Received: 27 April 2022 | Revised: 25 August 2022 | Accepted: 6 October 2022

DOI: 10.1002/asmb.2722

RESEARCH ARTICLE

WILEY

Optimisation of nuclear reactor primary coolant design and maintenance parameters

Mark James Wootton[1] | Ying Zhou[2] | John D. Andrews[1] | Roger Smith[2] | John Arul[3] | Gopika Vinod[4] | Hari Prasad Muruva[4] | Vipul Garg[4]

[1]Faculty of Engineering, University of Nottingham, Nottingham, UK

[2]Department of Mathematical Sciences, Loughborough University, Loughborough, Leicestershire, UK

[3]Reactor Engineering Group, Indira Gandhi Centre for Atomic Research, Kalpakkam, India

[4]Reactor Safety Division, Bhabha Atomic Research Centre, Mumbai, India

Correspondence
Mark James Wootton, Faculty of Engineering, University of Nottingham, University Park, Nottingham NG7 2RD, UK.
Email: m.j.wootton@bham.ac.uk

Funding information
Engineering and Physical Sciences Research Council, Grant/Award Number: EP/R021988/1

Abstract

An optimisation methodology is presented using the primary coolant circulation system of a nuclear reactor as its case study, the purpose of which is to find combinations of selected design and maintenance parameters to maximise the reactor safety and minimise monetary expenditure. The parameter space was sampled by a Monte Carlo method and Petri net modelling was used to predict the performance of each of these options. The optimal solutions were then extracted from the data via computation of the Pareto front, with further analysis conducted on parameter sets of interest.

KEYWORDS

design optimisation, high performance computing, Monte Carlo, nuclear power safety, Pareto front analysis, Petri net modelling, risk and reliability engineering

1 | INTRODUCTION

Two important considerations in the design of nuclear reactor systems are cost and safety, with the objective being to respectively minimise or maximise these parameters. By their nature, these two objectives are in conflict as little to improve the safety of a design can be done without some increase in expenditure. To address this issue this work presents the use of models of component failure, maintenance, and repair, in conjunction with an optimisation methodology. The primary coolant circulation system of a generic nuclear reactor with modern design features is used as a case study, from which a number of design options are identified. Their parameter space is explored via Monte Carlo sampling, whereby corresponding models are created and simulated to produce the performance of each configuration.

Risk assessment is traditionally performed using a combination of fault tree analysis and event tree analysis. The former has its origins in the 1960s at Bell Laboratories with the work of Watson.[1] A fault tree is used to calculate the rate of occurrence of a system failure mode (*top event*) by systematically breaking the cause of the event down into lower levels of complexity, using the Boolean logic gates *AND* and *OR*, as seen in Figure 1. This development of the failure logic continues until each branch is terminated by an event, such as a component failure mode, for which failure and repair data is available. These are referred to as *basic events*. A set of basic events is called a *minimum cut set* if the occurrence of all of its members would cause the top event, and if the top event would not occur if any one member was removed from the set.

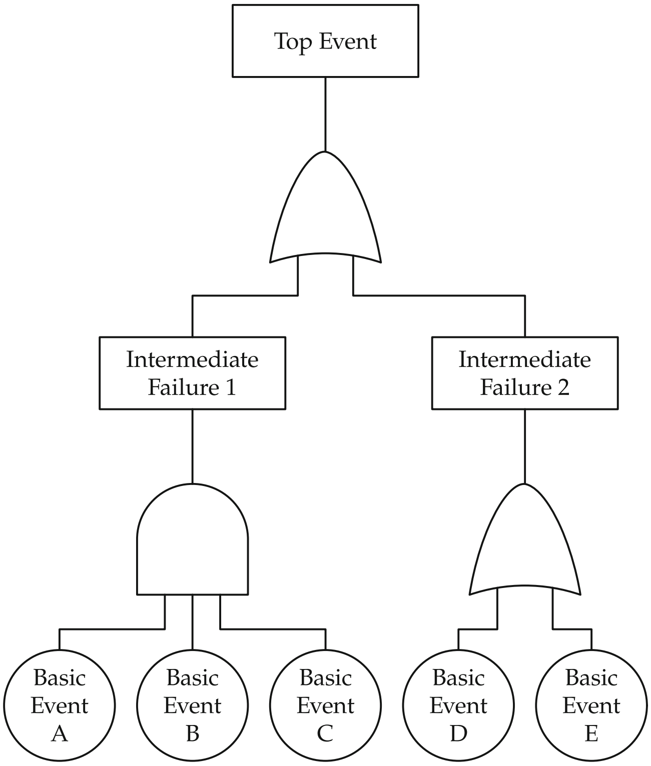

FIGURE 1 In this simple fault tree, the top event is caused by either one of two intermediate failures. The first requires all of the basic events on its branch to occur, whereas the second is caused by either basic event below it. Therefore, the minimum cuts sets are {A,B,C}, {D}, and {E}.

Quantitative analysis was developed by Vesely[2], termed kinetic fault tree theory, in which the frequency of the top event is calculated from the minimal cut sets and the failure rate and likelihood of each basic event. Additionally, calculation of importance measures[3] gives the relative significance of the basic events and minimum cut sets and thus enables prioritisation to be given to the most safety critical failure modes. The traditional process of analysing a large fault tree still requires approximations to be made however. An advanced methodology providing a fast and accurate means to quantify fault trees through the use of binary decision diagrams was developed in 1993 by Rauzy,[4] whose efficiency and accuracy was sufficient to help alleviate this problem.[5]

Rasmussen led a team of people to produce the WASH 1400 report in 1975.[6] This performed a risk assessment on a nuclear plant and introduced the event tree methodology. By contrast to fault trees, an event tree starts with a single *initiating event* and progresses, from left to right across a page, towards the set of all possible resulting outcomes. This is an inductive approach which tracks the functionality and failure of system safety responses along branching pathways, as is illustrated in Figure 2.

The rate at which each outcome occurs is given by the product of the frequency of the initiating event and the probability of every event appearing along its route. The probabilities of the failure events are commonly obtained from fault tree analysis and it is then trivial to calculate the probabilities of success.[7] However, when basic events are shared between multiple fault trees, the resulting dependency has to be accounted for and can be achieved using binary decision diagrams.[8]

Despite being mature and well developed methodologies, neither fault tree nor event tree analysis is suitable for the requirements of this work. When applying them, it is necessary to assume independence of the occurrence of basic events, and commercial implementations typically only permit the use of constant failure rates, preventing the modelling of processes such as component infant mortality or ageing and wear-out. Likewise, component repair times are unrealistically constrained to exponential distributions, and it is infeasible to model complex maintenance strategies or repair processes. At present, there exists no variation of fault tree or event tree that addresses all these concerns, which is of particular concern in relation to the increasing failure rates associated with ageing, the neglect which reduces the accuracy of the resulting safety assessment.

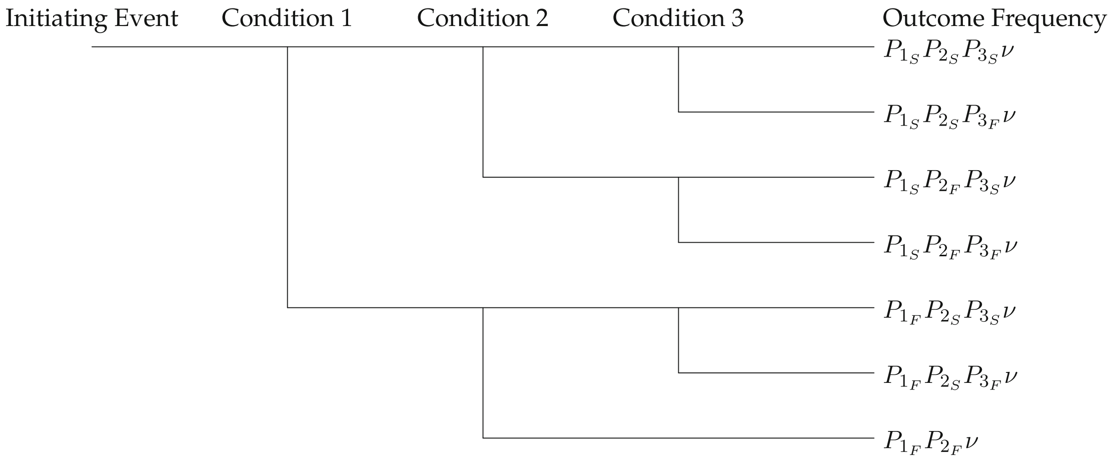
Initiating Event
Condition 1
Condition 2
Condition 3
Outcome Frequency
$P_{1_S}P_{2_S}P_{3_S}\nu$
$P_{1_S}P_{2_S}P_{3_F}\nu$
$P_{1_S}P_{2_F}P_{3_S}\nu$
$P_{1_S}P_{2_F}P_{3_F}\nu$
$P_{1_F}P_{2_S}P_{3_S}\nu$
$P_{1_F}P_{2_S}P_{3_F}\nu$
$P_{1_F}P_{2_F}\nu$

FIGURE 2 This event tree has three possible branching points following the initiating event, with the failure of both the first and second condition rendering the third irrelevant. The probabilities of success and failure for each condition are respectively P_{1_S}, P_{2_S}, and P_{3_S}, and P_{1_F}, P_{2_F}, and P_{3_F}, where $P_{n_S} = 1 - P_{n_F}$, with the initiating event occurring at a frequency of ν.

This work was conducted using the Petri nets methodology[9,10] to overcome the limitations of event tree and fault tree methods to model these system characteristics. As this research aims to explore the effectiveness of alternative maintenance strategies for the system, Petri nets are significantly better suited in comparison to the traditional methodologies. Any fault tree or event tree structure can be replicated as a Petri net, but its potential enables more details and less restrictive assumptions to be included in the assessment. For example, a Petri net can include both multiple occurrences of the same basic events and different failure mode states for a component, and cyclic behaviours and multiple concurrent or synchronous processes can be modelled easily. Any probability distribution can be used to define the time to failure or time to repair for a component. The component can also exist in more than two binary conditions (working and failed) and can therefore progress through a sequence of degraded states until it achieves the failed condition. Where these condition states can be monitored, complex, condition based maintenance, strategies can be developed and their performance modelled. It is also possible to model the shift between multiple modes of operation, such as running in its normal configuration for a fixed period of time before commencing the shutdown sequence, or engaging the use of a emergency system in reaction to a disruptive fault.

The major cost associated with Petri net modelling is the necessity to execute large batches of simulations in its evaluation, with potentially many iterations being required to reach convergence. Although Petri net modelling can be computationally expensive compared to the fault tree and event tree methodologies, this can be justified by its facilitation of high fidelity representations of system dynamics, which would not be possible using the traditional techniques. The methodology has already been applied in civil nuclear energy contexts,[11-22] and prior works employ Petri nets for various purposes. For example, some works[11,12,14,15,19-22] describe or develop operational procedures or processes in nuclear systems using Petri nets, among which a number were directly concerned with safety, risk or reliability.[19-22] More indirect applications can also be found, where Petri nets were created to describe nuclear reactor subsystems, but the evaluation was not performed by simulating the models. Instead the Petri net represented an intermediate form, which was converted for assessment by an alternative methodology, such as Markov chains,[16] reachability graphs,[17] or both.[18]

Where this work differs in its aim, is that it seeks to tie its Petri net models into an optimisation process. Their ability to represent arbitrary system configurations makes them well suited for this end. While examples of optimisation in Petri net models exist in the literature from other industrial contexts,[23-25] there is an unrealised opportunity to develop its use for reactor safety engineering.

Monte Carlo sampling[26,27] of the optimisation parameter space is used in this work to generate data, from which the Pareto optimal[28,29] configurations are extracted. The advantage of this methodology is the simplicity of its implementation and execution, and that the system is sampled in an even and unbiased way, such that the full extent of the Pareto front is seen. Although other methodologies exist that converge faster than the Monte Carlo algorithm, this advantage would then be lost.

