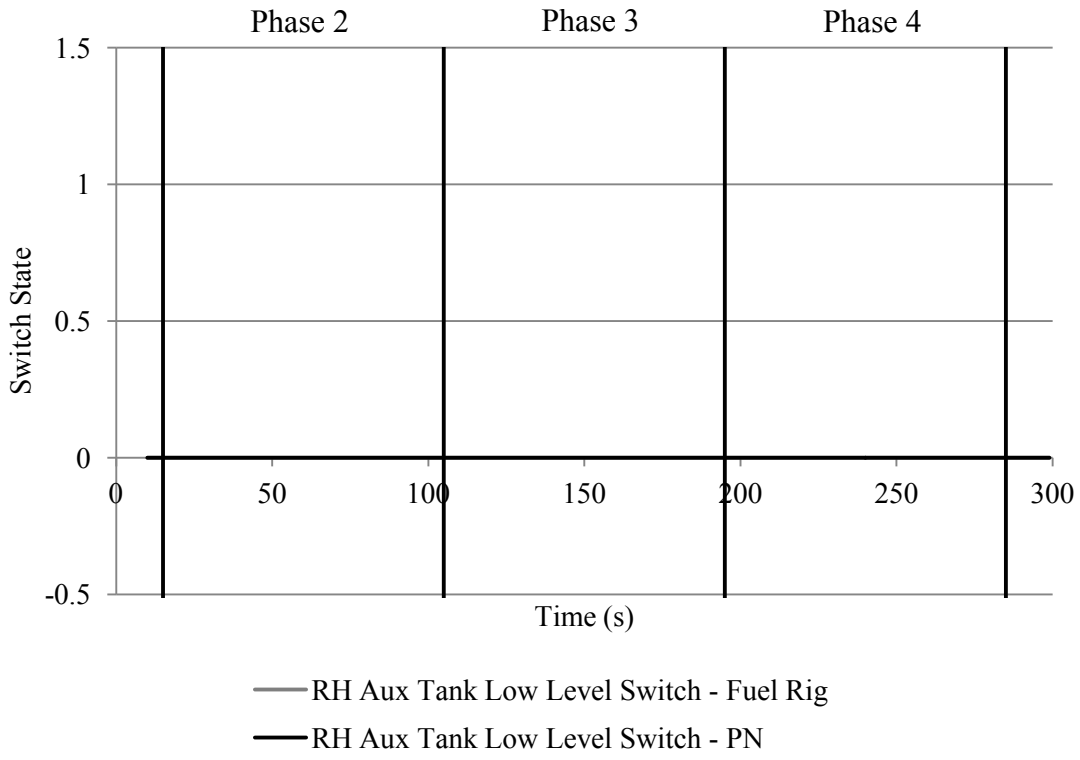

Figure 5.13: 'Clean' fuel rig arrangement - RH auxiliary tank low level switch state

Table 5.2: SD of fuel rig variables in fault free mission

Fuel Rig Variable	Tolerance	SD
LH Auxiliary Tank Level (LHAT)	1.500cm	0.601cm
RH Auxiliary Tank Level (RHAT)	1.500cm	0.443cm
LH Wing Tank Level (LHWT)	1.500cm	0.407cm
RH Wing Tank Level (RHWT)	1.500cm	0.522cm
LH Flow Rate (LHFR)	0.30L/min	0.10L/min
RH Flow Rate (RHFR)	0.30L/min	0.23L/min
LH Fuel Flow Pressure (LHFP)	9,000Pa	858Pa
RH Fuel Flow Pressure (RHFP)	9,000Pa	1,792Pa
LH Wing Tank High Level Switch (LHWTHLS)	0.1	0.0
RH Wing Tank High Level Switch (RHWTHLS)	0.1	0.0
LH Wing Tank Low Level Switch (LHWTLLS)	0.1	0.0

Continued on next page

Fuel Rig Variable	Tolerance	SD
RH Wing Tank Low Level Switch (RHWTLLS)	0.1	0.0
RH Auxiliary Tank High Level Switch (RHATHLS)	0.1	0.0
RH Auxiliary Tank Low Level Switch (RHATLLS)	0.1	0.0

5.3.1 Fuel Rig Sensor Output Anomalies

A number of inconsistent fuel rig behaviours can be identified from the above figures. These behaviours can also be seen when faults are injected into the fuel rig and therefore the cause of these inconsistencies will be evaluated now.

In Figures 5.1, 5.2 and 5.3, the initial tank levels of the LH and RH wing tanks and auxiliary tanks are not the same. A specific initial tank level was not prescribed prior to the period of testing undertaken. In retrospect while it would be preferable to use a consistent initial tank level in terms of comparing behaviour and performance, the results do not suggest this has had a significant effect. Furthermore the refilling process in use on the system is not conducive with achieving high levels of accuracy. It is likely therefore, that achieving a consistent initial tank level would have required a relatively high amount of resources, which would not have provided a significant benefit.

All of the curves on the tank level figures show a small increase in level at the start of phase 3, when the engine and auxiliary pumps are both active, and a decrease at the end of phase 3. This behaviour is representative of the peristaltic pump effects described in Section 4.10.1.

It can be seen in Figures 5.4 and 5.5 that the flow rates recorded from the LHS and RHS of the fuel rig are markedly different. There are a number of potential causes for this variation. The peristaltic pumps installed on the system contain rotational motors that operate in such a way as to create distorted and irregular flow. This could cause variations in recorded flow rate values. There could also be measurement variances at the flow meters themselves. However, these faults may not cause the constant variation seen between the two sets of flow rate values. More likely is an issue with the inconsistent use of components such as pumps on the system. As one pump will likely have been used more than another, the operating performance of the two will be different. Another

possible cause can be identified from the conversion of the electrical sensor outputs to a flow rate value. These conversion factors were determined when the system was initially constructed. Over time the replacing of components, general use of the system and any minor defects could lead to these factors being inaccurate. Considering all of the possible sources listed above, it is likely that a combination of these factors is having an effect on the recorded results. In spite of this variation, Figures 5.4 and 5.5 show that the PN has accurately modelled the behaviour of the flow rate variable on both sides of the system.

Figures 5.6 and 5.7 show that the flow pressure variables exhibits a similar variation in behaviour as is seen with the flow rate variable. Again the output from the LHS of the system is lower than that recorded from the RHS of the system. While no further causes to this variation beyond those discussed previously have been identified, the consistent pattern suggests the inconsistent use of components and inaccurate conversion factors are having the greatest influence on the recorded values. The figures also show that the PN model has accurately represented the flow pressure outputs from both sensors.

5.4 First Order Failure Modes

5.4.1 Overview

The results from using the fault verification technique to evaluate first order fuel rig faults will now be considered. Table 5.3 gives an overview of these results, all of which are evaluated in more detail below. The failure modes considered by the work are listed in the first column of Table 5.3, the fault codes are defined in Table 3.1. Each section of the table shows which system variables have been affected by certain failure modes. The system variables, as listed in column four, are all affected in some manner when any individual failure mode listed in the same section occurs. The second column shows whether the fault verification technique has been able to confirm the presence of the failure mode from the SD values when the arising is injected into the fuel rig and included in the PN model. The third column shows whether the fault verification technique has been able to identify the failure mode as false, from the SD values, when the arising is not present in the system but is included in the PN model.

Considering the first section of Table 5.3, it can be seen that if the RH engine pump fails off or the RH engine isolation valve, RH TPLV isolation valve or RH wing tank isolation

valve is blocked or failed closed it will cause the behaviour of the RH wing tank level, RH flow rate and RH flow pressure variables to change from that seen when no faults were present in the fuel rig. Each of the four failure mode occurring on their own will affect the behaviour of all three variables listed. The outputs recorded from the remaining system variables exhibit similar behaviour to that shown in the respective figures in Section 5.3. In the case of every failure mode listed in the first section of column one, the fault verification technique has been able to confirm the presence of an arising as true (column two) and false (column three) where appropriate. The third section in Table 5.3 shows that this will not always be the case as the fault verification technique has failed to identify the RH level sensor failed stuck arising as false.

Table 5.3: First order failure modes results overview

Failure Mode(s)	Genuine Arising All SD Within Limit	False Arising >0 SD Outwith Limit	Variable(s) Affected
Eng IV B/FC,	✓	✓	RHWT
TPL-V IV B/FC,	✓	✓	RHFR
WT IV B/FC,	✓	✓	RHFPR
Eng Pump FO	✓	✓	
AT IV B/FC,	✓	✓	RHAT
Aux Pump FO	✓	✓	RHWT
LS FH,	✓	✓	RHWT
LS FL,	✓	✓	
LS FS	✓		
FS FH,	✓	✓	RHFR
FS FO,	✓	✓	
FS FS	✓	✓	
FP FH,	✓	✓	RHFP
FP FO,	✓	✓	
FP FS	✓		
WT HLSw FOn,	✓	✓	RHWTHLSW
WT HLSw FOff,	✓		
WT HLSw FS	✓		

Continued on next page

Failure Mode(s)	Genuine Arising All SD Within Limit	False Arising >0 SD Outwith Limit	Variable(s) Affected
WT LLSw FOn,	✓	✓	RHWTLLSW
WT LLSw FOff,	✓		
WT LLSw FS	✓		
Eng Pump D	✓	✓	
Aux Pump D	✓		RHAT
			RHWT
WT Lk	n/a	n/a	RHWT
AT Lk	n/a	n/a	RHAT

Both leak failure modes are not evaluated using the same process as the other failure modes and therefore cannot be assessed in the second and third columns. The leak failure modes are considered in Section 5.4.9.

5.4.2 Isolation Valve Failure Modes

A detailed analysis of the failure mode 'Right-Hand Engine Isolation Valve Blocked/Failed Closed' will now be considered. As the vast majority of the fuel rig failure modes under consideration have been analysed using the same process the remaining failure modes will be analysed in an abbreviated form.

5.4.2.1 RH Engine Isolation Valve Blocked/Failed Closed

In the event of the RH engine IV becoming blocked or failing closed the flow path from the RH wing tank to the RH engine will be lost. It would be expected that the effect of this fault would be seen in the RH wing tank level, the RH flow rate and RH fuel flow pressure variables. Figure 5.14 shows the tank levels recorded from the LH and RH wing tank level sensors on the fuel rig over the course of the phased mission when this fault occurs. Also plotted are the tank levels predicted by the PN model. Evaluating both wing tanks should allow the effect of the fault on the RH wing tank level to be directly compared to the LH wing tank level, which should be behaving normally. The clean RH

wing tank level behaviour, and that of all the variables, can be seen in Section 5.3.

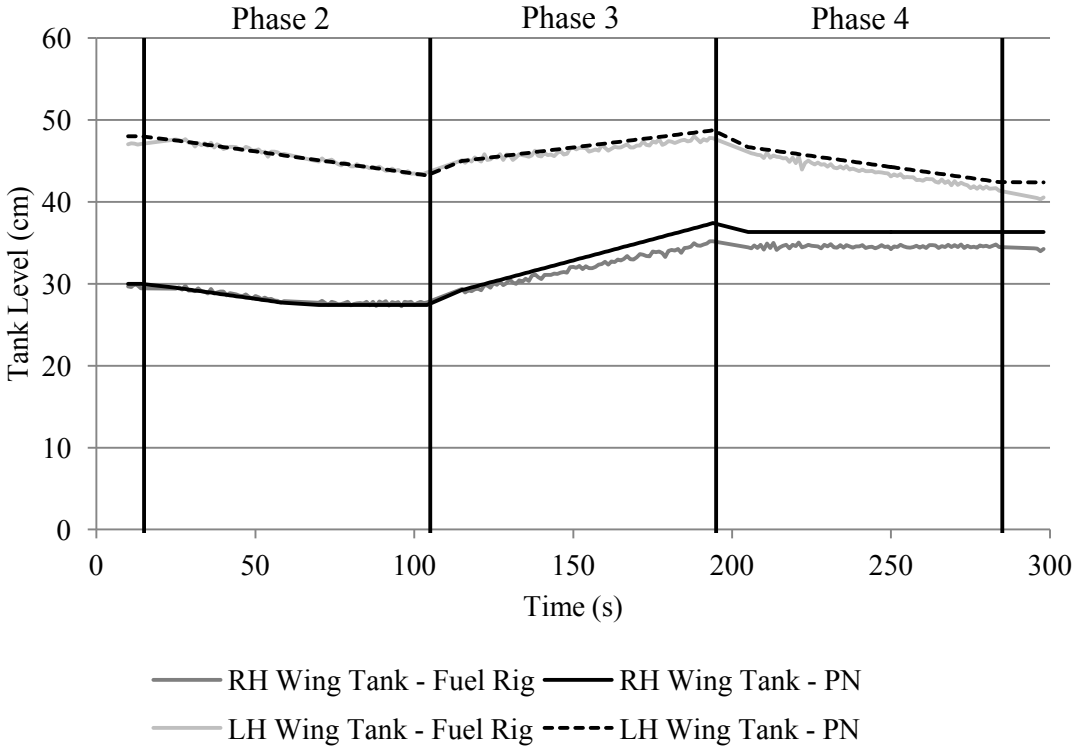

Figure 5.14: RH engine IV blocked - Wing tank levels

The effect of the RH engine IV fault can be seen in the tank level variables plotted in Figure 5.14. From 60 seconds onwards the RH wing tank level never decreases as the fault prevents any flow from leaving the tank. By comparison the LH wing tank level continues to decrease in the remainder of phase 2 and all of 4. In phase 3 both the LH and RH wing tank levels increase due to the input flow from the auxiliary tanks. However, the RH wing tank level increases at a greater rate than the LH wing tank level, as there is only flow into the RH wing tank but flow both into and out of the LH wing tank. The fault has therefore had a visible effect on the behaviour of the RH wing tank level variable in phases 2 – 4.

Applying the SD technique to both the LH and RH wing tank level data sets produces SD values of 0.496cm and 0.901cm respectively. Both of these values are within the tolerance limit of 1.500cm for tank level variables. The PN model has therefore accurately represented the behaviour of a RH engine IV blockage in the wing tank level variables.

Although both of the wing tank SD values were within the tolerance limits, from Figure 5.14 it can be seen that in phase 3 the recorded RH tank level values increase at a lower

rate than the PN values. Over a greater period of time this could cause the RH wing tank SD value to exceed the tolerance limit. The lack of conformity could be a result of several factors. System noise is likely to be one of these factors. Issues created by inconsistent equipment use as discussed in Section 5.3.1 could also be a cause. Another factor could be due to the pipe configuration from the auxiliary tanks to the wing tanks, which is not the same on the two sides of the system. The longer series of piping feeding the RH wing tank would require a slightly higher auxiliary pump rating to achieve the same flow rate, yet both auxiliary pumps were run at the same setting. Individually or collectively, these factors could cause the variation between the recorded and predicted wing tank level values.

Figures 5.15 plots the flow rate measured by the sensor on the LHS of the system between the TPLV and the engine. Figure 5.16 plots the flow rate measured from the RHS of the system.

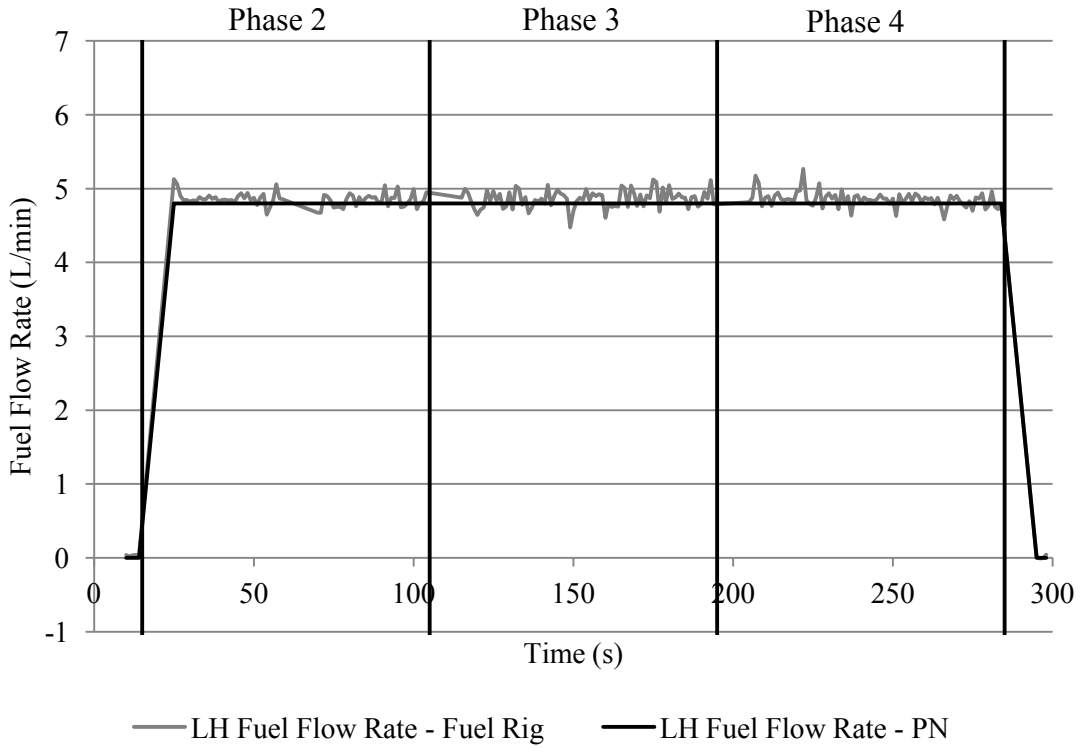

Figure 5.15: RH engine IV blocked - LH flow rate

There is a good level of similarity between the predicted and recorded flow rate values shown in Figures 5.15 and 5.16. The effect of the IV fault can be seen in the RH flow rate variable, which rapidly falls from approximately 6L/min to 0L/min after 60 seconds

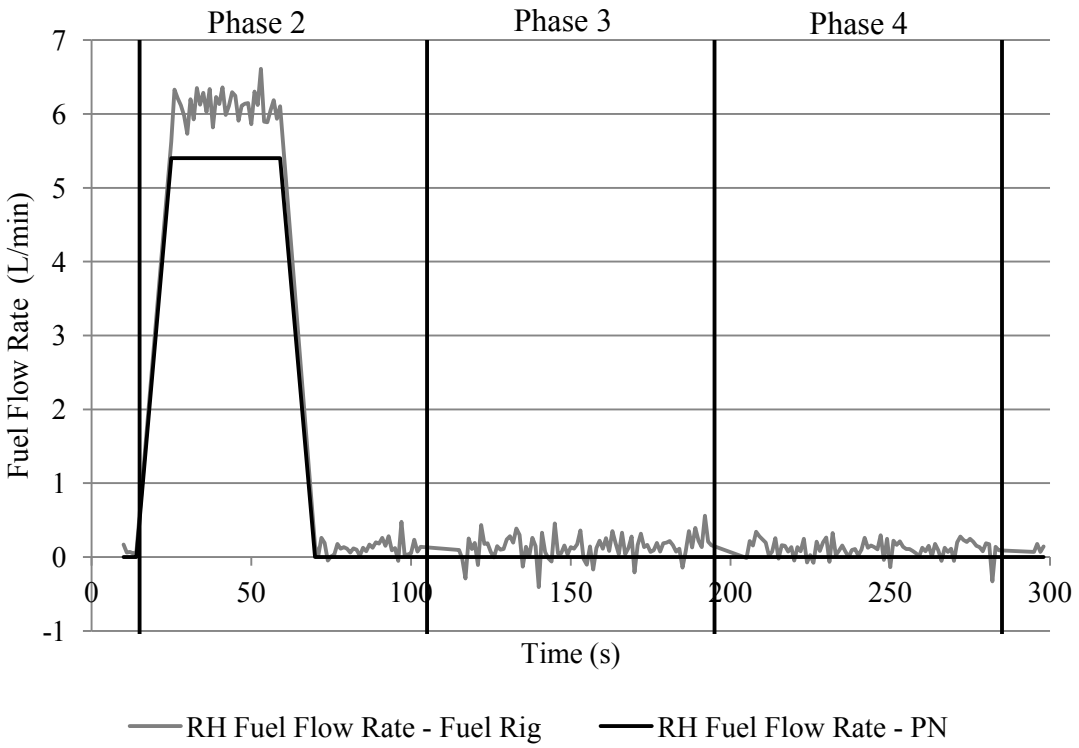

Figure 5.16: RH engine IV blocked - RH flow rate

and remains there for the remainder of the mission. The LH flow rate remains relatively constant at 4.8L/min in phases 2 – 4 and at 0L/min in phases 1 and 5, reflecting the engine pump demand in those phases. Applying the SD technique to the two data sets produces a SD value of 0.11L/min for the LH flow rate variable and 0.25L/min for the RH flow rate. Both of these values are within the flow rate tolerance of 0.30L/min. The cause of the variation between the flow rates recorded from the two sides of the fuel rig is discussed in Section 5.3.1.

The fuel flow pressures recorded from the fuel rig and predicted by the PN model over the course of the mission, where the RH engine IV is blocked/failed closed, are displayed below. The results from the LH side of the system are shown in Figure 5.17 and from the RH side of the system in Figure 5.18

Figures 5.17 and 5.18 show that while on the LHS the flow pressure remains relatively constant from phases 2 through 4, the RH flow pressure variable increases from approximately 25,000Pa to more than 300,000Pa at the time the engine IV fault occurs. This represents an increase of more than 1,200% in the recorded flow pressure. The cause of this significant increase is the IV blockage. As the blockage occurs downstream of the RH

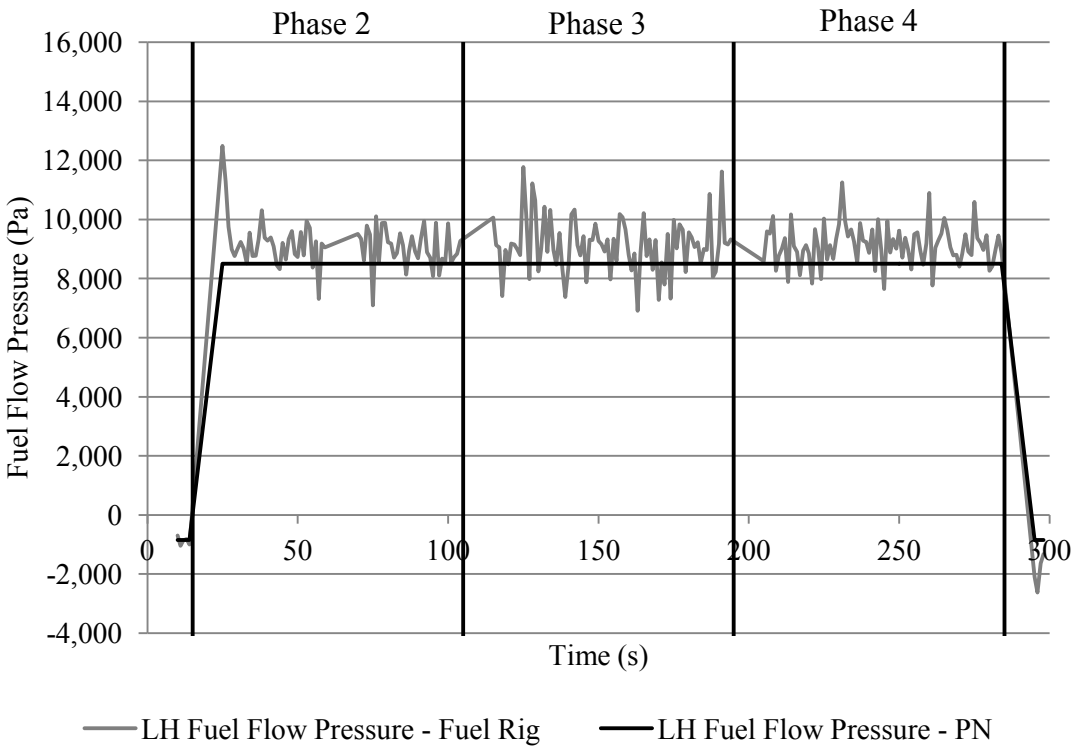

Figure 5.17: RH engine IV blocked - LH flow pressure

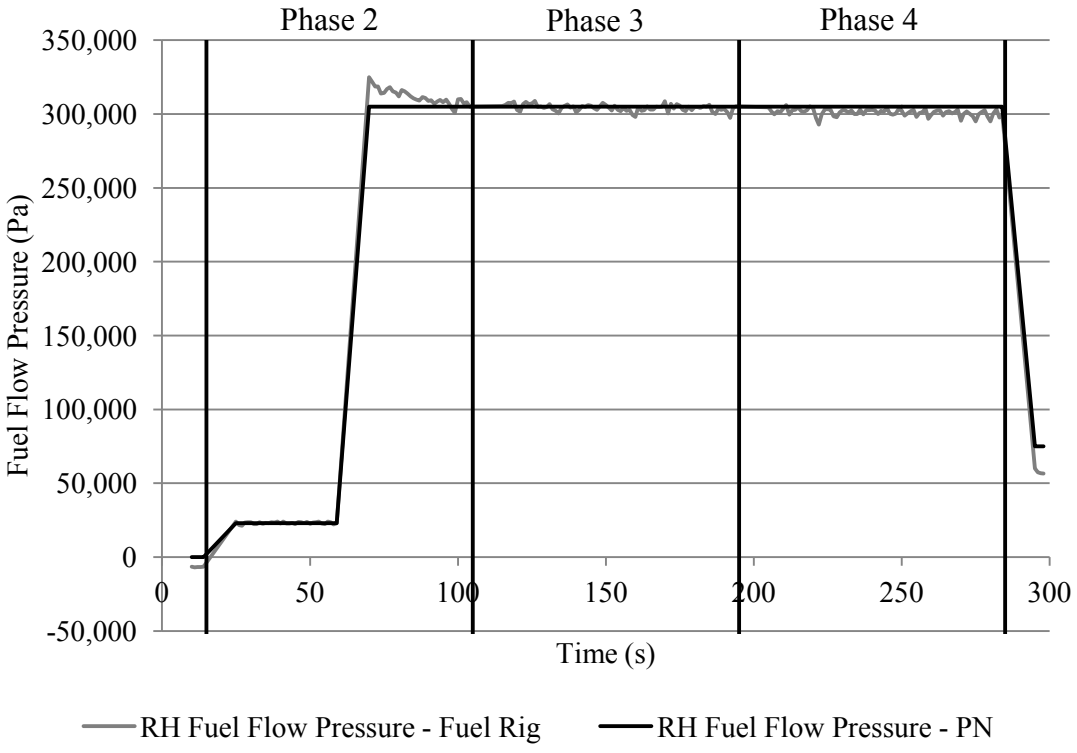

Figure 5.18: RH engine IV blocked - RH flow pressure

low pressure engine pump, fuel and air is trapped at the location of the pressure sensor. When the pump is on these trapped materials quickly become highly pressurised which causes the large values plotted in Figure 5.18 to be produced. By comparison a blockage upstream of the engine pump would not see any air or fuel trapped at the pressure sensor and so only a blockage downstream of the engine pump can cause the effect seen in the fuel flow pressure outputs of Figure 5.18. Although unique flow rate behaviour has been recorded from each side of the fuel rig, the PN model has produced an accurate prediction of both variable outputs in the presence of the fault.

Analysing the fuel flow pressure data sets using the SD technique produces results of 821Pa and 4,950Pa for the LH and RH sides of the fuel rig respectively. Both of these results are within the fuel flow pressure tolerance limit of 9,000Pa.

Figure 5.19 displays the LH and RH auxiliary tank level values as recorded from the fuel rig and predicted by the PN model.

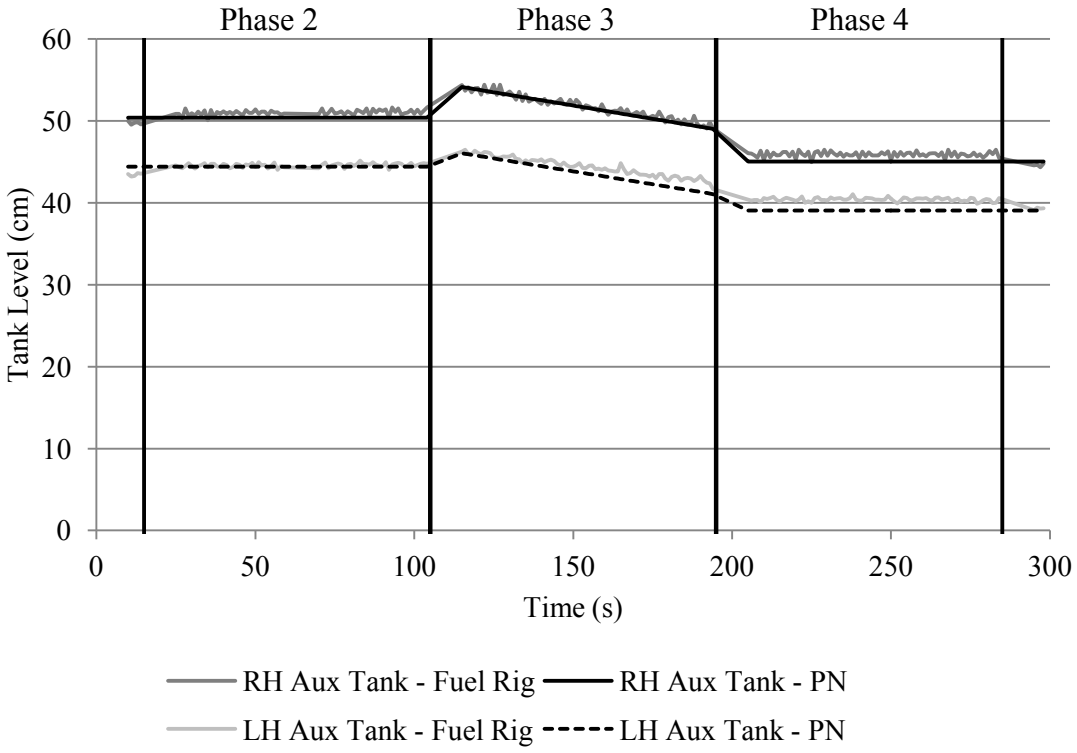

Figure 5.19: RH engine IV blocked - Auxiliary tank levels

Figure 5.19 shows that, accounting for the tank level adjustments due to the peristaltic pump effects as described in Section 4.10.1, the recorded and predicted auxiliary tank levels are very similar on both sides of the fuel rig system. The only true change in the auxiliary

tank levels occurs in phase 3, when a demand is applied to the auxiliary pumps. In the remaining phases the tank level does not change. The RH engine IV fault therefore has no visible effect on the performance of the auxiliary tank level variables.

The SD of the LH auxiliary tank level data sets is 0.612cm. The SD of the RH auxiliary tank level data sets is 0.499cm. Both of these values are within the tank level tolerance limit of 1.500cm.

The SD values for a range of system variables are listed in Table 5.4. The RH wing tank level, RH flow rate and RH fuel flow pressure variables all displayed behaviour that was different from the behaviour expected when no fault was present. The PN model that included the fault however, accurately predicted this recorded behaviour. As a result when the 'RH Engine IV Blocked/Failed Closed' arising is present in the fuel rig and modelled in the PN model all of the SD values are within the tolerances for the respective variables, as shown in column three. This indicates that the expected behaviour of the system with the fault present is similar to that exhibited by the fuel rig. Therefore it is possible to confirm that the fault verification technique can verify the occurrence of the 'RH Engine IV Blocked/Failed Closed' arising. The high and low level switch results have not been listed as their behaviour is unaffected by presence of the fault during the course of the mission. Their behaviour is consistent with that shown in Section 5.3.

Table 5.4: SD of fuel rig variables - RH engine IV fault

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.612cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.499cm	0.480cm
LH Wing Tank Level	1.500cm	0.496cm	0.437cm
RH Wing Tank Level	1.500cm	0.901cm	**4.246cm**
LH Flow Rate	0.30L/min	0.11L/min	0.10L/min
RH Flow Rate	0.30L/min	0.25L/min	**2.13L/min**
LH Fuel Flow Pressure	9,000Pa	821Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	4,951Pa	**107,294Pa**

To demonstrate how the fault verification technique can also be used to identify false

arisings consider column four of Table 5.4. This column shows the SD values of the fuel rig variables when the engine IV fault is included in the PN model simulation but is not present in the fuel rig itself. It can be seen that the SD values of the RH wing tank level, RH flow rate and RH fuel flow pressure variables all exceed the respective tolerance values. These results show that should the 'RH Engine IV Blocked/Failed Closed' arising be falsely generated when the system is operating normally, the fault verification technique would correctly filter the arising. Figures 5.20, 5.21 and 5.22 show how the recorded and predicted curves of the RH wing tank level, RH fuel flow rate and RH flow pressure contrast when the RH engine IV fault is only included in the PN model. It should be noted that the variables which exceed the SD limit when the arising is false are the same as those whose behaviour changes when the arising is genuine.

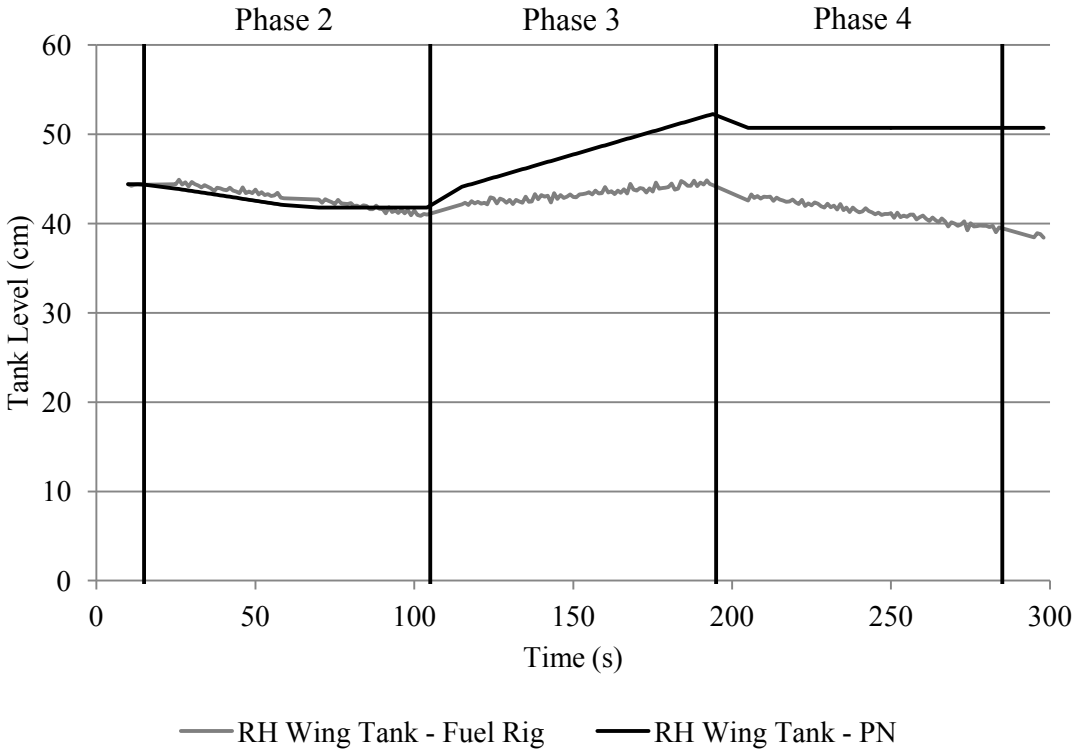

Figure 5.20: RH engine IV blocked falsely diagnosed - RH wing tank levels

Further sections will be considered as follows; graphs will only be presented where the variable behaviour in the presence of a fault differs from that where no faults are present. Graphs of the fuel rig variables where no faults are present in the system are shown in Section 5.3. Table 5.3 presented an overview of these results.

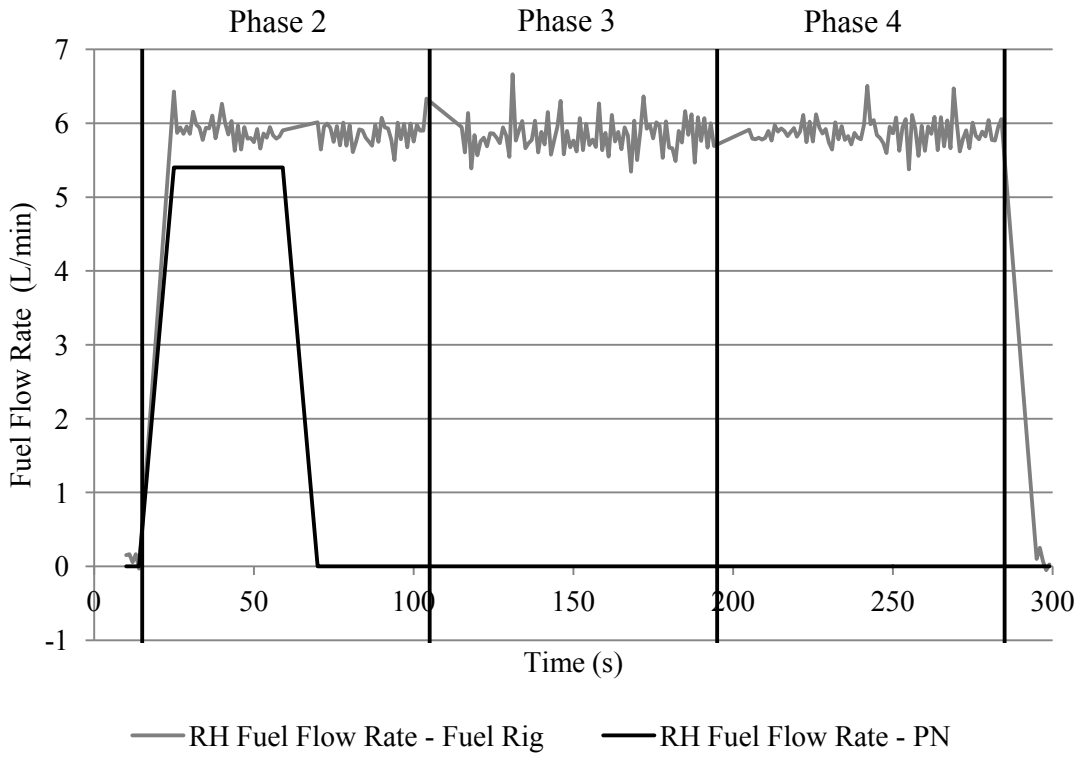

Figure 5.21: RH engine IV blocked falsely diagnosed - RH fuel flow rate

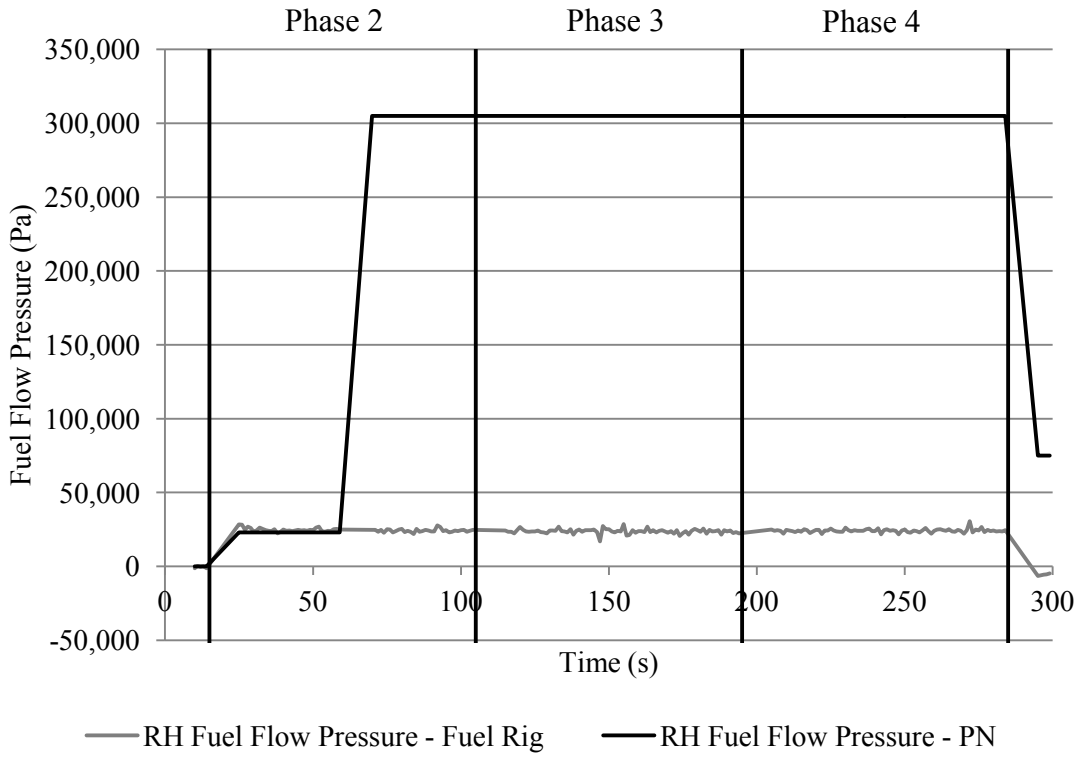

Figure 5.22: RH engine IV blocked falsely diagnosed - RH fuel flow pressure

5.4.2.2 RH Triple Port L-Valve Isolation Valve Blocked/Failed Closed

If a TPLV IV becomes blocked or fails closed the flow path from the wing tank to the IV and all flow paths downstream of the IV will be terminated. The effect of the blockage/failure on the RH TPLV IV should be visible in the RH wing tank level, RH flow rate and RH fuel flow pressure variables. Figure 5.23 shows the RH wing tank levels recorded from and predicted for the fuel rig system during the five phase mission under consideration.

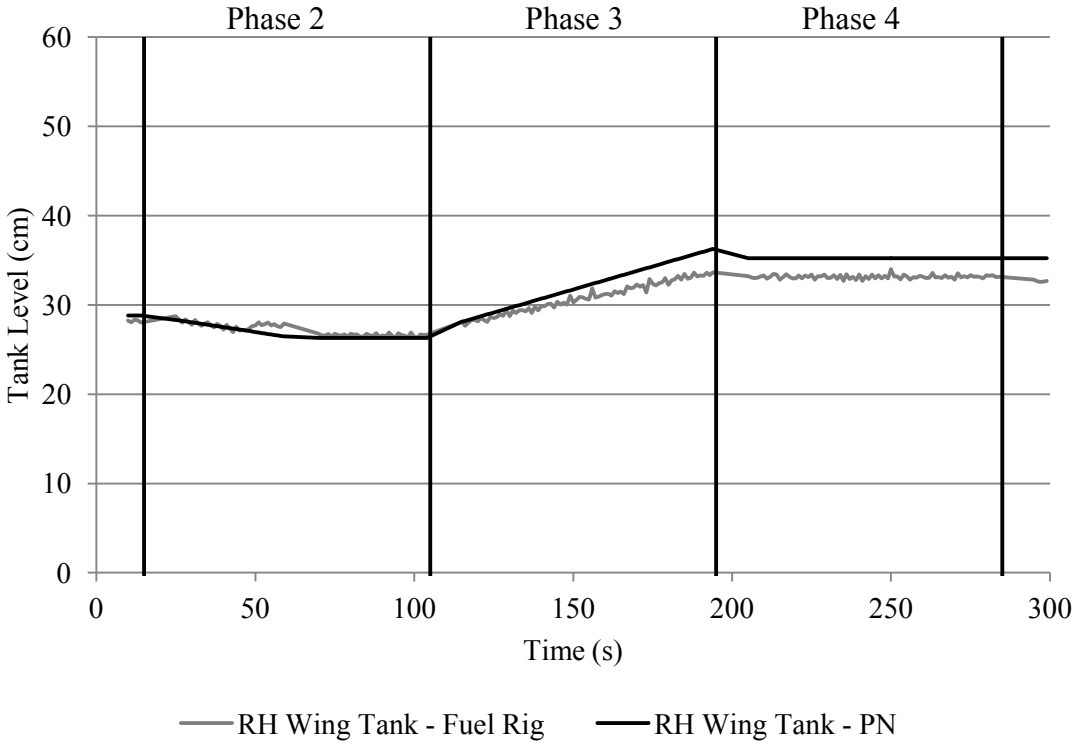

Figure 5.23: RH triple port L-valve IV blocked - RH wing tank levels

Figure 5.23 shows that from the time of the fault occurring in phase 2 the only change in the RH wing tank level is an increase in phase 3 due to the flow in from the auxiliary tank. The location of the fault, downstream of the wing tank, prevents any further fuel from leaving the tank. Comparing this figure to that from Section 5.3 it can be seen that the effect of the fault can be seen in the RH wing tank variable. This behaviour is also the same as that seen when there was a blockage in the RH engine IV, as described in Section 5.4.2.1. As both of these IVs are on the same section of the system and both are downstream of the wing tank, this result is consistent with what would be expected. The potential causes of the variation between the fuel rig and PN tank level gradients in phase 3 were identified and discussed in Section 5.4.2.1. The LH wing tank level variables is

unaffected by the fault.

Figure 5.24 shows the flow rate recorded by the RH flow rate sensor and as predicted by the PN model throughout the mission.

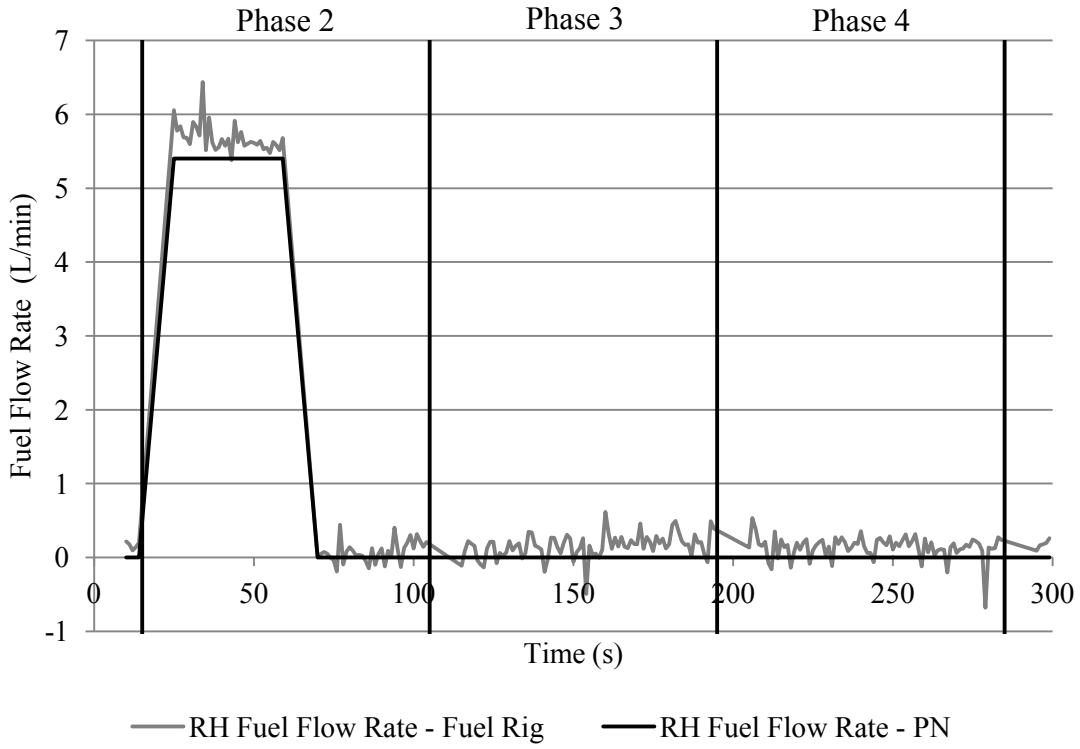

Figure 5.24: RH triple port L-valve IV blocked - RH flow rate

The effect of the fault on the fuel rig system can be clearly seen in the RH flow rate variable. Once the fault is injected into the system, there is no flow downstream of the wing tank and therefore no flow past the flow rate sensor. As a result, the recorded flow rate falls to approximately zero. The PN system model has predicted this performance well. The LH flow rate variable is not affected by the TPLV fault and its behaviour is in line with that shown in the relevant figure of Section 5.3.

Figure 5.24 also provides a good example of the level of noise in the system and the fuel flow rate variable. After the fault has been injected the flow rate recorded from the fuel rig should be 0L/min. However it can be seen that the flow rate fluctuates between approximately ± 0.5L/min. This occurs despite of the fact that there is no flow in the system at this time. Furthermore there shouldn't be any reverse flow in the system and so negative flow rate values are unexpected and can therefore be attributed to noise. The effects of noise can be seen both prior to and after the fault occurring. The fault

verification technique accounts for this level of noise in the tolerances that it applies to variables. It follows therefore that a system with smaller amounts of noise could utilise narrower tolerances.

Figure 5.25 shows the fuel flow pressure values recorded at the RHS of the fuel rig system and predicted by the PN model.

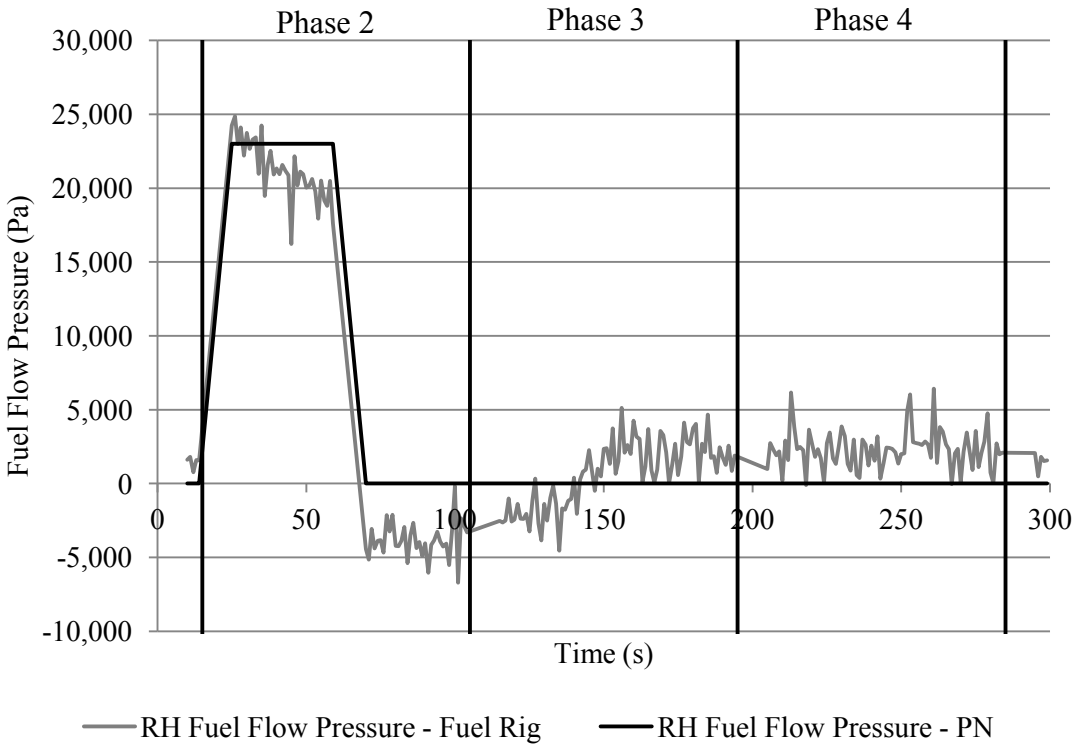

Figure 5.25: RH triple port L-valve IV blocked - RH fuel flow pressure

Figure 5.25 shows that the TPLV IV blockage causes a significant change in the behaviour of the fuel flow pressure values on the fuel rig. The blockage has prevented any fuel flow passing through the low pressure engine pump and, as a result, past the flow pressure sensor. The fuel flow pressure values therefore fall to around zero. The figure also shows that the fuel rig PN model has predicted the behaviour of the fuel flow pressure variable in the presence of the TPLV IV fault well. The presence of noise in the flow pressure variable can be clearly seen in Figure 5.25 and the observations that were applied when considering noise in the flow rate variable can also be applied here. The LH flow pressure sensor outputs have not been affected by the TPLV IV fault and the behaviour of this variable can be seen from the graph in Section 5.3.

Table 5.5 summarises all of the SD values for the fuel rig variables when the TPLV IV

is blocked/failed closed. It can be seen that when correctly diagnosed all of the SD values are within the relevant tolerances, including those variables whose behaviour has changed as a result of the fault. An arising created in this scenario would therefore be correctly verified as a genuine fault. The table also shows the SD results when the fault is included in the PN model but is not present in the system itself - a scenario representative of a false arising. In this case the RH wing tank level, RH flow rate and RH fuel flow pressure variables all exceed their respective tolerances. As a result any TPLV blocked/failed closed arising falsely generated by the health management system would be filtered.

Table 5.5: SD of fuel rig variables - RH TPL-valve IV fault

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.531cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.510cm	0.480cm
LH Wing Tank Level	1.500cm	0.454cm	0.437cm
RH Wing Tank Level	1.500cm	1.104cm	**4.270cm**
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.17L/min	**2.13L/min**
LH Fuel Flow Pressure	9,000Pa	1,170Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	2,748Pa	**9,247Pa**

5.4.2.3 RH Wing Tank Isolation Valve Blocked/Failed Closed

Should a wing tank IV become blocked or fail closed, the flow paths downstream of the respective wing tank IV will be terminated. When the RH engine pump is on, the effect of the failure mode will be seen in the RH wing tank level, RH flow rate and RH fuel flow pressure variables. Figure 5.26 shows the RH wing tank level variable over the course of the mission when the wing tank IV fault occurs.

Figure 5.26 shows that the shape of the fuel rig and PN tank level curves are very similar to those shown in Section 5.4.2.1 and in Section 5.4.2.2. This is due to the fact that the location of all three faults is in a similar section of the system and all of the faults

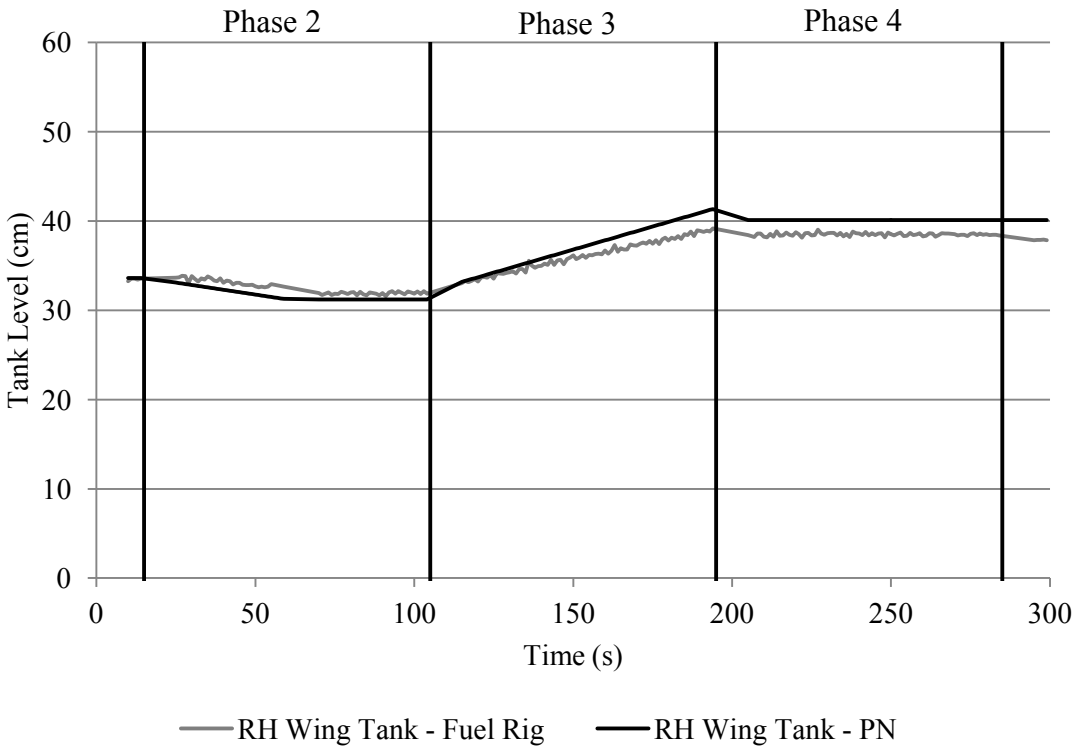

Figure 5.26: RH wing tank IV blocked - RH wing tank level

prevent fuel from leaving the RH wing tank. The observations made in these previous sections regarding the wing tank level therefore also apply here.

