Risk modelling of ageing nuclear reactor systems

Wootton, Mark James; Andrews, John D.; LLoyd, Adam; Smith, Roger; Arul, A. John; Vinod, Gopika; Prasad, M. Hari; Garg, Vipul

DOI:
10.1016/j.anucene.2021.108701

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (Harvard):
Wootton, MJ, Andrews, JD, LLoyd, A, Smith, R, Arul, AJ, Vinod, G, Prasad, MH & Garg, V 2022, 'Risk modelling of ageing nuclear reactor systems', *Annals of Nuclear Energy*, vol. 166, 108701. https://doi.org/10.1016/j.anucene.2021.108701

Annals of Nuclear Energy 166 (2022) 108701

Contents lists available at ScienceDirect

Annals of Nuclear Energy

journal homepage: www.elsevier.com/locate/anucene

Risk modelling of ageing nuclear reactor systems

Mark James Wootton [a,*], John D. Andrews [a], Adam L. Lloyd [b], Roger Smith [b], A. John Arul [c], Gopika Vinod [d], M. Hari Prasad [d], Vipul Garg [d]

[a] Faculty of Engineering, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom
[b] Department of Mathematical Sciences, Loughborough University, Loughborough, Leicestershire LE11 3TU, United Kingdom
[c] Reactor Engineering Group, Indira Gandhi Centre for Atomic Research, Kalpakkam 603 102, India
[d] Reactor Safety Division, Bhabha Atomic Research Centre, Trombay, Mumbai 400 085, India

ARTICLE INFO

Article history:
Received 16 March 2021
Received in revised form 2 August 2021
Accepted 6 September 2021
Available online 21 September 2021

Keywords:
Petri nets
Nuclear reactors
Reliability engineering
Failure modelling
Ageing components

ABSTRACT

A nuclear reactor is expected to function for extensive periods, during which, coolant circulation and core reactivity must always be maintained safely. Understanding the risks associated with the operation of such systems requires proper consideration of ageing components and the effects of preventative maintenance. The traditional methodologies, such as Fault Trees and Event Trees, have limitations in their abilities to model ageing processes and complex maintenance strategies. Petri Nets have been used in this research as a more suitable alternative. A case study reactor is presented to demonstrate this capability. Petri Nets were developed for five key subsystems: primary coolant circulation, shutdown condensation, emergency core coolant injection, emergency shutdown, and control and monitoring, building a representation which considers their failure modes, reaction of the system to faults, and ongoing component maintenance actions. These models reveal statistics for the timing of failure of these subsystems and relative frequencies of outcome categories.

1. Introduction

The pressing climatological crisis demands the reorientation of electricity generation away from fossil fuels towards alternatives, including nuclear energy. The low rate of construction of new nuclear power installations from the mid-1980s onwards (IAEA, 2019) has resulted in a high average operating lifespan for existing reactors, often beyond their original design life. Rather than take on the political and economic costs associated with new build nuclear projects, the decision has often been made to extend the mission of existing power stations instead. Therefore, when considering the risks associated with a design, the assessment methodology must be able to capture the progressive increase in component failure rates as the reactor ages and also the complex asset management strategies used to control this process.

Risk assessment of modern industrial plant are commonly performed using a combination of fault tree and event tree methods (Rasmussen, 1975). H.A. Watson conceived the Fault Tree methodology at Bell laboratories in the 1960s (Watson et al., 1961). As seen in Fig. 1, a Fault Tree model is a graphical representation of how component failures, represented by basic events, can combine to cause the occurrence of a single specific undesired outcome, known the top event. Logic gates such as AND and OR are used to express how the fault propagation to system level occurs. Qualitative analysis of the fault tree structure yields the minimum cut sets, these being the smallest combinations of basic events required to cause the top event.

By the end of the decade, W.E. Vesely (Vesely, 1970) had progressed the quantification methodology with the development of a time-dependant analysis known as a Kinetic Fault Tree. This enables the computation of the top event probability or frequency from the probability and frequency of the component failures. Furthermore, importance measures (Ruijters and Stoelinga, 2015) reflecting the relative significance of the contributions from each component or cut sets may be found, allowing attention in the design process to directed to the occurrence of safety critical failures.

In 1993, A. Rauzy (Rauzy, 1993) provided an alternative means to quantify the fault tree using Binary Decision Diagrams (BDDs). By comparison to the Kinetic Fault Tree methodology, this method provides both an accurate and efficient methodology, eliminating the need for approximations in the traditional analysis method (Reay and Andrews, 2002).

Event Trees originate with the WASH-1400 report led by N.C. Rasmussen at the US Nuclear Regulatory Commission in 1975

* Corresponding author.
E-mail address: mark.wootton@nottingham.ac.uk (M.J. Wootton).

https://doi.org/10.1016/j.anucene.2021.108701
0306-4549/

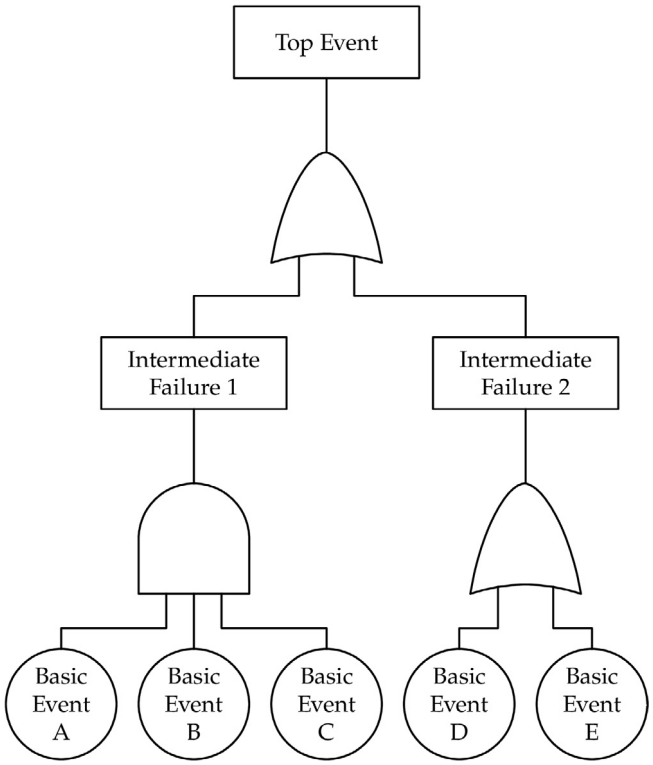

Fig. 1. A simple Fault Tree comprised of a system reliant on two components. Component 1 has two redundant subcomponents, both of which must fail to disable it. Component 2 can be disabled by either one of two basic events. The minimum cuts sets are {A,B,C}, {D}, and {E}.

(Rasmussen, 1975). Event Trees explore the consequences which can occur following the occurrence of an initial event, as illustrated in Fig. 2. They consider all possible responses of safety systems following the initiating event. In 2000, a methodology was developed by J.D. Andrews and S.J. Dunnett (Andrews and Dunnett, 2000) for Event Trees making use of Binary Decision Diagrams.

In the integrated Fault Tree and Event Tree analysis there is an assumption that the basic events occur independently. Also, in the majority of commercial codes implementing the analysis, there is an assumption that components fail with a constant failure rate. Hence their wear-out cannot be represented. Component repair times are also commonly limited to unrealistic exponential distributions and complex repair processes cannot be represented, meaning that the traditional assessment processes are limited in their ability to model ageing reactors. At present, there is no implementation of Fault Trees or Event trees that allows the representation of both non-constant failure rates and complex systems of repair and maintenance. Critically, neglecting to consider increasing failure rates, such as is found in an ageing component, can result in an underestimation of time-to-failure and thus also the overall system risk.

An alternative system failure modelling approach, which has the ability to overcome these deficiencies is the Petri Net method (Carl, 1962). Applications of this methodology have already occurred in the civil nuclear energy industry (Lee and Seong, 2004; Németh et al., 2009; Aldemir, 2013; Kachur and Shakhova, 2016; Ponciroli et al., 2016; Singh et al., 2016; Singh and Rajput, 2016; Kumar et al., 2019; Wootton et al., 2019). It has the potential to overcome the deficiencies in the Fault Tree and Event Tree methodologies when representing ageing systems and capturing restorative processes and asset management strategies. The trade-off for the advantages that Petri Nets offer is the requirement to perform simulations to evaluate the model where many iterations may be required to reach convergence, increasing computational expenditure. However, the extensive capabilities of Petri Nets justify this choice when a high fidelity representation of the dynamics of a system is desired. With a Petri Net, is it possible to model changes in system operating mode, e.g. from general operation to shutdown or emergency coolant injection. Likewise, a component in the model may have arbitrarily many states in which it can exist beyond simply functional or failed; for instance a stand-by state, any number of degrees of degradation, or several mutually exclusive or concurrent failure modes. Multiple components working together can be modelled, with the ability to capture their functional dependency, rather than being required to assume independence. Any number of different probability distributions can be used within a single model for the timing of its events, thereby facilitating the inclusion of diverse failure mechanisms and maintenance regimes. Furthermore, aside from allowing the representation of complex stochastic processes, Petri Net structures can also include features highly useful to reliability modelling, such as the development of concurrent and synchronous processes (Aubry et al., 2016).

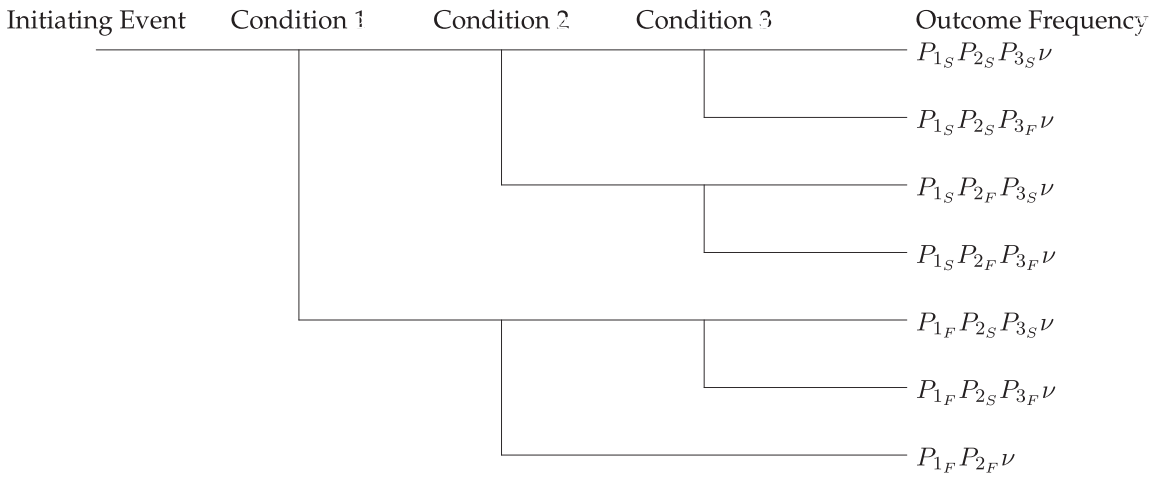

Fig. 2. A Event Tree with three possible branching points following the initiating event, in which the failure of both the first and second condition disables the opportunity to attempt the third. Each condition can succeed or fail with probabilities P_{1_S}, P_{2_S}, and P_{3_S}, and P_{1_F}, P_{2_F}, and P_{3_F}, where $P_{n_S} = 1 - P_{n_F}$, with the initiating event occurring at a frequency of ν.

The development of a methodology utilising such capabilities is the objective of the research presented in this paper. Its demonstration is achieved through the examination of a case study system described in Section 2. Existing Petri Net research regarding civil nuclear energy broadly falls into two categories. Some use Petri Nets as a language for the development and representation of specific operational processes or procedures (Lee and Seong, 2004; Németh et al., 2009; Kachur and Shakhova, 2016; Ponciroli et al., 2016), in other cases to drive a coupled physical simulation (Ponciroli et al., 2016; Wootton et al., 2019). With the exception of Wootton et al. 2019 (Wootton et al., 2019), these works did not perform risk or reliability assessment in their nuclear system. In other works, Petri Nets were used as a description of part of a nuclear system, but were evaluated by other means, with the model converted to another form, such as Markov chains (Singh et al., 2016), reachability graphs (Singh and Rajput, 2016), or both (Kumar et al., 2019), rather than by direct simulation. Consequently, some of the advantages of Petri Nets were eroded, such as the ability to represent time, or the quantification of uncertainty, with the latter demonstrated in this work. In many cases, the Petri Nets presented are relatively small, but greater complexity is achievable. The Petri Nets in this paper are simulated directly and illustrate how the methodology can facilitate probabilistic safety assessment of nuclear plants with estimation of statistical confidence.

2. Reactor case study

A generic nuclear reactor with modern design features, illustrated in Fig. 3, is used to demonstrate the methodology developed. The core itself is cooled by vertical coolant channels surrounding fuel rods, with a liquid heavy water moderator. Light water coolant enters the core via an inlet header and leaves through an outlet header. Circulation through the four steam separators on this loop is driven by thermosiphon action. From the steam separators, steam is extracted to drive turbines. Following re-condensation, the water is returned to the separators as a liquid by a set of feed pumps.

Under normal operating circumstances, reactor shutdown is scheduled after three years of operation in order to perform routine maintenance. In the event that a steam drum or one of the sections of pipe connecting it to the core (collectively referred to as a steam circuit) suffers rupture, the reactor may remain online, but maintenance is scheduled for six months following the break. However, should the functionality of a second steam circuit be lost, immediate shutdown is enacted. The loss of a third steam circuit constitutes a major fault, and to requires the use of the emergency core coolant injection system. Faults in either the inlet header, outlet header or the internal coolant channels of the core also activate the injection of emergency coolant.

Of the three feed pumps, two must be running at any one time, with the third on standby if either of the others fail. Additionally, the pumps in use are cycled to allow maintenance to be performed, running for six months at a time, staggered so as to come off-line separately. If a pump experiences an electrical or mechanical fault and no redundant pump is available, or the redundant pump fails to come online on demand, reactor shutdown becomes necessary to allow repairs to be made. Breakages in the pipes running from the steam separators to the turbine and faults in the turbine or condenser themselves also require the reactor to shutdown to rectify the problem.

Control rods are used to maintain a steady critical state in the reactor, and must be positioned appropriately to this end. Failure of the electric drive or the electric signal to the control rods results in inappropriate rod positioning. Hardware or software faults with the controlling automation system can also bring about this situation. If, because of such a misplacement, or for some other reason, a transient reactivity event occurs, a set of sensors is present to detect it. This consists of three neutron detectors (NDs) and three ion chambers (ICs). Both function with a two-of-three voting system, requiring that at least two sensors of the same set make a simultaneous reading to deliver a measurement from their block. This acts as a counter measure against the possibility that sensors can provide both false positive (detection of a non-existent increase or decrease in reactivity) and false negative (failure to detect a real increase or decrease in reactivity) readings. Failure to detect reactivity events will result in problems stemming from

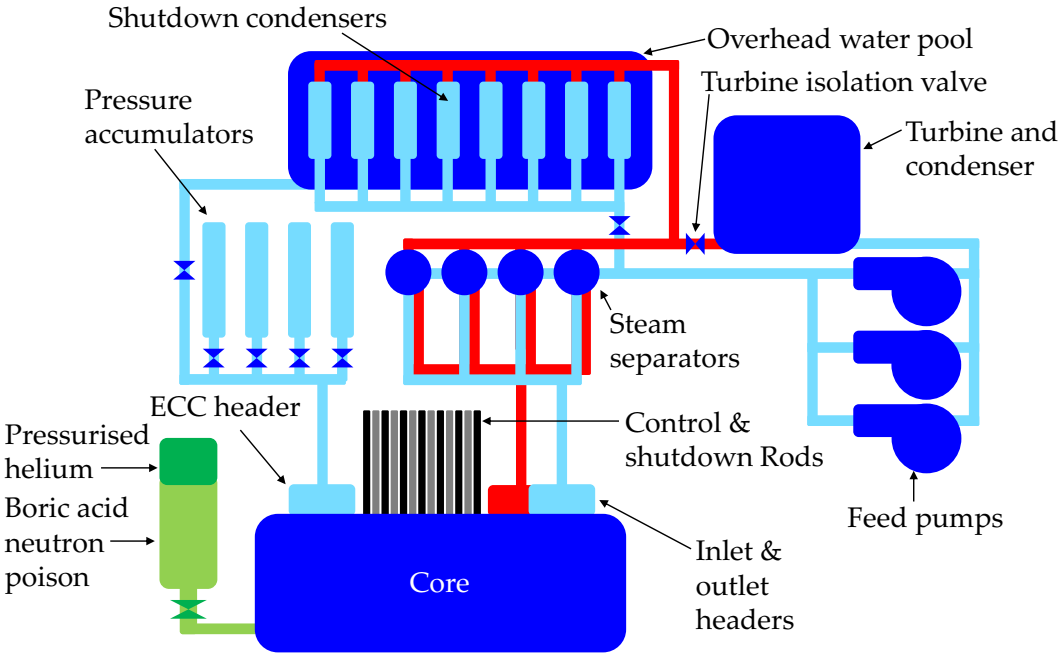

Fig. 3. Schematic overview of the case study reactor used in this work.

a supercritical or subcritical state, such as core overheating or loss of power output respectively.

To shutdown the reactor, the control rods are positioned to create a subcritical state and the turbine isolation valve closes to disconnect the turbine, with coolant instead running from the steam separators to eight shutdown condensers providing decay heat extraction using a large water tank as a heat sink. If the valve fails to shut, or subsequently opens once closed, a redundant valve is available. The event of either closing prematurely, forces the reactor to enter the shutdown process. Leakage from the tank requires that the reactor go directly to emergency shutdown. The isolation condensers themselves are arranged in pairs, with two parallel air operated valves (AOVs) separating each pair from the coolant circulation system during normal operation. The opening of one AOV is enough to engage the pair, the second valve providing redundancy. When the isolation valve closes to initiate shutdown, pressure builds in the isolation condensers until it is sufficient to open the AOVs. To complete the shutdown process, at least half of the isolation condensers must come into use and remain operational for its forty day duration. As well the potential to lose two isolation condensers if both valves of an AOV pair fail to open, individual condensers can be disabled by rupture or calcification once in active use.