2 | CASE STUDY

For our case study, a primary coolant system of a generic reactor system with an emphasis on passive safety is used. A schematic of the system is shown in Figure 3 with the relationship to the optimisable parameters illustrated. Pipes delivering low temperature coolant to the reactor are drawn in light blue, with the path of hot coolant from the core to the turbine shown in red.

The primary coolant extracts heat from the core through natural circulation between four steam separators by way of down-coming and returning pipes, which connect to the reactor vessel at a header. During normal operation steam is separated from the loop and directed to the turbine building to generate power, after which it is condensed and pumped back to the core loop by a set of feed pumps. To initiate shutdown, an isolation valve is used to disconnect the pipe running to the turbine, instead directing coolant into the shutdown condenser, which removes decay heat during the forty day shutdown period. This shutdown process may be initiated either at the end of the normal scheduled period for maintenance or in response to component failure requiring repair. However, for the event of a severe fault requiring emergency shutdown, a reservoir is on stand-by to inject high-pressurise coolant into the core coolant channels if adequate coolant pressure is lost. This is followed by low-pressure gravity-fed injection from an overhead tank, which eventually submerges the core over the course of 3 days.

During operation, maintenance can be performed on the isolation valves and feed pumps. The isolation valves become more likely to fail (on demand or otherwise) with each passing year, but at the end of their maintenance period, they are restored to their original condition, resetting associated failure modes. Two feed pumps are always required to be online to return adequate coolant, with any additional pumps kept ready to come into use in the event of a failure. The individual pumps are periodically replaced. When a pump is retired, one of the inactive pumps comes online to assume its capacity, with the subsequent replacement pump becoming available as a new redundant pump. The replacement of pumps is staggered to occur proportionately through the scheduled period.

For the purposes of the optimisation, there are a number of variable design parameters pertaining to the levels of component redundancies and maintenance actions. A description of each parameter, its permitted range of values, and any associated costs, given in terms of an arbitrary currency unit, ₲, are found in Table 1.

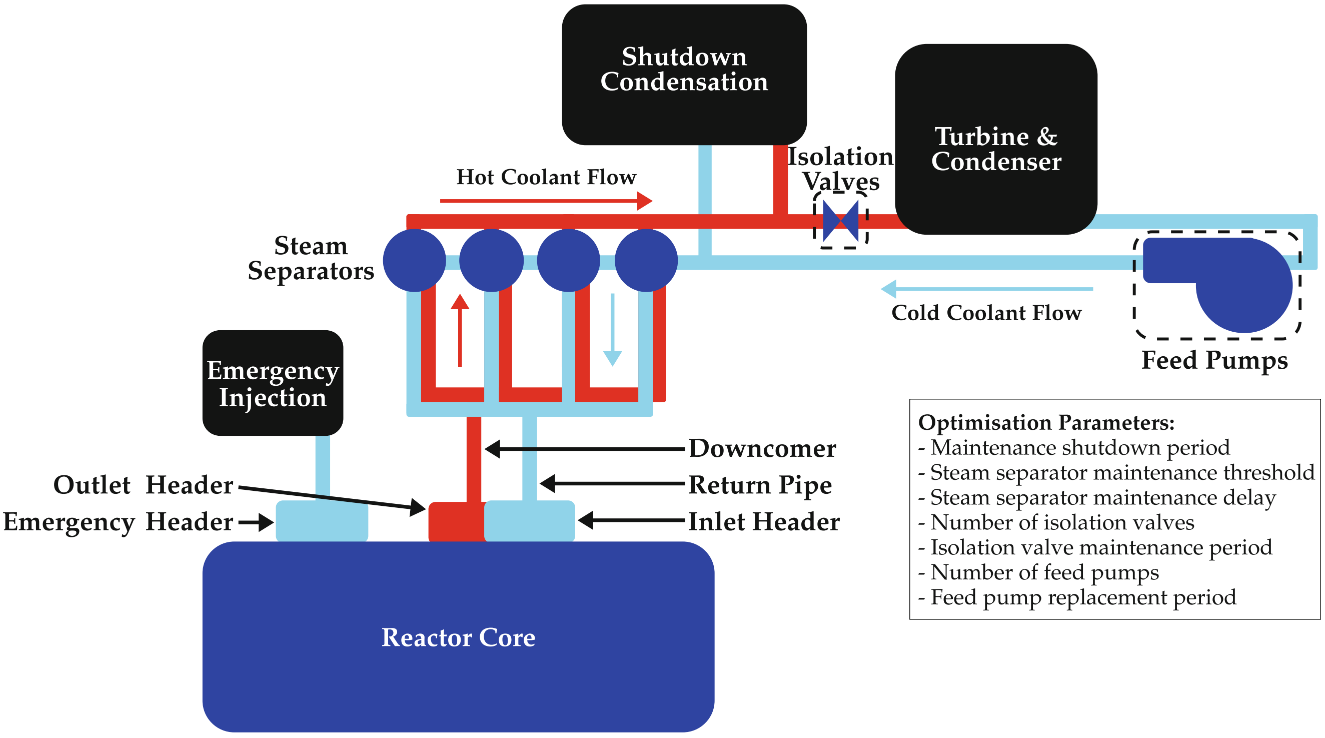

FIGURE 3 Schematic of reactor case study system with the variable design and maintenance parameters to be optimised. Red lines represent the flow of hot coolant from the core to the turbine and condenser and light blue lines represent its return to the core, with the direction of flow marked by arrows in the respective colours.

TABLE 1 Parameters for the configuration and maintenance of the primary coolant system, with their description, inclusive range of possible values, and associated costs in ₲ (arbitrary currency unit)

Parameter	Description	Range	Associated cost
m_t	**Maintenance shutdown period**—duration for which the reactor operators before routine maintenance if no need for early shutdown emerges	0.5 to 10 years	N/A
c_{ss}	**Steam separator maintenance threshold**—critical fraction of failed steam separators to schedule early maintenance	0 to 0.5[a]	N/A
m_{ss}	**Steam separator maintenance delay**—Time from critical fraction of failed steam separators to beginning of shutdown	0 to 10 years	N/A
n_{iv}	**Number of isolation valves**	1 to 10	New valve: 500 ₲
m_{iv}	**Isolation valve maintenance period**	0 to 10 years	Servicing procedure for one valve: 100 ₲
n_{fp}	**Number of feed pumps**	2 to 10	New pump: 5000 ₲
m_{fp}	**Feed pump replacement period**	0 to 10 years	Replacement pump: 5000 ₲

[a] Value rounded up to the nearest integer number of steam drums.

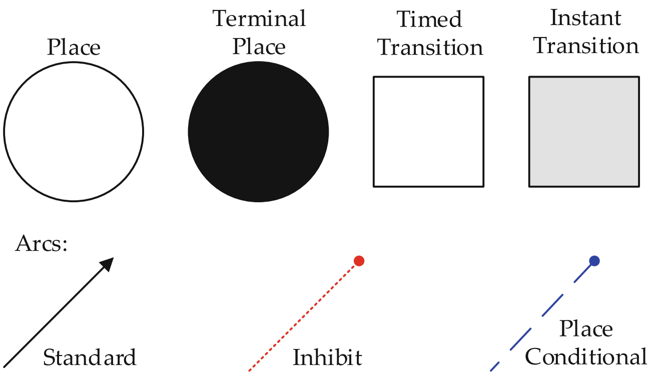

FIGURE 4 A key to the objects found in Petri net models.

These values are used to calculate the costs incurred during each batch of simulations with a particular set of parameters, including both the expenses of the initial set-up of the configuration and resulting from maintenance actions. System clock durations are tallied to produce the expenditure per unit time.

3 | PETRI NET METHODOLOGY

The Petri net modelling in this work was realised with *Macchiato*,[22] developed in-house at the University of Nottingham. The name given to the particular variation of Petri net methodology employed is *generalised stochastic Petri nets*.[30]

3.1 | Structure

A Petri net is a bipartite graph and can be used to model the evolving state of a dynamic system represented in its structure. Figure 4 shows a key to the representation of the Petri net elements as depicted in this work.

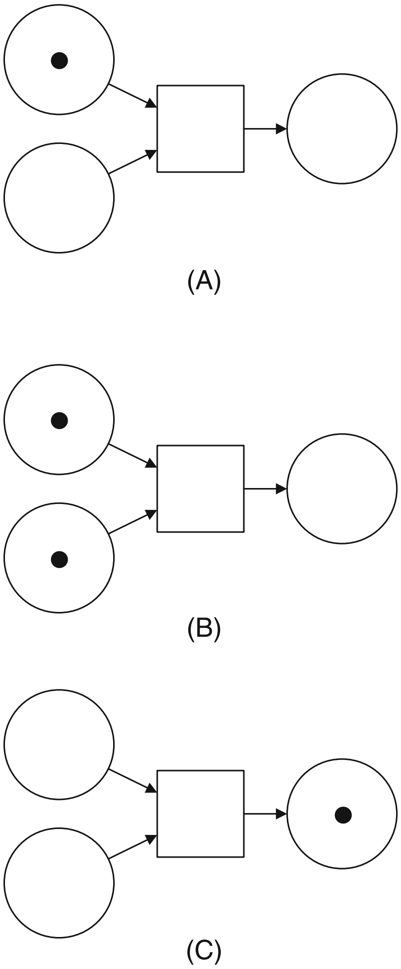

FIGURE 5 Illustration of the enabling conditions and firing of a transition. (A) The transition is not able to fire as not all arc weights are met; (B) the enabling conditions of the transition are met; (C) after firing, the incoming places have lost tokens and the outgoing place has gained one

The two objects found in a Petri net are called the *place* and the *transition*, drawn as circles and squares respectively, and the connections between them are *arcs*, drawn as arrows and referred to as *incoming* and *outgoing* relative to the connecting transition. When a place and transition are connected by both an incoming and an outgoing arc, the pair is drawn as a single double-headed arrow. An arc has a property known as *weight*, which takes a non-zero integer value. If the weight of an arc takes a value other than 1, it is marked adjacent to it. The role of the place is to hold *tokens*, visualised as black dots, with the state of the represented system at some moment being determined by marking of the places. The transitions control the placement and removal of tokens. At each step in the simulation, a transition is chosen to *fire*, and which transitions can and cannot do so is determined by the token markings of the places connected to it by arcs. A transition which is available to fire is said to be *enabled*, and for this to be so, every place connected to it by an incoming arc must hold a number of tokens greater than or equal to the weight of that arc. When a transition fires, it removes a number of tokens from each of the incoming places equal to the weight of the connecting arc. It then adds tokens to the places connected by outgoing arcs, each number of which being equal to the weights of those arcs. This process is illustrated in Figure 5. Transitions can be parameterised to fire instantaneously or following a delay, as described in more detail in Section 3.3, with these being referred to as *instant* and *timed* transitions, and coloured grey and white respectively.

In addition to the standard arcs described above, there also exist *inhibit arcs* and *place conditional arcs*, both of which can only appear as incoming arcs. A inhibit arc has the effect of disabling the transition to which it connects, preventing it from firing, see Figure 6. A place conditional arc applies a modifier to a timed transition, altering the delay between it becoming enabled and firing.

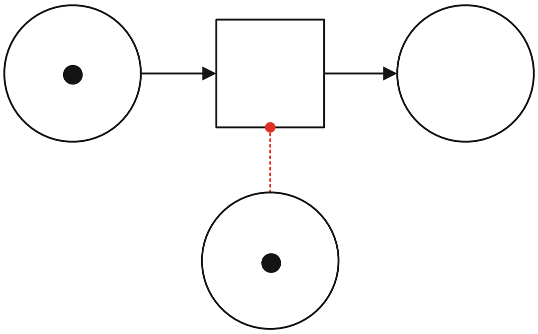

FIGURE 6 This transition cannot fire because the place connected by an inhibit arc holds a token.