Figure 5.27 shows that the RH fuel flow rate falls to approximately zero once the fault has been injected into the fuel rig. Again this behaviour is consistent with that shown where a fault occurred in the engine IV and TPLV IV. As with the observation made of the RH wing tank level variable, this is due to the location of the fault and its similarity to those already considered preventing any fuel flow from the wing tanks to the engines.

Figure 5.28 shows that at the time of the fault being injected into the fuel rig the fuel flow pressure value falls from approximately 23,000Pa to zero. This is a result of the fault preventing any flow through the engine pump and past the flow pressure sensor. The behaviour of the flow pressure variable is consistent with that produced when the previous IV faults were considered.

The behaviour of the LH wing tank level, flow rate and flow pressure variables have been unaffected by the presence of the wing tank IV fault in the system. Their sensor outputs are consistent with those shown in Section 5.3.

Table 5.6 summarises all of the SD results for the variables on the fuel rig when the

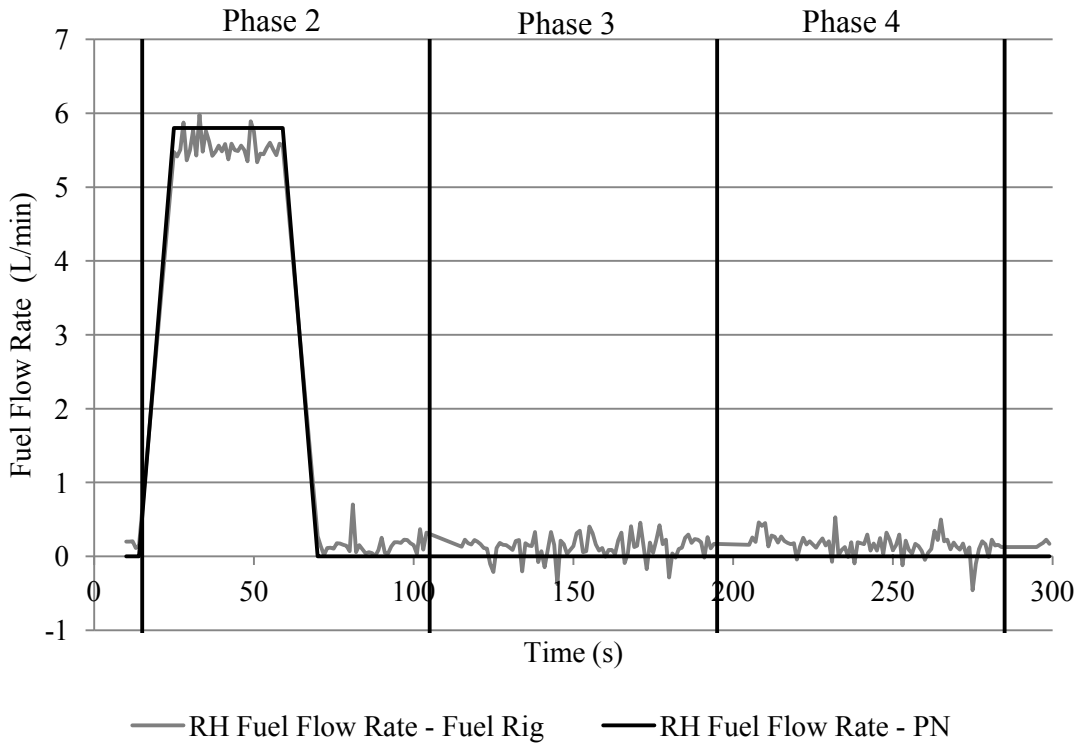

Figure 5.27: RH wing tank IV blocked - RH fuel flow rate

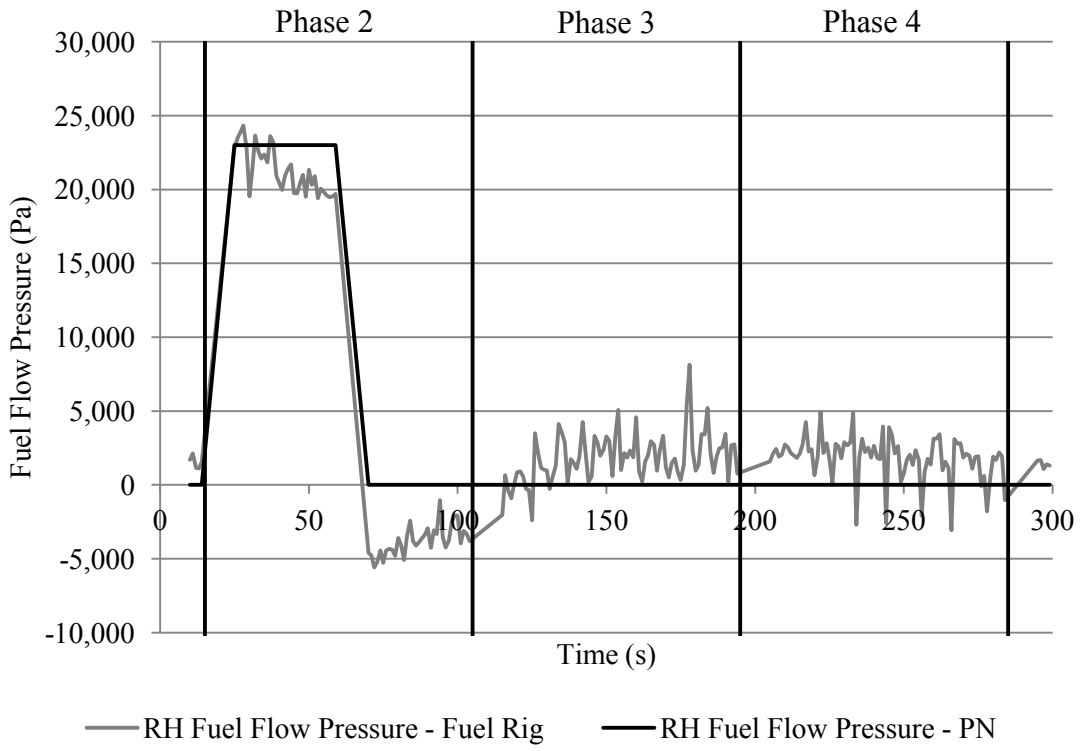

Figure 5.28: RH wing tank IV blocked - RH flow pressure

RH wing tank IV is blocked/failed closed. It can be seen that all of the SD values are within the tolerance limits. Also listed are the SD results produced when a false arising is generated citing a fault in the RH wing tank IV. The results shows that the RH wing tank level, fuel flow rate and fuel flow pressure variables all exceeded their tolerances and would therefore enable the arising to be correctly filtered.

Table 5.6: SD of fuel rig variables - RH wing tank IV fault

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.666cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.548cm	0.480cm
LH Wing Tank Level	1.500cm	0.730cm	0.437cm
RH Wing Tank Level	1.500cm	1.132cm	**4.294cm**
LH Flow Rate	0.30L/min	0.12L/min	0.10L/min
RH Flow Rate	0.30L/min	0.16L/min	**2.13L/min**
LH Fuel Flow Pressure	9,000Pa	883Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	2,774Pa	**9,247Pa**

5.4.2.4 RH Auxiliary Tank Isolation Valve Blocked/Failed Closed

When an auxiliary tank IV becomes blocked or fails closed, the flow path from the auxiliary tank to the wing tank will be terminated. Over the duration of the mission under consideration the effect of the failure mode will be seen in the output of the auxiliary tank level and wing tank level variables.

Figures 5.29 and 5.30 show the auxiliary tank levels and wing tank levels over the mission under consideration when the RH auxiliary tank IV is blocked/failed closed.

The effect of the RH auxiliary tank IV fault can be seen in Figure 5.29. The only change in the RH auxiliary tank level during the entire mission is due to the peristaltic pump effects seen in phase 3. The blockage in the auxiliary tank IV prevents any flow from leaving the tank, as would be expected in phase 3. By comparison the LH auxiliary tank level steadily decreases throughout phase 3, as fuel leaves the tank to replenish the LH wing tank.

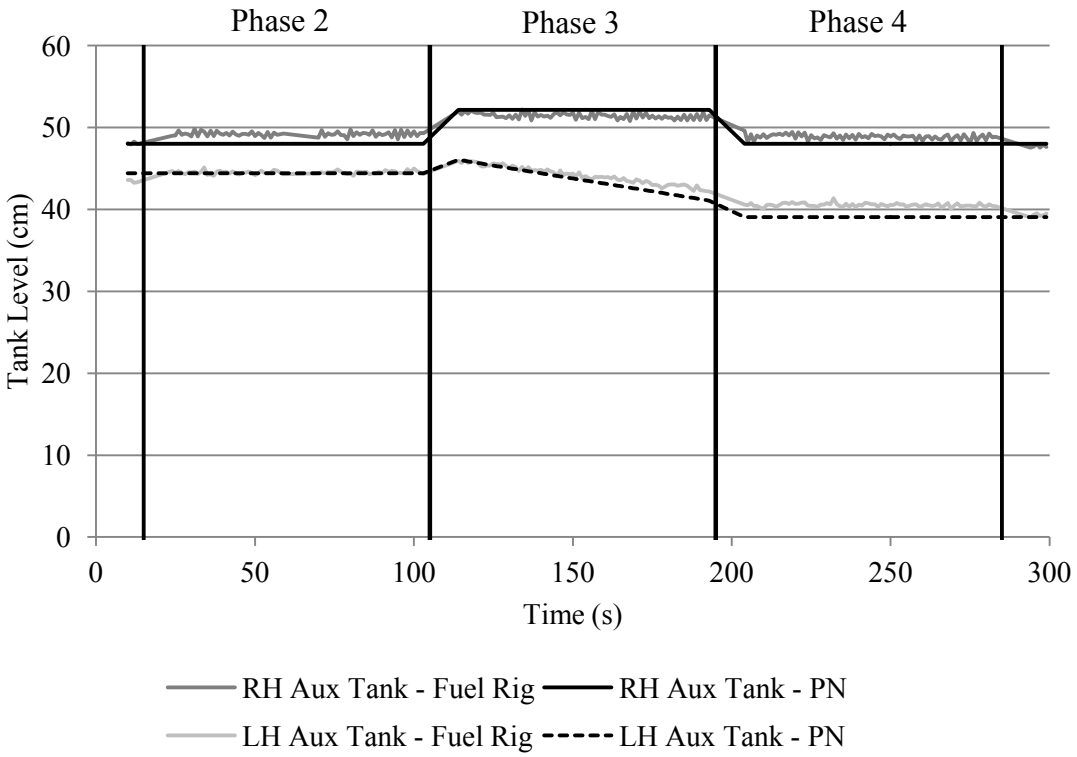

Figure 5.29: RH auxiliary tank IV blocked - Auxiliary tank levels

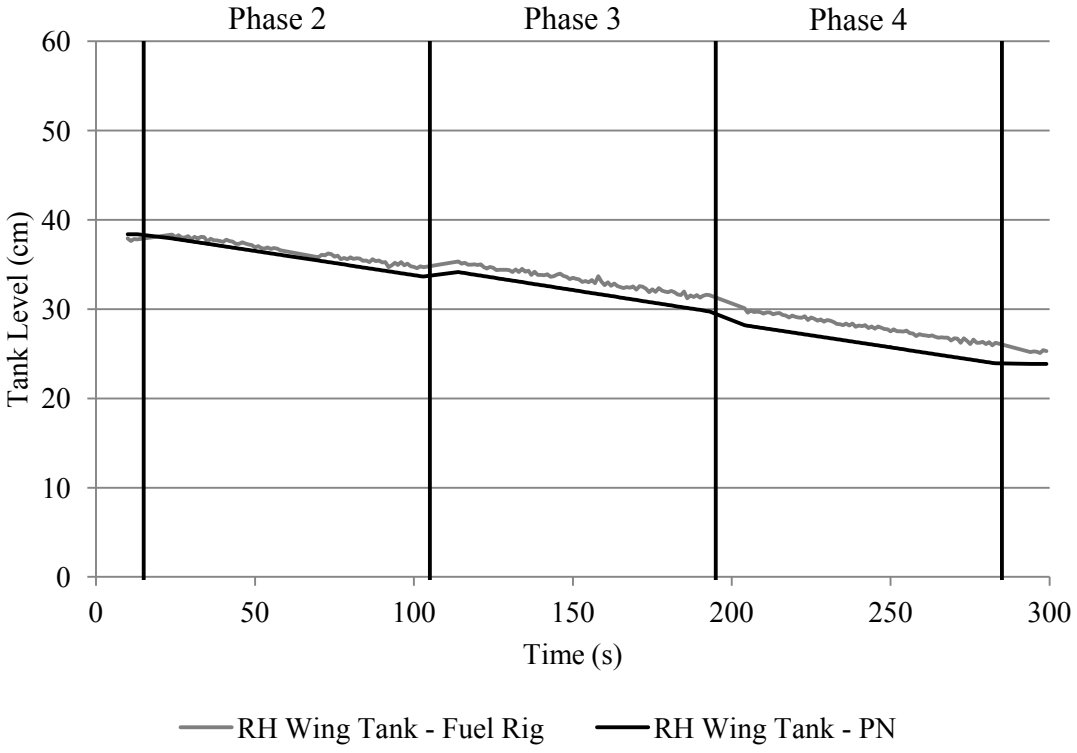

Figure 5.30: RH auxiliary tank IV blocked - RH wing tank level

In Figure 5.30 an increase in the tank level would be expected in phase 3, as shown in Figure 5.2. However, as the blockage prevents any fuel reaching the wing tank its level falls throughout the mission.

The RH auxiliary tank IV fault has affected none of the remaining fuel rig variables. This is primarily due to the fact that the RH wing tank level remains above zero throughout the mission and therefore both the LH and RH engines are supplied with fuel throughout the mission. The graphs of these variables are accurately represented by those shown in Section 5.3. The SD results for these variables are shown in Table 5.7. All of the variables are within their respective tolerances. Table 5.7 also shows the SD results when an arising related to the RH auxiliary tank IV is falsely generated. In this situation both the RH auxiliary and wing tank results exceed the specified tolerances and as a result the arising would be filtered.

Table 5.7: SD of fuel rig variables - RH auxiliary tank IV fault

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.673cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.888cm	**2.332cm**
LH Wing Tank Level	1.500cm	0.536cm	0.437cm
RH Wing Tank Level	1.500cm	0.581cm	**3.568cm**
LH Flow Rate	0.30L/min	0.07L/min	0.10L/min
RH Flow Rate	0.30L/min	0.15L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	728Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,833Pa	1,688Pa

5.4.3 Level Sensor Failure Modes

Three unique level sensor failure modes have been identified on the fuel rig system. These are; level sensor fails high, fails low and fails stuck. When the level sensor fails high, it constantly outputs the maximum tank level. Conversely, when the sensor fails low, it will constantly output a tank level of zero. Should the sensor fail stuck it will always output

the same tank level value. Table 5.3 showed that each level sensor fault only had an impact on the RH wing tank level variable. All of the remaining system variables should continue to operate as expected and as shown in Section 5.3.

5.4.3.1 RH Wing Tank Level Sensor Failed High

Figure 5.31 shows the RH wing tank level output when the fault is injected into the system after 60 seconds.

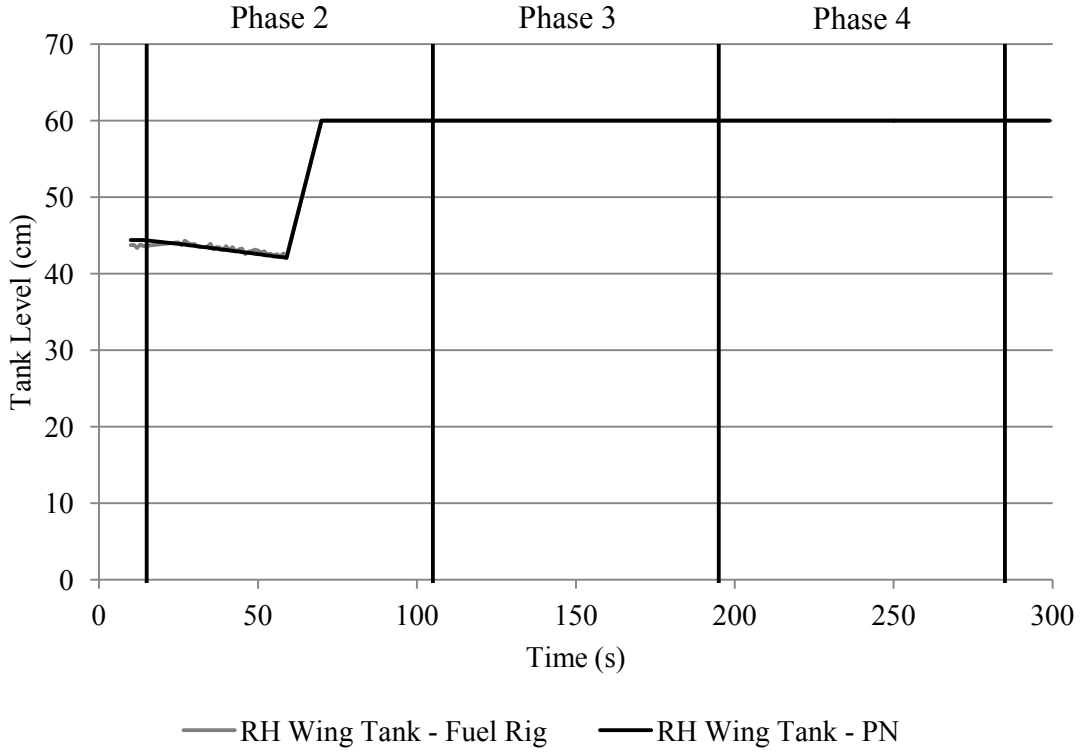

Figure 5.31: RH wing tank level sensor failed high - RH wing tank level

It can be seen in Figure 5.31 that the effect of the fault is to increase the tank level output to the maximum level of 60cm. This output does not show the actual fuel level in the RH wing tank and the performance of the system does not change as a result of the fault. It can be seen from the remaining system outputs that the system continues to operate normally - providing fuel to both engines. The figure also shows that the PN model has represented the effect of this fault very well.

Table 5.8 shows the SD values for a number of system variables. It can be seen that when the level sensor failed high fault is present in the fuel rig and PN model, all of the SD values are within the tolerance limits. When the fault is only included in the PN model, as

a result of a false arising, the SD of the RH wing tank level variable exceeds the tolerance limit. In this case, the arising would not be verified and would correctly classified as false.

Table 5.8: SD of fuel rig variables - RH wing tank level sensor failed high

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.810cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.496cm	0.480cm
LH Wing Tank Level	1.500cm	0.468cm	0.437cm
RH Wing Tank Level	1.500cm	0.158cm	**7.040cm**
LH Flow Rate	0.30L/min	0.09L/min	0.10L/min
RH Flow Rate	0.30L/min	0.19L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	820Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,373Pa	1,688Pa

5.4.3.2 RH Wing Tank Level Sensor Failed Low

Figure 5.32 shows the RH wing tank level over the course of the mission where the level sensor fails low. It can be seen from the figure that the fault has resulted in the reported wing tank level falling to 0cm after 60 seconds. Again, this output does not match the true state of the RH wing tank level and the operation of the system continues as normal. For example the RH flow rate graph, in the presence of this failure mode, shows a normal level of operation which indicates fuel is being provided to the RH engine. The RH flow rate graph is similar to that shown in Section 5.3.

The SD value for each variable is shown in Table 5.9. All of the SD values are within the respective tolerances when the failure mode is present in the fuel rig and included in the PN model. Also shown are the SD values produced when the 'RH wing tank level sensor failed low' fault is incorrectly diagnosed. It can be seen from the table that in this case, the SD of the RH wing tank level variable is significantly greater than the permissible tolerance limit. As a result, the fault verification technique would correctly filter the fault.

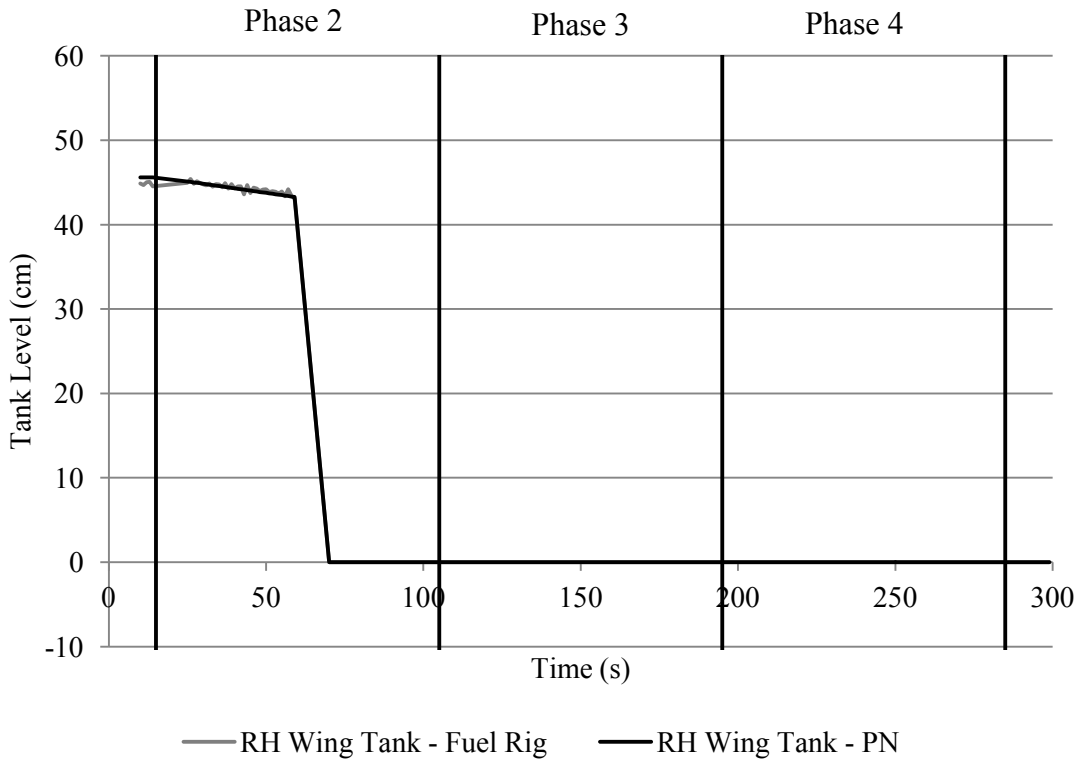

Figure 5.32: RH wing tank level sensor failed low - RH wing tank level

Table 5.9: SD of fuel rig variables - RH wing tank level sensor failed low

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.867cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.553cm	0.480cm
LH Wing Tank Level	1.500cm	0.535cm	0.437cm
RH Wing Tank Level	1.500cm	0.165cm	**15.493cm**
LH Flow Rate	0.30L/min	0.09L/min	0.10L/min
RH Flow Rate	0.30L/min	0.17L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	824Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,473Pa	1,688Pa

5.4.3.3 RH Wing Tank Level Sensor Failed Stuck

When the RH wing tank level sensor fails stuck, it will continuously output the value that was last output when the sensor was working correctly. Figure 5.33 shows the RH wing

tank level on the fuel rig, when the level sensor fails stuck.

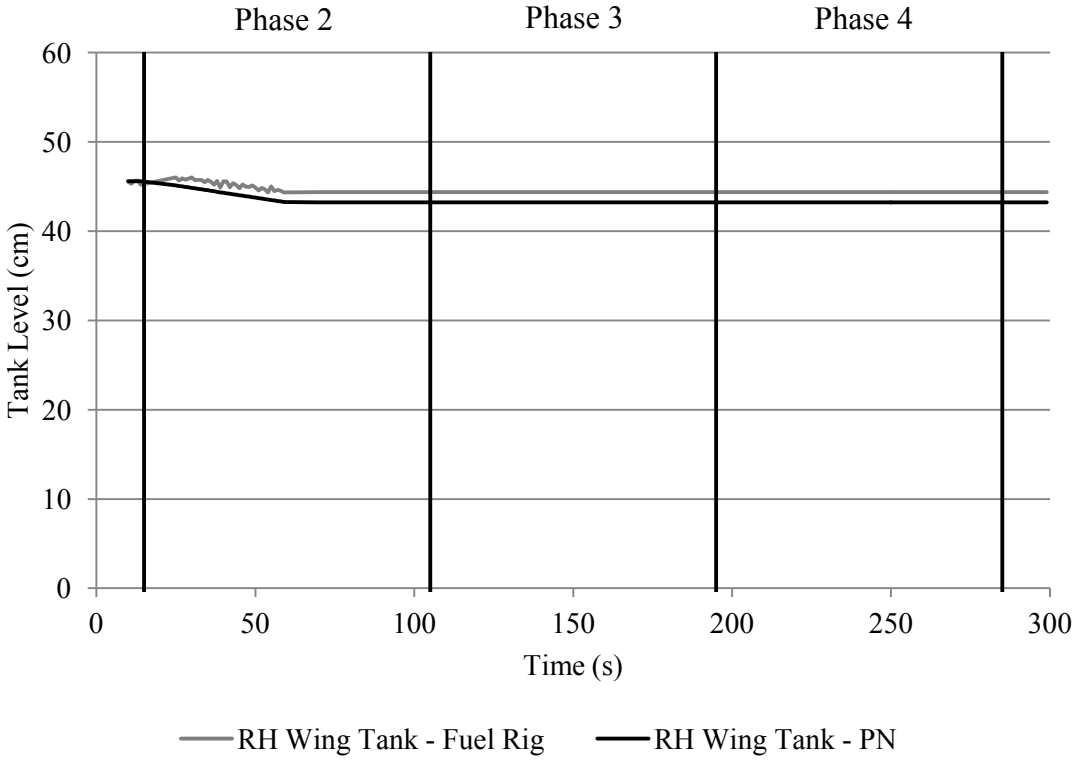

Figure 5.33: RH wing tank level sensor failed stuck - RH wing tank level

It can be seen in Figure 5.33 that the RH wing tank level recorded from the fuel rig remains constant at approximately 45cm from the time of the fault occurring throughout the remainder of the mission. The recorded value does also not change at the start and end of phase 3, where otherwise it would be expected to due to the control inputs and vibrational effects of the auxiliary pumps. The PN model has accurately reflected the behaviour of the RH wing tank level variable in the presence of the fault. None of the remaining fuel rig variables have been affected by the failure mode during the phased mission and their behaviour reflects that seen in the figures of Section 5.3.

The SD result of every variable is within the respective tolerance when the fault is present in the fuel rig and the PN. These SD results are shown in detail in Table 5.10. The table also displays the SD values that would be produced should a false arising be generated that stated the RH wing tank level sensor had failed stuck.

Considering the false arising results in column four of Table 5.10, it can be seen that none of the variables have exceeded their tolerance. This is unexpected, as when these SD values were calculated, the level sensor fault was only present in the PN model and not

Table 5.10: SD of fuel rig variables - RH wing tank level sensor failed stuck

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.769cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.492cm	0.480cm
LH Wing Tank Level	1.500cm	0.470cm	0.437cm
RH Wing Tank Level	1.500cm	0.202cm	1.309cm
LH Flow Rate	0.30L/min	0.09L/min	0.10L/min
RH Flow Rate	0.30L/min	0.16L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	885Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,435Pa	1,688Pa

the fuel rig. It would be expected that the SD of the RH wing tank level variable would exceed the tolerance limit. Figure 5.34 plots the RH wing tank level curves produced from the fuel rig and PN model, when the level sensor fault is falsely diagnosed.

Figure 5.34 shows the PN predicted tank level remains constant from the time of the fault occurring onwards. By comparison, the recorded tank level varies through the mission. Despite the fact that the fuel rig tank level value changes throughout the mission, the values remain relatively close to the equivalent PN tank level values. As a result the SD of the residual values of the two curves is 1.309cm, which is within the tolerance limit. This explains why the fault verification software would not filter the false arising.

It can be seen in Figure 5.34 that from approximately 250 seconds onwards the two tank level curves begin to diverge. Had the length of phase 4 been longer the difference in the tank levels would have continued to increase. As a result the RH wing tank level SD value would also be greater. These results indicate a limitation of the fault verification technique where short phases and missions do not allow enough time for some variables to diverge sufficiently and for the resultant SD value to exceed the variable tolerances. In its current form however, the fault verification technique is unable to filter out a false arising of the 'RH Level Sensor Failed Stuck' fault.

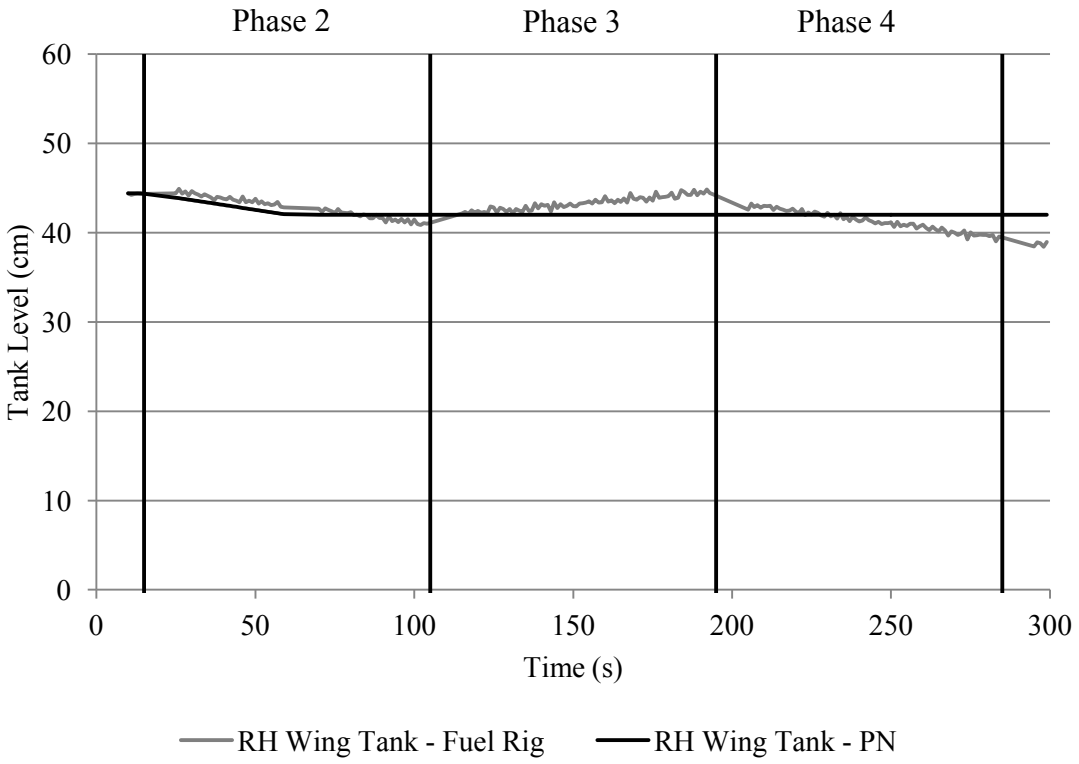

Figure 5.34: RH wing tank level sensor failed stuck false arising - RH wing tank level

5.4.4 Fuel Flow Rate Sensor Failure Modes

The fuel flow rate sensor on the fuel rig system can fail in a number of different ways. Those are failed high, failed off and failed stuck. Each of these failure modes will be considered individually.

5.4.4.1 RH Fuel Flow Rate Sensor Failed High

Should a fuel flow rate sensor fail high, the output from the sensor will increase to its maximum value and remain at that value. Figure 5.35 shows the RH fuel flow rate output from the fuel rig and PN, when the RH fuel flow rate senor fails high in the phased mission under consideration.

The figure clearly shows that the effect of the flow rate sensor fault is to increase the sensor output from approximately 5L/min to 60L/min. Having reached this value, the fuel flow rate then remains at this level for the remainder of the mission. The PN predicted fuel flow rate can be seen to correspond well with the sensor outputs recorded from the fuel rig system.

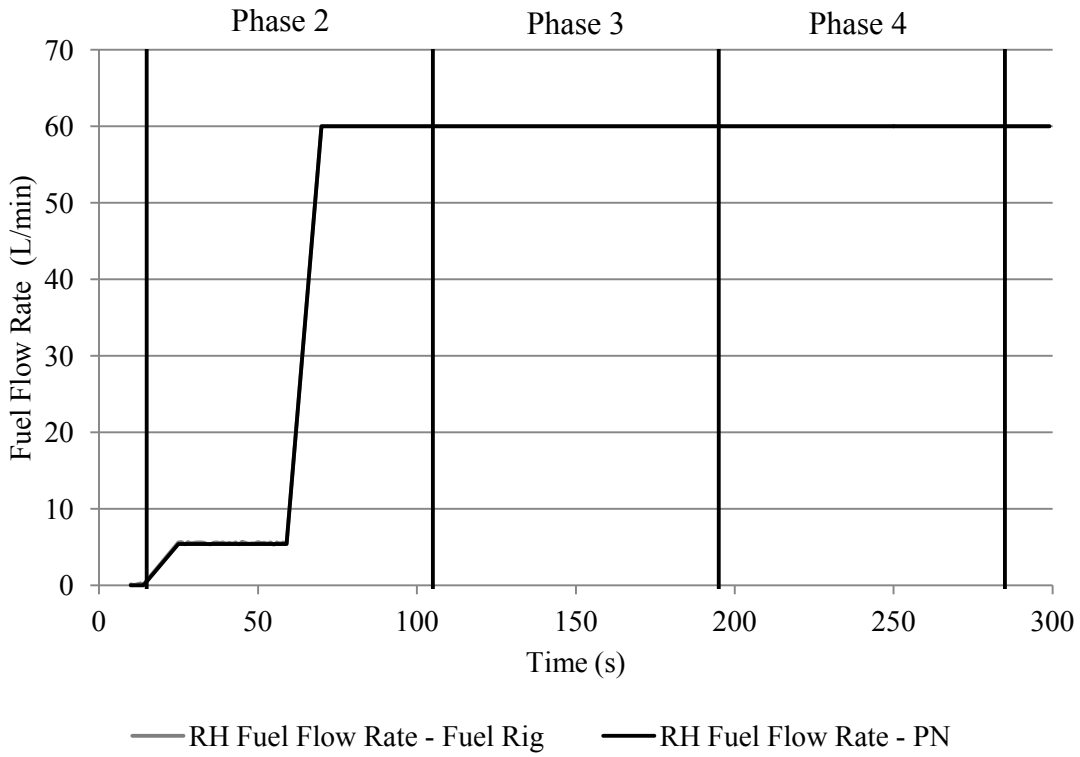

Figure 5.35: RH fuel flow rate sensor failed high - RH fuel flow rate

Table 5.11 lists the SD values determined when the 'RH Fuel Flow Rate Sensor Failed High' fault is correctly and incorrectly diagnosed.

Table 5.11: SD of fuel rig variables - RH fuel flow rate sensor failed high

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.646cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.567cm	0.480cm
LH Wing Tank Level	1.500cm	0.479cm	0.437cm
RH Wing Tank Level	1.500cm	0.567cm	0.506cm
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.07L/min	**2.13L/min**
LH Fuel Flow Pressure	9,000Pa	794Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,291Pa	1,688Pa

When correctly diagnosed, only the RH fuel flow rate variable behaviour changes yet all of the variables are within the tolerance limits. This indicates that the fault verification technique can verify this fault. When the fault is incorrectly diagnosed, the SD of the RH flow rate variable exceeds the tolerance limit thereby ensuring the arising is filtered as false.

5.4.4.2 RH Fuel Flow Rate Sensor Failed Off

The fuel flow rate sensor on the fuel rig may fail off as a result of a mechanical or electrical fault with the sensor. Should either of these circumstances occur, the fuel flow rate sensor will not produce an electrical output to the fuel rig computer. As a result, the fuel flow rate recorded by the system computer will only reflect the scaling and offset value applied to the electrical output of the flow rate sensor. Figure 5.36 shows the fuel flow rate values when the flow rate sensor fails off.

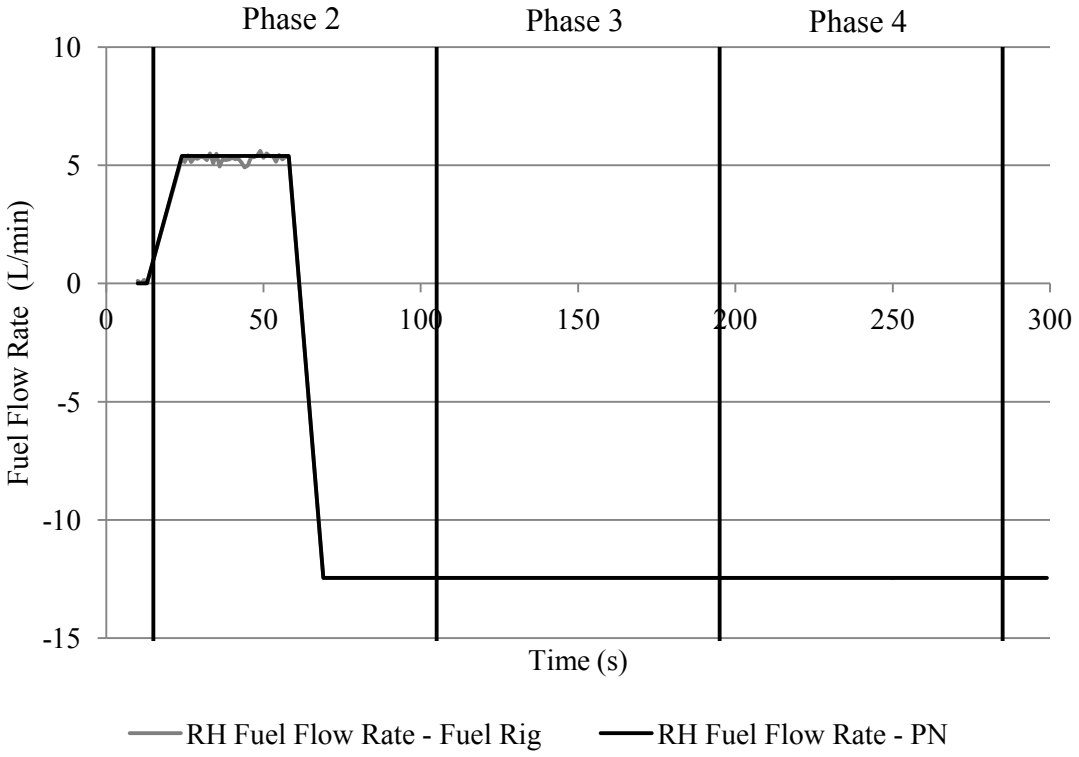

Figure 5.36: RH fuel flow rate sensor failed off - RH fuel flow rate

Figure 5.36 shows that after the fault occurs the fuel flow rate recorded by the system computer is approximately -12.5L/min. This value represents the scaling and offset figure that is applied to the electrical output from the fuel flow rate sensor to convert it to a flow

rate value. As the electrical output in the presence of the fault is zero, only the scaling and offset figure is recorded by the fuel rig. Figure 5.36 shows that the PN has modelled the effect of the fault on the flow rate variable output well.

None of the remaining fuel rig system variable outputs have been affected by the RH fuel flow rate sensor fault. Graphs of their outputs are similar in shape to those shown in Section 5.3. The SD of each variable is listed in Table 5.12. Also listed is the SD values produced when a false arising is produced that states the RH fuel flow rate sensor has failed off. Should this arising be generated, the SD of the RH fuel flow rate variable can be seen to exceed its tolerance resulting in the arising being filtered.

Table 5.12: SD of fuel rig variables - RH fuel flow rate sensor failed off

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.545cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.594cm	0.480cm
LH Wing Tank Level	1.500cm	0.466cm	0.437cm
RH Wing Tank Level	1.500cm	0.451cm	0.506cm
LH Flow Rate	0.30L/min	0.11L/min	0.10L/min
RH Flow Rate	0.30L/min	0.07L/min	**6.59L/min**
LH Fuel Flow Pressure	9,000Pa	1,262Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,572Pa	1,688Pa

5.4.4.3 RH Fuel Flow Rate Sensor Failed Stuck

When the fuel flow rate sensor becomes stuck, its output remains constant at the same value that was last output when the sensor was working correctly. Figure 5.37 shows the RH fuel flow rate output, when the flow rate sensor fails stuck during the phased mission.

Figure 5.37 shows that, as expected, the fuel flow rate recorded from the fuel rig has remained constant from the time of the fault occurring. The lack of noise in the recorded output and the fact that the fuel flow rate does not decrease in phase 5 when the engine pump rating falls to 0% are evidence that the sensor output has become stuck. The PN can be seen to model this behaviour accurately.

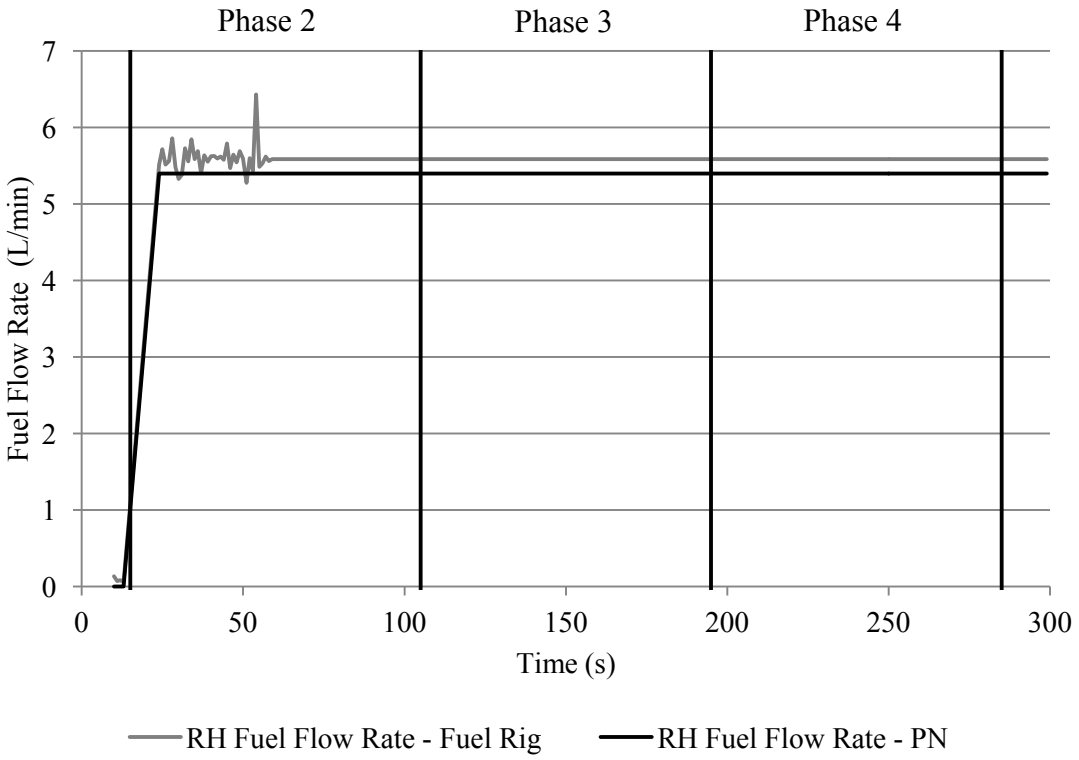

Figure 5.37: RH fuel flow rate sensor failed stuck - RH fuel flow rate

The fuel rig output variables not discussed above have not been affected by the flow rate sensor fault. Their behaviour is consistent with that shown in the graphs of Section 5.3. When the fault has been diagnosed correctly the SD of each variable is within the specified tolerances as shown in Table 5.13. The table also lists the SD values produced when the fuel flow rate failed stuck fault is identified incorrectly. It can be seen that in this scenario the RH fuel flow rate variable has been exceeded. As a result the arising which details this fault would be filtered preventing unnecessary maintenance resources from being used.

Table 5.13 show that the RH fuel flow rate SD value is more than 10 times larger when diagnosed incorrectly compared to when it is diagnosed correctly. This is predominantly due to the fact that in phase 5, when diagnosed incorrectly, there is a significant variation in the fuel rig and PN flow rate values. As a result, the SD of the RH flow rate variable exceeds the SD tolerance. This outcome contrasts with that produced when the level sensor failed stuck. As was shown in Section 5.4.3.3, when the level sensor failed stuck there was only a small difference between the two sets of results and this kept the RH wing tank level SD within the tank level tolerance.

Table 5.13: SD of fuel rig variables - RH fuel flow rate sensor failed stuck

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.638cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.530cm	0.480cm
LH Wing Tank Level	1.500cm	0.501cm	0.437cm
RH Wing Tank Level	1.500cm	0.567cm	0.506cm
LH Flow Rate	0.30L/min	0.09L/min	0.10L/min
RH Flow Rate	0.30L/min	0.07L/min	**0.85L/min**
LH Fuel Flow Pressure	9,000Pa	898Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,314Pa	1,688Pa

5.4.5 Flow Pressure Sensor Failure Modes

The flow pressure sensors on the fuel rig can fail in one of three ways; fail high, fail off or fail stuck. The consequences of these failure modes occurring are the same on the flow pressure sensors as they are on the fuel flow rate sensors. For example, if the flow pressure sensor fails high it will constantly output its maximum possible value.

5.4.5.1 RH Flow Pressure Sensor Failed High

The RH flow pressure sensor output over the course of the mission is shown in Figure 5.38. The RH flow pressure sensor is failed high after 60 seconds. Figure 5.38 shows that when the flow pressure sensor fails high the recorded output increases to 320,000Pa. This output value remains constant throughout the remainder of the mission. The PN model of the fuel rig system has accurately represented this behaviour, as can be seen in the figure.

The behaviour of all of the remaining output variables from the fuel rig can be seen in Section 5.3. They do not show any behaviour that illustrates the flow pressure sensor failure has affected them. The SD of every variable is listed in Table 5.14 below. It can be seen that, when diagnosed correctly, all of the SD values are within the tolerance limit and as a result the arising would be verified. The SD values produced as a result of the fault being diagnosed incorrectly are also listed in column four. As would be expected,

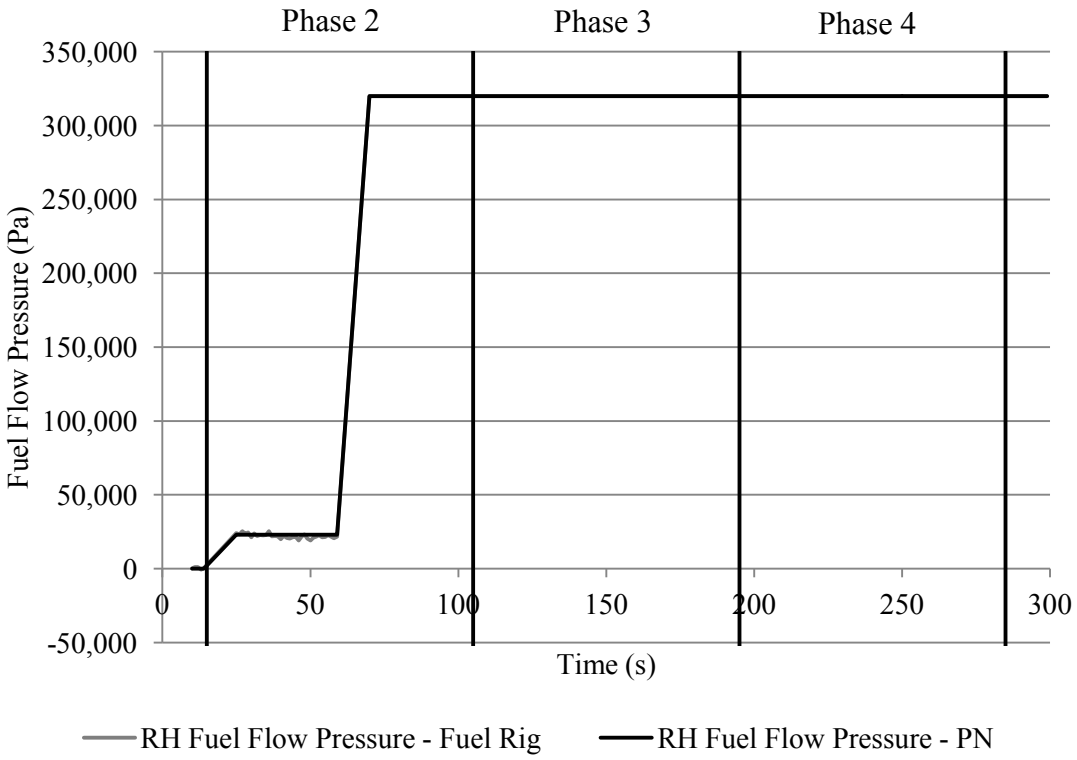

Figure 5.38: RH flow pressure sensor failed high - RH fuel flow pressure

Table 5.14: SD of fuel rig variables - RH flow pressure sensor failed high

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.659cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.530cm	0.480cm
LH Wing Tank Level	1.500cm	0.527cm	0.437cm
RH Wing Tank Level	1.500cm	0.510cm	0.506cm
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.17L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	844Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	673Pa	**111,240Pa**

only the RH flow pressure sensor variable has exceeded its tolerance when the sensor is falsely diagnosed as having failed high. The extremely large SD value is produced as a

result of the PN predicted flow pressure being 320,000Pa from the time at which the fault is injected in the model. By contrast the flow pressure value recorded from the fuel rig remains around 25,000Pa until phase 5 when it falls to zero.

5.4.5.2 RH Flow Pressure Sensor Failed Off

Figure 5.39 shows the fuel rig and PN flow pressure outputs when the flow pressure sensor fails off in phase 2 of the phased mission.

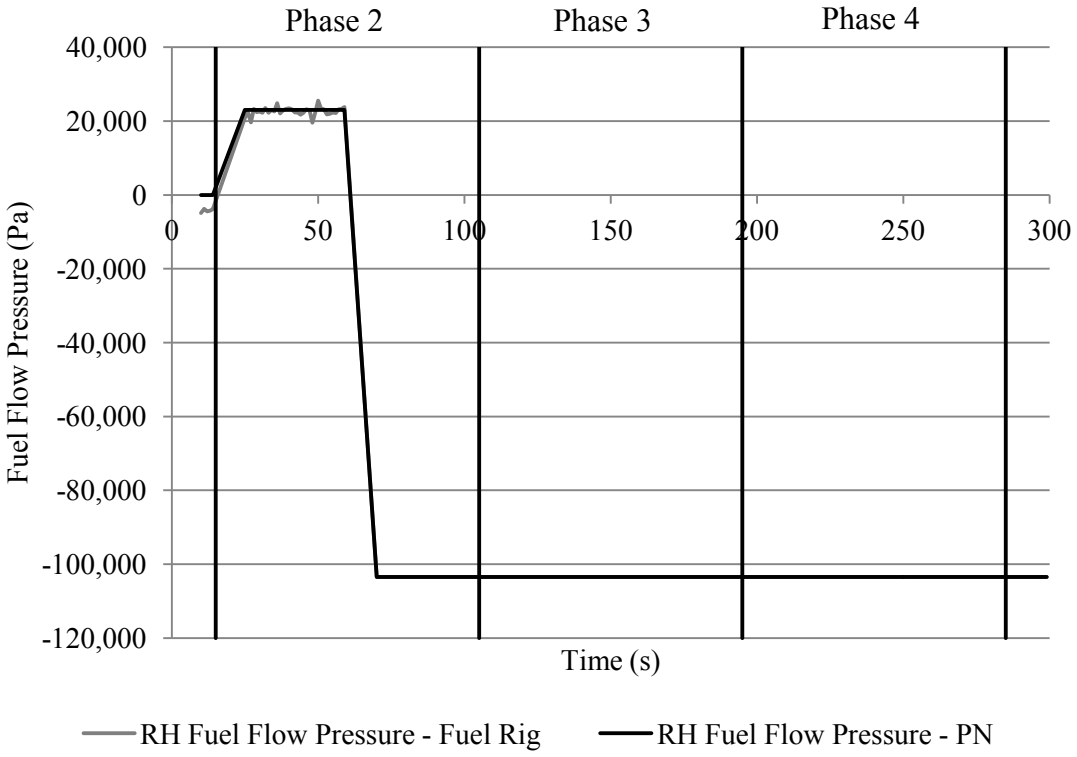

Figure 5.39: RH flow pressure sensor failed off - RH flow pressure

Figure 5.39 shows that when the flow pressure sensor fault occurs, the values recorded from the fuel rig fall to just under -100,000Pa. This value represents the scaling and offset figure that is applied to the electrical output from the flow pressure sensor. As was seen with the flow rate sensor when it failed off, with no sensor output to consider the fuel rig computer records only the scaling/offset factor. The PN model of the flow pressure behaviour matches well with the values recorded from the fuel rig system.

Graphs of all the other fuel rig variables show system behaviour that would be expected with no faults present. This can be seen in Section 5.3. The effect of the flow pressure sensor fault is therefore only seen in the flow pressure sensor variable. The fuel rig SD

results are listed in Table 5.15. All of the results are within the respective tolerances when the sensor fault is correctly diagnosed. Also shown are the SD values when the flow pressure sensor is incorrectly diagnosed as having failed off. Should this scenario occur the RH flow pressure variable can be seen to exceed its tolerance and as a result the arising would be filtered.

Table 5.15: SD of fuel rig variables - RH flow pressure sensor failed off

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.487cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.928cm	0.480cm
LH Wing Tank Level	1.500cm	0.563cm	0.437cm
RH Wing Tank Level	1.500cm	0.458cm	0.506cm
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.18L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	1,467Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	752Pa	**46,890Pa**

5.4.5.3 RH Flow Pressure Sensor Failed Stuck

The output from the RH flow pressure sensor when it fails stuck during the phased mission is shown in Figure 5.40. It can be seen on Figure 5.40 that when the flow pressure sensor fails stuck, the fuel rig output remains constant at the last value output when the sensor was working correctly. As the sensor is stuck, there is no noise in the recorded sensor output. Also, the sensor output does not change in any of the subsequent phases of system operation. The PN predicted output of the RH flow pressure sensor can be seen to be very similar to that recorded from the fuel rig.