In the event of a loss of core pressure, four pressure accumulators stand ready to perform high pressure coolant inject through an emergency inlet header for a period of half hour. This is immediately followed by low pressure injection lasting three days and at the end of this process the core is completely submerged in water. Low pressure injection is driven by gravity and sources its water from the same overhead tank as that used as a heat sink by the shutdown condensers. There are four pressure accumulators in total, and the equivalent of half their collective capacity is required for the high pressure injection period. When demand is placed on a pressure accumulator, an AOV must open, followed by a rupture disk. If either of these fail, the coolant will not be released. In the event that a rupture disk experiences a partial failure and does not open fully, it is assumed that half of its pressure accumulator contribution is available. Once the pressure accumulators are online, the connection between them and the core must remain intact without rupture until the beginning of low pressure injection. Additionally, the rupture disks can break prematurely, and in doing so, render their pressure accumulator unavailable on demand. As with the steam separators, if one such fault occurs, shutdown for maintenance is scheduled for six months after its emergence. In the event that two failures occur, the reactor shutdown is immediate.

The system has two independent processes for rapid emergency reactor shutdown, respectively Shutdown System One and Shutdown System Two (SDS-1 and SDS-2). In SDS-1, forty shutdown rods, suspended above the core, de-latch automatically and fall into the core. This is activated on a trip signal, generated in the event that either the coolant pressure drops too low, or the temperature reaches the Currie point of a magnet in the mechanism. Once released, each rod falls into its slot under gravity. If thirty-eight of the rods correctly insert, SDS-1 is considered successful. Otherwise, SDS-2 activates and uses pressurised helium to rapidly inject the neutron poison, boric acid, into the moderator. Providing the helium release mechanism operates correctly, the boric acid must pass through one of three valves to reach the core – a pressure activated valve, a reactor trip signal operated valve, and a manually operated valve. As these are in parallel, the successful opening of any one valve is sufficient. If any of these valves opens prematurely, the unintended presence of boric acid in the moderator will result in an unplanned shutdown. Leakage from the supply of either helium or boric acid will require the reactor to immediately go to shutdown to allow for restorative actions to be taken.

3. Timed Petri Net methodology

3.1. Overview

Generalised Stochastic Petri Nets (Carl, 1962; Balbo, 2007) with atomic firing are used to model the system dynamics in this research. They consist of four basic objects; the *place*, the *transition*, the *arc*, and the *token*, as illustrated in Fig. 4. Places, drawn as circles, represent information about the current state of the system, such as the condition or availability of components and resources. The status of the system at any point in the lifetime simulation is represented by locating tokens in the places. The number of tokens, drawn as black dots, in a place is referred to as its *marking*. A transition, drawn as a square, updates the markings of places connected to it by arcs in order to represent events which change the state of the system, such as failures and repairs, as well as performing simple logical operations. In this way the dynamics of the system operation are propagated through the Petri Net. The arcs are either *input arcs* or *output arcs*, which determines whether tokens are taken or given to a place when the transition *fires*. The orientation of the arrows on the arcs denotes the direction of token motion. The number of tokens travelling along each arc is indicated by an associated *weight* (known as its multiplicity) for the arc. Before a transition fires, it first needs to be enabled. For a transition to be enabled each input place must contain at least its arc multiplicity of tokens. The transition then fires after a delay time associated with the transition. This delay time can be specified or sampled from a probability distribution. On firing, the multiplicity of tokens is extracted from each of the input places for the transition and the multiplicity of tokens is added to each of the output places. Should an arc not have an associated multiplicity then the default multiplicity of 1 is assumed. The requisite token placement required to enable a transition must persist uninterrupted until it fires, otherwise it will return to an unenabled state and its previous time to fire discarded. If multiple transitions are scheduled to fire at the same moment, one is chosen at random with uniform probability. The scheduled firing time of all other transitions is preserved when one transition fires if the requisite places markings continue to be satisfied.

Inhibit arcs can also be featured on a Petri net. These are represented by a red arc with a dotted line from a place to a transition. The arrowhead on the arc is replaced by a circle. In the event that the place contains at least its arc multiplicity of tokens then the transition it is connected to is prevented from firing. An example of the transition firing process is seen in Fig. 5.

As depicted in Fig. 6, other features used include the place condition arc and the voting transition. The place conditional arc is drawn with a blue dashed line with a circular end. A place conditional arc applies a modifier to the delay between a transition becoming enabled and its firing based on its weight and the number of tokens held at the connected place. The voting transition has a threshold associated with it, which specifies the number of

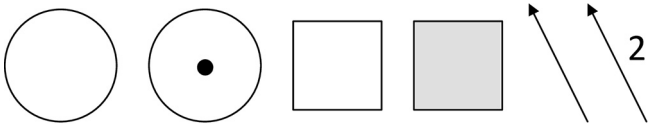

Fig. 4. Key to standard Petri Net objects. From left to right: empty place, place with one token, timed transition, instant transition, standard arc, standard arc with weight of two.

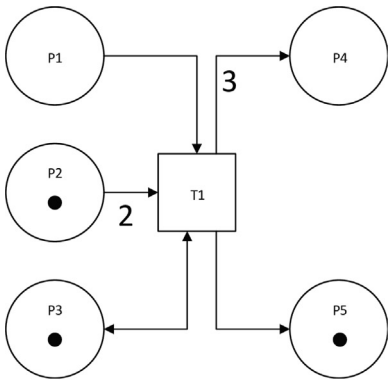

(a) Net before T1's token requirements are met.

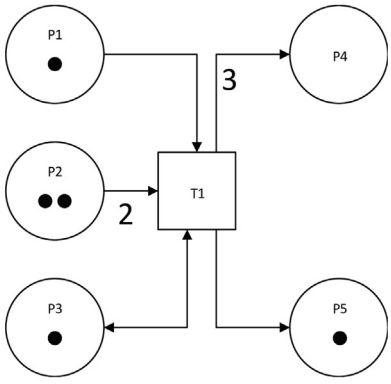

(b) T1 enabled.

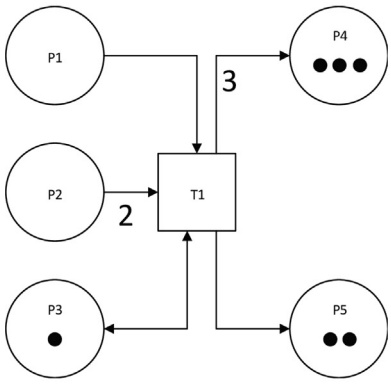

(c) After some duration, T1 fires.

Fig. 5. Illustration of Petri Net transition firing.

incoming arcs it requires to become enabled, in contrast to the standard transition which requires all to be satisfied. When the voting transition fires, all places of incoming arcs whose weight is met lose tokens, with the rest left unaltered. A voting transition is drawn as a three-dimensional box, with its incoming arcs drawn with alternating length dashed line and its threshold written below its identifying label. It interacts with place conditional and inhibit arcs in identical fashion as normal transitions.

In this work, the software used for Petri Net modelling is called *Macchiato* and was developed in-house at the University of Nottingham.

3.2. Analysis of the Petri Net

The analysis of the Petri net is carried out using Monte Carlo simulation (Metropolis and Ulam, 1949). This method performs

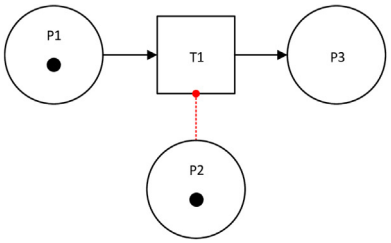

(a) The inhibit arc from P2 prevents T1 from firing, despite the requirements of the arc from P1 being met.

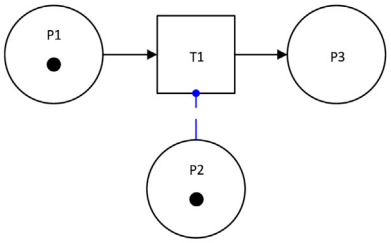

(b) This timed transition is enabled and the wait calculated is altered by the token applied to the place conditional arc.

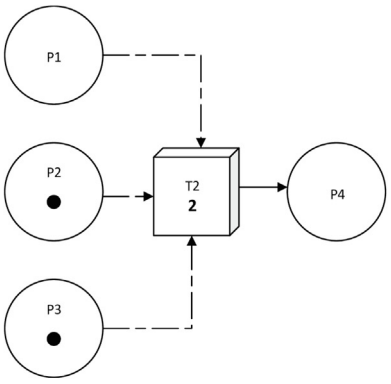

(c) A voting transition with threshold equal to two. This transition is enabled.

Fig. 6. Examples illustrating the extended Petri Net objects used in this work.

experiments on the computer which replicates potential life histories of the plant taking random samples from the probability distributions representing transition firing delay times. When enough simulations have been performed to give convergence, a statistical analysis can be carried out to provide performance metrics for the time durations and number of occurrences of any places on the network. The network places are selected to give the performance variables required by the study.

3.3. Firing delay calculation

A fixed delay transition fires after a set duration, a, and a transition with a uniform distribution gives even weighting to values within a range, parametrised by u, giving the probability density function, $f(t;u)$, for time to fire, t, such that:

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

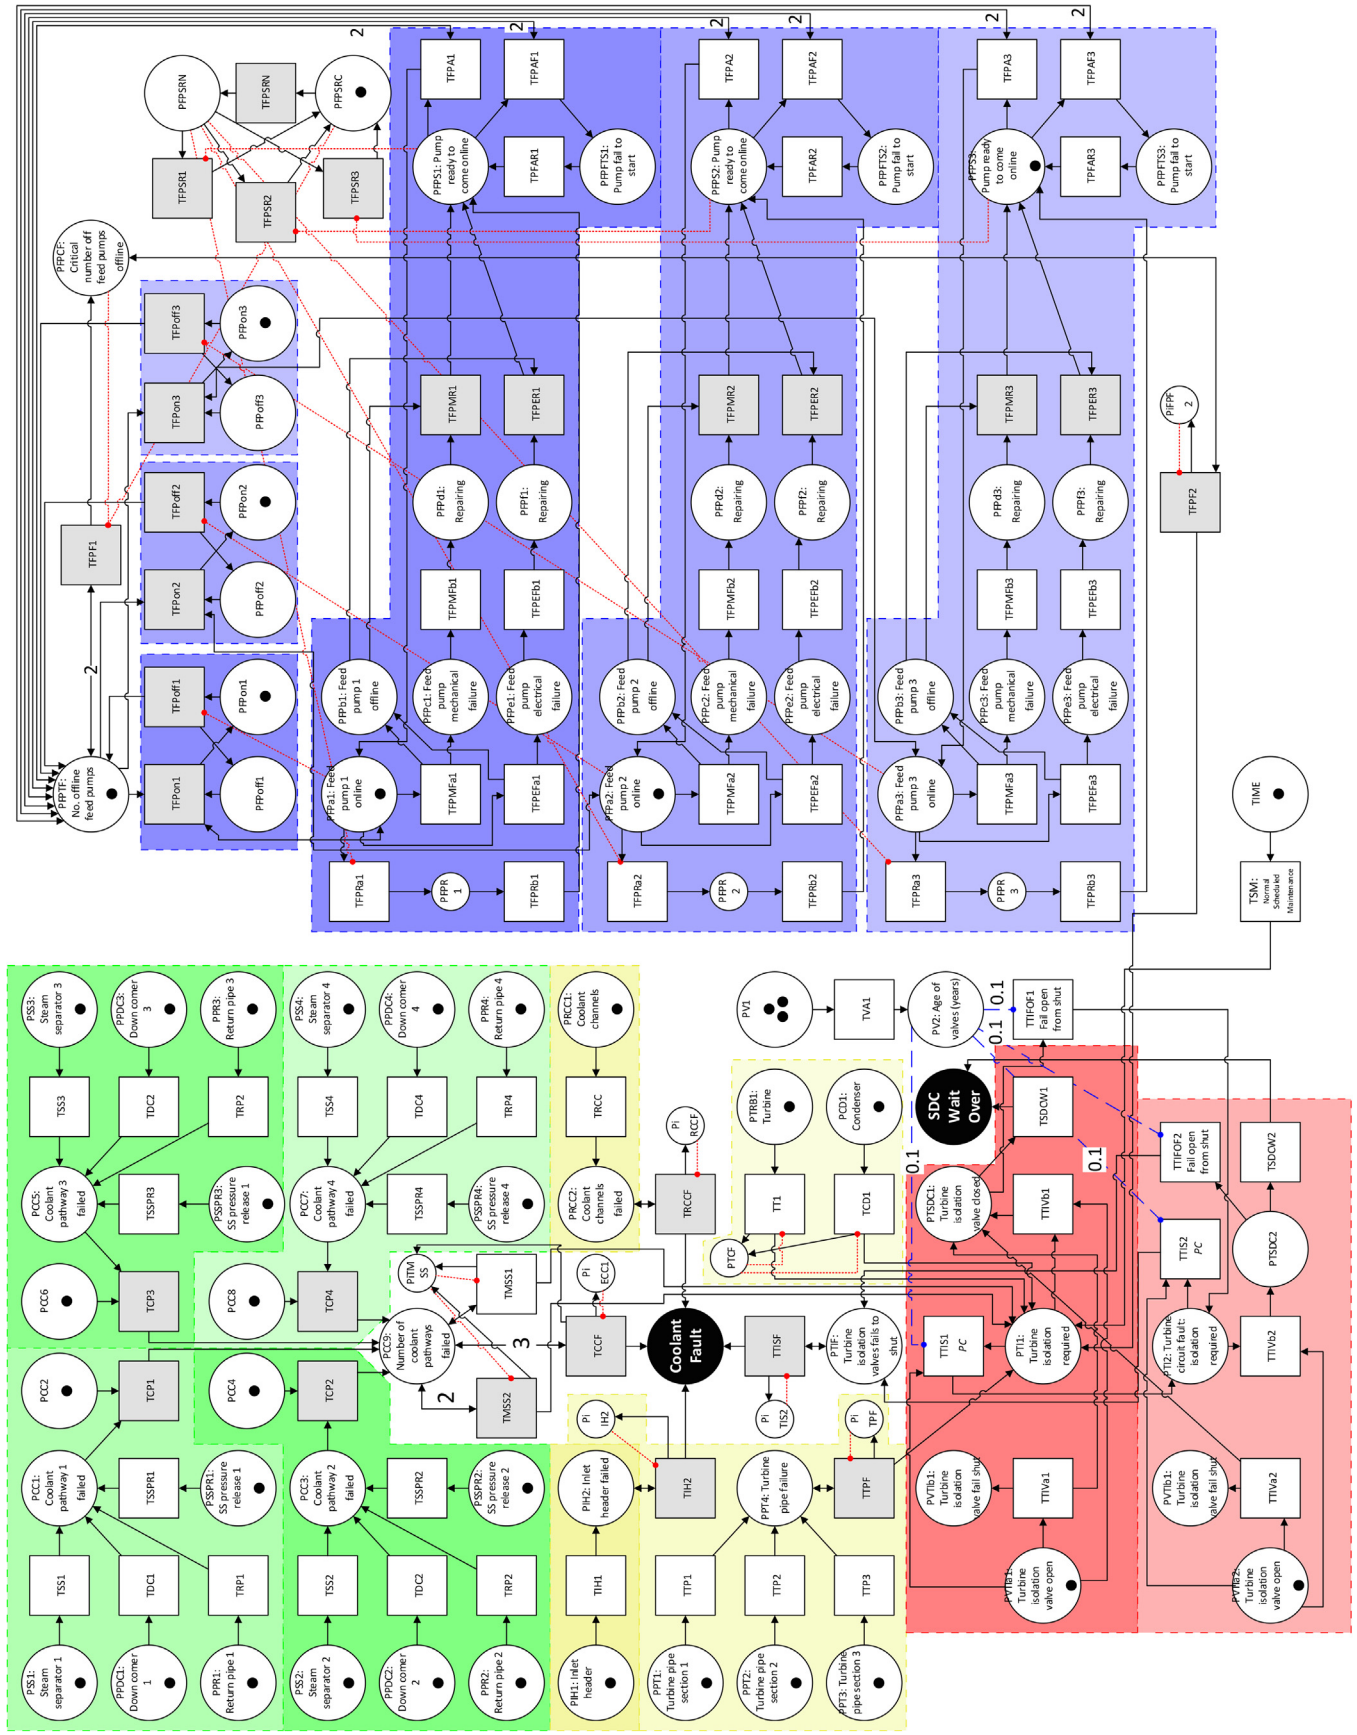

Fig. 7. Petri Net model depicting the primary coolant system. Steam separators are highlighted in green, feed pumps in blue, isolation valves in read, and miscellaneous pipes in yellow.