3.2 | Simulation algorithm

The successive firing of transitions updates the marking of the Petri net, which represents the dynamics of the system. At each step, every transition is inspected to determine whether its enabling conditions are met. Transitions that are now enabled that were not enabled on the previous step are added to a list and a firing time is calculated according to their parameters. Any transition that was enabled on the previous step, but whose enabling conditions are no longer satisfied, is removed from the list. If that transition is re-enabled at a later step, it retains no memory of its former firing time.

When the assessment of the status of the transitions is complete, the list of enabled transitions is used to choose the next transition to fire. First, it is established whether any instant transitions are enabled. If so, one of them is chosen at random. If not, the transition whose calculated firing time is closest to the current clock is selected. In the event that multiple transitions are scheduled to fire at the same time, one is again chosen at random. Having selected a transition, the token marking of the places is updated in accordance with its firing, the simulation step is advanced by one, and the system clock is set to the scheduled time of firing. This continues until a terminal place, marked with black fill, receives a token, at which point the simulation ends. The algorithm described here is shown in flowchart form in Figure 7.

Given the stochastic nature of the Monte Carlo simulation routine it is necessary to perform many iterations, until convergence is reached, at which point statistical data may be extracted from the body of results.

3.3 | Transition firing delay

Four distributions are used for firing times in the modelling presented in the article, referred to as "delay", "uniform", "cyclic", and "Weibull". These distributions are adequate to capture the salient feature of the system. Generally, the delay and cyclic distributions are used to control the timing of system actions, for example, the duration of a process and the regular occurrence of maintenance respectively, while the uniform (when paired with a delay) and Weibull distributions produce the probabilities associated with failure events, for example, the likelihood that an action will be performed successfully and the time at which a component will experience a failure mode respectively.

A delay transition fires after a fixed duration a and a transition with a uniform distribution fires at a time between 0 and u after becoming enabled, such that its probability density, $f(t; u)$ is given by,

$$f(t;u) = \begin{cases} \frac{1}{u} & \text{for} \quad t \in (0,u] \\ 0 & \text{otherwise} \end{cases}. \tag{1}$$

A cyclic transition fires after becoming enabled at a system clock value equal to the next integer multiple of its parameter, c, optionally offset by a second parameter, ω. For example, a cyclic transition with $c = 2$ h and $\omega = 0$ h would fire at system clock values of 0, 2, 4, 6 and so forth whereas with $c = 2$ h and $\omega = 1$ h, the firings would occur at 1, 3, 5, 7 and so forth assuming that in both cases the transition was enabled at those times.

The Weibull distribution[31,32] is commonly used for modelling failure times in reliability engineering, and generalises the exponential distribution for non-constant failure rates. It has two parameters η and β, which are respectively the scale parameter and the shape parameter. The former characterises the point at which approximately two-thirds of a population

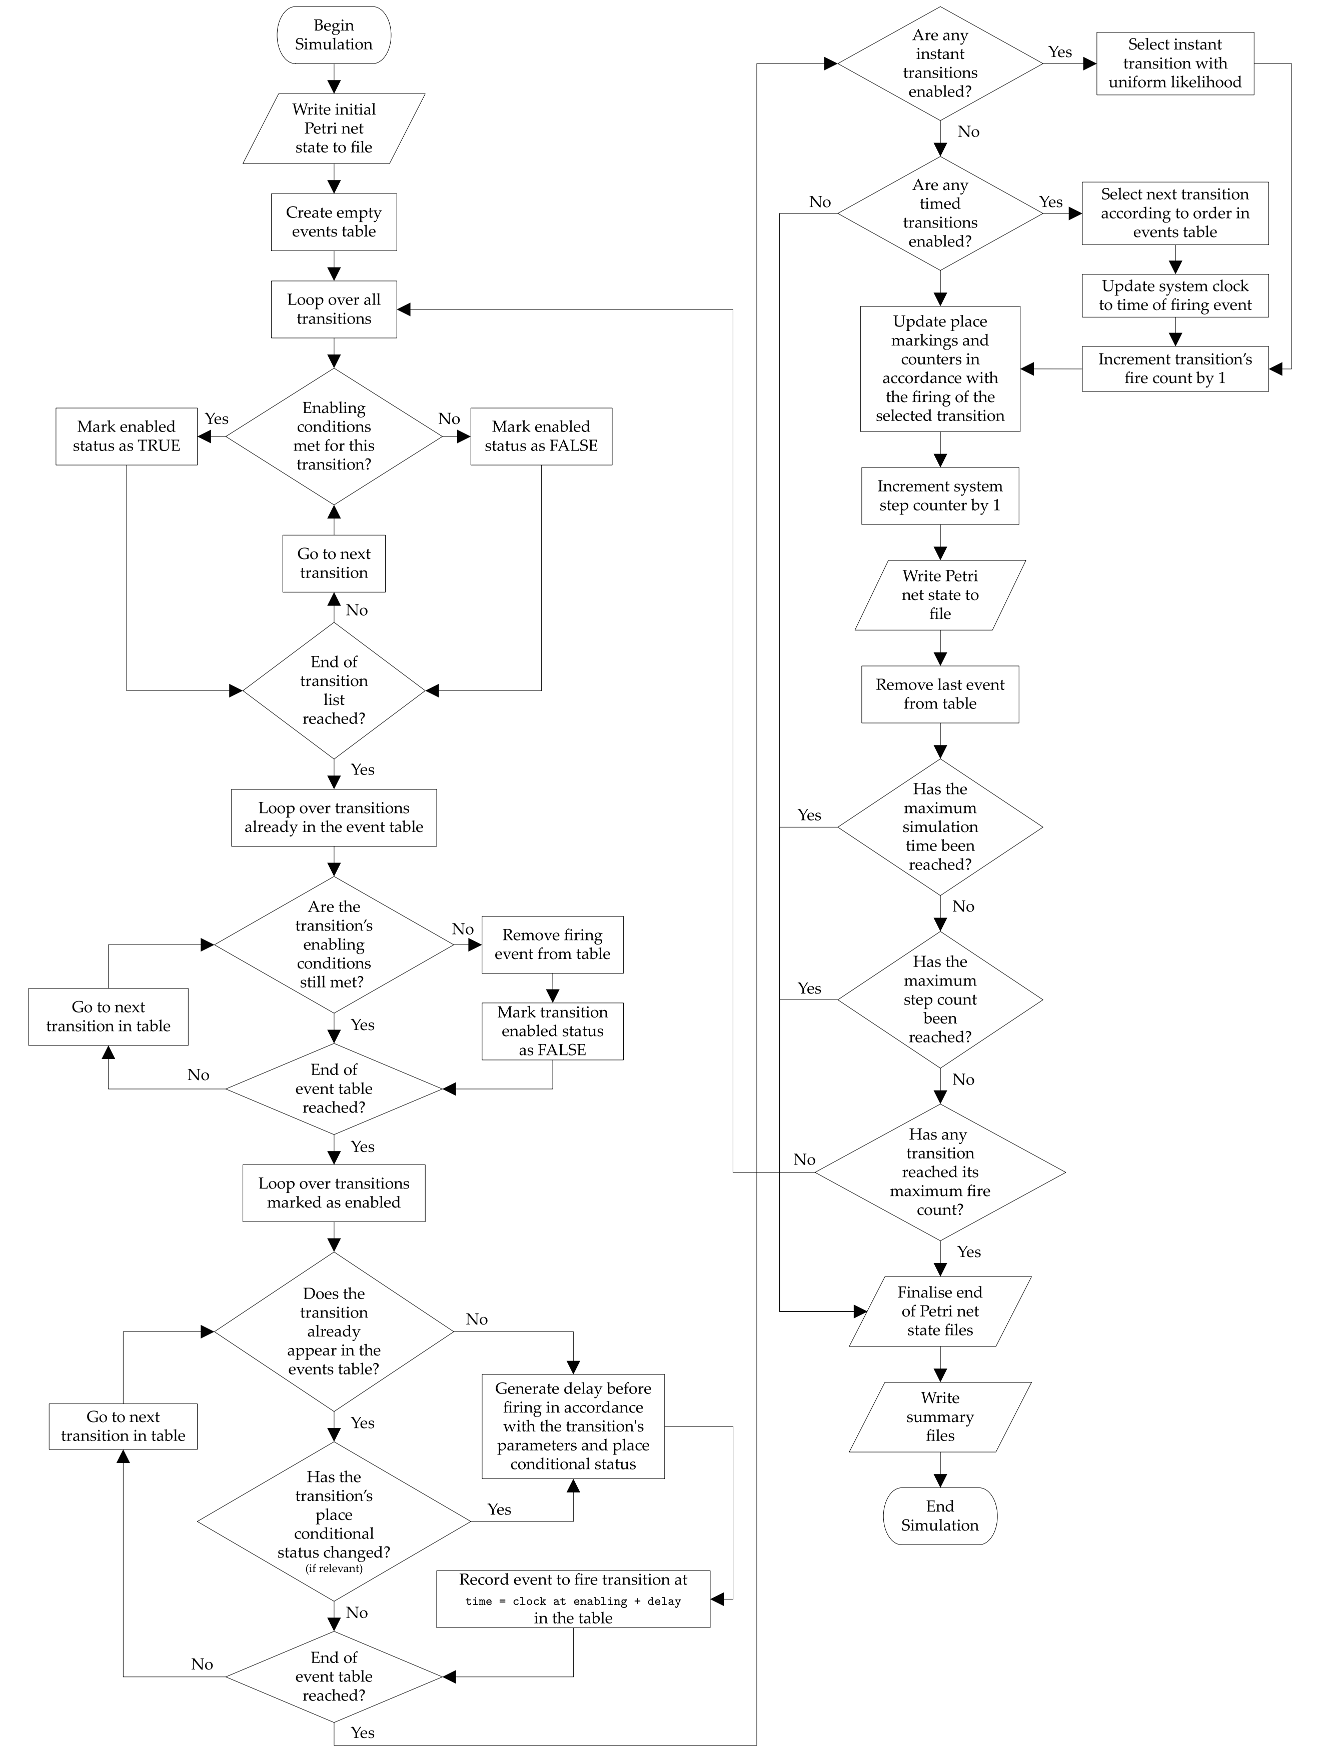
Begin Simulation
Write initial Petri net state to file
Create empty events table
Loop over all transitions
Enabling conditions met for this transition?
Yes
Mark enabled status as TRUE
No
Mark enabled status as FALSE
Go to next transition
No
End of transition list reached?
Yes
Loop over transitions already in the event table
Are the transition's enabling conditions still met?
No
Remove firing event from table
Mark transition enabled status as FALSE
Go to next transition in table
Yes
No
End of event table reached?
Yes
Loop over transitions marked as enabled
Does the transition already appear in the events table?
No
Generate delay before firing in accordance with the transition's parameters and place conditional status
Go to next transition in table
Yes
Has the transition's place conditional status changed? (if relevant)
Yes
Record event to fire transition at time = clock at enabling + delay in the table
No
No
End of event table reached?
Yes
Are any instant transitions enabled?
Yes
Select instant transition with uniform likelihood
No
No
Are any timed transitions enabled?
Yes
Select next transition according to order in events table
Update system clock to time of firing event
Update place markings and counters in accordance with the firing of the selected transition
Increment transition's fire count by 1
Increment system step counter by 1
Write Petri net state to file
Remove last event from table
Yes
Has the maximum simulation time been reached?
No
Yes
Has the maximum step count been reached?
No
No
Has any transition reached its maximum fire count?
Yes
Finalise end of Petri net state files
Write summary files
End Simulation

FIGURE 7 The algorithm used by Macchiato to integrate a Petri net model.

is expected to have failed, and the latter controls the evolution of the failure rate, such that $\beta < 1$ is indicative of infant mortality, $\beta = 1$ yields a constant failure rate, and $\beta > 1$ produces an ageing effect, giving rise to a probability density function, $f(t; \eta, \beta)$, of

$$f(t; \eta, \beta) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1} \exp\left(-\left[\frac{t}{\eta}\right]^{\beta}\right), \tag{2}$$

where t is the time of failure.

