



Politecnico  
di Torino

ScuDo  
Scuola di Dottorato - Doctoral School  
WHAT YOU ARE, TAKES YOU FAR



Doctoral Dissertation

Doctoral Program in Energetics (35<sup>th</sup> cycle)

# **ARC-class tokamaks: Preliminary design, system modeling and safety assessment**

## **Primary cooling system, fuel cycle and machine learning for safety applications**

By

**Samuele Meschini**

\*\*\*\*\*

**Supervisor(s):**

Prof. Massimo Zucchetti, Supervisor

Dr. Raffaella Testoni, Co-Supervisor

Eliana De Marchi, Co-Supervisor

**Doctoral Examination Committee:**

Dr. Marco Riva, Referee, Murmuration Research Institute

Dr. Luigi Di Pace, Referee, RINA Consulting CSM

Dr. Sara E. Ferry, Plasma Science and Fusion Center, MIT

Prof. Fabio Giannetti, Sapienza University of Rome

Prof. Francesco Laviano, Polytechnic University of Turin

Politecnico di Torino

2023

## **Declaration**

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Samuele Meschini

2023

\* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

*To my family*

## Acknowledgements

I would like to acknowledge and express my gratitude to the people who played a fundamental role during these three years of doctorate.

First, a special thank to Dr. Raffaella Testoni, who managed to coordinate our efforts during the Covid pandemic, and has always found time for fruitful discussions and interesting conversations. I am also grateful for the opportunity to take part in satellite research activities that greatly improved my understanding of safety issues in tokamaks.

I would like to thank my supervisor, Professor Massimo Zucchetti, for giving me the opportunity to carry out my research in a dynamic team, and for his mentoring during the PhD.

I would also like to thank Professor Dennis Whyte, Dr. Sara Ferry, and Dr. Kevin Woller for their supervision during my visiting period at MIT PSFC. Those five months were fundamental to my professional growth.

My doctoral scholarship has been funded by Eni S.p.A. I would like to thank all the colleagues from Eni who provided feedback to my work.

To all my friends and colleagues who took part in this journey, thank you. Thanks to Andrea, for improving my everyday logistical skills, and to Roberta, for sharing pleasant lunches during the lockdown. Thank you to Stefano, Davide, and Nicolò for the coffee breaks at PSFC and for our endless brainstorming on a variety of scientific topics. Thanks also to Nicolò, Camilla, Alessandro and Sofia for introducing me to the extraordinary world of medicine and biology, and for the beautiful time we had in Boston. A special thank to my friends from the Grugliasco fencing team, with whom I shared my passion for sport in the last nine years. Thanks to Alberto for keeping it challenging, and to Giulio, without whom it would have been impossible to continue practicing fencing. Thanks to Gabriele, Simone, Giuseppe Francesco, Raffaella, and

Antonio for making the work environment more enjoyable with scientific outreach activities and sport initiatives. Thanks to all my friends from Narni (there are too many to mention, and I am glad there are so many) who have always found time for a reunion in my beloved hometown.

I would like to give special thanks to my parents, Laura and Francesco, who taught me much more than they think, and to my brother Marco. It has always been a pleasure to come back home, recharging me even during the toughest periods of my PhD.

Lastly, I would like to express my gratitude to Sara for her unconditional love and patience. Her dedication and enthusiastic approach have been a constant source of inspiration for me during these years.

## Abstract

Fusion power plants will play a crucial role in sustainable energy production if their technological maturity and economic feasibility are demonstrated in the next decades. In this regard, ensuring safety during nominal operations, and expected or unexpected transients is of paramount importance for the development (and deployment) of fusion reactors. Lack of operational experience and data, absence of detailed design, and a safety framework still under development make preliminary safety analysis of fusion reactors extremely challenging. Nevertheless, preliminary safety analysis can be carried out alongside the design process to produce a robust, safety-oriented design. This work focuses on the methodological and modeling part of a preliminary safety analysis for fusion reactors. Modeling is a persistent need throughout the whole safety analysis, since most of the safety-critical systems in fusion reactors has not been previously modeled. The methodological part deals with strengthening the deterministic approach by integration with a probabilistic framework and the investigation of machine learning (ML) algorithms for nuclear safety. The ARC (Affordable, Robust, Compact) tokamak is exploited as case study. ARC is an innovative, compact tokamak design that features demountable, high-temperature superconducting magnets, a replaceable vacuum vessel, and a molten salt liquid immersion blanket (LIB). The work presents a preliminary design, modeling, and safety analysis of two critical systems, the fuel cycle and the primary cooling loop. Components of the fuel cycle contain the largest radiological source terms in the plant, i.e. tritium inventories that span from grams to kilograms. The time dependent tritium inventories are quantified by a dynamic, system-level model of the fuel cycle implemented in MATLAB Simulink®. The fuel cycle analysis targets minimization of tritium inventories and reactor self-sufficiency simultaneously. Results show that the start-up inventory can be lower than 1 kg and the required TBR to ensure self-sufficiency within the achievable range ( $TBR = 1.05 - 1.10$ ). Accidents occurring in the primary cooling system (i.e., the FLiBe loop) are a major pathway for the release

of tritium. The FLiBe loop is indeed part of the fuel cycle itself because FLiBe works as coolant, breeder and tritium carrier in ARC. A considerable amount of tritium is therefore found in the FLiBe loop and related components. A preliminary design of the primary cooling system is proposed by means of a functional analysis, and the most likely accidental scenarios are identified by an FMEA and simulated by a thermal-hydraulics model of the primary cooling loop. The model is developed in OpenModelica and exploits OMPython to be effectively coupled with machine learning algorithms. The safety analysis is carried out by integrating deterministic analysis in a probabilistic framework, which allows to deeply investigate the space of possible accidents that might occur in a fusion reactor. The results from the thermal-hydraulic simulations are used to build a dataset of safe and accidental transients. Lastly, ML algorithms are applied to the safety analysis. Supervised and un-supervised models are trained on the dataset built by the thermal-hydraulic model, and used to perform tasks of increasing complexity (clustering, identification of the end states, prediction of thermal-hydraulic variables and remaining time before system failure). Results show that ML models can be effectively used for safety purposes, they can execute tasks in times that are compatible with on-line monitoring, and with relative errors on the target variables as low as  $10^{-5}$ . Nevertheless, many criticalities are identified during the process (from the synthetic data generation to the interpretation of the results). A framework for the development and strengthening of ML models for nuclear safety is therefore proposed. The outcomes of the work can be summarized as follows. (1) An ARC-class tokamak minimizes the radiological hazard related to tritium inventories thanks to its compact size, and an optimal operational region has been identified by the fuel cycle model developed in this work. (2) The functional analysis and the FMEA allowed to draw a conceptual layout of the primary cooling system, and the thermal-hydraulic model of the primary cooling system can be exploited to build a database of safe and accidental transients. (3) Machine learning algorithms (especially deep neural networks) are effective tools to strengthen the safety analysis and to support operations and decision making. To further improve their applicability to nuclear safety analysis a framework to compensate the lack of data and the development of numerical simulation tools compatible with ML integrations is proposed.

# Contents

<b>List of Figures</b>	<b>xii</b>
<b>List of Tables</b>	<b>xxiii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivations . . . . .	1
1.2 Nuclear Safety . . . . .	2
1.2.1 Safety Methodologies . . . . .	3
1.2.2 Nuclear hazards . . . . .	4
1.3 ARC-class reactors . . . . .	5
<b>2 ARC primary cooling system</b>	<b>8</b>
2.1 Review of relevant experiments and operations . . . . .	9
2.1.1 Molten salts . . . . .	9
2.1.2 Liquid metals . . . . .	10
2.1.3 Design and safety consideration for ARC PCS . . . . .	11
2.2 Functional analysis of the PCS . . . . .	12
2.3 Failure modes and effects analysis of the PCS . . . . .	13
2.4 Development of a thermal-hydraulics model of the PCS . . . . .	20
2.4.1 PCS layout . . . . .	22
2.4.2 Model description . . . . .	25

---

2.4.3	Verification and benchmark . . . . .	34
2.5	Integration of the thermal-hydraulics model within a probabilistic framework . . . . .	35
2.5.1	Development of the failure injection engine . . . . .	38
2.5.2	Python wrapper and preprocessing . . . . .	40
2.5.3	Post-processing . . . . .	41
2.6	Results from the integrated deterministic and probabilistic model . .	42
2.7	Discussion . . . . .	46
<b>3</b>	<b>Machine learning for nuclear safety applications</b>	<b>51</b>
3.1	Overview of machine learning models . . . . .	52
3.1.1	Supervised algorithms . . . . .	53
3.1.2	Unsupervised algorithms . . . . .	54
3.1.3	Reinforcement learning . . . . .	55
3.2	Criticality and challenges in ML for nuclear safety applications . . .	55
3.2.1	Data availability . . . . .	55
3.2.2	Domain shift . . . . .	56
3.2.3	False positives and false negatives . . . . .	56
3.2.4	Evaluation metrics . . . . .	57
3.3	A machine learning pipeline for nuclear safety applications . . . . .	60
3.3.1	Solve data availability and domain coverage issues . . . . .	61
3.3.2	A common framework to build database for nuclear safety applications . . . . .	62
3.4	Applications of machine learning algorithms to ARC PCS . . . . .	63
3.4.1	Prediction of the end states of the PCS . . . . .	64
3.4.2	Prediction of safety variable evolution during PCS transients	70
3.5	Discussion . . . . .	86

<b>4</b>	<b>ARC fuel cycle</b>	<b>93</b>
4.1	Tritium self-sufficiency, tritium safety and fuel cycle . . . . .	93
4.2	Methodology . . . . .	95
4.2.1	ARC fuel cycle layout . . . . .	95
4.2.2	Fuel Cycle Model . . . . .	99
4.3	Tritium Self-Sufficiency in an ARC-class fusion power plant . . . . .	108
4.3.1	ARC-class tokamak model parameters . . . . .	109
4.3.2	Sensitivity analysis . . . . .	112
4.3.3	Tritium burn efficiency . . . . .	115
4.3.4	Outer fuel cycle residence time . . . . .	115
4.3.5	Tritium extraction efficiency . . . . .	119
4.3.6	Direct internal recycling fraction . . . . .	121
4.3.7	Availability factor . . . . .	123
4.3.8	Doubling time . . . . .	126
4.3.9	Reserve time . . . . .	128
4.3.10	Non-radioactive loss fraction . . . . .	129
4.3.11	Implications for ARC-class FPP design and operations . . . . .	131
4.4	Discussion . . . . .	132
4.4.1	Comparison of D-T FPP designs . . . . .	133
4.4.2	Limitations to the maximum TBE in a D-T FPP . . . . .	138
4.4.3	Potential effect of T trapping in solid components . . . . .	140
4.4.4	Reducing uncertainty in fuel cycle modeling and design . . . . .	140
4.4.5	Doubling time, construction time, and FPP penetration into the electricity market . . . . .	141
<b>5</b>	<b>Conclusions and future work</b>	<b>143</b>
5.1	Summary . . . . .	143

---

5.2	Recommended future work . . . . .	148
5.2.1	Primary cooling system and thermal-hydraulics . . . . .	148
5.2.2	Machine learning and nuclear safety . . . . .	149
5.2.3	Fuel cycle and tritium transport . . . . .	149
	<b>References</b>	<b>151</b>
	<b>Appendix A Summary of the FMEA of ARC PCS</b>	<b>168</b>
	<b>Appendix B Machine learning for safety applications</b>	<b>178</b>
	<b>Appendix C ARC fuel cycle model</b>	<b>187</b>
	<b>Appendix D List of Acronyms</b>	<b>196</b>

# List of Figures

2.1	Primary cooling system model. The main components depicted are the VV section, the molten salt pumps (only one shown), the blanket tank, the heat exchangers (only one shown). An enthalpy sensor at the heat exchanger outlet is connected to the electric heaters controller in the VV section. The VV section and the blanket tank are thermally connected to take into account conduction from the outer VV structural material to tank wall. A ramp is used to simulate pump failures. Internal power generation from neutron heating is accounted in the tank. A constant mass flow rate component sets He flow in the heat exchanger. The sink pressure components provide a reference pressure for both FLiBe and He flows. . . . .	23
2.2	Layout of ARC cooling system. Flow directions in VV channels and LIB tank are depicted. FLiBe enters the VV channels from external pipes that pass through the LIB tank. The VV channels end in the tank itself. Reproduced from [1]. . . . .	24
2.3	Poloidal cross section of an ARC-class tokamak. The main layers are depicted. A thin tungsten FW protects the VV structural material. The VV is placed inside the LIB tank. An outer shield surround the outer tank walls. Reproduced from [2] . . . . .	25
2.4	Temperature of VV and blanket materials as a function of the convective heat transfer coefficient ( $\gamma$ ) between tank wall and FLiBe inside the tank. First wall, inner structural material and beryllium are almost unaffected by variations of $\gamma$ . . . . .	29

---

2.5	Simplified representation of a VV channel. A constant heat flux is applied to the FW. Internal power generation due to neutron interactions is considered in all the materials. Electric heaters can provide additional power on demand (note that the heaters will be likely placed upstream of the VV channel, before the pipes enter the LIB tank). Purple: FW; Grey: Structural material; Blue: FLiBe; Green: Beryllium. . . . .	30
2.6	VV sector Modelica model. Starting from the top, a constant heat flux is applied to the first wall. The first wall is connected to the inner structural material through a heat transfer port. The VV channel exchanges heat with both the inner structural material and the beryllium layer. Two flow ports connect the channel input and output with the loop components. Electric heaters are connected to the channel to provide additional power on demand. The beryllium layer is connected to the outer structural material that exchanges heat with the FLiBe tank. Internal power generation is considered in all the components shown in the figure. Note that both red and orange lines refer to heat transfer connections. . . . .	31
2.7	Heat exchanger model, including the control logic and the additional components required to trigger the PI controller failure. The primary coolant components are in blue, secondary coolant components are in gray. . . . .	33
2.8	Comparison of FLiBe temperature at channel outlet computed analytically and by the VV sector model. . . . .	36
2.9	Comparison between the blanket inlet and outlet FLiBe temperature computed by Kuang et al. [3] and by the present model. A short transient occurs at the beginning of the simulation before the solution converges at $T_{in} = 800$ K and $T_{out} = 900$ K. . . . .	37
2.10	Python wrapper developed to handle preprocessing, post-processing and simulation run. . . . .	40

- 2.11 PCS transient involving pump failure at  $t = 150$ s and HX PI controller failure at  $t = 200$ s. Relevant system variables are reported. Top left: FLiBe temperature at blanket inlet and outlet. Setpoint temperature for the heater controller and FLiBe freezing temperature are also shown. Top right: FLiBe and He mass flow rate in the heat exchanger. Bottom left: heaters power delivered to FLiBe. Heaters are triggered when FLiBe temperature drops below the setpoint to avoid FLiBe freezing. Bottom right: FW, VV, and blanket materials temperature. . . . . 44
- 2.12 FLiBe loop transient involving pump failure at  $t = 150$ s. The HX PI controller does not fail in this transient. Relevant system variables are reported. Top left: FLiBe temperature at blanket inlet and outlet. Setpoint temperature for the heater controller and FLiBe freezing temperature are also shown. Top right: FLiBe and He mass flow rate in the heat exchanger. Bottom left: heaters power delivered to FLiBe. Heaters are triggered when FLiBe enthalpy drops below the setpoint to avoid FLiBe freezing. Bottom right: VV and blanket materials temperature. . . . . 45
- 2.13 Silhouette coefficient for different number of clusters. The silhouette coefficient is a cluster validity measure that allows to identify the optimal number of clusters to categorize the data. In that case a minimum is found for 10 clusters. . . . . 47
- 2.14 Clusters for the transients represented on a suitable space. The failure times have been normalized before the clustering procedure. Cluster 2, 3 and 6 are located on the right area of the plot. Cluster 0 contains less than 10 elements located on the top of the plot. . . . . 48
- 2.15 Number of transients for each end state and cluster. Note that the total number of transients belonging to a cluster is different among the clusters. Cluster 5 is lowly populated, whereas cluster 1 shows the largest number of objects. Cluster 0, 2, 3, and 6 contain transients ending in a safe state only. . . . . 49

- 
- 3.1 Receiver Operating Characteristic curve. The red dotted line is the ROC curve of a random classifier (the probability of getting a TP or a FP is the same). The ROC curve of a good classifier (Classifier A) and a worse classifier (Classifier B) are depicted. A perfect classifier achieves a unitary TP rate for an FPR = 0 (green dot). . . . . 59
- 3.2 Comparison between ROC curve and Precision-Recall curve for a strongly imbalanced dataset (99% of the samples belonging to the same class). While the ROC curve shows a good behaviour of the classifier, the Precision-Recall curve highlights the extremely low precision of the model when REC increases. . . . . 60
- 3.3 Flowchart showing the collaborative research approach. Each user contributes to a shared database, which solves both the lack of data and the domain coverage issue. Users can also use different codes, provided they output coherent results with the other codes. . . . . 62
- 3.4 A simple workflow to integrate ML algorithms in nuclear safety analysis. Coupling a probabilistic routine with a deterministic code increases the domain coverage. Once the database is built, a ML pipeline exploits the huge amount of data to train ML models. . . . 63
- 3.5 ROC for two logistic regression algorithms using different scoring metrics. Results refer to the test set. The blue curve shows the ROC of the Logit trained to maximize the accuracy, whereas the orange curve shows the Logit trained to maximize the AUC. Maximizing the AUC leads to an algorithm that identifies better TP (incipient accidents), which is the goal of such an algorithm applied to prediction of end states. Each point on the curves are associated to a classification threshold. Two thresholds (one for low TP rate and one for high TP rate) are depicted. . . . . 65
- 3.6 Confusion matrix of Logit 1 (a) and Logit 2 (b) for the test set. Values on the matrix diagonal represent TN and TP. The lower left corner shows the FN, while the upper right corner FP. Despite the lower accuracy, Logit 2 achieves appreciably higher performances in executing the task (i.e., identify TP and minimize FN). . . . . 66

- 
- 3.7 Safe end states and fault end states as a function of the normalized failure times of the pump, PI HX controller and heater controller. The (linear) decision boundary of the logistic regression model is depicted (light blue plane). A good separation of the two classes is achieved, even though a linear boundary has limited capabilities. . . . 66
- 3.8 Safe end states (a) and fault end states (b) represented separately to highlight the separating capability of the linear decision boundary (light blue plane). Most of the safe end states fall on one side of the boundary, but many fault end states lay on the same side of the boundary, limiting the performance of the model. . . . . 67
- 3.9 Comparison between logistic regression models and MLP for both training and test set, using the ROC AUC as scoring metric. The MLP achieves  $AUC = 1$  (i.e., perfect classifier) both on the train and the test set (train set not shown for simplicity). . . . . 68
- 3.10 Comparison between logistic regression models and MLP for both training and test set, using REC as scoring metric. The MLP strongly outperforms the logistic regression model, achieving a high precision even when  $REC \rightarrow 1$ . A random classifier would have achieved  $PREC = 0.21$  with this dataset. . . . . 69
- 3.11 Confusion matrix for the MLP for the test set. The MLP achieves almost a perfect classification accuracy, with a just a FN over the test set. . . . . 69
- 3.12 Pandas dataframe of a random sample from the dataset showing the input variables (the input set has been already reduced from the initial set reported in Table (3.1)). Each row is a snapshot of the system at time  $t$ . . . . . 72

3.13	Input set provided to the LSTM and output (target variables) returned by the model. The input is passed at each timestep and the target variables are computed at the same timestep. $\Delta p$ : pressure increase across the pump; $T_{in,sec}$ : inlet temperature (at the HX) of the secondary coolant; $\frac{dT_{in,sec}}{dt}$ : time derivative of the inlet temperature of the secondary coolant; $\frac{dm_{FLiBe}}{dt}$ : time derivative of the FLiBe mass flow rate; $\dot{Q}_{FW}$ : thermal power to the FW; $\dot{Q}_{heaters}$ : thermal power delivered by the electric heaters; $\frac{dm_{sec}}{dt}$ : time derivative of the secondary coolant mass flow rate; $\dot{m}_{sec}$ : secondary coolant mass flow rate; $\dot{m}_{FLiBe}$ : FLiBe mass flow rate; $h_{out,bl}$ : FLiBe enthalpy at the blanket outlet. . . . .	73
3.14	Failure time distributions within the whole dataset. The distribution is exponential as explained in Section 2.5.1. . . . .	76
3.15	Input distribution at $t = 100$ s (a) and $t = 1000$ s (b) for fixed failure magnitudes. As the components fail, the input variables spread over a larger range. Only the most relevant input are shown. . . . .	78
3.16	Prediction of accident evolution (zoom on the initial transient on the left). The failure times and the safety threshold temperature (1250 K) are depicted. Note that the model does not receive as input the failure times, but it manages to successfully predict the strong variation in temperature when a failure takes place. . . . .	79
3.17	Outlet blanket temperature distribution at 5000 s over the dataset. . .	79
3.18	Running loss on the training set as a function of the iterations during the first epoch. The loss quickly decreases down to $\sim 10^{-7}$ . . . . .	80
3.19	Prediction of accident evolution. The failure times and the system failure time are depicted. Note that the model does not receive as input the failure times, but it manages to successfully predict the TTF. . . . .	81
3.20	Prediction of the temperature evolution for an outlier scenario. The model fails to correctly predict the temperature evolution because it has been trained on a very different dataset. . . . .	82

3.21	Input distribution at $t = 100$ s (a) and $t = 1000$ s (b) for randomly sampled failure magnitudes. As the components fail, the input variables spread over a larger range. Only the most relevant input are shown. . . . .	88
3.22	Outlet blanket temperature distribution at 5000 s over the dataset that includes the failure magnitudes in the input set. The temperature spans over the whole possible range for FLiBe, from the melting temperature ( $T_{melt,FLiBe} = 732$ K) to the boiling temperature ( $T_{boil,FLiBe} = 1703$ K). . . . .	89
3.23	Transient ending in a safe state. The temperature increases smoothly up to a steady-state value. The predicted value slightly deviates from the real value immediately after each failure, but the model manages to improve the prediction soon after. . . . .	89
3.24	Transient ending in a cold failure. The temperature first increases, and then starts decreasing after the HX controller failure. The predicted value slightly deviates from the real value immediately after each failure, but the model manages to improve the prediction soon after. . . . .	90
3.25	Transient ending in a hot failure. The temperature first increases slowly, and starts increasing sharply after the HX controller failure. . . . .	90
3.26	Transient ending in a safe state. The TTF is $10^4$ s by definition. The predicted TTF oscillates during the initial transient due to the presence of three component failures. However, once the last failure occurs, the model correctly predict a value close to $10^4$ s, meaning that the system is not expected to fail. . . . .	91
3.27	Transient ending in a cold failure. The model predicts that a system failure will happen after the third component failure, but the TTF prediction is largely overestimated. . . . .	91
3.28	Transient ending in a hot failure. The TTF oscillates during the initial transient. Nevertheless, the model correctly predict the TTF after the third failure occurs, with a small underestimation. . . . .	92

- 4.1 The block diagram shows the high-level fuel cycle of an ARC-class FPP. The outer fuel cycle (OFC) breeds tritium via  $\text{Li}(n,T)$  reactions in the blanket and coolant channels. This tritium is passed to the inner fuel cycle (IFC), where it is processed into fuel. Only the components that have been modeled in this analysis are shown in the figure. The DIR system is represented by a dotted path to indicate that it may or may not be present in the IFC. The dotted path from the heat exchanger to the detritiation system is present if the chosen heat exchanger has an active tritium removal system (e.g. a double-walled heat exchanger with purge gas flow). The inclusion of DIR and the decision to include T extraction in the heat exchanger (HX) are examples of ARC design decisions which have yet to be finalized. 96
- 4.2 Evolution of the storage inventory during plant operations. (a): Ideal scenario ( $AF = 1$ ) with no reserve inventory. The minimum storage inventory is zero at the minimum point. (b): Realistic scenario ( $AF < 1$ ) with a reserve inventory of 1 kg. The minimum storage inventory equals the reserve inventory at the minimum point.  $I_{\text{storage}}$  decreases from  $t = 0$  to  $t = t_{\text{min}}$  because the bred tritium is not immediately available as fuel. The minimum point is reached when the tritium bred at  $t = 0$  has passed through the OFC, the ISS, and reached the storage system. For  $t > t_{\text{min}}$  the storage inventory increases linearly due to the constant inflow of bred tritium (assuming no interruption to plant operation). The values reported in the plots do not refer to any particular design, and were chosen to emphasize the contribution of the reserve inventory to the storage inventory. . . . . 104
- 4.3 A basic fuel cycle layout. All IFC components except the storage system are collapsed in a single block with residence time  $\tau_{\text{IFC}}$ . Similarly, all the components in the OFC are collapsed in a single block with residence time  $\tau_{\text{OFC}}$ . The tritium flows exchanged by the systems are depicted by the arrows. . . . . 105