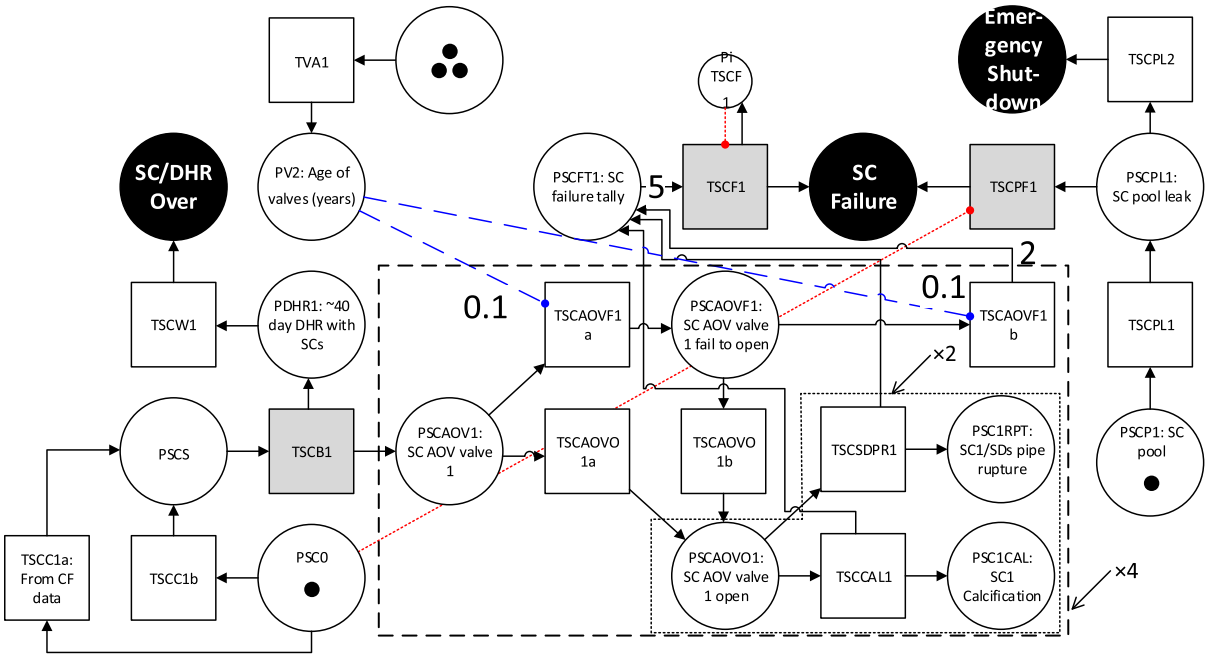

Fig. 8. Petri Net model for the shutdown condenser subsystem used for decay heat removal. Note the repeated sections indicated within the dashed and dotted boxes.

A cyclic distribution fires when the simulation clock reaches the next whole multiple of its parameter, c. A second parameter, ω, allows one to offset these firings. For example, for two transitions required to fire once an hour, thirty minutes apart, these would be parameterised with $c = 1$, $\omega = 0$, and $c = 1$, $\omega = 0.5$.

The Weibull Distribution (Papoulis and Unnikrishna Pillai, 2002; Jiang and Murthy, 2011) is used widely in reliability engineering for modelling component failures. Its two-parameter form is characterised by the scale parameter, η, which indicates the time at which approximately two-thirds of the population of components will fail, and a shape parameter, β, which indicates if infant mortality, random failures or wear-out is taking place. For $\beta < 1$, failure rate decreases with time. With $\beta > 1$, the failure rate grows with time, representative of an ageing process. If $\beta = 1$, the failure rate is constant. Eq. (2) gives the probability density function, $f(t;\eta,\beta)$, of the Weibull Distribution:

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

The Log-Normal Distribution, widely used to model repair times, is given by the result of the exponential function applied to the result of a normal distribution (Weisstein, 2019), such that the resulting variable is normally distributed in magnitude (Dennis and Patil, 1987). As seen in Eq. (3), for a normal distribution with mean and standard deviation given by μ and σ, the probability density function, $f(t;\mu,\sigma)$, is:

$$f(t;\mu,\sigma) = \frac{\exp\left(-\frac{1}{2}\left[\frac{\ln(t)-\mu}{\sigma}\right]^2\right)}{t\sigma\sqrt{2\pi}}. \tag{3}$$

When a transition is connected to one or more place conditional arcs, a modifying factor, P, is calculated,

$$P = 1 + \sum_i W_i N_i, \tag{4}$$

where W_i is the weight of the ith place conditional arc and N_i is the number of tokens on the corresponding place. A parameter (a, u, c, η, or μ, as relevant) from the transition's probability distribution is then modified by dividing it by P, with these parameters being chosen to scale the time to fire of a transition in inverse proportion with P. This system was chosen to allow the description of arbitrary mechanisms of time to fire alteration with respect to one or many place markings. Uniquely among arc types, the place conditional arc is not limited to an integer weight.

4. Subsystem models

The five subsystems isolated for study are the circulation of primary coolant, the shutdown condensers, the emergency injection of core coolant, the emergency shutdown mechanisms, and the reactor control and reactivity monitoring equipment. For each, a Petri Net model has been created and these are presented in section 4.1 to 4.5, with illustrations in Figs. 7–11 and corresponding parameters in Tables 1–5. Every object is assigned a unique code to identify it and are referred to by such in the text, with places having a code commencing with "P" and those of transitions with "T", with the exception of terminal places, which are labelled with the outcome that they represent. Throughout this work, wherever relevant, parameters are expressed in terms of hours.

The structure of these Petri Nets follows from the description of the system outlined in Section 2, being developed in consultation with the knowledge and experience of reactor system engineers at IGCAR and BARC. An individual Petri Net structure does not require a long time to develop, but the process to compile and combine the necessary data (e.g. failure rates) can be time consuming, and the model may go through multiple iterations as part of the discussion with experts. As such, any one Petri Net and its parameters might typically be the product of several days worth of work.

Fig. 9. Petri Net model for emergency core coolant injection process. The high pressure injection section is highlighted yellow and the low pressure section in green. Note the repeated sections indicated within the black dashed box.

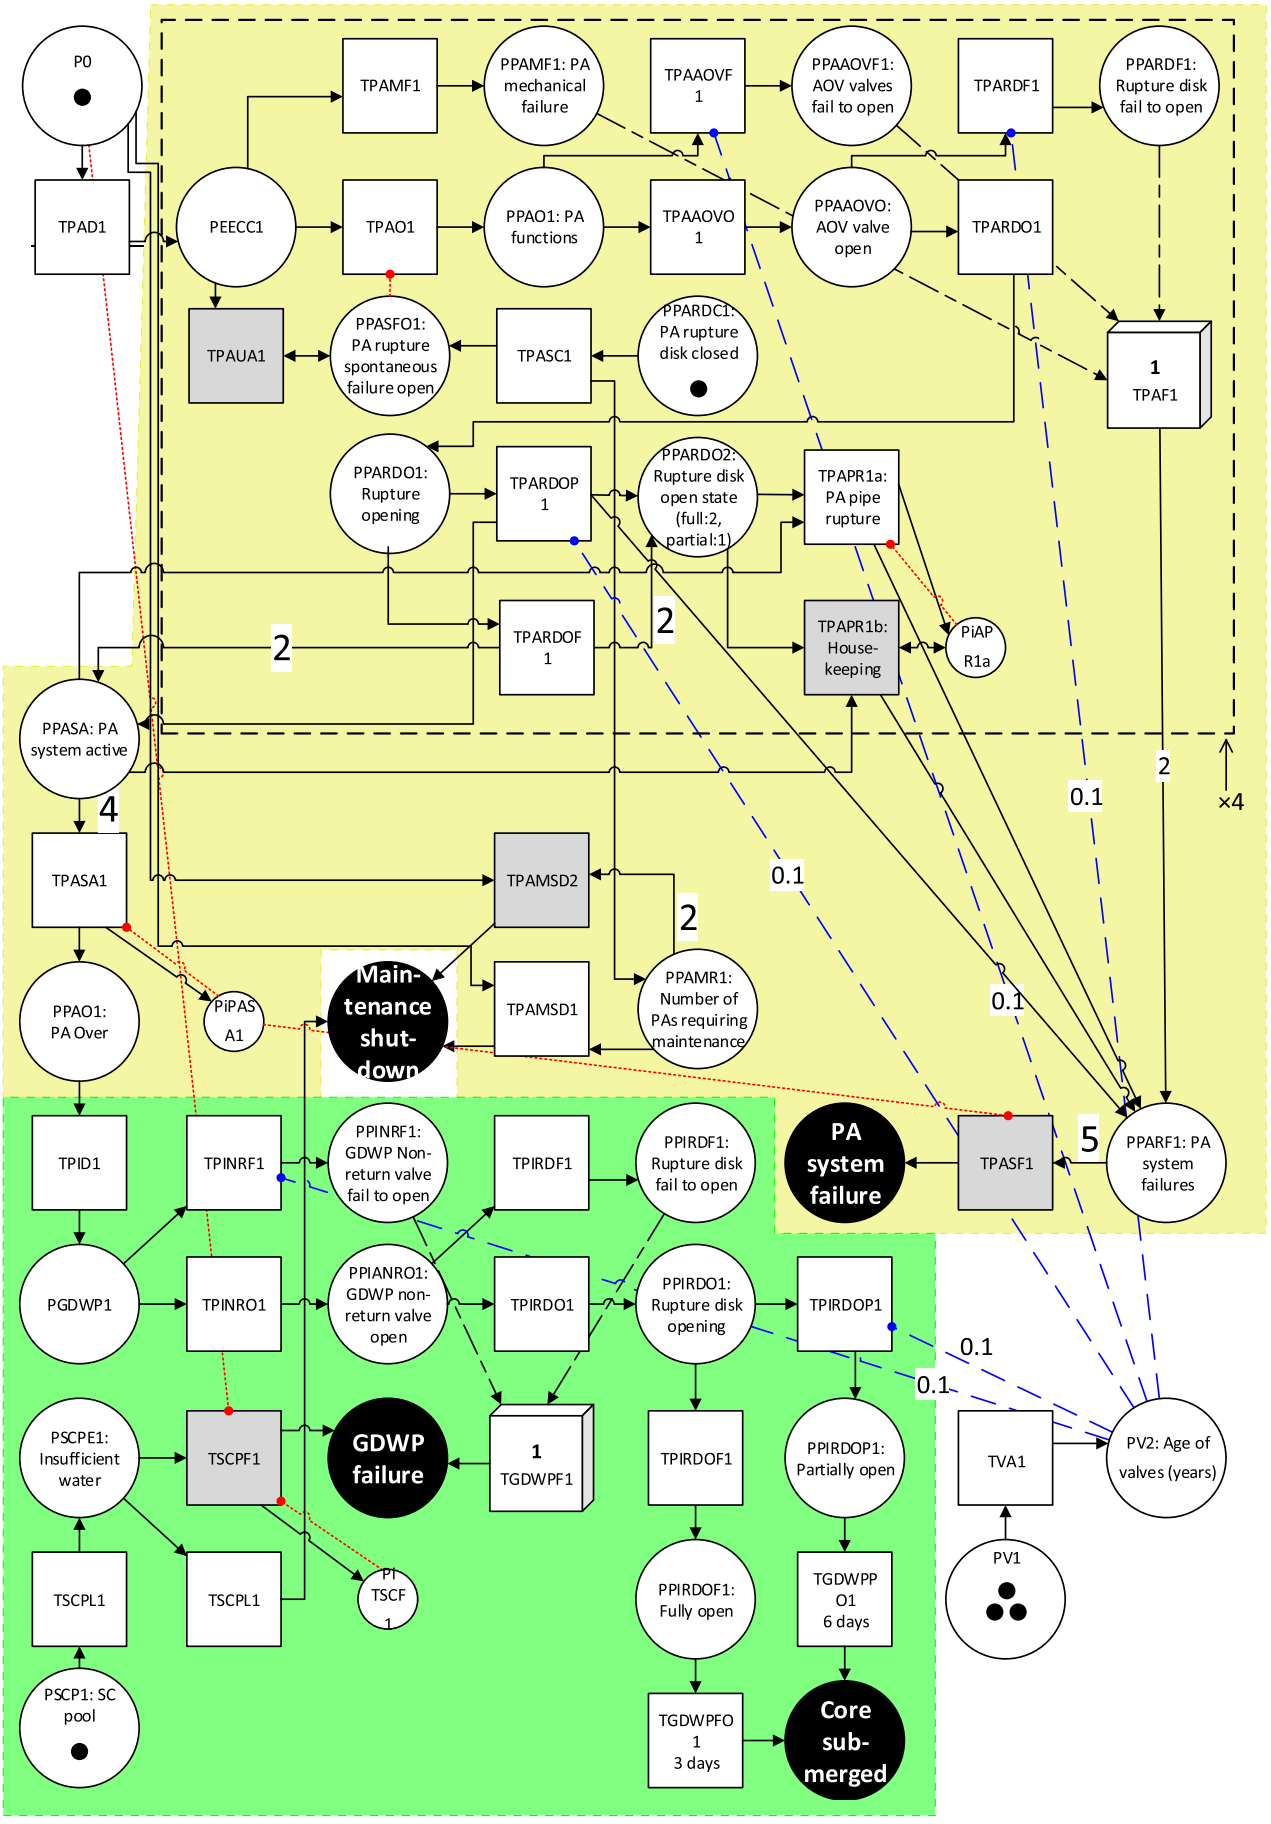
P0
TPAD1
PEECC1
TPAMF1
PPAMF1: PA mechanical failure
TPAAOVF 1
PPAAOVF1: AOV valves fail to open
TPARDF1
PPARDF1: Rupture disk fail to open
TPAO1
PPAO1: PA functions
TPAAOVO 1
PPAAOVO: AOV valve open
TPARDO1
TPAUA1
PPASFO1: PA rupture spontaneous failure open
TPASC1
PPARDC1: PA rupture disk closed
1 TPAF1
PPARDO1: Rupture opening
TPARDOP 1
PPARDO2: Rupture disk open state (full:2, partial:1)
TPAPR1a: PA pipe rupture
TPARDOF 1
TPAPR1b: House-keeping
PiAP R1a
PPASA: PA system active
TPASA1
TPAMSD2
TPAMSD1
PPAMR1: Number of PAs requiring maintenance
Main-tenance shut-down
PPAO1: PA Over
PiPAS A1
TPID1
TPINRF1
PPINRF1: GDWP Non-return valve fail to open
TPIRDF1
PPIRDF1: Rupture disk fail to open
PA system failure
TPASF1
PPARF1: PA system failures
PGDWP1
TPINRO1
PPIANRO1: GDWP non-return valve open
TPIRDO1
PPIRDO1: Rupture disk opening
TPIRDOP1
PSCPE1: Insufficient water
TSCPF1
GDWP failure
1 TGDWPF1
TPIRDOF1
PPIRDOP1: Partially open
TVA1
PV2: Age of valves (years)
TSCPL1
TSCPL1
PI TSCF 1
PV1
PPIRDOF1: Fully open
TGDWPP O1 6 days
PSCP1: SC pool
TGDWPFO 1 3 days
Core sub-merged
2
4
2
0.1
0.1
0.1
2
×4
5
0.1
0.1