When a transition has one or more place conditional relations, a factor P is calculated, such that

$$P = 1 + \sum_{i} W_i N_i \ , \tag{3}$$

where W_i is the weight of the ith place condition arc of the transition and N_i is the corresponding token count. The resulting value is used to apply the modification by dividing a timing parameter (a, u, c, or η as relevant) by P. Unlike that of standard arcs and inhibit arcs, the weight of a place conditional arc may take an non-integer value. The purpose of this mechanism is to allow representation of arbitrary schemes of time to fire alteration with respect to any number of place markings. For example, this work contains a type of valve which becomes progressively more likely to fail to shut on demand with each passing year since maintenance was performed. This is achieved with a transition which periodically adds a token to a place until maintenance occurs. A place conditional arc connects the place to the transition that determines whether the valve functions correctly or otherwise, such that when it is enabled by a token arriving on the place that represents demand, the probability of failure generated corresponds to the elapsed duration into the maintenance cycle.

4 | PETRI NET MODEL

An example of one of the Petri net models generated in this work is shown in Figure 8 with a component configuration of two turbine isolation valves and three feed pumps. Sections of the Petri net that vary with respect to the optimisation parameters are highlighted. In Table 2, the function of each transition is briefly summarised, along with the relevant parameters used in the timed transitions.

The structure of each model represents a system of the type described in Section 2. While the parameters chosen were selected on the basis of expert opinion to be a realistic reflection of genuine nuclear reactors, their principle purpose is to facilitate the demonstrate of the methodology presented in this work, and not to provide advice about any specific real world design.

The reactor primary coolant system begins in full working order. The Petri net models the emergence of component failure modes and maintenance actions, running until one of two possible outcomes are reached, these being the safe shutdown of the reactor or a problem that requires the invocation of the emergency coolant injections systems. These are respectively labelled as “Safe Shutdown” and “Coolant Fault”. Safe shutdown is reached following the end of shutdown condensation, which is requested either when the end of the maintenance shutdown period is reached, see TSM, or if a fault arises, for example, a burst pipe or a critical number of failed feed pumps. If a sufficiently severe fault occurs in a part of the system that cannot be isolated, such as a rupture in the internal coolant channels of the core or the loss of more than half of the steam separator circuits, shutdown condensation is insufficient to maintain safe core temperatures, necessitating the use of the injection system and resulting in a “Coolant Fault” outcome. In this model, this outcome is considered a failed state for the primary coolant circulation system.

The top left section of the Petri net models the four steam circuits. The four failure modes are the rupture of the steam separator, rupture of the downcomer, rupture of the return pipe, and spurious opening of the pressure release valve, with their occurrences respectively represented by the transitions TSS[1, 2, 3, 4], TDC[1, 2, 3, 4], TRP[1, 2, 3, 4], and TSSPR[1, 2, 3, 4]. The failure of a circuit is recorded by TCP[1, 2, 3, 4], adding a token to PCC9. When the first token arrives there, it enables TMSS1, setting a countdown, at the end of which reactor shutdown is requested. The length of this countdown is parameterised by m_{ss}. However, if additional steam circuit failures arise, this period is cut short when the token count at PCC9 reaches the threshold set by c_{ss}, with the firing of TMSS2 requesting immediate shutdown.

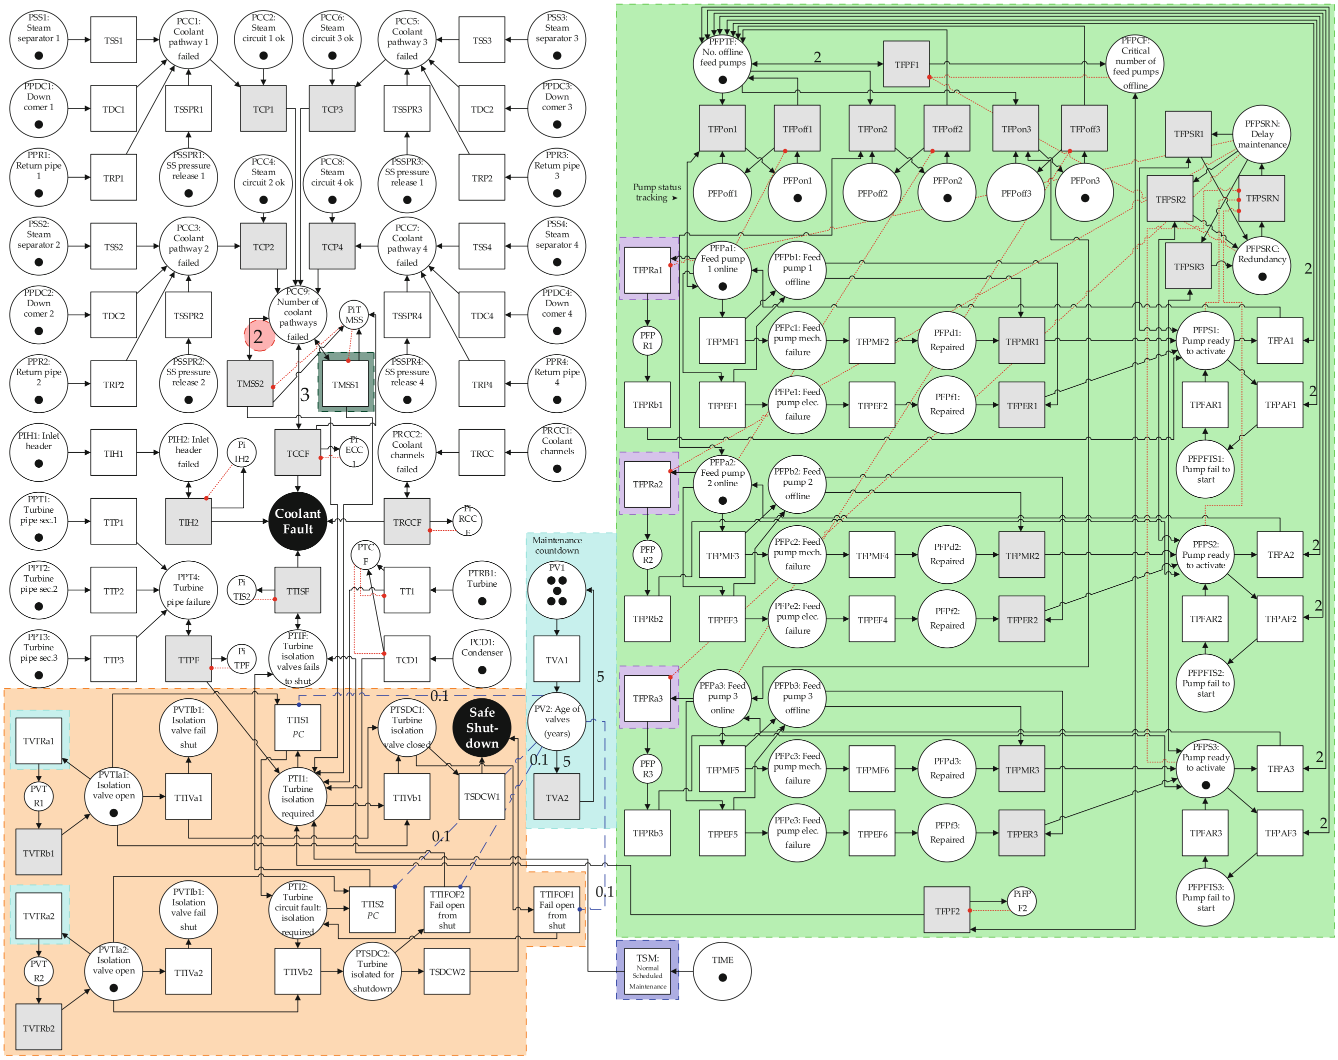
PSS1: Steam separator 1
TSS1
PCC1: Coolant pathway 1 failed
PCC2: Steam circuit 1 ok
PCC6: Steam circuit 3 ok
PCC5: Coolant pathway 3 failed
TSS3
PSS3: Steam separator 3
PPDC1: Down comer 1
TDC1
TSSPR1
TCP1
TCP3
TSSPR3
TDC2
PPDC3: Down comer 3
PPR1: Return pipe 1
TRP1
PSSPR1: SS pressure release 1
PCC4: Steam circuit 2 ok
PCC8: Steam circuit 4 ok
PSSPR3: SS pressure release 1
TRP2
PPR3: Return pipe 3
PSS2: Steam separator 2
TSS2
PCC3: Coolant pathway 2 failed
TCP2
TCP4
PCC7: Coolant pathway 4 failed
TSS4
PSS4: Steam separator 4
PPDC2: Down comer 2
TDC2
TSSPR2
PCC9: Number of coolant pathways failed
PiT MSS
TSSPR4
TDC4
PPDC4: Down comer 4
PPR2: Return pipe 2
TRP2
PSSPR2: SS pressure release 2
2
TMSS2
3
TMSS1
PSSPR4: SS pressure release 4
TRP4
PPR4: Return pipe 4
PIH1: Inlet header
TIH1
PIH2: Inlet header failed
Pi IH2
TCCF
Pi ECC 1
PRCC2: Coolant channels failed
TRCC
PRCC1: Coolant channels
PPT1: Turbine pipe sec.1
TTP1
TIH2
Coolant Fault
TRCCF
Pi RCC F
PPT2: Turbine pipe sec.2
TTP2
PPT4: Turbine pipe failure
Pi TIS2
TTISF
PTC F
TT1
PTRB1: Turbine
PPT3: Turbine pipe sec.3
TTP3
TTPF
Pi TPF
PTIF: Turbine isolation valves fails to shut
TCD1
PCD1: Condenser
Maintenance countdown
PV1
TVA1
5
PV2: Age of valves (years)
0.1
5
TVA2
0.1
TVTRa1
PVT R1
TVTRb1
PVTIa1: Isolation valve open
PVTIb1: Isolation valve fail shut
TTIVa1
TTIS1 PC
PTI1: Turbine isolation required
PTSDC1: Turbine isolation valve closed
TTIVb1
Safe Shut-down
TSDCW1
0.1
TVTRa2
PVT R2
TVTRb2
PVTIa2: Isolation valve open
PVTIb1: Isolation valve fail shut
TTIVa2
PTI2: Turbine circuit fault: isolation required
TTIS2 PC
TTIFOF2 Fail open from shut
TTIFOF1 Fail open from shut
0.1
TTIVb2
PTSDC2: Turbine isolated for shutdown
TSDCW2
TSM: Normal Scheduled Maintenance
TIME
PFPTF: No. offline feed pumps
2
TFPF1
PFPCF: Critical number of feed pumps offline
TFPon1
TFPoff1
TFPon2
TFPoff2
TFPon3
TFPoff3
TFPSR1
PFPSRN: Delay maintenance
Pump status tracking ➤
PFPoff1
PFPon1
PFPoff2
PFPon2
PFPoff3
PFPon3
TFPSR2
TFPSRN
TFPRa1
PFPa1: Feed pump 1 online
PFPb1: Feed pump 1 offline
TFPSR3
PFPSRC: Redundancy
2
PFP R1
TFPMF1
PFPc1: Feed pump mech. failure
TFPMF2
PFPd1: Repaired
TFPMR1
PFPS1: Pump ready to activate
TFPA1
TFPRb1
TFPEF1
PFPe1: Feed pump elec. failure
TFPEF2
PFPf1: Repaired
TFPER1
TPFAR1
TFPAF1
2
PFPFTS1: Pump fail to start
TFPRa2
PFPa2: Feed pump 2 online
PFPb2: Feed pump 2 offline
PFP R2
TFPMF3
PFPc2: Feed pump mech. failure
TFPMF4
PFPd2: Repaired
TFPMR2
PFPS2: Pump ready to activate
TFPA2
2
TFPRb2
TFPEF3
PFPe2: Feed pump elec. failure
TFPEF4
PFPf2: Repaired
TFPER2
TPFAR2
TFPAF2
2
PFPFTS2: Pump fail to start
TFPRa3
PFPa3: Feed pump 3 online
PFPb3: Feed pump 3 offline
PFP R3
TFPMF5
PFPc3: Feed pump mech. failure
TFPMF6
PFPd3: Repaired
TFPMR3
PFPS3: Pump ready to activate
TFPA3
2
TFPRb3
TFPEF5
PFPe3: Feed pump elec. failure
TFPEF6
PFPf3: Repaired
TFPER3
TPFAR3
TFPAF3
2
PFPFTS3: Pump fail to start
TFPF2
PiFP F2

FIGURE 8 The Petri net model, with two isolation valves and three feed pumps. Areas varying with respect to the optimisation parameters are highlighted, pertaining in accordance with colour to m_t, c_{ss}, m_{ss}, n_{iv}, m_{iv}, n_{fp}, and m_{fp}. See Table 2 for parameters. Small size places are used for miscellaneous house-keeping tasks such as controlling the number of times a transition can fire.