- 4.4 (a) Normalized TBR and relative sensitivity index coefficient (RIC) show the sensitivity of  $TBR_r$  to the eight parameters listed in Table 4.3. (b) Normalized  $I_{\text{startup}}$  and RIC show the sensitivity of  $I_{\text{startup}}$  to the same eight parameters. To normalize  $TBR_r$  and  $I_{\text{startup}}$ , all design parameters were set to the mean value reported in Table 4.3. Each design parameter has been varied in  $[\mu - \sigma, \mu + \sigma]$  range, where  $\sigma$  is the range reported in Table 4.3, choosing 7 equally-spaced points. Both  $TBR_r$  and  $I_{\text{startup}}$  sharply increase as TBE decreases. TBE shows the second largest (by absolute value) RIC for  $TBR_r$  and the largest RIC for  $I_{\text{startup}}$ . With the exception of AF, the RIC for each design parameter has the same sign for  $TBR_r$  and  $I_{\text{startup}}$ , meaning that the minimization of  $TBR_r$  and  $I_{\text{startup}}$  can be achieved without tradeoffs. AF considerations are explained in detail in Section 4.3.7. . . . . 113
- 4.5  $TBR_r$  and  $I_{\text{startup}}$  as a function of TBE, for different processing times and doubling times. For  $t_d = 2$  y (left column),  $t_p$  has a small impact on  $TBR_r$  when TBE is high ( $> 4\%$ ). For low TBE, tritium self-sufficiency becomes difficult to practically achieve unless  $t_p$  is quite low. For the longer  $t_d$  of 5 y (right column), engineering and design constraints on the blanket are relaxed: tritium production rates can be lower and longer  $t_p$  is acceptable. The plots also show that for very low TBE  $< 0.5\%$ ,  $I_{\text{startup}}$  and  $TBR_r$  sharply increase to impractical values. . . . . 116
- 4.6  $TBR_r$  and  $I_{\text{startup}}$  for different residence times in the breeding zone.  $\tau_{\text{BZ}}$  has effectively no impact on  $TBR_r$  and  $I_{\text{startup}}$  for  $\tau_{\text{BZ}} < 1$  day. At longer  $\tau_{\text{BZ}}$ ,  $TBR_r$  and  $I_{\text{startup}}$  increase slightly. . . . . 117
- 4.7 Evolution of relevant system inventories for (a)  $\tau_{\text{BZ}} = 1.25$  h and (b)  $\tau_{\text{BZ}} = 10$  d. The slower dynamics and the shift of the minimum time for  $\tau_{\text{BZ}} = 10$  d is depicted. . . . . 118

- 4.8 Steady-state tritium inventory in the breeding zone for different residence times. Longer  $\tau_{BZ}$  is associated with larger T inventories in the blanket, which presents maintenance and regulatory challenges. Higher values of TBE are associated with smaller tritium inventories in the breeding zone, as the fuel is being used more efficiently and  $TBR_r$  is lower. . . . . 118
- 4.9 Steady-state tritium inventory in the TES for different values of  $\eta_{TES}$  and TBE. Higher efficiency in the TES and in the fueling system corresponds to lower TES tritium inventory. . . . . 120
- 4.10  $TBR_r$  and  $I_{startup}$  as a function of the direct internal recycling fraction. DIR is more advantageous when operating at low tritium burn efficiency. . . . . 122
- 4.11  $TBR_r$  as a function of availability factor (AF). The dashed lines identify the AF for two different combinations of ( $t_{pulse}$  and  $t_{off}$ ). . . 124
- 4.12 Evolution of the tritium inventories in the tritium extraction system, isotope separation system, and tritium storage system. These are the most relevant tritium inventories in the plant. At fixed  $TBR_r$ , a FPP with AF = 50% takes almost twice the time to double the initial inventory with respect to a FPP with AF = 100%. . . . . 125
- 4.13  $TBR_r$  for different doubling times  $t_d$ . A short  $t_d$  requires high  $TBR_r$ . For  $t_d = 1$  y, tritium self-sufficiency is highly unlikely if TBE < 1%. If a longer  $t_d$  is acceptable, tritium self-sufficiency is much easier to achieve. For  $t_d = 5$  y, tritium self-sufficiency is possible even if TBE = 0.5%. The bottom plot zooms in on  $TBR_r$  in the region of interest ( $TBR_r < 1.3$ ). . . . . 127
- 4.14  $TBR_r$  and  $I_{startup}$  for an ARC-class FPP as a function of tritium reserve time.  $I_{startup}$  increases linearly with  $t_{res}$ . The  $TBR_r$  increases accordingly for fixed  $t_d$  as a higher inventory must now be doubled. 128
- 4.15  $TBR_r$  for different non-radioactive loss fractions  $\epsilon$ . For  $\epsilon > 10^{-3}$  the  $TBR_r$  increases exponentially for any value of TBE. Tritium self-sufficiency is impossible unless TBE > 5%. For  $\epsilon < 10^{-3}$  tritium self-sufficiency becomes achievable, especially for TBE > 2%. For  $\epsilon < 10^{-5}$  no significant changes are observed if TBE > 1%. . . . 130

- 
- 4.16 Relative sensitivity index coefficient (RIC) of  $TBR_r$  (a) and  $I_{\text{startup}}$ (b) of ARC, STEP-EBB and STEP-Li [4]. No significant differences are found between ARC, STEP-EBB and STEP-Li as they have similar fuel cycle designs. . . . . 134
- 4.17  $I_{\text{startup}}$  with and without a reserve inventory at the beginning of the operations (FC scenario) for ARC, STEP-EBB and STEP-Li [4].  $I_{\text{startup}}$  of high-power-output FPPs is significantly increased if the reserve inventory is included at the beginning of the operations. . . . 137
- 4.18  $TBR_r$  and fusion power density ratio as a function of TBE for an ARC-class FPP with baseline parameters ( $\eta_{\text{He}} = 0.5$ ,  $t_d = 2y$ ). The fusion power reduction is approximately 30% at TBE = 10% for  $\Sigma = 1$ , and 10% for  $\Sigma = 3$ . TBE < 10% is therefore necessary to avoid excessive power reduction at  $\Sigma < 3$ , while 10% < TBE < 15% might be acceptable for higher  $\Sigma$ . . . . . 139
- C.1 Detailed layout of ARC IFC considering the DIR loop. Tritium, PEG, helium and impurity flows are shown in different colors. . . . 190
- C.2 ARC fuel cycle layout. Only the components that have been modeled are shown in the figure. The tritium flows between components are depicted. For the definition of each tritium flow see Tab. C.3. . . . . 191

# List of Tables

1.1	Nominal parameters describing an ARC-class fusion power plant [5].	7
2.1	Functional analysis for ARC PCS. The functional requirements, desing parameters and components to fulfil the fuctions are identified.	14
2.2	List of relevant failure modes in ARC PCS. The second column provides a short description of each failure mode. The components that are susceptible to a failure mode are reported in the third column.	16
2.3	Severity criteria in case of failure for economics and plant operations, workers and population, and environment. . . . .	17
2.4	Risk matrix resulting from the FMEA. Green: low risk; Yellow: medium risk; Orange: severe risk; Red: high risk. . . . .	18
2.5	Postulated initiating events and frequency of occurrence according to the results from the FMEA. . . . .	19
2.6	Structural and functional materials of the VV and the LIB. . . . .	26
2.7	Input variable modified by the wrapper. The transients that involve those variables and the related components are reported as well. . .	41
2.8	Safety limits considered during post-processing. For structural and functional materials, the safety limit has been taken equal to the melting point. For the inner and outer structural materials (Inconel 718), the lowest value in the melting temperature range (1644 K - 1700 K) has been used. . . . .	42
3.1	Thermal-hydraulics variables selected as input for the DNN model. .	71

3.2	DNN model data. . . . .	74
4.1	Design parameters for the baseline ARC-class FPP fuel cycle analyzed in Section 4.3. . . . .	109
4.2	Tritium residence time in the fuel cycle components and systems. Nominal value for the baseline case and the range used in the parametric analysis are reported. . . . .	111
4.3	Mean value $\mu$ and range $\sigma$ of the parameters used in the Sec. 4.3.2 sensitivity analysis. Each parameter was varied across the specified range, centered on the mean value, in increments of (range)/6. . . .	112
4.4	Tritium inventories in FC components at steady-state, for the baseline design parameters. The tritium inventory for the separation membrane downstream of the TES is not reported because it is negligible. The storage system inventory does not reach a steady-state, and the start-up inventory is reported as maximum value (assuming that tritium is removed at each $t_d$ ). . . . .	132
4.5	$TBR_r$ and $I_{startup}$ for an ARC-class FPP, assuming baseline design parameters, an ambitious operating scenario ( $t_d=1$ y), and a moderately ambitious operating scenario ( $t_d=2$ y). This allows to consider how parameters change if we can implement significant improvements in fuel cycle (FC) technologies, in plasma operations (PO), or in both. The advanced-FC scenario considers $AF = 90\%$ , $f_{DIR} = 0.7$ , $t_p = 1$ h, and $TBE =$ baseline case (0.02, based on the current state-of-the-art). The advanced-PO scenario considers $TBE = 0.1$ and all other parameters unchanged from the baseline scenario. The “both” scenario assumes that both FC and PO advances have been implemented. . . . .	133
4.6	Comparison of D-T FPP designs for the baseline case and the three possible advanced scenarios. The points refer to the “moderate” scenarios ( $t_d = 2$ y). . . . .	135
4.7	ARC design parameters, required TBR and start-up inventory for a “worst-case” scenario. . . . .	136

---

C.1	Key components in the outer fuel cycle of a DT-burning fusion power plant, along with a non-comprehensive list of examples. . . . .	188
C.2	Key IFC components and a non-exhaustive list of associated technologies. . . . .	189
C.3	Mathematical definition of the terms appearing in ARC fuel cycle model (Fig. C.2). . . . .	194
C.4	Tritium flow rate fractions for the fuel cycle model. The flow rate fractions from the FLiBe tank to the other components have been estimated according to the FLiBe flow rates specified in [3] ( $\sim 66\%$ of coolant flow rate to the main chamber, $\sim 33\%$ to the divertors). An equal distribution from the main chamber to the FW/VV channels and the FLiBe tank was assumed (i.e., $f_{5-3} = f_{5-1}$ ) . . . . .	195

# Chapter 1

## Introduction

### 1.1 Motivations

Nuclear fusion brings an invaluable promise for clean energy production. A growing scientific interest has been observed in the last decades, pushed somehow by the rise of fusion startups. This acceleration unravelled topics that have been poorly investigated in the past.

Since its discovery in the 20's most of the works in nuclear fusion focused on understanding physics fundamentals [6], from the very basic fusion reaction [7, 8] to plasma physics in experimental devices [9–11]. Part of the recent projects in nuclear fusion shifted the goals from academic research to energy production. The improvements required to step out from the experimental field to the industrial scale introduced cross-field challenges that were absent, or minor, in experimental devices, such as components manufacturing, RAMI (Reliability, Availability, Maintainability, Inspectability), economics, safety, and fuel management, to cite few. It is mandatory for scientists, researchers, engineers and, more in general, practitioners in the fusion field to tackle these new challenges. This work focuses on a subset of these issues, identifying and facing the challenges that must be tackled in the next decades for the successful development of fusion devices for energy production.

The remainder of this section is devoted to the introduction of the leading topic of the work, namely nuclear safety, of ARC-class fusion power plants (FPPs), and on the three topical area addressed, i.e., the primary cooling system, the fuel cycle, and machine learning for safety applications. The work was carried out in tight

collaboration with Eni, and the analysis presented in Chapter 4 was developed at MIT PSFC.

Chapter 2 and Chapter 3 are motivated by a methodological research question: is it possible to strengthen safety analysis by applying modern techniques and computational tools? If so, can we use this approach for FPP safety? Results have to be intended as qualitative, yet required to investigate this topic in detail. Specifically, Chapter 2 describes the preliminary design and safety analysis of the molten salt (FLiBe) circuit (i.e., the primary cooling system), a thermal-hydraulics model and the code architecture. Chapter 3 presents a broad discussion on machine learning for nuclear safety applications. Specific ML techniques are applied to the results from Chapter 2, and a framework for a robust establishment of machine learning in nuclear safety is depicted. Chapter 4 is dual with Chapter 2 and focuses on an ARC-class FPP fuel cycle. FLiBe works indeed as primary coolant, tritium breeder and tritium carrier, resulting in an overlapping of the reactor cooling system and part of the fuel cycle. A preliminary layout of the fuel cycle is proposed, and a system-level model is built and implemented in MATLAB Simulink<sup>®</sup> to perform safety and tritium self-sufficiency analyses. Preliminary models for fusion fuel cycle have been already developed and investigated, so Chapter 4 is mostly motivated by the need of quantitative results to assess reactor safety and self-sufficiency. Chapter 5 summarizes all the results obtained in the previous chapters and provide a unified, holistic view on the topics addressed in this work. An R&D pathway is proposed to reduce the uncertainties connected to ARC-class reactor design and safety, to improve tritium transport analysis, and to foster machine learning into nuclear safety.

Each chapter can be considered self-consistent, with an introductory section that presents the state of the art and the main challenges for the specific chapter topic.

## 1.2 Nuclear Safety

Hazards are ubiquitous. Yet, nuclear technologies have been historically associated to extreme hazard. The huge amount of radioactive material in a nuclear (fission) reactor<sup>1</sup> is indeed a serious hazard for the population and the environment. The

---

<sup>1</sup>For sake of simplicity, *nuclear reactor* and *nuclear power plant (NPP)* refer in this work to their conventional meaning, i.e. nuclear *fission* technologies. *Fusion reactors* and *fusion power plants* are used when referring explicitly to fusion technologies.

Chernobyl accident manifested to the population the consequences of a nuclear accident, and the need for strict and third-party safety regulation and control. Besides Chernobyl, which can be considered a failure (from a safety standpoint) to the extent of operators' capability to override safety protocols and safety equipment, nuclear power plants report the lowest safety risk among industrial activities. This achievement has been made possible by detailed safety analysis, conservative design and technological innovation.

### 1.2.1 Safety Methodologies

Historically, safety analyses have been carried out from a deterministic perspective. Deterministic safety analysis (DSA) has evolved along the years from the first applications that considered design basis accidents (DBA) with limited computational (sometimes analytical [12]) resources to modern analyses that investigate beyond design basis accidents (BDBA) with best-estimate codes [13]. The WASH-1400 report challenged the deterministic approach showing that most of the reactor shutdown were driven by minor failures, which were less dangerous but more frequent than DBA [14]. Probabilistic safety analysis (PSA) was therefore developed as complementary tool to strengthen safety analysis [15]. To date, both DSA and PSA are used for nuclear power plants safety analysis, but their application rarely intersect in an integrated analysis.

Integration between the DSA and PSA has been proposed, encouraged and investigated by many institutions [16] and authors [17]. This integration usually leverages (at least conceptually) on artificial intelligence (AI), which is in all respects a probabilistic tool. As for the huge improvements brought by the introduction of advanced numerical codes in DSA, AI can provide key advantages that are not possible with different tools. However, the following points has to be addressed when investigating AI applications in nuclear safety:

- Is the introduction of AI needed?
- Are AI tools (models, algorithms) mature for safety applications?
- How can we practically tackle safety issues by AI?
- Is this approach scalable to the nuclear industry?

The first point has been somehow answered in the previous paragraph. AI tools can indeed improve safety analysis. They can be considered not mandatory, as PSA was not mandatory when developed, but they might bring huge advantages in terms of reduction of conservatism, accident prediction, and handling system complexity. The remaining three questions are addressed in Chapter 3 by focusing on a subset of AI, i.e. machine learning (ML).

### 1.2.2 Nuclear hazards

A second big question arises when dealing with nuclear fusion safety. The previous discussion clearly referred to the nuclear fission sector. The experience, approaches and tools can be transferred to fusion safety to a certain extent based on the obvious link between the two, i.e. they are both nuclear technologies. It is instructive to consider similarities and differences between nuclear fusion hazards and nuclear fission hazards<sup>2</sup>. There are three main radiological sources in a fusion reactor<sup>3</sup>: the plasma, the activated materials, and tritium.

A fusing plasma can be considered a neutron source of 14 MeV neutrons, which are one order of magnitude more energetic than the 2 MeV fission neutrons. This introduces extremely harsh challenges for structural and functional materials that must withstand considerable neutron fluences. However, from a purely radiation safety perspective, which in this case translates to shielding requirements, the hazard is comparable, if not lower, to fission reactors. Magnets lifetime sets indeed strict requirements on the neutron current exiting the shield, in such a way that the neutron fluxes outside the reactor (i.e., outside the cryostat/biological shield) are already quite low. Additional shielding thickness outside the cryostat can be foreseen if necessary, but the primary shield thickness is constrained between the neutron source and the magnets. The primary shield design has been already investigated in a preliminary work by Segantin et al. [2], while more detailed analyses would require coupling shielding calculations with irradiation damage estimations on the magnet coils.

High fluences lead also to strong materials activation, becoming a limiting factor for key fusion reactor components, like vacuum vessels and blankets. This issue has been already deeply investigated for ARC-class reactors by Segantin et

---

<sup>2</sup>Here and in the following, *hazards* refer to nuclear (i.e., radiological) hazards only.

<sup>3</sup>*Fusion reactors* are considered D-T tokamaks in this work if not differently specified.

al. [18], who concluded that Nickel alloys are not suitable for this reactor class, while other commercial alloys (e.g., V-Cr-Ti alloys) can drastically reduce activation. Activated materials can also be found as dust in fusion reactor [19]. Dust production and transport is quite a complex topic that involves plasma edge physics, plasma-material interactions, and transport of macroscopic dust formation during disruptive events. Numerical codes that try to model these phenomena have been developed (e.g. DTOKS [20], DUSTT [21], etc.), but most of the results on this topic come from post-mortem analysis of samples gathered from operating tokamaks. An assessment of radioactive dust in an ARC-class reactor (as in any conceptual design for commercial fusion reactors) is therefore unfeasible at the present time.

The last huge hazard in a fusion reactor is tritium, which has never been addressed before for ARC-class FPPs. Loosely speaking, the amount of tritium circulating in a fusion reactor will be in the order of 100g-10kg, which are amounts much larger than those found in CANDU reactors (i.e., our deepest source of experience for tritium handling). Tritium safety presents many layers of complexity. On a very fundamental level, tritium transport is still scarcely investigated experimentally, and key properties to model and assess tritium transport in a fusion reactor are highly uncertain [22, 23]. From the point of view of a fusion power plant, the amount of tritium in the plant is unknown. The former issue motivates tritium experiments (which are beyond the scope of this work) in fusion-relevant environments, whereas the latter issue is deeply investigated in Chapter 4.

### 1.3 ARC-class reactors

The conceptual design of an ARC reactor was proposed in 2015 by Sorbom et al. [24]. Further analyses targeted specific design areas, such as heat exhaust management [3], neutronics and tritium breeding [25, 26], thermo-fluid dynamics and heat transfer [27], and safety [28]. ARC is a compact D-T tokamak that leverages on demountable high temperature superconductor (HTS) magnets to achieve a high fusion power density (525 MWth in  $\sim 140 \text{ m}^3$ ) with a compact size ( $R=3.3 \text{ m}$ ,  $a = 1.1 \text{ m}$ ). Exothermic reactions in FLiBe (mostly  ${}^6\text{Li}(n,T)\text{He}$ ) and auxiliary heating eventually radiated by the plasma raise the total thermal power up to 630 MWth. The primary cooling system is based on FLiBe flowing through vacuum vessel (VV) channels and divertor channels, with an inlet temperature of  $\sim 800 \text{ K}$  and

an outlet temperature of  $\sim 900$  K [3]. Multiple options for the power cycle have been investigated, from a classical Brayton cycle to more advanced concepts like supercritical CO<sub>2</sub> cycle [24, 29]. The relatively low fusion power minimizes the tritium inventory in the plant. A fusion power of 525 MWth corresponds indeed to a  $\sim 80$  g/day tritium consumption (at full power), which translates to a tritium fueling rate of 46 mg/s for a tritium burn efficiency (TBE) of 2%. The required start-up inventory ranges from  $\sim 100$ g to few kilograms depending on the fuel cycle design and tritium self-sufficiency targets (Chapter 4). The relevant parameters of the conceptual design are summarized in Table 1.1. Three additional features make ARC design innovative: a replaceable vacuum vessel, demountable coils, and a liquid immersion blanket (LIB). The replaceable VV, coupled with demountable coils, aims at achieving quick maintenance. Such a goal is of paramount importance in light of the low component lifetime and long mean time to repair (MTTR) expected for commercial fusion reactors [30], and RAMI has been identified as the Achille's heel of fusion power plants [31]. Although RAMI is outside the scope of this work, the preliminary design of the primary coolant loop and fuel cycle have been developed by prioritizing system RAMI. The LIB brings instead specific advantages for safety. The idea behind the LIB is simple: the VV is surrounded by molten FLiBe contained in a tank, minimizing structural materials. Tritium breeding is therefore optimized by reducing neutron parasitic absorption from structural materials, the liquid blanket cannot be damaged by radiations, the amount of solid waste is minimized, and the large mass of FLiBe provides an effective heat sink in case of severe accidents ( $T_{boil,FLiBe} = 1703$  K).

Specific details on ARC design are reported in the relevant Chapters.

<b>Parameter</b>	<b>Symbol</b>	<b>Value</b>	<b>Units</b>
Fusion power	$P_{\text{fus}}$	525	MW <sub>th</sub>
Toroidal magnetic field	$B_0$	9.2	T
Peak on-coil magnetic field	$B_{\text{max}}$	23	T
Major radius	$R_0$	3.3	m
Plasma semi-minor radius	$a$	1.13	m
Plasma elongation	$k$	1.84	-
Plasma volume	$V_p$	141	m <sup>3</sup>
Average plasma temperature	T	14	keV
Average plasma density	$n$	$1.3 \cdot 10^{20}$	m <sup>-3</sup>

Table 1.1 Nominal parameters describing an ARC-class fusion power plant [5].

# Chapter 2

## ARC primary cooling system

This section investigates ARC primary cooling system (PCS), i.e. the FLiBe loop. The PCS is a critical system during both nominal operations and accidental scenarios. In nominal operations the PCS transfers the fusion power from the blanket to the power cycle. In accidental scenarios the huge mass of FLiBe in the PCS provides a large heat sink even if FLiBe circulation is stopped. Also, failures in the PCS are likely to be initiating events (IE) as in NPP [32].

The goal of this section is to develop a conceptual, safety-oriented design of the PCS, and to build a dynamics thermal-hydraulic model to analyse transients in the PCS. Operational experience with molten salts is quite limited, especially at industrial scale. Therefore, molten salt experiments and molten salt-based design are reviewed first. Liquid metal experiments that can be relevant for molten salt blankets and ARC PCS are included in the review as well (Section 2.1). A functional analysis is carried out to systematically identify the functions that the PCS has to fulfil (Section 2.2). A Failure Mode and Effects Analysis (FMEA) is build on top of the functional breakdown and the components identified by the functional analysis (Section 2.3). A preliminary layout based on the same components is proposed, and a thermal-hydraulics model is developed exploiting the Modelica language [33] (Section 2.4). The model is integrated in a probabilistic framework by taking advantage of OMPython API, and the model capabilities are showcased by simulating unanticipated transients due to multiple components failures (Section 2.5).

## 2.1 Review of relevant experiments and operations

The two major sources of information for the molten salt PCS design are the Generation IV molten salt reactor (Gen-IV MSR) licensing basis [34] and the Molten Salt Reactor Experiment (MSRE) [35–37]. Relevant information from the safety point of view can instead be found in the (liquid metal) ITER Tritium breeding module [38] and Pb-Li experimental setups (see Section 2.1.2 for a detailed list of references). Crucial results for FLiBe behaviour in fusion-relevant conditions, tritium transport, tritium handling and extraction from FLiBe are expected from the LIBRA experiment, which is under development at MIT-PSFC [39].

### 2.1.1 Molten salts

The fuel in the MSRE was a mixture of LiF, BeF<sub>2</sub>, ZrF<sub>4</sub> and UF<sub>4</sub>, while the (secondary) coolant is FLiBe. Temperature and pressure ranges are similar to ARC coolant. The MSRE reported a fuel temperature in the range 908–935 K, a coolant temperature from 825 K to 866 K [35] and a design pressure of 3.4 bar, while ARC FLiBe temperature lays in the range 800 K – 900 K [24] and the design pressure is few bar<sup>1</sup> [24, 3]. Experiences from non-nuclear components can be directly applied to ARC, such as pumps, drain and storage tanks, chemical process units, salt melters, heat exchangers. Operational experiences for the MSRE [35] reported no leakages at any time, whilst small internal leakages of oil in the pump bowl are reported. Heaters experienced few failures that did not affect reactor operation, but their performance were highly reduced by air leakages. Yearly testing of instrumentation found faults in sensors. Finally, the hazard analysis reported the following relevant events: loss of flow, loss of heat sink, decay heat removal, freeze valve and flange failures, excessive wall temperature, corrosion, salt spillage and beryllium release from a leak.