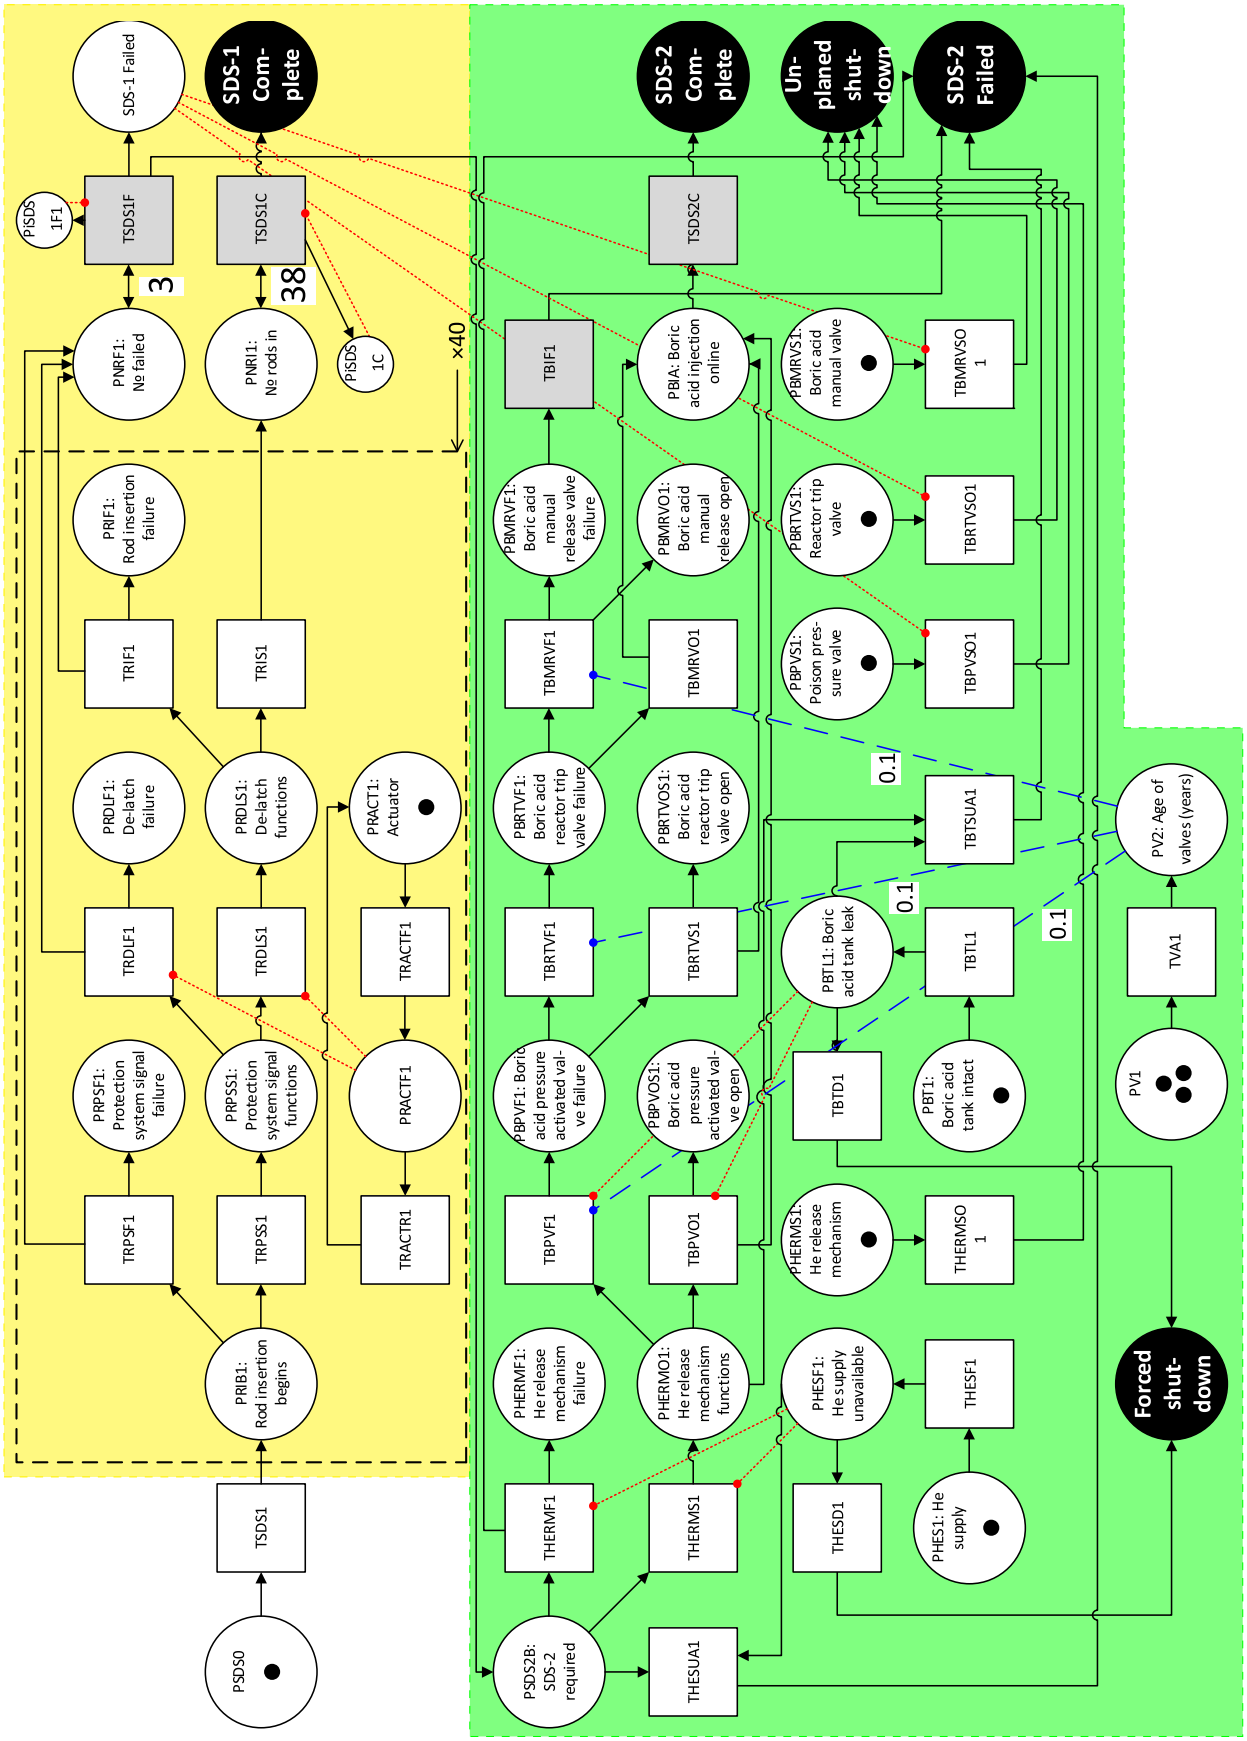

Fig. 10. Petri Net model for the emergency reactor shutdown processes. The shutdown rods are highlighted in yellow and the neutron poison injection in green. Note the repeated sections indicated within the dashed box.

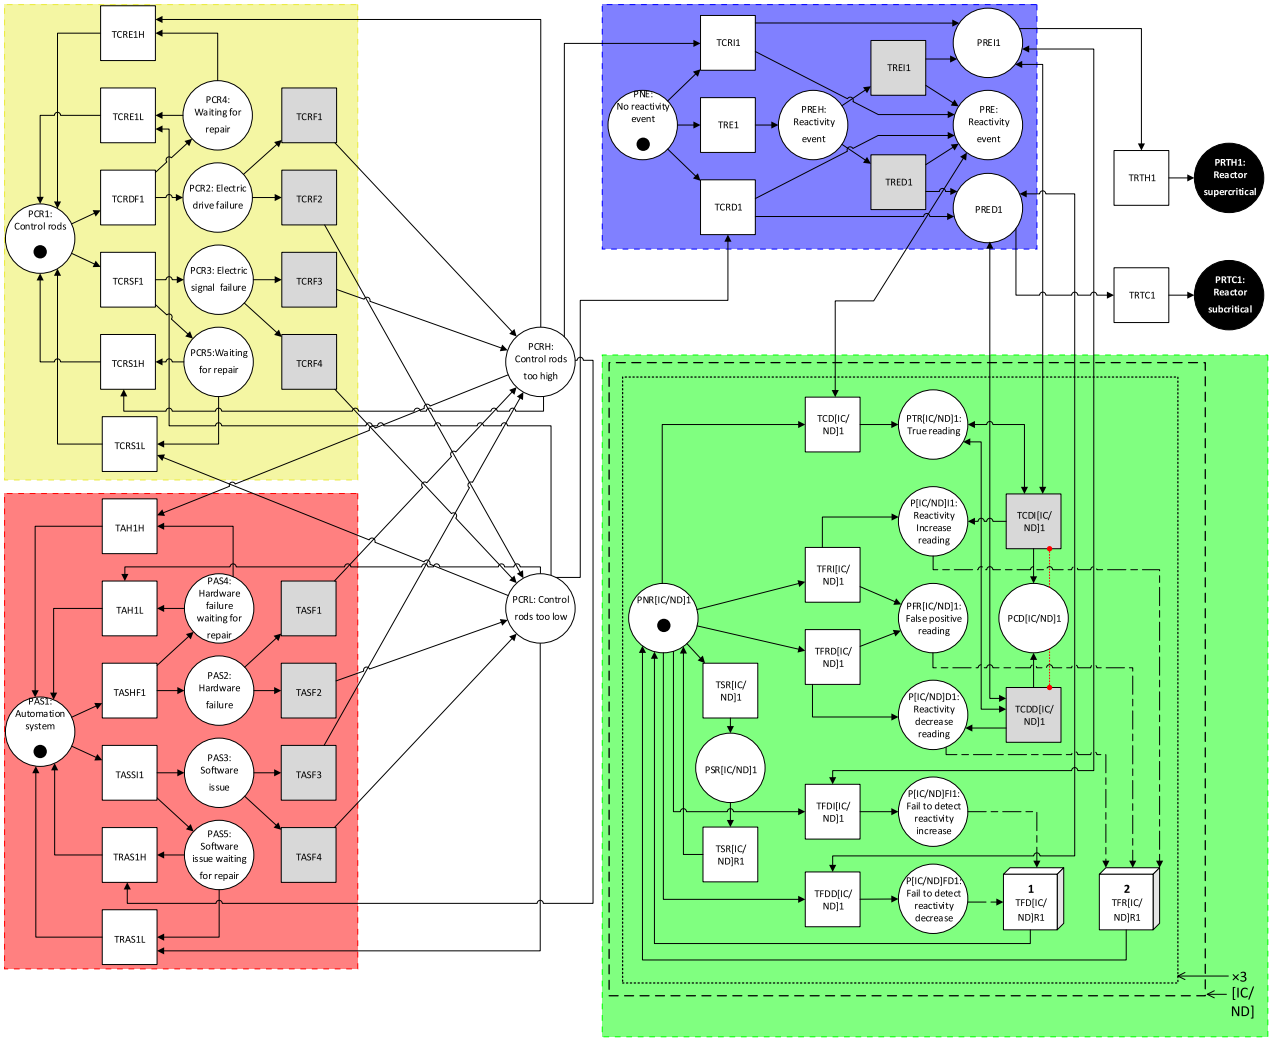

Fig. 11. (a) Control rod hardware (yellow), automation system (red), and ion chambers (IC) and neutron detectors (ND) in the reactor controls system Petri Net. Note the repeated sections indicated within the dashed and dotted boxes. The green section describes the operation of the ion chambers and neutron detectors, and the blue section processes the resulting reactivity events. This figure continues overleaf. Sensor outcomes and signal processing for the reading from the ion chambers (IC) and neutron detectors (ND) in the Petri Net model for the reactor controls systems, continued from previous page. Note the repeated sections indicated within the dashed box.

4.1. Primary coolant circulation

The Petri Net model concerning initiating events in the primary coolant circulation system is found Fig. 7 and the parameters associated with its transitions are found in Table 1. There are two possible outcomes of each simulated iteration – the emergence of a failure event that requires urgent intervention from the emergency coolant injection system (hereafter referred to as an initiating event), or the safe shutdown of the reactor through the shutdown condensers, such that repairs may be performed, respectively labelled "Coolant Fault" and "SDC Wait Over". If no faults have occurred within a three year period, the reactor is shutdown for routine maintenance, which is represented by the transition, TSM.

The green section of the Petri Net covers the four circuits between the reactor headers and the steam separators, including four failure modes, rupture of the down comer pipe to the core, the returning pipe, and the steam separator itself, as well as its pressure release valve, respectively TDC(1–4), TRP(1–4), TSS(1–4), and TSSPR(1–4). Any of these failures disables the circuit, putting a token on PCC(1/3/5/7), which causes TCP(1–4) to record the failure as an additional token at PCC9. The initial token at PCC(2/4/6/8) ensures that the loss of a circuit can only be recorded once. TSMSS1 fires after a delay of six months to initiate shutdown for repairs, requiring only one token at PCC9. Two tokens are needed to fire the instant transition, TMSS2, which calls for immediate shutdown in the event of two circuits losses. The number of tokens on PCC9 is conserved when either of these transitions fires. Therefore, the place continues to track the number of failed circuits during shutdown, such that TCCF fires if a third fails in this period to indicate a major coolant fault.

The feed pumps are modelled in the blue section. The simulation begins with pumps 1 and 2 online and 3 on standby, as indicated by the initial token placements for PFPa(1–3). For a running pump, transitions TFPMF(1/3/5) and TFPEF(1/3/5) respectively represent mechanical and electrical failures, with TFPMF (2/4/6) and TFPEF(2/4/6) likewise representing the repair processes. When repairs are complete, TFPMR(1–3) and TFPER(1–3) mark PFPS(1–3) to indicate that the pump is ready for use again. Transitions TFPRa(1–3) are the repair processes, taking a pump offline to services it before TFPRb(1–3) returns it to availability. The transitions TFPSR(1–3) and places PFPSRC and PFPSN track whether a pump is available to be brought into active use, and as two pumps are always required to be online, the inhibit arcs from PFPSRN to TFPSR(1–3) prevent a pump from being serviced when

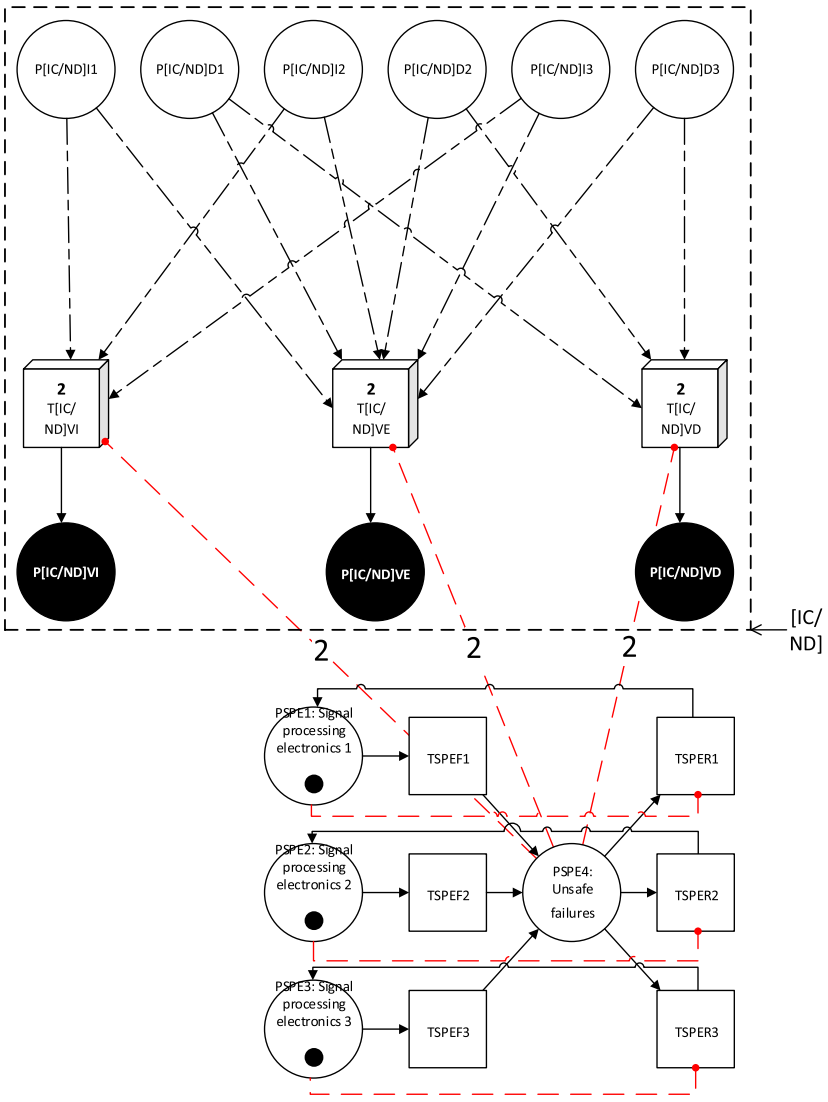

Fig. 11 (*continued*)

there is no available replacement in working order. There is a chance that when being brought into online, a pump will fail to start, and one cannot attempt to use it again until repaired, with this possibility described by TFPAF(1–3), PFPFTS(1–3), and TPFAR (1–3), where TFPA(1–3) is the normal behaviour of activation of a pump. TFPon(1–3), TFPoff(1–3), PFPoff(1–3), PFPoff(1–3), record which pumps are online, recording the total number of offline pumps on PFPTF. If this reaches two with no replacement pumps available (note the inhibit arc from PFPSRC), TFPF1 fires, setting motion a demand for reactor isolation and shutdown.

The red section pertains to the isolation valves. The arrival of a token at PTI1 indicates demand for reactor shutdown. If PVTIa1 still has its initial token, an attempt to close the first isolation valve is made, with TTIS1 representing the possibility that it fails to do so. This becomes progressively more likely with each passing year of operation, which is represented by the place conditional arc from PV2, with TVA1 incrementing its tokens once a year. A successful closure of the first isolation valve, TTIVb1, puts a token on PTSDC1, beginning the shutdown condenser process. If the valve fails to shut or opens due to failure once closed, see TTIFOT1, the second valve must be used, and a token is placed on PTI2 to indicate this demand. The Petri Net structure for the second valve is identical to the first, except that the outcome of failed closure of this valve is an emergency coolant fault. There also exists a possibility that either of the valves will close prematurely, instigating an unplanned shutdown, with TTIVa(1–2) representing this failure mode.

4.2. Shutdown condensers

The Petri Net modelling the shutdown condensers is shown in Fig. 8, with its parameterisation located in Table 2. Its possible outcomes are the safe completion of the shutdown condensation decay heat removal process, its failure, and the requirement for emergency shutdown if a leak develops from the water pooled used as a heat sink.

Demand on the shutdown condensers is created by the closure of the isolation valve, and this Petri Net uses this event as its main starting point, with TSCC1a and TSCC1b respectively representing demand placed due the need for early shutdown (the paramterisation of which is derived from the results of the primary coolant Petri Net described in Section 4.1) and demand placed due the end of the three year operation period. TSCB1 puts a token on PDHR1 to set up a forty day timer *i.e.* TSCW1 and on each of

Table 1
Timed transition parameters for the primary coolant circulation Petri Net seen in Fig. 7.