The number of failed steam circuits reaching three constitutes a disruption to coolant circulation beyond that which can be managed within the primary coolant system, thus resulting in the "Coolant Fault" outcome.

Below the steam circuit section, the transitions TIH1 and TRCC model the failure of the reactor inlet header and the internal coolant channels, either of which causing an immediate "Coolant Fault" outcome. However, the turbine, the condenser, and the pipe to them from the reactor are placed beyond the turbine isolation valve, meaning that if a fault arises in one of these, the reactor can still be shut down by normal means. The transitions for those components are TT1, TCD1, and TTP[1, 2, 3].

The section within the orange box pertains to turbine isolation, with the total number of isolation valves being set by n_{iv}. When a request to shutdown the reactor is placed, a token is added to PTI1. The first of the isolation valves will shut, with the transition TTIVb1 firing. If the valve fails to shut, see TTIS1, or if it fails to an open state once shut, see TTIFOF1, reactor isolation is interrupted. If there is a redundant valve available, an attempt will be made to close it instead, and if that should also fail, the next redundant valve will be used, until no redundancy remains, at which point a token is placed at PTIF, indicating the failure of turbine isolation and resulting in a "Coolant Fault" outcome. Otherwise, once the turbine has been isolated for forty days, "Safe Shutdown" is achieved. A turbine isolation valve can also close spuriously, see TTIVa[1, … , n_{iv}], forcing the reactor to shut down before its normal scheduled time. The valves are serviced periodically, according to the parameter m_{iv}, which sets the firing delay of TVTRa[1, … , n_{iv}], the weight of the arcs connecting to TVA2, and the initial number of tokens at PV1. The former resets the time of occurrence of spurious valve closure and

TABLE 2 Parameters for the Petri net Model show in Figure 8

Transition(s)	Description	Distribution	Parameters	Source(s)
TSS[1, 2, 3, 4]	Failure of steam separator	Weibull	$\eta = 9.32\times10^{7}$, $\beta = 1.2$	33[b]
TDC[1, 2, 3, 4]	Failure of downcomer pipe	Weibull	$\eta = 4.38\times10^{8}$, $\beta = 1.0$	34,35
TRP[1, 2, 3, 4]	Failure of return pipe	Weibull	$\eta = 1.75\times10^{6}$, $\beta = 1.0$	34,35
TSSPR[1, 2, 3, 4]	Failure of pressure release valve	Weibull	$\eta = 1.00\times10^{5}$, $\beta = 1.0$	34,35
TCP[1, 2, 3, 4]	Steam separator coolant pathway becomes unavailable	Instant	N/A	N/A
TMSS1	Scheduled shutdown demand after critical threshold of steam separators reached	Delay	$a = m_{ss}$	N/A
TMSS2	Immediate shutdown demand at loss of half of steam separator capacity	Instant	N/A	N/A
TCCF	Failure of primary coolant loop due to loss of greater than half of steam separator capacity	Instant	N/A	N/A
TTP1	Turbine pipe section 1 failure	Weibull	$\eta = 1.75\times10^{8}$, $\beta = 1.0$	34,35
TTP2	Turbine pipe section 2 failure	Weibull	$\eta = 1.66\times10^{8}$, $\beta = 1.0$	34,35
TTP3	Turbine pipe section 3 failure	Weibull	$\eta = 1.00\times10^{8}$, $\beta = 1.0$	34,35
TT1	Turbine fault	Weibull	$\eta = 1.13\times10^{6}$, $\beta = 1.7$	33[b]
TCD1	Condenser fault	Weibull	$\eta = 1.07\times10^{6}$, $\beta = 1.2$	33[b]
TVTRa[1, …, n_{iv}]	Maintenance of isolation valves	Delay	$a = m_{iv}$	N/A
TVTRb[1, …, n_{iv}]	Restoration of isolation valves	Instant	N/A	N/A
TTIVa[1, …, n_{iv}]	Isolation valve fails shut from open	Weibull	$\eta = 1.51\times10^{9}$, $\beta = 1.1$	33[b]
TTIVb[1, …, n_{iv}]	Isolation valve closes on demand	Delay	$a = 3\times10^{-4}$	34
TTIS[1, …, n_{iv}]	Isolation valve fails to shut on demand	Uniform	$u = 0.13$	34
TTIFOF[1, …, n_{iv}]	Isolation valve fails open from shut	Weibull	$\eta = 1.75\times10^{5}$, $\beta = 1.0$	36[b]
TSDCW[1, …, n_{iv}]	Shutdown condensation period	Delay	$a = 960$	N/A
TVA1	Yearly counter for valve condition	Delay	$a = 8766$	N/A
TVA2	Valve condition reset	Instant	N/A	N/A
TFPMF[1, 3, …, $(2n_{fp}-1)$]	Mechanical failure of feed pump	Weibull	$\eta = 1.40\times10^{4}$, $\beta = 1.2$	33,34
TFPMF[2, 4, …, $2n_{fp}$]	Repair of mechanical pump failure	Delay	$a = 24.0$	34
TFPMR[1, …, n_{fp}]	Pump returned to availability	Instant	N/A	N/A
TFPEF[1, 3, …, $(2n_{fp}-1)$]	Electrical failure of feed pump	Weibull	$\eta = 1.06\times10^{4}$, $\beta = 1.2$	33,37
TFPEF[2, 4, …, $2n_{fp}$]	Repair of electrical pump failure	Delay	$a = 24.0$	34
TFPER[1, …, n_{fp}]	Pump returned to availability	Instant	N/A	N/A
TFPA[1, …, n_{fp}]	Feed pump successfully brought online	Delay	$a = 3\times10^{-4}$	34
TFPAF[1, …, n_{fp}]	Failure to start feed pump	Uniform	$u = 0.63$	34
TPFAR[1, …, n_{fp}]	Pump returned to availability	Delay	$a = 20.9$	34
TFPRa[1, …, n_{fp}]	Pump servicing begins	Cyclic	$c = m_{fp}$ [a]	N/A
TFPRb[1, …, n_{fp}]	Pump servicing ends	Delay	$a = 20.9$	34
TFPon[1, …, n_{fp}]	Record pump coming online	Instant	N/A	N/A
TFPoff[1, …, n_{fp}]	Record pump going offline	Instant	N/A	N/A
TFPSR[1, …, n_{fp}]	Record pump available for activation	Instant	N/A	N/A
TFPSRN	Record no pumps available for activation	Instant	N/A	N/A

(Continues)

TABLE 2 (Continued)

Transition(s)	Description	Distribution	Parameters	Source(s)
TFPF1	Record critical number of pump failures	Instant	$^{N}/_{A}$	$^{N}/_{A}$
TFPF2	Call for reactor shutdown due to feed pump critical failure	Instant	$^{N}/_{A}$	$^{N}/_{A}$
TSM	Call for reactor shutdown after normal operational period	Delay	$a = m_t$	$^{N}/_{A}$

Note: Times are given in hours where relevant.

[a] Values of ω are selected to be evenly staggered within feed the pump replacement period.

[b] Failure rate based on expert opinion.

the latter two cycle the number of tokens at PV2, which has a place conditional relationship with the other failure modes of the valves, increasing their likelihood later in the maintenance period, with TVA1 firing once a year.

The green area of the right half of the Petri net models the state of the feed pumps, the total number of which being controlled by $n_f m$. The two failure modes, being mechanical and electrical breakdowns, are represented by TFPMF[1, 3, … , $(2n_{fp}-1)$] and TFPEF[1, 3, … , $(2n_{fp}-1)$] respectively, with the repairs of which likewise being TFPMF[2,4, … , $2n_{fp}$] and TFPEF[2, 4, … , $2n_{fp}$]. TFPon[1, … , n_{fp}] and TFPoff[1, … , n_{fp}] are responsible for recording the state of each pump and when a pump fails, a token is added to PFPTF, which prompts the activation of a replacement if one is available. TFPSR[1, … , n_{fp}], TFPSRN, and PFPSRC track whether a redundant pump is available, and if fewer than two are running without a ready replacement, TFPF1 fires to request reactor shutdown. Otherwise, either TFPA[1, … , n_{fp}] or TFPAF[1, … , n_{fp}] will fire, respectively representing successful activation or a failure to start. In the latter case, TPFAR[1, … , n_{fp}] models the time taken to prepare for a subsequent attempt. TFPRa[1, … , n_{fp}] and TFPRb[1, … , n_{fp}] model the periodic maintenance of the pumps, respectively taking each offline in accordance with the parameter m_{fp}, and then returning it to a state ready for use.

5 | DATA GENERATION AND PARETO OPTIMISATION

A large body of data was generated for the optimisation process. The system designs to be explored during the optimisation are achieved by Monte Carlo sampling[26,27] of the parameter space. Each set of parameters is selected randomly within the ranges specified in Table 1. Having generated a set of optimisation parameters, the corresponding Petri net structure was generated, as discussed in Section 4, and was simulated 10^6 times. On the order of 10^5 iterations would be sufficient when typical system parameter values are used,[22] but to guarantee that an adequate body of data is collected even in the case of the more exotic parameter combinations, this larger sample size was used, with the total expenditure of computational time summing to 15,000 core hours on the Hydra HPC system.

A total of 4000 parameter sets were investigated in this fashion and the optimisation was subsequently performed via Pareto front[29] analysis with respect to the performance metrics of chief concern, namely the probability of safe shutdown and the monetary expenditure per unit time. The Pareto front is the line running through the set of data points which are Pareto optimal,[28] with an example given in Figure 9. A point is member of this set if it has the best possible score with respect to one performance metric that can be achieved for a given value of the other metric and vice versa, such that an improvement of either score can only be achieved to the detriment of the other metric. By examining the Pareto front, the optimal performance in terms of probability of safe shutdown can be achieved for a particular spending target, and counterwise, the most inexpensive parameter set for a given safe shutdown probability can be found.