The RH flow pressure output is the only variable that is affected by the flow pressure sensor fault. The behaviour of all the other variables on the system is the same as that when there are no faults present. This behaviour can be seen in the graphs of Section 5.3. Table 5.16 lists the SD values when the flow pressure sensor fault is correctly and

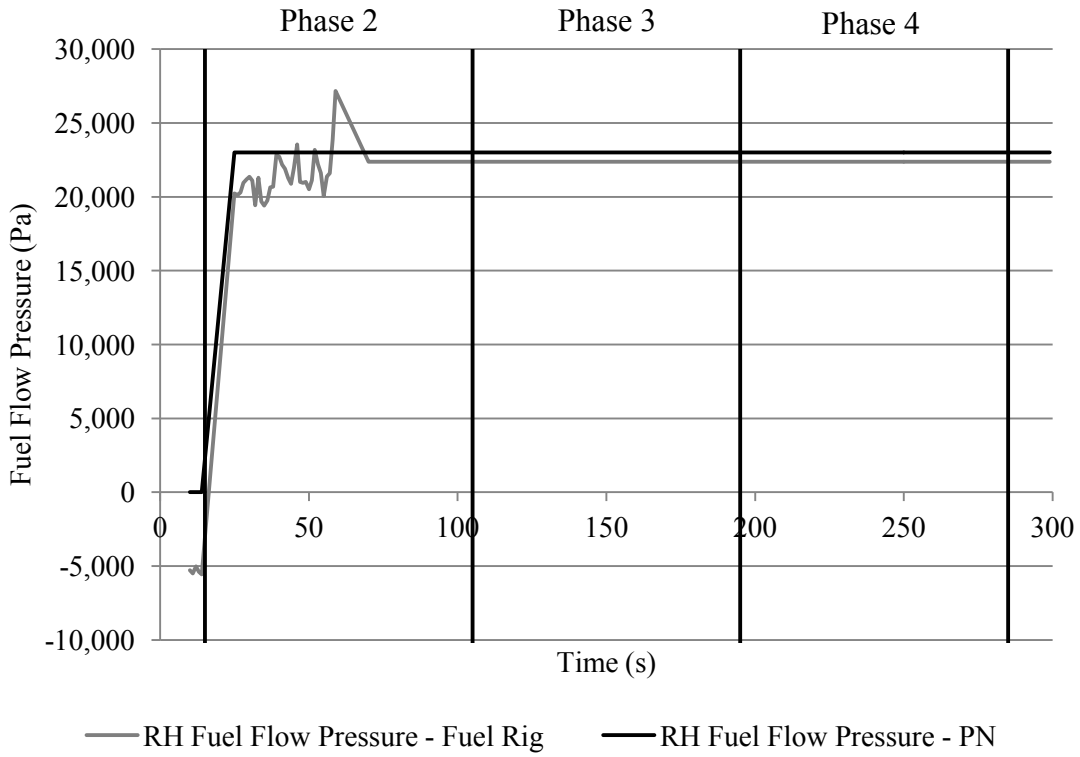

Figure 5.40: RH fuel flow pressure sensor failed stuck - RH fuel flow pressure

incorrectly diagnosed. When correctly diagnosed, the SD calculated for each variable is within the respective tolerance limit.

Table 5.16: SD of fuel rig variables - RH flow pressure sensor failed stuck

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.590cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.594cm	0.480cm
LH Wing Tank Level	1.500cm	0.565cm	0.437cm
RH Wing Tank Level	1.500cm	0.474cm	0.506cm
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.21L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	1,264Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	930Pa	4,474Pa

Table 5.16 shows that when the fault is falsely diagnosed, all of the SD values for the fuel rig output variables fall within their respective tolerances. It would be expected that, as with the other RH flow pressure sensor faults, the SD of the RH flow pressure variable would exceed the tolerance limit. However, this is not the case. Figure 5.41 shows the fuel rig and PN values of the RH flow pressure variable when the fault is falsely diagnosed.

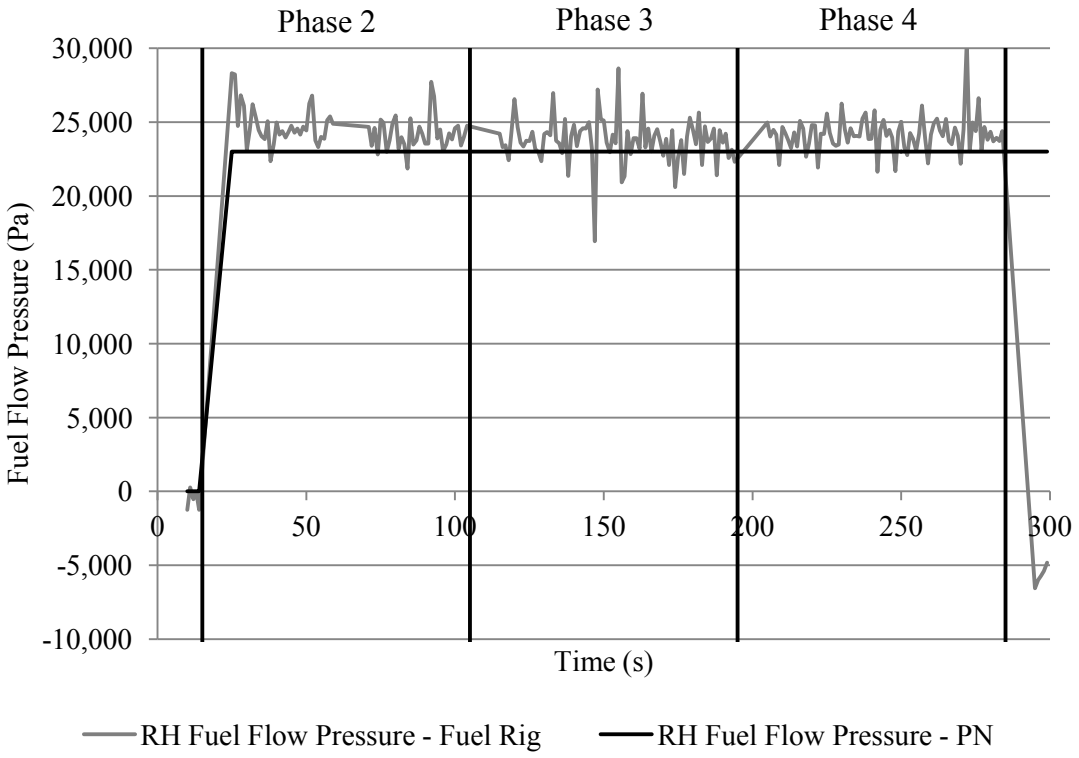

Figure 5.41: RH fuel flow pressure sensor failed stuck - RH fuel flow pressure

It can be seen in Figure 5.41 that the only major variation between the two curves appears in phase 5. In this phase the PN flow pressure remains at 23,000Pa while the fuel rig flow pressure falls to less than zero. However, as phase 5 is relatively short the SD of the entire mission remains below 9,000Pa. Were phase 5 to be extended, the SD of the RH fuel flow pressure variable would increase. Table 5.17 shows how the SD of the flow pressure variable increases as phase 5 is extended.

The table shows that phase 5 of the mission would have to be extended from 15 to 45 seconds in order for the flow pressure tolerance to be exceeded and the arising to be identified as false. In the current phased mission arrangement, the behaviour of the RH flow pressure outputs are too similar for too much of the mission an the tolerance too large for the fault to be identified as false using the techniques applied.

Table 5.17: Effect of increasing duration of phase 5 on RH fuel flow pressure SD

Phase 5 Length	RH Fuel Flow Pressure SD
15sec	4,474Pa
25sec	7,289Pa
35sec	8,916Pa
45sec	**10,194Pa**
55sec	**11,257Pa**

5.4.6 High Level Switch Failure Modes

High level switches are installed within the fuel tanks to provide an indicator of when the tank is full and, in association with the tank level sensors, to prevent overfilling. Three failure modes can occur; fail on, fail off and fail stuck. The high level switch output is a digital signal of either '0' or '1'. If the output is '0', this indicates that the switch is on and the fuel level in the tank is equal to or greater than that of the switch height. If the output is '1' the fuel level is below that of the switch height and the switch is off. The high level switches are located at the top of the 60cm tall fuel tanks on the fuel rig.

5.4.6.1 RH Wing Tank High Level Switch Failed On

If the high level switch fails on, the output from the switch will be '0'. Figure 5.42 shows the output from the RH wing tank high level switch when it fails on during the mission.

Figure 5.42 shows that the impact of the high level switch failing on is to change the switch output from '1' to '0'. The latter switch output, therefore, reports that the fuel tank is full. Figure 5.43 shows the RH tank level over the course of the mission. It can be seen in Figure 5.43 that the RH wing tank level is always below 60cm during the mission and therefore, the high level switch output does not represent the true state of the system after the fault occurs. The PN predicted behaviour of the RH wing tank high level switch output matches exactly with that recorded from the fuel rig during the mission. As a result the SD of the RH wing tank high level switch variable is zero and within the tolerance limit. In the presence of the RH wing tank high level switch fault, only the behaviour of the RH wing tank high level switch variable is affected. The behaviour of all the other system variables is similar to that shown in Section 5.3.

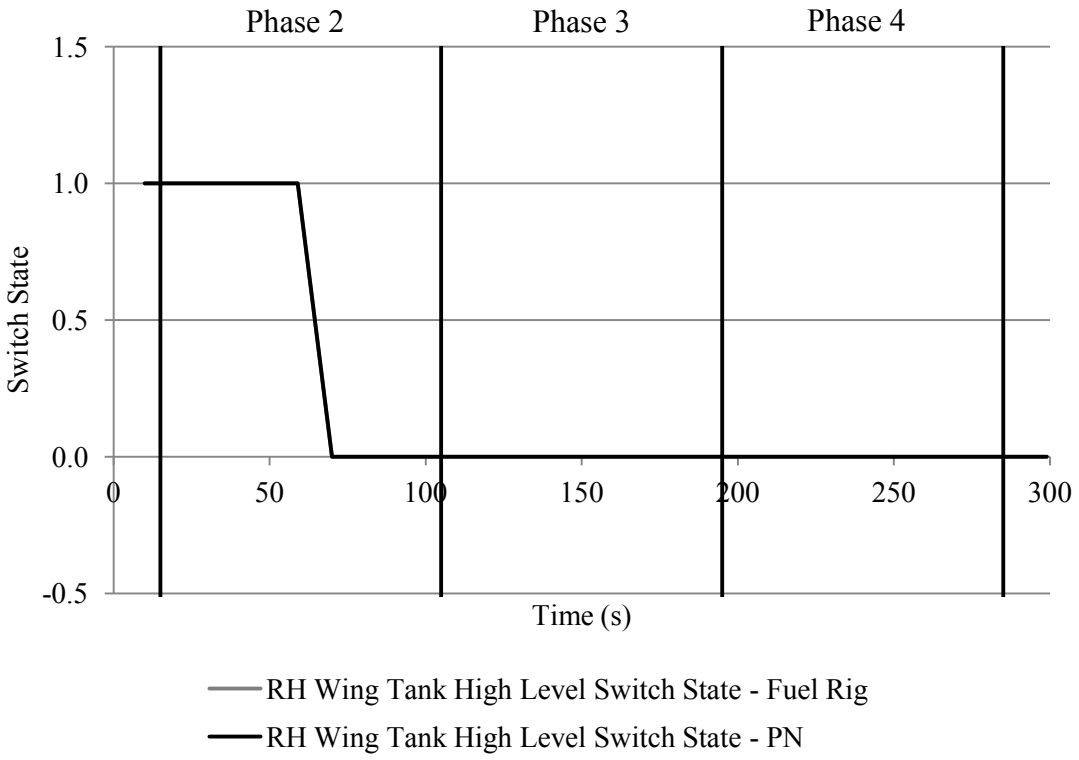

Figure 5.42: RH wing tank high level switch failed on - High level switch state

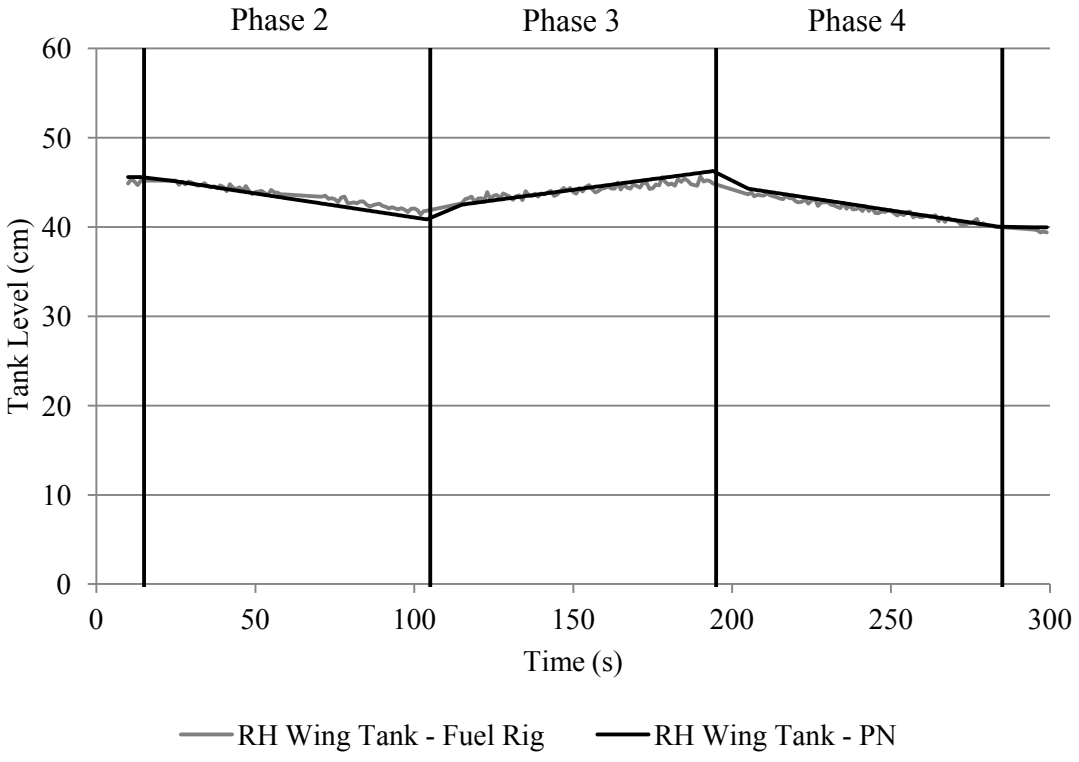

Figure 5.43: RH wing tank high level switch failed on - RH wing tank level

Table 5.18 shows that when the high level switch fault is correctly diagnosed, all the SD values of the high and low level switch variables are within the tolerance limits. As none of the other variables are affected by the fault, their SD results have been omitted for brevity.

Table 5.18: SD of fuel rig variables - RH wing tank high level switch failed on

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Wing Tank High Level Switch	0.1	0.0	0.0
RH Wing Tank High Level Switch	0.1	0.0	**0.37**
LH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Aux Tank High Level Switch	0.1	0.0	0.0
RH Aux Tank Low Level Switch	0.1	0.0	0.0

Column four of Table 5.18 shows that when the high level switch fault is falsely diagnosed, the SD of the RH wing tank high level switch variable exceeds its tolerance. When falsely diagnosed, the fuel rig output of the high level switch does not change state. However the fault will be present in the PN model and therefore, the PN high level switch output does change state. This results in the SD value exceeding the high level switch tolerance limit. As a result the false arising is correctly identified and can be filtered.

5.4.6.2 RH Wing Tank High Level Switch Failed Off

When the high level switch fails off, its output will constantly be '1' or off. It was shown in Section 5.3 that in normal operation the RH wing tank high level switch is off throughout the entire mission. Injecting the fault in the fuel rig therefore has no effect on the behaviour of any system variables. The occurrence of this fault creates a hidden failure that would not be revealed, until the RH wing tank level reaches the level of the high level switch.

As none of the system variables are affected by this failure mode, all of the graphs in Section 5.3 give an accurate representation of the system behaviour in the presence of the fault. The SD values calculated from the graphs of all the switch variables are listed in Table 5.19. All of the SD values are within the tolerances listed for the respective variables.

The SD values found, when the high level switch fault is incorrectly diagnosed, are also listed. It can be seen that again none of the SD values exceed the tolerances. Therefore, if an arising included the failure mode 'RH Wing Tank High Level Switch Failed Off', the fault verification technique would verify it as genuine whether it was true or false.

Table 5.19: SD of fuel rig variables - RH wing tank high level switch failed off

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Wing Tank High Level Switch	0.1	0.0	0.0
RH Wing Tank High Level Switch	0.1	0.0	0.0
LH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Aux Tank High Level Switch	0.1	0.0	0.0
RH Aux Tank Low Level Switch	0.1	0.0	0.0

5.4.6.3 RH Wing Tank High Level Switch Failed Stuck

If the high level switch fails stuck, it will continue to output the last correct switch state prior to the component failing. When the high level switch fails stuck in the phased mission it's output will remain '1' or off. However, in normal operation the high level switch output will be off throughout the mission. Therefore the switch failing stuck has no effect on the fuel rig system or any of its variables, it is a hidden fault. As none of the fuel rig system variables are affected by the failure mode, their behaviour can be seen from the graphs in Section 5.3. The SD values produced in the presence of this fault are the same as those shown in Table 5.19.

5.4.7 Low Level Switch Failure Modes

Low levels switches are used on the fuel rig to indicate, when the fuel level in the tanks approaches zero. They are used with the tank level sensors to monitor the fuel level in the tanks and prevent pumps from 'running dry'. The operation, output signals and failure modes of the low level switches are the same as those produced/experienced by the high

level switches.

5.4.7.1 RH Wing Tank Low Level Switch Failed Off

Section 5.3 showed that in normal operation the RH wing tank low level switch was on throughout the phased mission. The switch therefore produced a constant '0' output. If the low level switch were to fail off, the switch output would change to '1'. Figure 5.44 shows the RH wing tank low level switch output during the phased mission when the low level switch fails off.

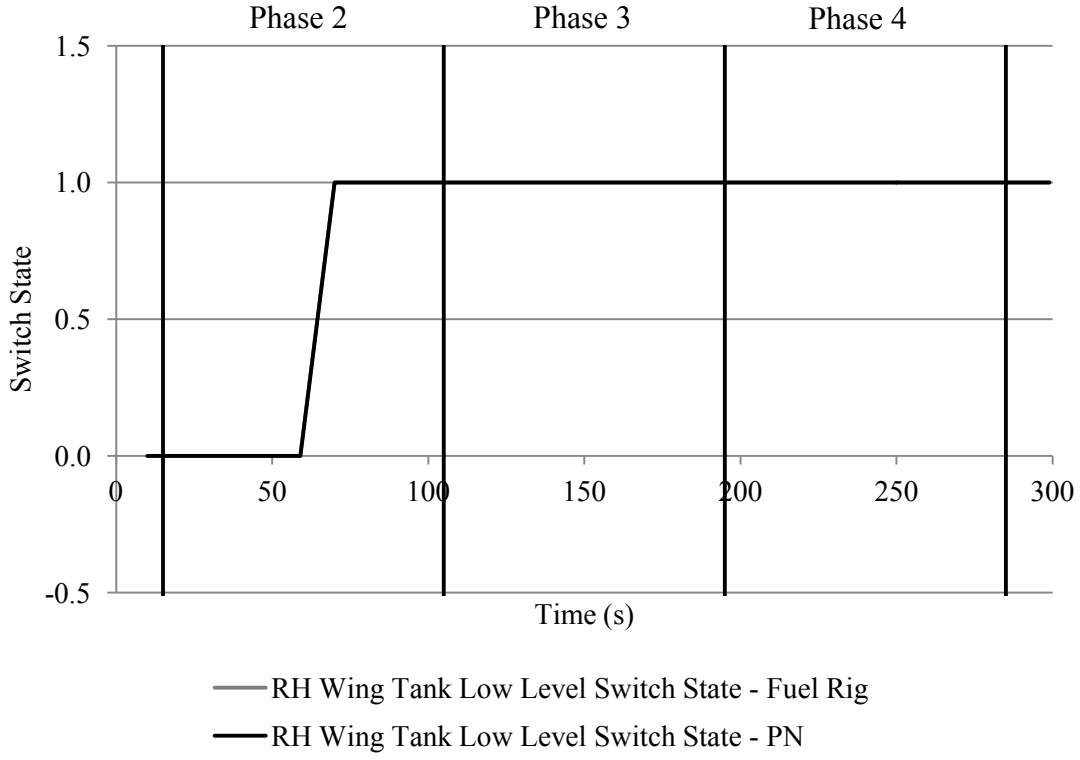

Figure 5.44: RH wing tank low level switch failed off - Low level switch state

Figure 5.44 shows that when the low level switch fails off, the output from the switch changes from '0' to '1'. The switch remains off from the time of the fault occurring till the end of the mission. Figure 5.45 shows the RH wing tank level over the course of the same mission. It can be seen from Figure 5.45 that the tank level remains between approximately 40 and 50cm throughout the mission. The low level switch failure therefore has no effect on the actual tank level in the RH wing tank.

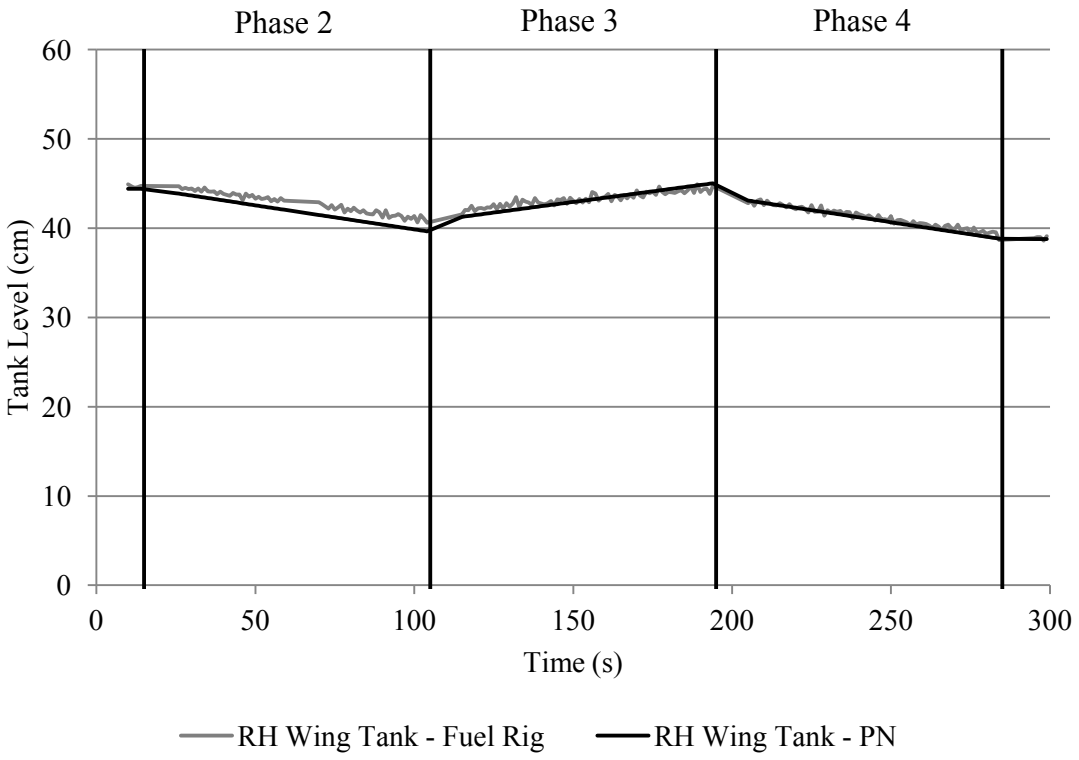

Figure 5.45: RH wing tank low level switch failed off - RH wing tank level

The PN predicted behaviour of the RH wing tank low level switch state can be seen to match very well with that recorded from the fuel rig system. As a result when the fault is diagnosed correctly, the SD of the RH wing tank low level switch variable is zero. As the failure mode also has no effect on any of the other system variables, all other SD values also fall within the respective tolerance limits. Graphs of these variables behaviour are similar to those shown in Section 5.3. Also listed are the SD values produced as a result of an incorrect diagnosis of the low level switch failing off. It can be seen that should this scenario occur, the RH wing tank low level switch variable will exceed its tolerance. Therefore, the fault verification technique would be able to correctly identify both a genuine and false arising in this case.

5.4.7.2 RH Wing Tank Low Level Switch Failed On

When the RH wing tank low level switch fails on, it will continue to produe an output of '0' irrelevant of the actual tank level. Nonetheless, as shown in Section 5.3 the expected output of the low level switch is already '0' throughout the mission. Any occurrence of the fault 'low level switch failed off' in the phased mission under consideration would therefore

Table 5.20: SD of fuel rig variables - RH wing tank low level switch failed off

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Wing Tank High Level Switch	0.1	0.0	0.0
RH Wing Tank High Level Switch	0.1	0.0	0.0
LH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Wing Tank Low Level Switch	0.1	0.0	**0.37**
RH Aux Tank High Level Switch	0.1	0.0	0.0
RH Aux Tank Low Level Switch	0.1	0.0	0.0

be hidden. Evaluating the fuel rig system outputs when this fault is injected reveals this to be the case. The graphs in Section 5.3 therefore give an accurate representation of the system behaviour in the presence of this fault.

The SD values of the switch variables on the system are presented in Table 5.21. All of the SD values calculated when the fault is correctly diagnosed are within the tolerance limits. Also listed are the SD values produced as a result of the failure mode being incorrectly diagnosed. As can be seen from these results however, no variables exceed their tolerance limit, as the behaviour of the system with the fault present is the same as that without it. This means that the fault verification technique would correctly verify a genuine arising but would also incorrectly verify a false arising.

5.4.7.3 RH Wing Tank Low Level Switch Failed Stuck

In the event of the low level switch failing stuck, it will continuously output that it is either in an on or off state. The state output by the switch is dependent upon what state was last output when the switch was working correctly. In the case of the phased mission under consideration, it has been shown in Section 5.3 that the RH wing tank low level switch state remains constant throughout the mission. As a result any occurrence of the low level switch failing stuck during the mission would have no effect on the fuel rig system or its outputs. Section 5.3 therefore presents an accurate representation of the fuel rig behaviour. The results shown in Table 5.21 also represent the SD values that are produced when the low level switch fails stuck in the phased mission.

Table 5.21: SD of fuel rig variables - RH wing tank low level switch failed on

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Wing Tank High Level Switch	0.1	0.0	0.0
RH Wing Tank High Level Switch	0.1	0.0	0.0
LH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Wing Tank Low Level Switch	0.1	0.0	0.0
RH Aux Tank High Level Switch	0.1	0.0	0.0
RH Aux Tank Low Level Switch	0.1	0.0	0.0

5.4.8 Pump Failure Modes

The low pressure pumps on the fuel rig system can fail in two different modes; pump failed off and pump degraded. If the pump fails off, it will not produce any output regardless of the input demand. If the pump has become degraded, its operational capacity will be reduced and the level of degradation will determine its maximum operational capability.

5.4.8.1 RH Engine Pump Failed Off

If the RH egine pump were to fail off, the flow path from the RH wing tank to the RH engine would be lost. The failure mode would have an impact on the RH wing tank level, the RH flow rate and the RH flow pressure variables. The effect of this failure mode is therefore the same as that caused when the RH wing tank IV is blocked/failed closed and when the RH triple port L-valve IV is blocked/failed closed. The effect of the RH engine pump failing off can therefore be seen in Section 5.4.2.2 and Section 5.4.2.3. The result of applying the fault verification technique is also the same, genuine faults are verified and false faults are filtered.

5.4.8.2 RH Auxiliary Pump Failed Off

Should the RH auxiliary pump fail off the flow path from the RH auxiliary tank to the RH wing tank will be lost. This means that no fuel will be added to the wing tank in phase 3, when both auxiliary pumps should be on. The effect of this failure mode is the

same as that produced when the RH auxiliary tank IV is blocked/failed closed. The result on the fuel rig system of the auxiliary pump failing off can therefore be seen in Section 5.4.2.4. The fault verification technique can successfully verify genuine faults and identify false faults when an arising includes the failure mode 'RH Auxiliary Pump Failed Off'.

5.4.8.3 RH Auxiliary Pump Degraded 50%

If the RH auxiliary pump is degraded by 50%, its maximum output will be 50%. Table 5.1 lists the demand of the RH auxiliary pump in phase 3 at 75%. Due to the degraded pump, the fuel rig system will be unable to meet this demand, when it is requested. The effect of the failure mode should, therefore, be seen in the RH auxiliary tank level and RH wing tank level. Figure 5.46 shows the LH and RH auxiliary tank levels over the course of the mission.

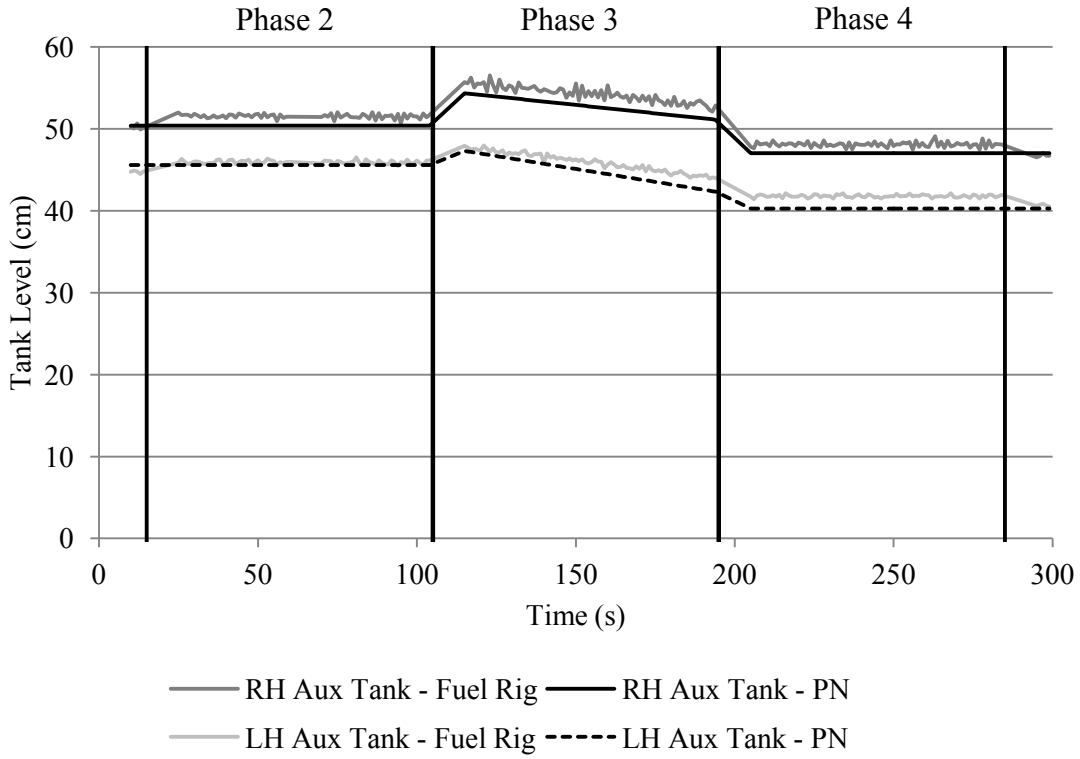

Figure 5.46: RH auxiliary pump degraded 50% - Auxiliary tank levels

On close inspection of Figure 5.46, it can be seen that in phase 3 the RH auxiliary tank level falls at a lower rate than LH auxiliary tank level. This is true when considering both the tank level values recorded from the fuel rig and those predicted by the PN model. This is due to the failure mode, which limits the operation of the RH auxiliary pump to

50% while the LH auxiliary pump operates at 75%. The effect of this fault should also be, therefore, visible in the RH wing tank level. Figure 5.47 shows the RH wing tank level of the fuel rig during the mission.

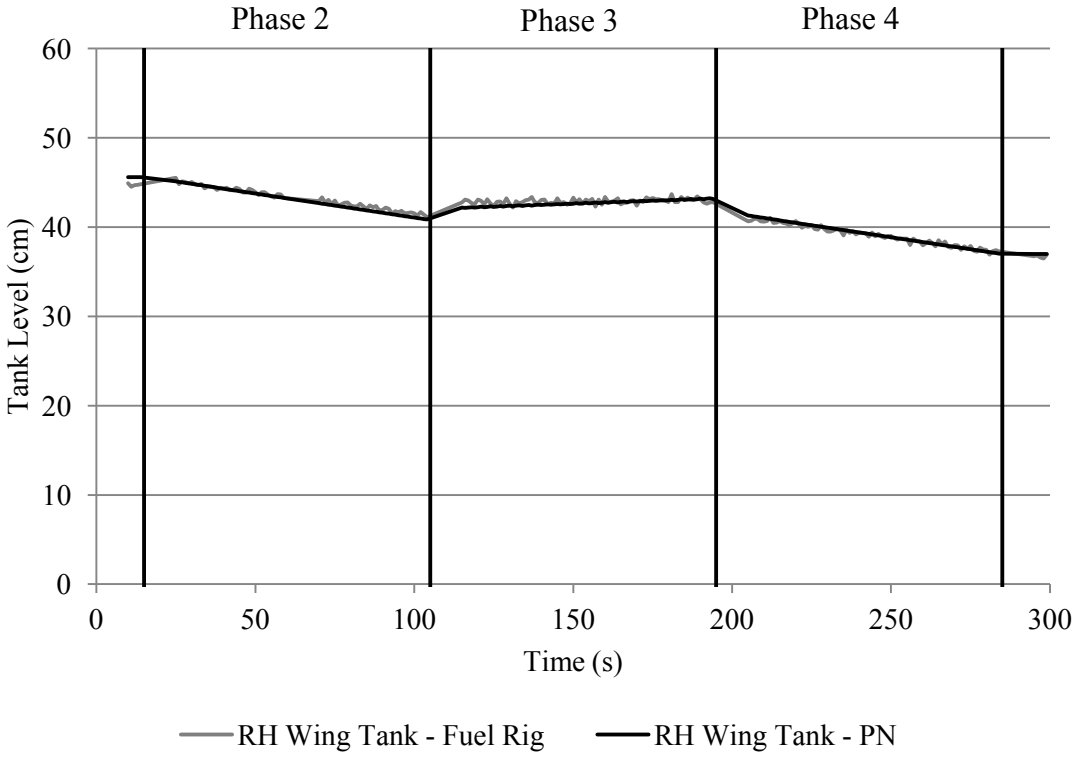

Figure 5.47: RH auxiliary pump degraded 50% - RH wing tank level

Comparing the behaviour of the RH wing tank level variable in Figure 5.47 to that in Figure 5.2 ,when there was no fault present in the system, the effect of the fault is clear. In Figure 5.47 the RH wing tank level shows only a small increase in phase 3. In Figure 5.2 the increase in tank level is much larger, approximately 5cm.

When the auxiliary pump is degraded to 50% capacity, it should only have the ability to provide the RH wing tank with enough fuel to replace that being removed by the RH engine pump, which is also operating at 50% demand. Figure 5.47 shows a slight increase in the RH wing tank level, which shows that a greater amount of fuel enters the tank than leaves it in phase 3. The reasons for this small variation in component performance have been discussed in Section 5.4.2.1.

All of the remaining system variables have not been affected by the auxiliary pump fault. Graphs of their behaviour over the course of the mission are therefore similar to those in Section 5.3. The SD values for a number of system variables are listed in Table

5.22. It can be seen that when diagnosed correctly, all of the SD values are within the tolerance limit. The fault verification technique would therefore be able to verify the arising correctly.

Table 5.22: SD of fuel rig variables - RH auxiliary pump degraded 50%

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.621cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.507cm	0.929cm
LH Wing Tank Level	1.500cm	0.541cm	0.437cm
RH Wing Tank Level	1.500cm	0.353cm	1.044cm
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.19L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	901Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,725Pa	1,688Pa

When the failure mode is falsely diagnosed, none of the SD values exceed their tolerances and therefore the fault verification technique cannot filter the arising as false. It would be expected that the RH auxiliary tank level and RH win tank level variables would have SD values that exceed their tolerances. However the effect of the fault on these variables does not change their behaviour enough to cause the SD tolerance to be exceeded. Figure 5.48 shows the RH wing tank level, when the auxiliary pump degraded 50% fault is falsely diagnosed. It can be seen, that without the fault present on the fuel rig, the wing tank level increases as expected in phase 3 by approximately 5cm. The PN tank level, by comparison, only increases by 1 – 2cm. However, the difference between recorded tank levels and those predicted by the PN model with the fault present are not large enough to create a SD which exceed the tank level tolerance. Without reducing the tank level tolerance it would not, therefore, be possible to filter this auxiliary pump fault were it falsely diagnosed.

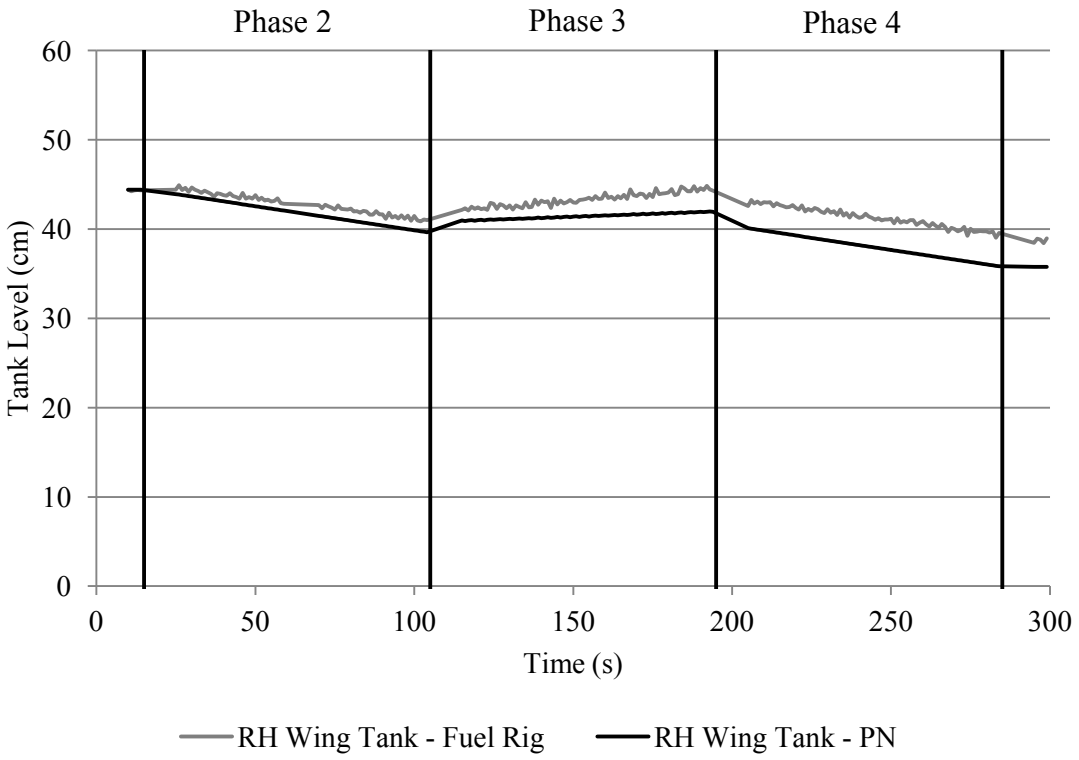

Figure 5.48: RH auxiliary pump degraded 50% false arising - RH wing tank level

5.4.8.4 RH Engine Pump Degraded 50%

Should the engine pump become degraded by 50%, it will only be able to operate at a maximum of 50% of its capability. Any occurrence of this fault in the phased mission being undertaken will not have an effect on the fuel rig variables, as the engine pumps are only operated at a maximum of 50%. Table 5.1 shows that the engine pumps are operated at a demand of 50% in phases 2, 3 and 4, while they are off in phases 1 and 5. The failure mode 'engine pump degraded' will therefore be hidden, should it occur.

As the failure mode is hidden, the graphs in Section 5.3 display the behaviour that would be expected from the fuel rig variables. Table 5.23 contains the SD values calculated from the data of the fuel rig variables, when the fault was injected into the fuel rig and modelled in the PN. In this scenario all of the SD values are within the tolerance limits, which means the arising would be verified. The table also lists the SD values calculated when the engine pump degraded fault was falsely diagnosed. In this scenario the failure mode is present in the PN model results but not in the fuel rig variable outputs. However, as the fault is hidden the behaviour of the PN predicted behaviour is the same as that recorded and therefore none of the SD values exceed the respective tolerances. As a result,

a false arising specifying that the engine pump had degraded up to 50% would go unfiltered in the phased mission under consideration.

Table 5.23: SD of fuel rig variables - RH engine pump degraded 50%

		Standard Deviation	
		Fault Diagnosed	Fault Diagnosed
Fuel Rig Variable	Tolerance	Correctly	Incorrectly
LH Auxiliary Tank Level	1.500cm	0.612cm	0.633cm
RH Auxiliary Tank Level	1.500cm	0.430cm	0.480cm
LH Wing Tank Level	1.500cm	0.458cm	0.437cm
RH Wing Tank Level	1.500cm	0.514cm	0.506cm
LH Flow Rate	0.30L/min	0.10L/min	0.10L/min
RH Flow Rate	0.30L/min	0.20L/min	0.20L/min
LH Fuel Flow Pressure	9,000Pa	865Pa	872Pa
RH Fuel Flow Pressure	9,000Pa	1,957Pa	1,688Pa

5.4.9 Tank Leak Failure Modes

Tank leak failures are unique in that the extent to which they impact on the system depends on the location and size of the leak. A large leak in the base of a tank, for example, will have a greater impact on a system than a small leak near the top of a tank. Two types of tank leaks have been considered on the fuel rig; a smaller leak in the base of the RH wing tank and a larger leak in the side of the LH auxiliary tank. The verification technique that will be applied to these tank leaks is detailed in Section 4.10.2.

5.4.9.1 RH Wing Tank Base Leak

In order to apply the leak fault verification technique, to the RH wing tank leak it was necessary to place an additional flow rate sensor between the RH auxiliary tank and the RH wing tank. This meter will provide a measure of the flow rate into the RH wing tank, while the flow rate meter next to the RH engine pump measures the flow rate out of the wing tank.

Applying the process described in Section 4.10.2, the first stage in the verification

process is to find the tank levels during the mission from the flow rate outputs. Figure 5.49 shows the tank level recorded from the level sensor and as determined from the flow rate sensors.

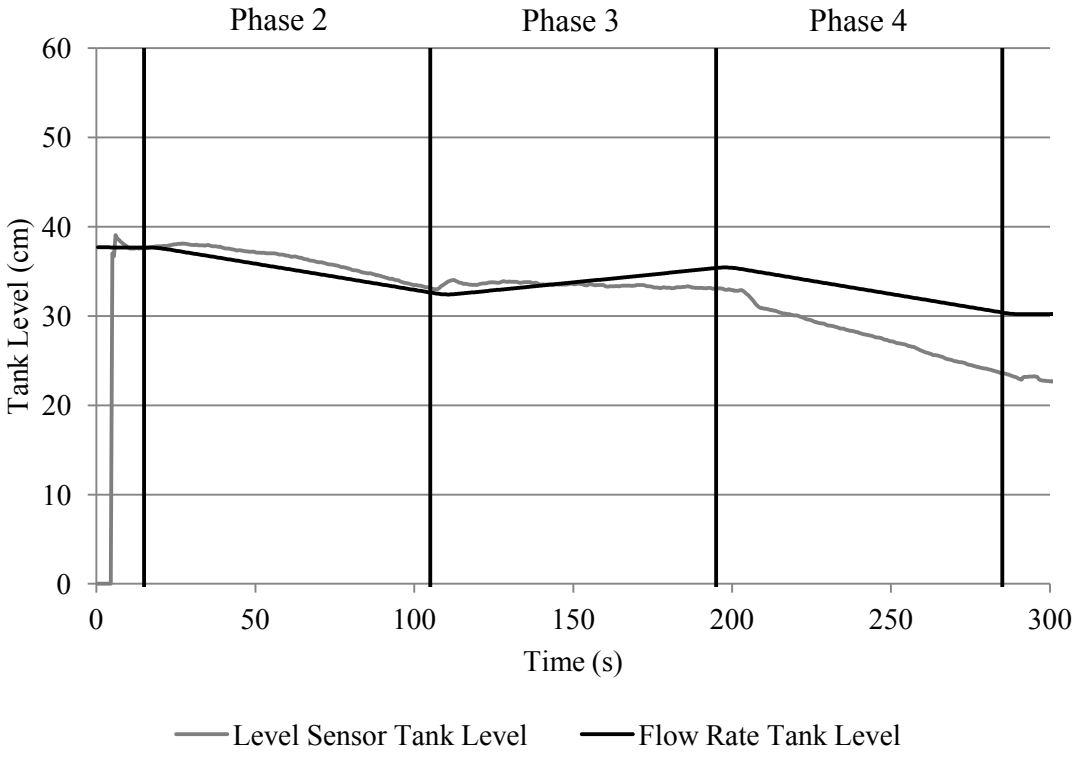

Figure 5.49: RH wing tank base leak - RH wing tank level

Figure 5.49 shows that the level sensor and flow rate tank level variables have different behaviours throughout the mission. This is due to the leak effects, which are only displayed in the level sensor curve. When the leak occurs, i.e. after 60 seconds, the gradient of the tank levels recorded by the level sensor decreases. By comparison the flow rate determined tank level stays constant throughout phase 2. The gradient of the tank levels recorded by the level sensor are also lower in phase 4 compared to the tank levels found from the flow rates. In phase 3, the flow into the wing tank creates a positive gradient in the flow rate curve but this is not seen in the level sensor curve because of the leak effects.

Knowing the time at which the leak occurred from the arising details, the gradient of each tank level curve prior to and after the leak occurring is identified. Table 5.24 lists these gradients and the resultant gradient residual values prior to and after the fault.

The gradient residual prior to the arising time is 0.0104cm/sec. The gradient residual after the arising is -0.0193cm/sec. There is therefore a decrease in the gradient residual

Table 5.24: RH wing tank leak - Tank level curve gradients

	Prior to Arising	After Arising
Level Sensor Tank Level Curve Gradient	-0.0479cm/sec	-0.0783cm/sec
Flow Rate Sensor Tank Level Curve Gradient	-0.0583cm/sec	-0.0590cm/sec
Gradient Residual	0.0104cm/sec	-0.0193cm/sec

of 0.0297cm/sec. It was stated in Secion 4.10.2 that a decrease in the gradient residual of at least 0.0190cm/sec was necessary in order to verify a leak. As this value has been exceeded a leak can be confirmed.

The final stage in the analysis of leak faults is to determine the location of the leak. The verification technique considers the fuel rig data at 15 second timesteps in order to accurately determine the location of the leak. Applying the technique as described in Section 4.10.2 the gradient residual over one of the 15 second timesteps will have to be greater than -0.0086cm/sec in order to confirm the leak height (0.0104-0.0190 = -0.0086). Table 5.25 lists the gradient residuals determined at every valid timestep in the mission from the time of the leak appearing.

Given that none of the interval residual values was greater than -0.0086cm/sec prior to the end of the mission, it was not possible to determine the exact height at which the leak occurred. However, knowing the tank level, from the wing tank level sensor, at the end of the mission it is possible to state the leak is present between the base of the tank and a height of 22.7cm.

5.4.9.2 LH Auxiliary Tank Side Leak

The second leak failure mode considered is that of a leak in the side of the LH auxiliary tank. In order to assess this fault it was necessary to place a flow rate meter at the exit from the LH auxiliary tank in order to monitor the flow out of this tank and into the LH wing tank. As there is no flow into the auxiliary tanks during the phased mission, there is no need for a flow rate sensor at the input to the auxiliary tank. The leak verification technique is applied in the same way as it was in Section 5.4.9.1.

Figure 5.50 shows the LH auxiliary tank levels during the mission as determined from the level sensor and flow rate outputs. The auxiliary pumps only experience a demand

Table 5.25: Residual values after arising - RH wing tank leak

Interval (sec)	Interval Residual (cm/sec)
35 – 60	0.0104
80 – 90	-0.0193
90 – 105	-0.2623
105 – 120	Phase Change
120 – 135	-0.0173
135 – 150	-0.0489
150 – 165	-0.0502
165 – 180	-0.0470
180 – 195	-0.0446
195 – 210	Phase Change
210 – 225	-0.0281
225 – 240	-0.0347
240 – 255	-0.0323
255 – 270	-0.0613
270 – 285	-0.0295
285 – 300	Phase Change

during phase 3 of the mission. There should therefore only be a change in the auxiliary tank levels during phase 3. Figure 5.50 shows, however, that the leak has caused the tank levels recorded by the level sensor to fall during phase 2. The level sensor curve also falls in phase 3 and is relatively flat in phase 4, both as expected. Looking at the phase 3 section of the curve in more detail it can be seen there is a change in the gradient of the tank levels recorded from the level sensor at the midpoint of the phase, approximately 150 seconds. This is due to the tank level falling below the height of the leak. The tank levels determined from the flow rate outputs exhibit the expected auxiliary tank level behaviour as it only reflects the flow of fuel out of the tank through the pipe system when a demand is present.

The gradient of each tank level curve prior to and after the arising is shown in Table 5.26. The gradient residuals are also shown. Table 5.26 shows that the gradient residual decreased by 0.1738cm/sec as a result of the leak. The presence of a leak in the tank was

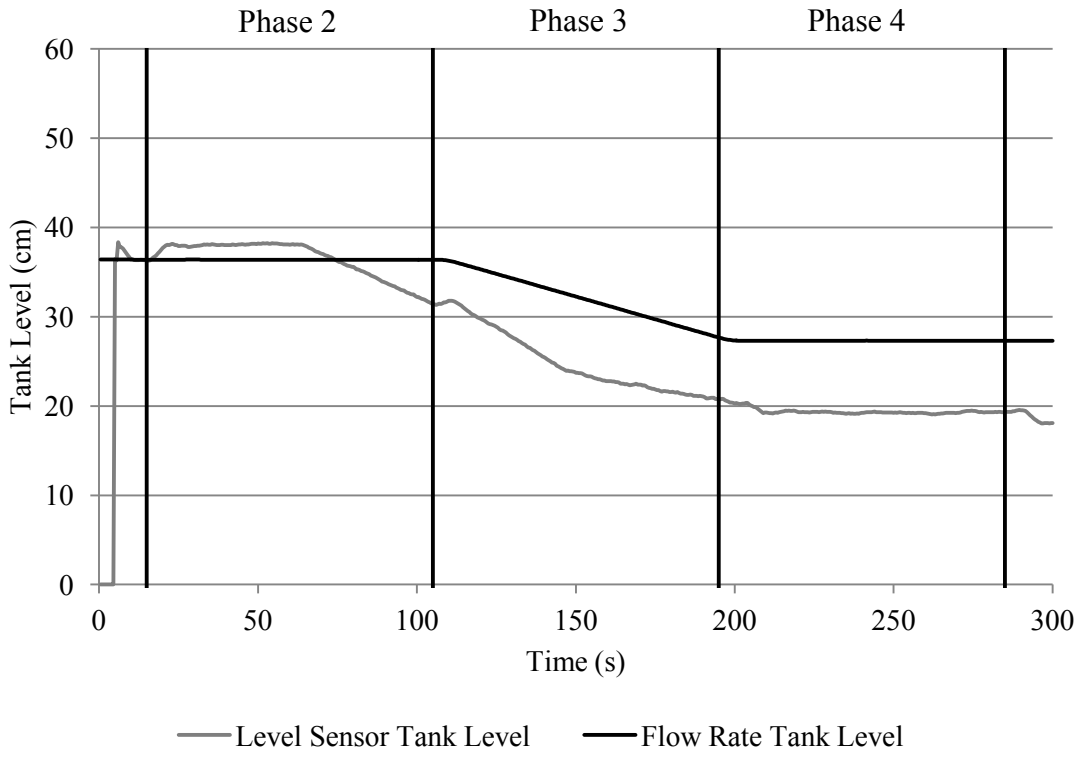

Figure 5.50: LH auxiliary tank side leak - LH auxiliary tank level

Table 5.26: RH wing tank leak - Tank level curve gradients

	Prior to Arising	After Arising
Level Sensor Tank Level Curve Gradient	-0.0015cm/sec	-0.1752cm/sec
Flow Rate Sensor Tank Level Curve Gradient	-0.0002cm/sec	-0.001cm/sec
Gradient Residual	-0.0013cm/sec	-0.1751cm/sec

therefore verified (0.1738 ¿ 0.0190). Table 5.27 lists the interval gradient residuals found at 15 second timesteps after the arising time. In order to confirm that the tank level has fallen below the height of the leak, an interval residual of greater than -0.0203cm/sec is required (-0.0013-0.0190 = -0.0203).

Table 5.27 shows that between 150 and 165 seconds the gradient residual was 0.0190cm/sec. This residual is greater than -0.0203cm/sec and therefore indicates that the tank level had fallen below the leak height. Using the level sensor outputs the leak height was determined to be 23.6cm. Looking again at Figure 5.50 it can be seen that the time interval and leak height determined match well with the point at which the level

Table 5.27: Residual values after arising - LH auxiliary tank leak

Interval (sec)	Interval Residual (cm/sec)
35 – 60	-0.0013
80 – 90	-0.1751
90 – 105	-0.1592
105 – 120	Phase Change
120 – 135	-0.1127
135 – 150	-0.0865
150 – 165	0.0190

sensor tank level gradient changed noticeably.

5.5 Identifying Genuine Faults Among Multiple Arisings

Due to the number of arisings that can be generated over a short period of time by complex systems, it is possible that several arisings will be generated at the same time. This section will consider the scenario, where four arisings are generated at the same time but only one is genuine, i.e. the other three arisings are false. The variable comparison aspect of the fault verification technique applied is consistent with that applied throughout this chapter.

Using the five phase mission, the fault 'RH Flow Pressure Sensor Failed Off' has been induced in the fuel rig after 60 seconds. This failure mode has been included in the health log input file along with the 'RH Flow Pressure Sensor Failed Stuck', 'RH Wing Tank Isolation Valve Blocked' and 'RH Flow Sensor Failed Stuck' failure modes which are all false. Including two failure modes of the same component will ensure the fault verification software can distinguish between them. The effect of the genuine fault should be seen in the wing tank level, fuel flow rate and flow pressure variables on the RHS of the system, as described in Section 5.4.5.2. Figure 5.51 shows the RH wing tank level behaviour predicted by the PN model when each of the failure modes are included individually. The fuel rig data plotted on the graph represents that with the 'RH Flow Pressure Sensor Failed Off' failure mode present. Figures 5.52 and 5.53 show equivalent graphs of the RH fuel flow rate and flow pressure.