A more detailed analysis on plant hazards and initiating events (IEs) was carried out recently for the Gen-IV MSR. A quite comprehensive list of IEs for a generic molten salt reactor was defined by a group of experts in 2019 [34]. The list includes 140 IEs for the following systems: Reactor and primary fuel salt system, Drain tanks,

---

<sup>1</sup>The precise value of the PCS working pressure has not been defined yet because of the lack of a detailed design (i.e., of the total pressure losses in the PCS). FLiBe boiling temperature at 1 bar is quite large  $T_{\text{FLiBe,boil}} = 1703 \text{ K}$ , so operations at few bar are expected.

Fuel processing system and Reactor building. Among these IEs, 61 are relevant for ARC and about 50 are relevant for the FLiBe circuit. Most of the relevant IEs involve loss of boundaries integrity, loss of heat sink and loss of flow. Those IEs are used to integrate the result from the FMECA. Similar IEs categories are identified in [40] starting from the categories defined for LWRs in NUREG 800. To summarize, the main IEs categories relevant for ARC FLiBe loop are:

- Increase/Decrease in heat removal by the secondary system;
- Decrease in reactor coolant system flow rate;
- Increase/Decrease in reactor coolant inventory.

### 2.1.2 Liquid metals

ITER Tritium Breeding Module safety reports are a valuable source of information for what concerns classical equipment failure modes in a fusion environment. Failure modes like external/internal leakages and rupture are described for the FW and VV, and the failure frequency is quantified [38].

Additional safety experiences are available from experiments with liquid metals, mainly LBE (Lead-Bismuth Eutectic) and Pb-Li. Foletti [41] reported failures of flow meters and pressure transducers in the experimental loops at CR ENEA Brasimone. Oxide buildup was the leading cause of those failures. The TRITEX experimental facility [42] Pb-Li loop was run for about 13,000 hours between 1988 and 1996 [43]. Operational experiences give a good insight on possible failure modes of components and instrumentation. The drain tank was continuously heated to avoid thermal shock issues if the coolant had to be dumped in the tank. Catch pans were located below the experimental device to gather possible coolant leakage - the importance of catch pans and drain lines heating is confirmed by the operational experience by Drobyshev et al. [44] during the operations of liquid metal loops. A non-water fire suppression system was used to avoid steam explosions.

Oil leakage from the pumps shaft was reported many times as a failure event. Gromov et al. [45] reported an accident occurred with Pb-Bi liquid metal coolant. The oil from the shaft seal leaked in the coolant, where the high temperatures pyrolyzed the oil in into other hydrocarbons. The resulting impurities coated the

heat transfer surfaces and plugged the tubes, greatly reducing the heat removal from the core. The Dounreay reactor reported six trips in 5 years due to pump oil system and other auxiliary systems [46]. Instrumentation and their penetrations constitute a notable source of failures [47] as well. Failure of welds from classical instrumentation have been reported for coolant leakages in many reactors.

### 2.1.3 Design and safety consideration for ARC PCS

The review on design and operating experiences of the previous systems allows to define key safety goals for ARC. The first consideration applies to the coolant itself. In the hot leg, FLiBe has the highest tritium concentration of the whole loop. A rupture in this section of the piping system could lead to a large tritium release. Even small leakages from instrumentation welds and flanges constitutes a serious hazard that must be considered in the safety assessment. Beryllium and fluorine dispersion must also be accounted in FLiBe release events. A possible solution is the discharge of FLiBe in a safety tank to limit dispersion of hazardous elements. In a MSR, the drain tank main safety function is to avoid criticality accidents by removing the fuel from the core. This is clearly not a concern for a fusion reactor. However, drain tanks must be present to discharge the coolant in case of severe LOCAs, especially if the failed pipe section cannot be efficiently isolated. Furthermore, an emergency cooling system must be foreseen if the coolant is drained from the loop. Such an additional cooling system greatly increases the whole system complexity. Alternatively, the drain tank can be connected to the main coolant loop, and the blanket decay heat can be removed by a salt-to-air radiator such as in MSRE. Multiple cooling lines must be foreseen in this second case to bypass the fault pipe.

FLiBe pumps should be designed to avoid lubricant oil leakages in the molten salt flow, or the resulting impurities could prevent valves closure and deteriorate the heat transfer mechanism in the heat exchanger. Hence, FLiBe purity must be continuously monitored and guaranteed. This is a function that can be easily integrated in the redox control unit of the PCS that regulates Be content in FLiBe. Even though the main components exploit corrosion resistant materials, instrumentation welds and flanges could be damaged by oxides build-up, resulting in FLiBe leakages. Catch pans should be present below the main components and the piping to avoid FLiBe unconfined release. FLiBe is not flammable, but its high temperature could start a fire due to the interaction with other components and materials in the reactor building.

Also the concrete structure can be seriously damaged by FLiBe high temperature: the mechanical properties of concrete are dramatically reduced at high temperature, showing more than a 50% reduction for the compressive and tensile strength at 800 K [48]. A metallic liner could be envisioned to protect the concrete from direct contact with FLiBe.

Finally, FLiBe freezing risk should be minimized. The immediate danger of FLiBe freezing is relatively low if the reactor is readily shut down. Mechanical stresses are not foreseen due to the volume reduction resulting in FLiBe phase transition. Possible plugs can be by-passed to remove the decay heat, but the subsequent startup could induce strong thermomechanical stresses on the pipes. FLiBe reheating should be carried on extremely slowly to ensure that the temperature increases uniformly in all the components. Electric heaters all along the pipes can fulfil this function. This means that the reactor downtime would be long even though pipe failures are avoided. Minimization of FLiBe freezing probability is therefore mandatory to improve the FPP reliability. Also, the FLiBe heating system should be oversized to face heaters failures: if one or more heaters fail, the remaining working heaters must be capable to increase the heating power to keep FLiBe temperature constant. Finally, the PCS pipes should be thermally isolated to reduce the risk of coolant freeze, but also to increase the reliability of the cryogenic system that is inevitably crossed by the pipes.

## 2.2 Functional analysis of the PCS

The procedure to carry out a functional analysis can be considered complementary to a fault tree construction, with the goal to build a success tree instead. The process can be schematized as follows: system conditions are identified, and a hierarchy among system functions is developed starting from a top function. Each level below the top function lists the immediate sub-functions to fulfil the upper function, avoiding logical jumps. For instance, if the top function is “extract heat from blanket”, the immediate sub-function is “provide a coolant”, not “provide pumping power for FLiBe”. In this way, a detailed description of the functions, thus components, is obtained. Each function in the  $i$ -th level must not overlap with other functions of the same level. The hierarchy is stopped once the required level of detail is reached. For the case under consideration, steady-state, nominal operating conditions are chosen.

The top function is “remove power from blanket”. The physical boundary of the system is the FLiBe loop. The thermodynamic boundary conditions can be derived as follows.

ARC produces 525 MW<sub>th</sub> of fusion power, which increases up to 630 MW<sub>th</sub> due to exothermic reactions in the blanket. The working temperature of FLiBe is lower-bounded by FLiBe melting temperature ( $T_{melt} = 732$  K) and upper-bounded by the maximum operating temperature of VV materials. Ferrero et al. [27] showed that hotspots in FW and VV bring tungsten temperature above its recrystallization threshold, and Inconel 718 temperature above the creep limit. The last issue can be partially mitigated by replacing Inconel 718 with other commercially available alloys (e.g., V-Cr-Ti alloys [49]), but the highest FLiBe temperature to ensure materials integrity in an ARC-class FPP is around 900 K. One of the main advantages of a FLiBe PCS is the possibility to operate at low pressure ( $T_{boil,FLiBe} = 1703$  K at  $p_{FLiBe} = 1$  bar). The thermodynamic boundary conditions for the PCS main function (i.e., reactor cooling) are therefore  $P_{tot} = 630$  MW<sub>th</sub>,  $\Delta T_{FLiBe} \simeq 100$  K (with minimum and maximum operating temperatures  $T_{min} \simeq 800$  K and  $T_{max} \simeq 900$  K), and  $p_{FLiBe} = 1 - 5$  bar depending on the detailed layout of the system and the total pressure drops.

The results from the functional analysis are shown in Table 2.1.

## 2.3 Failure modes and effects analysis of the PCS

The FMEA is performed for the 22 main components identified in the PCS, for a total of 55 possible failures. A table summarizing the results of the FMEA can be found in Appendix A. The reactor is assumed to be operating at nominal conditions, and the components analysed by the FMEA are those that can lead to the loss of the heat removal function of the FLiBe loop. The possible failure modes and the causes leading to failure are derived from the past experiences described in Section 2.1, and from specific publications such as IAEA-TECDOC-478 [50], INEEL/EXT-98-00892 [51] and EGG-SSRE-8875 [52]. These database and publications refer to components that operated in different industrial environments (mainly NPP), which can be considered similar, but not equivalent, to FPP. The failure effects are evaluated locally, at system level and at plant level. The probability and severity categories are derived from MIL-STD-882E [53]. The severity category takes into

Table 2.1 Functional analysis for ARC PCS. The functional requirements, desing parameters and components to fulfil the fuctions are identified.

Functional requirement	Design parameter	Physical part
FR 0: Provide sufficient power extraction from the VV structure	DP 0: Blanket heat extraction system	-
FR 1: Provide a coolant in VV	DP 1: FLiBe molten salt	FLiBe
FR 1.1: Provide a plenum	DP 1.1: FLiBe tank in VV	Tank
FR 1.1.1 Provide continuous boundary for nominal operations	DP 1.1.1: Structural integrity	Tank wall
FR 1.1.1.1: Allow for displacement	DP 1.1.1.1: VV joints	VV joints
FR 1.1.1.2: Sustain high heat fluxes	DP 1.1.1.2: FW	FW and coolant channels
FR 1.1.2: Provide continuous boundary for accidental conditions	DP 1.1.2: Drain safety system	-
FR 1.1.2.1: Provide a safety tank	DP 1.1.2.1: Drain tank	Drain tank
FR 1.1.2.2: Provide connection to drain tank	DP 1.1.2.2: Valve	Freeze valve or temperature safety release valve
FR 1.2: Keep the salt at proper T and p conditions	DP 1.2: FLiBe loop managing system	-
FR 1.2.1: Provide information on FLiBe T	DP 1.2.1: T measurement	T sensor
FR 1.2.2: Provide heating if required	DP 1.2.2: Heaters	Heaters
FR 1.2.3: Provide information on FLiBe p	DP 1.2.3: p measurement	-
FR 1.2.4: Adapt to volume changes and control FLiBe pressure	DP 1.2.4: Pressurizer	Pressurizer
FR 1.2.5: Blow off FLiBe if pressure is too high	DP 1.2.5: Safety release valve	Rupture disk
FR 1.2.6: Blow off FLiBe if temperature is too high	DP 1.2.6: Safety release valve	Freeze valve or temperature safety release valve
FR 1.3: Produce salt	DP 1.3: FLiBe production system	-
FR 1.3.1: Melt or heat salt from storage	DP 1.3.1: Salt melter	Salt melter
FR 1.3.2: Provide an additional inventory of FLiBe	DP 1.3.2: FLiBe Storage	Storage tank
FR 2: Provide a flow path	DP 2: Piping system	-
FR 2.1: Provide connection to the VV	DP 2.1: Pipes	Pipes
FR 2.2: Provide inlet to the VV	DP 2.2: inlet valves	MOV
FR 2.3: Provide outlet to the VV	DP 2.3: outlet valves	MOV
FR 2.4: Ensure the system is working properly	DP 2.4: pressure and flow sensors	Sensors
FR 3: Provide pumping equipment	DP 3: Pumping system	-
FR 3.1: Provide head to the fluid	DP 3.1: FLiBe pressure/pumps	-
FR 3.1.1: Provide measurement of FLiBe pressure	DP 3.1.1: Pressure sensors	Sensors
FR 3.1.2.: Keep pressure to design values	DP 3.1.2: Pumps	Pumps
FR 3.2: Provide power to the pumping system	DP 3.2: Feeding system	-
FR 3.2.1.: Provide power from external sources	DP 3.2.1: Off-site power system	Electrical network
FR 3.2.2: Provide power from internal sources	DP 3.2.2: On-site power system	-
FR 3.2.2.1: Provide power from electricity conversion	DP 3.2.2.1: Diesel generator	Generator
FR 3.2.2.2: Provide power from batteries	DP 3.2.2.2: Battery system	Batteries
FR 4: Provide heat transfer mechanism	DP 4: DT between structure and coolant	-
FR 4.1: Provide low T FLiBe at inlet	DP 4.1: Remove heat from hot FLiBe	-
FR 4.1.1: Provide a system for heat extraction	DP 4.1.1: Heat exchange with secondary coolant	Secondary flow
FR 4.1.1.1: Provide isolation for HX	DP 4.1.1.1: Inlet valve	MOV
FR 4.1.1.2: Provide system isolation from HX	DP 4.1.1.2: Outlet valve	MOV
FR 4.1.1.3: Provide a mean for HT	DP 4.1.1.3: Heat exchanger	Shell and tubes
FR 4.1.1.4: Provide a tank if HX fails	DP 4.1.1.4: Secondary drain tank	Drain tank
FR 4.1.2: Estimate the amount of power to remove	DP 4.1.2: T measurement after HX	T sensor

account consequences on (1) plant economics; (2) workers and population; and (3) environment. Preventive and corrective measures are suggested for each failure, in order of decreasing effectiveness: remove the hazard through design selection, reduce risk through design alteration, incorporate engineered design features or devices. Warnings, procedures and training are not considered at this stage of design. The assigned severity of safety components failures is quite low if one considers that they should perform safety functions. However, the analysis considers nominal operating conditions, not accidental scenarios. The failure of a safety component reduces the availability of the reactor, but do not lead to accidental sequences as long as the reactor is operating at nominal conditions. The function that the system must carry on is heat removal, thus only the failure modes that affect this function are considered. For this reason, vital components like backup diesel generators do not appear in the FMEA. A brief description of the main failure modes for the different components is provided in Table 2.2.

Both qualitative and quantitative results are obtained by the FMEA. A risk matrix is built following the prescription in MIL-STD-882E for the probability category and severity category. Specifically, the probability categories are defined according to the frequency of occurrence:

- A if the event  $E$  is frequent, i.e.,  $P(E) > 5 \text{ y}^{-1}$ ;
- B if the event  $E$  is probable, i.e.,  $5 \cdot 10^{-2} < P(E) < 5 \text{ y}^{-1}$  ;
- C if the event  $E$  has a moderate probability of occurrence, i.e.,  $5 \cdot 10^{-3} < P(E) < 5 \cdot 10^{-2} \text{ y}^{-1}$ ;
- D if the event  $E$  has a low probability of occurrence, i.e.,  $5 \cdot 10^{-3} < P(E) < 5 \cdot 10^{-4} \text{ y}^{-1}$ ;
- E if the event  $E$  has a very low probability of occurrence, i.e.,  $P(E) < 5 \cdot 10^{-4} \text{ y}^{-1}$ .

The same criteria are used to define the probability category for each failure mode. The severity category can be defined according to the economics, workers and population, or environmental consequences. These criteria are shown in Table 2.3.

The risk matrix (Table 2.4) reports one high-risk failure only. One failure is classified as severe, six failures are classified as serious, 32 as medium and 16 as

Table 2.2 List of relevant failure modes in ARC PCS. The second column provides a short description of each failure mode. The components that are susceptible to a failure mode are reported in the third column.

<b>Failure mode</b>	<b>Description</b>	<b>Applies to</b>
External leakage	Leakage from the external boundary of a component	Walls, HX shell, pipes, valves, sensors
External rupture	Rupture of the external boundary of a component	Walls, HX shell, pipes, valves
Internal leakage	Leakage from the internal boundary of a component	Valves
Internal rupture	Leakage of the internal boundary of a component	Valves
Plugging	Reduction of the available flow cross section up to complete obstruction	Pipes
Fail to operate	Unable to fulfill the expected function	Sensors
Improper reading	Wrong output from a measurement	Sensors
Fail to start	Component fails to start from an off state	Pumps, heaters
Fail to run, given start	Component fails to keep functioning following a successful start	Pumps
Fail to remain in position	Component fails to stay in the original position	Process valves and safety valves
Spurious thaw	Melting of the frozen plug of the freeze valve	Freeze valve

Table 2.3 Severity criteria in case of failure for economics and plant operations, workers and population, and environment.

<b>Severity category</b>	<b>Economic/Plant consequences</b>	<b>Harmful effects and (workers population)</b>	<b>Harmful effect (environment)</b>
1	Monetary loss exceeding \$10M	Death, permanent total disability	Irreversible significant environmental impact
2	Monetary loss between \$1M and \$10M	Permanent partial disability, hospitalization of at least 3 personnel	Reversible significant environmental impact
3	Monetary loss between \$100k and \$1M	Injury resulting in one or more lost day(s)	Reversible moderate environmental impact
4	Monetary loss less than \$100k	Injury resulting not resulting in a lost day	Minimal environmental impact

low risk. The small number of failures in the serious and severe region is due to the reactor status considered. If the reactor had been considered under accidental scenarios, the risk would have been higher because of the more severe consequences and the component failure probability on demand. The reliability of this first design is confirmed by the first and the second row of the table: none of the failure falls in the A probability category, and just four failures fall in B. Instead, five failures lead to the worst consequences. Those failures are first wall and tank wall rupture, main pipes plugging and heat exchanger shell rupture. The consequences resulting in first wall or tank wall rupture can be mitigated to a certain extent only. In fact, while consequences to the workers and environment can be limited by additional safety barrier, the economic damage is severe. The expected cost of these two components is approximately \$100M and \$260M respectively [24]. Particular attention should be paid to the design of the heat exchanger and the choice of the secondary fluid. FLiBe low pressure is a desirable safety feature, but it leads to strong pressure difference in the heat exchanger. The heat exchanger tubes may experience strong mechanical stresses, and internal leakage is likely if the tubes lose their structural integrity. The secondary fluid would create a strong overpressure in the primary system, until the rupture disk in the expansion volume breaks. Another component that deserves special attention is the freeze valve, which has an interesting safety feature thanks to its passive mechanism. However, its fail-safe feature could lead to unexpected reactor

Table 2.4 Risk matrix resulting from the FMEA. Green: low risk; Yellow: medium risk; Orange: severe risk; Red: high risk.

		Severity			
		1	2	3	4
Probability	A	0	0	0	0
	B	0	1	3	0
	C	0	2	8	9
	D	1	0	3	5
	E	4	6	11	2

shutdown due to spurious thawing. Freeze valves present indeed a quite high failure rate that could negatively affect reactor availability [54]. Since the consequences of a drain tank safety valve failure are much less severe than those of a MSR (where the molten salt must be discharged to avoid super criticality), two standard safety valves in parallel can be envisioned to replace the freeze valve, resulting in a lower failure rate ( $7 \cdot 10^{-6} \text{ h}^{-1}$  for the freeze valve vs  $2.5 \cdot 10^{-5} \text{ h}^{-1}$  for the standard safety valves). Finally, sensors failures populate the low-risk area. A minimum redundancy of two has been assumed in the initial design. This makes the severity of a sensor failure quite low, because at least another sensor is expected to provide the required information. A possible criticality may arise from the improper reading of a sensor. If only two sensors are present, the erroneous reading is not filtered. This suggests the implementation of a voting logic with at least 3 sensors for the most important readings.

The initiating events triggered by a FLiBe loop component failure are identified from the FMEA. Also, IEs frequency is estimated based on the single components failure. This is an important point to be stressed. The FMEA considers single failures only. Therefore, IEs that arise from the failure of two or more components are overlooked. Even though the simultaneous failure of two components is usually less likely than the failure of a single component, a common cause failure could make those two probabilities comparable [55]. Hence, the IEs frequency obtained by the FMEA should be considered as not conservative. Table 2.5 summarizes the IEs annual frequency.

VV LOCA, tank LOCA and FLiBe loop LOCA annual frequency are quite low. The VV LOCA includes both VV leakage and rupture. In fact, even a small

Table 2.5 Postulated initiating events and frequency of occurrence according to the results from the FMEA.

IE category	Faults resulting in IE	Frequency ( $y^{-1}$ )
VV LOCA	First wall - Internal leakage First wall - Internal rupture	$1.68 \cdot 10^{-5}$
Tank LOCA	Tank - External rupture	$1.96 \cdot 10^{-6}$
FLiBe loop LOCA	Pressurizer - External rupture Temperature safety relief valve - External rupture Pipes (from blanket to TES) - Rupture Pipes (from TES to blanket) - Rupture Blanket inlet valve (MOV) - External rupture Blanket outlet valve (MOV) - External rupture Heat exchanger - Tube rupture HX inlet valve (MOV) - External rupture HX outlet valve (MOV) - External rupture	$9.83 \cdot 10^{-3}$
Small LOCA	Freeze valve - Spurious Thaw Pressurizer - External leakage Temperature safety relief valve - External leakage Blanket inlet valve (MOV) - External leakage Blanket outlet valve (MOV) - External leakage HX inlet valve (MOV) - External leakage HX outlet valve (MOV) - External leakage	$2.76 \cdot 10^{-1}$
LOFA	Pipes (from blanket to TES) - Plugging Pipes (from TES to blanket) - Plugging Blanket inlet valve (MOV) - Plugging Blanket outlet valve (MOV) - Plugging Pump (motor driven) - Fail to run, given start Heat exchanger - Plugging HX inlet valve (MOV) - Plugging HX outlet valve (MOV) - Plugging	$2.72 \cdot 10^{-1}$

leakage of FLiBe in the vacuum chamber could trigger a disruption, leading to severe consequences. Instead, the tank external leakage is not considered for tank LOCA. The “small LOCA” frequency and LOFA frequency are quite high. However, considering a small LOCA as IE is an arbitrary decision: on one hand, a small LOCA could hardly start a chain of events resulting in reactor damage and external radioactivity release; on the other hand, it may be considered an internal radioactivity release. The events evolution will depend in large measure by the PCS and multiple barriers design. The situation is different for LOFA. An expected LOFA in 5 years arises concern. The contribution to LOFA frequency of occurrence comes almost entirely from the pump failure. The second highest contribution derives from heat exchanger tubes plugging. The installation of a redundant pump and the possibility to divert the flow to a different heat exchanger reduces LOFA frequency by more than two order of magnitudes. Note that VV leakage and rupture could fall also in the LOVA (Loss of Vacuum Accident) category, but LOVA are outside the scope of the present analysis. These results set the bases for further, more detailed safety analyses. The postulated IEs frequency is comparable to that of other fusion blanket systems. However, the commonly used databases report components failure rates derived from operations in industrial or nuclear power plant environments. Specific assessment of components behaviour in fusion relevant environment should be carried out in the near future.

## **2.4 Development of a thermal-hydraulics model of the PCS**

The thermal-hydraulics model of ARC PCS has been developed in Modelica, which is an object-oriented programming language. Object-oriented modeling (OOM) is a well-suited approach for modeling the dynamics of complex systems and performing safety analysis. Notably, several system-level codes used for nuclear safety assessment are based on OOM, including RELAP-7 [56] and MELCOR [57]. The modularity and scalability features of OOM are particularly advantageous for conceptual designs, as once the basic objects are defined, multiple instances of these objects can be connected to build complex systems. Design parameters can be easily modified to perform parametric analyses, providing a wide coverage of the system

design space. Furthermore, objects can be easily manipulated from external routines, making the system a black box that can be integrated into broader models.

Relevant applications of OOM for safety analysis can be found in the work from Destino et al. [58] and Zio et al. [59]. Destino et al. employed a dynamic model of the superconducting magnet cryogenic cooling circuit of ITER to simulate many accidental scenarios involving the magnets under different operating conditions. Zio et al. exploited a 0-D thermal-hydraulics and neutronics model of the LBE-XADS [60] to carry out dynamic reliability analysis. In both cases, the authors exploited artificial intelligence algorithms to post-process the huge amount of data obtained.

The object-oriented model of the PCS presented in this section represents a first step towards a detailed safety assessment. The model is then embedded in a probabilistic framework by coupling it with a Monte Carlo routine for sampling component failures (Section 2.5). Such an approach allows for the explicit consideration of time-dependent interactions between physical phenomena, control logic, and component failures, and avoids analyst bias in defining the accident (or transient) sequences. Furthermore, the order, timing, and magnitude of multiple failures are considered, which is fundamental for assessing the evolution of transients and associated risks, as noted by Siu [61] and Aldemir et al. [62]. Accidental scenarios are not chosen based on expert judgment but arise naturally from the failure probability of components (or subsystems), making no *a priori* distinction between expected transients, minor accidents, and major accidents. Post-processing of the results determines whether a transient ends in a system failure or safe state.