Transition	Type	Parameter(s)	Ref.	Transition	Type	Parameter(s)	Ref.
TCD1	Weibull	$\eta = 1.07 \times 10^6$, $\beta = 1.2$	(Expert Opinion; Barringer & Associates, Inc., 2010)	TDC(1–4)	Weibull	$\eta = 4.38 \times 10^8$, $\beta = 1.0$	(IAEA, 1988; Reliability Eta Beta database, 2020)
TFPA(1–3)	delay	$a = 0.0003$	(IAEA, 1988)	TRP(1–4)	Weibull	$\eta = 1.75 \times 10^6$, $\beta = 1.0$	(IAEA, 1988; Reliability Eta Beta database, 2020)
TFPAF(1–3)	uniform	$u = 0.63$	(IAEA, 1988)	TRCC	Weibull	$\eta = 1.62 \times 10^6$, $\beta = 1.5$	(IAEA, 1988; Reliability Eta Beta database, 2020)
TFPAR(1–3)	delay	$a = 20.9$	(IAEA, 1988)	TSM	delay	$a = 26298.0$	$^N/_A$
TFPEFa(1–3)	Weibull	$\eta = 10600$, $\beta = 1.2$	(Smith, 1981; Barringer & Associates, Inc., 2010)	TSS(1–4)	Weibull	$\eta = 9.32 \times 10^7$, $\beta = 1.2$	(Expert Opinion; Barringer & Associates, Inc., 2010)
TFPEFb(1–3)	delay	$a = 24.0$	(IAEA, 1988)	TSSPR(1–4)	Weibull	$\eta = 100000$, $\beta = 1.0$	(IAEA, 1988; Barringer & Associates, Inc., 2010)
TFPF1	delay	$a = 0.0006$	$^N/_A$	TT1	Weibull	$\eta = 1.13 \times 10^6$, $\beta = 1.7$	(Expert Opinion; Barringer & Associates, Inc., 2010)
TFPMFa(1–3)	Weibull	$\eta = 14000$, $\beta = 1.2$	(IAEA, 1988; Barringer & Associates, Inc., 2010)	TTIFOF(1–2)	Weibull	$\eta = 1.75 \times 10^5$, $\beta = 1.0$	(Expert Opinion; Morris, 2019)
TFPMFb(1–3)	delay	$a = 24.0$	(IAEA, 1988)	TTIS(1–2)	uniform	$u = 0.13$	(IAEA, 1988)
TFPRa1	cyclic	$c = 4383.0$, $\omega = 1461.0$	$^N/_A$	TTIVa(1–2)	Weibull	$\eta = 1.51 \times 10^9$, $\beta = 1.1$	(Expert Opinion; Barringer & Associates, Inc., 2010)
TFPRa2	cyclic	$c = 4383.0$, $\omega = 2922.0$	$^N/_A$	TTIVb(1–2)	delay	$a = 0.0003$	(IAEA, 1988)
TFPRa3	cyclic	$c = 4383.0$, $\omega = 0.0$	$^N/_A$	TTP1	Weibull	$\eta = 1.75 \times 10^8$, $\beta = 1.0$	(IAEA, 1988; Reliability Eta Beta database, 2020)
TFPRb(1–3)	delay	$a = 20.9$	(IAEA, 1988)	TTP2	Weibull	$\eta = 1.66 \times 10^8$, $\beta = 1.0$	(IAEA, 1988; Reliability Eta Beta database, 2020)
TSDCW(1–2)	delay	$a = 960.0$	$^N/_A$	TTP3	Weibull	$\eta = 1.00 \times 10^8$, $\beta = 1.0$	(IAEA, 1988; Reliability Eta Beta database, 2020)
TIH1	Weibull	$\eta = 6.75 \times 10^7$, $\beta = 1.2$	(IAEA, 1988; Barringer & Associates, Inc., 2010)	TVA1	delay	$a = 8766.0$	$^N/_A$
TMSS1	delay	$a = 4383.0$	$^N/_A$	TVTRa(1–2)	delay	$a = 43830.0$	$^N/_A$

Table 2
Timed transition parameters for shutdown condenser subsystem Petri Net seen in Fig. 8.

Transition	Type	Parameter(s)	Ref.	Transition	Type	Parameter(s)	Ref.
TSCAOVF(1–4)(a/b)	uniform	$u = 0.278$	(IAEA, 1988)	TSCPL1	Weibull	$\eta = 1.53 \times 10^6$, $\beta = 3.0$	(Eide et al., 1990; Reliability Eta Beta database, 2020)
TSCAOVO(1–4)(a/b)	delay	$a = 0.000278$	(IAEA, 1988)	TSCPL2	delay	$a = 0.0$	$^N/_A$
TSCC1a	Weibull	$\eta = 19000$, $\beta = 1.0$, $\sigma_{\dot{x}} = 30.0$	Section 4.1	TSCSDPR(1–8)	Weibull	$\eta = 1.77 \times 10^6$, $\beta = 1.0$	(IAEA, 1988; Barringer & Associates, Inc., 2010)
TSCC1b	delay	$a = 26298.0$	$^N/_A$	TSCW1	delay	$a = 960.0$	$^N/_A$
TSCCAL(1–8)	Weibull	$\eta = 1.02 \times 10^7$, $\beta = 1.2$	(IAEA, 1988; Barringer & Associates, Inc., 2010)	TVA1	delay	$a = 8766.0$	$^N/_A$
				–	–	–	–

PDHR(1–4) to begin the process. TSCAOV1a and TSCAOV1b represent the opening of the AOVs when the requisite pressure is reached, while TSCAOVF1a and TSCAOVF1b represent their failure to open. If both fail, two tokens are placed on PSCFT1 to indicate the loss of a pair of shutdown condensers. TSCSDPR1 and TSCCAL1 respectively represent pipe rupture and condenser calcification during operation, each adding a token on PSCFT1. If the number of tokens at PSCFT1 reaches five, more than half of the isolation condensers are non-operational and TSCF1 fires to record the failure of the shutdown condensation process. The shutdown condensers rely on there being water in the pool to dump heat into. If a leak emergences, see TSCPL1, before demand on the condensers is placed, the reactor is at risk of being unable to enter the normal shutdown process, and therefore emergency shutdown is required to allow for the undertaking of repairs, represented by TSCPL2. If shutdown condensation has already begun, the loss of this water results in failure. As such, in the Petri Net, TSCPF1 can only fire if the initial token at PSC0 is removed due to the inhibit arc. Being an instant transition, TSCPF1 is given priority over TSCPL2 when both are available to fire.

4.3. Emergency core coolant injection

Fig. 9 depicts the Petri Net representing the process by which coolant is injected into the core to mitigate a loss of pressure. The four available outcomes are its completion having submerged the core, the failure of the pressure accumulators during high pressure injection, the failure of the low pressure gravity driven injection from the over head pool, and the development of a situation where the reactor must be shutdown to perform maintenance on emergency injection components.

For the purposes of this Petri Net, a demand is generated uniformly between the beginning of the simulation and 3 years 40 days by the transition TPAD1, which puts a token on each of PEECC(1–4). The yellow section represents the high pressure phase and in the green one finds the low pressure coolant injection. The

Table 3
Timed transition parameters for emergency coolant injection Petri Net seen in Fig. 9.

Transition	Type	Parameter(s)	Ref.	Transition	Type	Parameter(s)	Ref.
TGDWPFO1	delay	$a = 72.0$	$^N/_A$	TPARDO(1–4)	delay	$a = 0.000278$	(Expert Opinion; Eide et al., 1990)
TGDWPPO1	delay	$a = 144.0$	$^N/_A$	TPARDOF(1–4)	delay	$a = 0.000278$	(Expert Opinion; Eide et al., 1990)
TSCPL1	Weibull	$\eta = 1.53 \times 10^6$, $\beta = 3.0$	(Eide et al., 1990; Reliability Eta Beta database, 2020)	TPARDOP(1–4)	uniform	$u = 0.235$	(Expert Opinion; Eide et al., 1990)
TSCPL2	delay	$a = 1.0$	$^N/_A$	TPASA1	delay	$a = 0.5$	$^N/_A$
TPAAOVF(1–4)	uniform	$u = 1.54321$	(IAEA, 1988)	TPASC(1–4)	Weibull	$\eta = 87700$, $\beta = 1.0$	(Expert Opinion; Reliability Eta Beta database, 2020)
TPAAOVO(1–4)	delay	$a = 0.000278$	(IAEA, 1988)	TPINRF1	uniform	$u = 1.51$	(IAEA, 1988)
TPAD1	uniform	$u = 27258.0$	$^N/_A$	TPINRO1	delay	$a = 0.000278$	(IAEA, 1988)
TPAMF(1–4)	uniform	$u = 10300$	(IAEA, 1988)	TPIRDF1	uniform	$u = 3.09$	(Expert Opinion)
TPAMSD1	delay	$a = 4383.0$	$^N/_A$	TPIRDO1	delay	$a = 0.000278$	(Expert Opinion)
TPAO(1–4)	delay	$a = 0.000278$	(IAEA, 1988)	TPIRDOF1	delay	$a = 0.000278$	(Expert Opinion)
TPAPR(1–4) a	Weibull	$\eta = 1.67 \times 10^7$, $\beta = 1.0$	(IAEA, 1988)	TPIRDOP1	uniform	$u = 0.235$	(Expert Opinion)
TPARDF(1–4)	uniform	$u = 3.09$	(Expert Opinion; Eide et al., 1990)	TVA1	delay	$a = 8766.0$	$^N/_A$

Table 4
Timed transition parameters for emergency shutdown subsystem Petri Net seen in Fig. 10.

Transition	Type	Parameter(s)	Ref.	Transition	Type	Parameter(s)	Ref.
TBMRVF1	uniform	$u = 4.41$	(Miller et al., 1976)	THESD1	delay	$a = 2.0$	$^N/_A$
TBMRVO1	delay	$a = 0.000278$	(Miller et al., 1976)	THESF1	Weibull	$\eta = 1.12 \times 10^6$, $\beta = 3.0$	(Expert Opinion; Reliability Eta Beta database, 2020)
TBMRVSO1	Weibull	$\eta = 8.64 \times 10^5$, $\beta = 1.1$	(IAEA, 1988; Reliability Eta Beta database, 2020)	THESUA1	delay	$a = 0.25$	$^N/_A$
TBPVF1	uniform	$u = 4.41$	(Miller et al., 1976)	TRACTF (1–40)	Weibull	$\eta = 1060$, $\beta = 1.2$	(Smith, 1981; Reliability Eta Beta database, 2020)
TBPVO1	delay	$a = 0.000278$	(Miller et al., 1976)	TRACTR (1–40)	delay	$a = 24.0$	$^N/_A$
TBPVSO1	Weibull	$\eta = 8.64 \times 10^5$, $\beta = 1.1$	(IAEA, 1988; Reliability Eta Beta database, 2020)	TRDLF(1–40)	uniform	$u = 2.78$	(IAEA, 1988)
TBRTVF1	uniform	$u = 4.41$	(Miller et al., 1976)	TRDLS(1–40)	delay	$a = 0.000278$	(IAEA, 1988)
TBRTVS1	delay	$a = 0.000278$	(Miller et al., 1976)	TRIF(1–40)	uniform	$u = 9.26$	(Eide and Calley, 1993)
TBRTVSO1	Weibull	$\eta = 8.64 \times 10^5$, $\beta = 1.1$	(IAEA, 1988; Reliability Eta Beta database, 2020)	TRIS(1–40)	delay	$a = 0.000278$	(Eide and Calley, 1993)
TBTL1	Weibull	$\eta = 4.16 \times 10^7$, $\beta = 3.0$	(IAEA, 1988; Reliability Eta Beta database, 2020)	TRPSF(1–40)	uniform	$u = 1.07$	(Expert Opinion)
TBTSUA1	delay	$a = 0.25$	$^N/_A$	TRPSS(1–40)	delay	$a = 0.000278$	(Expert Opinion)
THERMF1	uniform	$u = 0.278$	(Expert Opinion)	TSDS1	uniform	$u = 27258.0$	$^N/_A$
THERMS1	delay	$a = 0.000278$	(Expert Opinion)	TVA1	delay	$a = 8766.0$	$^N/_A$
THERMSO1	Weibull	$\eta = 1.73 \times 10^5$, $\beta = 1.1$	(Smith, 1981; Reliability Eta Beta database, 2020)	–	–	–	–

pressure accumulator itself may function correctly or fail, seen as TPAO(1–4) and TPAMF(1–4) respectively, and while waiting for demand for the system to be placed, the rupture disk separating the pressure accumulator from the core may prematurely burst, see TPASC(1–4), rendering it unavailable on demand. Hence, an inhibit arc from PPASFO(1–4) suppresses TPAO(1–4), and TPAUA (1–4) removes the token from PREECC(1–4). The number of pressure accumulators failed before demand is recorded at PPAMR1, with maintenance scheduled six months following one failure, or immediately if a second fails. TPAAOVF(1–4) and TPAAOVO(1–4) represent the failure to open or successful opening of the AOV, with TPARDF1 and TPARDO1 respectively performing the same roles for the rupture disk. Either one token or two is given to PPARDO(2/4/6/8), depending on whether TPARDOP1 or TPARDOP2 fires. Two tokens indicate that the rupture disk opened fully, with a partial rupture being indicated by one token. The voting transition TPAF1 adds two tokens to PPARF1 for each failed valve. A partial disk rupture adds one. This tracks the equivalent high pressure injection capacity loss. For example, if all of the rupture disks partially opened or if two of the AOVs failed to open, there would remain sufficient capacity in both cases. Active pressure accumulator capacity is tallied at PPASA, fulfilling the requirements for TPASA1 to fire once the equivalent of two accumulators come online. TPASA1 fires after a half-hour delay to mark the end of high pressure injection. However, the pipes connecting each pressure accumulator to must remain intact in sufficient numbers for this duration, otherwise, its high pressure coolant contribution is lost midway upon pipe rupture, see TPAPR1a and TPAPR1b.

Once the pressure accumulators are exhausted, low pressure injection begins, wherein coolant flows under gravity from the over head pool into the core. First, the non-return valve must open, success and failure of this being TPINRO1 and TPINRF1. If open, the

Table 5
Timed transition parameters for reactor control and reactivity monitoring Petri Net seen in Fig. 11.

Transition	Type	Parameter(s)	Ref.	Transition	Type	Parameter(s)	Ref.
TAH1H	uniform	$a = 4.0$	$^N/_A$	TFRDIC(1–3)	Weibull	$\eta = 1.33 \times 10^5$, $\beta = 1.0$	(Expert Opinion)
TAH1L	uniform	$a = 4.0$	$^N/_A$	TFRDND(1–3)	Weibull	$\eta = 1.33 \times 10^5$, $\beta = 1.0$	(Expert Opinion)
TASHF1	Weibull	$\eta = 34500$, $\beta = 1.1$	(Smith, 1981; Barringer & Associates, Inc., 2010)	TFRIIC(1–3)	Weibull	$\eta = 1.33 \times 10^5$, $\beta = 1.0$	(Expert Opinion)
TASSI1	Weibull	$\eta = 100000$, $\beta = 1.0$	(Expert Opinion)	TFRICD(1–3)	Weibull	$\eta = 1.33 \times 10^5$, $\beta = 1.0$	(Expert Opinion)
TCDIC(1–3)	delay	$a = 0.0001$	(Expert Opinion)	TFRICR(1–3)	uniform	$a = 4.0$	(Expert Opinion)
TCDND(1–3)	delay	$a = 0.0001$	(Expert Opinion)	TFRNDR(1–3)	uniform	$a = 4.0$	(Expert Opinion)
TCRD1	delay	$a = 2.0$	$^N/_A$	TICVE	delay	$a = 0.0$	$^N/_A$
TCRDF1	Weibull	$\eta = 20000$, $\beta = 1.0$	(Expert Opinion)	TNDVE	delay	$a = 0.0$	$^N/_A$
TCRE1H	uniform	$a = 4.0$	$^N/_A$	TRAS1H	uniform	$a = 4.0$	$^N/_A$
TCRE1L	uniform	$a = 4.0$	$^N/_A$	TRAS1L	uniform	$a = 4.0$	$^N/_A$
TCRI1	delay	$a = 2.0$	$^N/_A$	TRE1	Weibull	$\eta = 87700$, $\beta = 1.0$	(Expert Opinion)
TCRS1H	uniform	$a = 4.0$	$^N/_A$	TRTC1	delay	$a = 6.0$	$^N/_A$
TCRS1L	uniform	$a = 4.0$	$^N/_A$	TRTH1	delay	$a = 6.0$	$^N/_A$
TCRSF1	Weibull	$\eta = 20700$, $\beta = 1.1$	(Expert Opinion; Barringer & Associates, Inc., 2010)	TSPEF(1–3)	Weibull	$\eta = 17100$, $\beta = 1.1$	(Expert Opinion; Barringer & Associates, Inc., 2010)
TFDDIC(1–3)	uniform	$u = 0.1$	(Expert Opinion)	TSPER(1–3)	delay	$a = 1.0$	$^N/_A$
TFDDND(1–3)	uniform	$u = 0.1$	(Expert Opinion)	TSRIC(1–3)	Weibull	$\eta = 1.43 \times 10^5$, $\beta = 1.0$	(IAEA, 1988; Barringer & Associates, Inc., 2010)
TFDIIC(1–3)	uniform	$u = 0.1$	(Expert Opinion)	TSRICR(1–3)	cyclic	$c = 4383.0$	$^N/_A$
TFDICD(1–3)	uniform	$u = 0.1$	(Expert Opinion)	TSRND(1–3)	Weibull	$\eta = 1.43 \times 10^5$, $\beta = 1.0$	(IAEA, 1988; Barringer & Associates, Inc., 2010)
TFDICR(1–3)	uniform	$a = 4.0$	(Expert Opinion)	TSRNDR(1–3)	cyclic	$c = 4383.0$	$^N/_A$
TFDNDR(1–3)	uniform	$a = 4.0$	(Expert Opinion)	–	–	–	–

Table 6
Results from the Petri Nets seen in Figs. 7–11, showing both overall operating durations and those of each of the possible outcomes. Additional results for SDS-2 are found in Table 7, and a break-down predictions relating to individual sensor readings in the control and monitoring system is found in Table 8.