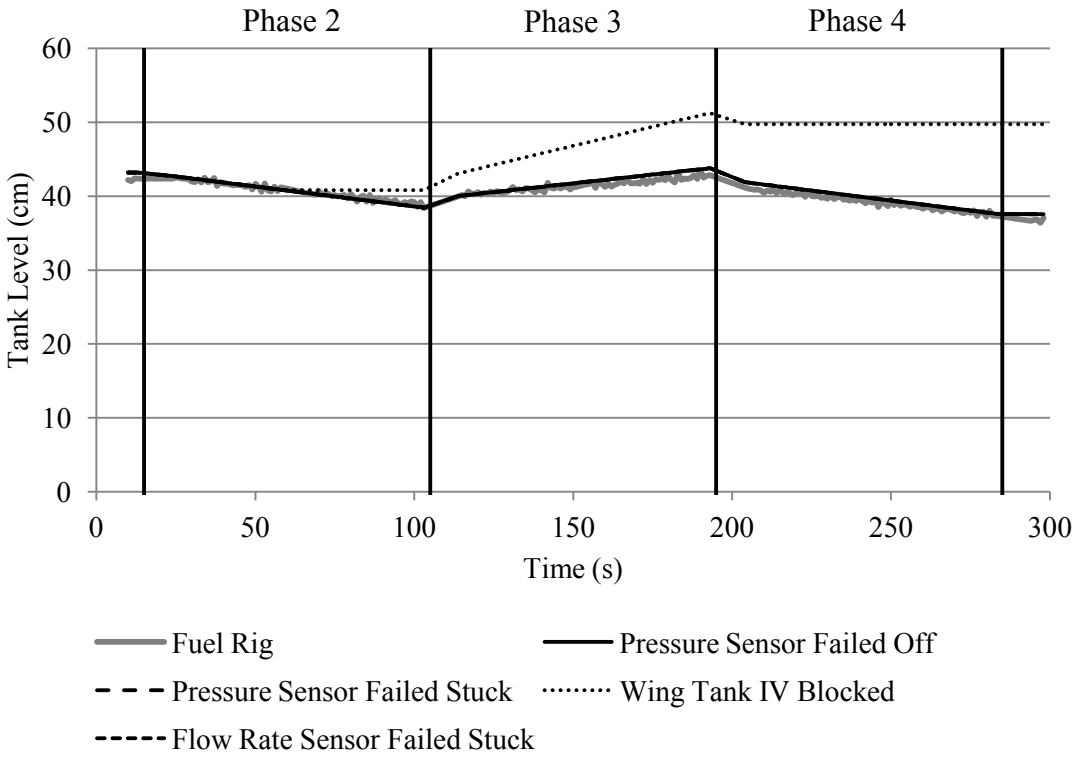

Figure 5.51: Multiple concurrent arisings - RH wing tank level

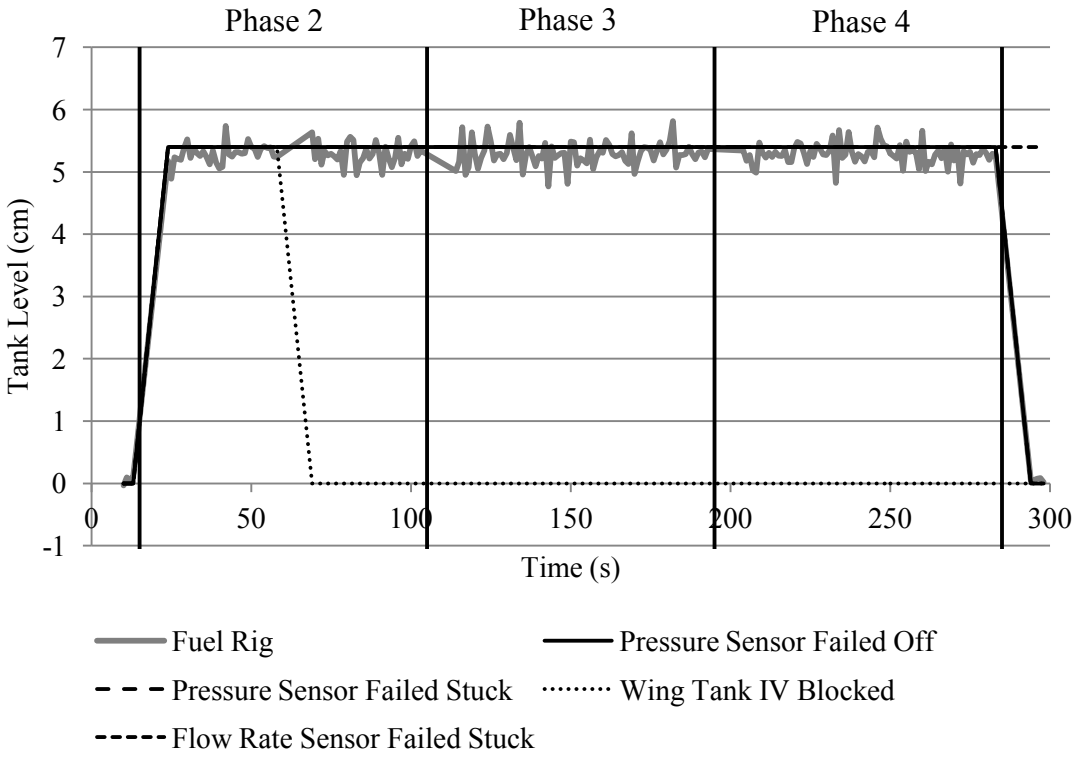

Figure 5.52: Multiple concurrent arisings - RH fuel flow rate

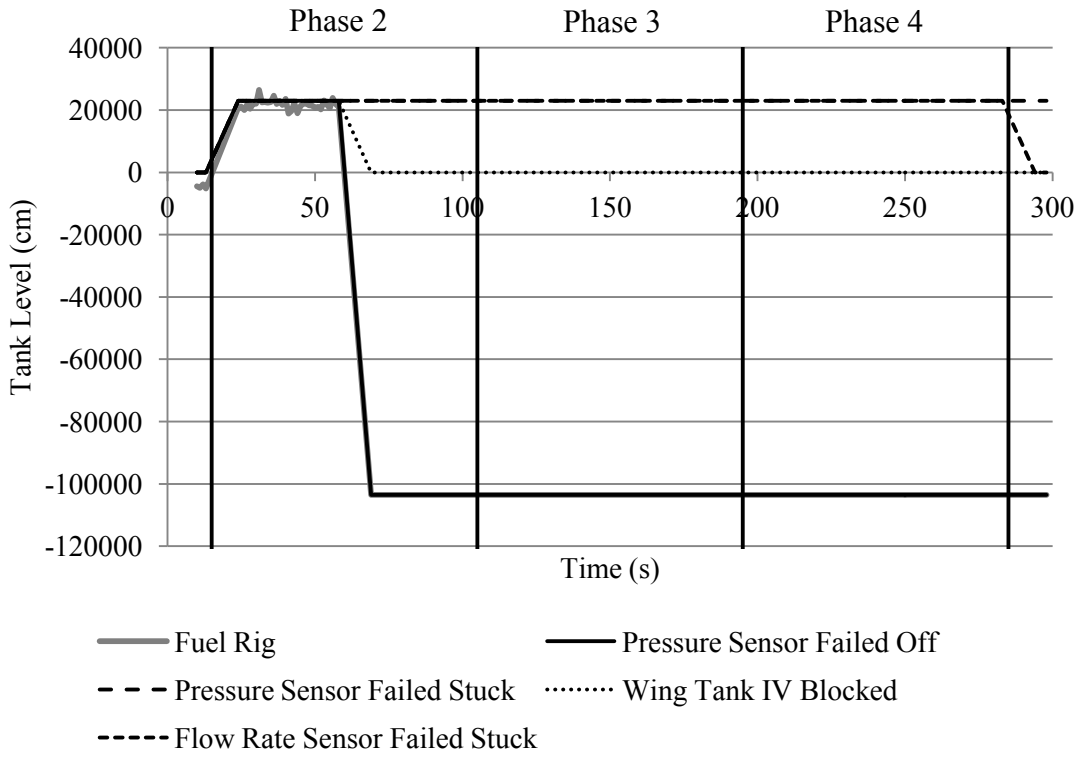

Figure 5.53: Multiple concurrent arisings - RH flow pressure

Figure 5.51 shows that the PN model predicts only one of the failure modes will produce a behaviour in the RH wing tank level variable, that significantly differs from that recorded from the fuel rig. The failure mode is 'Wing Tank Isolation Valve Blocked'. This is due to the fact that this is the only failure mode, which prevents the flow of fuel out of the RH wing tank. As a result, this is the only failure mode that causes the SD of the RH wing tank level variable to exceed the tank level tolerance. The PN predicted behaviour of the remaining failure modes produced similar behaviour to that recorded from the fuel rig. As a result the respective RH wing tank level SD values were also the same and within the SD limit.

Figure 5.52 shows that there was a deviation from the recorded fuel flow rate, when two of the failure modes were modelled in the fuel rig PN. The wing tank IV blockage prevents any flow from reaching the fuel flow rate sensor and, therefore, causes the fuel flow rate behaviour to differ. The flow sensor stuck failure mode also caused the behaviour of the fuel flow rate variable to deviate from that recorded in phase 5. The SD values determined for these failure modes both exceeded the fuel flow rate SD tolerance.

Figure 5.53 shows the flow pressure as recorded from the fuel rig and as predicted by

the PN models with the individual failure modes included. It can be seen that only one of the failure modes has an effect on the predicted RH flow pressure behaviour that results in it matching that recorded from the fuel rig. That is the failure mode 'Pressure Sensor Failed Off'. The three other failure modes produced a flow pressure curve that differs significantly from the recorded flow pressure curve. The SD of these three results also exceed the tolerance for the flow pressure variable by far

Table 5.28 lists the SD values for a range of system variables when each of the failure modes were considered by the PN software one by one.

Table 5.28: SD of fuel rig variables - Genuine fault among multiple arisings

	Failure Mode			
	Pressure Sensor	Pressure Sensor	Wing Tank	Flow Sensor
Fuel Rig Variable	Failed Off	Stuck	IV Blocked	Stuck
LH Aux Tank Level	0.494cm	0.494cm	0.494cm	0.494cm
RH Aux Tank Level	0.733cm	0.733cm	0.733cm	0.733cm
LH Wing Tank Level	0.507cm	0.507cm	0.507cm	0.507cm
RH Wing Tank Level	0.450cm	0.450cm	**4.247cm**	0.450cm
LH Flow Rate	0.11L/min	0.11L/min	0.11L/min	0.11L/min
RH Flow Rate	0.18L/min	0.18L/min	**2.07L/min**	**0.77L/min**
LH Flow Pressure	1,350Pa	1,350Pa	1,350Pa	1,350Pa
RH Flow Pressure	908Pa	**45,693Pa**	**37,283Pa**	**45,602Pa**

From Figure 5.53 alone it is possible to identify not only which failure modes are associated with false arisings but also which of the failure modes is genuine. Only one of the PN curves in this figures is similar to fuel rig curve, the pressure sensor failed off curve. The behaviour of the three other curves is different from 60 seconds onwards. Table 5.28 also provides numerical evidence that only the failure mode 'RH Flow Pressure Sensor Failed Off' produced SD values for all of the system variables that were within the tolerances. The fault verification process is, therefore, capable of identifying a genuine arising amongst several false ones. Given the large number of arisings that are generated by complex systems, the capability that has been demonstrated in this section has the potential to reduce the unnecessary use and wastage of valuable resources.

5.6 Second Order Failure Modes

Second order failure modes refer to situations where two faults are present on a system at one time. A number of second order failure modes have been considered on the fuel rig. Scenarios where both faults occur at the same time and scenarios where the faults occur at different times are considered. A second order failure mode that includes a leak is also evaluated.

When second order failure modes are considered by the fault verification technique, it first considers each fault individually and then together. The faults are considered individually first because it is possible that one or both of the faults may be false. All possible scenarios are therefore considered. The phased mission undertaken throughout this section is the same mission comprising 5 phases, as described in Section 5.2.

Three second order failure mode scenarios are presented below. They provide an insight as to the capability and limitation of the technique.

5.6.1 RH Fuel Flow Rate Sensor Failed High and RH Auxiliary Tank Low Level Switch Failed Off

In this scenario the failure modes are induced in the fuel rig at different times. After 60 second the RH fuel flow rate sensor is failed high. The RH auxiliary tank low level switch is then failed off after 135 seconds. It would be expected that the effects of these faults, as described in Section 5.4.4 and 5.4.7 respectively, would be seen simultaneously from 135 seconds onwards.

The fault verification technique first considers only the RH fuel flow rate sensor failing high in the PN model, i.e. it makes no account of the potential low level switch fault. This first order failure mode has already been considered and Figure 5.35 shows the resultant RH fuel flow rate graph. It can be seen in Table 5.3 that the low level switch fault does not affect the RH fuel flow rate and therefore the fuel rig curve will be the same when considering this second order fault as is shown in Figure 5.35. Figure 5.35 shows that the PN model has accurately predicted the behaviour of the fuel flow rate variable including the effect of the flow rate sensor fault which occurred after 60 seconds. As a result the SD of the RH fuel flow rate variable is 0.07L/min, which is within the tolerance for this variable.

Figure 5.54 shows the state of the RH auxiliary tank low level switch over the duration

of the same mission. As the PN model has not included the low level switch fault, it would be expected that the PN predicted low level switch behaviour will differ from that recorded from the fuel rig.

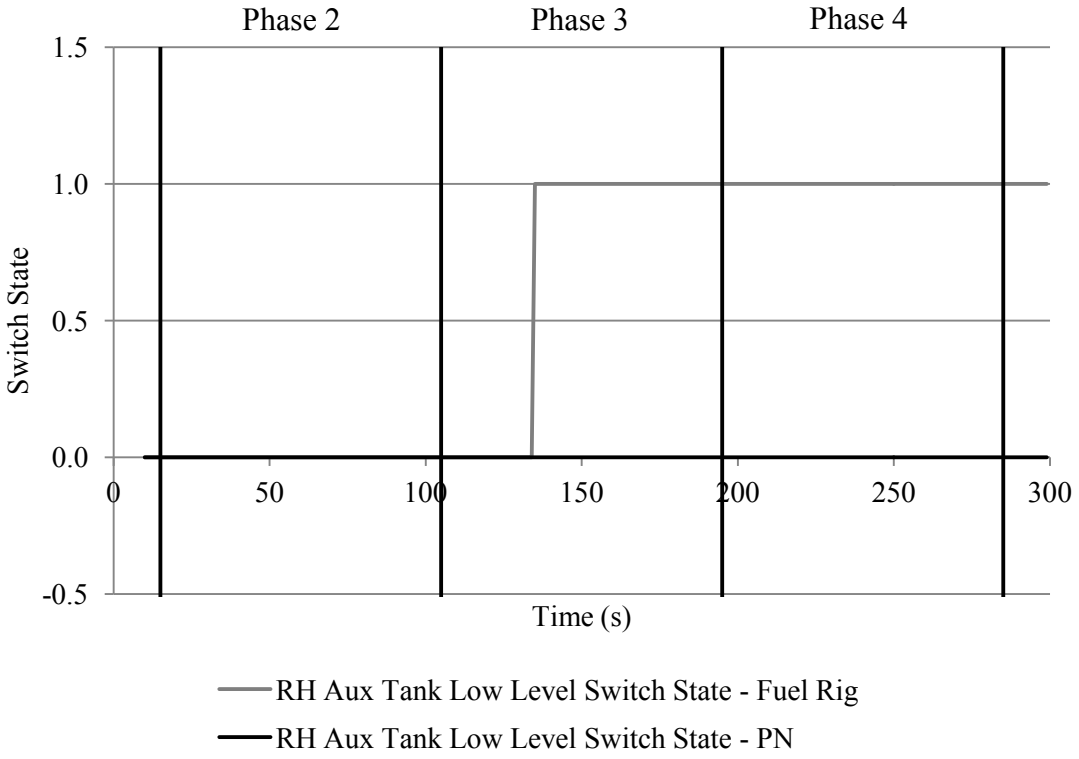

Figure 5.54: RH fuel flow rate failed high - RH auxiliary tank low level switch state

The effect of the auxiliary tank low level switch fault can be clearly seen on the fuel rig curve of Figure 5.54. The PN curve however does not indicate a change of state as the PN model did not include the low level switch fault. The SD of the RH auxiliary tank low level switch variable is therefore 0.49. This value exceeds the switch state tolerance of 0.1. All of the SD results of the remaining system variables were within the specified tolerances, when only the fuel flow rate sensor fault was considered.

Given that the SD of the RH auxiliary low level switch exceeded the tolerance limit, it can be concluded that either the fuel flow rate sensor has not failed or it is not the only component on the system that has failed. It can, however, be stated that the RH fuel flow rate sensor alone has not failed.

Having failed to successfully verify the first arising, the fault verification technique then considers the failure mode within the second arising – RH auxiliary tank low level switch failed off. The flow rate sensor fault is now ignored. Figure 5.55 and Figure 5.56

show the RH fuel flow rate and RH auxiliary tank low level switch state over the course of this mission.

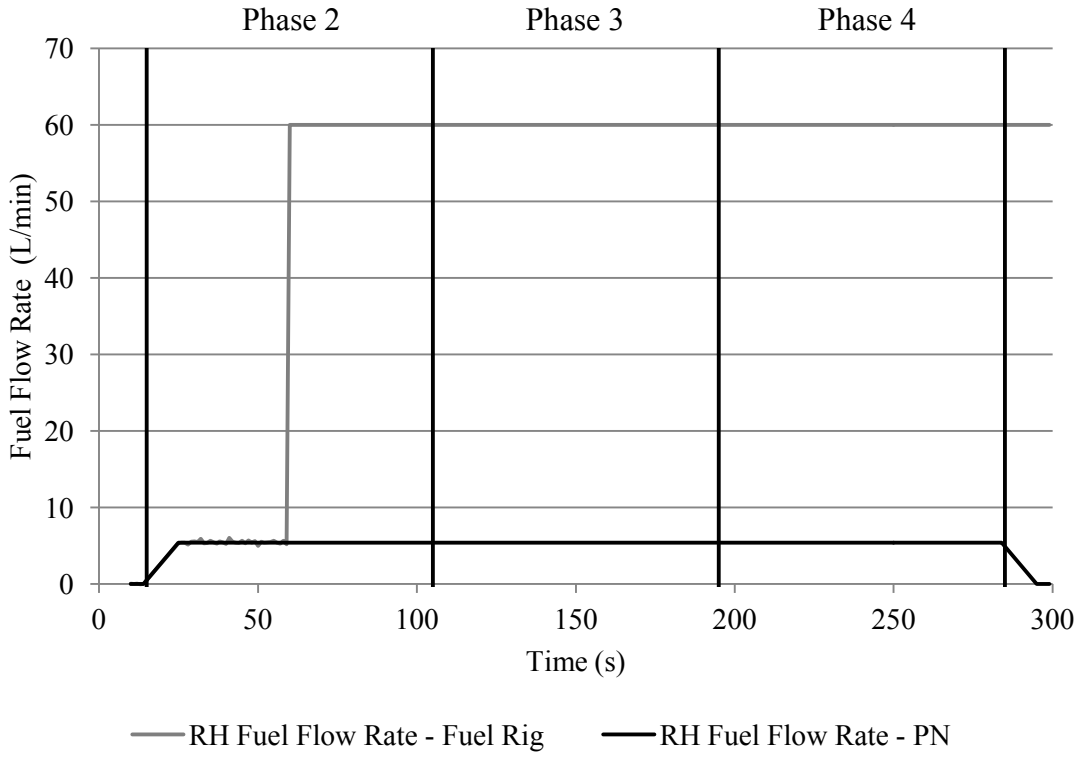

Figure 5.55: RH auxiliary tank low level switch failed off - RH fuel flow rate

Figure 5.55 shows that without consideration of the flow rate sensor fault, the PN fuel flow rate is significantly different from the fuel rig flow rate from 60 seconds onwards. The SD of the residual results in this figure is 20.39L/min. This SD result far exceeds the 0.3L/min tolerance of the fuel flow rate variable. In comparison to the fuel flow rate results, Figure 5.56 shows that the PN model has predicted the same behaviour in the RH auxiliary low level switch as has been recorded from the fuel rig. The SD of these results is therefore 0. However, as the SD of the fuel flow rate variable exceeded the permissible tolerance the second arising will also be filtered by the fault verification technique.

Having considered both of the arisings individually, they are now considered together. When both of the failure modes are included in the PN model, none of the SD tolerances are exceeded. The RH fuel flow rate behaviour, as taken from the fuel rig and PN data, produces the same graph as shown in Figure 5.35. The RH auxiliary tank low level switch behaviour is the same as that in Figure 5.56. The behaviour of all of all the remaining system variables is unaffected and the outputs of these variables are shown in Section 5.3.

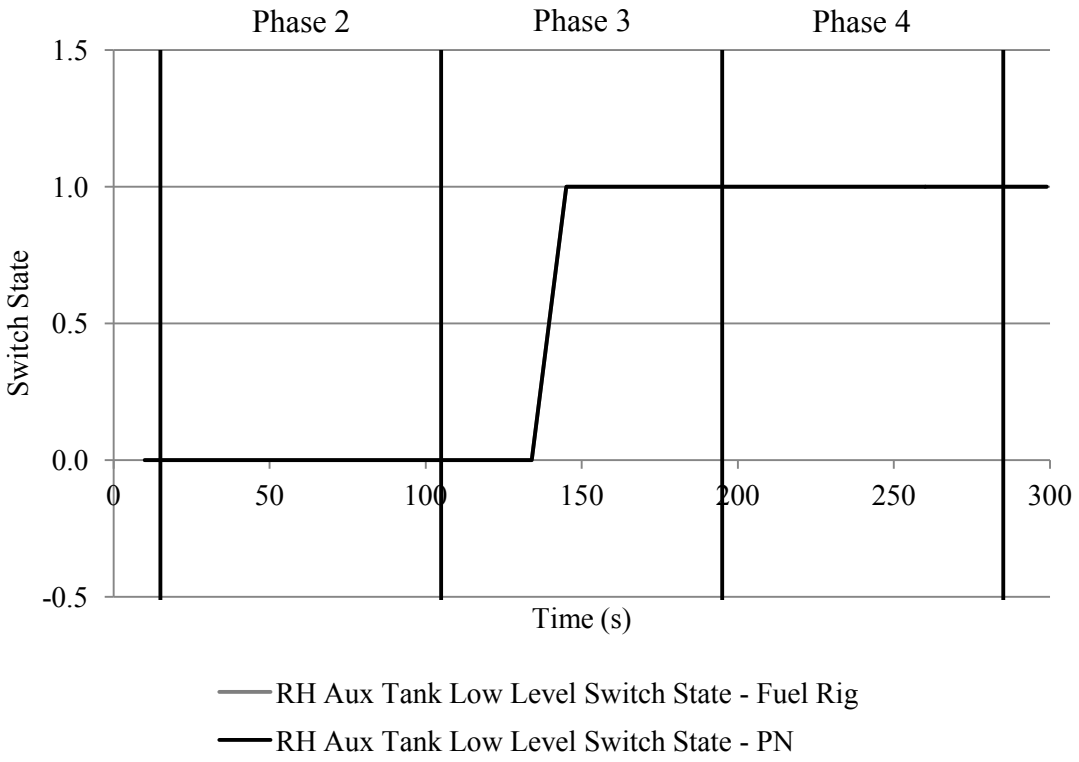

Figure 5.56: RH auxiliary tank low level switch failed off - RH auxiliary tank low level switch state

Table 5.29 shows the SD values determined for every variable on the system when the arising failure modes are considered independently and together.

The results of Table 5.29 confirm the fact that when considered individually, the failure modes are identified as false faults. The fault verification technique can determine this from the SD values in the first two columns, which exceed the tolerance limits. It is only when the faults are considered concurrently that all of the SD values fall within the tolerances. The fault verification technique has therefore correctly identified that a second order fault is present in the fuel rig system.

There is a small variation in some of the SD values for the same variables in Table 5.29, even though the same fuel rig data was considered each time. This is due to the fact that in each of the analysis considered above faults were injected into the PN model at different times. Immediately following any fault the initial data is ignored to allow the fuel rig system variables to settle. As a result different data points have been ignored in each analysis which has caused a small variation in the SD values.

Table 5.29: SD of fuel rig variables - Fuel flow rate sensor failed high and auxiliary tank low level switch failed off

	Failure Mode		
	Flow Rate Sensor Failed High	Aux Tank Low Level Sw Stuck	Both Failure Modes
LH Auxiliary Tank Level	0.492cm	0.492cm	0.497cm
RH Auxiliary Tank Level	0.616cm	0.610cm	0.619cm
LH Wing Tank Level	0.506cm	0.517cm	0.512cm
RH Wing Tank Level	0.510cm	0.530cm	0.518cm
RH Auxiliary Flow Rate	0.14L/min	0.13L/min	0.13L/min
RH Flow Rate	0.08L/min	**20.39L/min**	0.08L/min
LH Flow Pressure	995Pa	994Pa	1,008Pa
RH Flow Pressure	1,681Pa	1,679Pa	1,704Pa
RH Wing Tank High Level Sw	0.0	0.0	0.0
RH Wing Tank Low Level Sw	0.0	0.0	0.0
RH Aux Tank High Level Sw	0.0	0.0	0.0
RH Aux Tank Low Level Sw	**0.49**	0.0	0.0

5.6.2 RH Wing Tank Base Leak and RH Fuel Flow Rate Sensor Failed Off

It is possible for the fault verification technique to consider a leak and another fault within a system. It is necessary to carry out the leak fault verification process in order to determine both the size and the location of the leak. These values can then be used as an input to the PN model, so the behaviour of the system can be most accurately modelled. In this mission the leak was injected into the RH wing tank after 60 seconds and the RH fuel flow rate sensor was failed off after 135 seconds.

The leak verification process is first considered. Figure 5.57 shows the RH wing tank level over the course of the mission.

Figure 5.57 shows very unique tank level curves. The tank level values determined by the flow rate outputs increase significantly from 135 seconds onwards. This is due to the flow rate sensor fault, which produces an output of -12.5L/min. This is interpreted

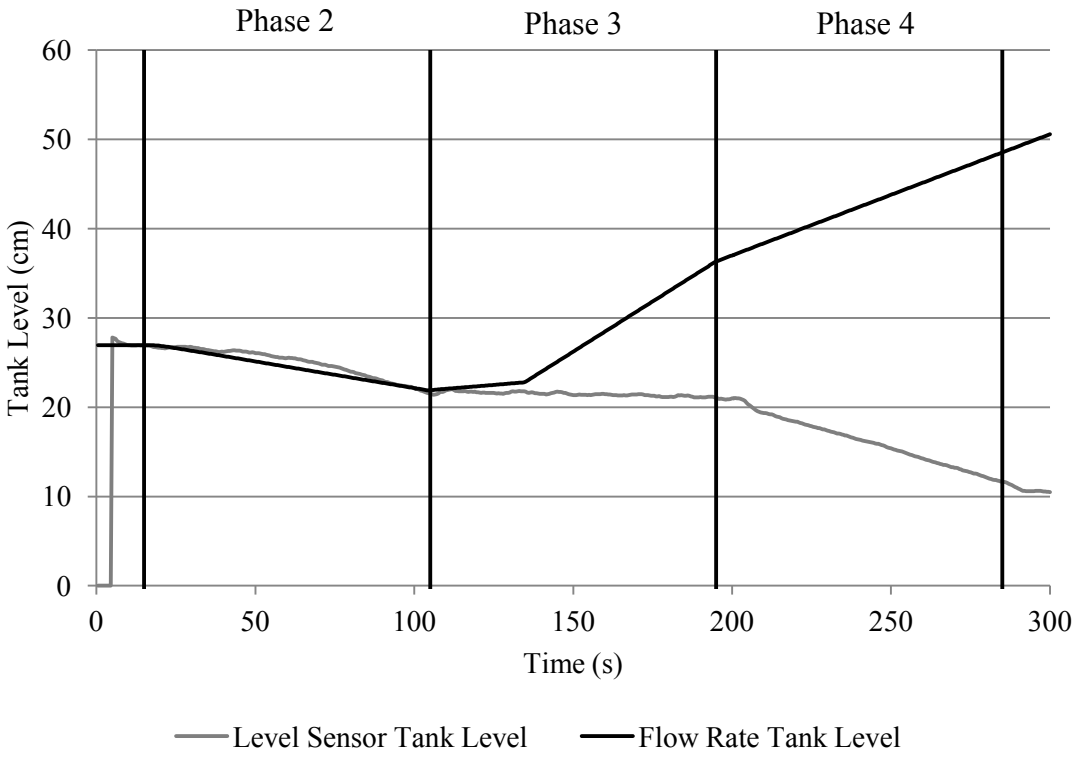

Figure 5.57: Wing tank leak and fuel flow rate sensor failed off - RH wing tank level

as reverse flow and therefore from 135 seconds onwards the wing tank is seen, incorrectly, to receive a large volume of fuel from the engine. By comparison, the tank level sensor curve is much closer to the behaviour expected in the presence of the wing tank leak fault. Applying the leak verification process, the tank level gradient residual changed from 0.0275cm/sec prior to the leak arising to -0.0470cm/sec after the arising. As this decrease in the tank level gradient residual is greater than 0.0190cm/sec, the presence of the leak is confirmed. The size of the leak is expressed in terms of its effect on the tank, in this case the leak size is 0.0745cm/sec.

In determining the location of the leak, gradient residuals were found at regular intervals after the arising time. None of these gradient residuals, however, indicated that the effect of the leak had disappeared. The location of the leak could therefore only be narrowed down to a height of between 0 and 10.5cm. When the leak fault is considered concurrently with the fuel flow rate sensor fault, the leak will be modelled in the base of the tank to consider the worst case scenario.

Having confirmed the presence of a leak, the fault verification technique next considers the fuel flow rate sensor fault individually. Figure 5.58 shows the RH wing tank level, as

recorded from the fuel rig and predicted by the PN model, when the wing tank leak is not considered.

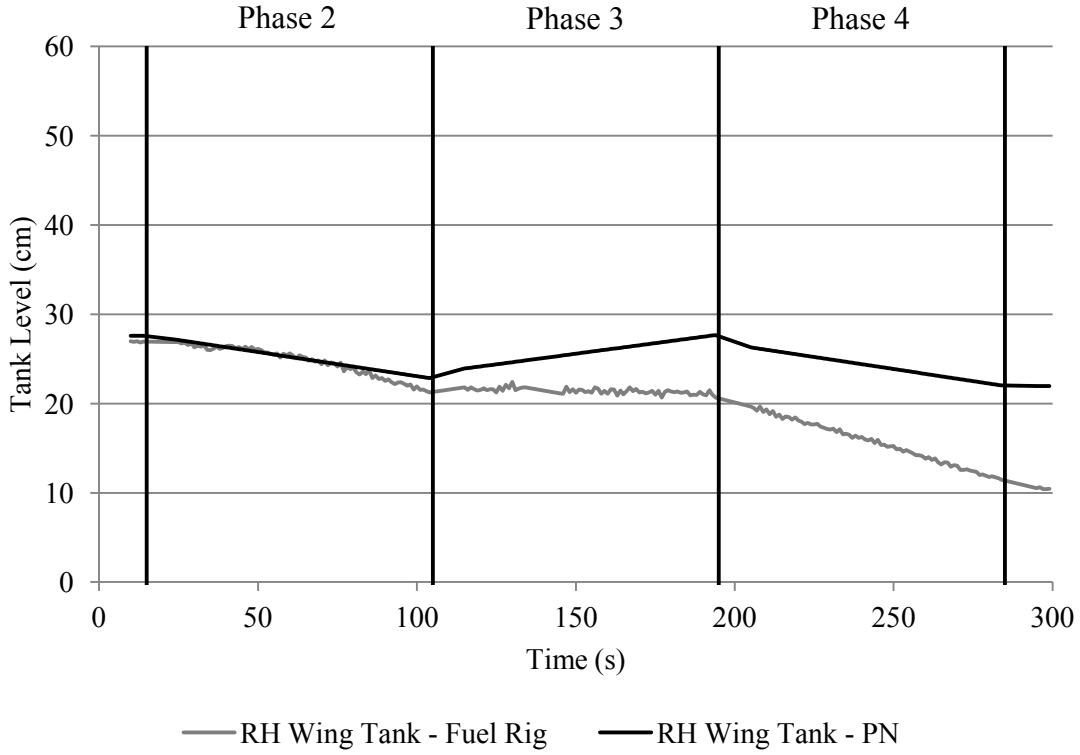

Figure 5.58: Wing tank leak and fuel flow rate sensor failed off - RH wing tank level

Figure 5.58 shows that without consideration of the wing tank leak the PN model has failed to accurately predict the behaviour of the RH wing tank level. As a result, the SD of the RH wing tank level variable is 3.709cm and exceeds the tolerance. The presence of the fuel flow rate fault alone in the system is therefore disproved.

The final step in the fault verification process is to consider the leak and fuel flow rate fault concurrently. Considering both of these faults in the same PN model, the predicted behaviour of the fuel rig system matches well with that recorded. Figure 5.59 shows the RH wing tank level variable when both faults are included in the PN model. It is only possible to achieve this set of results by including the leak fault in the PN model along with the flow rate sensor fault.

None of the SD tolerances have been exceeded and as a result the presence of both faults in the system is confirmed. The figures of the output variables as predicted by the PN model and recorded from the fuel rig when both failure modes are therefore similar to those displayed in Section 5.3. Table 5.30 lists the SD values of the fuel rig output

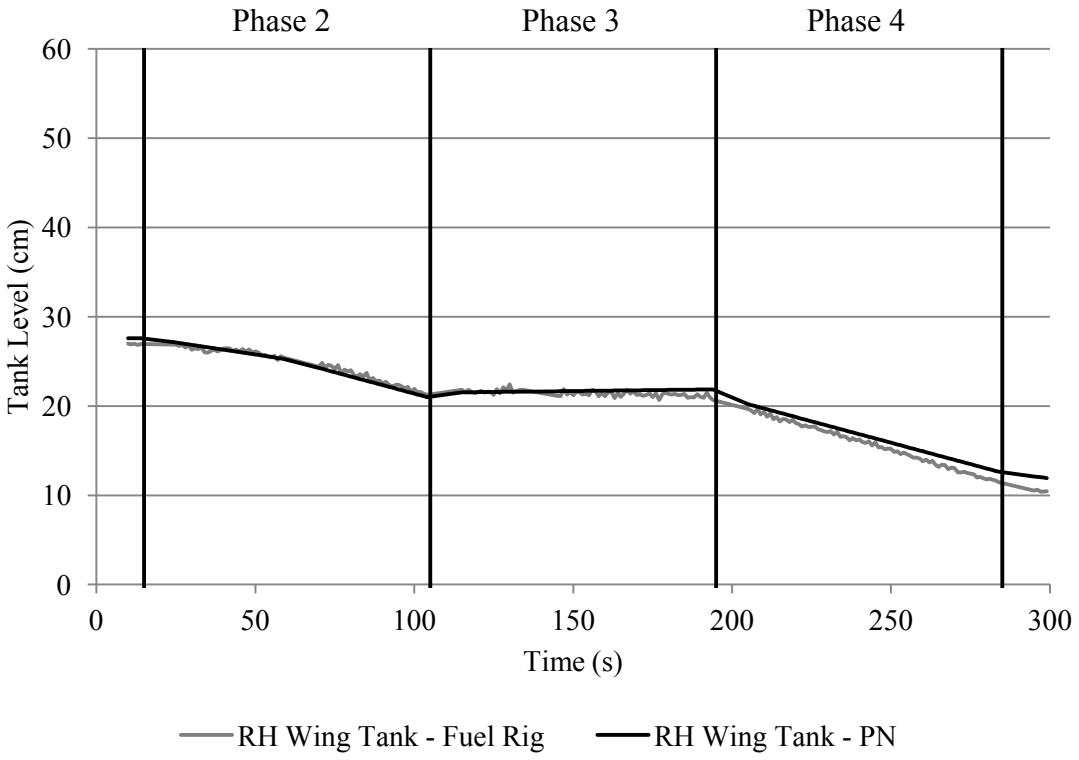

Figure 5.59: Wing tank leak and fuel flow rate sensor failed off - RH wing tank

variables, when only the fuel flow rate sensor fault is considered and when both faults are considered.

5.6.3 RH Flow Pressure Sensor Stuck and RH High Level Switch Failed On

In this final second order scenario, the RH flow pressure sensor is failed stuck after 60 seconds while the RH wing tank high level switch is failed on after 135 seconds. The fault verification technique first considers the flow pressure sensor fault then the high level switch fault and finally both faults together.

When the fault verification technique considers only the flow pressure sensor fault, the result of one system variable exceeds the SD tolerance limit. That variable is the RH wing tank high level switch. Figure 5.60 shows the predicted and recorded state of the RH wing tank high level switch during the mission. Figure 5.60 shows that the PN model has failed to represent the fault that occurs in the high level switch. As a result the SD of this variable is 0.49, which exceeds the tolerance limit of 0.1 for all the switch variables. Having exceeded this tolerance the fault verification technique considers the next arising.

Table 5.30: SD of fuel rig variables - Wing tank level sensor failed high and engine IV blocked/failed closed

	Failure Mode	
	Flow Rate Sensor Failed Off	Both Failure Modes
LH Auxiliary Tank Level	0.477cm	0.482cm
RH Auxiliary Tank Level	0.745cm	0.758cm
LH Wing Tank Level	0.585cm	0.595cm
RH Wing Tank Level	**3.709cm**	0.545cm
RH Auxiliary Flow Rate	0L/min	0L/min
RH Flow Rate	0.13L/min	0.13L/min
LH Flow Pressure	1,534Pa	1,562Pa
RH Flow Pressure	1,884Pa	1,894Pa

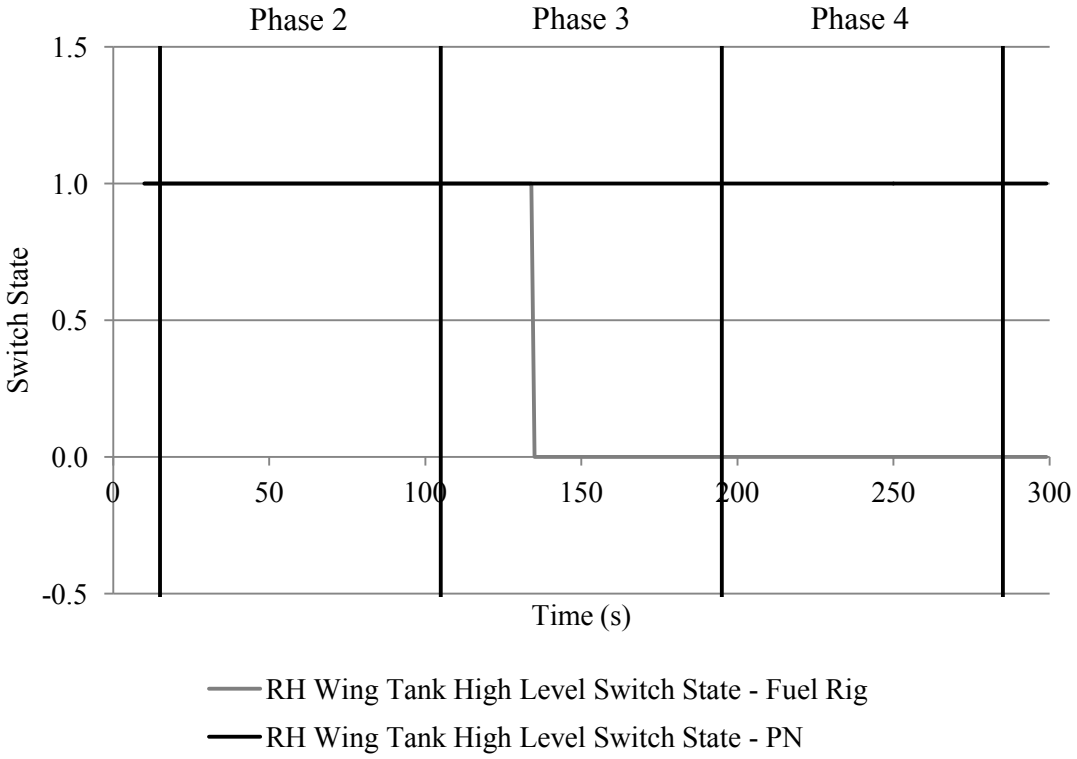

Figure 5.60: Flow pressure sensor stuck and high level switch failed high - RH wing tank high level switch

When considering only the high level switch failed on failure mode it would be expected that the PN model would fail to identify the flow pressure sensor fault. Figure 5.61 shows the flow pressure outputs in this scenario.

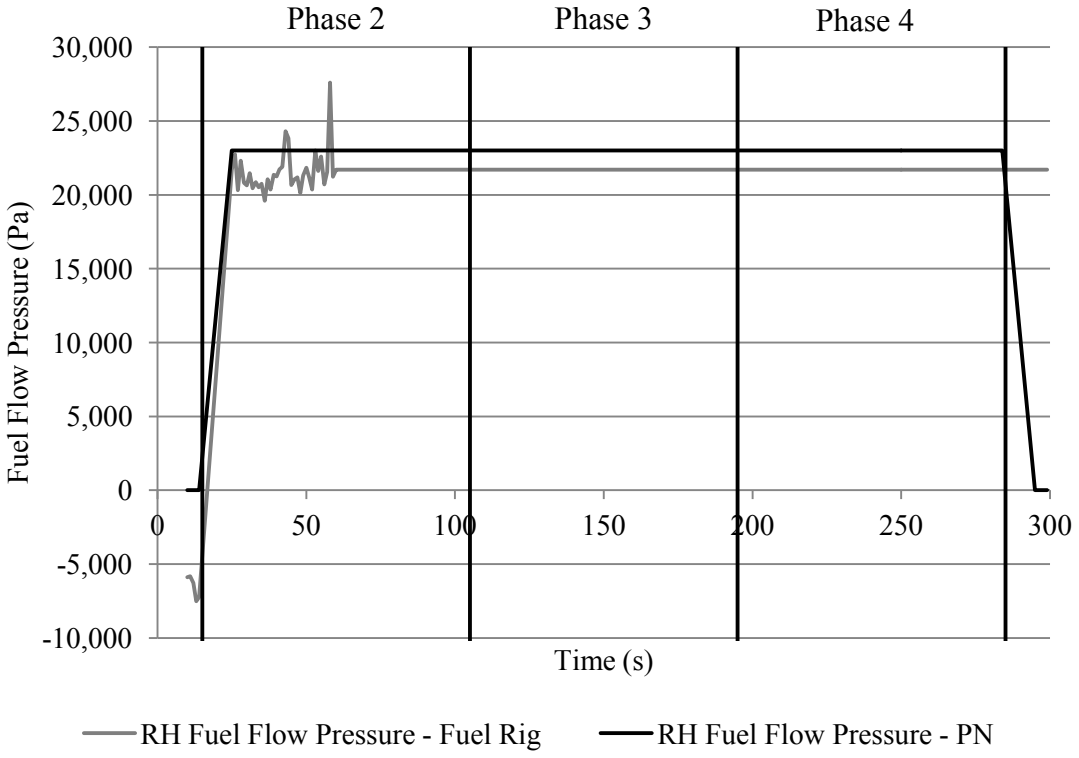

Figure 5.61: Flow pressure sensor stuck and high level switch failed high - RH flow pressure

Figure 5.61 shows that while the fuel rig and PN flow pressure curves are not identical during the mission, they are very similar over four of the five phases. There is a greater difference between the curves in phase 5, however, the duration of this phase is relatively short. Consequently the SD of these results, 3,434Pa, is within the tolerance limit for the flow pressure variable. The behaviour of the remaining variables on the system are not affected as a result of omitting the flow pressure sensor fault. Therefore, when considering the high level switch fault all of the SD results are within the specified tolerances. The presence of the fault 'RH Wing Tank High Level Switch Failed On' would, therefore, be incorrectly confirmed by the fault verification technique as the only fault present in the system.

As the process of evaluating second order faults has been initiated, the fault verification process will consider both faults concurrently, even though the high level switch fault was verified. When both faults are considered within a single PN model, the predicted system

outputs are very similar to those recorded from the fuel rig. All of the SD values are within the tolerances set for the respective variables. The greatest change, in comparison to the results that only considered the high level switch fault, occurs in the RH flow pressure variable. This change occurs as the PN model includes the RH pressure sensor fault, which means the PN predicted behaviour is more similar to that recorded in phase 5 of the mission. The SD values of the system variables, when each of the three failure mode combinations are considered by the PN model, are shown in Table 5.31

Table 5.31: SD of fuel rig variables - Fuel flow rate sensor failed high and auxiliary tank low level switch failed off

	Failure Mode		
	Flow Pres Sensor Stuck	Wing Tank High Level Sw Fail On	Both Failure Modes
LH Auxiliary Tank Level	0.467cm	0.407cm	0.409cm
RH Auxiliary Tank Level	0.487cm	0.468cm	0.475cm
LH Wing Tank Level	0.471cm	0.486cm	0.490cm
RH Wing Tank Level	0.406cm	0.464cm	0.464cm
RH Auxiliary Flow Rate	0.14L/min	0.14L/min	0.14L/min
RH Flow Rate	0.23L/min	0.23L/min	0.23L/min
LH Flow Pressure	1,217Pa	1,220Pa	1,240Pa
RH Flow Pressure	934Pa	3,434Pa	954Pa
RH Wing Tank High Level Sw	**0.49**	0.0	0.0
RH Wing Tank Low Level Sw	0.0	0.0	0.0
RH Aux Tank High Level Sw	0.0	0.0	0.0
RH Aux Tank Low Level Sw	0.0	0.0	0.0

The SD results shown in Table 5.31 show that only the 'Flow Pressure Sensor Failed Stuck' failure mode would be correctly identified as false by the fault verification technique. The first order failure mode, 'Wing Tank High Level Switch Failed On', and the second order failure mode would both be verified by the fault verification technique as all of the SD values were within the tolerance limits. The technique has therefore incorrectly identified the high level switch first order fault as genuine.

The results from this second order fault demonstrate that there are limitations to

the capability of the fault verification technique. This is especially true when there is a deviation between the actual and predicted system behaviour for only a short period of time. Nonetheless the technique has been able to correctly identify the fact that a second order fault is present within the system. Given the first and second order faults are evaluated immediately after each other, it would be possible for a human operator to check the software results and identify the anomaly. Furthermore the presence of the same high level switch fault as a first and second order fault would supersede its need to be physically checked or replaced separately.

5.7 Conclusion

This chapter has presented the results of applying the fault verification technique to a range of failure modes injected on the fuel rig system. All of the failure modes were injected onto the fuel rig while it was proceeding through a phased mission comprised of five phases. The mission lasted for 300 seconds. Prior to evaluating the failure modes, the PN predicted performance of the fuel rig system with no faults present was presented. These results showed a good level of similarity with those recorded from the fuel rig, which indicated that the PN model provided an accurate representation of the fuel rig system.

A wide range of first order failure modes have been considered in this work. They include blockages, sensor failures and leaks. In every case, the PN predicted behaviour of the fuel rig system in the presence of a failure mode has been very similar to that recorded from the fuel rig. The SD values of all the system output variables have been found to be within the tolerance levels set for each variable type. These results indicate that the PN model has accurately captured the behaviour of the fuel rig in the presence of the failure modes under consideration. The analysis also investigated the ability of the fault verification technique to identify falsely diagnosed failure modes. In the majority of cases the falsely diagnosed fault was correctly identified. However, two issues were discovered when trying to identify false faults. The first issue was caused by hidden failures. These faults did not change any of the system variable behaviours during the phased mission. As a result none of the SD values exceeded the respective tolerances and, therefore, it was not possible to identify the fault as false. The second issue was caused by false faults that only caused significant behavioural variation in short phases. In these cases, the PN predicted system behaviour that included the fault was very similar to that recorded from the fuel

rig without the fault for the majority of the mission. It was only in the final phase of the mission that a significant variation appeared. However the short final phase length meant this variation was often too small to increase the SD value above the tolerance limit.

The ability of the fault verification technique to identify a single genuine failure mode amongst several other false faults was also demonstrated. On complex systems, where many hundreds of arisings could be generated over a relatively short mission, the ability of the fault verification technique to correctly and quickly categorise faults as true or false is important. The analysis showed that the genuine fault that was injected on the fuel rig was correctly verified as true, while the remaining failure modes were proven to be false.

Finally a number of second order failure modes were evaluated. When considering second order failure modes, the fault verification technique analyses each fault individually and then both faults concurrently. The results of this process demonstrated that the fault verification technique was successful at verifying the presence of multiple faults on the fuel rig system. This included a scenario where a leak fault was one of the two faults present in the system. The results did however show that the second order technique is liable to produce inaccurate results when a significant variation in the fuel rig and PN variable behaviours is only apparent in short phases. This situation led to a first order fault being incorrectly verified.

The limitation of the fault verification technique when dealing with hidden faults has been identified. The short duration of the phased mission that was considered by this work can be partially attributed to this. A longer mission may have resulted in the fuel tanks either being completely full or empty at one point, which would have revealed the level switch failures. The fault verification technique also demonstrated a limited capability when dealing with certain false arisings. When the level sensor failed stuck, for example, the deviation between the predicted and actual tank level values was small enough that the tank level SD tolerance was not exceeded. Again, in a longer mission profile this issue may not be a concern.

CHAPTER 6

Software Operation

6.1 Petri Net Software Operation Overview

This chapter will provide detail of how the PN software is structured and executed. Consideration will be given to both the software and the input files that are required for the software to operate. The PN software has been written using C++ and can be split into four sub-routines. These sub-routines and their interactions are shown in Figure 6.1.

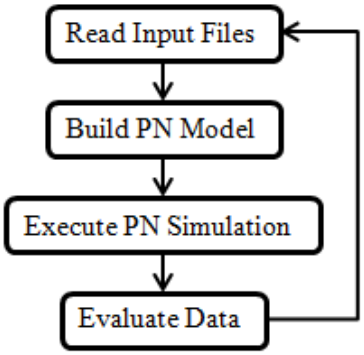

Figure 6.1: Petri net software process overview

The line on the RH side of Figure 6.1 illustrates the loop used when there are multiple arisings to consider. Multiple arisings are common on complex systems, especially during the start-up phase where sub-systems are activated at different times, which can cause numerous tolerances to be exceeded. Once all of the arisings have been evaluated, the software operation ends the process as will be discussed in detail in this chapter.

6.2 Input Files

The PN software has four input files; the PN file, the data log, the health log and the fault codes file. The PN file defines all of the PN constituent parts. An example of a PN file is shown in Figure 6.2. The PN file is separated into several sections; Places, Transitions, Initial Marking, Phase Places, Monitor and Compare.

```
PLACES
1       RH Tank Level
2       LH Tank Level
3       RH Tank Leak Size
.
.
.

TRANSITIONS
INPUTS  24      13:-1   OUTPUTS 26      DELAY   1
INPUTS  16:2    19:2    OUTPUTS 7:3     DELAY   1
INPUTS  17:2    19:2    OUTPUTS 8:3     DELAY   0
.
.
**1**

INITIAL MARKING
PLACES  1:30, 2:30, 5, 15...

PHASE PLACES
PHASE_1 5, 11, 13
PHASE_2 5, 15, 16
.
.
.

MONITOR
1       START   1       END     1       RH Tank Level
2       START   2       END     2       LH Tank Level
3       START   30      END     30      RH Flow Pressure
.
.
.

COMPARE
1 SD 1 1.5
2 SD 2 1.5
3 PRESSURE 3 9000
.
.
.
```

Figure 6.2: Petri net file example

The Places section of the PN file lists all of the place nodes in the PN and their descriptions. The Transitions section lists most of the transitions in the PN model. The only transitions not listed are those which inject the failure modes into the PN model. These transitions are selected individually as is described below. Every listed transition is displayed in terms of the respective input, output and inhibit place nodes. Edge weightings and the transition delay, if applicable, are also listed. Edge weightings are shown with

a ‘:x’ suffix to a place number, where x represents the weight of the edge. Inhibit edges are identified by a ‘:-1’ after the place number. At the end of the Transitions section is a placeholder seen as “**1**”. This code will be replaced with the phase change and failure mode transitions in a new PN file, once the data log and health log have been evaluated. Using this approach allows a single PN file to be used with any mission type, as the software will amend the PN file as appropriate. It also reduces the number of transitions that have to be included in the PN file thereby decreasing computational resources. The Initial Marking section of the PN file defines those PN places that are marked at the start of the simulation. The Phase Places section of the PN file defines which place nodes must be marked in each individual phase of system operation. The PN software can record the token count in place nodes of interest, such as those which represent flow rate or tank level. The Monitor section lists the place nodes which the PN software will record data from. Either a single place or a range of places can be monitored. The Compare section details which SD test each variable will be subject to. As values of different variables are represented by a unique number of tokens in the PN, a customised version of the SD technique is applied when evaluating each variable type. The ‘SD’ and ‘PRESSURE’ terms reflect the version of the SD technique applied to the tank level and flow pressure variables for example. The tolerance for each test is also specified.

All of the remaining input files are described below. The data log file contains all of the sensor outputs recorded from the fuel rig system over the course of the mission. The health log file lists all of the arisings that have been recorded from the fuel rig. This includes the diagnosed failure mode and the time at which the arising appeared. Finally, the fault codes file lists all of the failure modes on the fuel rig. This file also contains most parts of the transitions that are required to simulate each failure mode in the PN model. The transition delay time is added from the health log file when the new PN file is created. Figure 6.3 shows how the PN software uses these input files in the first sub-routine.

The input file sub-routine shown in Figure 6.3 requires that the four file names are input to the software. Subsequently the software will open each file and search for any and all required information automatically. From the data log the software finds the mission duration, the order of the phases in the mission and the respective phase numbers. The software opens the original PN file to find the places specific to each operational phase. These place numbers, along with the order of the phase numbers and their duration, will be used to list the operational phase changes in the Transition section of the new PN file.

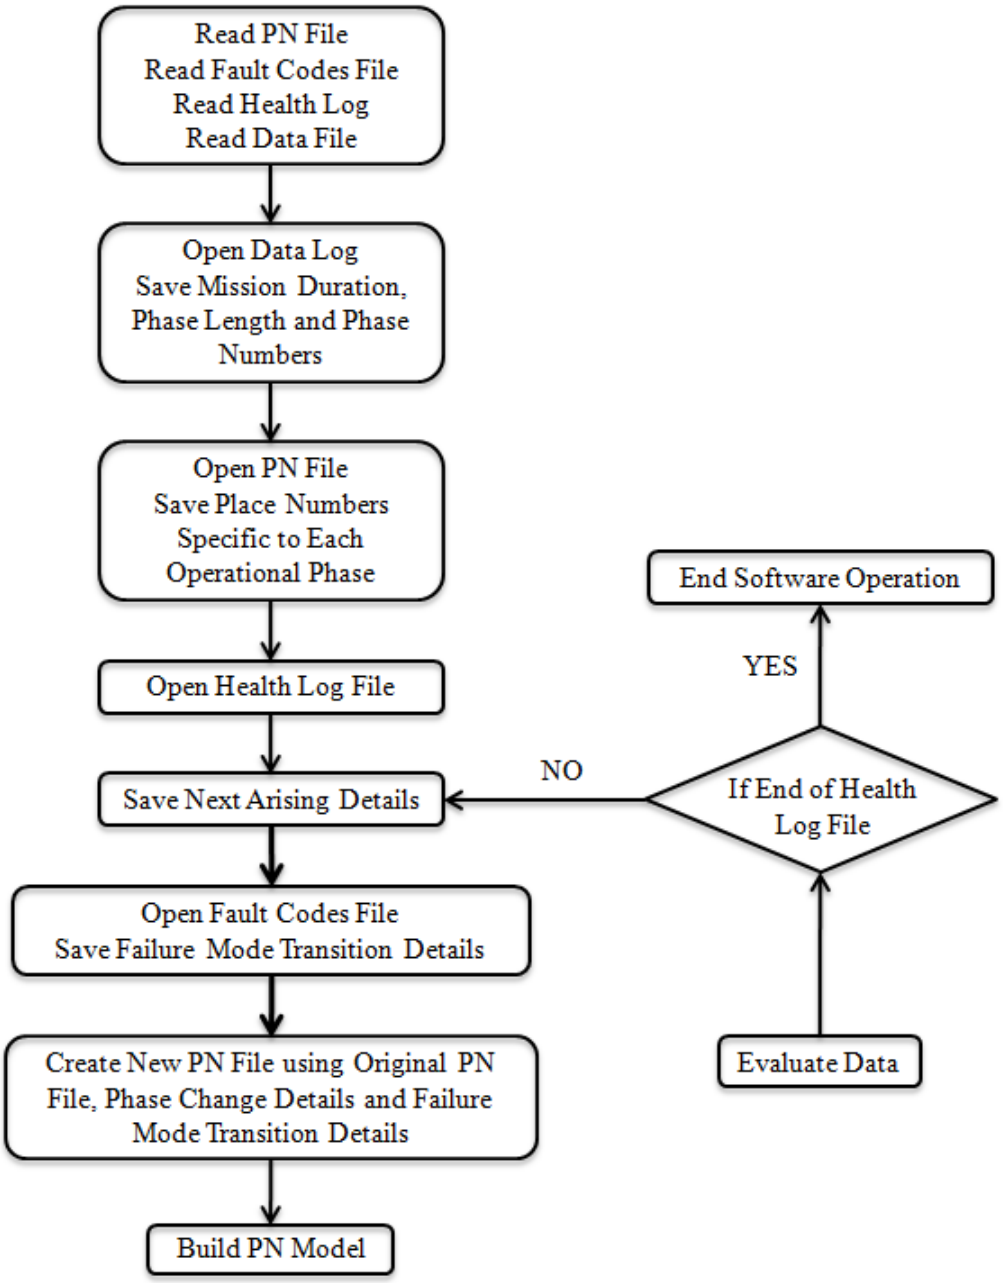

Figure 6.3: Petri net software input file sub-routine

The health log file is then opened to find the first arising. The failure mode code contained in the arising detail is matched with a code from the fault codes file. Having matched the respective fault codes, the correct failure mode transition can also be listed in the new PN file. The finalised new PN file is then input to the sub-routine, which builds the PN model in the software.