The number of possible scenarios to be simulated may become extremely large when many components, failure times, and failure modes are considered [63], requiring a trade-off between model accuracy and computational cost. The goals of the present model can be summarized as follows:

- Low computational cost to explore a wide range of possible transients, to compare different design options, and to find system criticalities;
- Easy integrability with external routines for future applications in an integrated deterministic and probabilistic safety assessment [17];
- Efficient post-processing to handle the huge number of results.

Since no experimental results are available for model validation, a verification against an analytical model and a benchmark against available results from detailed component-level simulations are carried out.

### 2.4.1 PCS layout

A brief description of the PCS based on Kuang et al. [3] and Ferrero et al. [27] works is provided here for readers' convenience. Cold FLiBe ( $\sim 800$  K) from the heat exchangers is pumped by molten salt pumps in the LIB tank, VV channels and divertor channels. The VV channel outlet is located in the LIB tank itself, meaning that FLiBe flowing through the VV channels is discharged in the tank. Similarly to the VV channels, divertor channels are fed with cold FLiBe from the HX, and their outlet is located in the tank itself. A separated coolant pipeline supplies cold FLiBe to the LIB tank from the HX [3]. All the pipes exit from the top of the LIB tank [3]. A layout of the LIB is shown in Figure 2.1. The PCS includes also components that belong to the fuel cycle (e.g., a tritium extraction system placed downstream of the LIB and upstream of the HX (Chapter 3), but those are not included in the thermal-hydraulics because they don't fulfil any heat transfer functions. Therefore, the FLiBe exiting the LIB is assumed to directly flow to the HX.

The conceptual layout of the PCS implemented in the OpenModelica Connection Editor [64] is shown in Figure 2.1.

The following simplifying assumptions have been made:

- 1) The divertors are not modeled due to lack of design specifications;
- 2) The VV sector and the LIB tank are connected in series;
- 3) The FLiBe pumps are placed between the VV sector and the LIB tank. The location of the pumps has not been identified yet but placing them inside the blanket tank seems extremely unlikely due to the engineering complexity of such solution.

Assumption 1) is dictated by the absence of a preliminary design of the divertor cooling system. While other PCS components have been designed specifically for this model (e.g., molten salt pumps and HX), divertor designs affect reactor




Fig. 2.1 Primary cooling system model. The main components depicted are the VV section, the molten salt pumps (only one shown), the blanket tank, the heat exchangers (only one shown). An enthalpy sensor at the heat exchanger outlet is connected to the electric heaters controller in the VV section. The VV section and the blanket tank are thermally connected to take into account conduction from the outer VV structural material to tank wall. A ramp is used to simulate pump failures. Internal power generation from neutron heating is accounted in the tank. A constant mass flow rate component sets He flow in the heat exchanger. The sink pressure components provide a reference pressure for both FLiBe and He flows.




Fig. 2.2 Layout of ARC cooling system. Flow directions in VV channels and LIB tank are depicted. FLiBe enters the VV channels from external pipes that pass through the LIB tank. The VV channels end in the tank itself. Reproduced from [1].

operations from many different perspectives (e.g., magnets location and magnetic configuration, particle exhaust, etc.) that their design can be approached in an integrated way only. This is clearly beyond the scope of this work. Assumption 2) is needed to simplify the fluid dynamics in the LIB. Ferreto et al. [27] showed that the peculiar layout of the LIB tank (i.e., the VV channels are fed by dedicated pipes, while the channels outlet discharges directly inside the tank, as shown in Figure 2.2) lead to a complex flow path that cannot be captured by a system-level model. Comparison of key figures of merit (Section 2.4.3) demonstrates that this approach preserves the essential fluid dynamics and the heat transfer features in the LIB, yet exploiting a much simpler and fast model. Assumption 3) provides a strong numerical stabilization of the model without affecting the results.

Besides these assumptions, additional components appear in Figure 2.1. A pressure sink is used to set the pressure in the closed circuit. An actuator sets the pump rotational speed. A sensor measures the enthalpy at the HX outlet. A mass flow rate source and another pressure sink are used to model the secondary coolant flowing in the heat exchanger. Section 2.4.2 provides a detailed description of the model components and their governing equations.




Fig. 2.3 Poloidal cross section of an ARC-class tokamak. The main layers are depicted. A thin tungsten FW protects the VV structural material. The VV is placed inside the LIB tank. An outer shield surrounds the outer tank walls. Reproduced from [2]

## 2.4.2 Model description

Basic models for channels, pumps and walls were imported from the *ThermoPower* library, which has been extensively validated [65, 66]. The VV sector and the HX control system models were built taking advantage of the standard components from *ThermoPower*. The LIB tank was modeled as a 0D component. FLiBe was modeled by exploiting a user-defined *medium* with temperature dependent properties. Table 2.6 shows the structural and functional materials of ARC VV and LIB, and the corresponding thickness [3]. Figure 2.3 shows the poloidal cross section of an ARC-class tokamak and depicts the main components.

The model utilizes the finite volumes method to discretize flow components, such as channels, as well as the structural and functional materials of the VV, HX and LIB tank. The discretization of these components is applied along the axial direction by 10 nodes. Taking advantage of the axial symmetry of the reactor, the model considers a slice of ARC blanket as representative of the conditions of the entire

Component	Material	Thickness [mm]
FW	Tungsten	1
Inner structural material (VV)	Inconel-718	10
Coolant	FLiBe	20
Neutron Multiplier	Beryllium	10
Outer structural material (tank wall)	Inconel-718	30

Table 2.6 Structural and functional materials of the VV and the LIB.

blanket. This slice encompasses all components located along the radial direction (FW, VV structural layers, VV channel, Be multiplier, and LIB tank), while it has a thickness of 2 cm in the rectified toroidal direction. In other words, the slice accounts for a single VV channel ( $D_{ch} = 2$  cm) and all the components that are intersected by the projection of the channel along the radial direction.

Each flow component solves mass, momentum, and energy equation:

$$\frac{dm}{dt} = \dot{m}_{in} - \dot{m}_{out} \quad (2.1)$$

$$\frac{L}{A} \frac{d\dot{m}}{dt} + (p_{out} - p_{in}) + \Delta p_{st} + \Delta p_{frict} = 0 \quad (2.2)$$

$$Al\rho_i \frac{dh_i}{dt} + \dot{m}_i(h_{i+1} - h_i) = \dot{Q}_{tot,i} \quad (2.3)$$

where  $m$  is the fluid mass in the component,  $\dot{m}$  is the average mass flow rate,  $\dot{m}_{in}$  is the mass flow rate entering the component, and  $\dot{m}_{out}$  is the mass flow rate exiting the component.  $A$  is the channel/tube cross sectional area,  $L$  is the total channel/tube length,  $p$  is the fluid pressure,  $\Delta p_{static}$  are the static pressure drops, and  $\Delta p_{frict}$  are the frictional pressure drops. The energy equation is solved in each control volume. Subscript  $i$  refers to the  $i$ -th volume,  $l = \frac{L}{N_{nodes}}$ ,  $h$  is the fluid enthalpy,  $\rho$  is the fluid density,  $\dot{Q}_{tot}$  is the sum of convective and conductive heat flux, and power generation inside the  $i$ -th volume of the component.

### Molten salt pump

The molten salt pump is described by a quadratic characteristic curve. The characteristic curve is derived from [67] and scaled down to meet ARC requirements. The pump rotational speed is provided as external input to simulate different transients involving pump failure. Since the FLiBe mass flow rate flowing in the blanket slice is much lower than the total FLiBe mass flow rate, the working point has been shifted to the slice mass flow rate. That is, if the pump provides a pressure head  $H_{nom}$  at  $\dot{m}_{nom} = \frac{\dot{m}_{FLiBe,tot}}{n_{pumps}}$ , the pump in the TH model provides the same pressure head  $H_{nom}$  at  $\dot{m}_{nom} = \dot{m}_{FLiBe,slice}$ .

### LIB tank

The complex geometry of the blanket makes it difficult to model the tank as a 1-D component. Large eddies are present in the central area of the tank, while small recirculation zones are found near sharp corners [27]. Therefore, FLiBe tank was modeled as a lumped component. A convective and conductive heat flux is exchanged with the outer wall of the VV. Radiative heat transfer is neglected due to the negligible contribution it gives at such high operating temperatures [68]. The volumetric power generation for FLiBe and blanket materials was obtained from OpenMC [69] 2-D neutronic simulations of ARC [25]. The volumetric power generation within the tank slice is  $\dot{q}_{vol} = 1.4 \text{ MW/m}^3$  (adapted from Segantin et al. [25]), for a total power generation in the slice  $\dot{Q}_{vol} = \dot{q}_{vol} V_{tank,slice} = 210 \text{ kW}$ , being the tank slice volume  $0.15 \text{ m}^3$ . This term considers also the exothermic reactions that take place in the blanket.

The heat transfer coefficient ( $\gamma$ ) between FLiBe, which is flowing in the tank, and the outer VV wall is set to  $\gamma = 1000 \text{ W m}^{-2} \text{ K}^{-1}$  under steady-state conditions, according to the findings of Ferrero et al. [27]. However, questions arise on the validity of a constant heat transfer coefficient during transients. Accidents such as pump failure can reduce the total coolant mass flow rate and, consequently, the convective heat transfer between the tank wall and the FLiBe inside the tank. Therefore, a parametric analysis was carried out to quantify the variation of the structural materials temperature for a wide range of  $\gamma$ . Results are shown in Figure 2.4, where no significant variations of the materials temperature is recorded, except for the outer VV wall. Moreover, the convective heat flux in the tank slice

( $Q'' = 9.2\text{kW}$  for  $\gamma = 1000\text{W m}^{-2}\text{K}^{-1}$ ) is considerably lower than the power generation inside the tank. Thus, for the nominal coolant mass flow rate,  $Q'' \ll Q_{\text{vol}}$ . In case of a lower mass flow rate caused by pump failure,  $\gamma_{\text{pump failure}} < 1000\text{W m}^{-2}\text{K}^{-1}$ , and  $Q'' \ll Q_{\text{vol}}$  holds true. Assuming a constant heat transfer coefficient leads to an overestimation of  $Q''$ , which is negligible in terms of FLiBe temperature. Nonetheless, this assumption can slightly underestimate the blanket material temperature. Notably, the VV structural material (Structural inner) always experiences higher temperatures than the blanket structural material (Structural outer), and potential failures are therefore driven by the inner structural layer. Thus, the assumption of a constant heat transfer coefficient in the tank does not have a significant impact on the present analysis.

From a hydraulic standpoint, the VV channel and the tank can be regarded as two parallel flow paths. The problem of calculating the pressure drop in such a complex geometry is shifted to the pressure drop calculation in the VV channel. Section 2.4.3 demonstrates that the results from the blanket slice are consistent with the CFD analysis results despite the strong modeling assumptions applied to the tank.

### VV sector

The VV sector includes the tungsten FW, the inner and outer Inconel 718 structural layer, the beryllium neutron multiplier and the FLiBe channel (Figure 2.5 and Figure 2.6). Solid materials were modeled with the *MetalWallFV* component from *ThermoPower* library. An internal power generation was added to the original definition of the component to properly account for the neutrons power deposition. The power deposition in VV materials was derived from Segantin et al. [25] as described in Section 2.4.2. The temperature drop along the radial direction is computed using the thermal resistance method:

$$q'' = U(T_{\text{bound}} - T_{\text{vol}}) \quad (2.4)$$

where  $U$  is the global heat transfer coefficient, which simply reduces to  $U = R_{\text{th}}^{-1} = (s/k_{th})$  for a plane wall, where  $s$  is the wall thickness and  $k_{th}$  is the thermal conductivity. The time-dependent energy balance equation is also solved in each solid volume. The channel takes advantage of the *Flow1DFV* component, which was extended to account for internal power generation and heat transfer with two




Fig. 2.4 Temperature of VV and blanket materials as a function of the convective heat transfer coefficient ( $\gamma$ ) between tank wall and FLiBe inside the tank. First wall, inner structural material and beryllium are almost unaffected by variations of  $\gamma$ .




Fig. 2.5 Simplified representation of a VV channel. A constant heat flux is applied to the FW. Internal power generation due to neutron interactions is considered in all the materials. Electric heaters can provide additional power on demand (note that the heaters will be likely placed upstream of the VV channel, before the pipes enter the LIB tank). Purple: FW; Grey: Structural material; Blue: FLiBe; Green: Beryllium.

adjacent components (e.g., inner Inconel 718 layer and Be multiplier). The heat transfer coefficient and friction factor are computed according to the Dittus-Boelter correlation [70] and the Colebrook correlation [71], which are valid for nominal operating conditions ( $Re \approx 10^4$ ,  $Pr \approx 25$ ,  $L/D = 350$ ) [72]. The lowest Reynolds number in the channel occurred for pump failure at 33% of nominal rotational speed, with  $Re = Re_{\min} = 2.9 \cdot 10^3$ . As such, the Colebrook correlation remained valid for the entire range of transients analyzed, whereas the Dittus-Boelter correlation was valid only for  $Re > 10^4$ . For transients leading to lower  $Re$ , the Gnielinski correlation is used, which is valid for  $3000 < Re < 5 \cdot 10^6$  [73].

### Electric heaters

The electric heaters (Figure 2.6) provide an additional power input to the FLiBe flowing inside the channels. Heaters power has been calibrated on steady-state operations. The total power delivered by the heaters is 100 kW, which allows to heat FLiBe in a smooth transient if its temperature drops below a critical value, set at 775 K (freezing temperature for FLiBe is 732 K).

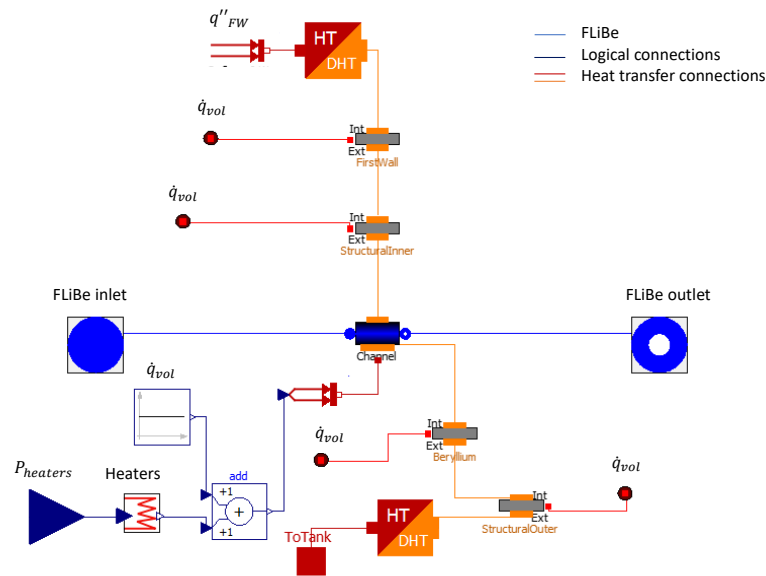


Fig. 2.6 VV sector Modelica model. Starting from the top, a constant heat flux is applied to the first wall. The first wall is connected to the inner structural material through a heat transfer port. The VV channel exchanges heat with both the inner structural material and the beryllium layer. Two flow ports connect the channel input and output with the loop components. Electric heaters are connected to the channel to provide additional power on demand. The beryllium layer is connected to the outer structural material that exchanges heat with the FLiBe tank. Internal power generation is considered in all the components shown in the figure. Note that both red and orange lines refer to heat transfer connections.

## Heat Exchanger

The heat exchanger was modeled using the HE component. The heat exchanger is a shell and tube type with helium flowing outside the tubes and FLiBe flowing inside the tubes. The nominal He mass flow rate is  $\dot{m}_{He} = 60.18 \text{ kg/s}$  according to the Brayton cycle configuration proposed by Segantin et al. [29]. The FLiBe mass flow rate entering the HX is  $\dot{m}_{FLiBe,HX} = \frac{\dot{m}_{FLiBe,tot}}{n_{HX}}$ , where  $n_{HX}$  is the number of heat exchangers, with the thermodynamic state determined by the outlet conditions from FLiBe tank. A PI controller regulates the helium mass flow rate to keep the FLiBe outlet temperature as close as possible to the nominal value. Controller gain ( $k_{gain} = 10^{-4}$ ) and time constant of the integrator block ( $t_{int} = 5 \text{ s}$ ) were calibrated around steady-state operations.

The HX model is shown in Figure 2.7 with the corresponding control logic system and secondary coolant circuit (components related to the secondary cooling circuit are represented in gray in the figure). The control logic of the HX is quite straightforward, but complications arise when implementing "artificial" logic to simulate controller failures. A PI controller receives the measurement from an enthalpy sensor at the FLiBe outlet of the HX. The PI controller is connected to a switch that is triggered in case of failure. Otherwise, the signal from the controller is received by the `SourceMassFlow` of the secondary coolant, which regulates the secondary coolant mass flow rate. In nominal conditions, only these components operate. A `booleanStep` triggers the switch when  $t = t_{PI, failure}$ , where  $t_{PI, failure}$  is the PI failure time computed by an external Monte Carlo routine (see Section 2.5). Note that  $t_{PI, failure} = \infty$  if no failures are sampled, i.e. the PI controller operates as expected. Once the switch is triggered by a PI controller failure a ramp modifies the controller output according to the failure magnitude of the PI. The failure magnitudes ranges between  $[0, 2]$ , so the secondary coolant mass flow rate ranges from zero to twice the nominal value. The ramp has a 10s duration to allow a smooth transition between the nominal value and the fault value, avoiding numerical errors. To enable this failure mode additional "artificial" logic is required. A mass flow rate sensor measures the secondary mass flow rate, and sends the output to a `triggeredSampler`. The `triggeredSampler` is normally opened (i.e., it does not affect nominal operations), but it is closed by the signal from the `booleanStep` when the PI controller fails. In this way the ramp receives the secondary mass flow rate at  $t = t_{PI, failure}$ .

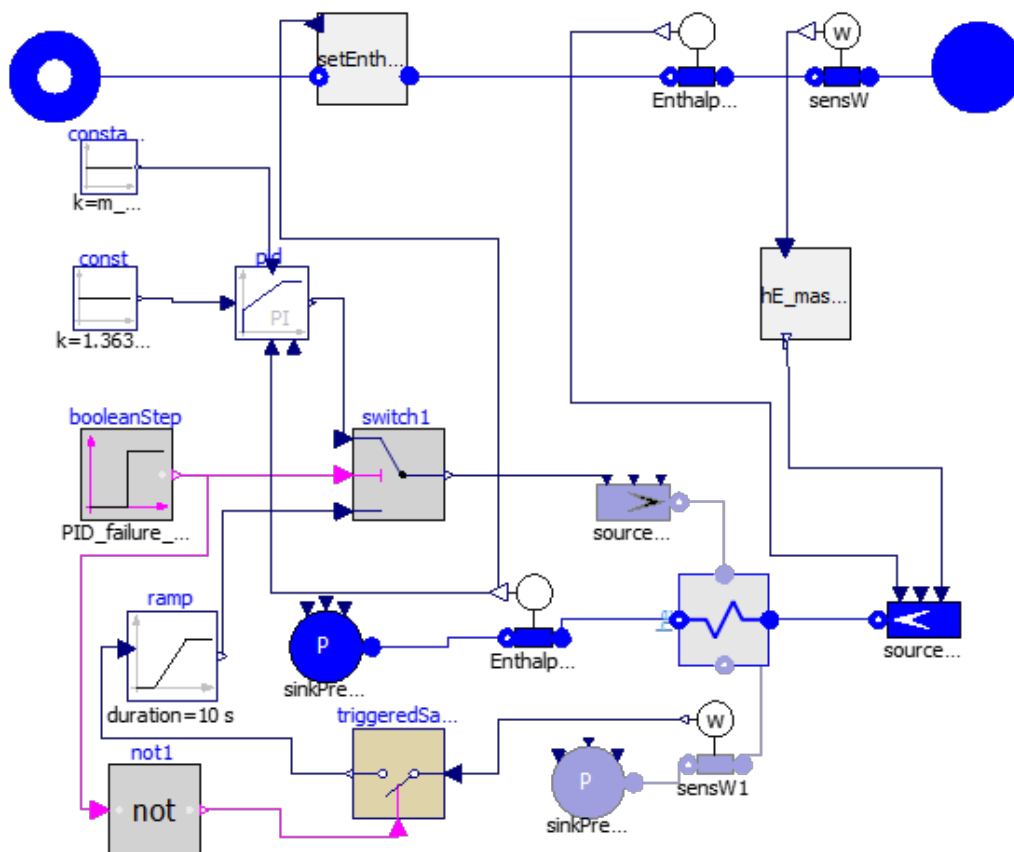


Fig. 2.7 Heat exchanger model, including the control logic and the additional components required to trigger the PI controller failure. The primary coolant components are in blue, secondary coolant components are in gray.

### 2.4.3 Verification and benchmark

A verification of the model was performed to verify the correct implementation of new components and new features in existing components. Most of the components from *Modelica Standard Library* and *ThermoPower* have been validated. Therefore, the verification was limited to the VV sector. An analytical case was considered. FLiBe flow in the channel was simulated numerically and the results were compared to the analytical solution. This procedure also allowed to check the validity of the FLiBe Medium model. The 1-D energy balance equation for a fluid flowing in a heated channel with internal power generation and constant wall temperature reads as:

$$A\rho c_p \frac{\partial T}{\partial t} = 2\pi R\gamma(T - T_{bound}) + \dot{q}_{vol}A - \dot{m}c_p \frac{\partial T}{\partial x} \quad (2.5)$$

where  $A$  is the cross-sectional area of the channel,  $R$  is the channel radius,  $\rho$  and  $c_p$  are the density and the specific heat capacity at constant pressure of the fluid,  $T_{bound}$  is the wall temperature and  $T$  is the fluid temperature. Axial conduction was neglected. The analytical solution ( $T(t, x)$ ) of the energy equation can be found by computing the Laplace transform, solving for  $x$  and performing the inverse Laplace transform:

$$T(t, x) = \frac{Q'v}{\dot{m}c_p}t - \frac{Q'v}{\dot{m}c_p}H\left(t - \frac{x}{v}\right) \cdot \left(t - \frac{x}{v}\right) + T_0 \quad (2.6)$$

where  $Q' = q''2\pi R + \dot{q}S$ ,  $v$  is the fluid velocity,  $T_0$  the initial temperature,  $S$  is the heat transfer surface, and  $H\left(t - \frac{x}{v}\right)$  is the Heaviside function. Figure 2.8 shows the comparison between the analytical and the numerical solutions. The numerical solution shows a knee instead of a sharp corner when the steady-state condition is reached due to numerical diffusion. Apart from this small difference, the two solutions agree.

Since results from more detailed blanket models have not yet been validated, the comparison between the current model and the 2-D CFD models ([3, 27]) should be considered as a benchmark rather than a validation. Nevertheless, the comparison provides a consistency check, which is mandatory due to the strong modeling assumptions that affect the blanket. The model predicts well the temperature variation through the blanket (Figure 2.9). A short transient occurs for the inlet temperature

during the first seconds of the simulation. The solver tries to match all the initialized variables until a convergence is reached. Hence, it has no physical meaning and can be neglected. After this transient, the inlet and outlet temperatures match those computed by Kuang et al. [3].

The friction pressure drop was compared to the results from Ferrero et al. ( $\Delta p_{\text{bl,ref}} = 0.96$  bar) [27]. Quantifying the pressure drop inside the tank is extremely difficult due to the complex geometry and the presence of large eddies. However, if the channel is considered as a parallel flow path, pressure drops in the channel equal the pressure drops in the tank. Local pressure drops were also considered because of the sharp turns of the VV. Four  $90^\circ$  turns ( $k_{90} = 0.7$  [71]) and four  $45^\circ$  turns ( $k_{45} = 0.4$  [71]) contribute to the total pressure drop. The four sections of VV channels were modeled with a  $2L = 14$  m long channel. The distributed pressure drops are  $\Delta p_{\text{dist}} = 0.67$  bar. Addition of the localized pressure drops results in a total pressure drop in the blanket of  $\Delta p_{\text{bl}} = \Delta p_{\text{dist}} + \Delta p_{\text{loc}} = 0.94$  bar. The relative error with respect to the pressure drop computed by Ferrero et al. [27] is  $\varepsilon_{\Delta p} = 2\%$ . Therefore, the model can be considered consistent.

## 2.5 Integration of the thermal-hydraulics model within a probabilistic framework

The model presented in Section 2.4 is purely deterministic. Accidental scenarios or minor transients can be investigated by defining system conditions. The PIEs identified in Section 2.3 can be used to set the boundary conditions of the simulation, and a deterministic analysis can be carried out. This is the approach followed for ITER [74] and DEMO [75] safety analysis.

A key issue of this approach is the domain coverage of accidental scenarios. The focus of this section is therefore the integration of the deterministic model into a probabilistic framework. Deterministic safety analysis is indeed a well established approach, and deterministic thermal-hydraulics codes have been developed and validated for this purpose. Integration of deterministic models in probabilistic routines is instead much less investigated in literature. The key steps to develop an integrated model are:




Fig. 2.8 Comparison of FLiBe temperature at channel outlet computed analytically and by the VV sector model.