Subsystem	Total Iterations	Average Duration [hours]	Outcome	Probability [%]	Outcome Average Duration [hours]	Percentiles [hours] 10th	90th
Primary Coolant Circulation	10^5	$(1.900 \pm 0.003) \times 10^4$	Coolant Faults	0.314 ± 0.018	$(1.31 \pm 0.04) \times 10^4$	3.88×10^3	2.41×10^4
			Safe Shutdown	99.686 ± 0.018	$(1.902 \pm 0.003) \times 10^4$	7.241×10^3	2.726×10^4
Shutdown Condensers	5.24×10^5	$(1.5193 \pm 0.0013) \times 10^4$	SDC/DHR Over	99.9998 ± 0.0002	$(1.5219 \pm 0.0013) \times 10^4$	2.9641×10^3	2.7258×10^4
			SDC Failure	0	$^N/_A$	$^N/_A$	$^N/_A$
			Emergency Shutdown	0.0002 ± 0.0002	9.36×10^3 (*One data point*)	$^N/_A$	$^N/_A$
Emergency Core Coolant Injection	10^5	$(1.096 \pm 0.002) \times 10^4$	Maintenance Shutdown	32.63 ± 0.15	$(1.049 \pm 0.003) \times 10^4$	4.950×10^3	1.801×10^4
			PA System Failure	0	$^N/_A$	$^N/_A$	$^N/_A$
			GDWP Failure	0.042 ± 0.007	$(1.30 \pm 0.14) \times 10^4$	1.52×10^3	2.61×10^4
			Core Submerged	67.33 ± 0.15	$(1.118 \pm 0.003) \times 10^4$	1.925×10^3	2.261×10^4
Emergency Shutdown (*See Table 7 for* SDS-2)	10^5	$(1.340 \pm 0.0020) \times 10^4$	SDS-1 Complete	96.77 ± 0.06	$(1.352 \pm 0.003) \times 10^4$	2.6562×10^3	2.4459×10^4
			SDS-2 Complete	0	$^N/_A$	$^N/_A$	$^N/_A$
			SDS Failure	0	$^N/_A$	$^N/_A$	$^N/_A$
			Forced Shutdown	0.0010 ± 0.0010	1.35×10^4 (*One data point*)	$^N/_A$	$^N/_A$
			Unplanned Shutdown	3.23 ± 0.06	$(9.64 \pm 0.11) \times 10^3$	1.764×10^3	1.899×10^4
Control & Monitoring (*See Table 8 for more details*)	5×10^5	$(1.2855 \pm 0.0018) \times 10^4$	Reactor Supercritical	0.0012 ± 0.0005	$(6.7 \pm 1.7) \times 10^3$	2.6×10^3	1.2×10^4
			Reactor Subcritical	0.0006 ± 0.0003	$(1.1 \pm 0.7) \times 10^4$	1.3×10^3	2.4×10^4
			IC Reactivity Increase Detection	24.97 ± 0.06	$(1.285 \pm 0.004) \times 10^4$	1.496×10^3	2.925×10^4
			IC Reactivity Decrease Detection	24.95 ± 0.06	$(1.286 \pm 0.004) \times 10^4$	1.480×10^3	2.904×10^4
			IC Contradiction	0.0024 ± 0.0007	$(1.5 \pm 0.4) \times 10^4$	4.0×10^3	2.5×10^4
			ND Reactivity Increase Detection	25.04 ± 0.06	$(1.281 \pm 0.004) \times 10^4$	1.485×10^3	2.903×10^4
			ND Reactivity Decrease Detection	25.03 ± 0.006	$(1.288 \pm 0.004) \times 10^4$	1.496×10^3	2.918×10^4
			ND Contradiction	0.0014 ± 0.0005	$(1.7 \pm 0.7) \times 10^4$	1.7×10^3	4.4×10^4

Note: IC = Ion Chambers, ND = Neutron Detectors.

valve is followed by the rupture dish. As before, the rupture disk may open fully, partially open, or fail completely, as captured by TPIRDO1, TPIRDOP1, and TPIRDF1 respectively. Full coolant flow allows the low pressure injection to be complete in three days, with partially flow, this is extended to six days. As low pressure injection relies on a reservoir of water, leakage from which will prevent its proper function. Leaks are represented by TSCPL1. If a leak occurs before demand, the reactor must be shutdown for cor-

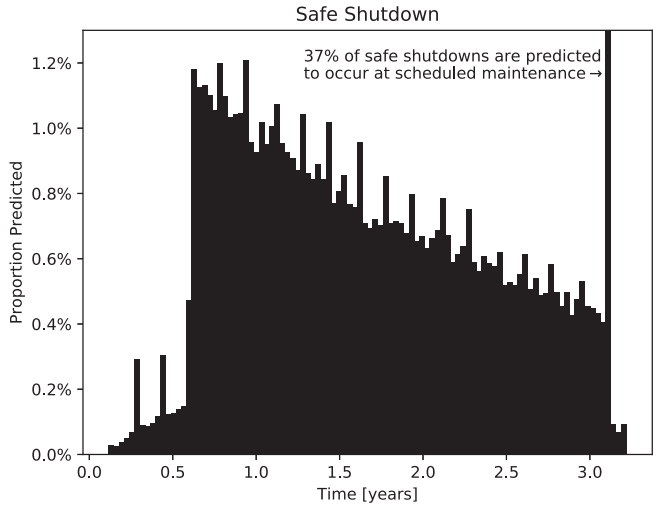

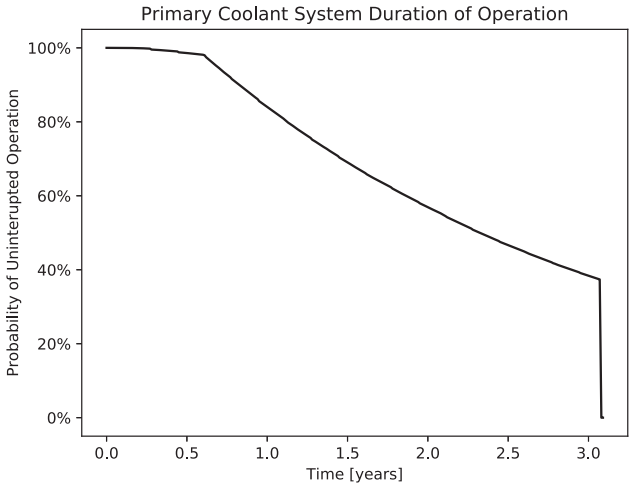

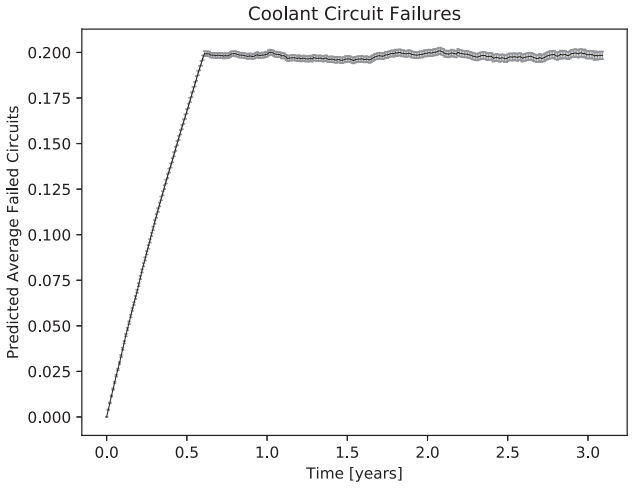

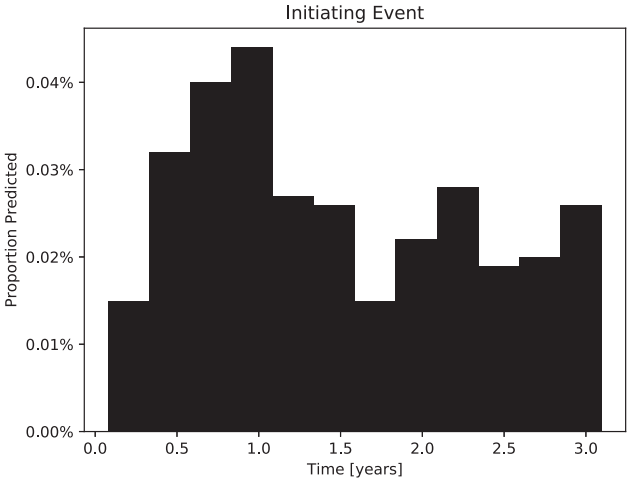

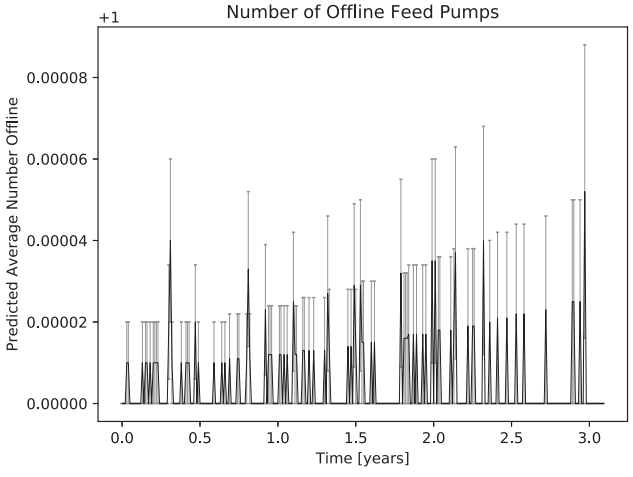

Fig. 12. Results from the primary coolant circulation system. (a) Safe Shutdown – Predicted distribution of operational duration of systems reaching a safe shutdown state in the primary coolant system. (b) Initiating Event – Predicted distribution of operational duration of systems reaching a critical failed state. (c) Primary Coolant System Duration of Operation – Probability of the primary coolant system operating uninterrupted over time (*i.e.* not encountering a critical failure or the need to shut down for repair). Note that reactor shutdown is scheduled to begin after three years of operation. (d) Coolant Circuit Failures – Predicted average number of failed steam separator coolant circuits over operating time. (e) Number of Offline Feed Pumps – Predicted average number of off-line feed pumps over operating time. Note that one pump should always be offline as a standby. (f) Critical Feed Pump Failure – Predicted average probability of encountering critical feed pump failure over operating time. (g) Turbine Isolation Demand – Probability of demand for turbine isolation over operating time. (h) Steam Circuit Failures – Probability of varying numbers of failed steam circuits during operation of the reactor. (i) Number of Failed Steam Circuits Causing Early Shutdown – Number of Failed Steam Circuits Causing Early Shutdown – Probability of early shutdown being caused by one or two steam circuit failures during operation of the reactor.

rective actions, thus TSCPL1 fires. If the leak happens during the coolant injection process, TSCPFP1 fires instead, marking failure of emergency core coolant injection.

All valves become more likely to fail on demand with each year that passes, and this is modelled by the place conditional arc from PV2 to those transitions representing valve failures.

4.4. Emergency shutdown systems

In the Petri Net for the emergency shutdown systems found in Fig. 10, the yellow and green sections respectively model the emergency shutdown rod insertion (SDS-1) and the boric acid neutron poison moderation injection process (SDS-2). Possible outcomes for this Petri Net are the shutdown of the reactor by either SDS-1 or SDS-2, the failure of both shutdown systems, the forced shutdown of the reactor for urgent maintenance, or an unplanned shutdown due to premature actions of components in the shutdown process.

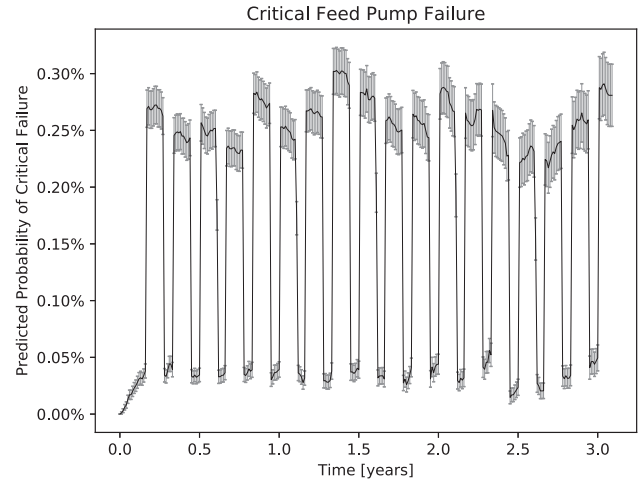

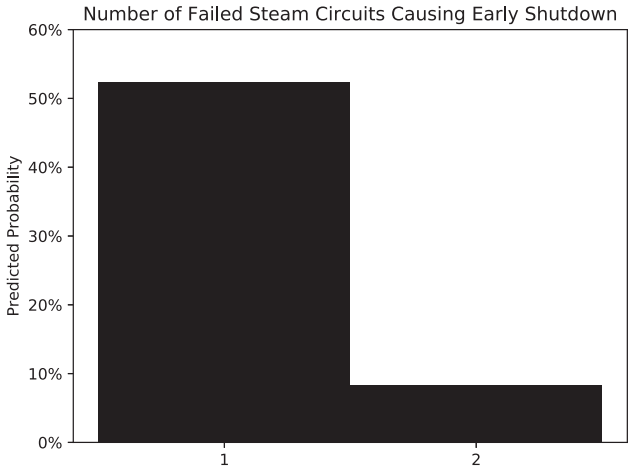

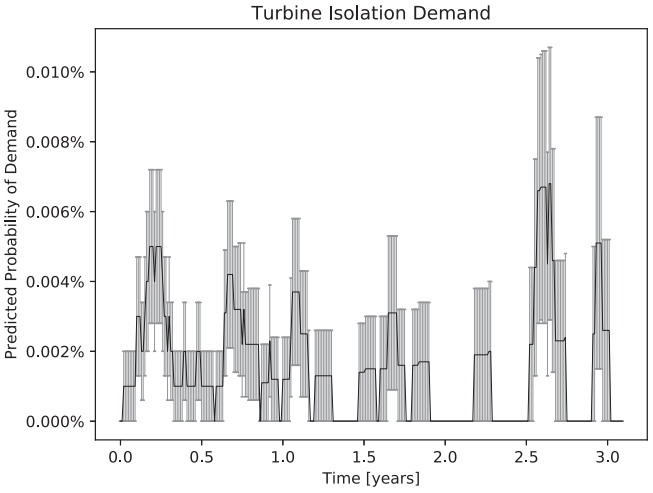

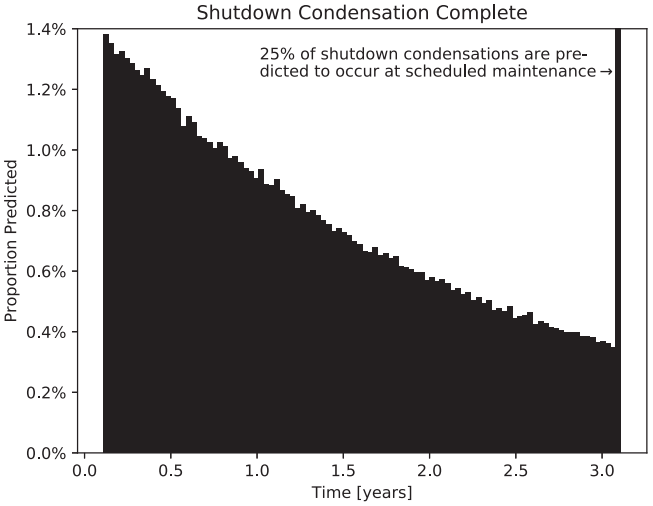

Fig. 13. Results from the shutdown condenser system. (a) Shutdown Condensation Complete – Predicted distribution of time of successful completion of the reactor shutdown condensation process. (b) Shutdown Condenser Failures – Predicted distribution of the number of shutdown condenser units failing during the shutdown condensation process.

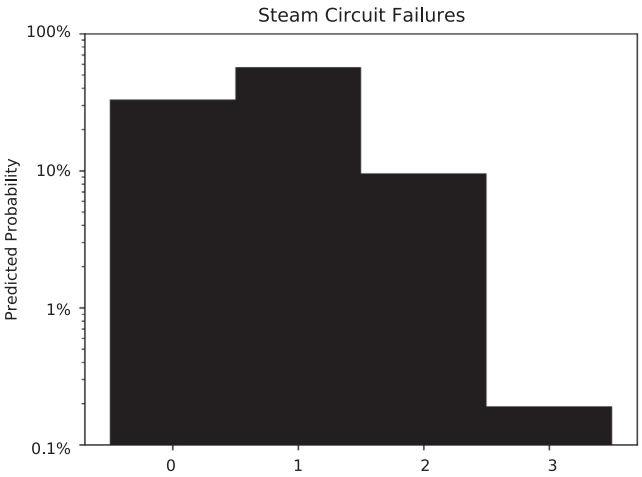

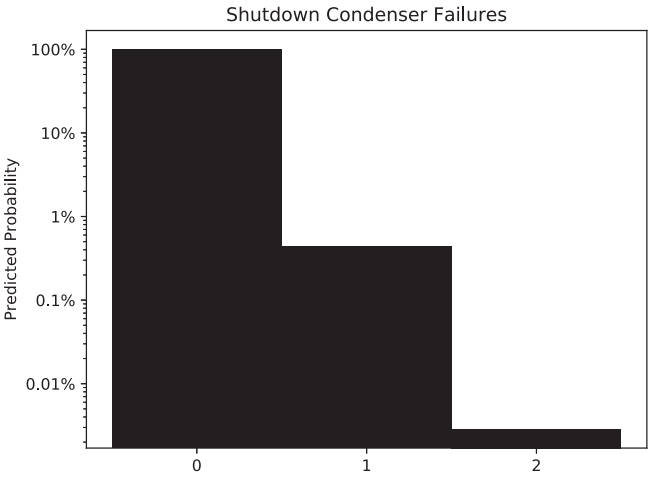

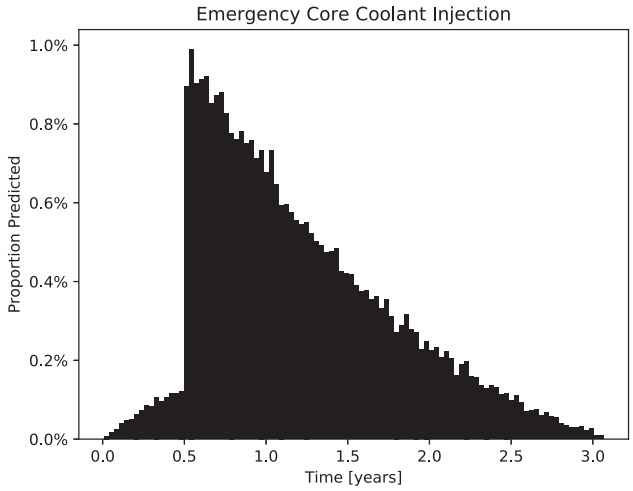

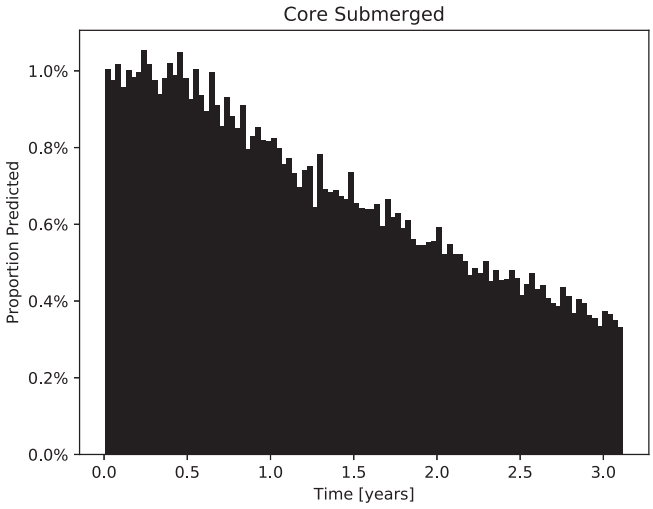

Fig. 14. Results from the emergency core coolant injection system. (a) Emergency Core Coolant Injection – Predicted distribution of operating time at which faults in the emergency core coolant injection system required the reactor to shutdown for maintenance. (b) Failure of Gravity Driven Low Pressure Coolant Injection – Predicted distribution of operational time at which the gravity driven injection of low pressure coolant failed. (c) Core Sumberged – Predicted distribution of operational time at which core submersion was completed by the emergency core coolant injection system. (d) Failed Pressure Accumulators – Predicted distribution of the number of pressure accumulators failing during demand for high pressure injection, representing the emergency core coolant injection system.