If the health log file contains more than one arising, then a loop in the code will return

the software process to the point shown in Figure 6.3, until all of the arisings have been considered. Each time a new arising is considered, the arising details have to be saved and a new match found in the fault codes file. An up-to-date new PN file can then be constructed and the remainder of the software process executed.

6.3 Computational Model

The software builds the computational model of the PN from the new PN file. An Object Orientated Programming approach is used when considering the Place, Transition and Monitor sections. Each of these variables is defined as a 'class', of which there are several 'instances'. A single PN place is an instance of the Place class, for example, and a single transition is an instance of the Transition class. This way all of the PN details can be stored in the computational model in an orderly and structured fashion. The sub-routine also uses the information listed in the Initial Marking sub-section of the PN file to add tokens to the relevant places prior to executing the PN simulation. Information in the Phase Places and Compare sub-sections of the PN file are not used in this sub-routine.

Once the PN model has been read-in and built by the software, it can be executed.

6.4 PN Simulation Execution

A sub-routine within the PN software titled 'nTimeSteps' allows the PN model to be executed for a defined time period or mission. Figure 6.4 shows the flow chart of the 'nTimeSteps' sub-routine.

The only input requirement of the 'nTimeSteps' sub-routine is to define the time step length and the mission duration. The mission duration is known from the first sub-routine where it was identified from the data log. The time step length can be any integer value; a step length of 1 second was applied throughout the work. The simulation time variable is also defined and set to zero by the software at the beginning of the operation. The time step length is the value by which the simulation time will be incremented once all of the PN transitions have been evaluated. This command can be seen on Figure 6.4 when the response to decision node 1 is 'NO'. Decision node 1 compares the current simulation time to the mission duration. If the simulation time is less than the mission duration, the simulation time is increased by the time step length and a counter, i, is initiated.

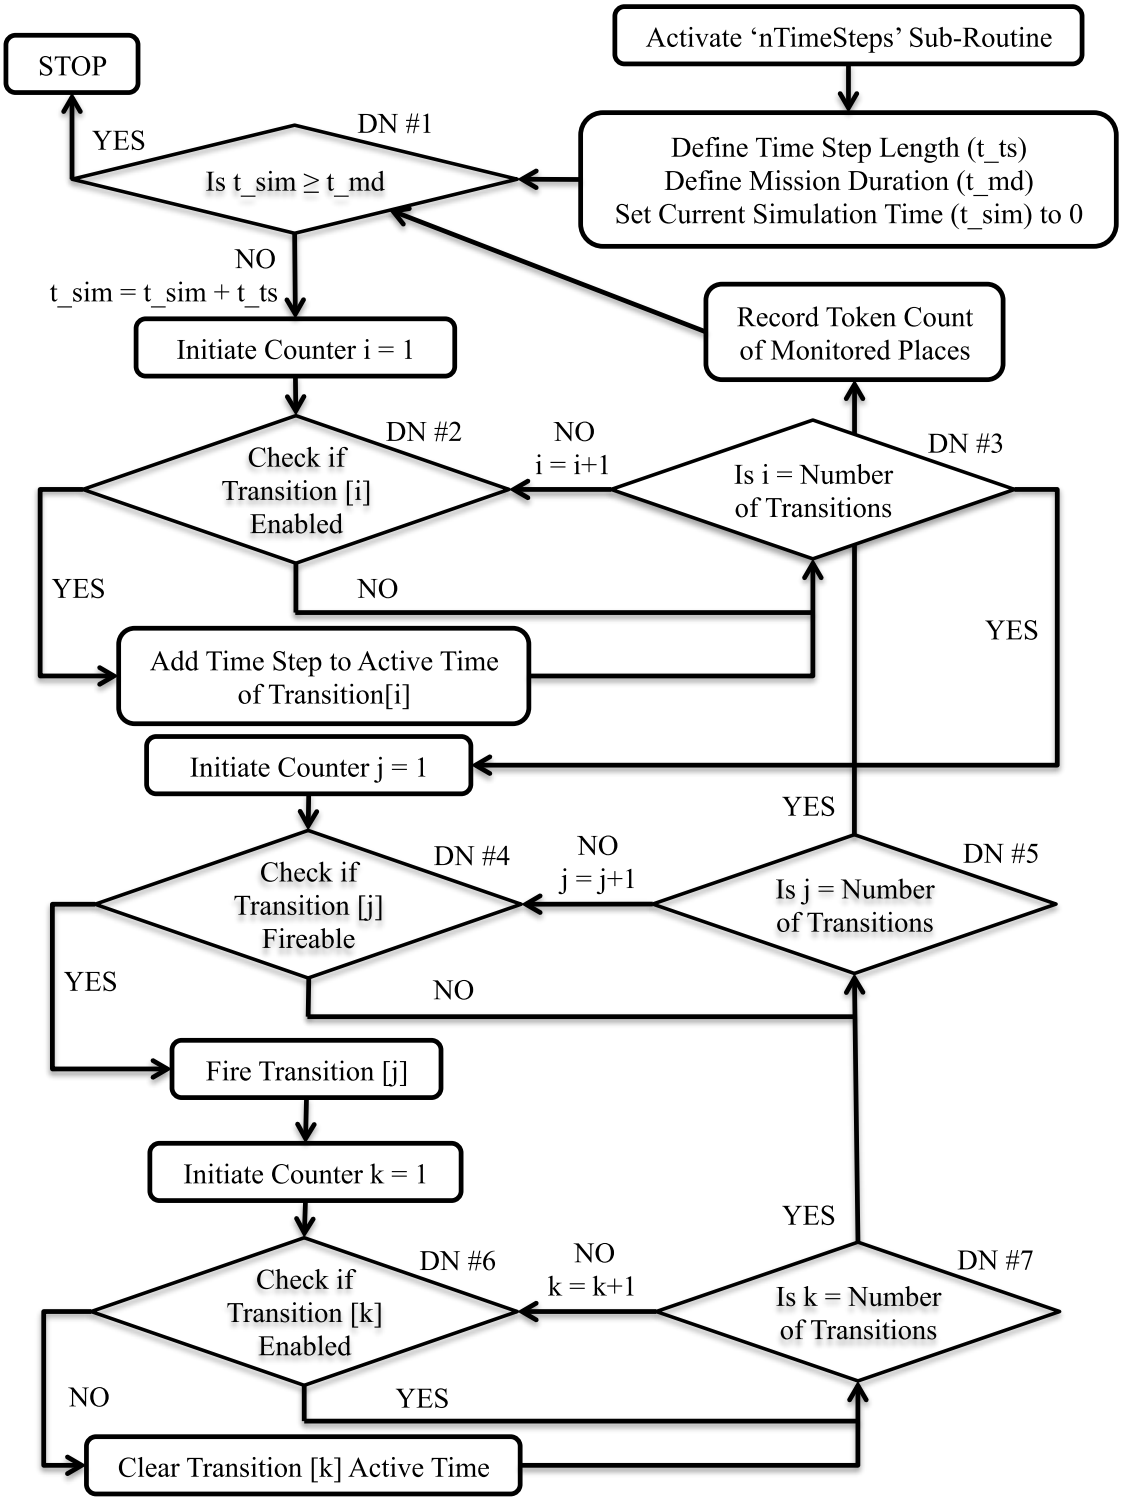

Figure 6.4: Petri net simulation sub-routine

Using this counter, all of the transitions in the PN are evaluated. Those transitions that are enabled have their active time increased by the time step length. Once all of

the transitions have been evaluated, the response to decision node 3 will be 'YES' and a further counter, j, is initiated. When decision node 4 identifies a transition which is fireable, that is the required number of input tokens are present and the transition delay has been satisfied, the software fires the transition and initiates counter k. This counter is used to evaluate if the effect of the transition firing has changed the enabled state of any other transitions. Those transitions that are no longer enabled have their active time reset to zero. Any transitions that remain enabled are not affected. Once all of the transitions have been checked, the output from decision node 7 will be 'YES'. The loop originating from decision node 4, in the lower half of Figure 6.4, then continues until counter j is equal to the number of transitions in the PN. At this time decision node 5 will produce a 'YES' output and the number of tokens in the places being monitored will be recorded. The process of evaluating the PN transitions is repeated, until the simulation time is no less than the mission duration, at which point the sub-routine terminates. The processes undertaken by the sub-routine in between each time decision node 1 makes a decision are known as a 'cycle'. A mission of duration 100 seconds, modelled with a 1 second timestep will therefore undertake 100 cycles.

6.5 Data Evaluation

The final sub-routine of the PN software is the comparison of the fuel rig data from the data log with the data recorded from the PN simulation. As was stated in Chapter 4, the first step in the comparison of the fuel rig variables is to convert the recorded number of tokens from the PN simulation into numerical values. This process is unique to each variable and is dependent on the value of the tokens representing each variable. Having determined the numerical value of a variable at every time step throughout the mission the SD comparison technique is applied as described in Chapter 4. Once all of the variables listed in the Compare sub-section of the PN file have been evaluated, the software will check to see if all of the arisings listed in the health log file have been evaluated. If they have, the program will end. If there are arisings that still require investigation, the software will return to the first sub-routine following the route shown in Figure 6.3.

6.6 Software Performance

The fuel rig PN model contains 239 places and 454 transitions. The PN software was used to simulate a 300 second phased mission of the fuel rig system. Every simulation included the analysis of fourteen variables. The analysis of a single failure mode using the software on a Samsung P510 laptop computer took between 5 and 10 seconds. The analysis of a second order failure mode therefore took approximately 20 to 30 seconds as three unique simulations were evaluated.

6.7 Conclusions

This chapter has described how the PN software, written in C++, was structured and executed. The software was designed to provide a high level of flexibility and reduced computational requirements. The ability to consider any type of phased mission without having to make changes to the software demonstrates the built-in flexibility. The short analysis times on a computer with relatively low processing power indicates that the software could be effectively applied to a real system.

CHAPTER 7

Airbus A340 Fuel System

7.1 Introduction

The BAE Systems fuel rig provided a means by which to test the PN modelling and fault verification techniques on a small scale system. In order to demonstrate the applicability of the PN modelling technique to a larger system, the fuel system of the Airbus A340 aircraft will be modelled using the PN technique. A number of first order faults and a second order fault will then be propagated through the model. The effects of these faults on the system will be captured using the PN software and the results evaluated. Successful application of the PN modelling technique to the Airbus fuel system will demonstrate the flexibility and applicability of the technique to both small and large scale systems.

7.2 System Description

7.2.1 System Operation

The Airbus A340's fuel system is comprised of eight fuel tanks. There are four inner tanks, two in each wing, two outer tanks, one in each wing, a centre tank in the body of the aircraft and a trim tank in the tail of the aircraft. Each inner tank feeds a collector cell which, in turn, feeds fuel to one of the aircraft's four engines. Each outer tank and the trim tank also have a connecting surge tank, which provides capacity for fuel expansion at high temperatures. Figure 8.1 shows the layout and piping arrangement of the A340 fuel system [34].

Figure 8.1 shows that each engine has a dedicated collector cell which provides the

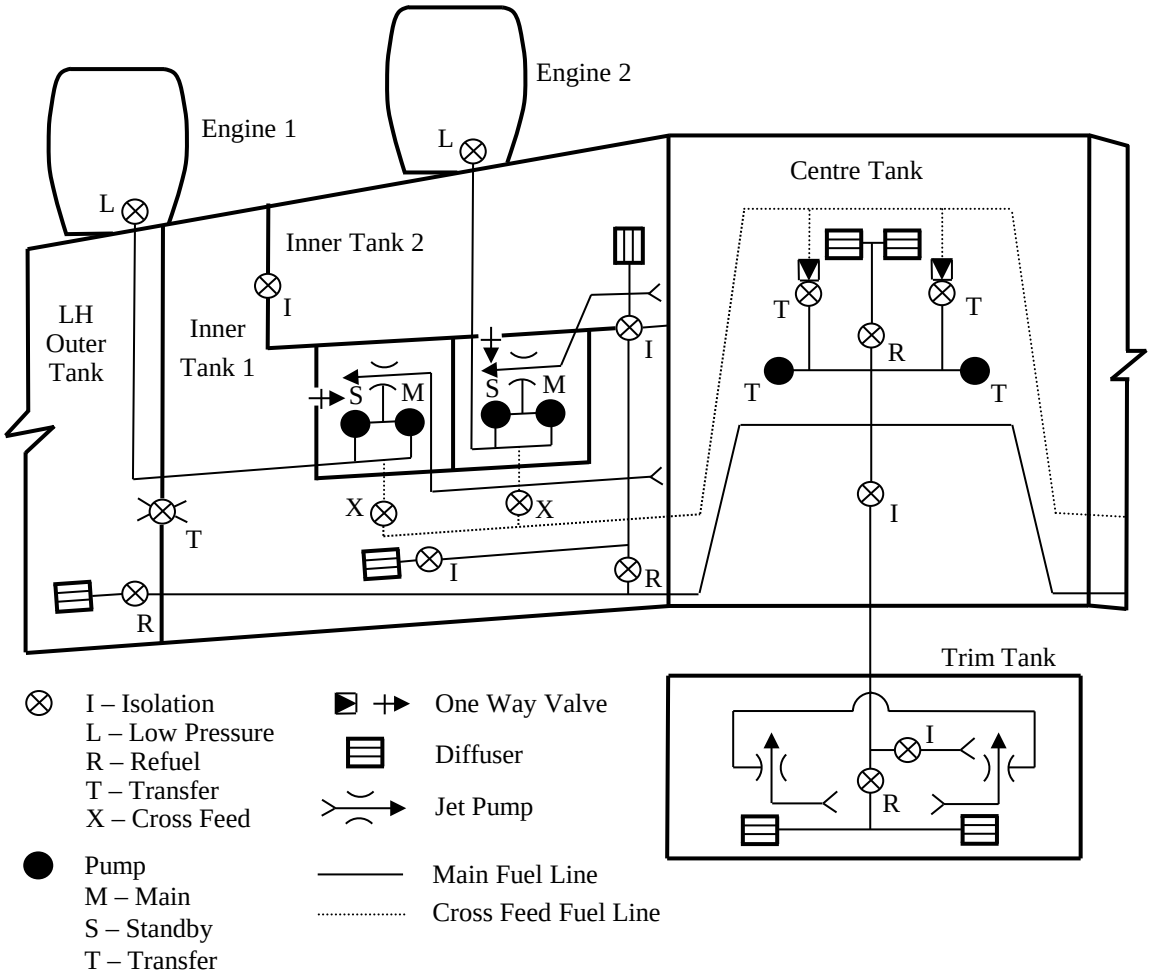

Figure 7.1: Airbus A340 fuel system

fuel supply. There are two pumps in each collector cell, 'Main' and 'Standby', which can deliver the fuel to the engine. The standby pump is a cold back up; it is only used if there is a fault with the main pump. Fuel is transferred from the inner tanks to the collector cells by jet pump. If this pump fails there is a one-way valve which can be opened to allow the transfer of fuel to the collector cell by gravity. There is also a series of piping which can transfer fuel from the collector cell to the centre tank in the case of a fault, such as an engine failure or a blockage in the supply pipe to the engine. The centre tank can supply fuel to all of the other tanks on the aircraft. Pumps in the centre tank and isolation valves on the system are used to control the transfer of fuel from the centre tank to its intended destination.

The transfer of fuel to and from the trim tank is used to control the centre of gravity point on the aircraft. Prior to take-off the trim tank is empty. Once the aircraft reaches

the end of its climb phase, fuel is pumped into the trim tank from the centre tank to optimise the centre of gravity location for cruise. The flight management system controls the forward transfer of fuel as the inner tank volumes fall and move the centre of gravity aft of its original position. Transfer of fuel to and from the trim tank is restricted below certain altitudes and within 75 minutes of landing. The trim tank is fitted with a main and standby jet pump to feed fuel forward to the centre tank for dispersion to the inner tanks.

The outer tanks are located at the furthest point from the centre of the aircraft. Fuel in these tanks can be transferred to the inner tanks by gravitational means. Table 8.1 lists the capacity of the fuel tanks on the A340 [34].

Table 7.1: Airbus A340 fuel tank volumes

Inner Tank	15,000L
Collector Cell	5,952L
Centre Tank	41,468L
Tail Tank	6,114L
Outer Tank	3,624L
Total Fuel Volume	138,638L

The cross feed piping in the A340 fuel system is in place to allow the re-distribution of fuel in the event of a fault such as an engine failure. Should engine 1 fail, for example, the cross feed valve shown below the collector cell would open allowing fuel to return to the centre tank. This fuel is then distributed to all of the tanks to maintain the ability of the aircraft to complete its mission. Fuel would continue to be supplied to all of the inner tanks to maintain the distribution of weight on the aircraft even though only three of the engines would be in use.

7.2.2 Fuel Usage

The distribution of fuel within the A340 fuel system is carefully controlled to ensure the centre of gravity position of the aircraft allows an efficient production of net lift and, as much as possible, reduces structural bending moments. The use of fuel from different tanks is therefore specifically prescribed in order to manage this distribution of weight.

An engine is always supplied with fuel from its collector cell, as shown in Figure 8.1. Fuel taken from this cell is automatically replaced with fuel from the related inner tank. The inner tanks are maintained full with fuel from the centre tank, until this has been emptied. The fuel volume in the inner tanks is then allowed to fall until the flight control and monitoring system initiates the forward transfer of fuel from the trim tank. The point at which this transfer occurs is determined by a number of factors including mission profile, atmospheric conditions, the aircraft weight and the distribution of that weight. It will be assumed when modelling the A340 fuel system that the forward transfer of fuel begins when the inner tank volumes have fallen to 75,000L (50%) and stops when they reach 80,000L (53%). Multiple forward fuel transfers may therefore occur until the trim tank is emptied. Fuel is then transferred in from the outer tanks when the inner tank volumes fall to 12% of their volume. Any remaining fuel in the inner tanks and collector cells would then be used before the system is empty. As mentioned previously, the trim tank is filled using fuel from the centre tank.

7.2.3 Operating Phase Flow Rates

The Airbus A340 is a commercial aircraft that operates in numerous phases of operation. Six unique phases have been identified for the purpose of modelling a phased mission. These phases are; taxi, take-off, climb, cruise, descend and approach. Table 8.2 shows the flow rates for these phases that have been identified from literature [35] [36].

Table 7.2: Airbus A340 mission flow rates

Taxi	0.20L/sec
Take-Off	1.90L/sec
Climb	1.30L/sec
Cruise	0.50L/sec
Descend	0.10L/sec
Approach	0.50L/sec

7.2.4 System Sensors

It is known that on the Airbus A340 there are a variety of sensors that monitor, among other variables, the fuel quantity, fuel level and fuel temperature. This work will consider only the volume of the fuel tanks on the system and flow rate at a number of chosen locations. The location of the flow rate sensors is shown on Figure 7.2.

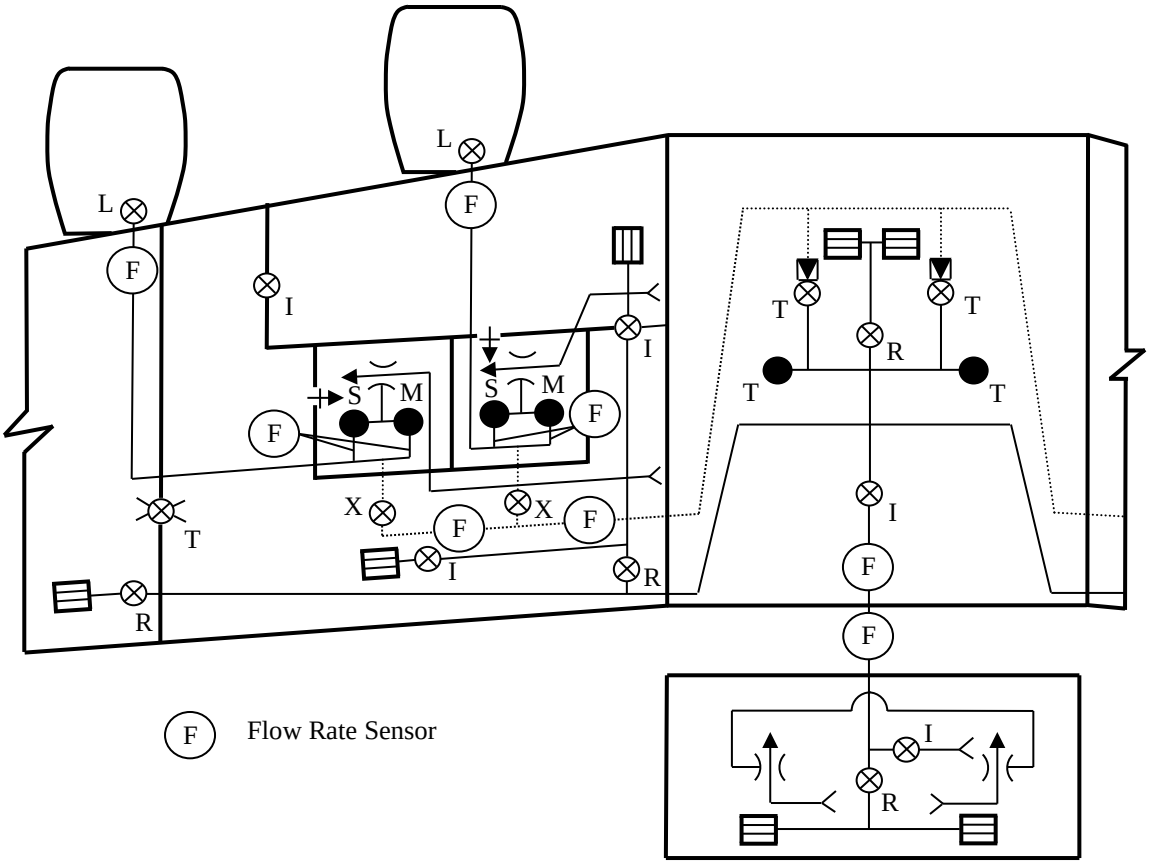

Figure 7.2: Location of flow sensors on A340 petri net model

7.3 Petri Net Model

7.3.1 Overview

The Airbus A340 fuel system has been modelled using the PN technique. The transition types used to model the fuel system are the same as those that were used to model the BAE Systems fuel rig; the standard transition type and the three specialist ones which were created. No further new transition or place types were developed. The A340 PN model

contains 208 place nodes and 451 transitions. Additional transitions would be created dependent on the mission profile and the number of faults that are to be injected into the model.

The complete model of the BAE Systems fuel rig contains more place (239) and transition (454) nodes, than the basic model of the more complex A340 fuel system. This is due to the fact that the A340 fuel system PN model includes only three failure modes. These faults have been specifically selected to demonstrate a variety of effects on the system. Had the fuel rig PN model only considered three faults that are equivalent to those modelled in the A340 PN, it would contain 127 place nodes and 203 transitions. The size of this basic fuel rig PN model is considerably smaller than the A340 fuel rig PN and indicates the greater level of complexity present in the A340 system.

7.3.2 Sub-Net Details

The Airbus A340 PN is presented in the form of a series of sub-nets. As with the fuel rig PN sub-nets, the order by which the sections of the PN are presented reflects the order by which they are listed in the PN software input file. Where there are two place descriptions associated with a single place in a sub-net figure, this represents two separate sub-nets where each place description is associated with one of the sub-nets.

7.3.2.1 Clear Flow Rate Places

The first parts of the PN that are listed in the input file are those that remove tokens from the flow rate and fuel used places, Figure 7.3 shows these sub-nets.

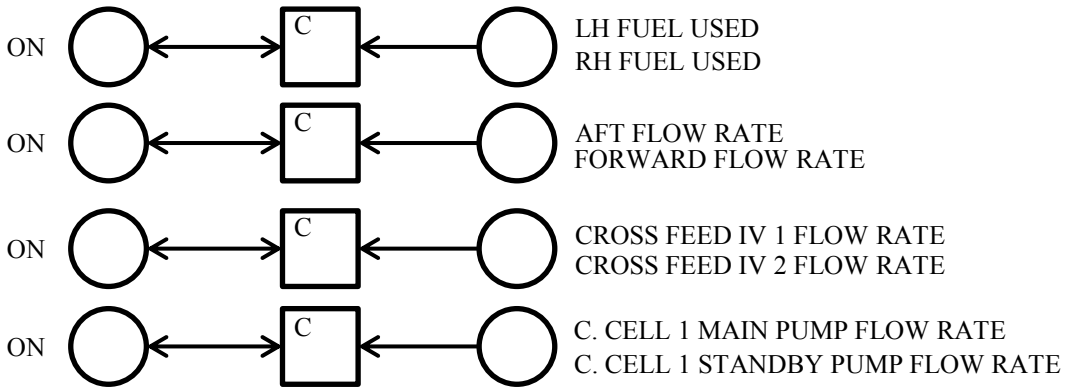

Figure 7.3: A340 petri net model flow rate and fuel used clear transitions

Figure 7.3 shows that, every time step, all of the tokens in the flow rate and fuel used places are removed. The fuel used places represent the fuel consumed by the engines on the LHS and RHS of the system respectively. The fuel used measure is necessary to indicate how much fuel has to be taken from the centre tank to replenish the spent fuel.

7.3.2.2 Fuel Pump State Changes

Figure 7.4 shows the sub-nets that control the state of the collector cell 1 pumps.

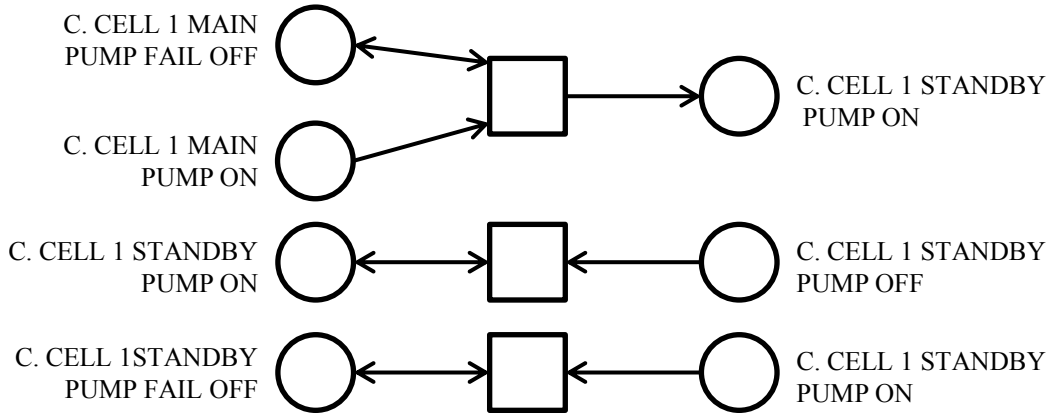

Figure 7.4: A340 petri net model collector cell pump states

It can be seen from the top sub-net in Figure 7.4 that when the main pump fails off, the standby pump comes online when the transition fires. The transition in the centre sub-net would then fire to update the PN marking. The transition in the bottom sub-net of the figure would fire if the standby pump was to fail. Similar sub-nets exist for all four of the collector cells in the A340 PN model. There are also equivalent sub-nets which control the operational state of the trim tank jet pumps. The flow of fuel from the inner tanks to the collector cells is controlled by a jet pump and, in the case of a failure, a one-way valve. Figure 7.5 shows how the PN models the failure of the jet pump.

Figure 7.5 shows that should the inner tank 1 jet pump fail, the inner tank 1 one-way valve is opened to maintain a supply of fuel to the collector cell. Equivalent sub-nets exist for all of the inner tank jet pumps and one-way valves.

7.3.2.3 Engine Fuel Feed Process - Taxi Phase

Figure 7.6 shows the sub-nets that remove tokens from collector cell 1 during the taxi phase of a mission.

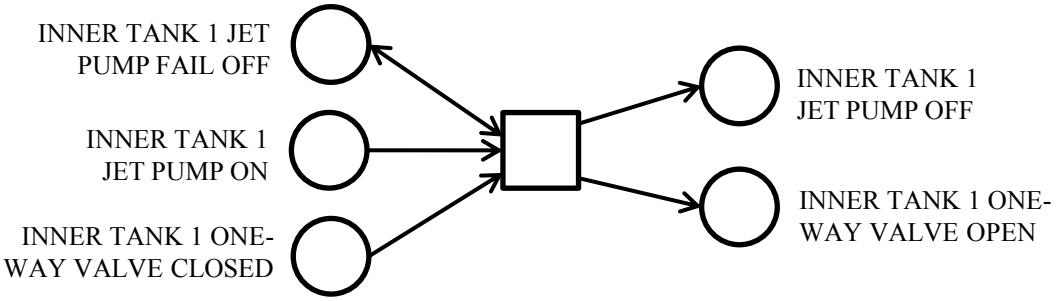

Figure 7.5: A340 petri net model inner tank jet pump states

The transition in the top sub-net of Figure 7.6 would fire when the system is operating in the taxi phase and the main collector cell pump is on. Firing the transition removes two tokens from the collector cell volume place and moves them to the fuel at collector cell main pump exit place. As the taxi phase flow fuel rate is 0.2L/sec, each fuel token represents 0.1L of fuel. The inhibit edges prevent this transition firing if there is a build-up of fuel at the collector cell main pump exit or if pipe section 1 is blocked. Pipe section 1 is assumed to be the section of pipe that stems vertically from the main collector cell pump in Figure 8.1. Pipe section 2 is the equivalent section of pipe that is connected to the standby collector cell pump. Pipe section 3 is the section of pipe that connects the outputs of pipe section 1 and 2 to engine 1. The central sub-net in Figure 7.6 moves the fuel tokens from the collector cell main pump exit to the standby pump exit in the taxi phase of operation. The bottom sub-net will move fuel from the collector cell, through the standby pump, to the standby pump exit in the taxiing operational phase. The transition in this sub-net will only fire if the collector cell standby pump is on. Figure 7.7 shows how the PN models the flow of fuel from the collector cell standby pump exit to the engine when the system is operating in the taxi phase.

The transition in the top sub-net of Figure 7.7 can fire to move fuel tokens from the standby pump exit place to the engine low pressure valve input place. The movement of these tokens will only occur if the main or standby collector cell pump is on. The lower sub-net moves these tokens from the low pressure valve input to the exit. In order for fuel to reach the low pressure valve exit either the main or standby collector cell pump must be on. In Figures 7.6 and 7.7, all of the sub-nets showed how the model represented the behaviour of the system in the taxiing operational phase. Equivalent sub-nets exist for all of the operational phases considered by this work, as listed in Table 8.2. The number of

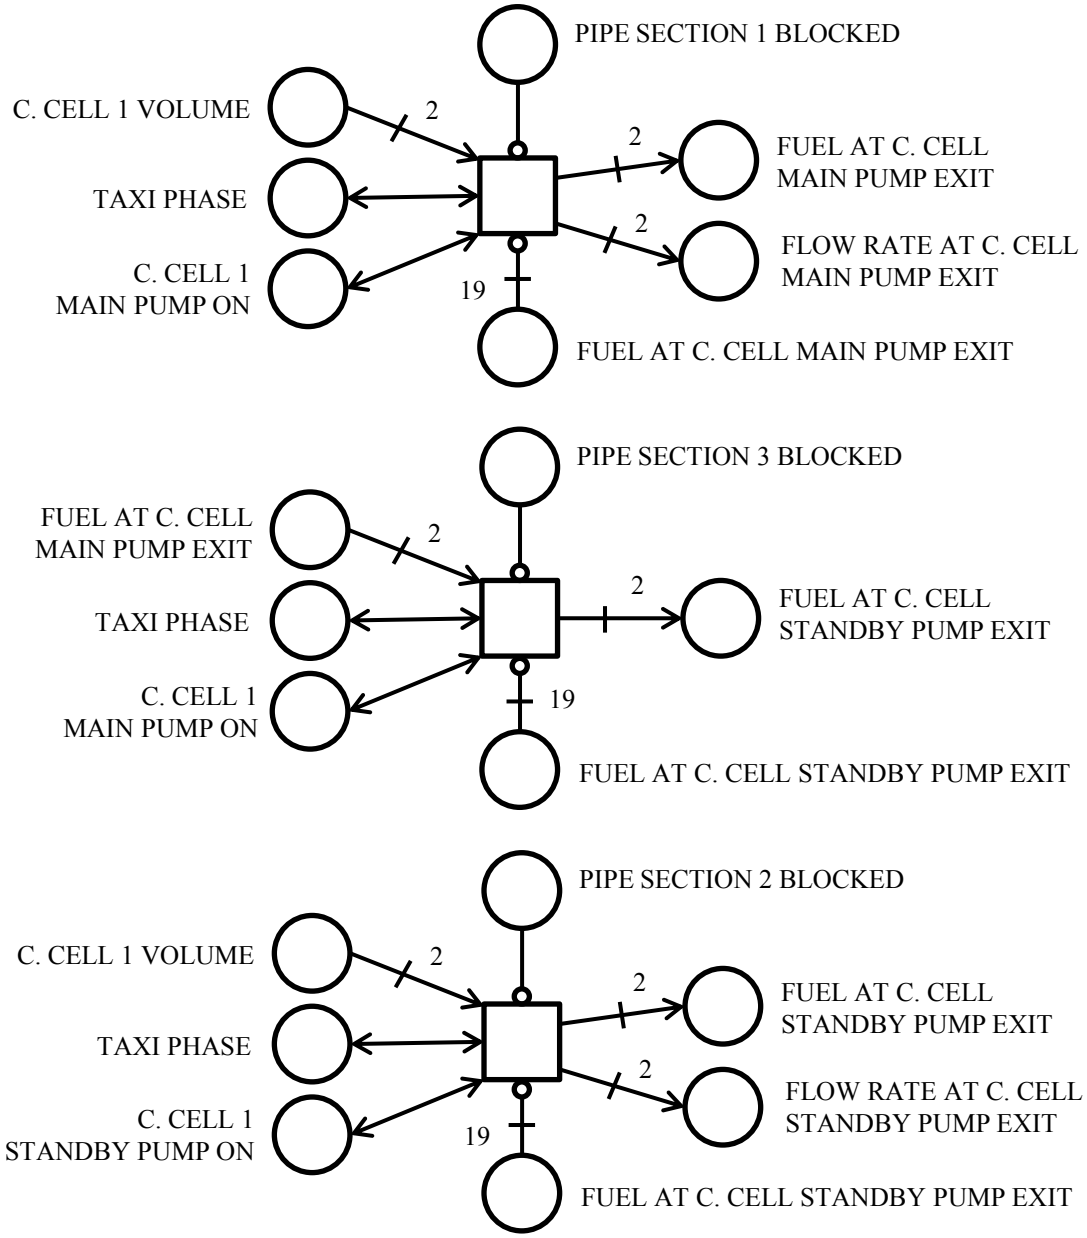

Figure 7.6: A340 petri net model collector cell output

fuel tokens moved during each operational phase is directly related to the flow rates listed in the same table. The weighting associated with the relevant inhibit edges in Figures 7.6 and 7.7 are changed such that no more than twenty fuel tokens can be present in the respective places in any phase of operation. This represents the assumed maximum capacity of the fuel pipes in the system.

Once fuel is removed from a collector cell, it is replaced with fuel from the inner tank. Figure 7.8 shows how this process is modelled in the taxiing phase.

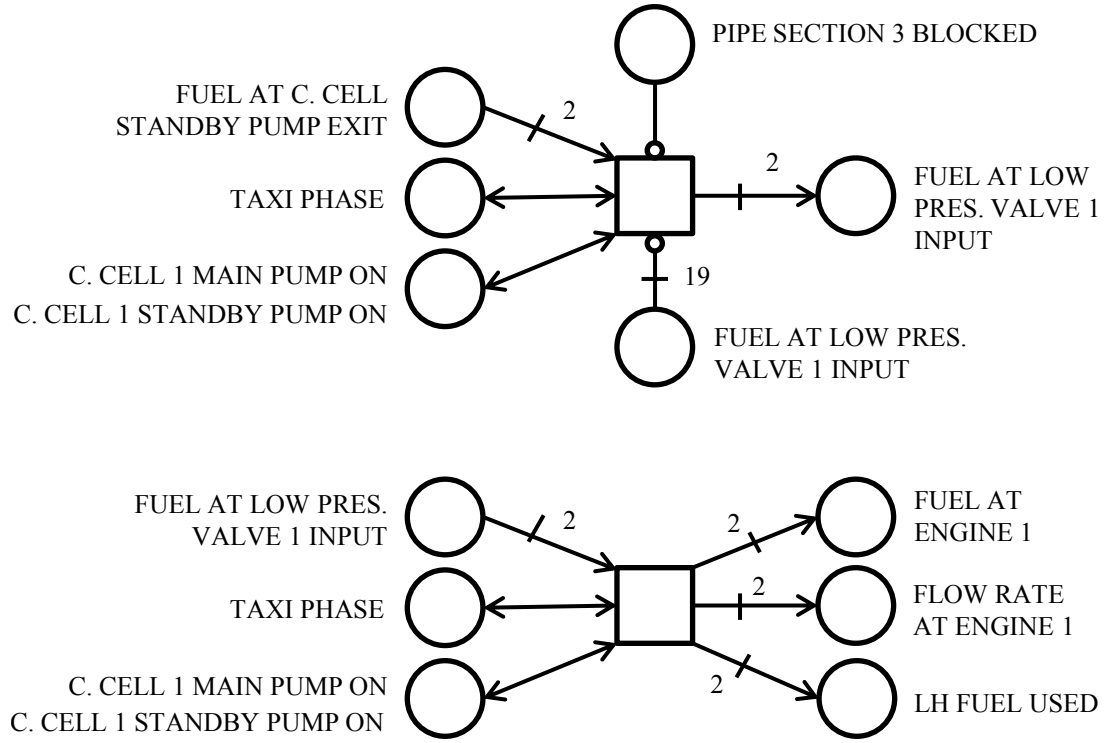

Figure 7.7: A340 petri net model engine 1 input

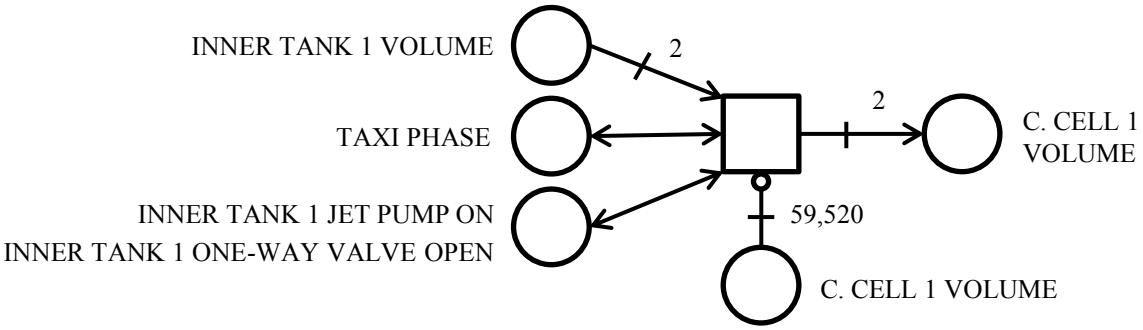

Figure 7.8: A340 petri net model collector cell refill

Figure 7.8 shows that, when the jet pump is on or one-way valve is open, two fuel tokens are taken from the inner tank and put into the collector cell to replace those lost to the engine. The sub-net shown in Figure 7.8 is only applicable in the taxiing phase, similar sub-nets have been constructed for all other operational phases considered and for all four of the inner tank/collector cell pairs in the system.

7.3.2.4 Cross Feed Fuel Flow Process

If fuel cannot reach an engine from a collector cell, due to a blockage for example, it will be fed back to the centre tank. In order for this to happen the cross feed isolation valve, centre tank transfer isolation valve and centre tank refuel isolation valve have to be opened. Figure 7.9 shows the sub-nets that model the opening and closing of these valves.

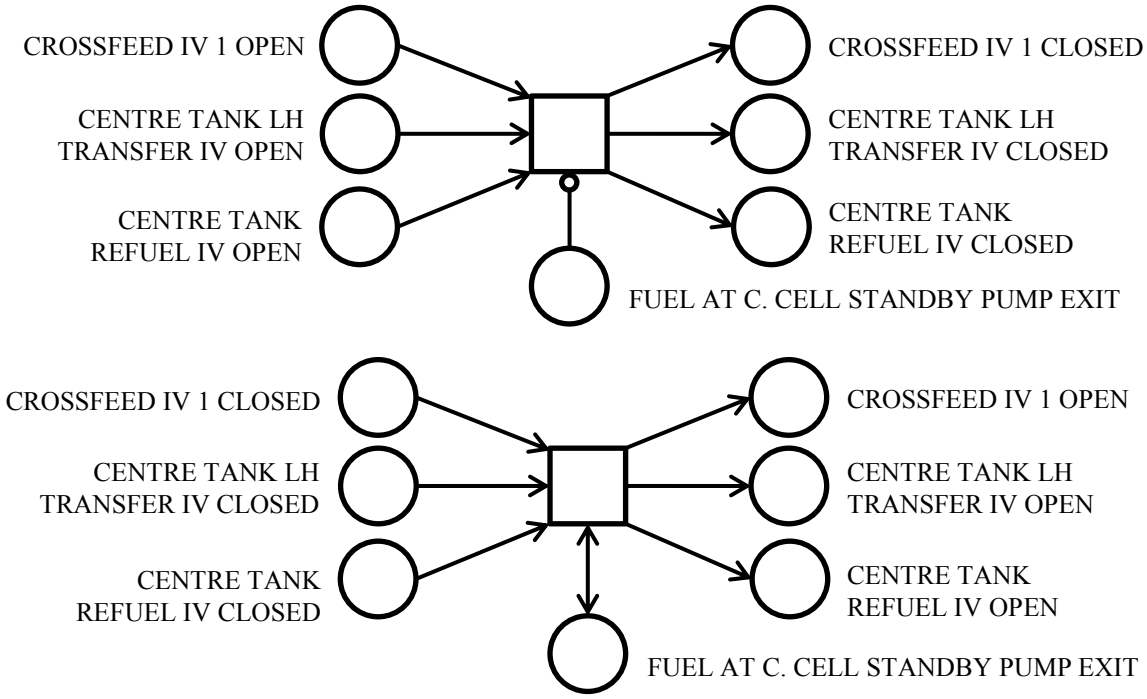

Figure 7.9: A340 petri net model cross feed valve states

The upper sub-net in Figure 7.9 closes the cross feed valves when there is no fuel at the standby pump exit. Considering the closing of the valves first ensures that the valves are not opened incorrectly or left open when no fuel is present at the collector cell standby pump exit. The lower sub-net opens the cross feed valves when a fuel token is present in the collector cell standby pump exit place. Once the cross feed valves have been opened, fuel can flow from the collector cell to the centre tank. Figures 7.10 and 7.11 show how this behaviour has been modelled in the A340 PN.

Figure 7.10 shows that the movement of fuel tokens from start of the cross feed pipe section at tank 1 to the centre tank is closely monitored. Each sub-net moves the fuel tokens along the cross feed section toward the centre tank. The second and third sub-nets in this figure also record the flow rates at the relevant points in the cross feed section of

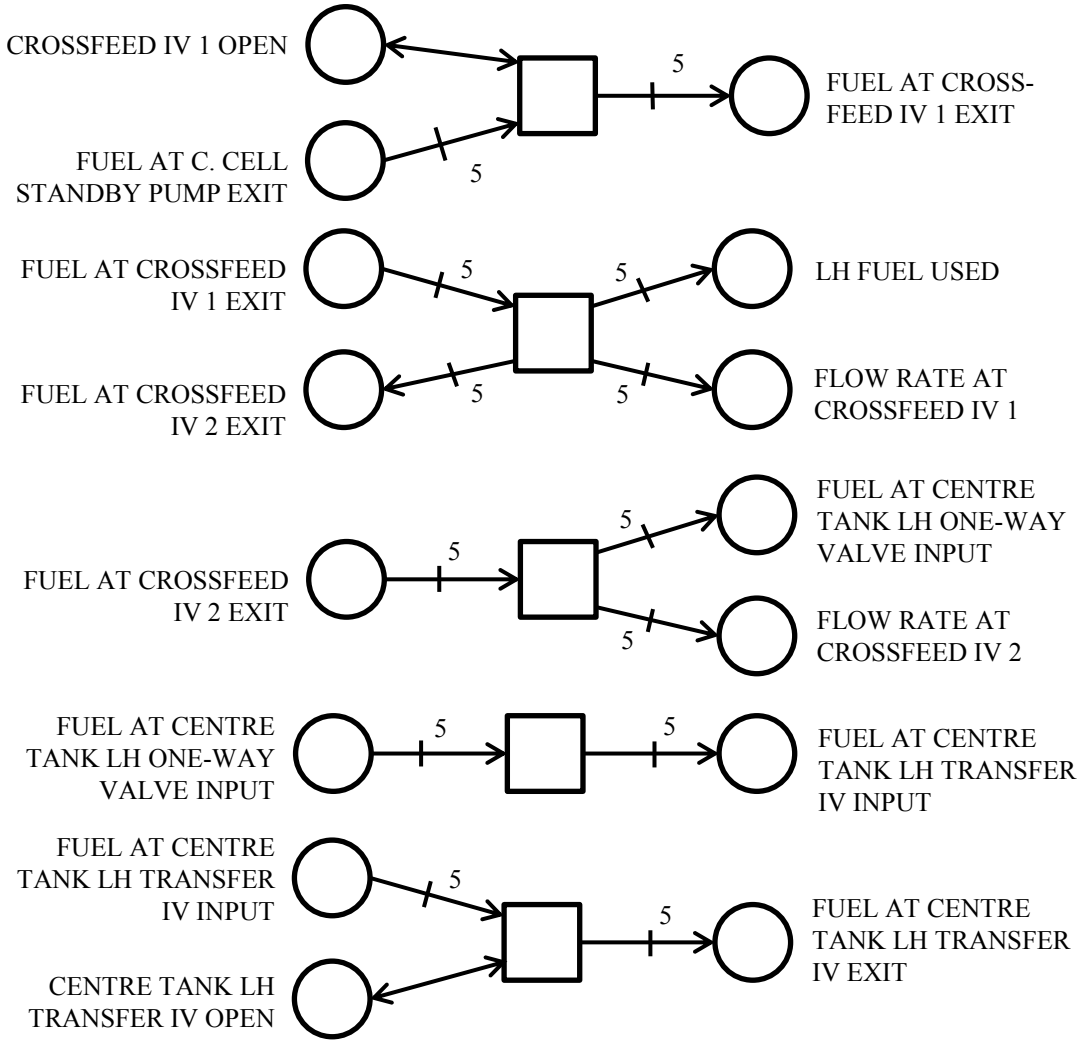

Figure 7.10: A340 petri net model cross feed fuel flow 1/2

the system as shown in Figure 7.2. In Figure 7.11, the flow of fuel through the centre tank refuelling pipes and into the centre tank itself is modelled. All of the sub-nets shown in Figures 7.10 and 7.11 model the behaviour of the fuel flow through the cross feed section of the system during the cruise phase. Similar sub-nets are also modelled for the remaining phases of operation, not shown here, where the flow rate is either higher or lower than that modelled for the cruise phase in Figures 7.10 and 7.11

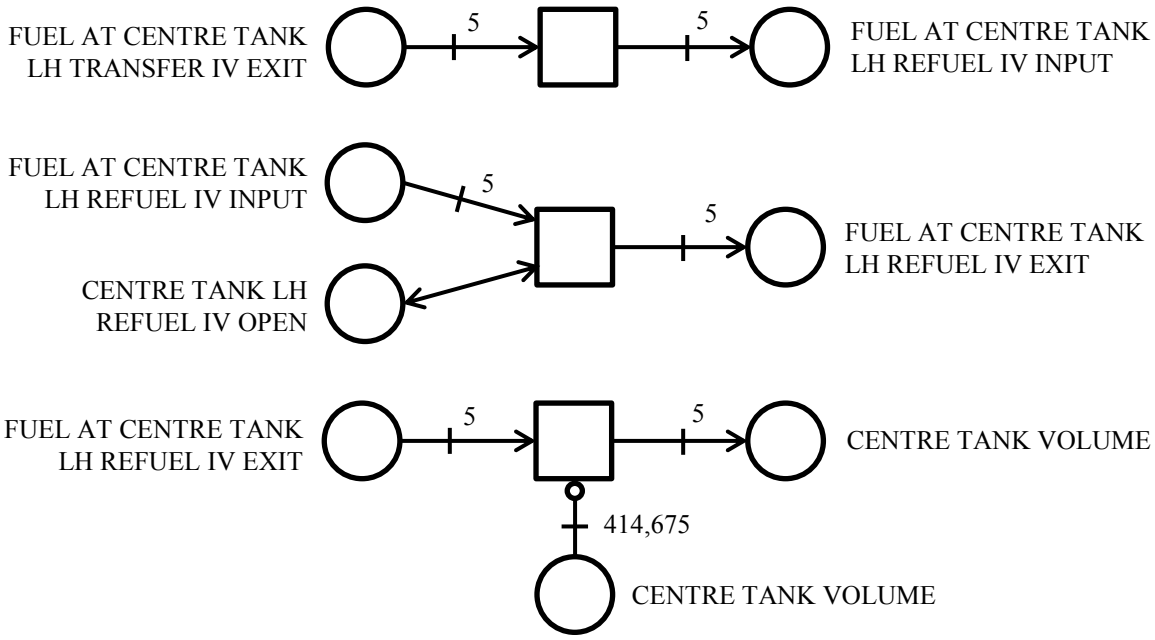

Figure 7.11: A340 petri net model cross feed fuel flow 2/2

7.3.2.5 Aft Fuel Transfer Process

Figure 7.12 shows the sub-nets that control the flow of fuel from the centre tank to the trim tank once the cruise phase of operation begins.

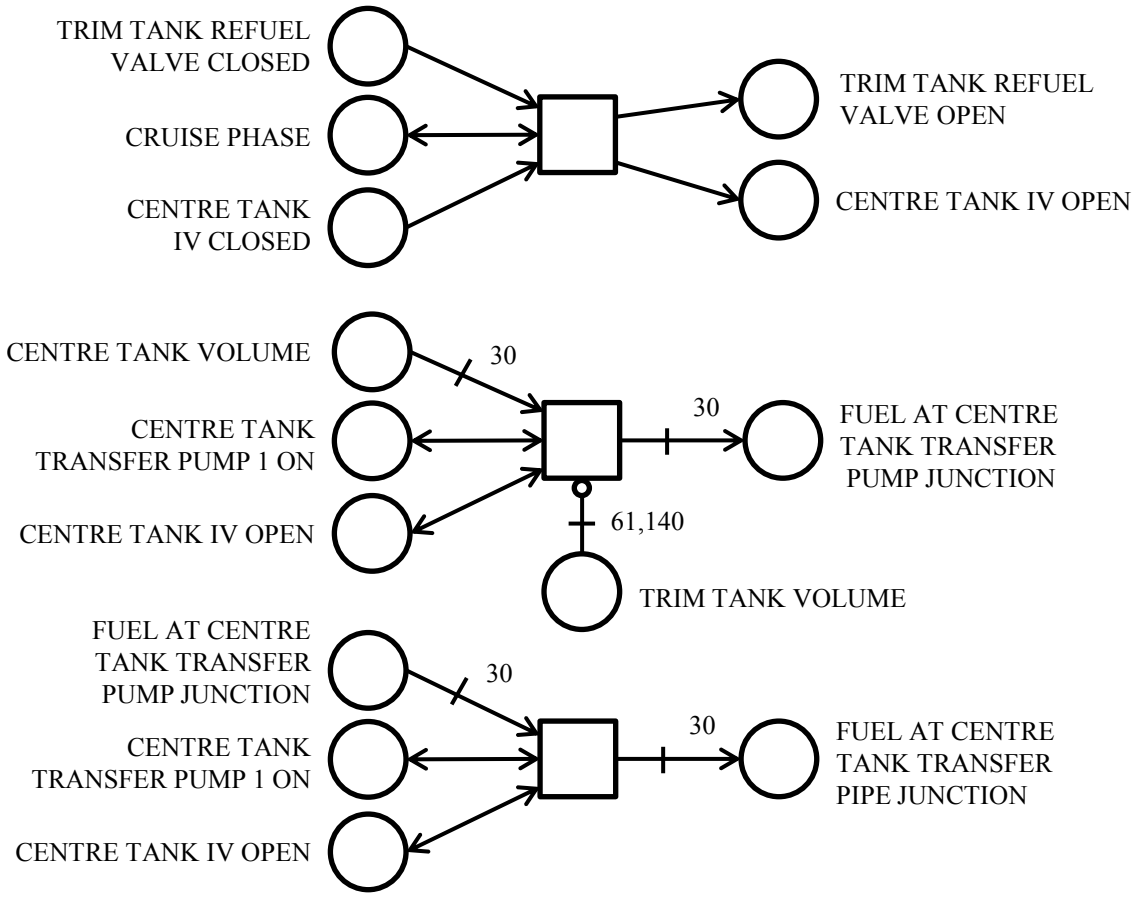

Figure 7.12: A340 petri net model aft fuel transfer 1/2

The top sub-net in Figure 7.12 opens the trim tank refuel valve and centre tank isolation valve when the cruise phase begins. The lower two sub-nets are then used to model the flow of fuel from the centre tank to the centre tank transfer pump junction, seen on Figure 8.1 as the junction immediately above the isolation valve in the centre tank. A flow rate of 30 fuel tokens per second was assumed for the purpose of modelling. The movement of fuel from the centre tank transfer pump junction to the trim tank is modelled by the top two sub-nets in Figure 7.13.