6 | RESULTS AND DISCUSSION

The results presented in this section are analysed using Spearman's rank correlation coefficient,[38] also known as Spearman's ρ, to find the strength of relation between each of the optimisation parameters and the outcome metrics, and between the latter themselves. Spearman's ρ gives the Pearson product-moment correlation coefficient[39] (PPMCC) of the rankings of every data point when ordered by its two sets of values. Unlike the PPMCC itself, Spearman's ρ is suitable

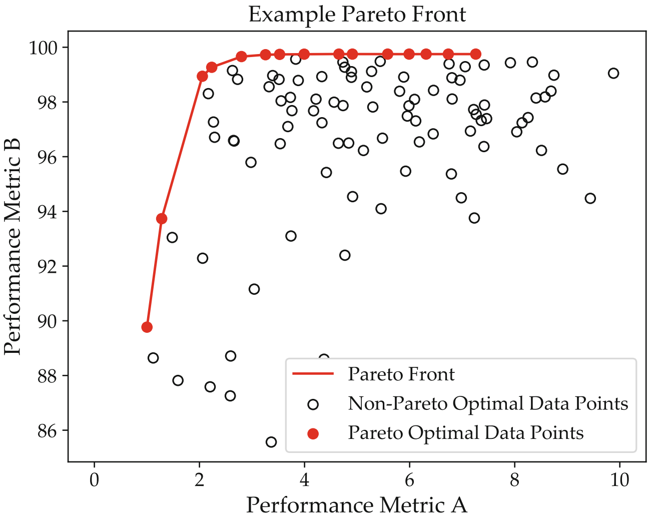

FIGURE 9 Example of Pareto front with 100 randomly generated values. The goal is to minimise performance metric A, while maximising performance B. Therefore, the Pareto optimal data points are found at the top and left of the plot.

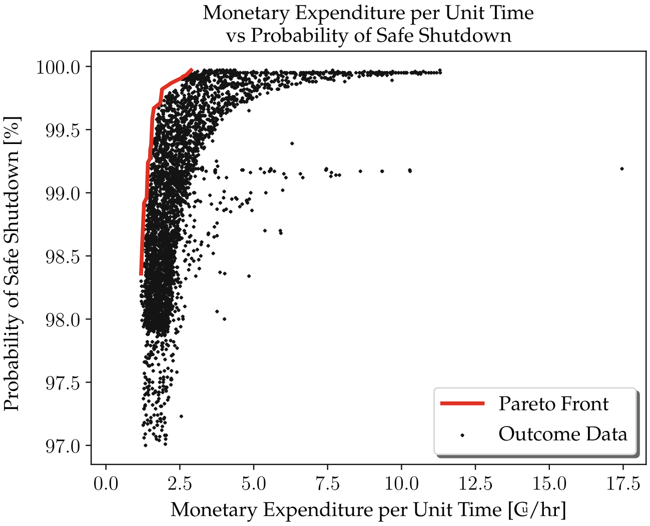

FIGURE 10 Comparison of the main performance metrics from the optimisation process, expenditure per unit time and reliability, with Pareto front drawn.

for non-linear correlation. It yields a value in the range −1 to 1, where −1 is perfect negative monotonic correlation, 0 is no monotonic correlation, and 1 is perfect positive monotonic correlation. The confidence intervals were computed following the method described by Ruscio 2008.[40] The *p*-value is also given, which is the probability that greater or equal correlation than the given value of Spearman's ρ would be observed in the case that the null hypothesis of no correlation were true. A *p*-value below 0.05 is typically used as the threshold to establish statistical significance.[41]

A scatter plot comparing the cost per hour and reliability results from the optimisation process is found in Figure 10, along with the Pareto front for highest reliability for lowest cost per unit time.

The set of Pareto optimal configurations has 27 members and the full list is given in Table 3. Safe shutdown probability and monetary expenditure per unit time are respectively seen to range from 98.36% to 99.97% and 1.187 to 2.878 Cu h^{-1}, and it is clear that increased expenditure yields improved safety by this measure, with the overall Spearman's rank correlation coefficient between these two values being 0.7682, with a 95% confidence interval from 0.7533 to 0.7823, and $p < 10^{-308}$. It is also seen that low numbers of isolation valves and feed pumps are generally favoured, as most Pareto optimal configurations include fewer than four of either of these. Notably however, the average duration of operation

TABLE 3 Full list of optimisation parameter sets found to be Pareto optimal with respect to safe shutdown probability and monetary expenditure per unit time, and their score with respect to all performance metrics.

Optimisation parameter							Performance metric				
Maintenance shutdown period (m_t) (years)	S.S.C. maintenance threshold (c_{ss})	S.S.C. maintenance delay (m_{ss}) (years)	№ isolation valves (n_{iv})	Isolation valve maintenance period (m_{iv}) (years)	№ feed pumps (n_{fp})	Feed pump replacement period (m_{fp}) (years)	Average operating duration (years)	Total cost [Cu]	Safe shutdown probability per unit monetary expenditure (%·Cu^{-1})	Safe shutdown probability (%)	Monetary expenditure per unit time (Cu h^{-1})
7.588	0.1074	6.069	3	2.734	3	0.2198	4.546	4.73×10^4	21.64	98.36	1.187
5.578	0.3327	6.85	3	5.102	3	0.4522	4.134	4.466×10^4	23.85	98.61	1.232
7.227	0.1983	1.955	4	8.462	3	0.4380	3.758	4.239×10^4	26.32	98.92	1.287
4.274	0.2892	6.720	8	4.779	3	0.3789	3.502	4.245×10^4	28.26	98.96	1.383
3.251	0.1963	8.33	3	2.126	3	0.3796	2.891	3.607×10^4	34.33	99.24	1.423
3.092	0.1851	3.062	3	7.371	3	0.8859	2.822	3.675×10^4	35.17	99.27	1.485
2.968	0.2497	3.962	5	5.912	3	0.3846	2.700	3.554×10^4	36.8	99.34	1.502
7.927	0.1577	0.7128	3	2.486	3	4.843	3.189	4.273×10^4	31.17	99.38	1.529
7.743	0.2488	0.3249	9	4.534	3	0.3493	2.773	3.815×10^4	35.91	99.59	1.569
6.454	0.1995	0.1784	3	6.872	3	3.557	2.658	3.777×10^4	37.5	99.67	1.621
5.135	0.1818	0.1749	10	8.659	3	5.779	2.499	4.022×10^4	39.91	99.71	1.837
3.490	0.2245	0.04059	6	3.692	3	4.123	2.067	3.442×10^4	48.3	99.82	1.900
1.868	0.00153	0.1290	5	3.565	3	8.969	1.510	2.937×10^4	66.13	99.87	2.218
3.246	0.1229	7.992	2	5.893	2	2.342	0.7350	1.656×10^4	135.9	99.91	2.571
6.764	0.2642	8.208	2	5.904	2	6.963	0.7347	1.657×10^4	136.0	99.91	2.572
3.198	0.2505	2.192	2	7.439	2	8.719	0.7347	1.656×10^4	136.0	99.91	2.572
6.504	0.07483	7.059	2	6.344	2	7.605	0.7348	1.656×10^4	136.0	99.91	2.571
8.563	0.09443	8.689	2	3.719	2	5.298	0.7351	1.656×10^4	135.9	99.91	2.57
8.746	0.1310	7.832	2	3.665	2	8.486	0.7350	1.657×10^4	135.9	99.91	2.571
8.706	0.1602	8.813	2	1.899	2	6.708	0.7353	1.657×10^4	135.9	99.91	2.571
8.092	0.2657	1.983	2	3.020	2	2.203	0.7354	1.657×10^4	135.9	99.91	2.57
3.214	0.1151	2.657	2	4.362	2	6.402	0.7339	1.656×10^4	136.1	99.91	2.574
6.496	0.3612	5.603	2	8.819	2	3.792	0.7346	1.657×10^4	136	99.91	2.572
7.210	0.1182	1.012	2	0.6579	2	6.381	0.7261	1.658×10^4	137.6	99.92	2.604
1.158	0.1538	5.349	2	2.813	2	1.443	0.6705	1.592×10^4	149.0	99.93	2.709
5.692	0.2499	0.1732	2	0.5178	2	2.347	0.6688	1.610×10^4	149.4	99.94	2.746
6.070	0.2497	0.02583	3	7.975	2	0.789	0.6449	1.627×10^4	155.0	99.97	2.878

Abbreviation: S.S.C, steam separator circuit.

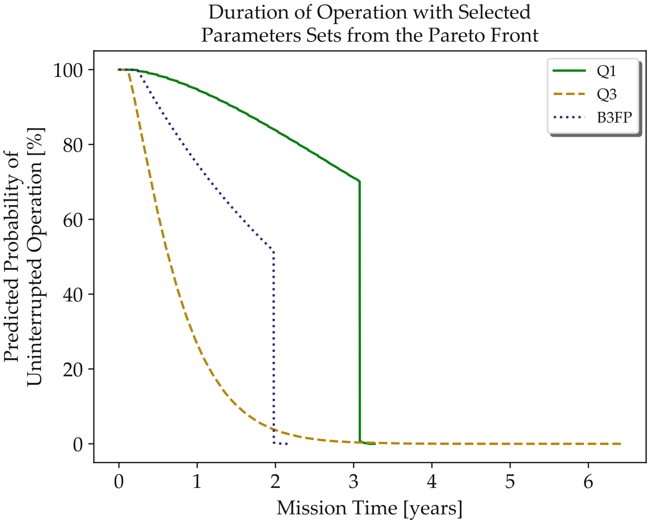

FIGURE 11 Probability of the primary coolant system operating uninterrupted with respect to time (i.e., not encountering a critical failure or the need to shut down for repair) for the first and third quartiles (Q1 & Q3) of safe shutdown probability on the Pareto front, see Table 3. Note that reactor shutdown is scheduled to begin after 2.968 and 6.352 years of operation for Q1 and Q3 respectively. The outcome with the best safety performance out of Pareto optimal configurations with three feed pumps (B3FP) is also included, with a scheduled maintenance time of 1.868 years.

declines substantially as safe shutdown probability increases. To examine what happens within primary coolant systems with such configurations, further simulations with parameter sets taken from the Pareto optimal values were performed, corresponding to the first, second, and third quartiles with respect to safe shutdown probability (as before, 10^6 iterations were performed). The first quartile has a value of 99.36% and both the second and third quartiles have values of 99.91% due a plateau in the Pareto front, where many parameter sets result in the same performance; for this reason the latter two will be referred to collectively as the third quartile hereon. The first quartile falls between two parameter sets, seen as the seventh and eighth rows in the body of Table 3. As this is the lower quartile, the parameters for the former were selected, that is, $m_t = 2.968$ years, $c_{ss} = 0.2497$, $m_{ss} = 3.962$ years, $n_{iv} = 5$, $m_{iv} = 5.912$, $n_{fp} = 3$, and $m_{fp} = 0.3846$ years. For the simulations representing the third quartile, the mean values for those parameter sets were used, such that $m_t = 6.353$ years, $c_{ss} = 0.1840$, $m_{ss} = 6.103$ years, $n_{iv} = 2$, $m_{iv} = 5.106$ years, $n_{fp} = 2$, and $m_{fp} = 5.852$ years. In Figure 11, the predicted durations of operation are seen.

There is a stark difference in the behaviour of these two parameters sets, with the probability of the system continuing to operate for a given duration dropping off rapidly for the third quartile results, with the mean duration being just 0.7346 years. It seems that with there being only two feed pumps, and therefore no redundancy, this system configuration can only operate until the first instance of a feed pump failure mode. Given that these occur quickly compared to other failure modes, see Table 2, the rest of the system is still in pristine condition and therefore the risk of unsafe shutdown (i.e., a "Coolant Fault" outcome) is low. The results from the first quartile parameters indicate a much more gradual decrease of the probability of the primary coolant system continuing to function with respect to time, such that there is a likelihood of 70% that the reactor will still be operating at the end of the scheduled maintenance shutdown period. Interestingly, the steam separator maintenance delay (m_{ss}) is set to a higher value than the maintenance shutdown period (m_t), meaning than no precautionary shutdown occurs after the maintenance threshold (c_{ss}) is met. However, with this configuration, there is a 0.66% probability of a coolant fault requiring the use of emergency mechanisms. This probability can be reduced to 0.13% that is, an 80% improvement, by increasing monetary expenditure by 48% from 1.502 to 2.218 Cu h^{-1}. This gives the best safety on the Pareto front for configurations with three feed pumps, although this does bring the average operating duration down by 44% from 2.700 years to 1.510 years. Additional simulations with this configuration were also performed and are also seen in Figure 11, with the relevant parameters being $m_t = 1.868$ years, $c_{ss} = 0.001530$, $m_{ss} = 0.1290$ years, $n_{iv} = 5$, $m_{iv} = 3.565$, $n_{fp} = 3$, and $m_{fp} = 8.969$ years. As expected, a faster drop off in the probability of uninterrupted operation is seen, but there remains a likelihood slightly greater than 50% that the scheduled maintenance time will be reached.