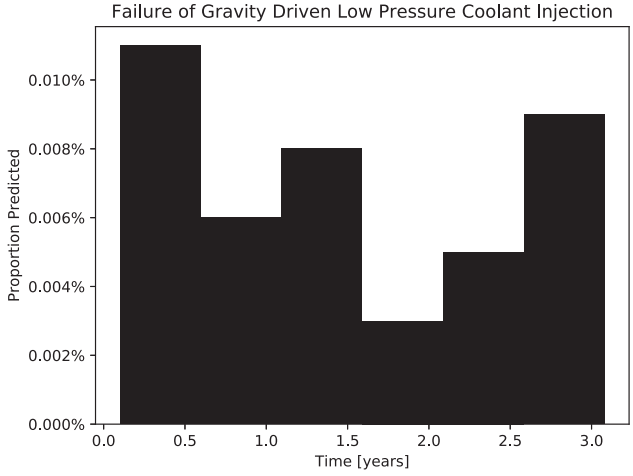

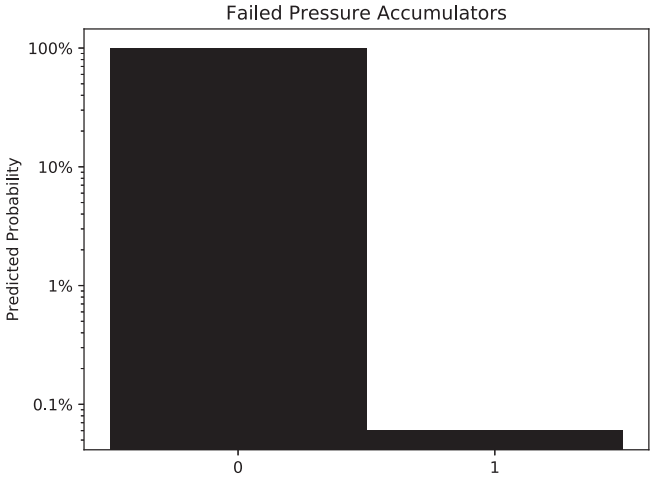

As with the emergency core coolant injection Petri Net, demand is generated at an arbitrary point in the operation period with uniform probability density by transition TSDS1, which puts a token on each of PRIB(1–40). To properly insert, an individual rod's protection system signal and de-latch mechanism must both operate correctly, and the rod must avoid jamming on its descent into its slot in the core, with the transitions representing those successes and failures respectively being TRPSS(1–40), TRDLS(1–40), and TRIS(1–40), and TRPSF(1–40), TRDLF(1–40), and TRIF(1–40). The actuator can also be disabled by failure before demand, see TRACTF1, with TRACTR1 representing its restoration. Every rod failure puts a token on PNRF1 and every correct insertion adds to PNRI1. SDS-1 is considered to have succeeded if 38 of the 40 rods insert fully. Therefore, if the token count at PNRI1 reaches that threshold, TSDS1 fires, indicating the success of SDS-1 and ending the simulation. Conversely if three or more tokens reach PNRF1, TSDS1F fires, marking the failure of SDS-1 and initiates demand for SDS-2.

The second shutdown system, SDS-2, injects the potent neutron poison, boric acid, directly into the core's heavy water moderator. The poison is driven by the releases of pressurised helium, the release mechanism for which may fail on demand. This is represented by THERMF1, and results in the failure of SDS-2. Otherwise, correct action is represented by THERMS1. The boric acid must then pass through of the parallel valves (pressure driven, reactor trip signal activated, and manually opened), and as such the correct opening of any one of which is sufficient alone for SDS-2 to be successful. If all three valves failure to open, SDS-2 fails. The transitions repenting the opening and failure of these valves are respectively, TBPVO1, TBRTVS1, and TBMRVO1, and TBPVF1, TBRTVF1, and TBMRVF1, the later having a place conditional relationship with PV2, increasing the risk of failure on demand with each year. The premature release of the helium pressure or opening of one of the boric acid release valve causes and unplanned shutdown resulting from undesired poison in the moderator, respectively represented as THERMSO1, TBPVSO1, TBRTVSO1, TBMRVSO1. If leaks develop in either the helium supply or the tank of boric acid, respectively THESF1 and TBTL1, the operators are forced to shutdown the reactor to enact immediate repairs.

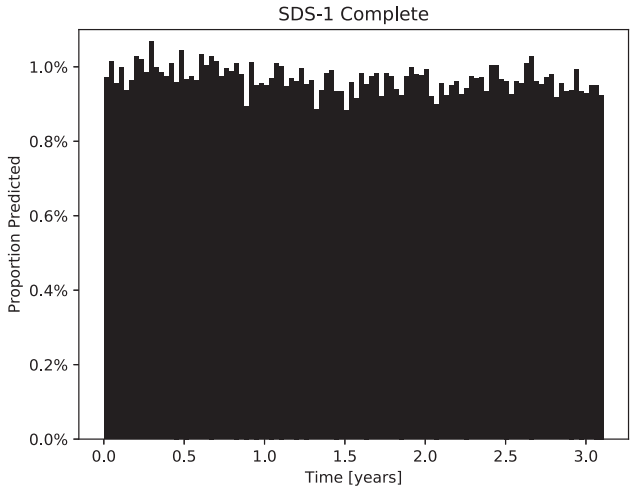

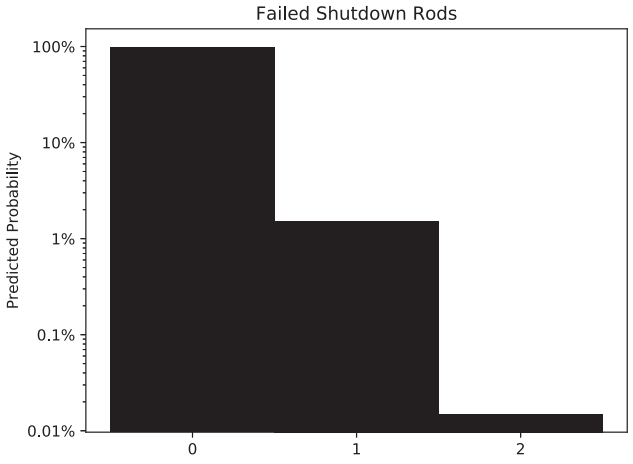

Fig. 15. Results from the emergency reactor shutdown systems. (a) SDS-1 Complete – Predicted distribution of operation time at which SDS-1 successfully completed following a call for reactor emergency shutdown. (b) Unplanned Shutdown – Predicted distribution of operational time at which an unplanned shutdown occurred as a result of premature action on the part of a component in the emergency shutdown system. (c) Failed Shutdown Rods – Predicted distribution of the number of shutdown rod failures during use of the emergency shutdown systems.

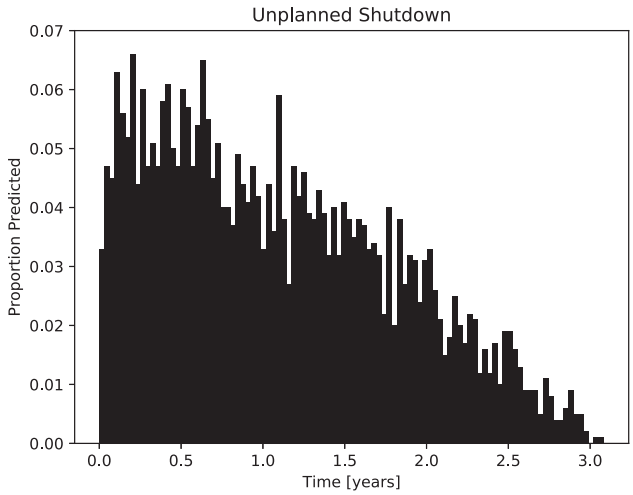

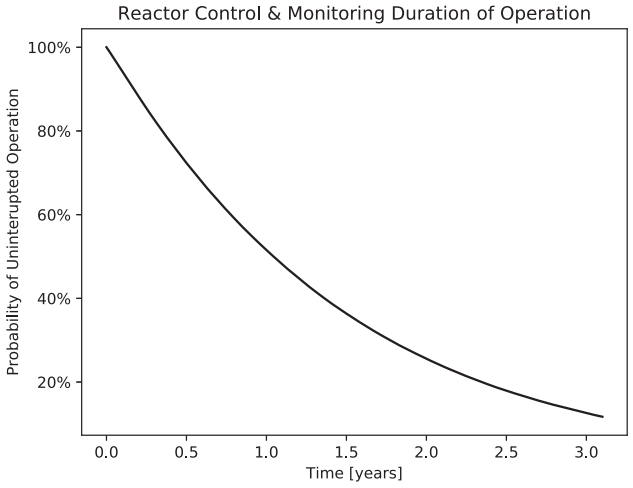

Fig. 16. Results from the reactor control and monitoring systems. (a) Reactor Control & Monitoring During of Operation – Probability of uninterrupted operation before the emergence of a critical fault in the reactor control and monitoring systems. This graph is truncated at the end of scheduled reactor shutdown at 3 years 40 days. (b) Overall Sensor Readiness – Probability of individual sensor readiness for all ion chambers and neutron detectors in the reactor control and monitoring system, where "sensor readiness" refers to a sensor being in a state of readiness to correctly report a change in reactor reactivity. The graph is truncated at the end of scheduled reactor shutdown after 3 years 40 days.

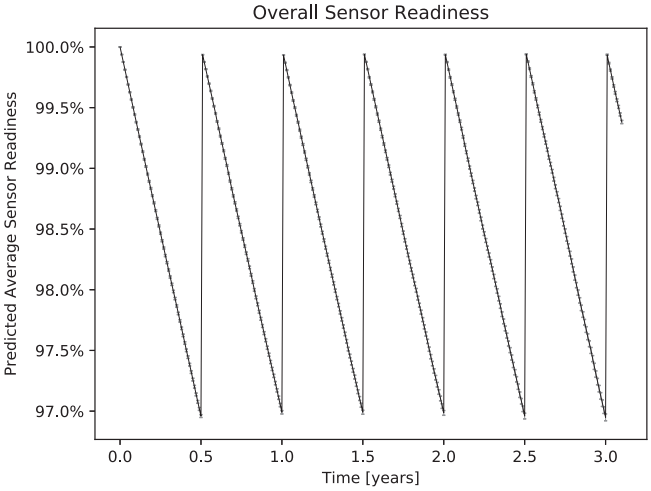

4.5. Reactor control & monitoring

This Petri Net, seen in Fig. 11, is concerned with control systems of the reactor and the sensors used to monitor its state. Outcomes for this Petri Net include the undetected emergence of a supercritical or subcritical state, and the detection by either the ion chambers or neutron detectors of such a problem, with the additional probability of contradictory results being produced by either sensor block.

The control rods can experience a fault in their drive system or signal controlling them, either of which will result in the rod being in an undesired position. Similarly, the rods will be incorrectly placed if a hardware or software failure emerges in the automation system. Therefore, each of these faults will result in the control rods being too high or too low. Note these outcomes are abbreviated, such that P(IC/ND) VI is the result of the ion chamber/neutron detector sensor block voting on a reactivity increase, P(IC/ND) VD

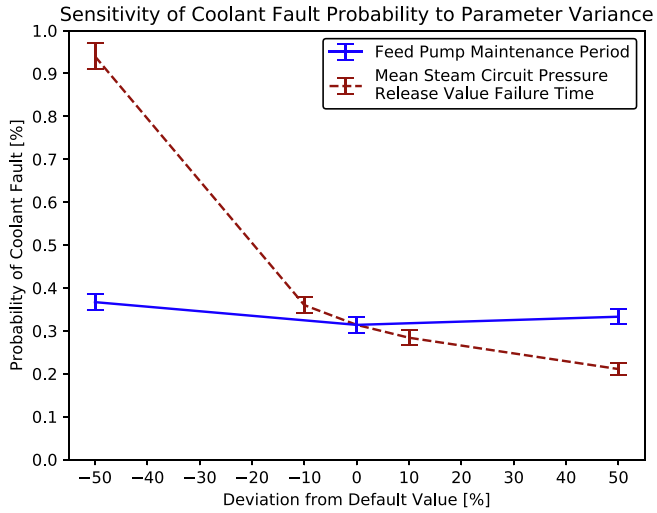

Fig. 17. Effect on the likelihood of a coolant fault from varying two parameters – feed pump service period, and mean time to failure of the steam separator pressure release valves.

Table 7
Probability of outcomes from demands on SDS-2.

Outcome	Proportion [%]
SDS-2 Complete	99.899±0.010
SDS-2 Failed	0.101±0.010

Table 8
Summary of probability of the possible readings given by all individual ion chambers and neutron detectors in the reactor control and monitoring system, as represented by the Petri Net in Fig. 11. Of false positives, (24.99±0.04)% were predicted to coincide with another in the same sensor block, such as to produce an overall false positive from the voting system.

Sensor Reading	Proportion [%]
True Positive	82.00±0.02
False Positive	17.96±0.02
False Negative	0.0446±0.0012

is the result of the ion chamber/neutron detector sensor block voting on a reactivity decrease, and P(IC/ND) VE is the result of the ion chamber/neutron detector sensor block voting producing a contradictory results – for example one sensor reading a reactivity increase and another reading a decrease.

A reactivity event can arise as a result of issues with the automation system, the control rods themselves, or from problems internal to the reactor, and the sections of the Petri Net dealing with each of these are respectively highlighted red, yellow, and blue on Fig. 11a. The automation system may suffer either hardware or software failure, and the control rods can develop a fault in their electric drive or ability to receive and respond to control signals. These are respectively represented by TASHF1, TASSI1, TCRDF1, and TCRSF1. When a such an issue arises, the result will be either the rods in a position higher than intended, or lower than intended, with transitions TCRF(1–4) and TASF(1–4) making the selection. There is a window of opportunity to effect restorative actions, represented by TAH1H, TAH1L, TRAS1H, TRAS1L, TCRE1H, TCRE1L, TCRS1H, and TCRS1L, but failure to act in time results in a reactivity event, recorded by the firing of TCRI1 and TCRD1. TRE1, represents the occurrence of a reactivity event resulting from an internal problem within the core, with TREI1 and TRED1 determining whether it is an increase or decrease in reactivity.

Once a reactivity event is ongoing, the goal of the sensor blocks is to detect it before it develops into a more serious problem. This is dealt with in the green section of the Petri Net. There are two sensor blocks, consisting of three ion chambers (IC) and three neutron detectors (ND). The structure shown represents one individual sensor and is therefore repeated as two sets of three. When a token is in its default, waiting for stimulation, a token is seen at PNR(IC/ND)(1–3). If a token appears at PRE to indicate a real event and the sensor functions as intended, TCD(IC/ND)(1–3) fires to record it, with TCDI(IC/ND)(1–3) and TCDD(IC/ND)(1–3) respectively categorising it as an increase or decrease in reactivity according to whether a token is at PREI1 or PRED1. Otherwise, the sensor may fail to detect the event, either because it has become temporarily unresponsive, with TSR(IC/ND)(1–3) and TSR(IC/ND)(1–3) representing its loss of responsiveness and return respectively, or because it fails on demand, represented by TFDI(IC/ND)(1–3) and TFDD(IC/ND)(1–3) for reactivity increases and decreases respectively. TFD(IC/ND) R(1–3) allows additionally opportunities for the sensor to respond after a delay. Temporary false positive readings are also a possibility, where false positive increases in reactivity are represented by TFRI(IC/ND)(1–3) and false positive decreases by TFRD(IC/ND)(1–3), and TFR(IC/ND) R(1–3) returns the sensor to its default state.

The section of the Petri Net found in Fig. 11b deals with the voting system. If any two sensors concurrently give the same reading, the operators of the reactor are supplied with that information such that action may be taken. T(IC/ND) VI fires for a reactivity increase measurement and T(IC/ND) VD for a decrease. If two sensors of a block are both producing a reading, but those readings are contradictory, T(IC/ND) VE fires to alert the operator that some problem has arisen, even if its exact nature is not known. These messages are dependant on the signal processing electronics, of which there are three sets. Providing that at least two are in normal working order, messages from the sensor block will be transmitted. The failure of one of the units in the signal processing electronics is represented by TSPEF(1–3), and the number of failures is recorded at PSPE4. TSPER(1–3) represent the repair of the units, putting tokens back to PSPE(1–3). If the number of tokens on PSPE4 reaches two, inhibit arcs to T(IC/ND) VI, T(IC/ND) VD, and T(IC/ND) VE are satisfied, preventing any signals being sent from the ion chamber or neutron detector blocks.