The top sub-nets in Figure 7.13 move the fuel tokens into the trim tank. This action requires that centre tank transfer pump 1 is on. The third sub-net in the figure closes the trim tank refuel valve and centre tank isolation valve when the volume of fuel in the trim

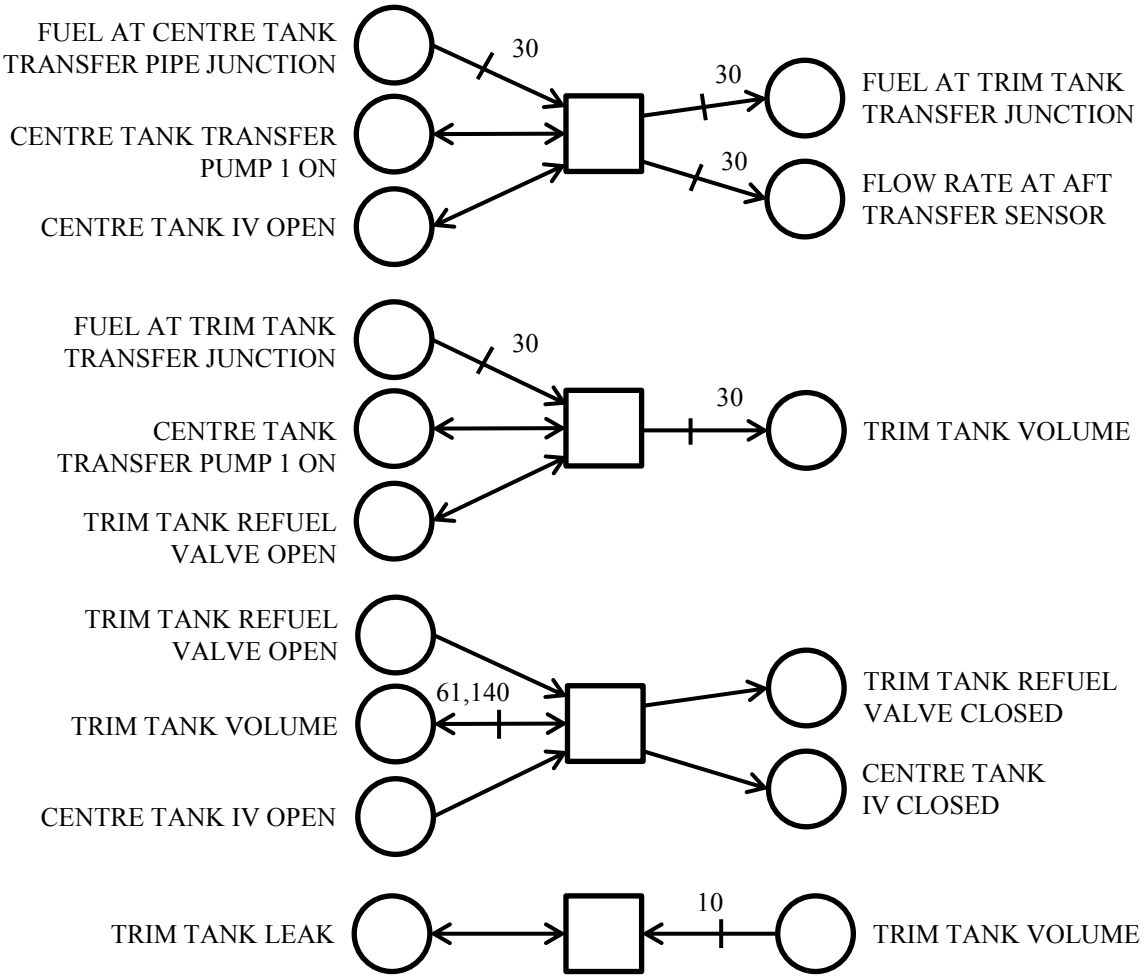

Figure 7.13: A340 petri net model aft fuel transfer 2/2 and trim tank leak

tank has reached its capacity. The final sub-net shown in Figure 7.13 models the effect of a leak in the base of the trim tank. Only a large leak size has been modelled as the current version of the A340 PN model is being used as a demonstrator. A large leak size has been modelled to demonstrate, most clearly, the effect of the fault on the system. A complete model would consider multiple leak sizes.

7.3.2.6 Inner Tank Refuel Process

When fuel is been removed from the inner tanks to replenish the collector cells, it is replaced with fuel from the centre tank, when possible. Figures 7.14, 7.15 and 7.16 model this process. Figure 7.14 contains the sub-nets that control the centre tank pumps, inner tank refuel valves and inner tank isolation valves.

The top sub-net in Figure 7.14 will turn off centre tank transfer pump 1 if the volume

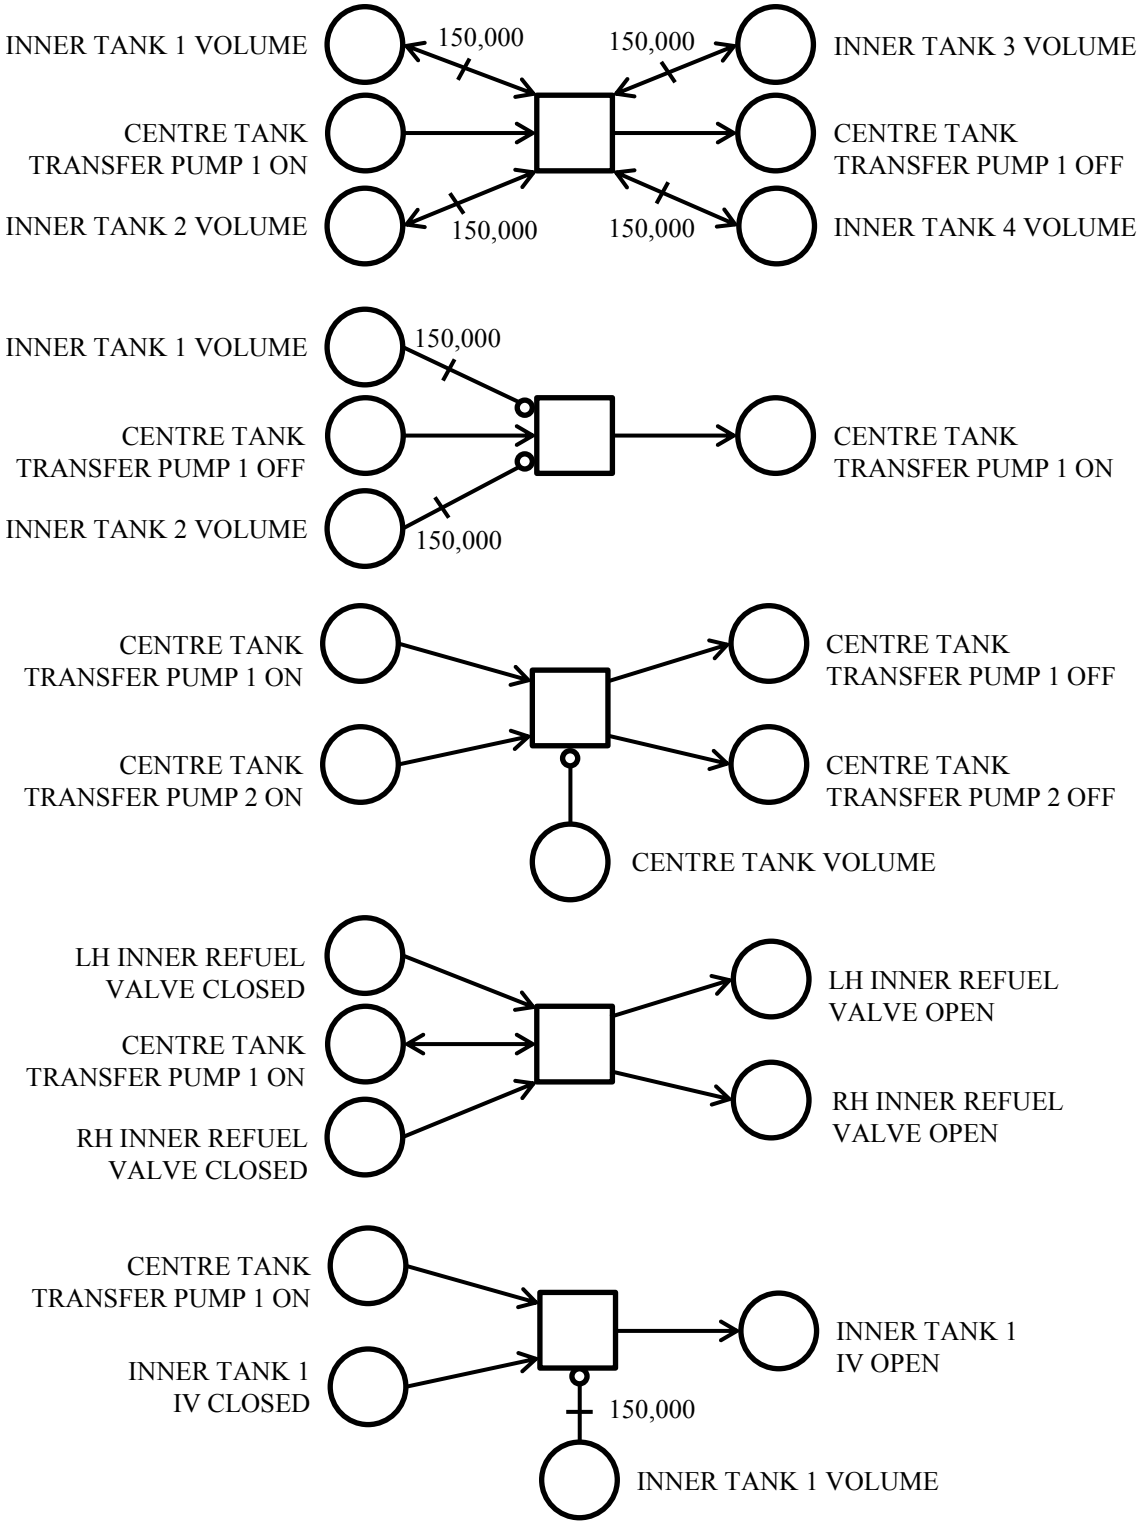

Figure 7.14: A340 petri net model inner tank refuelling 1/3

of fuel in all of the inner tanks is 15,000L; their maximum volume. A similar sub-net is in place for centre tank transfer pump 2. The second sub-net turns on centre tank transfer

pump 1 if the volume of fuel in inner tank 1 or 2 falls below 15,000L. An equivalent sub-net turns on centre tank transfer pump 2 if the volume of fuel in inner tank 3 or 4 is below 15,000L. The third sub-net ensures the centre tank transfer pumps are turned off if the centre tank level falls to zero. The fourth sub-net shown in Figure 7.14 open the LH inner refuel valve and RH inner refuel valve when centre tank transfer pump 1 is on. An equivalent transition will open these same valves if centre tank transfer pump 2 is on. The final sub-net shown in the figure opens the inner tank 1 isolation valve when centre tank transfer pump 1 is on. The inhibit place will prevent the isolation valve being opened if the inner tank is already full of fuel. An equivalent sub-net models the opening of the other inner tank isolation valves on the system.

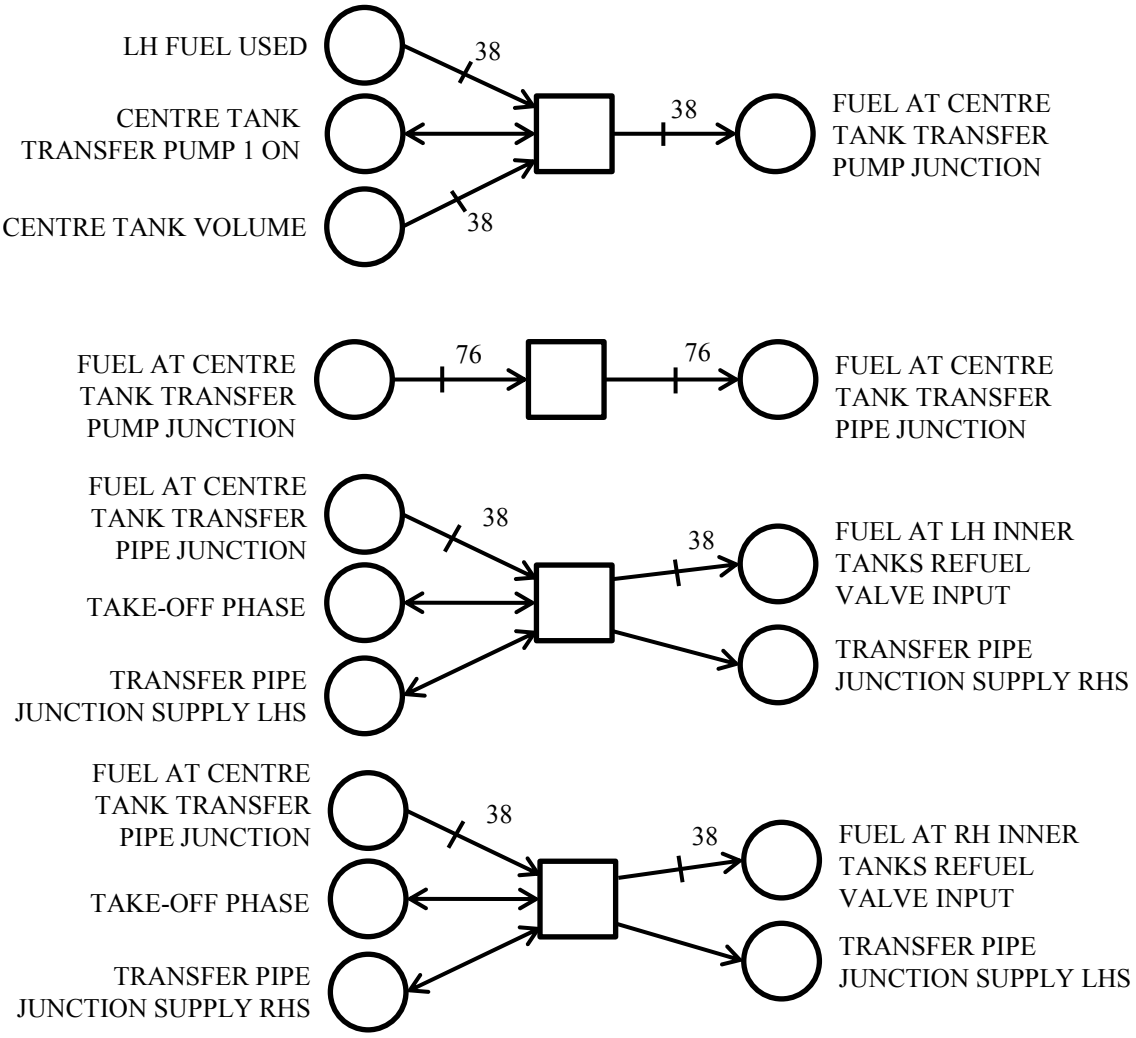

Figure 7.15: A340 petri net model inner tank refuelling 2/3

The sub-nets shown in Figure 7.15 move fuel from the centre tank to the inputs of the

refuel valves on the LHS and RHS of the system. In the first sub-net fuel is removed from the centre tank. The amount of fuel removed from the centre tank is the same as that which was used by the engines on the LHS of the system. The sub-net shown represents the fuel removed from the centre tank during the take-off phase of the mission. Similar sub-nets model this behaviour in every phase of operation. Equivalent sub-nets use centre tank transfer pump 2 to remove fuel from the centre tank to replace that used by the engines on the RHS of the system. The second sub-net moves the combined amount of fuel removed from the centre tank to the transfer pipe junction. The third and fourth sub-nets are then used to distribute this fuel to the LHS and RHS of the system equally. All of the sub-nets shown in Figure 7.15 are only applicable to the take-off phase of operation. Further sets of sub-nets are in place to model the remaining phases of operation.

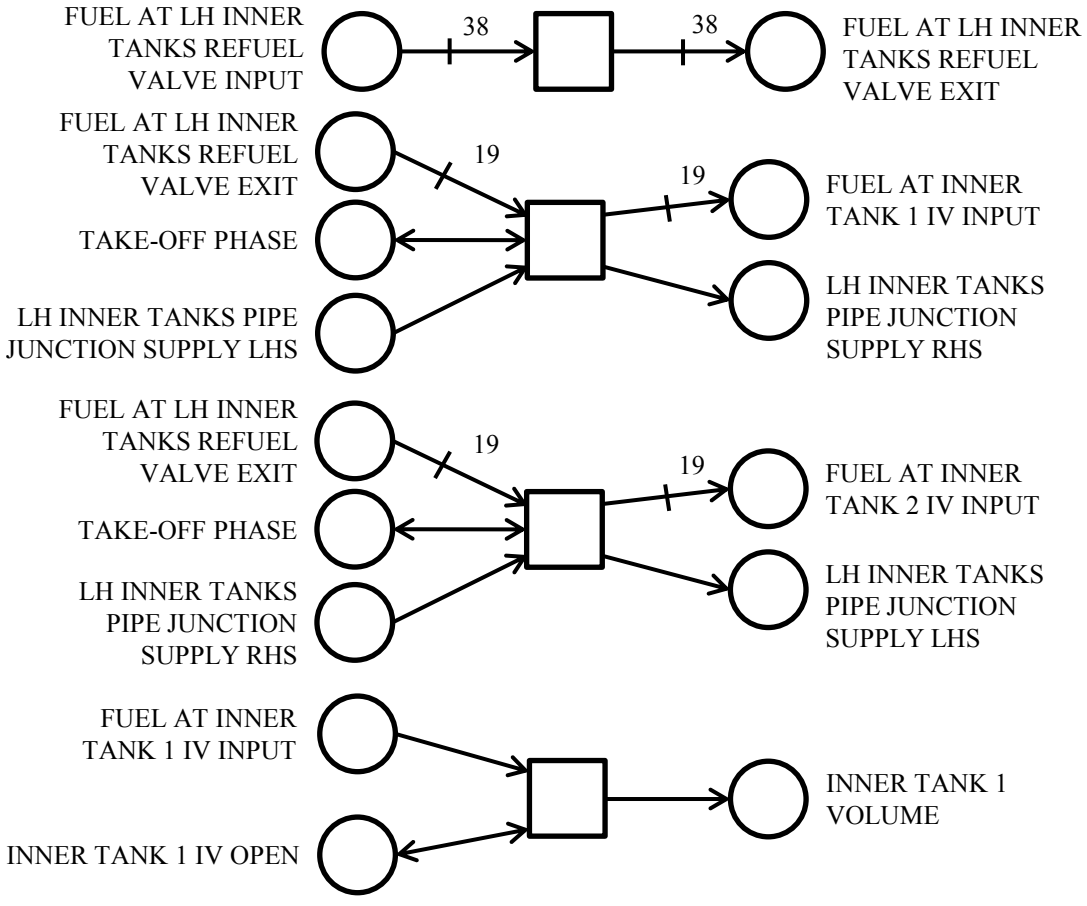

Figure 7.16: A340 petri net model inner tank refuelling 3/3

Figure 7.16 shows the sub-nets which control the movement of fuel tokens from the input of the inner tanks refuel valve on LHS of the system to the inner tanks themselves.

In the first sub-net fuel is moved from the input of the refuel valve to the exit. The second and third sub-nets split the fuel provided by the centre tank to inner tank 1 and inner tank 2 input pipes. The final sub-net moves the fuel tokens from the input of the isolation valve at inner tank 1 to the tank. Equivalent sub-nets for all of the tanks on the system and in every phase of operation are also present within the model but not shown for brevity.

7.3.2.7 Forward Fuel Transfer Process

The sub-nets which control the trim tank jet pumps, and therefore the forward transfer of fuel from the trim tank to the inner tanks, are shown in Figures 7.17, 7.18 and 7.19. Fuel is transferred forward from the trim tanks when the amount of fuel in the inner tanks falls below 75,000L. The transfer is stopped if the inner tank levels reach 80,000L.

The first sub-net in Figure 7.17 moves fuel tokens from the trim tank to the trim tank transfer junction if the trim tank isolation valve is open and either trim tank jet pump 1 or 2 is on. The second sub-net moves the fuel tokens to the centre tank transfer pipe junction. This requires the centre tank isolation valve be open and that one of the trim tank jet pumps be on. The third and fourth sub-nets split the fuel from the trim tank to supply all of the inner tanks on the system. These sub-nets are equivalent to the third and fourth sub-nets shown in Figure 7.15.

The sub-nets contained within Figure 7.18 control the trim tank jet pump, trim tank isolation valve and centre tank isolation valve. In the first sub-net the trim tank isolation valve is opened and trim tank jet pump 1 turned on when the fuel volume in inner tank 1 falls below 75,000L. In this sub-net the centre tank isolation valve is always open. The second sub-net in Figure 7.18 is similar to the first except that it opens the centre tank isolation valve in addition to changing the state of the trim tank isolation valve and trim tank jet pump 1. Both of these sub-nets are inhibited by the presence of any fuel in the centre tank, as the priority list requires fuel in this tank to be used first. The final sub-net in the figure closes the isolation valve and turns off the jet pump if the tank level in inner tank 1 reaches 80,000L. There are equivalent sub-nets in the model that use trim tank jet pump 2 as an input/output component should jet pump 1 have experienced a failure.

The first sub-net in Figure 7.19 closes the trim tank isolation valve and centre tank isolation valve when the trim tank is emptied. It also turns off trim tank jet pump 1. A similar sub-net closes the two isolation valves and turns off jet pump 2. The second sub-net in the figure closes the refuel valve and inner tank isolation valve on the LHS of

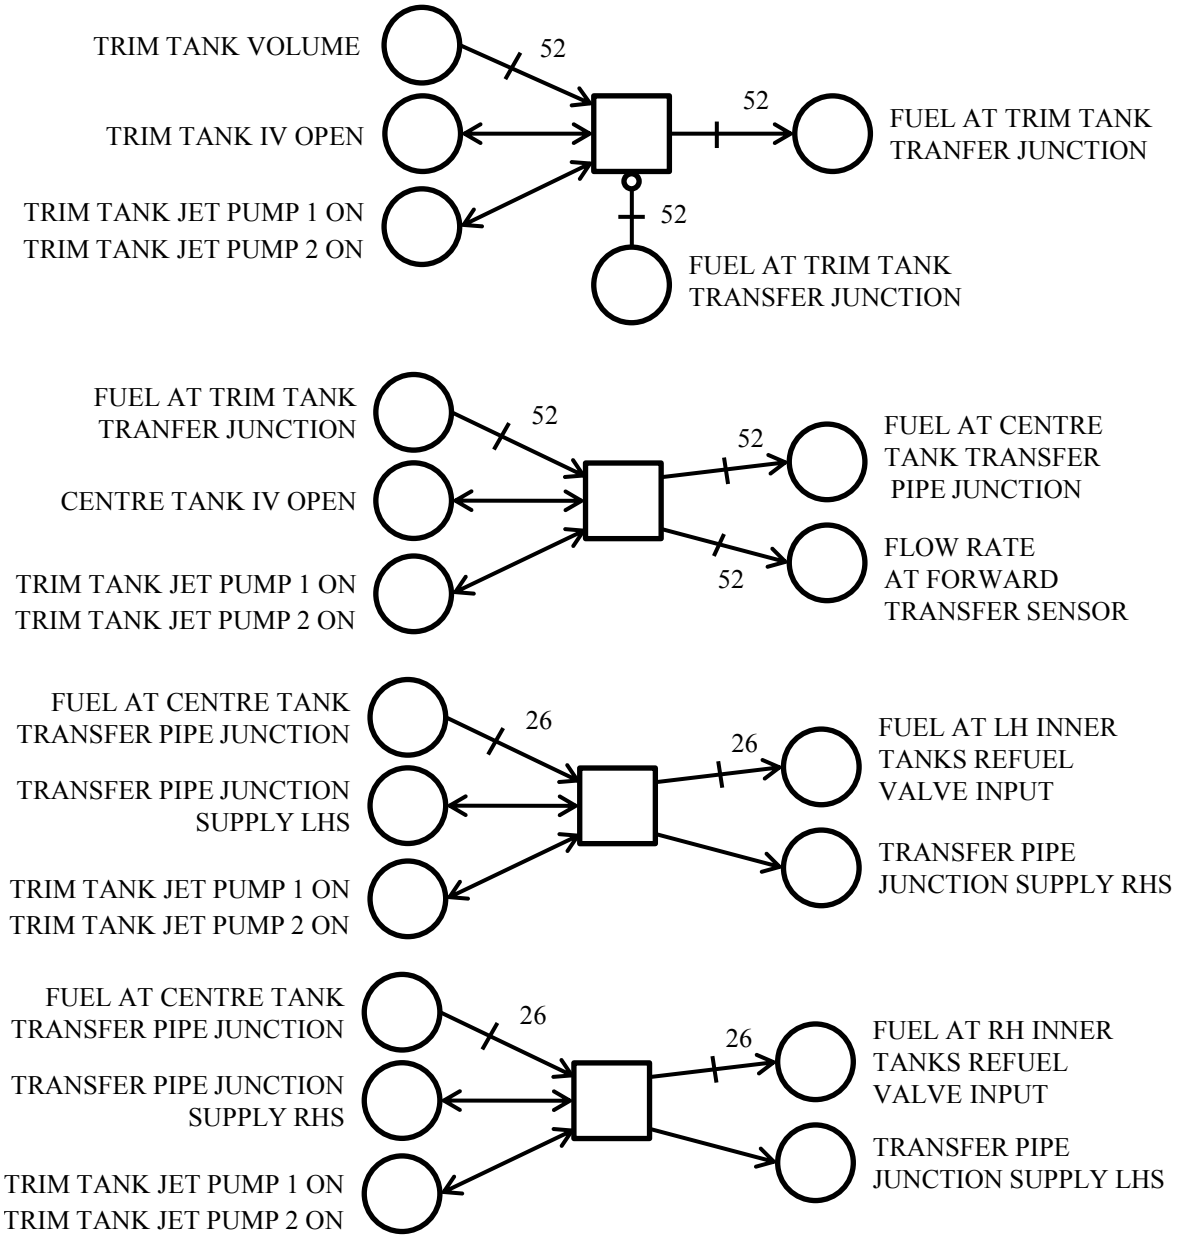

Figure 7.17: A340 petri net model forward fuel transfer 1/3

the system when the trim tank no longer contains fuel. An equivalent sub-net closes the respective valves on the RHS of the system.

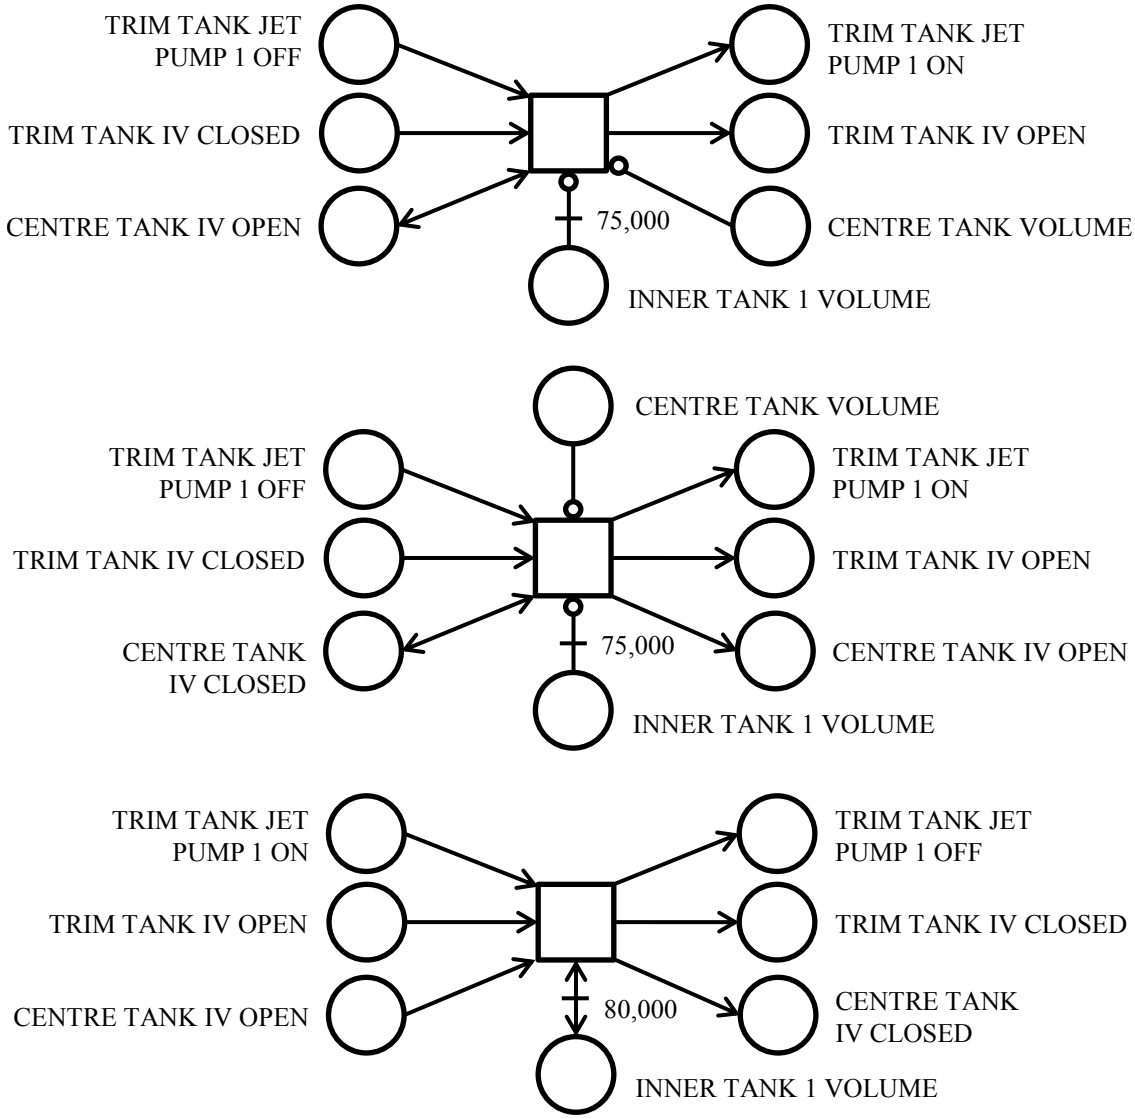

Figure 7.18: A340 petri net model forward fuel transfer 2/3

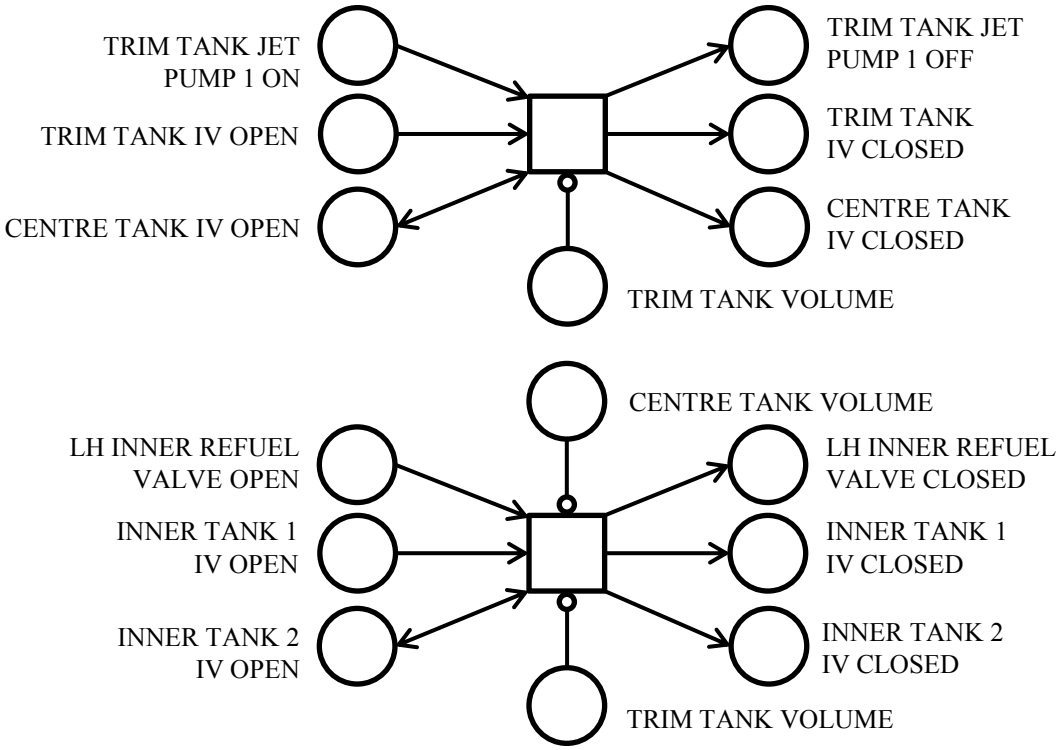

Figure 7.19: A340 petri net model forward fuel transfer 3/3

7.3.2.8 Outer Tank Fuel Flow Transfer

The final aspect of the fuel system's behaviour that is listed in the input file is the fuel transfer from the outer tanks to the inner tanks. Figures 7.20 and 7.21 display how this behaviour has been modelled in the A340 PN model.

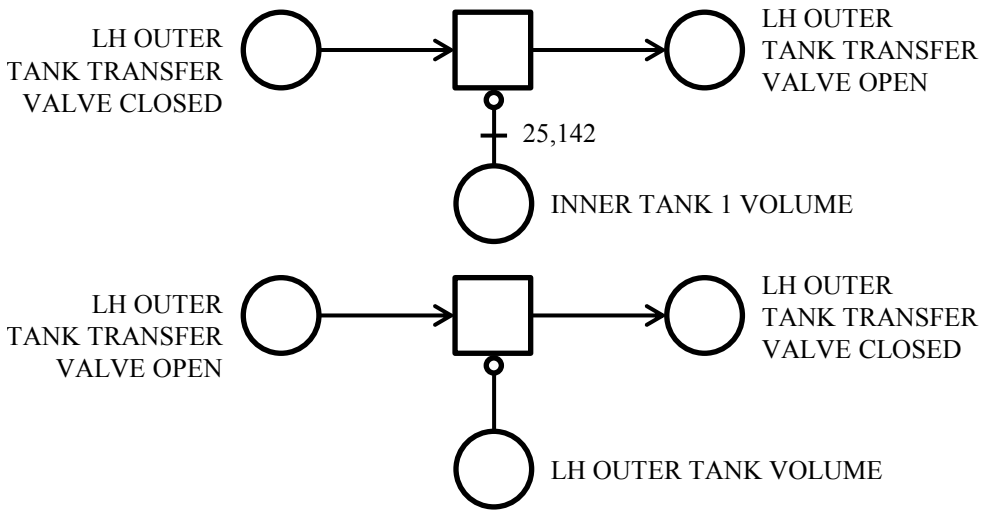

Figure 7.20: A340 petri net model outer tank fuel transfer 1/2

The upper sub-net in Figure 7.20 opens the LH outer tank transfer valve when the volume of fuel in inner tank 1 falls to 2,514L. An equivalent sub-net will open the RH outer tank transfer valve when the fuel volume in inner tank 4 falls to the same level. The second sub-net in the figure will close the transfer valve when the volume of fuel in the outer tank falls to zero. A similar sub-net models the same behaviour on the RHS of the system.

The sub-nets in Figure 7.21 model the flow of fuel from the LH outer tank to inner tanks 1 and 2. An equivalent set of sub-nets model the flow of fuel from the RH outer tank to inner tanks 3 and 4. When the LH outer tank transfer valve first opens, 52 fuel tokens are added to inner tank 1 every second and the same number are added to inner tank 2. The second sub-net in the figure models this behaviour. This sub-net requires that the inner tanks isolation valve, the valve linking the two tanks, be open. If the valve is closed all of the fuel leaving the outer tank will remain in inner tank 1, the bottom sub-net shows this behaviour. The top sub-net deals with the situation where there is not enough fuel tokens to add 52 to each of the inner tanks. When this situation occurs, it is assumed that inner tank 1 is favoured and receives all of the transferred fuel.

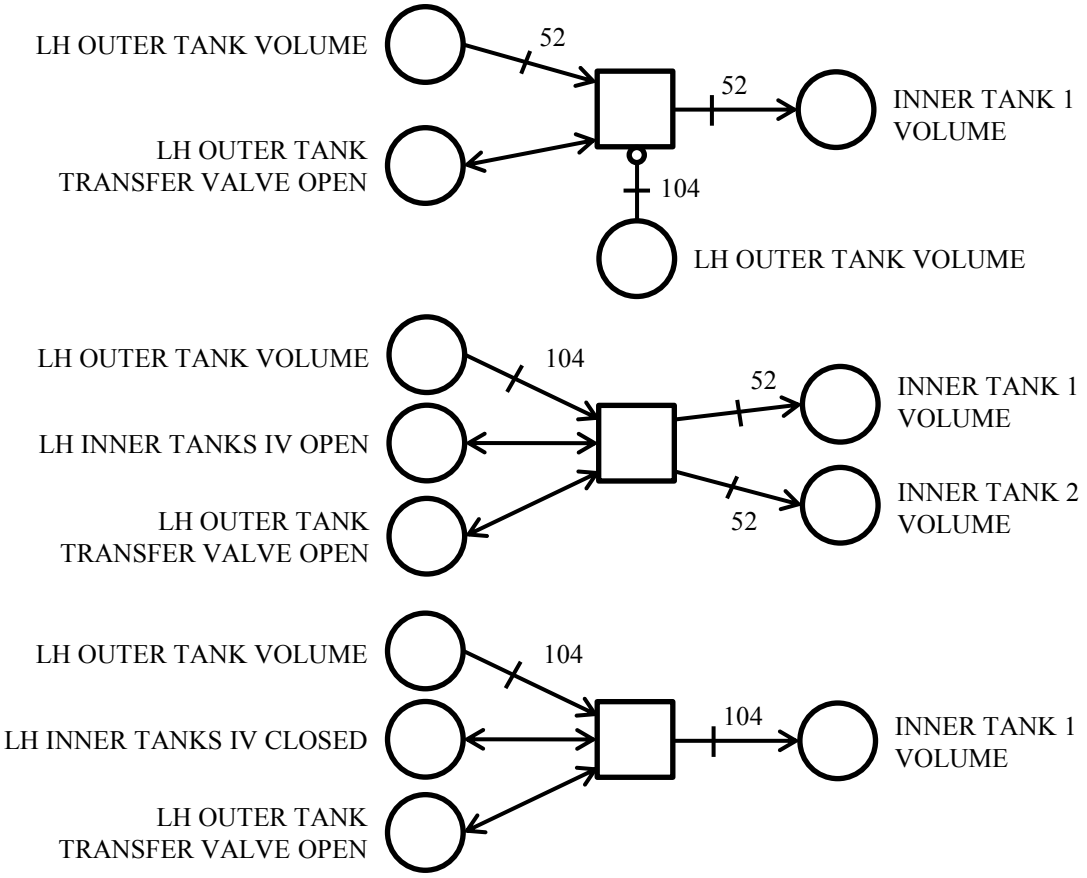

Figure 7.21: A340 petri net model outer tank fuel transfer 2/2

7.3.2.9 Fault Injection

Figure 7.22 shows part of the A340 PN model that would be used to inject a leak into the trim tank after 15,000 seconds. All faults injected into the model would be modelled in a similar fashion.

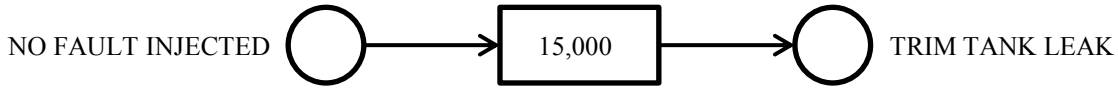

Figure 7.22: A340 petri net model fault injection

7.3.3 Model Accuracy

The A340 fuel system PN model was constructed using information from literature regarding the normal operating behaviour of the physical system [34]. The correctness of

the PN model was confirmed by replicating the system behaviour described in literature. In section 7.4.2 a phased mission is considered that is based on an actual A340 long range flight. The expected performance of the A340 system during this phased mission has been predicted using an aircraft analysis tool [35], which outputs fuel usage in each phase and overall. In the absence of official data to verify the A340 PN model, the outputs from this aircraft analysis tool have been used to confirm the accuracy of the PN model. A range of phased missions were considered that varied the mission duration and duration of individual phases. In every case, the PN model predicted behaviour matched well with that predicted by the aircraft analysis tool. This comparison has been used to verify the A340 PN model for the purpose of this work. Actual data from the A340 fuel system would be required to verify the PN model in order to use any results with confidence. Faults have also been modelled in the A340 PN. The effect of these faults on the system has been assumed to be similar to the respective effects on the fuel rig system. A more accurate representation of the effect of these faults would be required in a PN model that was to be used with the actual A340 fuel system.

7.4 Results

7.4.1 Phased Mission Description

The phased mission that has been identified for the purpose of modelling faults in the A340 fuel system is a typical long range flight that would be undertaken by an A340. The route is from New York's JFK International Airport to Dubai International Airport. The route is approximately 6,000 nautical miles and takes over 13 hours to complete. Table 8.3 shows the assumed duration of each phase in the flight.

7.4.2 Normal Operating Behaviour

The phased mission described above has been modelled using the Airbus A340 PN. Figure 7.23 shows the tank volumes predicted by the PN during the phased mission when no faults are present in the system. Only the tank volumes of inner tank 1 and collector cell 1 are shown on the figure, however the tank volume curves of all the other inner tanks and collector cells on the system are the same.

It can be seen in Figure 7.23 that from the start of the mission only the centre tank

Table 7.3: Airbus A340 phased mission phases

Taxi	5min
Take-Off	1min
Climb	19min
Cruise	731min (12h11m)
Descend	23min
Approach	3min
Taxi	5min
Total	787min (13h7m)

volume falls as it replaces any fuel used from the collector cells and inner tanks. At the start of the cruise phase the trim tank is also filled from the fuel in the centre tank. Once the centre tank is empty, the inner tank volumes fall until the forward transfer of fuel from the trim tank begins. Two forward transfers can be seen on the figure. Finally, towards the end of the mission the outer tanks empty of fuel, which is transferred to the inner tanks. At the end of the mission the collector cells remain full of fuel and the inner tanks contains approximately 3,500L of fuel. This is expected and represents the reserve fuel carried to deal with any unexpected diversions or queuing that may be experienced while waiting to land. Figures 7.24 and 7.25 show a selection of flow rates predicted by the PN of the A340 fuel system for the phased mission under consideration.

The flow rates from the main pump in collector cell 1 and at engine 1 can be seen to vary in Figure 7.24 as the system progresses through different phases of operation. The majority of the mission occurs at cruise and it can be seen that, in this phase, the flow rate is recorded at 0.50L/sec as specified in Table 8.2. The standby pump flow rate remains at zero throughout the mission as it is always off. Identical fuel rates have been predicted at equivalent locations in the other engine feed pipes. In Figure 7.25 the aft fuel transfer to the trim tank can be seen near the start of the mission and the two forward fuel transfers create spikes at approximately 8.5 hours into the mission.

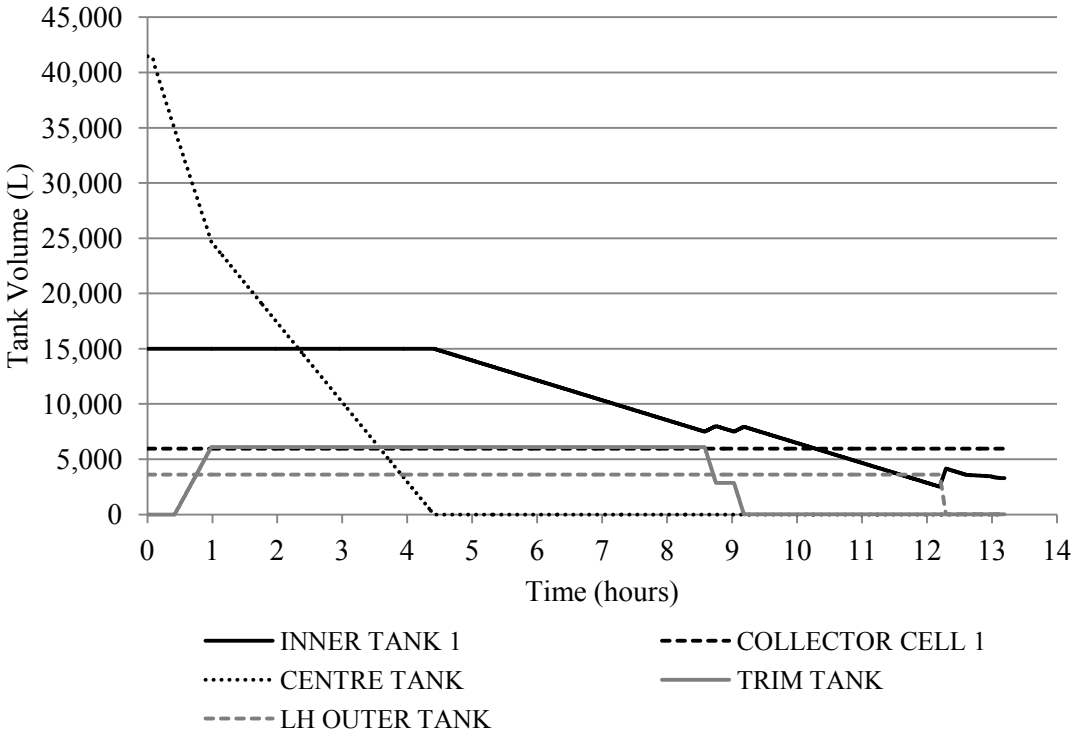

Figure 7.23: Fault-free A340 arrangement - Tank volumes

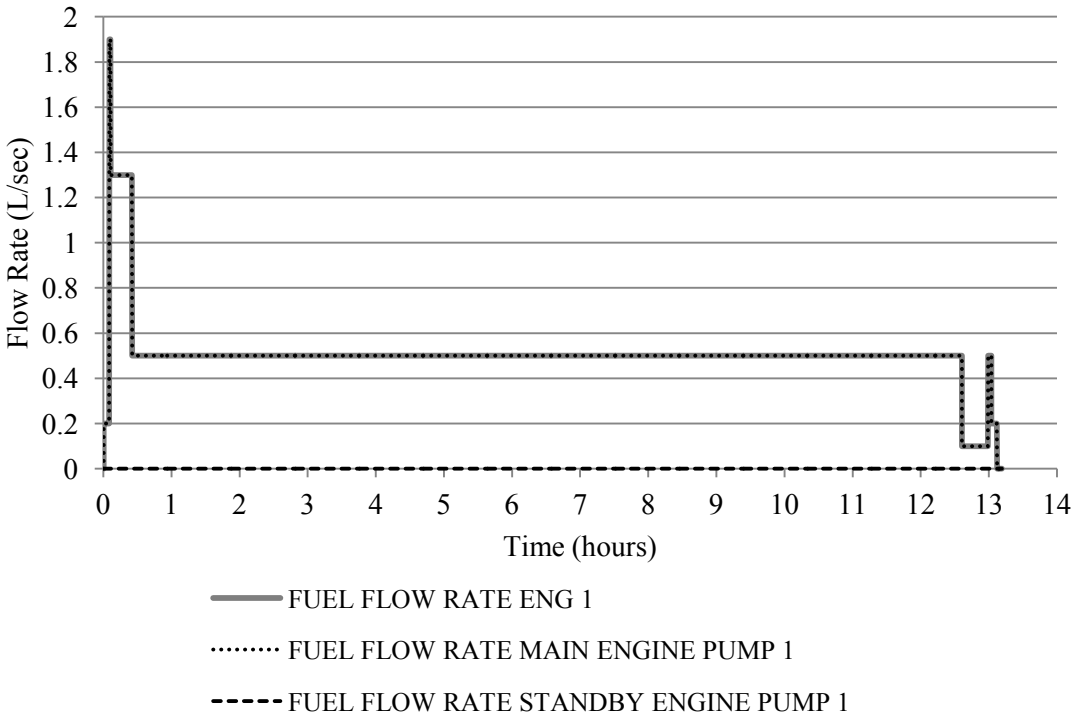

Figure 7.24: Fault-free A340 arrangement - Engine 1 flow rates

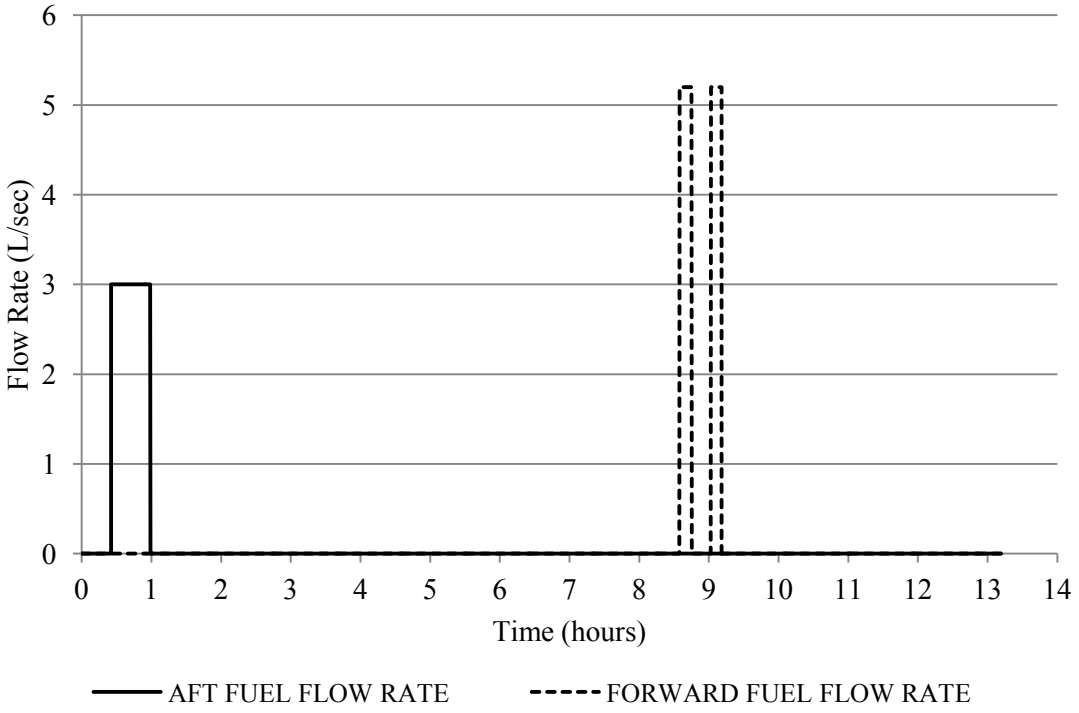

Figure 7.25: Fault-free A340 arrangement - Trim tank transfer flow rates

7.4.3 Collector Cell 1 Main Pump Failed Off

The first failure mode that will be considered is the failure of the main pump in collector cell 1. This fault was selected as there is redundancy in this part of the system which should result in the key measurements of system performance remaining unchanged, that is fuel reaches the engines in the amount demanded. The effect of the fault should still be seen in the flow rate sensors on the A340 PN model. The concept of redundancy was not exhibited in the fuel rig system and this fault therefore demonstrates the ability of the PN technique to model this capability.

If the main pump in collector cell 1 fails off, the standby pump, which is a cold backup, should come online to ensure engine number 1 continues to receive a supply of fuel. The tank volume curves should therefore be unaffected by the fault. Figure 7.26 shows the flow rates at the outputs of both collector cell 1 pumps and at the entrance to engine number 1. The fault is injected into the PN model after approximately 4 hours.

Figure 7.26 shows that the flow rate curve at the engine 1 entrance is the same as that shown in Figure 7.24. The main and standby pump curves, however, are different. It can be seen that after 4 hours the main pump flow rate falls from 0.5L/sec to zero. At the

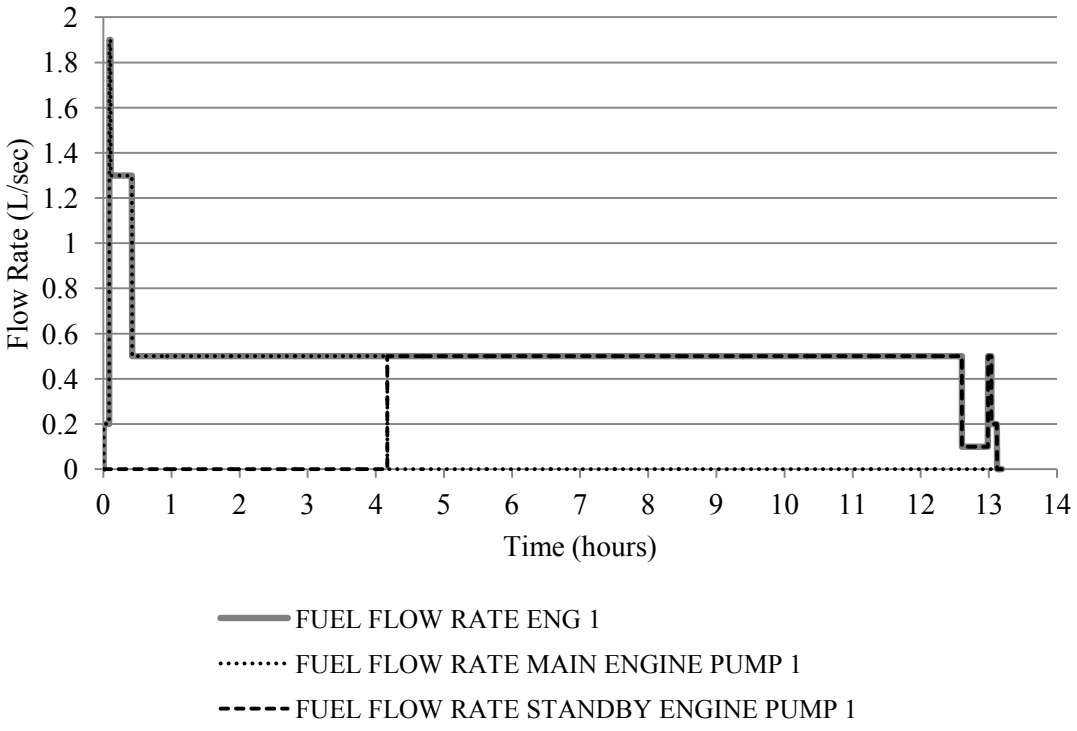

Figure 7.26: A340 collector cell 1 main pump fail off - Engine 1 flow rates

same time the standby pump flow rate increases from zero to 0.5L/sec. The model has assumed that the standby pump reacts immediately to the failure of the main pump, on a physical system it is likely that there would be a short delay. However, without data from the physical system to reference, it has been assumed that no delay is present. As the flow rate to engine 1 does not change as a result of the fault, the tank volume readings predicted by the PN model are the same as those shown in Figure 7.23.

The results shown in Figure 7.26 are in line with those that were predicted above when the effect of the pump failure was considered. This result shows that the A340 PN model is correctly modelling the effect of the pump failure, assuming the effects shown are accurate.