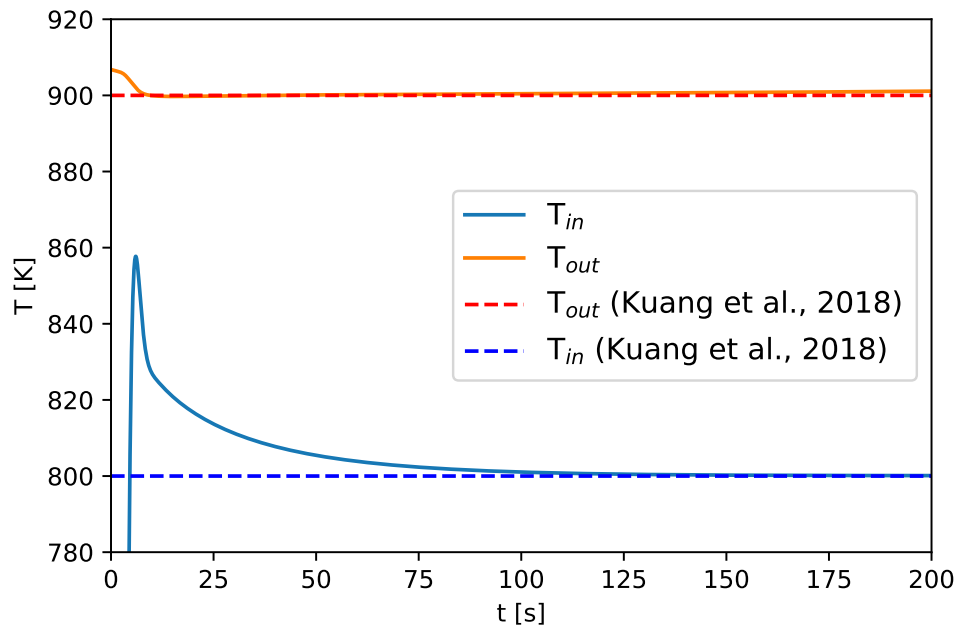


Fig. 2.9 Comparison between the blanket inlet and outlet FLiBe temperature computed by Kuang et al. [3] and by the present model. A short transient occurs at the beginning of the simulation before the solution converges at  $T_{in} = 800$  K and  $T_{out} = 900$  K.

- Gather probabilistic (failure) data for system components. This has already been done in Section 2.3;
- Develop a probabilistic routine embedding the deterministic model. This task can be subdivided in:
  - Develop a Monte Carlo (MC) routine to randomly sample failure times and failure magnitudes of the components. This is called failure injection engine in the remainder of this work;
  - Develop a wrapper that interfaces the failure injection engine with the TH model;
  - Adapt the TH model in such a way that simulations can be automated by the wrapper. The input deck of the TH model has to be accessible from external routines (the wrapper).
- Identify and develop suitable tools for post-processing. The huge amount of data retrieved by this approach cannot be processed in a classic way (i.e., by visual inspection of the analyst). Data mining techniques are required to extract relevant information from the resulting database of simulations.

Readers familiar with deterministic TH codes (e.g., MELCOR) can appreciate the enormous effort required to integrate such codes in a probabilistic framework, not to mention the challenges related to build an automated pipeline. Issues related to the integration of existing codes with probabilistic routines are discussed in Section 2.7, after the results from the integrated model are presented. The development of an automated pipeline for safety analysis is instead discussed in Chapter 3, where machine learning algorithms are introduced and the concept of pipeline is clarified.

### **2.5.1 Development of the failure injection engine**

The failure injection engine is a MC routine specifically developed to sample failure times and failure magnitudes. Component failure time can be randomly sampled from the corresponding probability density function (PDF). The failure rate provided by standard database (e.g., WASH-1400 and IAEA-TECDOC-478 [14, 50]) is a constant (mean) value. An exponential PDF with constant failure rate is therefore considered for the system components. The failure magnitude of the components is a

## 2.5 Integration of the thermal-hydraulics model within a probabilistic framework 39

less investigated parameter. Components are usually assumed to fail completely (i.e., a binary logic is used to describe component failures), such as in fault tree analysis [76] or DBA analysis [32]. However, components can fail partially, and the failure magnitude quantifies the loss of functionality of the component. For example, a pump failing with a failure magnitude equal to 50% experiences a rotational speed reduction by 50%. The failure magnitude can be a discrete or a continuous variable, depending on the level of details sought and the available data. Since quantitative data is not available in literature, a discrete failure magnitude (e.g., 0%-33%-66%-100%) is assumed for the system components. Failure magnitudes are assumed to belong to a uniform distribution.

From a mathematical point of view, the failure time and the failure magnitude are random variables defined by the corresponding PDF. Random variables can be sampled by MC codes by means of the direct or indirect methods [77]. The direct method can be applied straightforwardly when the cumulative density function (CDF) of a given PDF can be inverted. Since failure times and failure magnitudes are modeled by an exponential distribution or by a uniform distribution, the random variables can be easily sampled.

The failure probability of a component with constant failure rate  $\lambda$  is:

$$f(t) = \lambda e^{-\lambda t} \quad (2.7)$$

and the corresponding CDF is:

$$F(t) = 1 - e^{-\lambda t} \quad (2.8)$$

the failure time  $t_{fail}$  can be sampled by taking advantage of a continuous random variable  $\xi$  and inverting Eq. (2.8):

$$t_{fail} = -\frac{\ln(1 - \xi)}{\lambda} \quad (2.9)$$

where  $\xi$  is uniformly distributed. A similar approach can be used to sample the failure magnitude  $m_{fail}$  by using a discrete random variable  $\xi'$ :

$$m_{fail} = \xi' \quad (2.10)$$

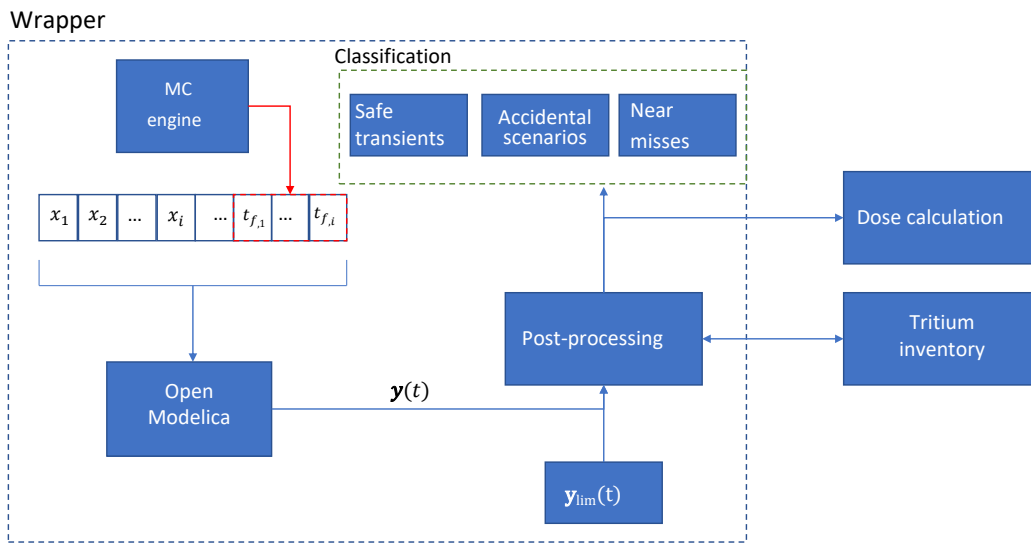


Fig. 2.10 Python wrapper developed to handle preprocessing, post-processing and simulation run.

where  $\xi'$  is sampled from a (discrete) uniform distribution with probability mass function:

$$f(x) = \frac{1}{n} \quad (2.11)$$

where  $n = 4$  is the number of possible failure magnitudes that can be assigned to the component failure (i.e., 0%-33%-66%-100%).

## 2.5.2 Python wrapper and preprocessing

The wrapper handles all the simulation runs. At the beginning of each simulation, the failure times and the failure magnitudes are sampled by the failure injection engine. Multiple failures are allowed to occur. These inputs must be passed to the TH input deck in a suitable format. OpenModelica features a python API (OMPpython) that makes this process straightforward. A flowchart describing the operations performed by the wrapper is shown in Figure 2.10.

Simulations start prior to the first failure that is sampled to allow the model setting on a steady state before the first failure takes place. When a component fails,

## 2.5 Integration of the thermal-hydraulics model within a probabilistic framework 41

Table 2.7 Input variable modified by the wrapper. The transients that involve those variables and the related components are reported as well.

Component	Input variable	Transient
Heat exchanger	$m_{He}(t)$	Secondary cycle failure
Heat exchanger	Controller state	Controller failure
Plasma	$q''_{FW}(t)$	Reactor shutdown or plasma transient
MS pumps	$n_{pump}(t)$	Pump failure
Electric heaters	$P_{heater}(t)$	Heater failure
Electric heaters	Controller state	Controller failure

the process variables associated to that component are modified in the TH model according to the failure magnitude. The safety related quantities for each component are reported in Table 2.7.

Variables of interest for the safety analysis (failure times, process variables, materials temperature, controllers state, etc.) are tracked and saved at the end of each simulation. A database of transients involving the PCS is built through this procedure. Transients can end in a system safe state (i.e., no safety variables cross the safety threshold) or in a system fault state (i.e., at least one safety variable crosses the safety threshold). The identification of the end state is carried out during the post-processing phase (Section 2.5.3).

### 2.5.3 Post-processing

The main goal of the post-processing task is to classify the transients according to their end state and to extract salient features from the database.

A single simulation takes  $\sim 1$  min to run on a standard computer (Intel® i-3® CPU @ 3.2 GHz). A wide range of possible transients can therefore be explored in a reasonable amount of time. Results from about 1400 simulations were collected and post-processed to obtain a comprehensive view of system transients. First, the transients were represented in a suitable space to perform a clustering procedure. Choosing the failure times as basis seems natural. Each transient is therefore identified by three variables: pump failure time, HX PI controller failure time and heaters controller failure time. This choice of basis intrinsically provides the order of component failures. This feature is quite advantageous because it does not require

Table 2.8 Safety limits considered during post-processing. For structural and functional materials, the safety limit has been taken equal to the melting point. For the inner and outer structural materials (Inconel 718), the lowest value in the melting temperature range (1644 K - 1700 K) has been used.

Variable	Temperature limit	Units	Reference
$T_{FW}$	3695	K	33
$T_{Str,Inner}$	1644	K	34
$T_{Be}$	1560	K	35
$T_{Str,outer}$	1644	K	34
$T_{FLiBe,boil}$	1703	K	36
$T_{FLiBe,melt}$	732	K	36

more advanced techniques of machine learning to infer the failure order from other variables (e.g., process variables) that are less correlated with the failure order. Clustering was performed with a k-means algorithm. The number of clusters  $k$  was chosen by inspection of the silhouette coefficient [78].

## 2.6 Results from the integrated deterministic and probabilistic model

Thousands of transients were simulated thanks to the low computational cost of the model. This section reports results from transients involving multiple failures only. Single failure transients are not reported because of their low interest, being extensively analysed in classic DSA.

Figure 2.11 shows the results from a single simulation. A pump failure occurs at  $t = 150$  s, reducing the pump rotating speed by 66% of its nominal value. FLiBe mass flow rate starts decreasing accordingly. The HX PI controller reduces the secondary coolant velocity (i.e., mass flow rate) to avoid excessive cooling of FLiBe and its consequent freezing. However, the regulation of  $\dot{m}_{He}$  is too slow and the FLiBe temperature drops below the setpoint temperature ( $T_{setpoint} = 800$  K), triggering the heater switch on. Heaters start delivering thermal power to FLiBe to raise its temperature above the setpoint, and switch off once this task is accomplished. The secondary coolant mass flow rate is still changing to match the new steady state set by FLiBe flow reduction, but the HX PI controller fails at  $t = 200$  s. This stops

the regulation of  $\dot{m}_{He}$ , which sets on a constant value that is too low to balance the power production in the blanket. The heat transfer coefficient in the HX on the secondary side does not allow a to exhaust  $P_{fus} = 630 \text{ MW}_{th}$  from FLiBe, which starts increasing its temperature. The transient proceeds until the melting temperature of the Inconel 718 is reached.

The transient presented in Figure 2.11 ends in a system fault state, and is classified as *hot failure*. This scenario highlights a crucial feature of the integrated deterministic and probabilistic approach. The fault end state is reached because of the failure of *multiple* components. A single failure would have not led to a fault state. Figure 2.12 exemplifies this concept. The same transient is simulated, but only the pump failure takes place. The evolution of the process variables is the same until the heaters are switched off. The HX PI manages to regulate the secondary coolant mass flow rate. FLiBe temperature at the HX outlet is kept at 800 K by the HX PI controller, while the blanket outlet temperature increases up to 1025 K due to the reduced mass flow rate. Structural and functional material temperature is kept below the safety limit, therefore the transient ends in a safe state.

Despite visual inspection of the results from single simulation can be useful to understand the system dynamics, such an approach cannot be used to process thousands of simulations (i.e., to obtain a significant domain coverage). The results from  $\sim 1400$  simulations were therefore gathered to build a database, and the aggregated data were post-processed by a k-means algorithm. The number of cluster  $k$  was chosen based on the minimum value of the silhouette coefficient, which is show in Figure 2.13. The minimum is achieved for 10 clusters.

The variables that build up the basis of the space where the results are projected were normalized prior to the clustering procedure. Figure 2.14 shows the clusters obtained when projecting the transients in the space defined by the normalized heaters failure time, normalized HX PI controller failure time, and normalized pump failure time. The clustering procedure alone does not provide relevant insights. Plotting the end states of the different clusters facilitates the interpretation of Figure 2.14, as depicted in Figure 2.15. The largest cluster, Cluster 1, comprises transient events with multiple failures that occur in close proximity to one another (mean failure time for the HX PI controller is 314 s, for heaters is 268 s, and for the pump is 66 s). As the time interval  $\Delta t$  between failures is relatively short in this cluster, the end states are primarily dependent on the order of failures. If the heater controller

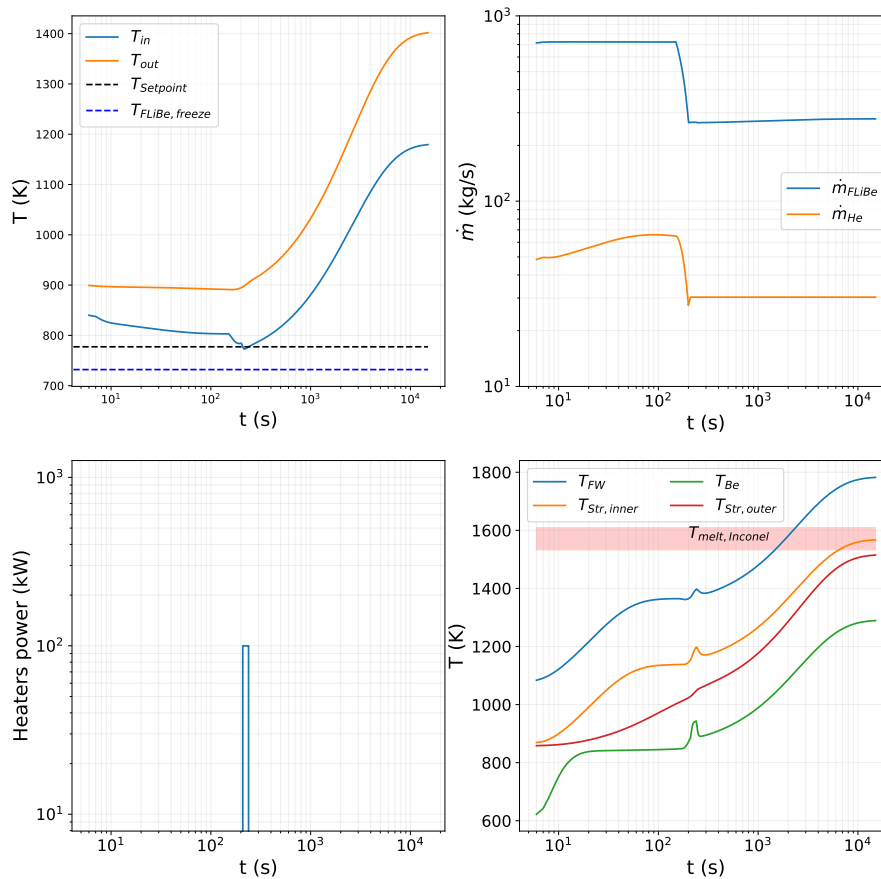


Fig. 2.11 PCS transient involving pump failure at  $t = 150$  s and HX PI controller failure at  $t = 200$  s. Relevant system variables are reported. Top left: FLiBe temperature at blanket inlet and outlet. Setpoint temperature for the heater controller and FLiBe freezing temperature are also shown. Top right: FLiBe and He mass flow rate in the heat exchanger. Bottom left: heaters power delivered to FLiBe. Heaters are triggered when FLiBe temperature drops below the setpoint to avoid FLiBe freezing. Bottom right: FW, VV, and blanket materials temperature.

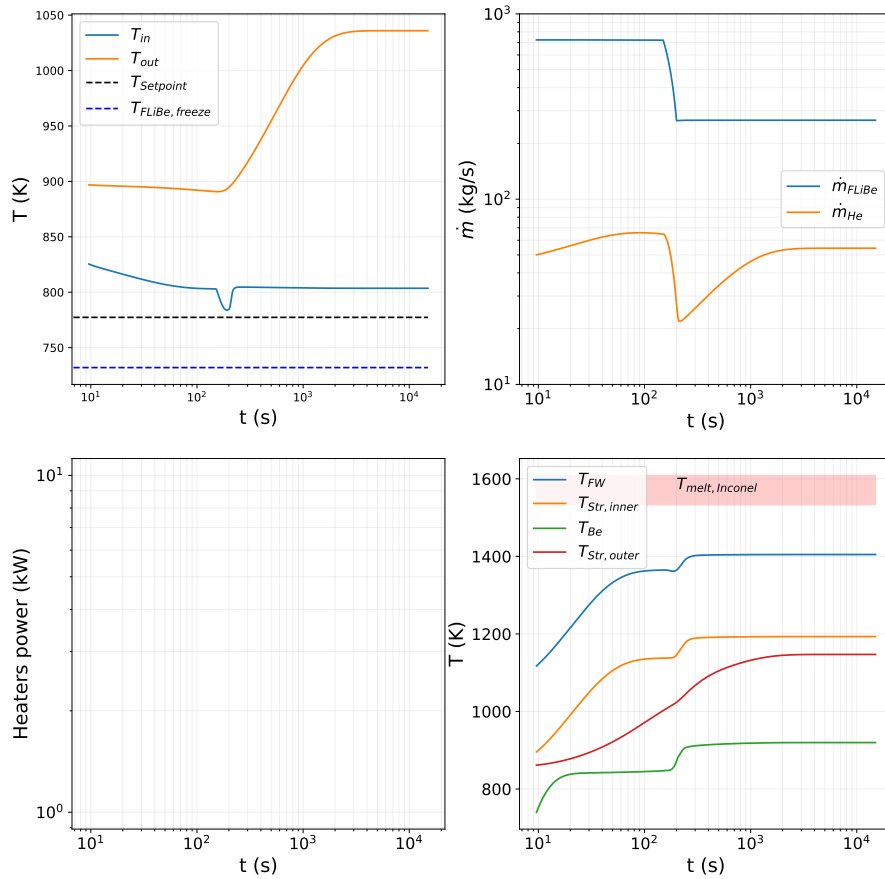


Fig. 2.12 FLiBe loop transient involving pump failure at  $t = 150$  s. The HX PI controller does not fail in this transient. Relevant system variables are reported. Top left: FLiBe temperature at blanket inlet and outlet. Setpoint temperature for the heater controller and FLiBe freezing temperature are also shown. Top right: FLiBe and He mass flow rate in the heat exchanger. Bottom left: heaters power delivered to FLiBe. Heaters are triggered when FLiBe enthalpy drops below the setpoint to avoid FLiBe freezing. Bottom right: VV and blanket materials temperature.

fails prior to the pump and PI controller, the heaters are activated. Subsequently, if the pump fails, followed by the HX controller, the heaters may produce sufficient power to prevent the freezing of FLiBe. However, if the heaters turn off when they are in use (e.g., after a pump failure), the end state will most likely be a cold failure. Alternatively, if the heaters are activated following both pump and HX controller failures, the scenario becomes comparable to the one depicted in Figure 2.11, with additional power provided by the heaters. In this case, a hot failure will rapidly ensue. Clusters 0, 2, 3, and 6 are characterized by the absence of failures. Cluster 0 contains a few points of little interest, whereas the remaining clusters provide useful information. In Clusters 2, 3, and 6, the HX controller failures take place with a large positive value of  $\Delta t$  between other failures. If the HX controller has sufficient time to regulate the helium mass flow rate subsequent to the pump and heater failures, a new steady state will be reached, and the loss of regulation for the system will be negligible. Conversely, if the HX controller fails with a large negative value of  $\Delta t$ , the regulation of helium mass flow rate will not be available, resulting in probable failure. Indeed, none of the transients in Clusters 0, 2, 3, or 6 show this feature. The other clusters consist of transients with intermediate conditions. The end state is more quantitatively dependent on  $\Delta t$  between different failures, and the end states in these clusters are mixed.

## 2.7 Discussion

The previous section presented an application of the integrated deterministic and probabilistic model to transients involving the PCS. The transient conditions were randomly generated to limit any bias on the part of the analyst. The results suggest several considerations. The first consideration relates to the model itself. Two primary goals have been achieved: the model has low computational cost and shows small errors when compared to detailed blanket models ([3, 27]). Blanket models are referred to because the blanket is the most critical component from a modeling perspective. Many other components belong to standard equipment in the nuclear field and have already been modeled (e.g., in the ThermoPower library). Modeling the VV as a 1-D component and the LIB tank as a 0-D component required strong simplifying assumptions. Section 2.4.3 demonstrated that those assumptions are acceptable from a thermal-hydraulics viewpoint. However, the




Fig. 2.13 Silhouette coefficient for different number of clusters. The silhouette coefficient is a cluster validity measure that allows to identify the optimal number of clusters to categorize the data. In that case a minimum is found for 10 clusters.




Fig. 2.14 Clusters for the transients represented on a suitable space. The failure times have been normalized before the clustering procedure. Cluster 2, 3 and 6 are located on the right area of the plot. Cluster 0 contains less than 10 elements located on the top of the plot.

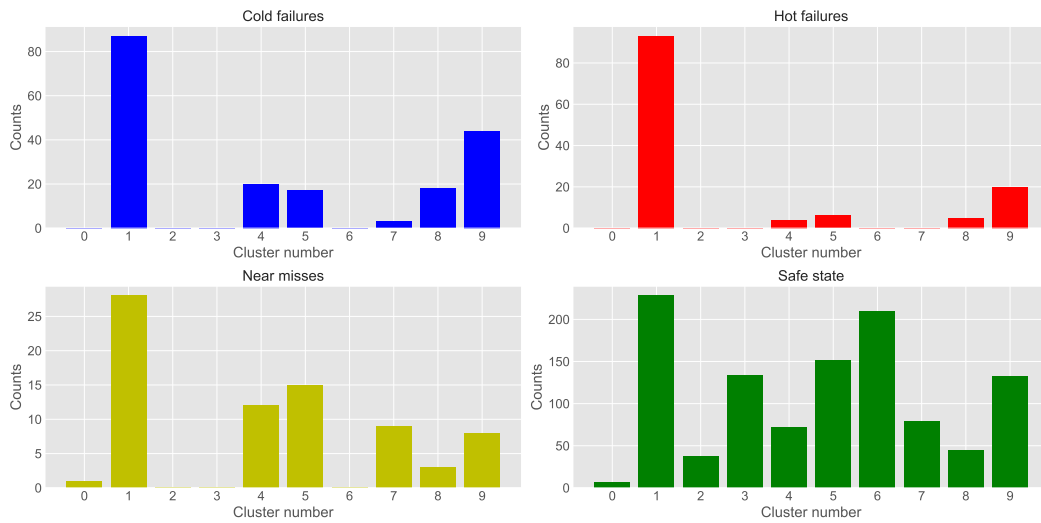


Fig. 2.15 Number of transients for each end state and cluster. Note that the total number of transients belonging to a cluster is different among the clusters. Cluster 5 is lowly populated, whereas cluster 1 shows the largest number of objects. Cluster 0, 2, 3, and 6 contain transients ending in a safe state only.

model is not entirely self-consistent because the heat transfer coefficient in the tank is provided as an external input based on component-level simulations. This approach is valid as long as the flow regime matches that of the component-level simulations, i.e., for an average FLiBe velocity above 0.5 m/s. Under extreme conditions, such as those that occur during severe accidents, these assumptions may be invalidated. This is related to the second main point of the discussion. The model has been developed to be as general as possible to investigate a wide range of possible transients, which is a primary goal in the pre-conceptual design stage. Severe accidents are a specific category of transients that require a specific treatment. Validated codes (e.g., MELCOR) might be used for the analysis of severe accidents in a mature stage of the design. For instance, a Loss of Flow Accident (LOFA) simulation requires a detailed knowledge of the system layout to compute the pressure drops for natural circulation analysis. This is clearly not possible at this design stage. The model can be extended to properly describe LOFA, but the primary goal remains the analysis of minor transients that are usually overlooked by deterministic analysis.