TABLE 4 Spearman's rank correlation coefficient between each design parameter and five performance metrics, with *p*-value given

			95% confidence interval		
Performance metric	**Optimisation parameter**	**Spearman's ρ**	**Lower**	**Upper**	***p*-value**
Average duration	Maintenance shutdown period (m_t)	0.7839	0.7698	0.7972	$<10^{-308}$
	S.S.C. maintenance threshold (c_{ss})	0.0782	0.0473	0.1090	7.337×10^{-7}
	S.S.C. maintenance delay (m_{ss})	0.0892	0.0583	0.1200	1.571×10^{-8}
	Isolation valves (n_{iv})	−0.0134	−0.0444	0.0176	0.3976
	Isolation valve maintenance period (m_{iv})	−0.0026	−0.0336	0.0284	0.8710
	Feed pumps (n_{fp})	0.3011	0.2720	0.3296	1.366×10^{-84}
	Feed pump replacement period (m_{fp})	−0.0295	−0.0605	0.0015	0.06181
Total cost	Maintenance shutdown period (m_t)	0.5197	0.4951	0.5434	1.287×10^{-275}
	S.S.C. maintenance threshold (c_{ss})	0.0560	0.0250	0.0868	0.0003974
	S.S.C. maintenance delay (m_{ss})	0.0637	0.0327	0.0945	5.613×10^{-5}
	Isolation valves (n_{iv})	0.0783	0.0474	0.1091	7.135×10^{-7}
	Isolation valve maintenance period (m_{iv})	−0.0300	−0.0610	0.0010	0.0576
	Feed pumps (n_{fp})	0.7717	0.7571	0.7856	$<10^{-308}$
	Feed pump replacement period (m_{fp})	0.0273	−0.0037	0.0583	0.08375
Safe shutdown probability	Maintenance shutdown period (m_t)	−0.7238	−0.7399	−0.7068	$<10^{-308}$
	S.S.C. maintenance threshold (c_{ss})	−0.0844	−0.1151	−0.0535	9.042×10^{-8}
	S.S.C. maintenance delay (m_{ss})	−0.1138	−0.1444	−0.0830	5.245×10^{-13}
	Isolation valves (n_{iv})	0.1897	0.1594	0.2197	1.014×10^{-33}
	Isolation valve maintenance period (m_{iv})	−0.0014	−0.0324	0.0296	0.9295
	Feed pumps (n_{fp})	−0.2685	−0.2975	−0.2390	5.201×10^{-67}
	Feed pump replacement period (m_{fp})	0.0305	−0.0005	0.0614	0.05387
Safe shutdown probability per unit monetary expenditure	Maintenance shutdown period (m_t)	−0.5286	−0.5521	−0.5044	6.9×10^{-287}
	S.S. maintenance threshold (c_{ss})	−0.0573	−0.0882	−0.0264	0.0002857
	S.S. maintenance delay (m_{ss})	−0.0659	−0.0968	−0.0350	3.015×10^{-5}
	Isolation valves (n_{iv})	−0.0730	−0.1038	−0.0420	3.834×10^{-6}
	Isolation valve maintenance period (m_{iv})	0.0297	−0.0013	0.0606	0.06039
	Feed pumps (n_{fp})	−0.7638	−0.7781	−0.7487	$<10^{-308}$
	Feed pump replacement period (m_{fp})	−0.0262	−0.0572	0.0048	0.09728
Monetary expenditure per unit time	Maintenance shutdown period (m_t)	−0.6910	−0.7086	−0.6726	$<10^{-308}$
	S.S. maintenance threshold (c_{ss})	−0.0398	−0.0707	−0.0088	0.01189
	S.S. maintenance delay (m_{ss})	−0.0491	−0.0800	−0.0181	0.001913
	Isolation valves (n_{iv})	0.1112	0.0804	0.1418	1.76×10^{-12}
	Isolation valve maintenance period (m_{iv})	−0.0364	−0.0673	−0.0054	0.02138
	Feed pumps (n_{fp})	0.2253	0.1952	0.2548	3.444×10^{-47}
	Feed pump replacement period (m_{fp})	0.0698	0.0389	0.1006	9.898×10^{-6}

Note: 4000 parameter sets were tested, each sampled 10^6 times.

Abbreviation: S.S.C, steam separator circuit.

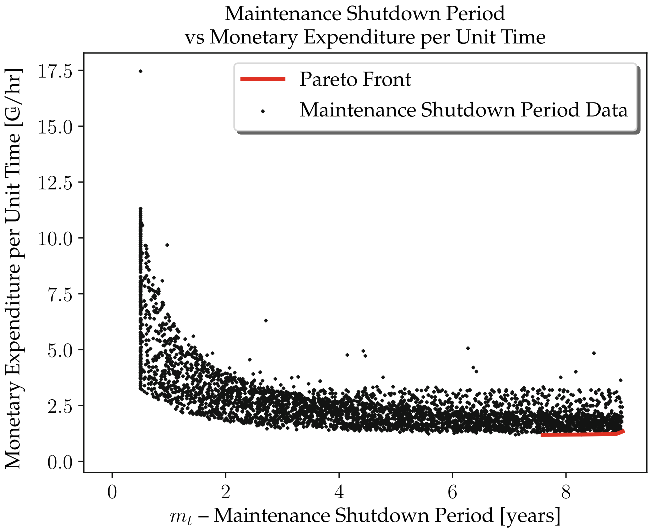

FIGURE 12 Expenditure per unit time with respect to maintenance shutdown period, with the Pareto front drawn.

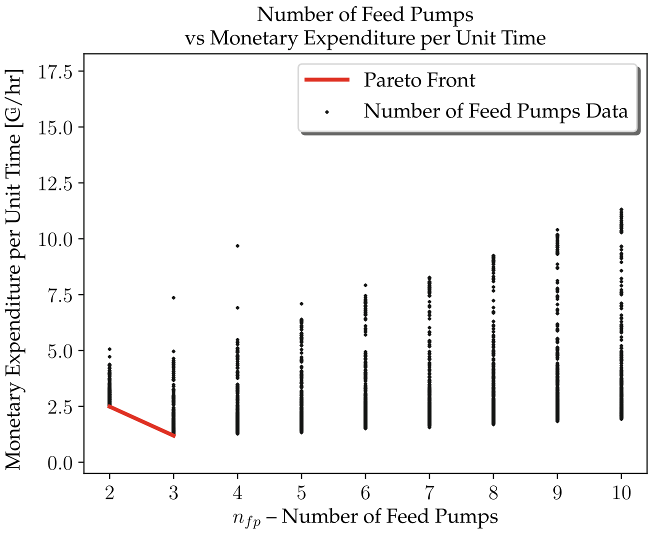

FIGURE 13 Expenditure per unit time with respect to the number of feed pumps with the Pareto front drawn.

To identify the parameters to which the performance metrics are most sensitive, the Spearman's rank correlation coefficients were calculated and are found in Table 4. While not all parameters are highly correlated with the performance outcomes, most Spearman's ρ values have sufficiently small p-values for statistical significance. The Spearman's ρ values of greatest magnitude are found for the maintenance shutdown period, the number of feed pumps, and the number of isolation valves, and in Figures 12–17 these design optimisation parameters are plotted against monetary expenditure per unit time and against safe shutdown probability, with relevant Pareto fronts included where applicable.

It is seen in Figure 12 that monetary expenditure per unit time falls with increasing maintenance shutdown period, making longer scheduled run times more cost effective. However, the curve has become quite flat once the upper end of its permitted range is reached. By contrast, in Figure 15 the deleterious effect on safety from increasing the maintenance shutdown period is seen, although it is most strongly expressed in the visible bands of concentrated non-optimal data points.

The number of isolation valves has little impact on the total expenditure per unit time, as seen in Figure 14, but the benefits to safety from having some measure of redundancy are obvious in Figure 17, although there appears to be no further gains from increasing the number beyond three.

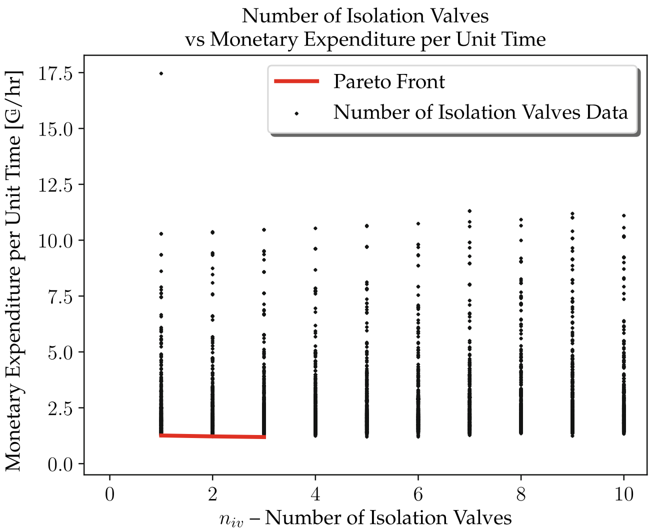

FIGURE 14 Expenditure per unit time with respect to the number of isolation valves, with the Pareto front drawn.

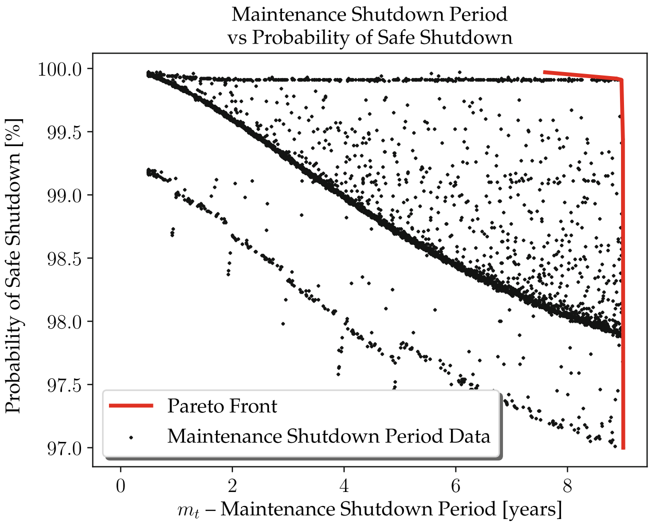

FIGURE 15 Probability of safe reactor shutdown with respect to maintenance shutdown period, with the Pareto front drawn.

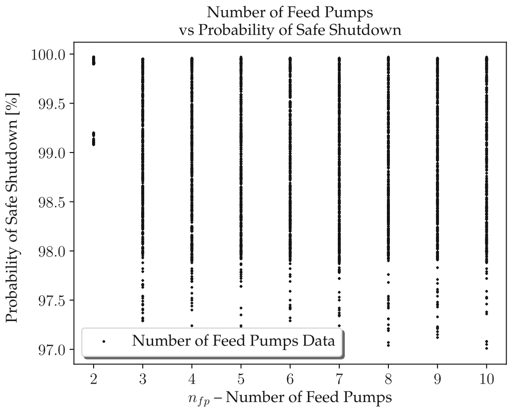

FIGURE 16 Probability of safe reactor shutdown with respect to the number of feed pumps.