7.4.4 Trim Tank Leak

The next first order failure mode that will be considered is that of a large leak in the trim tank. This type of failure mode has been considered as it is a key fault on all fuel systems. This is especially true on aircraft where the loss of fuel through a leak could prevent the aircraft from successfully completing its mission. Considering a large leak will also ensure

the effect of the fault can be visibly identified from sensor outputs. The leak is assumed to be in the base of the trim tank, such that it would always cause a loss of fuel, if there is fuel present. The fault is again injected into the model after 4 hours, which is after the trim tank has been filled with fuel from the centre tank. The effect of the leak fault should be seen in both the tank volume and flow rate graphs. Figure 7.27 shows the tank volumes during the phased mission as predicted by the A340 PN model.

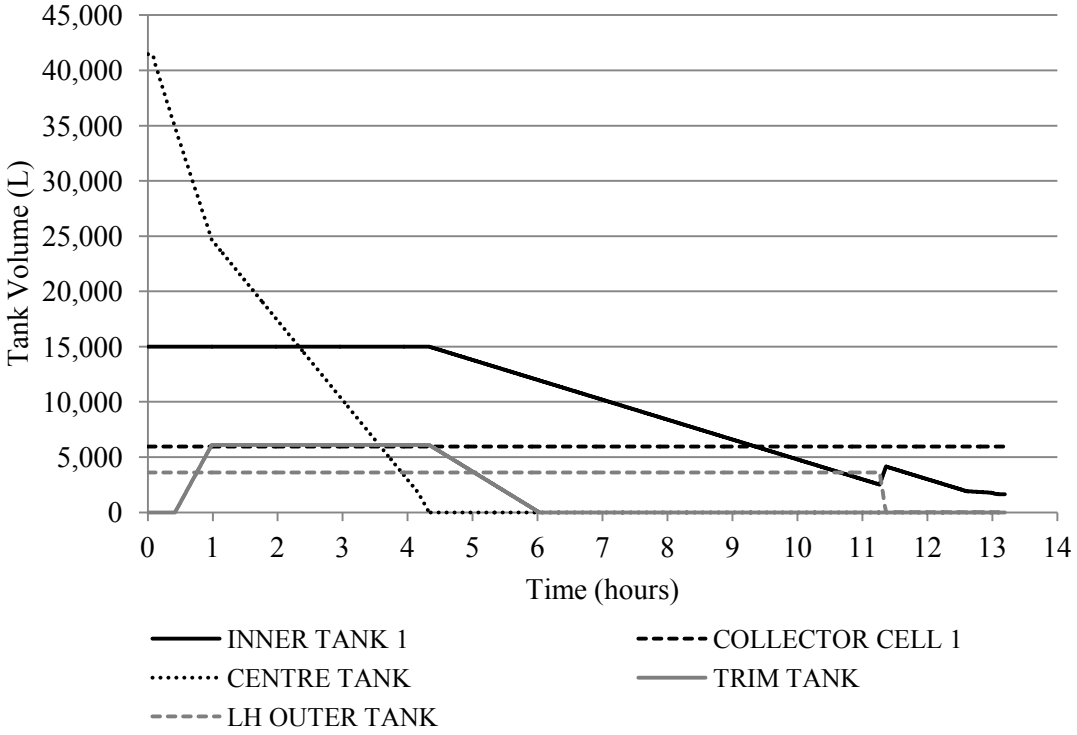

Figure 7.27: A340 trim tank leak - Tank volumes

Figure 7.27 shows that the leak fault has affected the behaviour of the centre tank, trim tank, outer tank and inner tank. Only the behaviour of the collector cell has been unaffected by the fault during the course of the mission. Once the leak fault occurs in the trim tank, fuel from the centre tank is sent to replenish the lost fuel. As a result, the centre tank empties earlier than when no fault is present. Once the centre tank can no longer replenish the trim tank, the leak fault causes the volume of the trim tank to fall to zero. With no fuel in the trim tank, the inner tank volume falls until it receives an input from the outer tank. This transfer occurs earlier than expected as a result of the leak, which has prevented any forward transfers of fuel from the trim tank. The volume of fuel in the inner tanks at the end of the mission, approximately 1,500L, is also lower

as a result of the leak fault. Nonetheless, the constant presence of fuel in the inner tanks means that the collector cells remain full throughout the mission. The other inner tanks and collector cells on the aircraft show identical behaviour to that seen in Figure 7.27. As both the inner tanks and collector cells have a volume of fuel stored at the end of the mission, all of the fuel flow rates demanded by the engines on the system were satisfied. The engine flow rate graph is therefore the same as that shown in Figure7.24.

The results of considering a leak in the trim tank of the Airbus A340 fuel system demonstrate that the PN technique can be used to effectively model leaks and their effects on large system. The success of this application assumes that the behaviour of the system in the presence of the leak has been accurately represented.

7.4.5 Collector Cell 1 Main Pump Failed Off and Engine 1 Pipe Section 3 Blocked

The situation where two faults occur on the Airbus A340 fuel system in the same mission will now be considered. The two faults are collector cell 1 main pump fails off and engine 1 pipe section 3 blocked. These faults have been chosen as they will create a flow of fuel in the cross feed section of the system. This will demonstrate the ability of the PN technique to model the flow of fuel in a system in multiple directions at the same time.

When the two faults occur simultaneously there will be no fuel supply to engine number 1. However, in order to maintain an even distribution of fuel in the system, fuel from collector cell 1 is pumped through the cross feed system back to the centre tank. Fuel is then supplied to all four inner tanks on the system. By transferring the fuel in this manner the fuel weight on the aircraft will be evenly distributed. It is assumed that the behaviour of the fuel system, in terms of the flow rate demands, does not change as a result of the faults. Figures 7.28 and 7.29 show the tank levels and flow rates on the fuel system as predicted by the PN model in the presence of the two faults. The main pump fault is injected after approximately 1 hour and the blockage fault is injected after 1.5 hours.

The behaviour of multiple variables shown on Figure 7.28, have changed as a result of the second order failure mode. In every case it is the blockage of the engine 1 feed pipe, which reduces the amount of fuel burned by the system, which has caused the respective behaviours to change. The centre tank volume falls at a lower rate as fuel is only burnt by three engines, with one engine's fuel consumption being recycled. As a result the time at

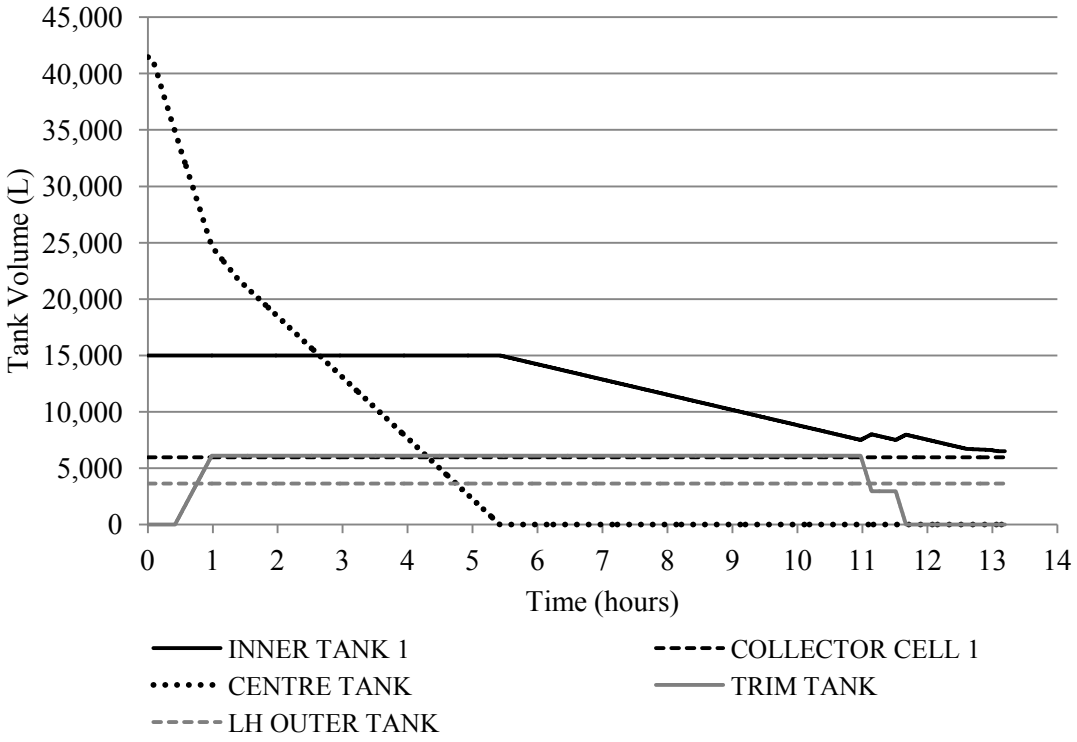

Figure 7.28: A340 collector cell 1 main pump fail off and engine 1 feed pipe blocked - Tank volumes

which the centre tank is emptied is delayed in comparison to the equivalent point on Figure 7.23. The trim tank curve also shows that the forward transfer of fuel is delayed compared to the behaviour seen in Figure 7.23. Again, this is due to the reduced fuel demands in the system. Figure 7.28 shows that the outer tank level stays constant throughout the mission. In Figure 7.23, the outer tanks were emptied after approximately 12 hours and therefore the blockage fault has also affected this variable. Finally, the inner tank level curve behaviour has also been affected by the fault. The tank remains full for a longer period of time as the emptying of the centre tank is delayed. Also, at the end of the mission it contains a larger volume of fuel as the overall system fuel demand is reduced.

The effect of both faults can be seen in Figure 7.29. Immediately prior to 1 hour, the curve illustrating the flow rate from the main collector cell pump falls to zero while the standby pump flow rate increases to 0.5L/sec at the same time. This concurrent decrease in the main pump flow rate and increase in the standby pump flow rate means that the flow rate to engine 1 does not change. However, when the blockage fault is injected after 1.5 hours, the flow rate to engine 1 falls to zero as fuel cannot reach the flow rate sensor or

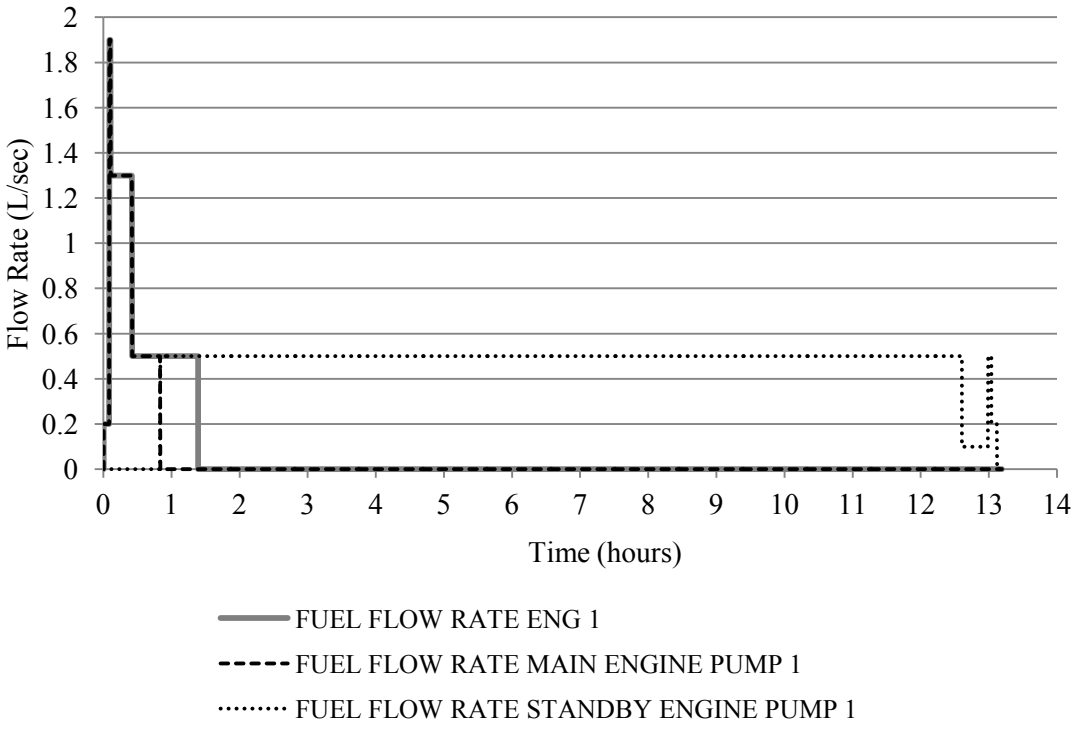

Figure 7.29: A340 collector cell 1 main pump fail off and engine 1 feed pipe blocked - Engine 1 flow rates

the engine. When the blockage fault occurs, the standby pump flow rate does not change as fuel is directed through the cross feed system to the centre tank for recycling. Figures 7.30 and 7.31 shows the other flow rates recorded from the A340 fuel system.

From Figure 7.30 it can be seen that the forward transfer of fuel from the trim tank first occurs after approximately 11 hours. This is delayed compared to the behaviour seen in Figure 7.25 where the forward transfer begins after approximately 9 hours. Finally, Figure 7.31 shows that from the time of the blockage fault onwards the fuel from the collector cell standby pump flows through the cross feed system. It enters the centre tank where it is immediately used to fill the inner tanks of the systems. As a result the centre tank remains empty from 5.5 hours onwards. It can be seen that the flow rates recorded from the standby system match those that would be expected from the engine flow rate sensor and reflect the demand applied to the collector cell pumps.

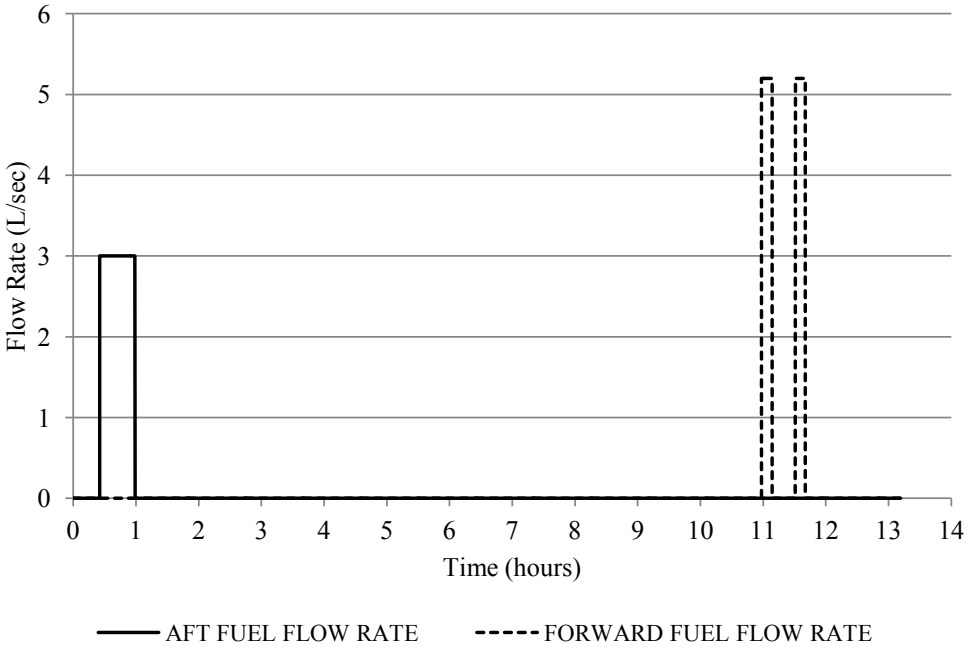

Figure 7.30: A340 collector cell 1 main pump fail off and engine 1 feed pipe blocked - Trim tank transfer flow Rates

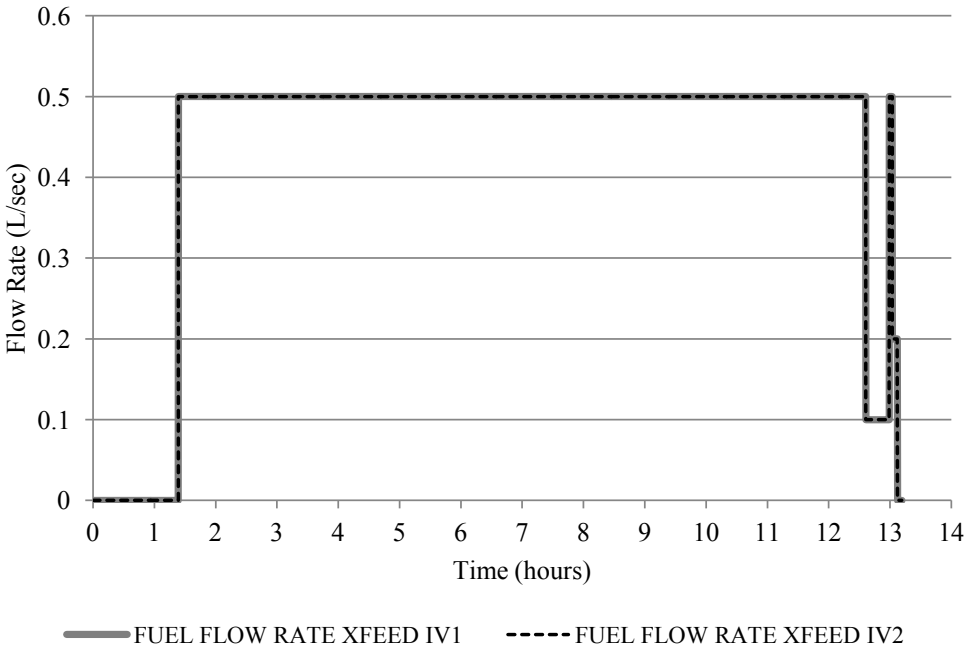

Figure 7.31: A340 collector cell 1 main pump fail off and engine 1 feed pipe blocked - Cross feed flow rates

7.5 Application to a Physical System

The work presented in this chapter has shown how the PN technique can be used to model the fuel system of a large commercial airliner. Using the PN model and the fault verification technique outlined in Chapter 4, the faults diagnosed on an Airbus A340 could also be assessed. However, several steps would have to be taken before the A340 PN model could be considered suitable for use with the physical system.

A complete PN model of the A340 fuel system, including all possible failure modes, would produce a PN of such size that it would present many challenges in terms of its construction and modification. It is likely therefore, that either a heirarchical approach would be required or the use of a coloured PN approach would be necessary in order to model the system accurately and with confidence. A heirarchical appoach would allow the fuel system to be broken down into smaller sections that could be easily checked and verified. It would also reduce the complexity with which the system is modelled. Alternatively using the coloured PN approach would decrease the size of the PN, however it would not necessarily reduce the complexity. In addition to using more advanced PN modelling approaches, the PN model of the A340 fuel system would have to account for a higher level of detail than it has been possible to consider in this work. The performance of the A340 fuel system will be affected by variables such as altitude, climate, gravitational forces and the aircraft's mass. None of these variables have been accounted for in the A340 PN model. The need to include the effect of these variables in the PN model further necessitates that the heirarchical/coloured PN approach is used in order to accurately account for this information. Another part of the PN model that lacks detail is the sensor types that have been considered. Due to limited available information on the A340 fuel system, only flow sensors and level sensors were modelled in the PN. On the actual A340 system there are many more sensor types including flow pressure, temperature and multiple types of level sensors. All of these sensor types would have to be modelled in order to accurately represent the effect of failure modes and the behaviour of the system. The range of failure modes that must be included in the A340 PN model must also be expanded to cover all possible faults that could be experienced by the system. While the work done with the fuel rig provides a basis for the behaviour of the system in the presence of some faults, data from the A340 fuel system would be required in order to accurately model these faults in the PN. Data from the A340 fuel system would also be required to indicate

the subtle component interaction effects that are often present in systems and would have to be included in the PN model. Examples of these interactions include variable start-up times, how the system deals with redundancy and any effect on the system created by electronic or mechanical components.

Once the detail with which the A340 system has been increased in the PN model, it would be necessary to undertake a period of thorough testing on a mock-up or test version of the system. The fuel rig provided a means of physically testing the PN model, and fault verification technique, and a similar system would be required to test the A340 PN model. This testing would verify the PN model of the A340 fuel system, in the same way that the fuel rig system was used to verify the fuel rig PN. The testing process would have to replicate the behaviour of the system in any and all operating conditions. All of the failure modes that could affect the system would also have to be replicated in order to ensure that the PN model is correctly representing their effects. The fault verification aspect of the technique would also have to be tested at this stage to ensure that faults are being correctly verified where genuine, and filtered where false. Once the PN model and fault verification capability has been verified in the testing phase, the technique can be considered for integration into the A340 aircraft. Integration of new sub-systems is a significant investment in both time and money for a company. For this reason, undertaking a thorough testing regime is critical to enable a smooth integration phase. Integration includes testing the fault verification capability with the fuel system on each aircraft and with the higher level aricraft systems. It also includes physically fitting the new equipment onto the aircraft. There is a risk that the verification software will not work as designed or expected if it is not integrated correctly.

This section has provided an overview of how the A340 PN model could be developed to make it suitable for use with the A340 aircraft. The development of the PN model can be broken down into three stages; increased detail, testing and integration. Sucessfully completing each of these stages would be necessary in order to use the A340 PN model and fault verification software with confidence.

7.6 Conclusion

The aim of this chapter was to demonstrate the applicability of the PN technique to a large scale system. The Airbus A340 fuel system has been used for this purpose. The

system was modelled using the PN technique and verified, with no faults present, using performance data from a variety of literature sources and an aircraft analysis tool. A number of failure modes were also modelled in the PN and their effects on the system were assumed from the equivalent fault effects on the fuel rig. These faults were then propagated using the PN software. The result of propagating two first order failure modes and a second order failure mode showed that the PN model had predicted the expected system behaviour. The PN predicted system behaviour was measured using tank volume and flow rate outputs at the chosen sensor locations. In the case of every failure mode the behaviour that was expected of the fuel system was also seen in the PN outputs.

The A340 model considers more complex behaviour than the fuel rig PN model as it includes a sophisticated fuel feed arrangement and also multi-directional fuel transfers. The A340 PN model also accounts for the fuel weight distribution in the system, a feature not accounted for in the fuel rig PN model. While the A340 PN model has demonstrated that large systems could be modelled using the PN technique, the current PN model is of limited value due to its simplicity. As only limited information is available with which to model the fuel system, several aspects of the system design have been omitted. This includes sensor types, factors which affect the operational behaviour of the system (altitude, climate, etc) and the majority of faults that could occur in the system. Section 7.5 presented a detailed look at how the current PN model could be developed for use with the A340 aircraft. However, in its current form the PN model can only provide a simplified example of how the PN technique can be used to model large scale systems and propagate faults through the system.

The A340 PN model, in addition to the fault verification technique developed previously, could be used on a physical version of the A340 fuel system to verify faults. However, in order to be used effectively, the PN model would require several developments. The detail of the model would have to be increased to take account of environmental effects, a greater number of failure modes and more sensor types. The set-up would also have to be suitably tested. Developing the fault verification technique to operate in real time would also significantly improve the capability of the overall technique to verify faults.

CHAPTER 8

Sensor Value Calculation and Sensor Selection

8.1 Introduction

The PN software described in Chapter 6 and used to verify arisings from the fuel rig system, could also be used in the process of system design. By modelling potential system designs using PN models and then propagating failures through the model, the behaviour of the system due to these failures can be predicted. Of particular interest at this stage is to identify where to place sensors in the system. Sensors must be carefully positioned to enable system faults to be identified and diagnosed effectively. However, a designer must also take into consideration limitations such as cost, weight and space. The value of a sensor and its location in a system can depend on what sensor information is the most desirable. The sensors that identify the greatest number of failure modes, for example, may be different from the sensors that identify the failure modes which will have the greatest effect on the system. The aim of this chapter, therefore, is to assess the value of sensors at different locations on the system and provide some guidance as to how such information can support design decisions.

8.2 Assessing Sensor Value

8.2.1 Generic Deviation of Output Measured by Sensor

The value of a sensor, in the diagnosis of faults on a system, is directly related to how the sensor responds when a fault occurs. If a sensor output changes when a fault occurs, such information can be useful in diagnosing the fault. By comparison, if the sensor output shows no change when a fault occurs, the sensor is of less value. If no other sensor outputs change as a result of the fault, it would exist without the knowledge of a system operator. Such situations should be avoided, as the fault could change the behaviour of the system unexpectedly or the effects of the fault could become more severe, thereby increasing the risk of a serious incident occurring. A sensor, whose output changes as a result of a fault, is therefore more valuable than a sensor whose output does not change.

The first stage in assessing the value of a sensor is to determine whether its output will vary as a result of a failure occurring in a system. Consider the fuel rig system from the wing tanks to the engine tanks, as shown in Figure 8.1. The system has four pipe sections, six isolation valves, two pumps and four tanks. Fourteen possible flow rate sensor locations have been identified, denoted as Fi, and the predicted output at these sensor locations will be considered in the presence of the failure modes listed in the first column of Table 8.1. There are no flow sensor locations or faults modelled in the cross feed pipes. The fuel rig system and the flow rate sensors at the proposed locations have been modelled using the PN technique.

If the presence of a fault in the system does not change the sensor output, a value of '0' is shown in the respective cell of the table. If the fault causes the sensor output to immediately deviate, a value of '2' is shown. If the fault could have a delayed effect on the sensor output, dependent on the mission profile and time of the fault occurring, a value of '1' is shown in the cell. The expected system behaviour is flow from the LH wing tank to the LH engine tank and from the RH wing tank to the RH engine tank.

Table 8.1: Effect of failure modes on flow rate sensors F1-F7

Failure Mode	F1	F2	F3	F4	F5	F6	F7
Pipe W Blocked	2	2	2	2	2	2	2
Pipe X Blocked	2	2	2	2	2	2	2

Continued on next page

Failure Mode	F1	F2	F3	F4	F5	F6	F7
Pipe Y Blocked	0	0	0	0	0	0	0
Pipe Z Blocked	0	0	0	0	0	0	0
Valve 1 Blocked	2	2	2	2	2	2	2
Valve 2 Blocked	2	2	2	2	2	2	2
Valve 3 Blocked	2	2	2	2	2	2	2
Valve 4 Blocked	0	0	0	0	0	0	0
Valve 5 Blocked	0	0	0	0	0	0	0
Valve 6 Blocked	0	0	0	0	0	0	0
Pipe W Leak	0	2	2	2	2	2	2
Pipe X Leak	0	0	0	2	2	2	2
Pipe Y Leak	0	0	0	0	0	0	0
Pipe Z Leak	0	0	0	0	0	0	0
Pipe W Rupture	2	2	2	2	2	2	2
Pipe X Rupture	2	2	2	2	2	2	2
Pipe Y Rupture	0	0	0	0	0	0	0
Pipe Z Rupture	0	0	0	0	0	0	0
LH Engine Pump Fail Off	2	2	2	2	2	2	2
RH Engine Pump Fail Off	0	0	0	0	0	0	0
LH Engine Pump Degraded	2	2	2	2	2	2	2
RH Engine Pump Degraded	0	0	0	0	0	0	0
LH Wing Tank Leak	1	1	1	1	1	1	1
RH Wing Tank Leak	0	0	0	0	0	0	0
LH Engine Tank Leak	0	0	0	0	0	0	0
RH Engine Tank Leak	0	0	0	0	0	0	0

Table 8.2: Effect of failure modes on flow rate sensors F8-F14

Failure Mode	F8	F9	F10	F11	F12	F13	F14
Pipe W Blocked	0	0	0	0	0	0	0

Continued on next page

Failure Mode	F8	F9	F10	F11	F12	F13	F14
Pipe X Blocked	0	0	0	0	0	0	0
Pipe Y Blocked	2	2	2	2	2	2	2
Pipe Z Blocked	2	2	2	2	2	2	2
Valve 1 Blocked	0	0	0	0	0	0	0
Valve 2 Blocked	0	0	0	0	0	0	0
Valve 3 Blocked	0	0	0	0	0	0	0
Valve 4 Blocked	2	2	2	2	2	2	2
Valve 5 Blocked	2	2	2	2	2	2	2
Valve 6 Blocked	2	2	2	2	2	2	2
Pipe W Leak	0	0	0	0	0	0	0
Pipe X Leak	0	0	0	0	0	0	0
Pipe Y Leak	0	2	2	2	2	2	2
Pipe Z Leak	0	0	0	2	2	2	2
Pipe W Rupture	0	0	0	0	0	0	0
Pipe X Rupture	0	0	0	0	0	0	0
Pipe Y Rupture	2	2	2	2	2	2	2
Pipe Z Rupture	2	2	2	2	2	2	2
LH Engine Pump Fail Off	0	0	0	0	0	0	0
RH Engine Pump Fail Off	2	2	2	2	2	2	2
LH Engine Pump Degraded	0	0	0	0	0	0	0
RH Engine Pump Degraded	2	2	2	2	2	2	2
LH Wing Tank Leak	0	0	0	0	0	0	0
RH Wing Tank Leak	1	1	1	1	1	1	1
LH Engine Tank Leak	0	0	0	0	0	0	0
RH Engine Tank Leak	0	0	0	0	0	0	0

Table 8.1 shows that the majority of faults that occur on the LH side of the system have an immediate effect on the flow rate sensors F1 - F7. The only exceptions to this are the pipe W leak, pipe X leak and LH wing tank leak faults. If pipe W or X experiences a leak, the effect will only be seen in the flow rate sensors downstream of the fault. The

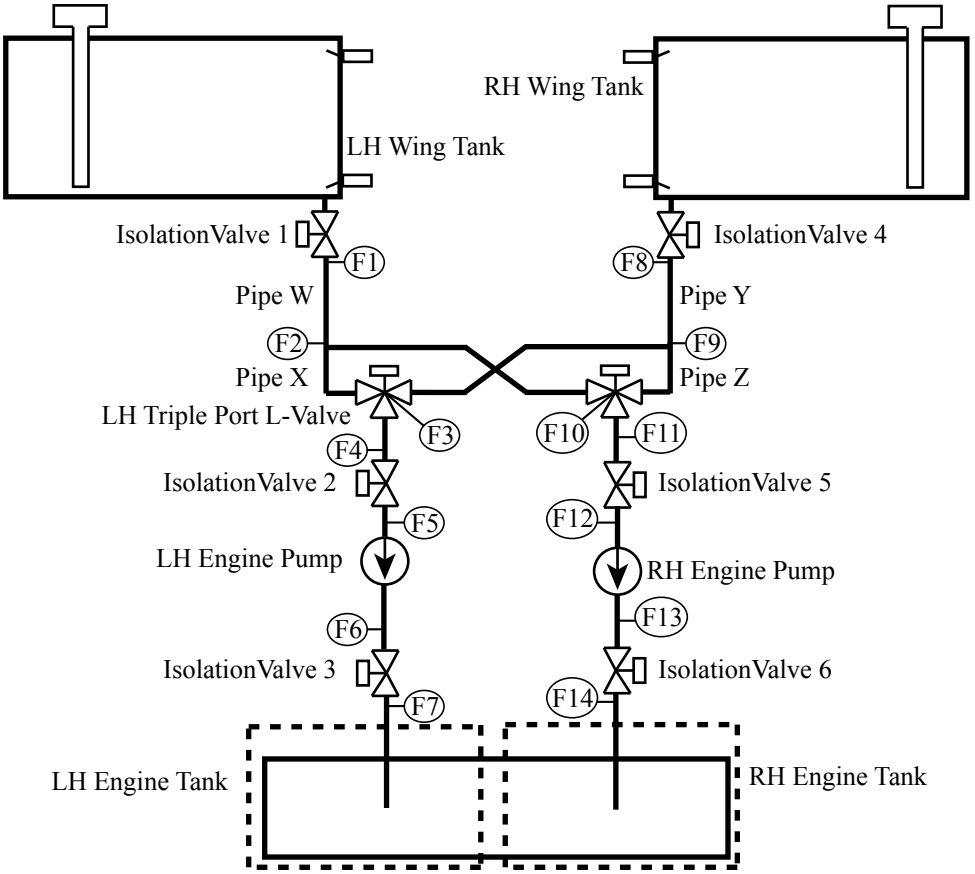

Figure 8.1: Fuel rig system with proposed flow rate sensor locations

output from the upstream sensors, F1 (F2 and F3) will not change. In the case of the LH wing tank leak, the effect of this fault will only be seen in the sensor outputs if the tank empties of fuel before the end of the mission. As the effect of the fault will not be immediate, a value of 1 is shown in the respective cells. None of the faults that occur on the RH side of the system have an effect on the flow rate sensors F1 - F7. The results of Table 8.2 show a similar pattern to that seen in Table 8.1. As would be expected, only faults from the RH side of the system have affected the output of sensors F8 - F14. The only faults that do not have an immediate effect on sensors F8 - F14 are pipe Y leak, pipe Z leak and RH wing tank leak. The reasons for this are the same as those identified for the equivalent faults on the LH side of the system.

Tables 8.1 and 8.2 provide a system designer with a high level analysis of which sensors will deviate as a result of faults occurring in the fuel rig system. These results would be of particular value if the design requirements specified knowledge of which sensors exhibited the greatest number of fault induced deviations. Twenty-four of the twenty-six deviations

could be identified by selecting one flow rate sensor from F4, F5, F6 or F7 as well as one from the group of F11, F12, F13 and F14. Selecting F4 and F11, for example, twenty-four of the twenty-six faults would cause one of the sensor outputs to deviate. Table 8.1 shows that the F4 output deviated as a result of twelve different failure modes while Table 8.2 shows that the output from F11 deviated as a result of twelve further unique failure modes. It should be noted that the sensors in the first group listed above are all located below the triple port L-valve on the LH side of the system while all of the sensors in the second group are located below the triple port L-valve on the RH side of the system. As no pipe leaks are considered downstream of these valves, the output from the sensors in each group is the same irrespective of the leak or sensor location considered. By contrast, it can also be seen in the above tables that upstream of the triple port L-valves, the sensor output recorded is dependent on the sensor and leak location. It would not be possible to identify all twenty-six failure modes using any combination of the flow rate sensors as the engine tank leak faults are downstream of all the sensor locations considered and do not change any of the sensor outputs when present.

The above work has shown how the value of a sensor can be simply assessed. However, as no account is taken of the type of sensor deviation that occurs, little information is available to aid in the diagnosis of the specific fault that has occurred.

8.2.2 Particular Deviation of Output Measured by Sensor

By considering the type of deviation that a failure mode causes in a sensor output, more information can be extracted and used to define sensor value. A sensor whose output falls to zero in the presence of all failure modes, for example, offers limited help to the fault diagnosis process. In this scenario, while it is clear that a failure is present, it is impossible to identify which fault has occurred. However, a sensor that is affected by failure modes in different ways can reduce the possible number of true failure modes thereby making the diagnosis process more efficient by reducing the potential number of failure modes that have to be considered for further evaluation.

Consider again the fuel rig system shown in Figure 8.1. On this occasion the flow paths in the system are from the LH and RH wing tanks to the RH engine tank. The failure modes considered are the same as those listed in the first column of Table 8.1. Table 8.3 shows the sensor outputs from the F1 and F10 sensors in the presence of each of the

failure modes. Three possible outputs have been identified; the flow rate sensor value will stay constant (N), fall to zero (Z) or fall to a value greater than zero (L). A value of '1' indicates whether a particular scenario occurs.

Table 8.3: Sensor deviations of F1 and F10

	F1			F10		
Failure Mode	N	L	Z	N	L	Z
Pipe W Blocked	0	0	1	0	1	0
Pipe X Blocked	1	0	0	1	0	0
Pipe Y Blocked	1	0	0	0	1	0
Pipe Z Blocked	0	0	1	0	0	1
Valve 1 Blocked	0	0	1	0	1	0
Valve 2 Blocked	1	0	0	1	0	0
Valve 3 Blocked	1	0	0	1	0	0
Valve 4 Blocked	1	0	0	0	1	0
Valve 5 Blocked	0	0	1	0	0	1
Valve 6 Blocked	0	0	1	0	0	1
Pipe W Leak	1	0	0	0	1	0
Pipe X Leak	1	0	0	1	0	0
Pipe Y Leak	1	0	0	0	1	0
Pipe Z Leak	1	0	0	1	0	0
Pipe W Rupture	0	0	1	0	1	0
Pipe X Rupture	1	0	0	1	0	0
Pipe Y Rupture	1	0	0	0	1	0
Pipe Z Rupture	0	0	1	0	0	1
LH Engine Pump Fail Off	1	0	0	1	0	0
RH Engine Pump Fail Off	0	0	1	0	0	1
LH Engine Pump Degraded	1	0	0	1	0	0
RH Engine Pump Degraded	0	1	0	0	1	0
LH Wing Tank Leak	0	0	1	0	1	0

Continued on next page

	F1			F10		
Failure Mode	N	L	Z	N	L	Z
RH Wing Tank Leak	1	0	0	0	1	0
LH Engine Tank Leak	1	0	0	1	0	0
RH Engine Tank Leak	1	0	0	1	0	0
Total	16	1	9	10	11	5

Table 8.3 shows that by considering the sensor output deviation shape, it is possible to reduce the potential number of failure modes that have occurred and caused such a deviation. If the F10 output falls to zero, for example, it is known that one of five possible failure modes has occurred; pipe 5 blocked, valve 5 blocked, valve 6 blocked, pipe Z rupture or pump 2 failed off. By comparison, if the type of deviation was not considered, sixteen possible failures would have to be considered; eleven 'fall to a value greater than zero' plus five 'fall to zero'. If the F1 sensor falls to a value greater than zero, then only the fault 'pump 2 degraded' can be present in the system. However, the results in the table also show that when the F1 sensor outputs are considered, there are sixteen hidden failure modes, while there are also ten hidden failure modes when F10 is considered. Table 8.4 shows the number of failure modes that cause the respective sensor deviations to occur for all of the sensor locations on the system and the resultant sensor ranking.

Table 8.4: Summary of sensor deviations

Flow Rate Sensor	N	L	Z	Rank
F1	16	1	9	=8
F2	15	2	9	=6
F3 - F7	26	0	0	=10
F8	16	1	9	=8
F9	15	2	9	=6
F10	10	11	5	5
F11 - F14	9	12	5	=1

Table 8.4 shows the highest ranked sensors, those with the greatest value, are F11 -

F14. These are the highest ranked sensors as they identify the greatest number of failure modes, i.e. the number of hidden failures is lowest - nine. These sensors are located at the end of the active flow paths in the system and therefore any faults that propagate from upstream have an effect on these sensors. The lowest value sensors are those from F3 - F7. As no flow passes by these sensors their output never changes from that expected and all twenty-six faults are hidden. Table 8.4 also shows that the F1 and F8 sensors have an identical distribution of sensor deviations. The same observation can be made of the outputs at the F2 and F9 locations. This result is expected given the system arrangement and the symmetrical positioning of the respective sensors.

While the F11 - F14 sensors recorded a form of deviation as a result of seventeen failure modes, there remain nine hidden failures. Given the success of combining sensor outputs to identify a greater number of deviations in Section 8.2.1, it is logical to consider the result of evaluating sensor pairs to reduce the number of hidden faults predicted. From Section 8.2.1, it was shown to be effective to consider a sensor from each side of the system when selecting sensor pairs. Using the results the Table 8.4 as a guide, the following sensor pairs will be considered; F1 and F8, F1 and F9, F1 and F10, F1 and F11. Pairs that include F2 and F3 in place of F1 are also considered. Twelve sensor combinations are therefore evaluated. Every pair of sensors will have nine possible outputs, as each sensor outputs will be 'N', 'Z' or 'L'. Table 8.5 shows the results of evaluating sensors F1 and F10 as a pair. In the column header the first letter represents the output from F1 and the second letter indicates the F10 output.

Table 8.5: Sensor deviations of F1 and F10 sensor pairs

Failure Mode	N,N	N,L	N,Z	L,N	L,L	L,Z	Z,N	Z,L	Z,Z
Pipe W Blocked	0	0	0	0	0	0	0	1	0
Pipe X Blocked	1	0	0	0	0	0	0	0	0
Pipe Y Blocked	0	1	0	0	0	0	0	0	0
Pipe Z Blocked	0	0	0	0	0	0	0	0	1
Valve 1 Blocked	0	0	0	0	0	0	0	1	0
Valve 2 Blocked	1	0	0	0	0	0	0	0	0
Valve 3 Blocked	1	0	0	0	0	0	0	0	0
Valve 4 Blocked	0	1	0	0	0	0	0	0	0

Continued on next page

Failure Mode	N,N	N,L	N,Z	L,N	L,L	L,Z	Z,N	Z,L	Z,Z
Valve 5 Blocked	0	0	0	0	0	0	0	0	1
Valve 6 Blocked	0	0	0	0	0	0	0	0	1
Pipe W Leak	0	1	0	0	0	0	0	0	0
Pipe X Leak	1	0	0	0	0	0	0	0	0
Pipe Y Leak	0	1	0	0	0	0	0	0	0
Pipe Z Leak	1	0	0	0	0	0	0	0	0
Pipe W Rupture	0	0	0	0	0	0	0	1	0
Pipe X Rupture	1	0	0	0	0	0	0	0	0
Pipe Y Rupture	0	1	0	0	0	0	0	0	0
Pipe Z Rupture	0	0	0	0	0	0	0	0	1
LH Engine Pump Fail Off	1	0	0	0	0	0	0	0	0
RH Engine Pump Fail Off	0	0	0	0	0	0	0	0	1
LH Engine Pump Degraded	1	0	0	0	0	0	0	0	0
RH Engine Pump Degraded	0	0	0	0	1	0	0	0	0
LH Wing Tank Leak	0	0	0	0	0	0	0	1	0
RH Wing Tank Leak	0	1	0	0	0	0	0	0	0
LH Engine Tank Leak	1	0	0	0	0	0	0	0	0
RH Engine Tank Leak	1	0	0	0	0	0	0	0	0
Total	10	6	0	0	1	0	0	4	5

Table 8.5 shows that by evaluating the F1 and F10 sensor outputs as a pair, the number of hidden failures, shown in the first results column, is ten - the same number present when the F10 output is considered on its own. The hidden failure modes are also the same. Evaluating these sensors as a pair does therefore not offer an improvement in terms of being able to identify a larger number of faults present in the system. However, of the sixteen revealed failures, the maximum number of possible faults that could cause any one set of sensor outputs is six, which occurs when the F1 output is constant and the F10 output falls to a value greater than zero - the second results column. Three other unique sensor outputs are produced and these can be caused by one, four and five

failure modes respectively - as shown in the fifth, eighth and ninth columns respectively. F10 was able to identify the same number of failure modes when considered individually, however, on average a greater number of failure modes would have to be evaluated in the diagnosis process. Combining sensor pairs is therefore of greater overall value compared to using individual sensor outputs, as fewer failure modes would have to be evaluated thereby making the diagnosis process more efficient.

By evaluating the remaining sensor pairs in the same way as was shown in Table 8.5, it will be possible to rank the sensor pairs in terms of their value. The primary factor by which the sensor pairs will be assessed is the number of faults which can be identified as a result of a sensor deviation, i.e. the revealed failures. The higher the number of faults that cause at least one sensor deviation, the higher the value of the sensor pair. The SD of the revealed faults will then be used to rank the sensor pairs. Considering Table 8.5, the SD is measured using the values in the 'Total' row excluding the first column of hidden failures. In the case of Table 8.5, a SD value of 2.56 is calculated. A small SD value can be considered advantageous, as it will indicate that a small number of possible failure modes are associated with each set of sensor outputs thereby improving the efficiency of the diagnosis process. Finally the number of unique sensor outputs (USOs) will be used to rank the results. A USO will be present if a set of sensor outputs can be caused by only one failure mode. Column five in Table 8.5 shows an example of a USO. The greater the number of USOs, the more efficient the diagnosis process will be and therefore the greater value a sensor pair has. Table 8.6 shows the sensor pairs as ranked using the above method. The F11 sensor is also included on its own as, from Table 8.4, it identified the greatest number of faults of all the single sensor arrangements.

Table 8.6 shows that three sensor pairs and sensor F11 identified seventeen revealed failure modes. For that reason, they are the top ranked results. All three of the top ranked sensor pairs also include F11. This suggests that if a system designer was only interested in knowing that a fault had occurred, installing F11 would offer the same level of detail as any of the top ranked sensor pairs. However by using F1/F2 and F11 together, as opposed to F11 on its own, the fault diagnosis process could be made more efficient as there would be fewer potential faults to investigate. This conclusion can be made as the SD values of the top ranked sensor pairs are lower than that of F11 individually. As the sensor pair F2 and F11 can identify seventeen failure modes and it has the lowest SD value, 2.53, it is the top ranked result by sensor value.

Table 8.6: Sensor value rankings

Sensor(s)	No. of Revealed Faults	SD	USOs	Rank
F2 and F11	17	2.53	0	1
F1 and F11	17	2.80	1	2
F3 and F11	17	4.36	0	=3
F11	17	4.36	0	=3
F2 and F9	16	2.00	3	5
F2 and F10	16	2.33	0	6
F1 and F10	16	2.56	1	7
F3 and F10	16	4.04	0	8
F1 and F9	15	2.10	2	=9
F2 and F8	15	2.10	2	=9
F1 and F8	14	2.19	1	11
F3 and F9	11	3.16	0	12
F3 and F8	10	3.15	1	13

The results in Table 8.6 also show that the fewest faults are identified by the sensor pairs where one sensor, F3, is located in a section of the system where no flow is present. The lowest ranked sensor pairs identified ten and eleven failure modes respectively. As the F3 sensor output does not change, these sensor pairs are therefore reliant on deviations being recorded from the F8 and F9 sensors respectively. Given the limited data available to these sensor pairs, it would be expected that they perform poorly compared to sensor pairs where both sensors are located in areas of the system where flow is present.

If a system designer was primarily interested in being able to immediately identify which failure mode had occurred in the system, the number of USOs should be used as the initial ranking factor. Given the results shown in Table 8.6, the sensor pair F2 and F9 would be the top result as three USOs are recorded.

8.2.3 Measured Sensor Deviation with Weighted Failure Mode Effects

Faults can affect systems in different ways and the effect of one fault may be greater than the effect of another. A pipe leak, for example, will have a smaller effect on a system than

a pump failing off. Sensors that can identify faults which have a greater effect on a system are therefore somewhat advantageous. This section will explain how the magnitude of a fault's effect on a system can be included in the calculation of sensor value.

Assessing the effect of a fault on a system is an objective process that will require a level of system knowledge from a designer. The faults considered for the fuel rig system have been assessed a value between 1 and 5, where 5 indicates the highest level of effect. Table 8.7 shows the effect values assigned to the faults being considered on the fuel rig.

Table 8.7: Sensor deviations of F1 and F10

Failure Mode	Effect Value
Pipe Blockage	3
Valve Blockage	3
Pipe Leak	1
Pipe Rupture	4
Pump Fail Off	4
Pump Degraded	2
Wing Tank Leak	5
Engine Tank Leak	1

Pipe and valve blockages have been given an effect value of 3, as a blockage could prevent flow reaching one, or both, of the engines but the system arrangement could be changed to allow the mission to be completed, i.e. flow could be delivered from the alternative side. A pipe leak has an effect value of 1, as the system can still operate in the same arrangement albeit at a reduced level of performance. The pipe rupture and pump failed off faults have an effect value of 4, as the number of alternative system arrangements will be limited, which could prevent the mission from being completed. An effect value of 2 is given to a pump degradation fault to reflect the minor effect on the system, but also its potential to degrade further. A leak in the LH/RH wing tank has an effect value of 5 to represent the worst case scenario, where there is insufficient fuel to complete the mission irrelevant of the system arrangement. Finally, a LH/RH engine tank leak is given an effect value of 1 as fuel has reached the engine and only small amounts of fuel could be lost.

The process of analysing the sensor outputs in terms of the deviation shape is first

applied in the same manner as described in Section 8.2.2. However, instead of using a '1' to indicate whether the sensor shape is present, the effect value is used. Table 8.8 shows how the pipe blockage rows now appear, when there are flow paths from the LH and RH wing tanks to the RH engine tank.

Table 8.8: Sensor deviations of F1 and F10

	F1			F10		
Failure Mode	N	L	Z	N	L	Z
Pipe W Blocked	0	0	3	0	3	0
Pipe X Blocked	3	0	0	3	0	0
Pipe Y Blocked	3	0	0	0	3	0
Pipe Z Blocked	0	0	3	0	0	3

The extended version of Table 8.8 therefore shows, which faults cause a deviation in which sensor outputs and the effect value of these faults. Evaluating this data includes determining the same variables as were listed in Table 8.6, namely the number of faults that cause a sensor deviation, the SD of the revealed failures and the number of USOs. The ability of each sensor and sensor pair to identify high effect failure modes is also determined. A high effect failure mode will be considered as one that has an effect value of 4 or 5. If one of these faults occurs, it is likely that the system arrangement will have to be changed, the mission profile changed or the mission will fail. Of the twenty-six faults considered, eight have an effect value of 4 or 5. The sum total of sensors or sensor pairs that show a deviation when such faults occur, are expressed as a percentage of the eight high effect faults. The ranking of the sensors will therefore now be based on four factors. The order in which these factors will be evaluated is as follows: faults identified, high effect percentage, SD and number of USOs. Table 8.9 lists the sensor pairs and F11 using the ranking system outlined.

Table 8.9 shows only a small amount of change in the ranking of the sensors compared to that shown in Table 8.6. The sensor pairs F2 and F10, F3 and F10, F2 and F9, F1 and F10 appear in a different order in Table 8.9, as a result of their ability to identify high effect faults. The sensor pairs F2 and F9, F1 and F10 only exhibit a deviation in 50% of high effect fault cases whereas the other sensor pairs deviate in 75% of cases. As a result the sensor pairs F2 and F10, F3 and F10 are ranked above the F2 and F9, F1 and

Table 8.9: Sensor value rankings with fault effect rating

Sensor(s)	No. of Revealed Faults	High Effect %	SD	USOs	Rank
F2 and F11	17	75%	2.53	0	1
F1 and F11	17	75%	2.80	1	2
F3 and F11	17	75%	4.36	0	=3
F11	17	75%	4.36	0	=3
F2 and F10	16	75%	2.33	0	5
F3 and F10	16	75%	4.04	0	6
F2 and F9	16	50%	2.00	3	7
F1 and F10	16	50%	2.56	1	8
F1 and F9	15	75%	2.10	2	=9
F2 and F8	15	75%	2.10	2	=9
F1 and F8	14	75%	2.19	1	11
F3 and F9	11	75%	3.16	0	12
F3 and F8	10	75%	3.15	1	13

F10 pairs. All of the remaining sensor pairs and the F11 entry have the same ranking as in Table 8.6 as they all identified 75% of the high effect fault cases. On a system where there are more high effect failures to consider, the change in the rankings could be more significant.

As was mentioned in Section 8.2.2, the order by which the sensors are ranked in the table is determined by the sorting factors. In Table 8.6 and 8.9 the priority was given to the number of faults that cause a deviation observed by a sensor or group of sensors. This measurement was therefore used to initially rank the sensors. Dependant on what factors are of the greatest importance to a system designer, the sensors can be ranked using the factors discussed above in any order. Alternatively additional factors can be determined and then applied in the ranking process. Possible factors that could be considered include accounting for the probability of a failure mode occurring or the reliability of the sensor working correctly.

8.3 Conclusion

This chapter has shown a process by which the potential positioning of sensors in a system can be assessed to determine a sensor scheme for the specified design requirements. The positioning of sensors in a system can have an effect on the number and type of faults that can be identified. Failure to identify faults can have effects from a reduced system performance all the way up to a serious incident, such as a catastrophic failure. Assessing and selecting sensor locations is therefore an important stage in any system design. The placement of sensors to improve the safety of the system must, however, be evaluated against the cost and weight of any sensors, as well as the space available in a system to position sensors. There is no value in placing so many sensors that the cost or weight of the resultant system is impractical.

Modelling the placement and outputs of sensors on the fuel rig system was effectively achieved using the PN technique. The technique developed to rank the sensor positioning provides several levels of detail. At the highest level, the technique uses a generic sensor deviation to identify those sensors that can identify the presence of a fault in the system. This level, while lacking in detail, does offer a simple means of quickly evaluating the value of a sensor. However, it is not possible using this approach to diagnose a fault in the system. The next level of the analysis considered the type of the deviation measured by a sensor as a result of the failure mode. Three possible sensor outputs were considered in the chosen example, however more levels could be added if observed on the system. Using the PN modelling technique, integrating these further levels would be straightforward. Sensor pairs were also considered at this stage, based on the results of the individual sensor values. The sensor combinations were ranked by assessing the number of faults that caused a deviation in the sensor pair outputs, a SD measure of the faults identified and the number of faults that produced a unique set of sensor outputs. The ranking gave priority to the number of faults that caused a deviation in the sensors in the example, but any of the variables mentioned above could be prioritised depending on the design requirements. Finally, the effect of each fault on the system was given a weighting, where faults that caused a greater effect on the system were given a higher weighting. This factor was then included in the ranking of the sensors. The aim of this ranking system was to give extra weight to those sensors that can identify more of the high effect faults.

The results of applying the above technique to the BAE Systems fuel rig system appear

to be successful. The ranking system provides a flexible, yet robust, means of analysing the sensor outputs and data is available to justify the rankings provided. The technique also allows a comparison to be made between a range of sensor combinations thereby providing a system designer with the opportunity to make an informed decision. Further work in this area could include automating the analysis process, taking account of a greater number of operational modes and accounting for occurrence of failure mode probabilities. In order to address any cost, weight or space limitations that may be present in the system requirements, an optimisation algorithm could also be developed to rank the sensors while satisfying these requirements.

CHAPTER 9

Conclusions and Future Work

9.1 Introduction

The aim of this thesis was to develop a process by which arisings generated on complex systems could be analysed in order to determine whether they were true or false. The motivation for this work stems from the fact that on many complex system a large number of arisings are generated, many of which are known to be false. However, no efficient method exists by which a structured and automated analysis of all arisings can be undertaken. This thesis has developed a technique to satisfy this requirement and the main conclusions of this work are presented in this chapter.

9.2 Conclusions

9.2.1 System Modelling Technique

Three modelling techniques, the decision table, digraph and PN techniques, were evaluated in detail in order to identify the most suitable modelling technique for the purpose of fault propagation. The strengths and weaknesses of each of the modelling techniques were established from their use in literature. The decision table technique provides a componentistic modelling approach that could be systematically applied to large, complex systems. However, the technique lacks the ability to model complex component interaction such as the reverse propagation of faults. The digraph technique can be used to effectively capture the global system behaviour as it models the functional details of a system. However, the detail and flexibility with which systems can be modelled is limited

by the finite number of states that can be used to describe component relationships and fault disturbances. The PN technique allows systems to be modelled both in detail and with a great deal of flexibility. It can be applied modularly and as such encapsulates the advantages of both the decision table and digraph techniques. An identified weakness of the PN technique is that system models can become very large and difficult to interpret. To confirm these observations, each technique was applied to model an example hot water system, operating in two phases. The ability of each technique to propagate faults through the respective system models was then assessed, by comparing the results determined by the model to those that would be expected from the observed behaviour of the simple system. Of the faults considered, only the PN model correctly predicted the behaviour of the system in every case. The digraph model failed to identify seven of ninety-nine expected symptoms as a result of the limitation imposed on describing the relationships between variables. The decision table model failed to identify twenty-five expected symptoms, it also identified several symptoms that were not expected. The inability of the technique to model the reverse propagation of faults is the cause of most of these errors.