In this regard, coupling the model with a pre-processing and a post-processing routine becomes mandatory due to the large number of possible transients. The most significant modeling effort relates to triggering a component failure in a Modelica

model. For some components, this task is straightforward. Simulating a pump failure means switching the rotational speed of the pump from the nominal value to a new value given by the failure magnitude computed at the pre-processing stage. Instead, simulating the HX controller failure is far more complex. Additional "fictitious" control logic is required to simulate the proper behavior of the controller and to avoid internal variable loops that cannot be solved by the solver. Nevertheless, object-oriented modeling has proven to be extremely effective in handling these situations.

The last point involves post-processing. A large amount of results does not translate into useful information unless the data is correctly post-processed. Many data mining techniques can be implemented to extract relevant information. The end states were classified by comparison with safety limits, and the transients were clustered using a k-means algorithm. The interpretation of the outcomes is somewhat arbitrary. A different number of clusters would not have shown the features described in Section 2.6. For instance, choosing nine clusters would likely have merged the clusters 2 and 3, but the reasoning would have been the same. Instead, by using four clusters, a clear distinction between "safe" and "faulty" clusters would not have been possible. A similar reasoning applies to the choice of the space in which the transients were represented.

The database built with the model and the workflow presented in this section has been used to train more advanced ML algorithms. Chapter 3 deals with ML for nuclear safety applications starting from the results of Chapter 2.

# Chapter 3

## Machine learning for nuclear safety applications

Machine learning (ML) is revolutionizing most of the industrial fields. The potential of ML algorithms in the nuclear industry is enormous, for both nominal operations or accidental conditions. Yet, ML algorithms have found little room in nuclear safety if compared to other safety-critical sectors, and ML can be considered quite a novel topic in the nuclear field. As such, the goal of this Chapter is to investigate possible ML applications with an end-to-end approach, i.e. from data collection (and synthetic data production) to real-world applications. Since big data is missing in the nuclear field, and real-world applications can only be simulated, strong emphasis is put on the required workflow to successfully implement ML algorithms for nuclear safety. ARC PCS is used as test bench. Quantitative results are obtained to showcase the potentiality of ML algorithms. The scope is then broadened to propose a pathway to foster ML development for nuclear safety, leveraging on the more modern computational tools of nuclear fusion if compared to other sectors (e.g., nuclear fission).

The Chapter is structured as follows: Section 3.1 provides an overview of the main ML algorithms. Section 3.2 presents the main criticalities that are encountered when implementing ML models in nuclear safety. Applications of these algorithms and criticalities for autonomous driving, aeronautics, rotating machinery, energy production and nuclear plants are reported in Appendix B. Despite these fields are not strictly relevant for ARC-class tokamaks, Section 3.1, Section 3.2 and Appendix

B may be prodromic to understand the discussion in Section 3.3 and the results of Section 3.4. Appendix B presents also more in depth the main criticality of ML for safety applications. It includes both high-level discussions on ML and safety, and low-level discussions, more technical-related. Strategies and guidelines for the correct implementation of ML algorithms in the nuclear fields are presented in Section 3.3. Section 3.4 describes the implementation of ML algorithms for ARC safety analysis, by leveraging on the integrated model developed in Chapter 2. Specifically, a logistic regression algorithm and a multi-layer perceptron (MLP) is exploited to predict the end-state of TH transients in ARC PCS, whereas a deep neural network (DNN) is built and used to predict the evolution of process variables during the same transients. Section 3.5 summarizes the key points for a correct implementation and use of ML algorithm in nuclear safety.

### 3.1 Overview of machine learning models

Machine Learning algorithms are rapidly spreading in many different domains, such as computer vision [79, 80], automotive [81–83], finance [84, 85], and safety [86, 87]. The capability of ML algorithms to find patterns among data, execute tasks without being specifically programmed to do so and predict future trends makes ML a field with enormous potential. Among the numerous subareas of ML, Deep Learning (DL) [88] and Reinforcement Learning (RL) [89] have gained momentum in the last years. Deep learning allows to automate and merge tasks in a single model, overcoming limitations that affected previous algorithms. Such a breakthrough in ML is evident from the results from ImageNet competition (one of the most important competition in computer vision) [90], where DL algorithms greatly outperformed classical ML algorithms (e.g., shallow neural networks). Reinforcement learning breaks the ML paradigm of training the algorithm with a fixed dataset, and exploits trial-by-error interactions with the environment for training. Despite its novelty and complexity, also RL has found application in the safety field [91, 92].

The huge potential of ML models arises from the capability of the algorithm to learn to execute a task from data. A useful and widely accepted definition of learning in this context can be found in [93]: “A computer program is said to learn from experience  $E$  with respect to some class of tasks  $T$  and performance measure  $P$  if its performance at tasks in  $T$ , as measured by  $P$ , improves with experience

E”. Hence, a ML model is developed to perform a task, it requires the definition of a metric to quantify its performance and it gains experience from data. Two broad categories of machine learning approaches can be identified: supervised and unsupervised learning. RL may be considered a standalone category that does not fit in the previous categories.

### 3.1.1 Supervised algorithms

Supervised and unsupervised learning differ in the training approach, which reflects on the task that may be executed by the algorithm and, ultimately, on the application. Supervised algorithms rely on a set of paired input data  $(X, Y)$ . Here,  $X \in R^{(N,D)}$  is a set of  $N$  arrays  $x_i$  each of dimension  $D$ , where  $D$  is the number of input features.  $Y \in R^{(N,C)}$  is a set of  $N$  arrays  $y_i$  each of dimension  $C$ , where  $C$  is the number of target variables (i.e., the variables we want our model to output). The goal of the training phase is to find the optimal parameters  $\theta$  of a model  $h$  such that the performance of the model, as measured by a user-defined loss function, are maximized for the training dataset  $(X, Y)$ . The model is then used to compute the target variables for arrays  $x_j$  not belonging to the training dataset. Common applications of supervised algorithms are classification and regression.

A classification task requires the model to predict a target variable  $Y$  that is qualitative, and  $C$  identifies the number of classes. The possible classes are given by the user, and the goal of the algorithm is to assign the correct class label to each input. For 2-classes problems the output may be binary (i.e., 0 or 1), whereas for multi-class problems a score is assigned to each class and the class with the highest (or lowest) score is output. The most common algorithms for classification tasks are logistic regression [94, 95], support vector machine [96, 97], decision tree [98, 99], k-nearest neighbour [100, 101], random forest [102] and deep neural networks [88, 103]. A detailed treatment of ML models for classification can be found in [104]. A logistic regression algorithm<sup>1</sup> and a MLP are used in this Chapter to predict the end state of ARC PCS failures.

A regression task requires the model to predict one or more target variables  $Y$  that are quantitative. These may be, for instance, process variables as a function of

---

<sup>1</sup>Note that a logistic regression algorithm is not by itself a classifier. However, it can be exploited as a classifier by choosing a proper cutoff value, and assigning a class to a sample based on the probability of that sample to belong to a class. This is what is done in Section 3.4.1.

time. The goal of the model is to estimate the value of  $Y$  for unseen input  $X$  that may lie between training input values (interpolation) or outside the region defined by the training set (extrapolation). The most common algorithms for regression tasks are linear regression [105], support vector regression [106], gaussian process regression [107] and deep neural networks [108]. A DNN is exploited in this Chapter to predict the evolution of process variables during ARC PCS transients.

### 3.1.2 Unsupervised algorithms

Unsupervised algorithms work on a set of input data  $X$ , without any target variable  $Y$ . Differently from supervised algorithms, unsupervised algorithms do not predict any output variable  $Y$ . Their goal is instead to find correlations or patterns by exploring the feature space defined by  $X$ . Common tasks for unsupervised learning are clustering and anomaly detection.

A clustering task consists in finding clusters in the features space. The number of clusters and the membership of any input to a specific cluster is not known a priori. This makes clustering quite different from classification, where the true labels of any input (during the training phase) are known. Most common algorithms for clustering are K-means [109], hierarchical clustering methods [110], density-based clustering [111], gaussian mixture models [112] and deep neural networks coupled with clustering algorithms [113]. In-depth reviews of clustering methods can be found in [114] and [115]. A K-means was indeed used to find cluster for the transients analysed in Chapter 2.

Anomaly detection aims to find data, or patterns in data, that are somehow deviating from the normal behaviour. A simple example of an individual anomaly is a point in the feature space that lies far from high-density regions. Such a point may be classified as anomalous with respect to the other data points. Commons algorithms for anomaly detection are self-organizing maps [116], K-means [117], and one-class support vector machine [118]. The reader can found a detailed review of anomaly detection algorithms in [119]. Anomaly detection algorithms are not exploited in this work because they are usually used when dealing with real-world data (not synthetic data) from sensor measurements.

### 3.1.3 Reinforcement learning

Reinforcement learning is based on the idea of letting the model (the agent) interact with the environment. There is not a fixed training dataset as in supervised learning. The agent and the environment start in a generic system state  $s_i$ , and the actions of the agent may modify the system state to a new state  $s_j$ . The training phase of the agent is therefore an exploration of the environment within a feedback loop. The performances of the agent are measured by a reward function, which is used by the agent itself to find the optimal, or nearly optimal, set of actions (policy) to execute the given tasks. These features make RL algorithms extremely flexible and suitable for complex tasks.

## 3.2 Criticality and challenges in ML for nuclear safety applications

Several issues must be tackled and overcome before ML algorithms can be introduced in nuclear safety. Some of them are presented in this section, namely data availability, domain shift, false positives/false negatives, and evaluation metric. A possible pathway to solve these issues is then presented in Section 3.3. Additional topics that are relevant for the discussion, but are not directly applicable to the case under study (ARC PCS) are interpretability of ML models, lack of self-awareness, adversarial error, and formal verification. These are discussed in Appendix B.

### 3.2.1 Data availability

Data availability is a mandatory requirement for ML, since ML algorithms are data-driven tools. Lack of data may be attributed to the approach to the deterministic safety analysis in the nuclear field. In fact, data from nuclear accidents is quite limited (only three major accidents occurred in the last 50 years). Safety analysis mostly relies on computer code simulations, and few cases are analysed in a DSA (i.e., DBA and BDBA). Uncertainty quantification is performed by best-estimate codes, which allow to quantify the uncertainty with a prescribed confidence, still minimizing the amount of simulations needed. This approach is clearly highly effective for the purpose of a DSA, but it is not suitable for data-driven methods. Furthermore,

the domain coverage is strongly limited to the DBA and BDBA investigated. The nuclear safety field is therefore characterized by small datasets, which are also highly unbalanced. While a small dataset is a showstopper for ML algorithms, unbalanced datasets are more subtle. An unbalanced dataset contains data belonging to a small region of the domain of interest. Algorithms trained with these datasets may show high accuracy, yet failing to fulfil their task when the input does not fall within the highly populated region of the domain. This introduces the need for a comprehensive domain coverage, which can be achieved by integrating deterministic and probabilistic tools as described in Section 2.5. Domain coverage is also strictly related to the domain shift issue.

### 3.2.2 Domain shift

A major assumption in ML is that the training data belongs to the same probability distribution of operational data (i.e., the real-world data that the model will process when deployed). For instance, the training dataset may contain data representing the nominal flow regime of the reactor coolant. In accidental conditions the flow regime may drastically change, transitioning from turbulent flow to natural circulation for a Loss Of Flow Accident (LOFA). A model trained only with nominal condition data would fail to provide the desired output during an accidental scenario.

Out-of-distribution data may appear even if care is taken in considering the accidental scenario domain because accidents are rare events. The available data may not be representative of the whole spectrum of possible accidental scenarios that may occur. Using a ML model for a safety-relevant task during an accidental scenario (e.g., predicting the plant end-state) may lead to incorrect predictions if the accidental scenario differs from the ones used to train the model. Such a behaviour may be experienced even if the training data and the operational data look similar, at least superficially, as demonstrated [120]. The generalization capability (i.e., domain shift robustness) of the model is therefore a leading issue for safety applications.

### 3.2.3 False positives and false negatives

Wrong predictions from ML algorithm are extremely critical in nuclear safety. Let's suppose to exploit a classifier to predict the end state of unanticipated transients,

whose task is to identify the transients that would end in a fault state. A false positive is the misclassification of a safe transient, which is labelled as accidental scenario. The reactor might be shutdown as safety measure, leading to an economic losses but no safety consequences. A false negative is much more dangerous, because safety measures might not be undertaken even though the transient is an incipient accidental scenario. ML algorithms for nuclear safety must therefore strongly penalize misclassification by exploiting a suitable metric.

### 3.2.4 Evaluation metrics

#### Accuracy

The choice of the evaluation metric is crucial for a correct understanding and evaluation of a ML model. The accuracy is the most common method to estimate the performance of a classifier. Accuracy is defined as the ratio of correct classifications to all the classifications. For a binary classifier it simply reads as:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (3.1)$$

where TP is the amount of true positive, TN of true negative, FP of false positive, and FN of false negative. For nuclear safety applications, the accuracy might not be a suitable evaluation metric. Since accidents are rare events, we might expect the available dataset to be strongly unbalanced on the "safe" side (i.e., most of the cases in the dataset refer to nominal operations, or to unexpected transients that end in a safe state). Let's assume that 99 cases are safe, and the model is developed to correctly identify (i.e., classify) accidental scenarios where the system ends in a fail state. A classifier trained with this dataset would show a very high accuracy (close to 100%), yet failing to correctly identify the cases of interest. Furthermore, the accuracy is not threshold-invariant. A crucial choice when developing classifiers is the threshold at which a sample is classified in a given class. For instance, one of the feature of interest in nuclear safety might be the coolant temperature, and the threshold for classifying a transient as safe can be set close to the boiling point of the coolant. However, the choice of *how close to the boiling point?* is somehow arbitrary, and this choice strongly affects the accuracy of the classifier. In this regard, a threshold-invariant evaluation metric might be more suitable for safety applications.

### Receiver Operating Characteristic curve

The Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC) can be a better evaluation metric than accuracy, though it can fail as well if applied to the wrong problems [121]. The AUC of the ROC provides richer information than the accuracy, and is indeed more suitable in decision making contexts [122]. An ROC curve is shown in Figure 3.1. The ROC curve is drawn on a TPR (True Positive Rate) vs FPR (False Positive Rate) plot, where TPR and FPR are defined as follows:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (3.2)$$

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (3.3)$$

Maximizing the TPR means reducing the amount of positive samples classified as FN. This is particularly important when dealing with end state prediction. Say the fault end state is the positive class, then a FN is the most dangerous prediction because the algorithm is considering the transient as a safe transient, while the real end state would be fault. While the interpretation of the ROC, and of the AUC, is more complex than the accuracy, the ROC AUC results in a more effective tool. A random classifier would have a ROC curve as a 45° line between (0,0) and (1,1). The ROC curve of a good classifier lies above the random classifier curve, meaning that the probability of classifying a TP correctly is higher than the probability of wrongly classifying a TN as FP. Each point in the ROC identifies an operating point of the classifier with a given threshold. Changing the threshold shifts the operating point along the curve, and allows designers and analysts to set a suitable threshold for their purposes. This introduces qualitatively the concept of cost of misclassification. As explained previously, the misclassification of an incipient accidental scenario as a safe transient can have tremendous consequences. Therefore, a classifier developed to detect incipient accidental scenario should operate close to  $P(\text{TP} = 1)$ , accepting a larger number of false positives.




Fig. 3.1 Receiver Operating Characteristic curve. The red dotted line is the ROC curve of a random classifier (the probability of getting a TP or a FP is the same). The ROC curve of a good classifier (Classifier A) and a worse classifier (Classifier B) are depicted. A perfect classifier achieves a unitary TP rate for an FPR = 0 (green dot).

### Precision-Recall curve

The Precision-Recall curve results in a better evaluation metric than the ROC AUC for binary classification tasks when the dataset is imbalanced [123]. The recall (REC) is defined as the TPR (previously defined), i.e.:

$$\text{REC} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (3.4)$$

while the precision (PREC) is defined as:

$$\text{PREC} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (3.5)$$

A high PREC means that the model can correctly identify TP still minimizing the FP (i.e., false alarms). Both REC and PREC are independent of the total number of TN, which is a desired feature for highly imbalanced dataset. Suppose to have 100 samples, with 99 samples being negative. A well-trained classifier would achieve a  $\text{REC} = \text{TPR} \sim 1$ . The ROC curve would provide a misleading information to the analyst because  $\text{FPR} \rightarrow 0$  for  $\text{TN} \gg \text{FP}$ . Such a result may lead to the wrong conclusion that the model is correctly identifying TP without evident drawbacks (i.e., few FPs). Conversely, PREC is independent of TN, and as the FP increases PREC

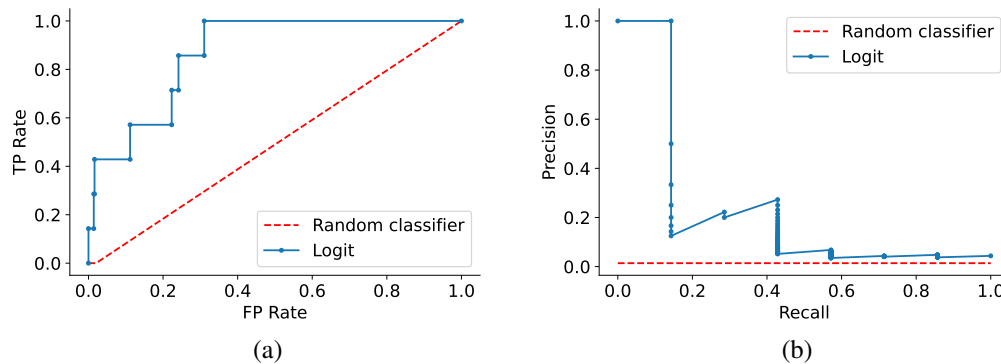


Fig. 3.2 Comparison between ROC curve and Precision-Recall curve for a strongly imbalanced dataset (99% of the samples belonging to the same class). While the ROC curve shows a good behaviour of the classifier, the Precision-Recall curve highlights the extremely low precision of the model when REC increases.

$\rightarrow 0$  (i.e., while the classifier can correctly identify accidents, many safe scenarios would be classified as accidents). Hence, both the metrics inform on the capability of the classifier to correctly identify accidental scenarios, but the Precision-Recall curve provides the additional information about the cost of a high performance in identifying TP, i.e. the fraction of negative scenarios wrongly classified as TP. Figure 3.2 shows a comparison between the ROC curve and the Precision-Recall curve for a strongly imbalanced dataset with 99% of the samples belonging to the same class. The ROC curve depicts good capabilities of the model (the TPR reaches 1 at FPR  $\sim 0.3$ ), but this happens at the expense of a very low precision.

### 3.3 A machine learning pipeline for nuclear safety applications

Section 3.3 highlighted many criticalities of ML when applied to safety-critical fields. A community effort is needed to overcome some of them in the nuclear (fusion) field, especially those related to data availability and domain shift. The former is indeed a mandatory requirement to exploit data driven models, while the latter can easily jeopardize a ML model when deployed on a real system. The goal of this section is to propose a framework to foster ML algorithms in nuclear safety. Section 3.4 can be regarded as a small-scale application of the framework described hereunder.

### 3.3.1 Solve data availability and domain coverage issues

Data can be retrieved from experimental setups, operating systems and plants, and simulations (in which case we refer to as *synthetic* data). Experimental setups are limited in number and are extremely case-specific, making their data extremely useful to test ML algorithms on real-world systems, but providing a very narrow domain coverage. Data from operating plants are clearly absent for FPP (while being present to some extent for experimental tokamaks) and limited for NPP. Synthetic data looks as the most promising option for short-term development of ML algorithm for nuclear safety.

Synthetic data for nuclear safety can be produced by deterministic codes that have been extensively validated (e.g. MELCOR, RELAP, TRACE). However, these codes usually lacks features that are crucial for integration with ML algorithms, specifically:

- Simple pre-processing. Huge amount of data can be produced only if simulations are automated, requiring simple pre-processing of the input files;
- Low computational cost. Even if simulations are automated, the computational cost must be kept low, otherwise databases cannot be built in reasonably short times;
- Simple post-processing. Results post-processing is fundamental to prepare suitable input files for the ML algorithms.
- Code flexibility (whenever the source code is available). Integration of these codes with probabilistic routines and ML algorithms might require low-level modifications.

Besides the computational cost, the open source model usually overcomes all these limitations. Open source codes benefit from active engagement of the community, are continuously improved and maintained, and put more emphasis on transparency, documentation and user experience. Pre-processing and post-processing are much more efficient and integrated with modern computational tools. The code is flexible because it is exploited by a broad community where users might have different needs. Therefore, open source codes presents many advantage over closed-group codes.

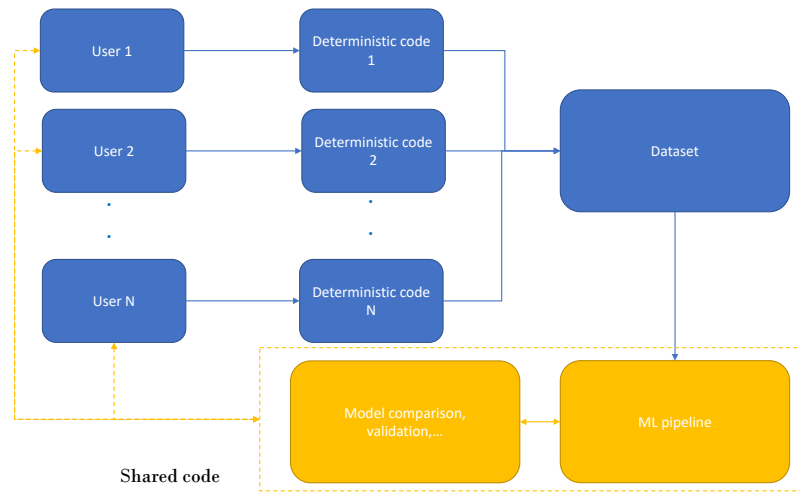


Fig. 3.3 Flowchart showing the collaborative research approach. Each user contributes to a shared database, which solves both the lack of data and the domain coverage issue. Users can also use different codes, provided they output coherent results with the other codes.

In this framework, open source is a subsection of a broader topic: collaborative research. Collaborative research allows to multiply the computational resources by increasing the number of users. Furthermore, each user exploits the codes for their purposes, increasing the domain coverage of the data (Figure 3.3). Nevertheless, such an approach requires a coherent and uniform framework to work.

### 3.3.2 A common framework to build database for nuclear safety applications

A critical step in Figure 3.3 is appending results to the same database from different user. This clearly requires rules and good practices, i.e. a common, formal framework. The FAIR (Findability, Accessibility, Interoperability, and Reusability) guiding principles have been formalized to allow stakeholders (universities, researchers, privates, industries) to share and reuse, in an effective way, data, algorithms and tools. This approach has been successfully exploited in other fields (e.g., computer vision), increasing the available data for data-driven models. These principles have been followed (when possible) for the development of the ML models in this work.

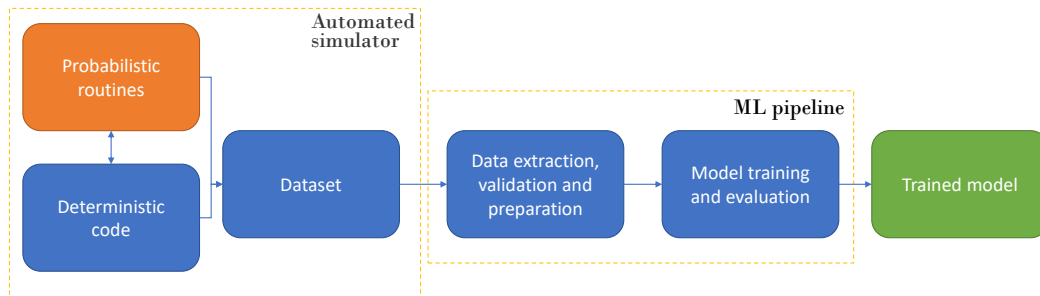


Fig. 3.4 A simple workflow to integrate ML algorithms in nuclear safety analysis. Coupling a probabilistic routine with a deterministic code increases the domain coverage. Once the database is built, a ML pipeline exploits the huge amount of data to train ML models.