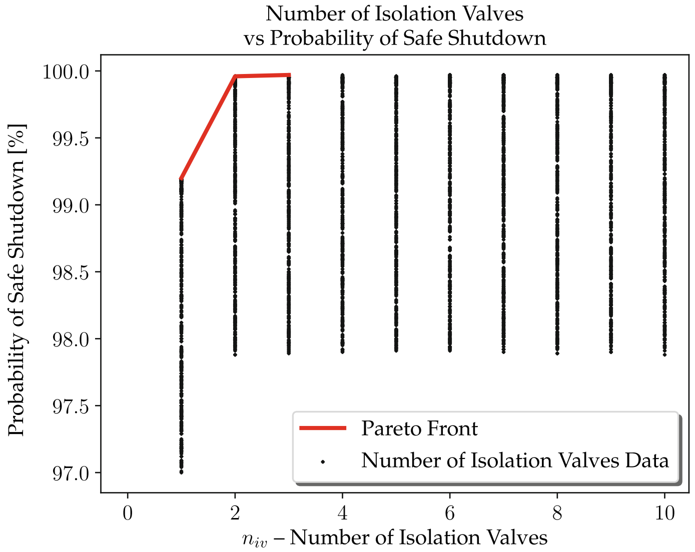

FIGURE 17 Probability of safe reactor shutdown with respect to the number of isolation valves, with the Pareto front drawn.

Some of the correlations with the number of feed pumps initially seem strange. In particular, the probability of safe shutdown appears negatively correlated with the level of redundancy, and it is also found that the average duration of operation is positively correlated with the parameter. Recalling that when examining the Pareto optimal points it was seen that the results where only two feed pumps were used had very short run times and high likelihoods of a safe outcome, it can be concluded that the same phenomenon is the cause of these results. Indeed, if Spearman's ρ is recalculated for only the 3504 parameter sets in which $n_{fp} > 2$, a value of −0.0106 is found, with a 95% confidence interval of −0.0437 to 0.0225 and $p = 0.5309$, that is, no correlation between the number of feed pumps and safe shutdown probability can be established with those results excluded. While they are required for the reactor to safely remain online, given that the feed pumps lie beyond the turbine isolation valve, their failure should not impede the success of the shutdown condensation process, and this is reflected in Figure 16; as such, the revised Spearman's ρ is not surprising. It is also seen that greater numbers of feed pumps increase monetary expenditure per unit time, with a clear anomaly for two feed pumps seen in Figure 13, due to the resulting short run times previously discussed.

7 | CONCLUSIONS

In this work, design and maintenance parameters for the primary coolant system of a nuclear reactor were sampled by a Monte Carlo method and used to generate Petri net models corresponding to those values. Simulation of the Petri nets generated data relating to the performance of the sampled configurations in terms of monetary expenditure, probability of safe shutdown, and duration of operation without interruption from a critical fault. The Pareto optimal parameter sets were then identified and discussed. The most balanced solution is arguably the Pareto optimal parameter set with safest performance of those which had three feed pumps. However, ultimately the preferred configuration is dependent on the priority of those responsible for the system (operators, regulators, etc.) as well as their degree of willingness to rely on emergency shutdown mechanisms. Regardless of what those priories might be, the methodology discussed in this work should serve as a useful means to extract the performance data critical to inform such decision making.

Future work should consider a greater array of cost factors. In particular, simulations should be conducted, in which an individual system is returned to use following maintenance shutdown and examined over a fixed long term period, with inclusion of the equivalent cost resulting from the loss of the supply of electricity from the power plant to the grid during these periods. Furthermore, there exist more complex methodologies for the iteration of system parameters, such as simulated annealing,[42] genetic algorithms,[43] or metaheuristics,[44] and these could be used to more efficiently converge upon the optimal configuration.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council as part of the NuRes project (Grant number EP/R021988/1) within the UK-India Civil Nuclear Collaboration program.

DATA AVAILABILITY STATEMENT

Research data are not shared.

REFERENCES

1. Watson HA. *Launch control safety study*. Technical report Bell Labs; 1961.
2. Vesely WE. A time-dependent methodology for fault tree evaluation. *Nucl Des Eng*. 1970;13:337-360.
3. Ruijters E, Stoelinga M. Fault tree analysis: a survey of the state-of-the-art in modeling, analysis and tools. *Comput Sci Rev*. 2015;15-16:29-62.
4. Rauzy A. New algorithms for fault trees analysis. *Reliab Eng Syst Saf*. 1993;40(3):203-211.
5. Reay KA, Andrews JD. A fault tree analysis strategy using binary decision diagrams. *Reliab Eng Syst Saf*. 2002;78(1):45-56.
6. Rasmussen NC. Reactor safety study: an assessment of accident risks in US commercial nuclear power plants (WASH-1400). Technical report. US Nuclear Regulatory Commission; 1975.
7. Xu H, Dugan JB. Combining dynamic fault trees and event trees for probabilistic risk assessment. Proceedings of the Annual Symposium Reliability and Maintainability, 2004 - RAMS; 2004; IEEE.
8. Andrews JD, Dunnett SJ. Event-tree analysis using binary decision diagrams. *IEEE Trans Reliab*. 2000;49(2):230-238.
9. Petri CA. *Kommunikation Mit Automaten (In German)*. PhD thesis. Technical University Darmstadt; 1962.
10. Schneeweiss WG. Petri net picture book (an elementary introduction to the best pictorial description of temporal changes). LiLoLe – Verlag GmbH (Publ. Co. Ltd.); 2004.
11. Lee SJ, Seong PH. Development of automated operating procedure system using fuzzy colored petri nets for nuclear power plants. *Ann Nucl Energy*. 2004;31(8):849-869.
12. Németh E, Bartha T, Fazekas C, Hangos KM. Verification of a primary-to-secondary leaking safety procedure in a nuclear power plant using coloured Petri nets. *Reliab Eng Syst Saf*. 2009;94(5):942-953.
13. Aldemir T. A survey of dynamic methodologies for probabilistic safety assessment of nuclear power plants. *Ann Nucl Energy*. 2013;52:113-124.
14. Kachur SA, Shakhova NV. Turbine generator status diagnostic system based on petri nets. *Nucl Energy Technol*. 2016;2(2):81-84.
15. Ponciroli R, Cammi A, Lorenzi S, Luzzi L. Petri-net based modelling approach for ALFRED reactor operation and control system design. *Prog Nucl Energy*. 2016;87:54-66.
16. Singh L, Rajput H. Safety analysis of life critical software systems: a case study of nuclear power plant. *IETE Tech Rev*. 2016;34(3): 333-339.
17. Singh LK, Vinod G, Tripathi AK. Early prediction of software reliability: a case study with a nuclear power plant system. *Computer*. 2017;49(1):52-58.
18. Kumar P, Singh LK, Kumar C. Performance evaluation of safety-critical systems of nuclear power plant systems; 2019; Nuclear Engineering and Technology.
19. Shukla DK, Arul AJ, Wootton MJ, Andrews J. Reliability analysis of a safety system using petri net and comparison with smart component methodology. Proceedings of PSA 2019 - International Topical Meeting on Probabilistic Safety Assessment and Analysis (PSA 2019); April 28 - May 3, 2019; Charleston, SC.
20. Wootton MJ, Andrews J, Lloyd AL, Smith R, Arul AJ, Vinod G, Prasad SH, Garg V. Petri nets and pseudo-bond graphs for a nuclear reactor primary coolant system. Proceedings of the 29thEuropean Safety and Reliability Conference; 2019:3831-3839.
21. Shukla DK, Arul AJ. An efficient Dijkstra's semaphore based Kout-of-N Petri net voting model: An application to shutdown actuation system reliability. *Ann Nucl Energy*. 2021;157:108218.
22. Wootton MJ, Andrews JD, Lloyd AL, et al. Risk modelling of ageing nuclear reactor systems. *Ann Nucl Energy*. 2022;166:108701.
23. Zimmermann A, Rodriguez D, Silva M. A two phase optimization method for Petri net models of manufacturing systems. *J Intell Manuf*. 2001;12(5/6):409-420.
24. Raghavan NRS, Roy D. A stochastic Petri net approach for inventory rationing in multi-echelon supply chains. *J Heuristics*. 2005;11:421-446.
25. Gam M, Lefebvre D, Nabli L, Telmoudi AJ. Optimization of maintenance patrols planning. Proceedings of the 29thMediterranean Conference on Control and Automation (MED); June 2021:1299-1304; IEEE.
26. Metropolis N, Ulam S. The Monte Carlo method. *J Am Stat Assoc*. 1949;44(247):335-341.
27. Hammersley JM, Handscomb DC. In: Cox DR, Hinkley DV, Rubin D, Silverman BW, eds. *Monographs on Statistics and Applied Probability 3, Monte Carlo Methods*. Chapman & Hall; 1964.
28. Gandibleux X, Ehrgott M. 1984- 2004 – 20 years of multiobjective meta-heuristics. But what about the solution of combinatorial problems with multiple objectives? In: Coello CAC, Aguirre AH, Zitzler E, eds. *Evolutionary Multi-Criterion Optimization, 3rd International Conference, EMO 2005 Guanajuato Mexico, March 2005, Proceedings (LNCS 3410)*. Springer; 2005:33-46.
29. Segovia-Hernández JG, Gómez-Castro FI. *Stochastic Process Optimization Using Aspen Plus*. CRC Press; 2017.

30. Balbo G. Introduction to generalized stochastic petri nets. In: Bernardo M, Hillston J, eds. *Formal Methodsfor Performance Evalutation – 7th International School on Formal Methods for the Design of Computer, Communication and Software Systems, SFM 2007, Bertinoro, Italy, May/June 2007, Advanced Lectures, Lecture Notes in Computer Science*. Vol 4486. Springer; 2007:83-131.
31. Papoulis A, Pillai SU. In: Shultz CF, Flomenhoft ML, eds. *Probability, Random Variables and Stochastic Processes*. 4th ed. McGraw Hill; 2002.
32. Jiang R, Murthy DNP. A study of Weibull shape parameter: properties and significance. *Reliab Eng Syst Saf*. 2011;96:1619-1626.
33. Barringer & Associates, Inc. Weibull reliability database for failure data for various components. www barringer1 com / wdbase htm. Accessed November 2018.
34. IAEA. IAEA-TECDOC-478 Component Reliability Data for Use in Probabilistic Safety Assessment. Technical report International Atomic Energy Agency; 1988.
35. Reliability Eta Beta database. www reliabilityetabeta com Accessed January 2020.
36. Morris S. Failure rate estimates for mechanical components. www reliabilityanalyticstoolkit appspot com / mechanical _ reliability_data. Accessed January 2020.
37. Smith DJ. Practical methods for engineers including reliability centred maintenance and safety-related systems. *Reliability, Maintainability and Risk*. 7th ed. Elsevier Butterworth-Heinemann; 1981.
38. Spearman C. The proof and measurement of association between two things. *Am J Psychol*. 1904;15(1):72-101.
39. Puth M-T, Neuhäuser M, Ruxton GD. Effective use of Pearson's product-moment correlation coefficient. *Anim Behav*. 2014;93:183-189.
40. Ruscio J. Constructing confidence intervals for Spearman's rank correlation with ordinal data: a simulation study comparing analytic and bootstrap methods. *J Modern Appl Stat Methods*. 2008;7(2):416-434.
41. Andrade C. The P value and statistical significance: misunderstandings, explanations, challenges, and alternatives. *Ind J Psychol Med*. 2019;41(3):210-215.
42. van Laarhoven JMP, Aarts EH. Simulated annealing. *Simulated Annealing: Theory and Applications*. Springer; 1987:7-15.
43. Pattison RL, Andrews JD. Genetic algorithms in optimal safety system design. *Proc Inst Mech Eng E J Process Mech Eng*. 1999;213(3):187-197.
44. Talbi E-G. *Metaheuristics: From Design to Implementation*. John Wiley & Sons; 2009.

How to cite this article: Wootton MJ, Zhou Y, Andrews JD, et al. Optimisation of nuclear reactor primary coolant design and maintenance parameters. *Appl Stochastic Models Bus Ind*. 2022;1-21. doi: 10.1002/asmb.2722