The results from the application of each technique to model the hot water system showed that the PN technique most accurately represented the behaviour of the system in the presence of the faults considered. To confirm these results and aid the development of an automated PN simulation software, the PN technique was used to model a more complex example system; a tank level control system. The system mission contained more operational phases (five) than the hot water system (two), included two feedback loops and contained multiple sensor types that were not present on the hot water system. After successfully modelling the behaviour of the tank level control system in the presence of a larger range of failure modes, such as multiple leak sizes and partial pipe blockages, it was concluded that the PN technique would be used as the modelling technique of choice.

9.2.2 Fuel Rig System and PN Model

The BAE Systems Fuel Rig system was used to aid the development of the fault verification technique. The fuel rig is an experimental facility that can be configured to represent the fuel system of an aircraft. A wide range of faults were physically injected into the system or artificially represented using the computational input. There were number of sensors on the fuel rig including level sensors, flow rate sensors and flow pressure sensors. The

output from these sensors was recorded and used in the fault verification process. The output from the fuel rig sensors was also used when developing the fuel rig PN model and confirming its accuracy. By recording the behaviour of the fuel rig when operating without faults present, the PN model was developed to accurately represent the fuel rig system behaviour. Faults were then injected into the fuel rig and its sensor outputs recorded again to inform the further development of the model in capturing the effect of different faults on the system. In order to accurately capture the fuel rig system behaviour, several custom transitions were developed. These new transitions were used to clear tokens from places (i.e. flow rate places), represent an 'IF' statement within a transition and to create single firing transitions.

The fuel rig PN model contains 239 place nodes and 454 transitions. In order to display key aspects of the PN effectively, smaller sub-nets are shown. These show a small number of place and transition nodes from within the larger PN.

9.2.3 Fault Verification Techniques

Six data comparison techniques were considered for the purpose of comparing the fuel rig sensor data and PN simulation results when evaluating arisings from the fuel rig system. Two techniques were taken from literature; the SD and DTW techniques. The point-by-point, delta, binary and time techniques were specifically developed for the fuel rig system.

The techniques were evaluated by comparing the same set of the level sensor outputs from the fuel rig system with data from the PN model. Two scenarios were considered; one with no fault present in the fuel rig or the PN model, and one with a fault present in the fuel rig and modelled in the PN. The results of the tests showed that the binary and time techniques failed to offer a suitable level of detail for comparison. The point-by-point and delta techniques both proved to be highly susceptible to noise, which had a significant effect on the accuracy of results that were produced. The SD technique was chosen as the most suitable comparison technique, as it proved to be more robust and easier to extend to further variables, while also having lower computational requirements compared to the DTW technique.

Having chosen the SD comparison technique, a period of testing and analysis allowed a set of tolerances to be established for each of the variables on the fuel rig system. In

addition, a custom analysis technique was developed to deal with tank leak arisings that, if verified as true, could provide additional information in the form of the leak size and leak location in the vertical axis of the fuel tank.

9.2.4 Fuel Rig System Results

Using the PN software, the ability to verify the presence of faults in the fuel rig, using the recorded sensor outputs and PN model outputs, was demonstrated. All of the revealed failure modes were correctly verified when modelled in the PN. A number of failure modes that remained hidden were also considered. In normal operation, these faults would not be diagnosed as they do not change the behaviour of the system. However, as they did not affect the PN output when included in the system model, they were technically 'verified' by the PN software.

First order fault were also modelled in the PN and then analysed using fuel rig data that was fault free so that a false arising could be simulated. For the majority of failure modes considered, the fault verification technique correctly identified when the fault was false. However, there were issues with both hidden faults and faults that had only a small effect on the fuel rig variables. As hidden failures did not change the behaviour of the system, the fuel rig data always matched well with the PN predicted data that included the fault. As a result the SD values were all within the tolerance limits and the fault was verified incorrectly. For the failure modes which only became revealed in the short, final operational phase or did not significantly change the behaviour of the system throughout the mission, the fuel rig and PN data sets were sufficiently similar and the SD values did not exceed the tolerance limits.

A further scenario considered the ability of the fault verification technique to identify a single, genuine fault when four faults were listed in the health log. The genuine fault was correctly identified and the remaining arisings were correctly determined to be false. Finally a number of second order failure modes were considered. The results showed that the technique could verify the presence of two genuine faults on the system simultaneously. However, as it was identified earlier, the technique had issues when one of the faults became revealed for a short period of time. In this case, a genuine fault was incorrectly categorised as false.

The results of applying the fault verification technique to the fuel rig system indicated

that the aim of this research has been fulfilled. The technique is able to consider arisings generated by the fuel rig system and assess whether the fault is true or false. There are, however, still some limitations that relate to issues created by hidden faults and faults that become revealed in phases of short duration, where false faults can be incorrectly verified as being true and further work is needed.

9.2.5 Software Operation

The PN software has been developed to implement the PN model and the fault verification technique. Using an input file containing a system PN model, the software can be used to simulate phased missions of any duration. The software can automatically record the number of tokens present in a single place or over a range of places throughout the mission. This allows variables, such as tank level or flow rate, to be evaluated after the simulation.

The PN software includes the process by which the PN simulation data is compared to data recorded from the fuel rig, or any other system. The comparison of data is customised to the type of variable being considered and how that variable is represented in the PN model. For example, the flow rate and flow pressure places in the fuel rig PN model can store the same number of tokens but the tokens in each of the places represent different values. This parameter specific information is accounted for in the PN software.

Using a laptop computer the PN software can simulate a 300 second, five phase mission of the fuel rig system and compare fourteen variables in less than 10 seconds.

9.2.6 Airbus A340

The PN software was used to model the behaviour of the Airbus A340 fuel system. A seven phase mission was simulated using the A340 PN model which is considerably larger and contains more complexity than the BAE Systems fuel rig system PN model. The A340 PN model contains 208 place nodes and 451 transitions. Only three failure modes were modelled in this PN. Had the fuel rig PN considered three equivalent failures it would have contained 127 place nodes and 203 transitions. The larger size of the A340 PN model, when considering normal operating behaviour and a similar set of faults, is indicative of its greater complexity. The normal operating behaviour of the A340 PN model was developed and verified using information from literature and an aircraft analysis tool. The behaviour of the A340 system in the presence of faults was an assumption based on the effect of the

same faults on the fuel rig. In all cases, the PN predicted behaviour of the system in the presence of the faults matched that which would be expected from the physical system.

Modelling and simulating the behaviour of the A340 fuel system demonstrated that the PN technique can be used to model large and complex systems. The consideration of multiple failure modes also shows that the fault verification technique could be applied on an industrial scale. There are, however, a number of developments that would be necessary in order for the A340 PN model to be suitable for use in industry. These developments include increasing the model detail by considering all possible failure modes, modelling a greater range of sensors and accurately representing complex component interactions. The use of a more detailed version of the PN modelling technique, such as heirarchical PNs, would likely have to be used to achieve this level of detail without making the PN model overly complex. The model and software must also be thoroughly tested on a physical mock-up and integration testing with the A340 fuel system would also be necessary.

9.2.7 Sensor Value and Optimal Positioning

A technique has been developed by which the potential selection and arrangement of sensors in a system can be ranked in order to determine the most effective set-up to satisfy the design requirements. The technique uses PNs to model systems and the potential sensor locations. Faults are simulated in the PN model and the sensor outputs are recorded. The technique then considers the change in value measured by any sensor, as a result of the faults, and using this knowledge applies a ranking system to the sensors. The technique has also demonstrated how to consider the most effective sensor pairs in a system. The technique has been verified by application to the BAE Systems fuel rig. The results of applying the technique showed how flow sensors can be ranked according to multiple system design requirements, i.e. identify as many faults as possible, identify high consequence faults. The key advantage of the sensor value technique is its ability to rank the possible sensor locations during the design phase of system development. Using the PN technique, multiple analysis of system designs can be conducted accurately and quickly. This approach is more cost effective than a trial and error based technique applied during the manufacture or assessment stages of system development. While the technique provides a valuable capability, further development could allow the technique to take account of the probability of individual failure modes. Automating the analysis process would also make

the technique more appealing to potential users.

9.3 Future Work

The work undertaken within this thesis and the results produced have served to highlight a number of areas of potential future work. These are outlined below.

9.3.1 Fault Verification

9.3.1.1 Fault Assessment in Real Time

The research in this thesis has focused on arisings generated by PBITs. A retrospective approach to analysing the arisings could therefore be performed. There are however, two further built-in tests that would require consideration of arisings in real time; continuous and interrupted built-in tests. The continuous built-in tests are active throughout the operation of a system. The interrupted built-in tests can also be activated during a mission. Assessing arisings during the operation of a system would require development of the fault verification technique to consider the behaviour of the system up to the time of the arising. It would also be necessary to revisit the arising, if initially determined to be false, at phase changes as a fault may only then become revealed.

9.3.1.2 Tolerance Values

The tolerance values applied to the fuel rig variables when assessing the validity of a fault were determined from user experience and from the results generated when building the fuel rig PN model. To apply the fault verification technique to an industrial system, would require a systematic process to determine the tolerance values for each variable. The use of tolerance values/levels in industry is widespread. Companies have the resources to undertake many thousands of simulations recreating a huge range of conditions that would be experienced by a system. Given the results of these simulations, which may be based on the Monte Carlo method, an evidence based process can be used to establish tolerance levels. Safety factors that are specific to different industries can also be applied to give further confidence to the tolerance values applied.

9.3.1.3 Hidden Faults

Results from the fuel rig testing demonstrated issues created by hidden faults. This issue could be minimised or, at least reduced, in the future in a number of ways. The use of a wider range of sensor types is one possible solution. The more sensor types in place on a system, the greater the likelihood a disturbance will be registered in the event of a fault. The use of more sensors however, has to be balanced against the extra cost, weight and complexity they add to the system. A more effective solution may be to use different operational modes regularly throughout the mission of a system to 'flush out' hidden faults. Developing specialist BIT mechanisms to test the state of components/sub-systems, by changing the phase of operation locally, throughout a mission or on demand could also solve the issue. From an industrial perspective, any fault verification technique would be extensively tested. From this testing the conditions under which specific hidden faults may occur would be known. Given this information, action can be taken to make design changes, include further sensor types or modify the system's operational activity to minimise the effects of hidden faults.

9.3.1.4 Faults Appearing in Phases of Short Duration

The fault verification technique exhibited issues verifying genuine faults when the fault became revealed in phases of very short duration. On the fuel rig this was particularly prevalent when faults appeared in the final phase of operation, which lasted for fifteen seconds. This meant only five seconds of data was available to assess the legitimacy of the fault. While phases of such short duration may be rare in many complex systems, options are available to deal with the issue. Where faults only appear in very short operational phases, a specific fault verification test or unique tolerance limit could be applied. This test/tolerance limit would account for the fact that a small variation between the data sets is indicative of the presence of a fault. The risk of using either of these approaches is they will not be able to account for the presence of noise in the output variables or differentiate between noise and a fault. A less risky solution to the issue from an industrial perspective would be to intentionally extend the phase of operation where an arising has been generated in order to provide more data with which to assess the fault. Combining this with the use of the short phase verification tests and specialist tolerances would minimise the time by which the phase length would have to be extended by. Another

potential solution is to use the limited data available and provide users with a confidence level/measure that the fault is present. A user, or an algorithm, could then assess the arising type and the known level of confidence to decide whether to continue or abort the mission.

9.3.1.5 Efficiency Improvements

On systems where many thousands of arisings are generated, it is likely that several faults will be reported on multiple occasions. Developing a technique by which all of these repeated arisings can be assessed on one occasion, or once during every operational phase, would significantly reduce the computational requirements and improve the efficiency of the fault verification process.

9.3.1.6 Industrial Application Challenges

The fault verification technique has been validated on a system in a laboratory environment. Application to an industrial system would therefore show both the true capabilities and limitations of the technique. Identification of a suitable system would be the first stage in this process. Creating the PN model and customising and developing the fault verification software would then be necessary in order to provide a useable capability to an operator. Consideration of higher order failure modes, such as third order faults, may also be necessary. Components that can be repaired during the operation of a system may represent another area of work that has not be considered here, but would have to be accounted for in order to truly consider real world systems. Examples of these components, such as sensors or pipes, could be seen in power stations or chemical plants.

One key area of concern for using the fault verification technique on an industrial system is the requirement to model the system using PNs. As Schneeweiss noted [24], PNs are not widely used in industry. Combining the fault verification technique developed in this thesis with an automatic PN generating tool would increase the likelihood of application of the technique in industry. Some research undertaken at Loughborough University [37] has begun to look at the automatic construction of PNs from component and system descriptions. While future work would be required to combine these two techniques, the benefits to industry could be significant.

9.3.2 Sensor Value and Optimal Positioning

9.3.2.1 Probabilistic Approach

Work to determine sensor value in this thesis has only considered the consequence of certain failure modes occurring. A true risk analysis must also consider the likelihood of any failure mode occurring. This would require a probabilistic approach to be undertaken, whereby individual component failure rates are either researched or allocated. The sensor value could then be determined and sensor rankings adjusted to reflect the risk of the specific failure modes.

9.3.2.2 Further Sensor Types and Combinations

Flow rate sensors are the only sensor type to have been considered in this study. Other sensor types, including level sensors and flow pressure sensors could also be considered. Combining different types of sensor, as well as different sensor locations could also provide more detail to the fault diagnosis process and would be a valuable contribution.

9.3.2.3 Automation

The method developed to determine sensor value thus far is manual. In order to thoroughly and effectively analyse all possible sensor type and location combinations, an automated process will be required. The PN software produced as a part of this research is of value in terms of predicting the sensor outputs. The ranking of sensors, for given design requirements, can be automated by integrating the PN software with an optimisation algorithm.

References

[1] A. Birolini. *Reliability Engineering: Theory and Practice*. Springer Berlin Heidelberg, 2010.

[2] S. L. Pollack, H. T. Hicks, and W. J. Harrison. *Decision Tables: Theory and Practice*. Wiley, 1971.

[3] S.L. Pollack. How to build and analyze decision tables. Technical report, The RAND Corporation, November 1963. P-2829.

[4] S. L. Salem, G. E. Apostolakis, and D. Okrent. A new methodology for the computer-aided construction of fault trees. *Annals of Nuclear Energy*, 4(9-10):417–433, 1977.

[5] B. E. Kelly and F. P. Lees. The propagation of faults in process plants: 1. modelling of fault propagation. *Reliability Engineering*, 16(1):3–38, 1986.

[6] J. de Gelder and M. Steenhuis. A knowledge-based system approach for code-checking of steel structures according to eurocode 3. *Computers & Structures*, 67(5):347–355, 1998.

[7] A. Carpignano and A. Poucet. Computer assisted fault tree construction: a review of methods and concerns. *Reliability Engineering & System Safety*, 44(3):265–278, 1994. Special Issue On Advanced Computer Applications.

[8] A. Majdara and T. Wakabayashi. Component-based modeling of systems for automated fault tree generation. *Reliability Engineering & System Safety*, 94(6):1076–1086, 2009.

[9] J. Andrews and J. J. Henry. A computerized fault tree construction methodology. In *Proceedings of the Institute of Mechanical Engineers*, volume 211, pages 171–183, 1997.

[10] F. Witlox. Locational choice modelling using fuzzy decision tables. In *Biennial Conference of the North American Fuzzy Information Processing Society*, pages 80–84, 1996.

[11] J. Vanthienen, C. Mues, and A. Aerts. An illustration of verification and validation in the modelling phase of kbs development. *Data & Knowledge Engineering*, 27(3):337–352, 1998.

[12] J. Vanthienen and G. Wets. From decision tables to expert system shells. *Data & Knowledge Engineering*, 13(3):265–282, 1994.

[13] S. A. Lapp and G. J. Powers. Computer-aided synthesis of fault-trees. *IEEE Transactions on Reliability*, R-26(1):2–3, April 1977.

[14] L. M. Bartlett, E. E. Hurdle, and E. M. Kelly. Integrated system fault diagnostics utilising digraph and fault tree-based approaches. *Reliability Engineering & System Safety*, 94(6):1107–1115, 2009.

[15] J. Andrews and G. Brennan. Application of the digraph method of fault tree construction to a complex control configuration. *Reliability Engineering & System Safety*, 28(3):357–384, 1990.

[16] J. Andrews and J. Morgan. Application of the digraph method of fault tree construction to process plant. *Reliability Engineering*, 14(2):85–106, 1986.

[17] P. K. Andow. Difficulties in fault-tree synthesis for process plant. *Reliability, IEEE Transactions on*, R-29(1):2–9, April 1980.

[18] D. J. Allen. Digraphs and fault trees. *Industrial & Engineering Chemistry Fundamentals*, 23(2):175–180, 1984.

[19] C. Petri. *Kommunikation mit Automaten.* PhD thesis, Institut fur instrumentelle Mathematik, Bonn, 1962.

[20] J.L. Peterson. *Petri Net Theory and the Modeling of Systems.* Prentice-Hall, 1981.

[21] M. F. Russo and A. Sasso. Modeling, analysis, simulation and control of laboratory automation systems using petri nets: part 1. modeling. *Journal of the Association for Laboratory Automation*, 10(3):172–181, 2005.

[22] S.P. Chew, S.J. Dunnett, and J.D. Andrews. Phased mission modelling of systems with maintenance-free operating periods using simulated petri nets. *Reliability Engineering & System Safety*, 93(7):980–994, 2008.

[23] T. Murata. Petri nets: properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, April 1989.

[24] W. G. Schneeweiss. *Petri Nets for Reliability Modelling*. Life-Long-Learning, 1999.

[25] I. Mura and A. Bondavalli. Markov regenerative stochastic petri nets to model and evaluate phased mission systems dependability. *IEEE Transactions on Computers*, 50(12):1337–1351, 2001.

[26] H. Alla and R. David. A modelling and analysis tool for discrete events systems: continuous petri net. *Performance Evaluation*, 33(3):175–199, 1998.

[27] K. Jensen. *Coloured Petri Nets - Basic Concepts, Analysis Methods and Practical Use*. Springer, second edition, 1997.

[28] W. G. Schneeweiss. Tutorial: petri nets as a graphical description medium for many reliability scenarios. *IEEE Transactions on Reliability*, 50(2):159–164, June 2001.

[29] F. P. Garcia, D. J. Pedregal, and C. Roberts. Time series methods applied to failure prediction and detection. *Reliability Engineering & System Safety*, 95:698–703, 2010.

[30] J. B. Kruskal and M. Liberman. *The symmetric time-warping problem: from continuous to discrete*. Addison-Wesley, 1983.

[31] R. Niels. Allograph based writer identification, handwriting analysis and character recognition. Master's thesis, Donders Centre for Brain, Behaviour and Cognition, Radbound University Nijmegen, 2010.

[32] V. Atamuradov, F. Camci, S. Baskan, and M. Sevkli. Failure diagnostics for railway point machines using expert systems. In *The 7th IEEE International Symposium on Diagnostics for Electric Machines, Power Electronics and Drives*, 2009.

[33] V. Vuori, M. Aksela, and J. Laaksonen. Adaptive character recognizer for a hand-held device: Implementation and evaluation setup. 2000.

[34] Airbus Industrie. Fast - airbus technical digest. Technical Report 14, Airbus Industrie Product Support, 1 Rond-Point, 31707, Blagnac Cedex, France, February 2003.

[35] D. Simos. Piano aircraft analysis. http://www.lissys.demon.co.uk/PianoX.html, July 2013, Accessed: 18-07-13.

[36] European Environment Agency. Eea air pollutant emission inventory guidebook. http://www.eea.europa.eu/themes/air/emep-eea-air-pollutant-emission-inventory-guidebook/emep, June 2009, Accessed: 18-07-13.

[37] K. S. Stockwell and S.J. Dunnett. Application of a reliability model generator to a pressure tank system. In *Automation and Computing (ICAC), 2012 18th International Conference on*, pages 1–6, 2012.

APPENDIX A

Hot Water System Component Decision Tables

The decision tables shown below represent the steady state behaviour of the hot water system components.

A.1 Phase 1 Component Decision Tables

In Table A.1 the gas input and output states are either supply (S) or no supply (NS). The internal mode is either no blockage (NB) or blockage (B).

Table A.1: Gas pipe decision table

Row Number	Gas Input	Internal Mode	Gas Output
1	NS	–	NS
2	–	B	NS
3	S	NB	S

In Table A.2 the possible temperature inputs and signal outputs can be high (H), normal (N) or low (L). The temperature input could also be zero (0). The zero value represents atmospheric conditions which would be recorded in the event of the water pipe rupturing. The internal modes of the sensor are working (W), failed low (FL) and failed high (FH).

Table A.2: Temperature sensor decision table

Row Number	Temp. Input	Internal Mode	Signal Output
1	0	W	L
2	L	W	L
3	N	W	N
4	H	W	H
5	–	FL	L
6	–	FH	H

In Table A.3 the range of potential signal inputs are the same as the temperature sensor outputs; high (H), normal (N) and low (L). The internal modes are working (W), constant demand (CD), no demand (ND). The output signals to the control valve are either signal to open (SO) or signal to close (SC).

Table A.3: Controller decision table

Row Number	Signal Input	Internal Mode	Signal Output
1	L	W	SO
2	N	W	SC
3	H	W	SC
4	–	CD	SO
5	–	ND	SC

In Table A.4 the gas input and heat output will be either no supply (NS) or supply (S) and the internal mode will either be working (W) or failed off (FOff).

Table A.4: Pilot light decision table

Row Number	Gas Input	Internal Mode	Heat Output
1	NS	–	NS
2	–	FOff	NS
3	S	W	S

In Table A.5 the water input and output states are either no supply (NS) or supply (S). The possible internal modes of the valve are working (W), stuck open (SO) and stuck closed (SC).

Table A.5: Non-Return valve decision table

Row Number	Water Input	Internal Mode	Water Output
1	–	SC	NS
2	NS	–	NS
3	S	W	S
4	S	SO	S

In Table A.6 the input pipe pressure state will be one of zero (0), low (L), normal (N) or high (H). The internal mode of the valve will be working (W), stuck open (SO) or stuck closed (SC). The output pressure will therefore be either atmospheric (A) if the valve remains closed or above atmospheric (AA) if the valve is opened and pressurised gas leaves the system.

Table A.6: Pressure relief valve decision table

Row Number	Pipe Pres.	Internal Mode	Pres. Output
1	–	SC	A
2	0	W	A
3	L	W	A
4	N	W	A
5	H	W	AA
6	0	SO	A
7	L	SO	AA
8	N	SO	AA
9	H	SO	AA

In Table A.7 the water volume, temperature and pressure are all input states that can be one of zero (0), low (L), normal (N) or high (H). The internal modes of the tap are working closed (WC), stuck closed (SC) and stuck open (SO). The tap outputs are water output, measured as either supply (S) or no supply (NS), and water temperature,

described as either low (L), normal (N) or high (H).

Table A.7: Tap decision table

Row Number	Pipe Pres.	Pipe Temp.	Pipe Vol.	Internal Mode	Water Output	Water Temp.
1	–	–	–	SC	NS	–
2	–	–	–	WC	NS	–
3	0	0	0	SO	NS	–
4	L	L	L	SO	NS	–
5	L	N	L	SO	NS	–
6	L	H	L	SO	NS	–
7	H	H	N	SO	S	H
8	N	L	N	SO	S	L

A.2 Phase 2 Component Decision Tables

Tables A.8 shows the phase two decision table for the tap. The codes used to represent the input states, internal modes and output states have not changed from those defined with the phase 1 decision table.

Table A.8: Tap decision table

Row Number	Pipe Pres.	Pipe Temp.	Pipe Vol.	Internal Mode	Water Output	Water Temp.
1	–	–	–	SC	NS	–
2	0	0	0	WO	NS	–
3	L	L	L	WO	NS	–
4	L	N	L	WO	NS	–
5	L	H	L	WO	NS	–
6	N	N	N	WO	S	N
7	N	L	N	WO	S	L
8	0	0	0	SO	NS	–
9	L	L	L	SO	NS	–
10	L	N	L	SO	NS	–
11	L	H	L	SO	NS	–
12	N	N	N	SO	S	N
13	N	L	N	SO	S	L

APPENDIX B

Hot Water System Component Digraphs

B.1 Phase 1 Component Digraph Models

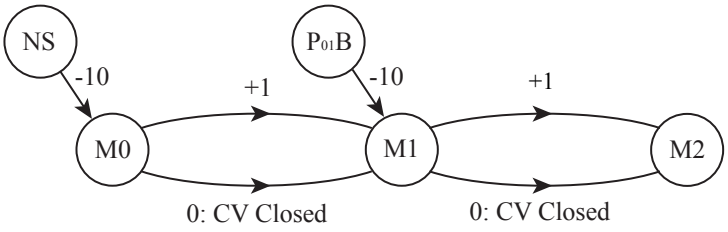

Figure B.1: Gas pipe digraph

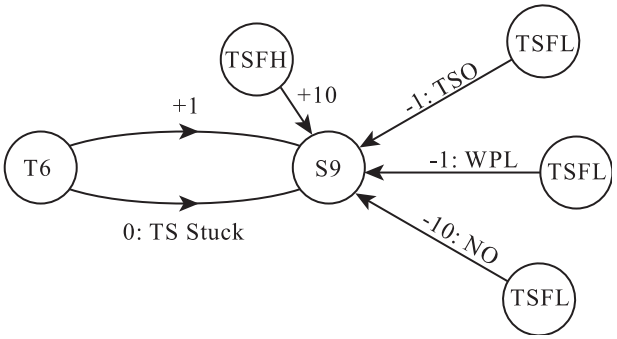

Figure B.2: Temperature sensor digraph - Phase 1

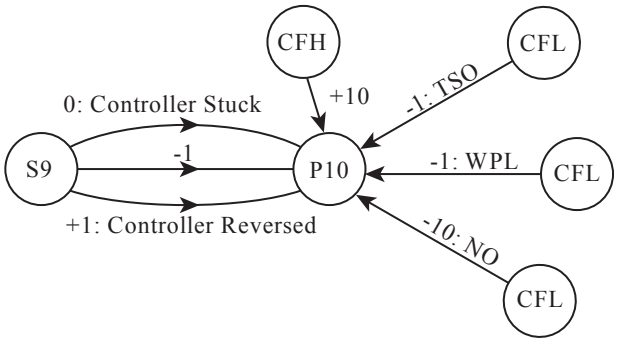

Figure B.3: Controller digraph - Phase 1

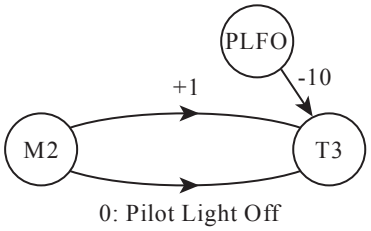

Figure B.4: Pilot light digraph

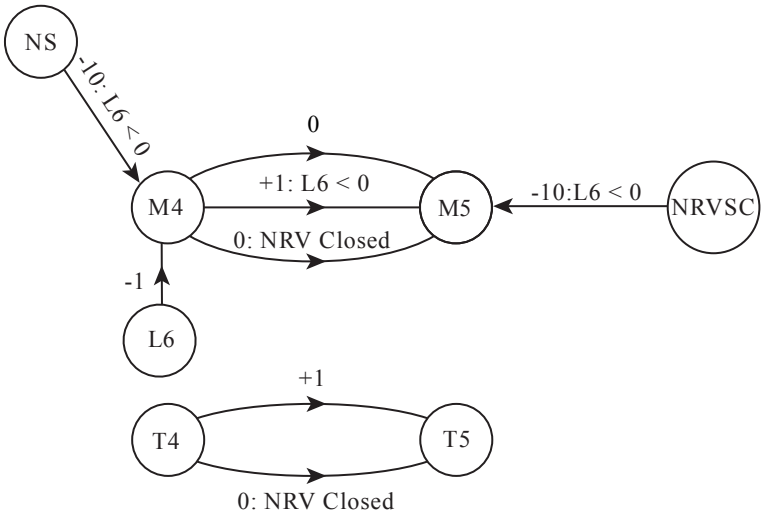

Figure B.5: Non-Return valve digraph - Phase 1

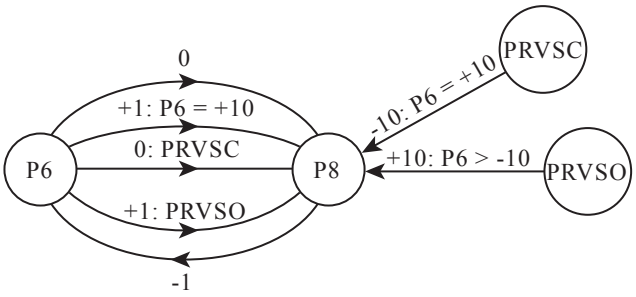

Figure B.6: Pressure relief valve digraph

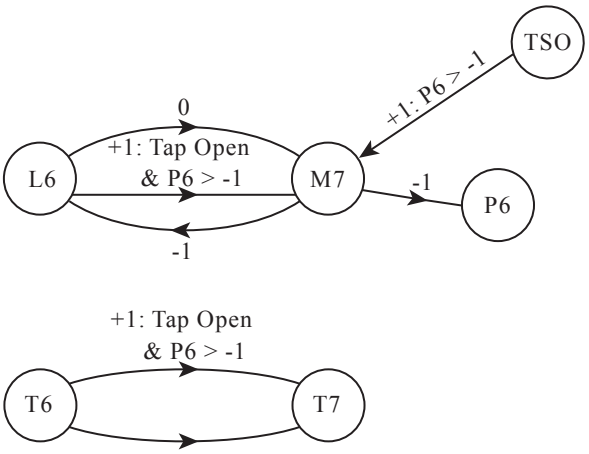

Figure B.7: Tap digraph - Phase 1

B.2 Phase 2 Component Digraph Models

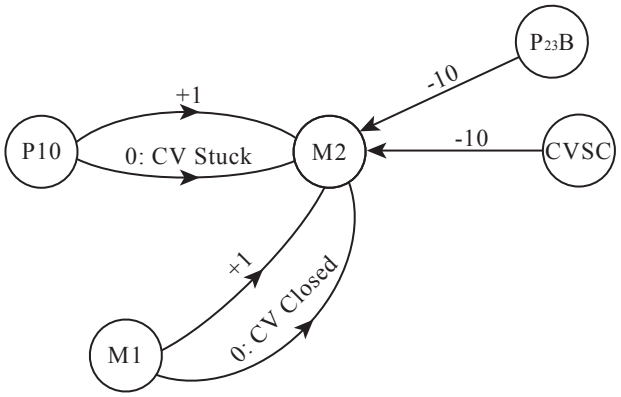

Figure B.8: Control valve digraph - Phase 2

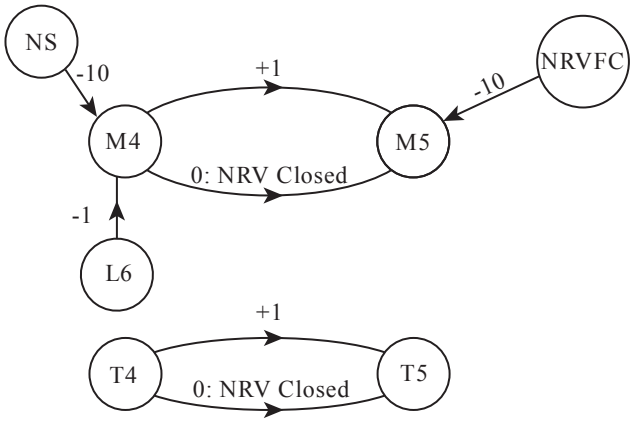

Figure B.9: Non-Return valve digraph - Phase 2

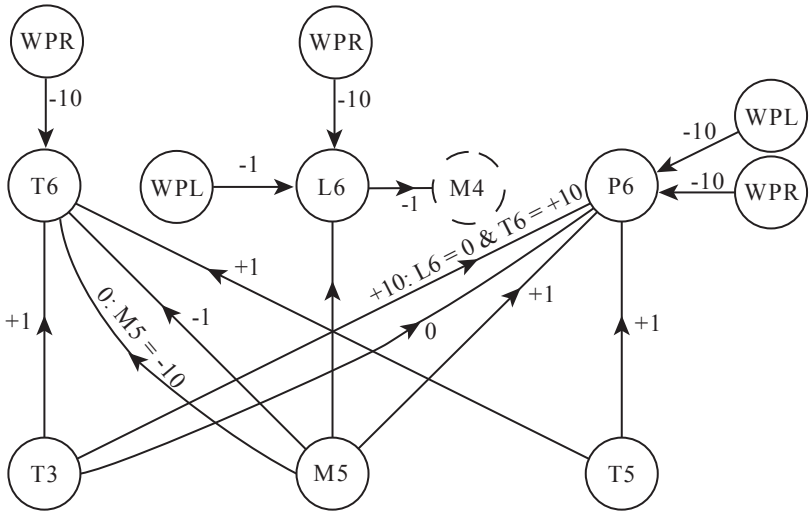

Figure B.10: Water pipe digraph - Phase 2

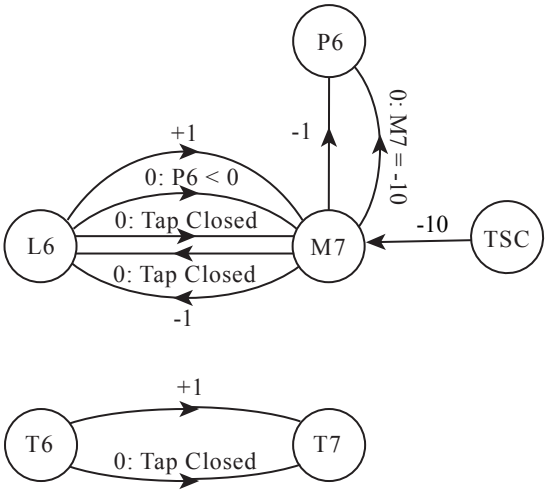

Figure B.11: Tap digraph - Phase 2

APPENDIX C

Hot Water System Petri Net Model

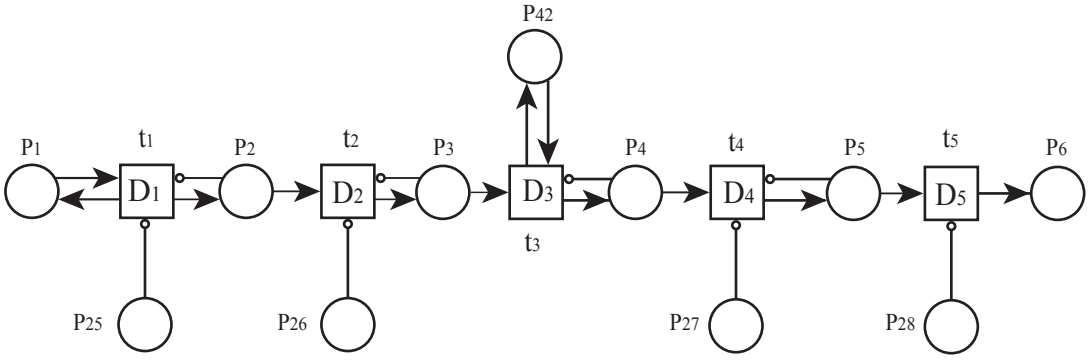

Figure C.1: Gas input petri net model

Table C.1: Gas input petri net place descriptions

Place No.	Description	Place No.	Description
1	Gas at Supply Point	25	No Supply
2	Gas at Entry to System	26	Pipe Blocked
3	Gas at Entry to Control Valve	27	Pipe Blocked
4	Gas at Exit from Control Valve	28	Pilot Light Off
5	Gas at Entry to Pilot Light	42	Control Valve Open
6	Heat out of Pilot Light		

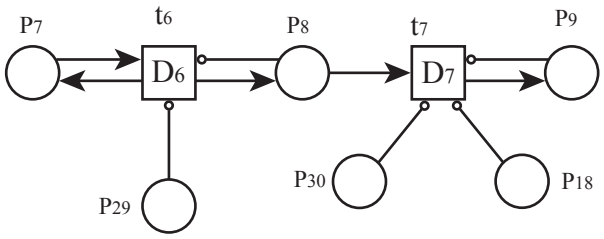

Figure C.2: Water input petri net model

Table C.2: Water input petri net place descriptions

Place No.	Description	Place No.	Description
7	Water at Supply Point	18	Normal Water Pipe Volume
8	Water into NRV	29	No Supply
9	Water out of NRV	30	NRV Failed Closed

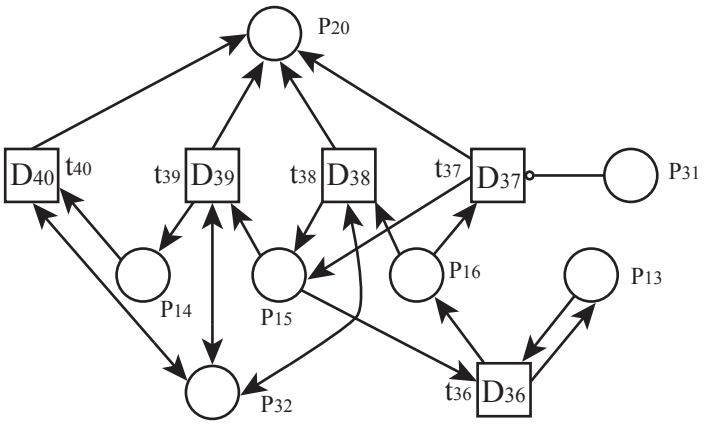

Figure C.3: Pressure relief valve petri net model

Table C.3: Pressure relief valve petri net place descriptions

Place No.	Description	Place No.	Description
13	High Water Pipe Temp	20	Lost Pressure
14	Low Water Pipe Pres	31	PRV Stuck Closed
15	Normal Water Pipe Pres	32	PRV Stuck Open
16	High Water Pipe Pres		

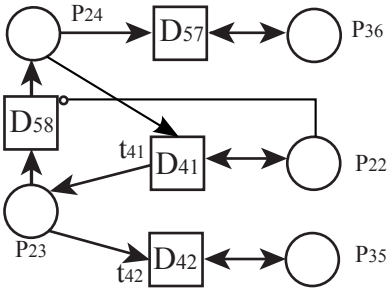

Figure C.4: Tap petri net model [1/2]

Table C.4: Tap petri net place descriptions [1/2]

Place No.	Description	Place No.	Description
22	User Demand	35	Tap Stuck Open
23	Tap Open	36	Tap Stuck Closed
24	Tap Closed		

Table C.5: Tap petri net place descriptions [2/2]

Place No.	Description	Place No.	Description
11	Low Water Pipe Temp	18	Normal Water Pipe Volume
12	Normal Water Pipe Temp	23	Tap Open
13	High Water Pipe Temp	37	Pressure at Tap
14	Low Water Pipe Pres	38	Water at Tap
15	Normal Water Pipe Pres	39	Heat at Tap
16	High Water Pipe Pres	40	Hot Water Out
17	Low Water Pipe Volume	41	Cold Water Out

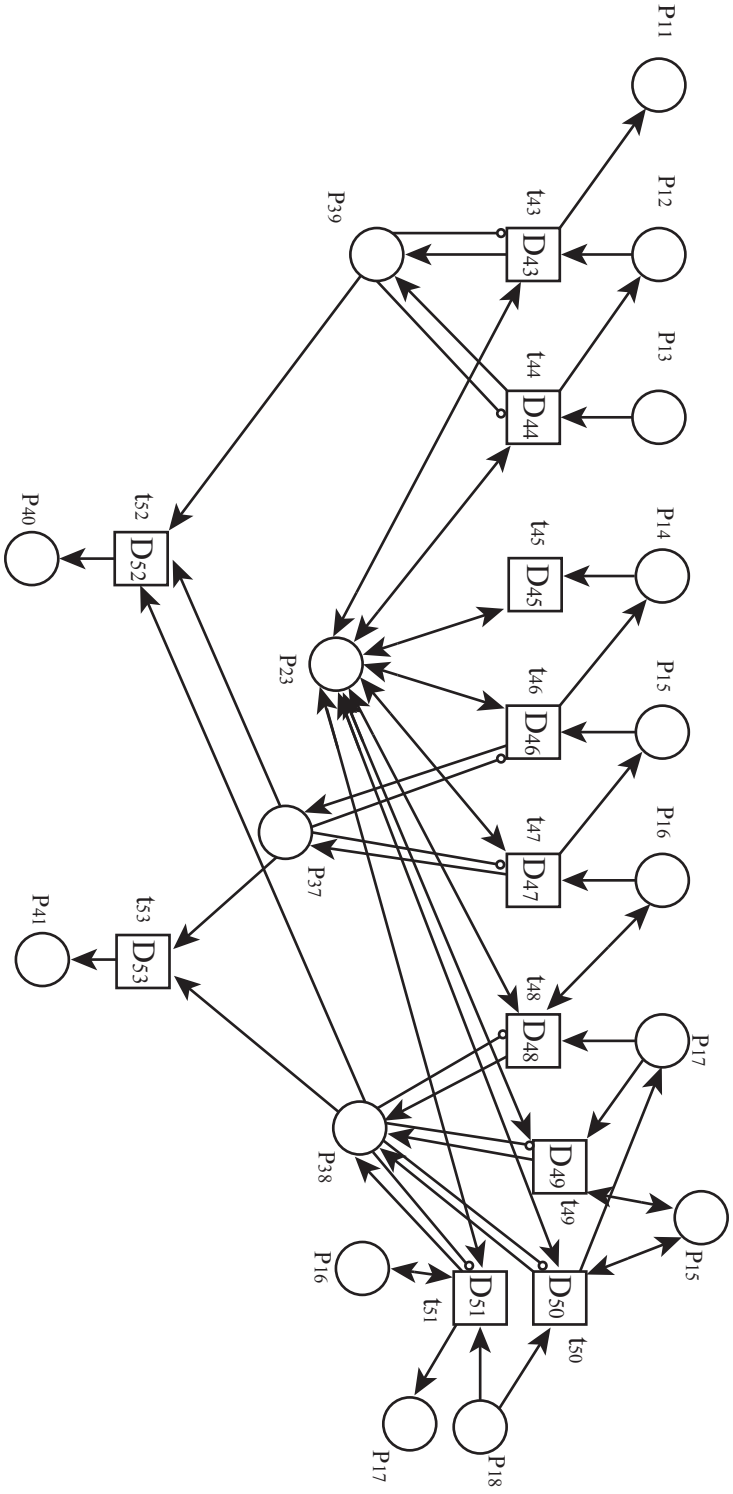

Figure C.5: Tap petri net model [2/2]

APPENDIX D

Tank Level Control System Petri Net Model

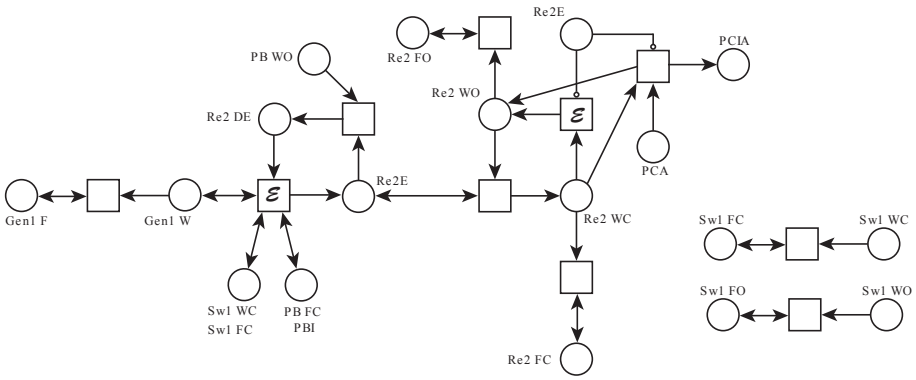

Figure D.1: Relay 2 powering up/down, opening and closing

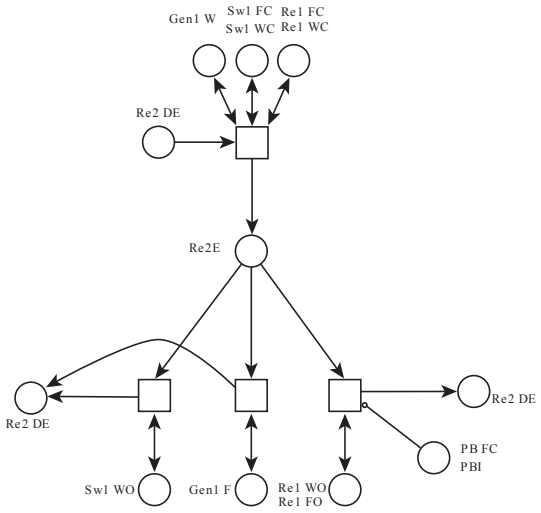

Figure D.2: Relay 2 powering down

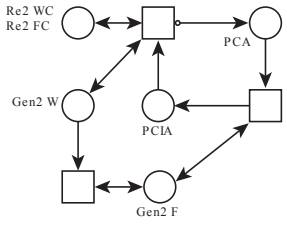

Figure D.3: Pump circuit powering up

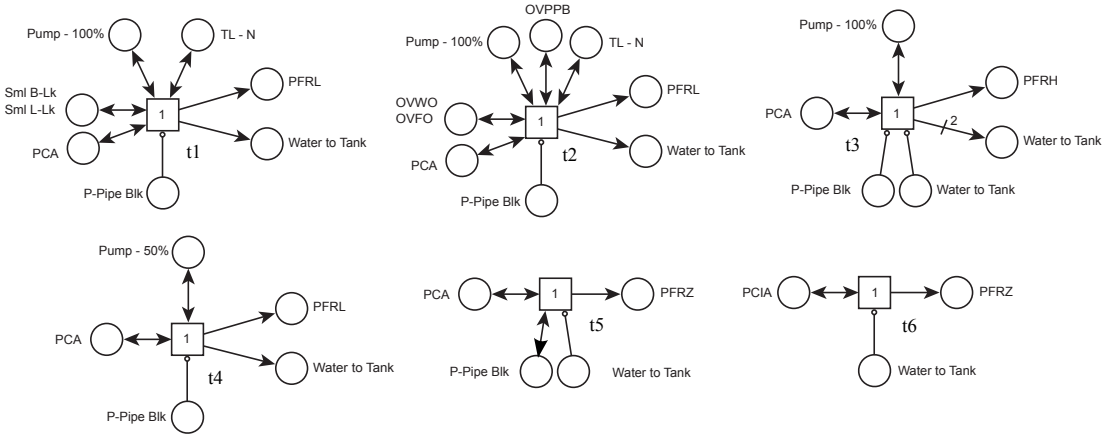

Figure D.4: Pump flow outputs

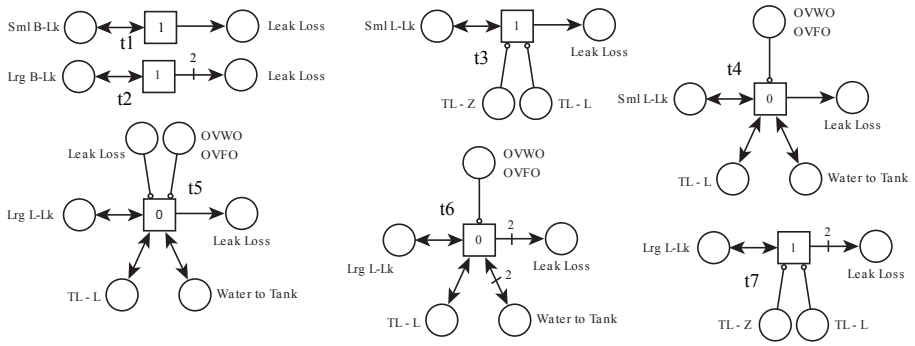

Figure D.5: Tank leak outputs

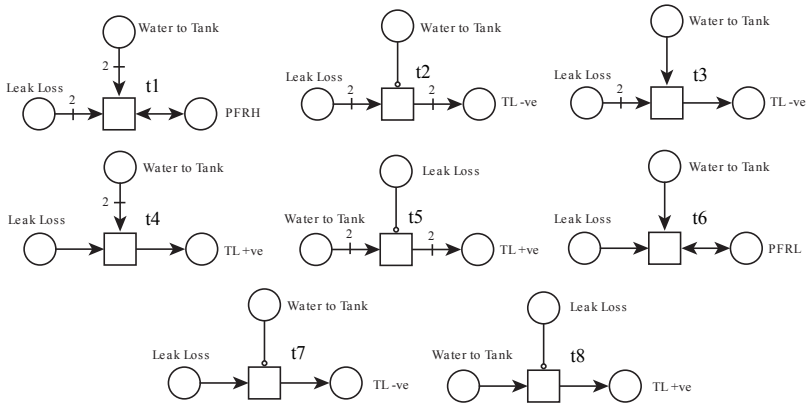

Figure D.6: Combined effect of pump and leak tank level changes

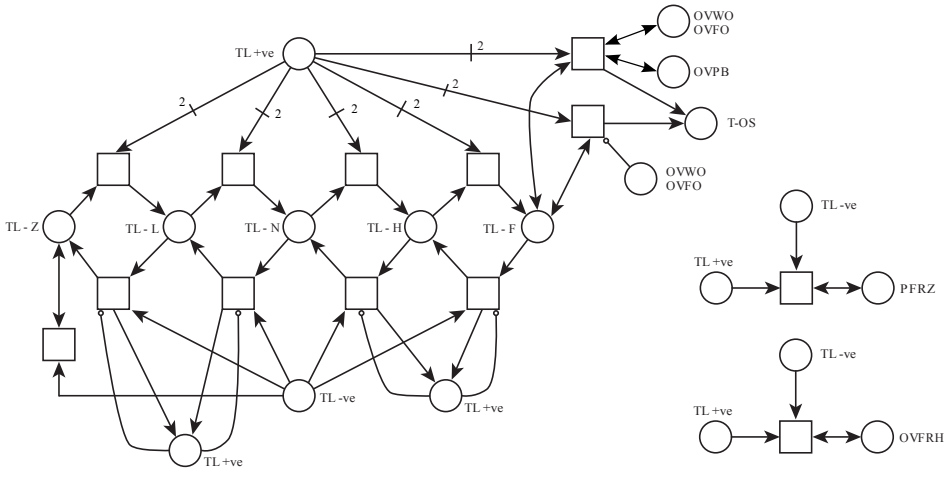

Figure D.7: Tank level changes

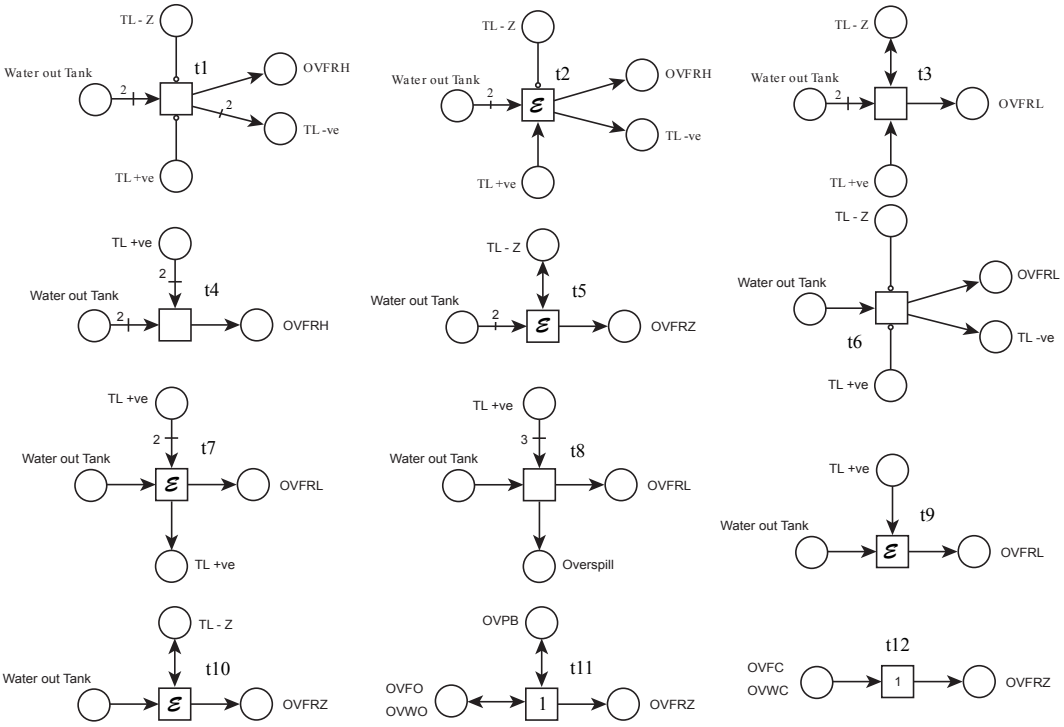

Figure D.8: Output valve outputs

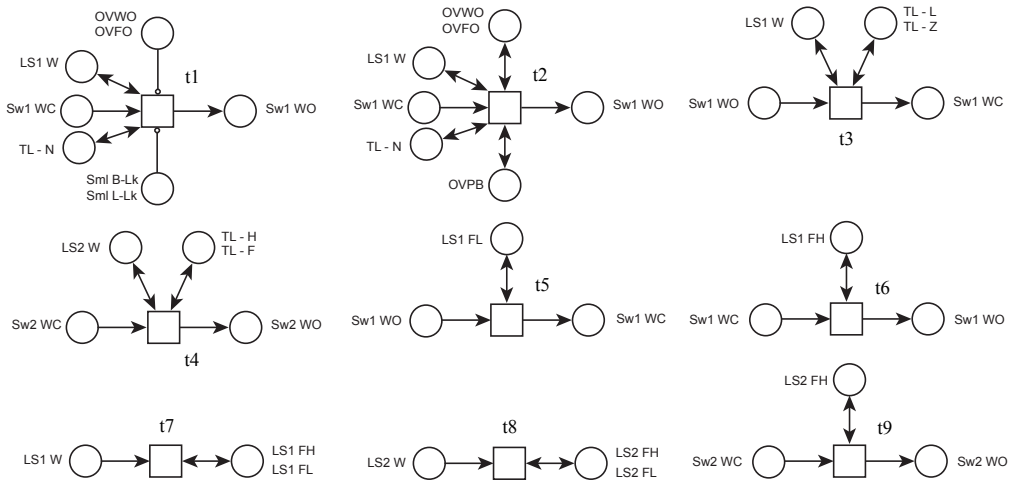

Figure D.9: Feedback sensor outputs

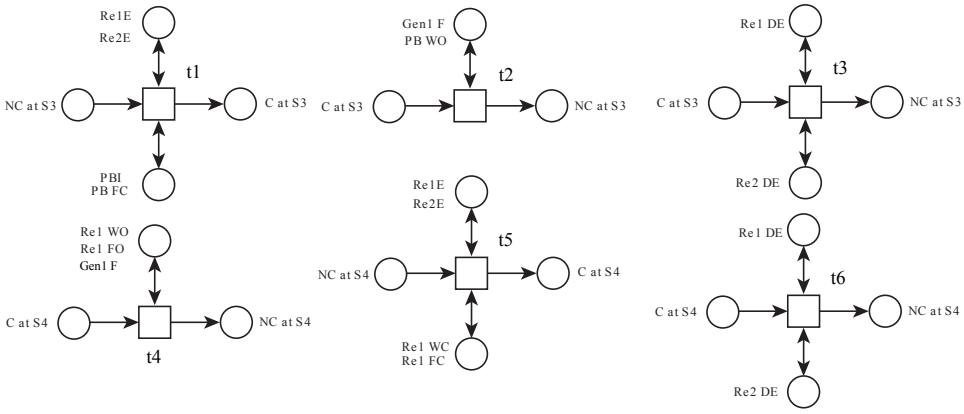

Figure D.10: Current sensors CS3 and CS4 outputs