Figure 3.4 shows a possible workflow for the integration of ML algorithms in nuclear safety analysis. The automated simulator has been already introduced and tested in Chapter 2, and the dataset built with the simulator is exploited to train the models presented in this Chapter. The ML pipeline is the main topic of Section 3.4.

## 3.4 Applications of machine learning algorithms to ARC PCS

This section showcases two different applications of ML for ARC safety assessment, one involving a classification task, the other involving a regression task.

A logistic regression model is developed and used to predict the end state of the PCS transients. Prediction of the end state requires a binary classifier (i.e., the end state can be safe or fault), which is trained on the dataset for the PCS transients (Chapter 2). Logistic regression exploits a linear decision boundary, which is shown to perform well, yet not achieving satisfactory performances for this application. A fully connected (FC) neural network (NN) is therefore developed to take advantage of the non-linear decision boundary to separate fault end states from safe end state.

A long short-term memory (LSTM) model is instead developed to predict the evolution of process variables from synthetic measurements in the PCS [124]. Synthetic measurements are here intended as signals from the sensors that were placed in the thermal-hydraulic model (Section 2.4). The LSTM is chosen among the possible algorithms for its capability to retain temporal information when dealing with time series (note that the process variables evolution is just time series from a data analysis standpoint), which is crucial for this kind of application.

### 3.4.1 Prediction of the end states of the PCS

The dataset contains 1435 samples, which were divided according to the 70/30 ratio in a training and in a test set. 1120 (78%) end states are safe, while the remaining (22%) are fault end states. The dataset can be considered slightly imbalanced. Hence, the advantages of the ROC AUC over the accuracy are shown first. Then, results using the recall as metric are presented and compared to the ROC AUC. A  $k$ -fold cross-validation was exploited during the training phase with  $k = 10$ . Grid search was used to find the optimal hyperparameters of the MLP, and the best model was used in the test phase. The Scikit-learn Python library was chosen for this application [125].

Each sample has 3 features, namely the pump failure time, PI controller failure time, and heater controller failure time. The features are scaled to have zero mean and unit variance. The coolant temperature is used to classify the end state according to the procedure described in Chapter 2. Cold failures and hot failure are both considered as a single class (i.e., a fault end state). Near misses are considered within the safe state class. In this way the problem is reduced to a binary classification problem, which eases the interpretation of the results.

Figure 3.5 shows the ROC for two logistic regression models (here called Logit 1 and Logit 2 for simplicity) using different scoring metrics (i.e., Logit 1 is trained using accuracy as scoring metrics, while Logit 2 is trained using the ROC AUC as scoring metrics). Since the goal of this algorithm is the prediction of the end states, the algorithm must be trained to detect as most TP possible (where a TP in this case is a fault end state correctly predicted), and minimizing FN (a fault end state wrongly predicted as safe end state), which means maximizing the TPR at expenses of a higher FPR. Figure 3.5 clarifies the concepts discussed in Section 3.2.4, i.e. the

better performance of the ROC AUC with respect to the accuracy. Logit 1 shows a higher accuracy on the test set ( $\mu = 0.86$  and  $\sigma = \pm 0.04$ ) but Logit 2 (accuracy  $\mu = 0.84$  and  $\sigma = \pm 0.04$ ) is more suited for this task. Using the AUC as scoring metric allows to drastically reduce the FN, at expenses of a lower accuracy in TN detection (compare the two confusion matrices in Figure 3.6a and Figure 3.6b for the same detection threshold of 0.47). Furthermore, ROC curves can be exploited to select the suitable threshold, which in this case can be set at 0.18 to avoid TP to be classified as FN (compare the TPR for a threshold of 0.47 and a threshold of 0.18 in Figure 3.5). From a practical standpoint this means accepting a slightly larger number of false alarms to minimize the probability that incipient accidents get undetected.

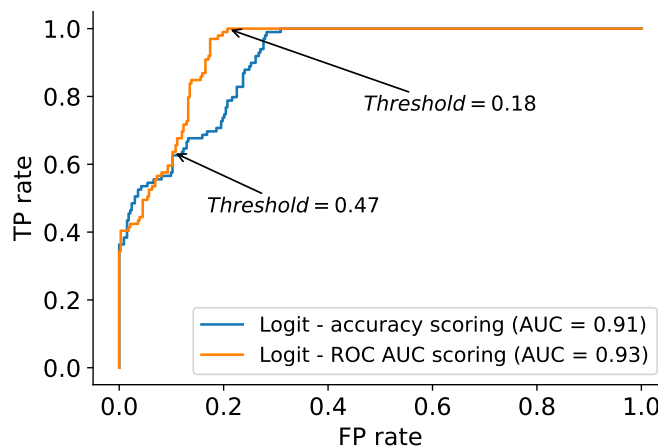


Fig. 3.5 ROC for two logistic regression algorithms using different scoring metrics. Results refer to the test set. The blue curve shows the ROC of the Logit trained to maximize the accuracy, whereas the orange curve shows the Logit trained to maximize the AUC. Maximizing the AUC leads to an algorithm that identifies better TP (incipient accidents), which is the goal of such an algorithm applied to prediction of end states. Each point on the curves are associated to a classification threshold. Two thresholds (one for low TP rate and one for high TP rate) are depicted.

A MLP algorithm is trained on the same dataset and the performances are compared to the logistic regression models. The best MLP model identified by the grid search has a hidden layer  $64 \times 2$ , exploits a *tanh* activation function, regularization strength  $\lambda = 10^{-3}$ , and the LBFGS solver. The MLP can learn a non-linear decision boundary, which grants improved classification capabilities to linear decision boundaries (Figure 3.7 and Figure 3.8). The accuracy achieved by

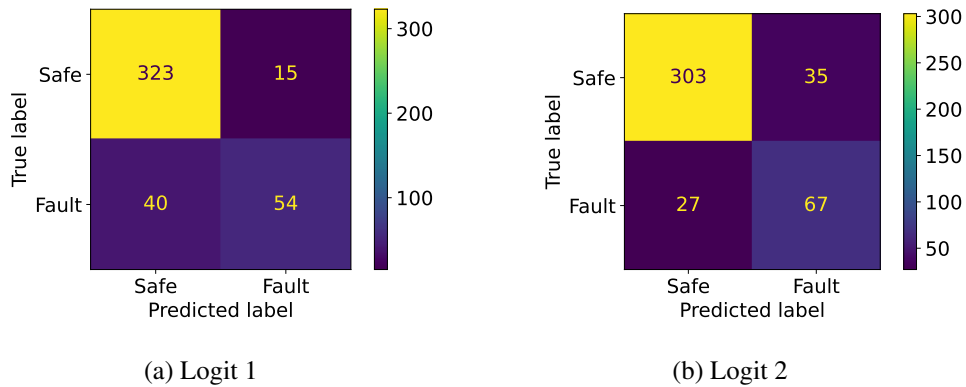


Fig. 3.6 Confusion matrix of Logit 1 (a) and Logit 2 (b) for the test set. Values on the matrix diagonal represent TN and TP. The lower left corner shows the FN, while the upper right corner FP. Despite the lower accuracy, Logit 2 achieves appreciably higher performances in executing the task (i.e., identify TP and minimize FN).

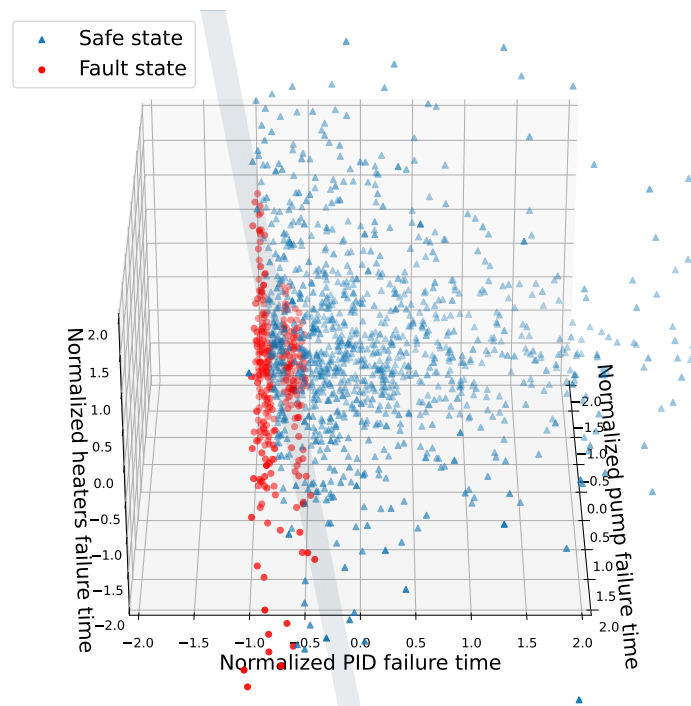


Fig. 3.7 Safe end states and fault end states as a function of the normalized failure times of the pump, PI HX controller and heater controller. The (linear) decision boundary of the logistic regression model is depicted (light blue plane). A good separation of the two classes is achieved, even though a linear boundary has limited capabilities.

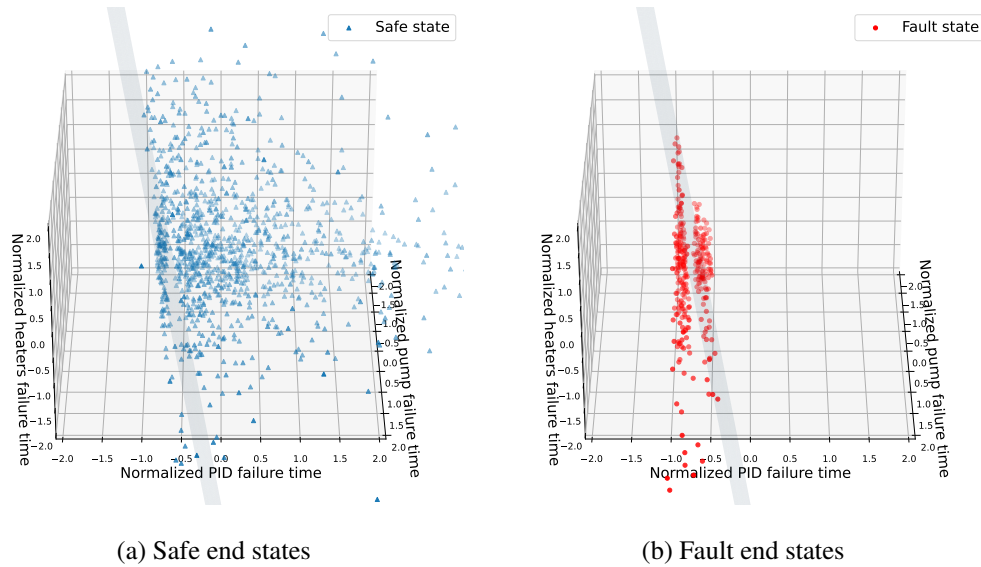


Fig. 3.8 Safe end states (a) and fault end states (b) represented separately to highlight the separating capability of the linear decision boundary (light blue plane). Most of the safe end states fall on one side of the boundary, but many fault end states lay on the same side of the boundary, limiting the performance of the model.

the MLP on the test set is indeed close to 100% (99.5%), with an AUC = 1 (Figure 3.9). Figure 3.10 shows the Precision-Recall plot. The different interpretation on models performance is evident when comparing Figure 3.9 with Figure 3.10. The precision of the logistic regression model quickly drops when  $REC > 0.2$  on the test set. Achieving  $REC \rightarrow 1$  comes at the expense of approximately  $PREC \sim 0.6$ . The MLP instead achieves very high REC values with just a small drop on PREC. A single scenario is misclassified as a TP, which would lead to a false alarm during system operations (Figure 3.11).

The MLP outperforms logistic regression models for this application. The computational time of a single prediction is negligible for practical applications (below 1 ms) for both the algorithms. Hence, end state prediction can be performed in real-time, provided that component failures can be detected during system operations. This does not seem a leading issue for ARC PCS because the large thermal inertia of the LIB provides enough time (i.e., from minutes to hours) for fault detection before the system fails. The recall works as a better performance metric than the accuracy and the ROC AUC for this particular case, even though this conclusion could be extended to many safety applications where the consequences of a misclassification can be dramatic. It should be noted that the dataset used for training and test the

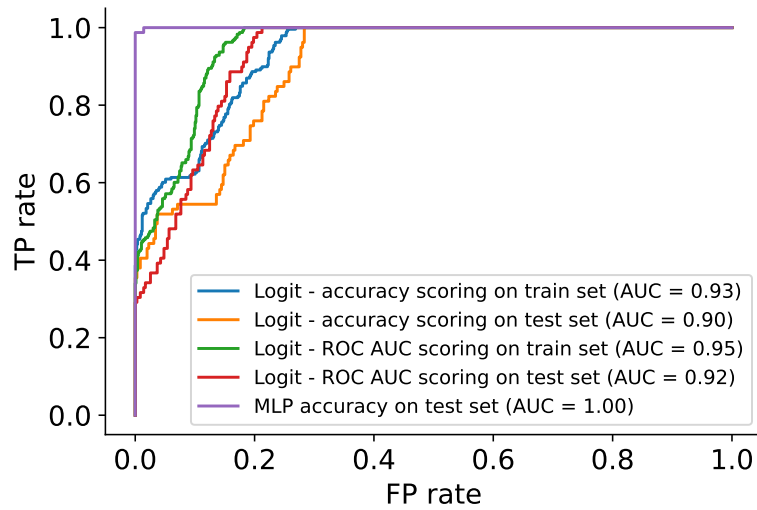


Fig. 3.9 Comparison between logistic regression models and MLP for both training and test set, using the ROC AUC as scoring metric. The MLP achieves AUC = 1 (i.e., perfect classifier) both on the train and the test set (train set not shown for simplicity).

algorithms is limited, although overfitting issues should have been prevented by the implementation of a k-fold cross-validation. Few outliers are present in the dataset, so the robustness of the model on edge cases was not assessed. Nevertheless, the dense region of data corresponds to multiple failures taking place at close times, i.e. those scenarios that are more likely to treat system safety. The outliers corresponds to scenarios where the failures occur with a long delay between one and another. Therefore, outliers should not be a major issue for this specific application.

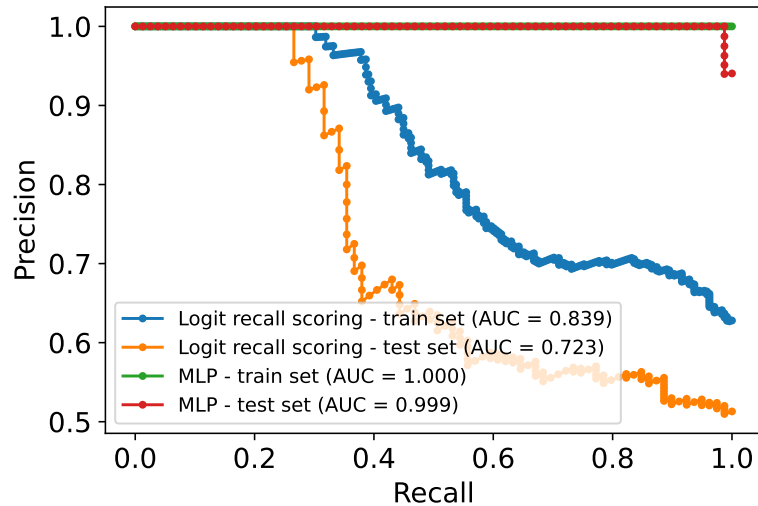


Fig. 3.10 Comparison between logistic regression models and MLP for both training and test set, using REC as scoring metric. The MLP strongly outperforms the logistic regression model, achieving a high precision even when  $REC \rightarrow 1$ . A random classifier would have achieved  $PREC = 0.21$  with this dataset.




Fig. 3.11 Confusion matrix for the MLP for the test set. The MLP achieves almost a perfect classification accuracy, with a just a FN over the test set.

### 3.4.2 Prediction of safety variable evolution during PCS transients

The capability to predict the evolution of safety variables is crucial during accidental scenarios. This task is much more complex than the previous one because of the inclusion of the time variable. The number of features (i.e., physical variables) is higher ( $\geq 9$ ), and each feature is tracked all along the simulation. The target variable(s) is a continuous variable, and the prediction must be performed at each timestep. The nature of the problem suggests that an LSTM is the most suitable algorithm (a detailed description of the LSTM is provided in AppendixB).

The previous task was quite straightforward from the model viewpoint. The choice of the input variables was natural, and exploiting just three input variables demonstrated the efficacy of the algorithms. For this task the choice is more arbitrary. The selection procedure of the input deserves a specific discussion, in line of the goals presented in Section 3.3.

#### Input data selection

The TH model developed in Chapter 2 can output any possible TH variable of interest for safety analyses. This is a strong advantage when training a ML model because different sets of input parameters can be tested to find the optimal combination that maximizes the accuracy (or any other metric is used) of the model. However, one must consider that a ML model when deployed on a real system would not have access to all the data that are made available by a computational model. Ideally, the ML model should be able to perform its task with the lowest possible set of input. The input data used to train the model has been chosen according to the following requirements:

- The set of input would not allow a deterministic code to compute the target variable quickly. Stated in another way, the ML model is able to predict the target variable much faster than a deterministic model. Clearly, if a deterministic model could compute the same target variable in less time, the ML model would be useless;
- There are no variables in the input set such that the target variable can be computed directly from a single input variable (e.g., if the blanket outlet

Variable
Heat flux on the FW surface
FLiBe mass flow rate
Time derivative of FLiBe mass flow rate
Inlet channel enthalpy
HX PI controller setpoint enthalpy
Channel pressure drop
Secondary coolant mass flow rate
Time derivative of the secondary coolant mass flow rate
HX inlet temperature of the secondary coolant
Heaters thermal power delivered to FLiBe
FW temperature
VV inner structural material temperature
Beryllium temperature
VV outer structural material temperature

Table 3.1 Thermal-hydraulics variables selected as input for the DNN model.

coolant enthalpy is the target variable, the blanket outlet coolant temperature must not be present in the input set);

- The input set is the smallest possible before the ML model starts failing its prediction. The threshold at which the ML model starts failing is arbitrary, but its choice does not lead to any loss of generality of this approach.

The case under study here can be summarized as follows. In a real-world situation these points would correspond to an accidental scenario where the data received by the instrumentation system is limited, and one or more safety variables cannot be measured directly. Also, no deterministic models can compute the safety variables quick enough to allow for a real-time monitoring.

The dataset contains  $5 \times 10^3$  samples, each one being a transient that lasts for  $10^4$  s. The variables selected for the input are listed in Table 3.1. Each row of the dataset is a snapshot of the system state at time  $t$  (Figure 3.12). Among those transients, 670 end in a fault end state. The dataset is therefore unbalanced, which makes the task of predicting the time-to-failure (TTF) harder. The dataset was split in  $4.5 \times 10^3$  training samples and  $5.0 \times 10^2$  test samples.

Two target variables are selected: the FLiBe blanket outlet enthalpy and the TTF. The FLiBe blanket outlet enthalpy is one of the most important safety variables

t[s]	T_inlet_gas[K]	dT_inlet_gas	dp[Pa]	dwdt_ch	Q_surf_FW[W]	P_heaters[W]	dwdt_gas	w_gas[kg/s]	w_flibe[kg/s]
0.0	555.000000	742.669354	97875.880325	-4.662564e-01	-70000.0	100000.0	-0.000000	10.000000	1.570000
2.0	713.356184	2.308791	118951.405447	1.433363e-08	-70000.0	100000.0	-0.000000	10.000000	1.469259
4.0	617.524525	-382.756448	118951.405558	2.611029e-15	-70000.0	100000.0	0.074473	19.098841	1.469259
6.0	438.700540	-8.681834	99809.180490	2.434108e-02	-70000.0	0.0	0.007967	48.489036	1.529229
8.0	435.423363	-0.063361	94621.997812	1.637331e-03	-70000.0	0.0	-0.000019	49.567265	1.547336
...	...	...	...	...	...	...	...	...	...
9992.0	355.159450	-0.000005	30160.018458	-1.202248e-11	-70000.0	0.0	0.000000	117.636684	0.739627
9994.0	355.159440	-0.000005	30160.018458	-1.001591e-11	-70000.0	0.0	0.000000	117.636684	0.739627
9996.0	355.159431	-0.000005	30160.018458	-8.013249e-12	-70000.0	0.0	0.000000	117.636684	0.739627
9998.0	355.159422	-0.000005	30160.018458	-6.007979e-12	-70000.0	0.0	0.000000	117.636684	0.739627
10000.0	355.159412	-0.000005	30160.018458	-3.998791e-12	-70000.0	0.0	0.000000	117.636684	0.739627

Fig. 3.12 Pandas dataframe of a random sample from the dataset showing the input variables (the input set has been already reduced from the initial set reported in Table (3.1)). Each row is a snapshot of the system at time  $t$ .

because FLiBe reaches the highest temperature at the blanket outlet (neglecting possible hot spots in recirculating regions), where boiling can occur. The TTF is here defined as the time before the system fails, according to the safety limits already defined in Section 3.4.1. In formula:

$$\begin{cases} \text{TTF}(t) = t_{s,f} - t & t \leq t_{s,f} \\ \text{TTF}(t) = 0 & t > t_{s,f} \end{cases} \quad (3.6)$$

where  $t_{s,f}$  is the system failure time. If instead the system does not fail, the TTF is set to an arbitrary large value, which corresponds to the simulation time:

$$\text{TTF}(t) = 10^4 \quad \forall t \quad (3.7)$$

This solution provides some improvements in terms of predicting capability of the model. The TTF is therefore a helpful variable that can be exploited by the operators to undertake safety actions by knowing the predicted time before the system completely fails. Figure 3.13 depicts the input of the smallest set that allows the model to perform well and the output of the model.

### ML pipeline and DNN model

An end-to-end ML pipeline has been built as described in Section 3.3. The simulator consists in the TH model coupled with the failure injection engine. The simulations

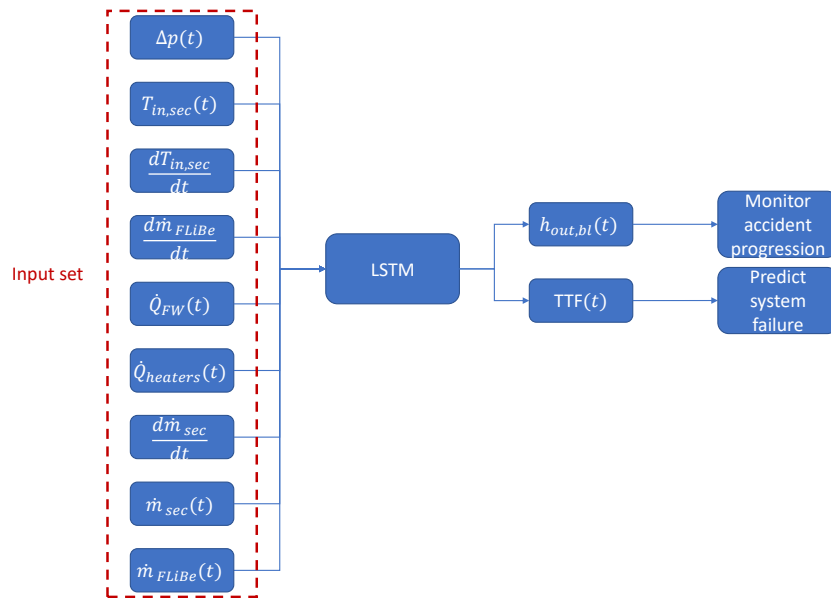


Fig. 3.13 Input set provided to the LSTM and output (target variables) returned by the model. The input is passed at each timestep and the target variables are computed at the same timestep.  $\Delta p$ : pressure increase across the pump;  $T_{in,sec}$ : inlet temperature (at the HX) of the secondary coolant;  $\frac{dT_{in,sec}}{dt}$ : time derivative of the inlet temperature of the secondary coolant;  $\frac{d\dot{m}_{FLiBe}}{dt}$ : time derivative of the FLiBe mass flow rate;  $\dot{Q}_{FW}$ : thermal power to the FW;  $\dot{Q}_{heaters}$ : thermal power delivered by the electric heaters;  $\frac{d\dot{m}_{sec}}{dt}$ : time derivative of the secondary coolant mass flow rate;  $\dot{m}_{sec}$ : secondary coolant mass flow rate;  $\dot{m}_{FLiBe}$ : FLiBe mass flow rate;  $h_{out,bl}$ : FLiBe enthalpy at the blanket outlet.

Variable	Value
Epochs	20
Learning rate	$8 \times 10^{-3}$
Number of features	Variable
Number of classes	1
Number of LSTM layers	2
LSTM hidden size	64
FC2 hidden size	8

Table 3.2 DNN model data.

are completely automated by taking advantage of OMPython. The results of the simulations are stored in a database as .csv files. The output of a single simulation can be quite heavy, so a `variableFilter` was used to store only the variables of interest (those reported in Table 3.1). The data is then loaded with a user-defined dataloader and splitted in training, validation, and test sets by PyTorch DataLoader.  $4.5 \times 10^3$  samples are used for training and  $5.0 \times 10^2$  samples are used for test. Each sample belongs solely to the training set or to the test set (i.e., the test scenarios are never seen by the model during training/validation). The input vector is scaled by means of a `StandardScaler`, while the output vector is normalized by a `MinMaxScaler`. These steps result in a completely automated pipeline, from the TH model preprocessing to the ML model preprocessing.