5. Results

Each sub-system was simulated from the start of its operation until it arrived at some terminal state which represented the safe shutdown of the system or the occurrence of a failed condition. After a significant number of simulations (iterations), enabling convergence in the predictions, the probability of arriving at each of the possible terminal states was determined and is reported in Table 6, with convergence considered to have been achieved when the error bounds of features of interest have reduced to be small relative to their absolute value. Along with the total iterations, the table also indicates the average time duration taken to arrive at each end state and the corresponding 10th and 90th percentiles. Where relevant, the confidence bounds given in the figures and tables of this work are quantified by the standard error of the mean, *i.e.* $\pm 1\sigma$ confidence bounds of approximately 68%. The quantification of the uncertainty attached to the results is helpful for two purposes: – Firstly, when the results of one simulation are used to determine parameters in another, the inclusion of the confidence bounds accounts for the uncertainty of the value introduced from another model, such as seen in Section 4.2 – Secondly, in a practical setting, error analysis is useful when assessing a specific system to highlight whether sufficient compu-

Table 9
Results from sensitivity analysis performed on the primary coolant circulation Petri Net seen in Section 4.1.

Variant	Average Duration [hours]	Outcome	Probability [%]	Outcome Average Duration [hours]	Percentiles [hours] 10th	90th
FPSP + 50%	$(1.912\pm0.003)\times10^4$	Coolant Faults	0.333±0.018	$(1.35\pm0.04)\times10^4$	4.47×10^3	2.40×10^4
		Safe Shutdown	99.667±0.018	$(1.914\pm0.003)\times10^4$	7.364×10^3	2.726×10^4
FPSP −50%	$(1.881\pm0.003)\times10^4$	Coolant Faults	0.367±0.019	$(1.40\pm0.04)\times10^4$	4.41×10^3	2.46×10^4
		Safe Shutdown	99.633±0.019	$(1.882\pm0.003)\times10^4$	7.104×10^3	2.726×10^4
SSPRV + 50%	$(2.090\pm0.002)\times10^4$	Coolant Faults	0.211±0.015	$(1.48\pm0.05)\times10^4$	4.66×10^3	2.43×10^4
		Safe Shutdown	99.789±0.015	$(2.091\pm0.002)\times10^4$	8.229×10^3	2.73×10^4
SSPRV + 10%	$(1.948\pm0.003)\times10^4$	Coolant Faults	0.284±0.017	$(1.33\pm0.04)\times10^4$	4.77×10^3	2.44×10^4
		Safe Shutdown	99.716±0.017	$(1.950\pm0.003)\times10^4$	7.498×10^3	2.726×10^4
SSPRV −10%	$(1.841\pm0.003)\times10^4$	Coolant Faults	0.360±0.019	$(1.34\pm0.04)\times10^4$	4.12×10^3	2.47×10^4
		Safe Shutdown	99.640±0.019	$(1.843\pm0.003)\times10^4$	7.043×10^3	2.726×10^4
SSPRV −50%	$(1.478\pm0.002)\times10^4$	Coolant Faults	0.94±0.03	$(1.13\pm0.02)\times10^4$	3.92×10^3	2.14×10^4
		Safe Shutdown	99.06±0.03	$(1.481\pm0.002)\times10^4$	6.05×10^3	2.726×10^4

Note: FPSP = Feed Pump Service Period, SSPRV = Steam Separator Pressure Release Valve.

tational power has been expended for the results to be considered reliable when making safety critical decisions relating to the design in question.

The results of the subsystem simulations are presented in Figs. 12–16, following the same order as that in which the models were presented in Section 4. Histograms in the figures depict the variation in the predicted durations to achieve each terminal state are found in Figs. 12a, 12b, 13a, 14a, 14b, 14c, 15a, and 15b. Figs. 12c and 16a show the respective predicted probabilities that the primary coolant system and the reactor control and monitoring system will be continuing to operate uninterrupted after a given period from the beginning of their operation, without shutting down or encountering a critical error. Predictions for five indicators of the operational health of the reactor across time are found in Figs. 12d, 12e, 12f, 12g, and 16b. These are respectively the average number of coolant circuits to have failed, the average number of feed pumps in an off-line state, the probability of encountering the critical threshold of feed pump failures, the probability of demand for turbine isolation, and the probability that each individual sensor in the ion chamber and neutron sensor blocks is in a ready state to produce an accurate reading of changes in reactivity status. The histograms found in Figs. 12h, 13b, 14d, and 15c respectively depict the predicted probability of a given number of failed steam circuits, shutdown condensers, pressure accumulators, and shutdown rods failing on demand. The histogram in Fig. 12i shows the probability that early shutdown will be caused by either one or two steam circuit failures within the operation of the reactor system.

6. Discussion

The simulations predict a high expectation that the primary coolant system will shut down safely when maintenance is required, with a low likelihood of requiring emergency intervention. The predicted average operational period until shutdown or failure was just under 2.2 years, with a 61.3% probability of requiring early maintenance before the end of the default three-year operating period. In Figs. 12a and 12c, the prominent spike in the probability of shutdown just after 0.6 years is the result of the delay between the emergence of a single steam separator circuit fault and the command to shut down to perform repair actions.

As seen in Fig. 12h, there is a 66.9% chance that one or more steam circuits will fail, 9.78% chance that two or more will fail, and a 0.191% chance that the critical threshold of three failed circuits will be reached. From Fig. 12i it is seen that there is a likelihood of 52.3% of a single steam circuit failure occurring without a second failure subsequently forcing immediate shutdown before the end of the six-month delay, and that there is an 8.29% percent chance of two steam circuit failures within six months of each other. Thus, once one steam circuit has failed, there is an 86.4% probability that a second will fail in that period.

Small spikes in the need for early shutdown probability are seen periodically in Fig. 12a, corresponding to the servicing of the coolant feed-pumps, wherein the failure to start the back-up pump may arise, which is reflected in both the likelihood of critical of feed pump failure and demand for turbine isolation, see Figs. 12f and 12g. Overall, critical feed pump failure is predicted to occur with 3.67% probability.

By inspecting the results from the model it was found that given a demand for turbine isolation, there is a 0.889% probability of subsequently encountered an initiating event during the shutdown condenser period, and a 0.944% chance of requiring the use of the back-up isolation valve due to undesired opening of the main valve. A fault more serious than that which could be handled by the primary coolant circulation system alone is predicted to occur in 0.314% of cases, at average time of 1.31×10^4 hours from the beginning of operation. Given that the reactor is predicted to be shut down on average once for every 1.90×10^4 hours, if repairs and maintenance take an average of two weeks, the resulting unavailability of primary coolant circulation would be 1.24%. Reducing repairs and maintenance to 24 h would bring the figure down to 0.0889%.

The simulations yielded a high probability of availability for the shutdown condensation system with a likelihood of approximately 0.0002% for the need to shut down to perform repairs within the given period. With regards failures during demand on the system, one or more failed shutdown condensers are predicted to be encountered with a probability of 0.437%, with the likelihoods of 0.434% and 0.00286% predicted for one and two failed condensers respectively, see Fig. 13b.

The emergency core cooling injection process had a high likelihood of requiring maintenance at a probability almost one in three, but given demand for the system, there was a 99.94% likelihood of successful core submersion. The low pressure injection phase represents the highest period of risk, see Fig. 14b, while even manageable disruptions to high pressure injection following demand are predicted to occur with a probability of only 0.0891%, as seen in Fig. 14d.

As visible in Fig. 15c, within the emergency shutdown systems, SDS-1 was seen to be effective at providing sufficient emergency shutdown capacity, with the failure of one and two shutdown rods

encountered at a probability of 1.51% and 0.0135% respectively. Disruptions from unplanned shutdown or forced shutdown, respectively being the premature action on the part of some component actioning undesired shutdown, and the emergence of a fault in the system that mandates immediate shutdown, are predicted to occur with respective probabilities of 3.23% and 0.001% within the operating period. A separate set of 10^5 simulations was conducted subsequently, to collect more data from SDS-2 under the condition of SDS-1 having failed, with these results presented in Table 7, and showing the probability of SDS-2 completing safely to be slightly less than 1000 times greater than its failure.

In the case of the control and monitoring systems, problems resulting from the reactor entering a supercritical or subcritical state are predicted to emerge in 0.0018% of cases, will all other systems expected to end with a detection of some variant, including spurious detections and contradictory detections. As seen in Table 8, although almost a fifth of the readings delivered by all sensors were predicted to be false, the 2 of 3 voting system means that of these, it is predicted that approximately three quarters would not coincide with another false reading, such as to cause the three sensors of the block to vote for a spurious result. In Fig. 16b, the predicted preparedness of sensors to make a reading is seen to periodically fluctuate with respect to time, mostly due to the failure of individual sensors between maintenance actions. On the basis of these probabilities, if an average repair time of two weeks could be achieved, the unavailability of the control and monitoring system would be 0.256%, and for a 24-h repair time, this would be 0.247%.

7. Sensitivity analysis

Computational safety analysis of a system has the advantage of enabling convenient exploration of *what-if?* scenarios. It is possible to reveal the sensitivity to alterations in parameters – for example, to consider the effects of replacing a component with an alternative with a different failure profile, or of changing the regime of maintenance actions. Using the primary coolant circulation model seen in Section 4.1 to demonstrate this, two such examples have been selected, namely the feed pump servicing period and the mean time to failure of the steam separator pressure release valve (chosen due to being the most common failure mode in simulations resulting in a coolant fault). Variants of model where the former is altered by $\pm$50% and the latter by $\pm$10% and $\pm$50% have been created, with simulation batch sizes of 10^5, the results of which are found in Table 9. These simulations were identical to the equivalent batch presented in Section 5 in all other respects.

As seen in Fig. 17, the variance in pump servicing period was slight in its impact, and more simulations would be required to confidently assert statistic significance. On the other hand, altering the mean time to failure of the steam separator pressure release valve was more substantial in its effect, particularly when reduced by 50%. These results would seem to imply that sourcing more reliable valves would be a more effective strategy to improve the overall system performance than alteration of the maintenance schedule.

8. Conclusions

To explore the application of the Petri Nets methodology to nuclear power generation, five models were constructed to depict various subsystems within a hypothetical power plant with typical modern design features, such as an emphasis on passive safety. With the parametrisation of the Petri Nets given, simulations predicted a high proportion of final states considered safe or otherwise preferable to more dangerous alternatives (for example, early shutdown of the reactor instead of a critical fault). With the two systems relating to everyday processes of the reactor, namely primary coolant fault circulation and the control and monitoring systems (respectively seen in Figs. 7 and 11), their anticipated failure rates can be quantified in terms of the occurrence of dangerous end states in proportion to the total time simulated across all their iterations. Thus, a coolant fault initiating event at a rate of $(1.65\pm0.09)\times10^{-7}$ per operational hour, and the emergence of an undetected reactivity event at a rate of $(2.3\pm1.3)\times10^{-9}$ per operational hour are predicted. Another notable general feature of the systems is a high probability of early failure, before the time limit for routine shutdown. If the period of scheduled whole-system maintenance was reduced from three years, the predicted occurrence of early shutdown could be reduced somewhat.

Potential avenues for future research include using these models as a starting point for building a single Petri Net for the whole system, and looking at expansions and augmentations to the methodology. For example, coupling a Petri Net with a physical model of the heat transfer processes resulting from a crisis. As it was also demonstrated in this work how analysis of the sensitivity of a system to variance in the value of its parameters can be performed, this could be extended to attempting to improve the predicted system reliability through the application of an optimisation routine to free parameters such as inspection and maintenance schedules or the selection and level of redundancy of key components.

CRediT authorship contribution statement

Mark James Wootton: Conceptualization, Methodology, Software, Writing - original draft, Visualization. **John D. Andrews:** Conceptualization, Methodology, Writing - review & editing, Funding acquisition. **Adam L. Lloyd:** Conceptualization. **Roger Smith:** Conceptualization, Supervision. **A. John Arul:** Conceptualization, Resources, Supervision. **Gopika Vinod:** Conceptualization, Resources, Supervision. **M. Hari Prasad:** Conceptualization, Resources. **Vipul Garg:** Conceptualization, Resources.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council as part of the DaMSSLE project [Grant No. EP/M018210/1].

References

Aldemir, Tunc, 2013. A survey of dynamic methodologies for probabilistic safety assessment of nuclear power plants. Annals of Nuclear Energy 52, 113–124.

Andrews, John D., Dunnett, Sarah J., 2000. Event-Tree Analysis Using Binary Decision Diagrams. IEEE Transactions on Reliability 49 (2), 230–238.

Aubry, Jean-François, Brinzei, Nicolae, Mazouni, Mohammed-Habib, 2016. Systems Dependability Assessment – Benefits of Petri Net Models. Wiley.

Gianfranco Balbo, 2007. Introduction to Generalized Stochastic Petri Nets. In Marco Bernardo and Jane Hillston, editors, Formal Methods for 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, pp. 83–131. Springer.

Barringer & Associates, Inc., Weibull Reliability Database For Failure Data For Various Components. URL:www.barringer1.com/wdbase.htm Accessed November 2018, Last edited: 2010.

Carl Adam Petri, 1962. Kommunikation mit Automaten (In German). PhD thesis, Technical University Darmstadt.

Dennis, Brian, Patil, G.P., 1987. Lognormal Distributions, Theory and Applications. Marcel Dekker New York.
Eide, S.A., Calley, M.B., 1993. Generic component failure database. Proceedings of PSA international topical meeting pp 1175, vol. 2.
Eide, S.A., Chmielewski, S.V., Swantz, T.O., 1990. Generic Component Failure Data Base for Light Water and Liquid Sodium Reactor PRAs. Idaho National Engineering Laboratory.
Expert Opinion.
IAEA, 1988. IAEA-TECDOC-478 Component Reliability Data for Use in Probabilistic Safety Assessment. Technical report. International Atomic Energy Agency.
IAEA. PRIS – Miscellaneous reports – Operational by Age. URL:www.iaea.org/PRIS/WorldStatistics/OperationalByAge.aspx Accessed Februrary 2019, Last editted: 2019.
Jiang, R., Murthy, D.N.P., 2011. A study of Weibull shape parameter: Properties and significance. Reliability Engineering and System Safety 96, 1619–1626.
Kachur, S.A., Shakhova, N.V., 2016. Turbine generator status diagnostic system based on petri nets. Nuclear Energy and Technology 2 (2), 81–84.
Pramod Kumar, Lalit Kumar Singh, Chiranjeev Kumar, 2019. Performance evaluation of safety-critical systems of nuclear power plant systems. Nuclear Engineering and Technology.
Seung Jun Lee, Poong Hyun Seong. 2004. Development of automated operating procedure system using fuzzy colored petri nets for nuclear power plants. Annals of Nuclear Energy, 31(8):849–869.
Metropolis, Nicholas, Ulam, Stanislaw, 1949. The Monte Carlo Method. Journal of the American Statistical Association 44 (247), 335–341.
Miller, C.F., Trojovsky, M., Hubble, W.H., Brown, S.R., 1982. Data summaries of Licensee Event Reports of valves at US commercial nuclear power plants, January 1, 1976 to December 31, 1980. Technical report, EG and G Idaho Inc, Idaho Falls (USA) – Prepared for U.S. Nuclear Regulatory Commission, 1982.
Morris Seymour, 2019. Failure Rate Estimates for Mechanical Components. URL: https://reliabilityanalyticstoolkit.appspot.com/mechanical_reliability_data Accessed: January 2020, Last Updated 2019.
Németh, E., Bartha, T., Fazekas, Cs., Hangos, K.M., 2009. Verification of a primary-to-secondary leaking safety procedure in a nuclear power plant using coloured Petri nets. Reliability Engineering and System Safety, 94(5):942–953.
Papoulis, Athanasios, Unnikrishna Pillai, S., 2002. Probability, Random Variables and Stochastic Processes. McGraw Hill 4^{th}, edition.
Ponciroli, Roberto, Cammi, Antonio, Lorenzi, Stefano, Luzzi, Lelio, 2016. Petri-net based modelling approach for ALFRED reactor operation and control system design. Progress in Nuclear Energy 87, 54–66.
Rasmussen, N.C., 1975. Reactor safety study: An assessment of accident risks in US commercial nuclear power plants (WASH-1400). US Nuclear Regulatory Commission.
Rauzy, Antoine, 1993. New algorithms for fault trees analysis. Reliability Engineering & System Safety 40 (3), 203–211.
Reay, Karen A., Andrews, John D., 2002. A fault tree analysis strategy using binary decision diagrams. Reliability Engineering & System Safety 78 (1), 45–56.
Reliability Eta Beta database. URL:www.reliabilityetabeta.com Accessed January 2020, 2020.
Ruijters, Enno, Stoelinga, Mariëlle, 2015. Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools. Computer Science Review 15–16, 29–62.
Singh, Lalit, Rajput, Hitesh, 2016. Safety Analysis of Life Critical Software Systems: a Case Study of Nuclear Power Plant. IETE Technical Review 34 (3), 333–339.
Singh, Lalit Kumar, Vinod, Gopika, Tripathi, A.K., 2016. Early Prediction of Software Reliability: A Case Study with a Nuclear Power Plant System. Computer 49 (1), 52–58.
Smith David J., 1981. Reliability, Maintainability and Risk. Practical Methods for Engineers including Reliability Centred Maintenance and Safety-Related Systems. Elsevier Butterworth-Heinemann, 2005, 7th ed., 1981.
Vesely, W.E., 1970. A Time-Dependent Methodology for Fault Tree Evaluation. Nuclear Design and Engineering 13, 337–360.
Watson, H.A. et al., 1961. Launch control safety study. Bell Labs.
Weisstein Eric W., 2019. Normal Distribution. From MathWorld–A Wolfram Web Resource. URL:mathworld.wolfram.com/NormalDistribution.html Accessed March 2019, Last edited: 2019.
Mark James Wootton, John Andrews, Adam L. Lloyd, Roger Smith, A. John Arul, Gopika Vinod, Shri Hari Prasad, Vipul Garg, 2019. Petri Nets and Pseudo-Bond Graphs for a Nuclear Reactor Primary Coolant System. Proceedings of the 29th European Safety and Reliability Conference, pp. 3831–3839.