The LSTM is built exploiting PyTorch `nn.Module`, which allow a greater control of the model architecture with respect to a simple `nn.LSTM`. The architecture of the DNN is the following: INPUT  $\rightarrow$  LSTM  $\rightarrow$  RELU  $\rightarrow$  FC1  $\rightarrow$  RELU  $\rightarrow$  FC2. Details for each layer are reported in Table 3.2. The mean squared error (MSE) was used as loss function. The MSE is indeed a robust loss function for regression problems that is commonly used in these applications. Larger deviations from the real values are strongly penalized by the MSE due to the presence of squared terms, which is a desired feature since the model is applied to safety variables.

The LSTM takes as input the 9 input variable at each timestep, and returns the FLiBe blanket outlet enthalpy and the TTF at the same timesteps (i.e., the model does not forecast the target variables at time  $t + 1$  given the system state at time  $t$ ).

Results of the DNN models are presented in the following. First, the DNN is trained and tested on accidental scenarios with fixed failure magnitude (0.4 for the

pump and 0 for the HX PI controller). Then, randomly generated failure magnitudes are considered, increasing the variability of transients in the dataset.

### **DNN models with fixed failure magnitudes in the dataset**

Two different DNNs were trained to predict the tank outlet FLiBe enthalpy and the time-to-failure. The distribution of the failure times is shown in Figure 3.14. The mean failure times for pump, PI controller and heaters are  $\mu_{pump} = 66\text{ s}$ ,  $\mu_{PI} = 315\text{ s}$  and  $\mu_{heaters} = 262\text{ s}$  respectively. Figure 3.15 shows the distribution of secondary coolant inlet temperature and mass flow rate, heat power delivered by the heaters, and the time derivative of the secondary coolant temperature at the inlet at two different times,  $t = 100\text{ s}$  and  $t = 1000\text{ s}$ . The two times have been chosen such that for  $t = 100\text{ s}$  only some failures have already taken place, whereas for  $t = 1000\text{ s}$  most of the failures have taken place. Results of the DNN models are shown in Figure 3.16. A fault scenario from the test set (#36) is chosen to showcase the prediction of the blanket outlet enthalpy and TTF. Scenario #36 has never been seen by the model during training. Note that in the following the blanket outlet enthalpy is converted into the outlet temperature for sake of simplicity.

The variables used as input for the DNN (Table 3.1) have been gradually reduced. The DNN is still able to predict correctly the two target variables even if the channel inlet enthalpy, the time derivative of the channel inlet enthalpy, and the secondary coolant outlet temperature are removed. Note that in this set of input variables there are no variables associated to the FLiBe temperature (or enthalpy) in the blanket. There is no easy way for the model to predict the outlet blanket temperature. Figure 3.17 shows the temperature distribution at 5000 s (i.e., in the mid of the transients, after the failures have occurred) over the dataset. FLiBe outlet temperature assumes values in the range 900 K - 1300 K, with a similar number of occurrences for  $T > 1050\text{ K}$ . The target variable is therefore spread over a wide range of possible parameters, requiring the ML model to learn to predict very different values depending on the input variables.

The three failures (pump, PI controller of the HX and PI controller of the heaters) that were randomly sampled for the scenario #36 are shown in 3.16. The model is not aware of these failures (i.e., the failure times are not passed as input), but it manages to correctly predict the evolution of the outlet temperature following each failure. This means that the model can be used to track the evolution of a safety variable

even if direct measurements are not available, and that an alarm can be triggered when the predicted safety variable overcomes the safety threshold.

The loss during the first epoch is shown in Figure 3.18. The model learns fast during the first epoch, and the loss quickly drops to  $\sim 10^{-7}$  after 150 iterations. The small pikes of the loss function may be related to the high learning rate chosen ( $8 \cdot 10^{-3}$ ), which was increased from the initial value of  $10^{-3}$  to speed up the training. However, no major oscillations are found in the running loss, meaning that the optimization of the model proceeded smoothly.

The MSE at the end of the training phase is  $2.3 \cdot 10^{-8}$ , while the MSE on the test set is  $5.4 \cdot 10^{-6}$ . The difference between the training and the test MSE is due to a possible overfitting of the model on the training data. However,  $\text{MSE} = 5.4 \cdot 10^{-6}$  is well below a maximum acceptable value for safety applications (e.g.,  $\text{MSE} = 0.1\%$  can be considered a maximum acceptable error on the outlet temperature, meaning that the absolute error is around 1 K).

The prediction of the TTF is reported in Figure 3.19. The TTF is an arbitrary defined variable with no direct physical meaning, nor direct correlation with one of the input variables (i.e., no one-to-one relationship with the input variables), but it is strictly correlated to the FLiBe enthalpy because the system is considered to fail when the enthalpy crosses a safety threshold. Nevertheless, the DNN model predicts the TTF with a reasonably low error. The MSE at the end of the training phase is  $\text{MSE} = 1.4 \cdot 10^{-6}$ , while on the test set the error is  $\text{MSE} = 5 \cdot 10^{-6}$ . Note that the prediction of the TTF is more challenging than the prediction of a process variable. The evolution of a process variable depends on the history of the system only (i.e.,




Fig. 3.14 Failure time distributions within the whole dataset. The distribution is exponential as explained in Section 2.5.1.

$T_{out}(t')$  depends on the values of the other system variables at  $t < t'$ . Instead,  $TTF(t')$  depends by definition on an event that occurs in the future for  $t' < t_{s,f}$ , and  $t_{s,f}$  is not available to the model as input. Therefore, at any time  $t' < t_{s,f}$  the model has to predict when the system failure will occur to compute the TTF. This is a remarkable results because it requires the model to correctly learn to predict system failure on long time intervals by knowing only the values of physical variables at previous times.

The training phase required 10 epochs to reach such a low error, i.e. 2x epochs if compared to the prediction of  $T_{out}$ . This can be explained by the definition of TTF itself. The same set of features (i.e., physical variables) at time  $t$  can be found in a safe transient and in a fault transient. Hence, the prediction of the TTF cannot be based on the instantaneous values of the system variables, but also on the history of these variables. The prediction is therefore more complex than  $T_{out}$  (even though also  $T_{out}$  prediction is somehow linked to the history of the input variables), and the model leverages on the capability of the LSTM to retain temporal information to perform its task.

Figure 3.20 shows the prediction for an outlier (i.e., a scenario whose features strongly differ from those of the scenarios in the training dataset). The failures time are 1000 s, 3300 s and 5100 s for the heaters PI controller, HX PI controller, and pump, respectively. These values can be considered to belong to a very different distribution from the failure times in the training dataset ( $\mu_{ft,pump} = 66$  s,  $\mu_{ft,HX} = 315$  s,  $\mu_{ft,heaters} = 262$  s). The model fails to correctly predict the evolution of the outlet blanket temperature. The temperature predicted at the beginning of the simulation is underestimated by 90 K. Then, it increases after the spurious switch on of the heaters, and the predicted value increases as well. However, the initial gap between the real and the predicted temperature remains all along the transient. The failure of the HX controller does not affect FLiBe temperature (the pump is still working, and the mass flow rate can be adjusted). Once the pump fails, the temperature sharply drops. The model tries to predict the temperature decrease, but it underestimates the temperature drop.

This simple case demonstrates the weakness of ML models when trained on a dataset that is not fully representative of the possible scenarios that may happen during operations. In such a situation, the model would have completely failed to perform its task, with possible consequences on the plant safety.




Fig. 3.15 Input distribution at  $t = 100\text{s}$  (a) and  $t = 1000\text{s}$  (b) for fixed failure magnitudes. As the components fail, the input variables spread over a larger range. Only the most relevant input are shown.




Fig. 3.16 Prediction of accident evolution (zoom on the initial transient on the left). The failure times and the safety threshold temperature (1250 K) are depicted. Note that the model does not receive as input the failure times, but it manages to successfully predict the strong variation in temperature when a failure takes place.




Fig. 3.17 Outlet blanket temperature distribution at 5000 s over the dataset.



Fig. 3.18 Running loss on the training set as a function of the iterations during the first epoch. The loss quickly decreases down to  $\sim 10^{-7}$ .




Fig. 3.19 Prediction of accident evolution. The failure times and the system failure time are depicted. Note that the model does not receive as input the failure times, but it manages to successfully predict the TTF.




Fig. 3.20 Prediction of the temperature evolution for an outlier scenario. The model fails to correctly predict the temperature evolution because it has been trained on a very different dataset.

### DNN models with randomly sampled failure magnitudes

The failure magnitude of the HX PI controller and the pump rotational speed were randomly picked in the intervals  $(0,1]$  and  $(0,0.8)$  respectively, where 0 stands for "no failure" and 1 is complete failure of the component. The mean failure times for pump, PI controller and heaters are the same as the previous case with fixed failure magnitude. Figure 3.21 shows the distribution of secondary coolant inlet temperature and mass flow rate, heat power delivered by the heaters, and the time derivative of the secondary coolant temperature at the inlet at two different times,  $t = 100$ s and  $t = 1000$ s. The two times have been chosen such that for  $t = 100$ s only some failures have already taken place, whereas for  $t = 1000$ s most of the failures have taken place. It should be noted that the larger variability in transients evolution due to a variable failure magnitude lead to a higher number of accidental scenarios in the dataset (1200 vs 670 for the fixed failure magnitudes case).

When the HX PI controller fails completely, the mass flow rate of the secondary coolant is reduced to zero. Including the failure magnitude in the simulations greatly increases the variability of the possible transients, making the predictive task much harder for the model. Figure 3.22 shows the temperature distribution at 5000 s over the dataset. The values are much more spread than those in Figure 3.17 due to the presence of the failure magnitudes in the input set. The secondary coolant mass flow rate can now increase or decrease, introducing also the cold failures in the dataset. On the one hand, this increases the number of transients that can end in a safe state by chance. If the heaters switch on spuriously, and the HX PI controller fails and increases the secondary coolant mass flow rate, then the system may reach a steady state within the minimum and the maximum allowable temperatures. On the other hand, the possible trajectories of the system are much more, and the ML model needs to be more flexible to correctly predict the evolution of the outlet blanket temperature.

Results of the DNN model are shown in Figures 3.23 to 3.25. Three scenarios (#5, #19 and #21) are chosen to showcase completely different evolutions of the system, and the corresponding prediction of the model. These three scenarios have never been seen by the model during training.

Figure 3.23 shows a transient that ends in a safe state. The DNN successfully predicts  $T_{out}$  during the initial transient. A small deviation from the real value

occurs after the HX PI failure. The model has been trained on a dataset where most of the times  $T_{out}$  changes significantly after the HX PI controller. Therefore, it overshoots the target variable soon after the failure, but it quickly reduces the predicted temperature to the real value.

Figure 3.24 shows instead a transient that ends in a cold failure.  $T_{out}$  increases after the pump failure and the heaters controller failure, but it then starts decreasing following the HX controller failure. Again, a small overshoot can be depicted after the HX controller failure. However, the DNN successfully predicts the evolution of  $T_{out}$  from the initial transient to the system failure, with a negligible delay between the predicted value and the real value.

Figure 3.25 shows a transient that ends in a hot failure. Remarkably, the DNN manages to correctly predict  $T_{out}$  evolution despite very few transients reach such a high  $T_{out}$  (Figure 3.22), and with such a fast dynamic. The pump and the heaters controller fail closely, but this does not affect the temperature prediction. Also, no deviations on the predicted temperature are present in this case. The predicted temperature crosses the safety threshold with a delay of few seconds with respect to the real value.

The results from Figures 3.23 to 3.25 can be generalized to the whole dataset. The first failure is usually handled by the remaining working components (e.g., the increase in thermal power to the FLiBe due to a spurious switch on of the heaters is balanced by an increase of the secondary mass flow rate, which keeps the FLiBe temperature within the nominal values). However, the variation of FLiBe temperature differs from case to case. In scenario #5, a small change occurs after the first and the third failure, with most of the variation taking place between the second and the third failure. Conversely, the temperature increases steadily and sharply since the first failure in scenario #21. If the model had learnt just to detect failures (which is alone a quite remarkable achievement), it would have been impossible to correctly predict the evolution of the FLiBe enthalpy due to the many different dynamic interactions that take place among the system components during these transients.

A similar reasoning applies if the model had learnt to predict the evolution of the FLiBe enthalpy based on the evolution of some input variable only. A similar change in a process variable can be observed in extremely different scenarios. Let's choose the secondary coolant mass flow rate  $\dot{m}_{sec}$  as example.  $\dot{m}_{sec}$  can be regulated by the (properly working) PI controller in the HX, for instance following the failure of

another component and the subsequent increase or decrease of the FLiBe temperature. However, the same variation of  $\dot{m}_{sec}$  can be caused by the failure of the PI controller. Other components might or might not operate to face this sudden change triggered by the PI controller failure. Hence, the information of the change of the input variable alone is not enough to correctly predict the evolution of the target variables.

We can conclude that the model has correctly learnt (at least from a conceptual point of view) to detect the component failures and to take into consideration the interactions between the input variables.

The prediction of the TTF is more complex in this case due to the presence of different failure magnitudes. The model still achieves a low error during the training phase. However, the dataset is strongly unbalanced, with 15% of accidental scenarios and 85% of safe transients. TTF for safe transients is predicted correctly, but the error associated to accidental scenarios is quite large for some scenarios. The TTF for scenarios #5, #19 and #21 is shown in Figures 3.26 to 3.28.

The model correctly predict the safe transient in scenario #5. As for the case without failure magnitude, the capability of the model to correctly predict the TTF for a safe transient even though three failures occur can be attributed to a proper learning of the model. The TTF prediction of the model is not solely based on the concept of failures, but it relies also on the physical variables passed as features to the model.

Scenario #19 can be considered a case where the model fails to predict the accident progression. This may sound strange since the temperature prediction was correct for scenario #19 (Figure 3.24). However, cold failures are a small subset of the fault scenarios, which are by themselves a small subset of all the possible scenarios in the dataset. Hence, the wrong prediction of the TTF for the cold failure can be attributed to the very small domain coverage of cold failures.

The TTF is instead correctly predicted in scenario #21. The initial prediction is close to  $10^4$  s because the majority of the scenarios in the dataset present this value (i.e., they are safe transients). At each failure the TTF rapidly changes, until the model correctly predicts the TTF from the third component failure onward.

## 3.5 Discussion

This Chapter targeted two key objectives: proposing and testing a framework to foster ML in the nuclear safety field, and develop and test innovative tools for more robust safety analyses.

Best practices and experience in different safety-critical fields have been reviewed to propose an effective framework. Despite being quite standard in other fields, such an approach has not been formalized yet in the nuclear industry. This is clearly well beyond the scope of this work, but the application of the proposed framework on a small case study gave promising results. Three qualitative considerations can be made based on the results of Section 3.4.

The prediction of the end states successfully achieved a 100% accuracy (and an acceptable Precision-Recall curve) with a single FN and a single FP in over 450 scenarios tested (Figure 3.11). This should not surprise. A binary classification task is among the easiest tasks a ML algorithm can perform. By selecting suitable features (the failure times) the classification works perfectly. The state-of-the-art ML algorithms are indeed developed for much hardest tasks, and this application is illuminating for its simplicity, yet providing valuable predictions during accidental scenarios. Stated in another way, ML algorithms have been exploited for so many years in other safety-critical industries that very advanced tools are already available for nuclear safety, shifting the effort from developing an effective ML model to building a framework where this model can be successfully implemented. One may argue that using the failure times as input features is unrealistic and ease the task of the ML model. However, the DNN model showed the capability to detect failures even if the failure time is not passed as input feature. Failure detection is therefore a realistic task, which can be performed by classical method or ML algorithms, and can provide the failure times to a ML model for end state prediction (see also the literature on this topic [126–130]). The DNN successfully predicted the evolution of a safety variable (FLiBe outlet blanket enthalpy) and the time before the system reaches a fault condition. This second variable is of particular interest because it provides a real-time measure of the available time before the system reaches a complete failure.

Besides the encouraging results achieved in these two applications, care must be paid when deploying ML models in safety-critical industries. The pipeline described

in Figure 3.4 allowed to explore a wide range of possible unexpected transients, building a large database of different scenarios. Nevertheless, the domain of possible transients is still huge, and the capability of the ML models to perform properly in a real-world situation is far from being demonstrated. The outlier case showed how the model can completely fail to predict the target variable evolution in a transient that differs quite a lot from those in the training dataset. An in-depth discussion on ML criticalities is presented in the Appendix B. In short, ML models cannot be employed at the present time to take critical actions, but can serve as an additional tool for decision making in critical situations, such as during accidental scenarios.

The resources required to develop robust ML algorithms are well beyond the possibility of single researchers and small teams. The computational time of both running integrated probabilistic and deterministic simulations and training deep ML models is extremely long (i.e., days). A community effort looks as the most promising solution for building more general databases and improving the robustness of the ML models.

Lastly, results from the ML models are provided without uncertainty quantification (UQ). The relevance of UQ is well documented in nuclear safety [131, 132], and ML models cannot be exempt by this procedure. Conversely, UQ must be addressed even more carefully because of the black box nature of ML models.




Fig. 3.21 Input distribution at  $t = 100$ s (a) and  $t = 1000$ s (b) for randomly sampled failure magnitudes. As the components fail, the input variables spread over a larger range. Only the most relevant input are shown.


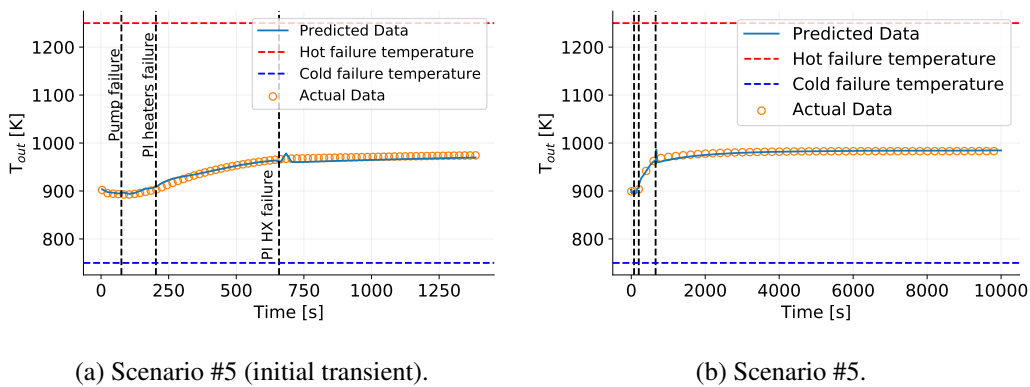


Fig. 3.22 Outlet blanket temperature distribution at 5000 s over the dataset that includes the failure magnitudes in the input set. The temperature spans over the whole possible range for FLiBe, from the melting temperature ( $T_{melt,FLiBe} = 732$  K) to the boiling temperature ( $T_{boil,FLiBe} = 1703$  K).



(a) Scenario #5 (initial transient).

(b) Scenario #5.

Fig. 3.23 Transient ending in a safe state. The temperature increases smoothly up to a steady-state value. The predicted value slightly deviates from the real value immediately after each failure, but the model manages to improve the prediction soon after.

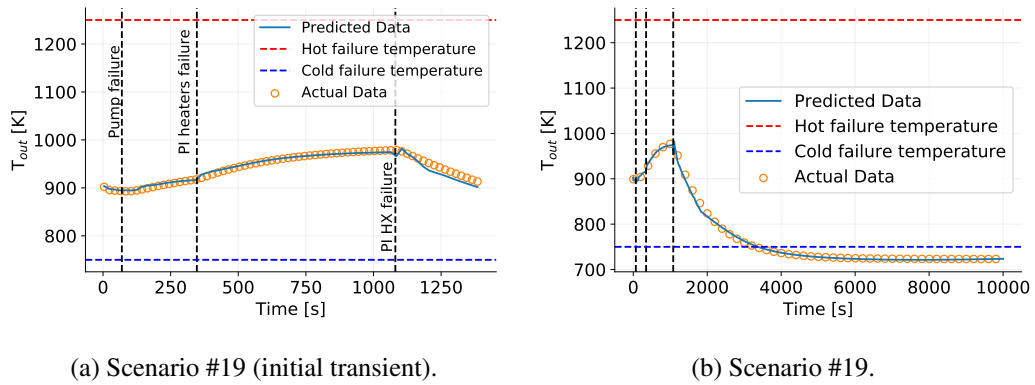


Fig. 3.24 Transient ending in a cold failure. The temperature first increases, and then starts decreasing after the HX controller failure. The predicted value slightly deviates from the real value immediately after each failure, but the model manages to improve the prediction soon after.






Fig. 3.25 Transient ending in a hot failure. The temperature first increases slowly, and starts increasing sharply after the HX controller failure.




(a) Scenario #5 (initial transient).




(b) Scenario #5.

Fig. 3.26 Transient ending in a safe state. The TTF is  $10^4$  s by definition. The predicted TTF oscillates during the initial transient due to the presence of three component failures. However, once the last failure occurs, the model correctly predict a value close to  $10^4$  s, meaning that the system is not expected to fail.



(a) Scenario #19 (initial transient).



(b) Scenario #19.

Fig. 3.27 Transient ending in a cold failure. The model predicts that a system failure will happen after the third component failure, but the TTF prediction is largely overestimated.




Fig. 3.28 Transient ending in a hot failure. The TTF oscillates during the initial transient. Nevertheless, the model correctly predict the TTF after the third failure occurs, with a small underestimation.

# Chapter 4

## ARC fuel cycle

### 4.1 Tritium self-sufficiency, tritium safety and fuel cycle

Tritium self-sufficiency refers to the ability of an FPP to breed, process, and use its own tritium (T) fuel while producing sufficient excess tritium to maintain the desired reserve inventory (in the event that some part of the fuel cycle breaks) and to contribute to  $I_{\text{startup}}$  for subsequent FPPs. Tritium self-sufficiency of FPPs [30, 133] is expected to be a significant challenge in FPP design and operation. Abdou et al. [30] showed that tritium self-sufficiency may be extremely challenging or even impossible for proposed DEMONstration power plant (DEMO) designs. For a fueling efficiency  $\eta_f \approx 25\%$  (obtained by extrapolating from DIII-D studies [134] to DEMO) and a burn fraction  $f_b = 0.36\%$  (the expected achievable value in ITER, assuming two pellet injectors with a maximum fueling rate of  $120 \text{ Pa}\cdot\text{m}^3/\text{s}$  [135]), the TBR required ( $\text{TBR}_r$ ) to achieve tritium self-sufficiency in DEMO will be 1.82. This is far above its maximum achievable TBR ( $\text{TBR}_a$ ), which is expected to range from 1.1-1.2 depending on the chosen blanket design [136]. Therefore, improvements in fuel cycle design are mandatory to lower  $\text{TBR}_r$  below 1.1 and achieve tritium self-sufficiency [30]. Relying on external tritium supplies is not a sustainable option, as global supply is highly constrained. Most of the world's tritium supply is produced by CANDU reactor operations [133, 137]. The available tritium supply from CANDU reactors is expected to peak at  $\approx 27 \text{ kg}$  in 2027 [30]. The peak may be higher (up to  $\approx 40 \text{ kg}$ ) if additional tritium is made available by the operating and planned heavy

water reactors (HWR) outside Canada (e.g., in Romania and the Republic of Korea [133, 137]).

No other significant tritium sources are expected to come online in the near future [133]. ITER plans to start D-T operations in the late 2030s and will consume approximately 12.3 kg of tritium during its planned experimental campaigns, leaving  $\approx 15$  kg of theoretically available tritium for the startup of FPPs that begin operation after 2050 [133, 137]. Therefore, minimizing start-up inventory and ensuring the robustness and efficiency of tritium breeding and extraction systems will be of paramount importance for the first non-ITER D-T fusion devices.

Tritium safety is strictly connected with tritium self-sufficiency. Minimizing the amount of tritium in components and system is the leading target for tritium safety. However, this task can be accomplished only if safety analysis is carried out in parallel to tritium self-sufficiency analysis. The first step is to quantify the tritium inventories in the components during nominal operations, and then minimizing the inventories without affecting tritium self-sufficiency. Luckily, tritium self-sufficiency and tritium safety are synergic goals: lower tritium inventories in the fuel cycle are favourable for tritium self-sufficiency and reduce the radiological hazard.

In this regard, the analysis presented in this Chapter aims at (1) assessing tritium self-sufficiency in ARC; (2) quantifying tritium inventories and tritium flows between fuel cycle components, with and without the implementation of Direct Internal Recycling (DIR); and (3) providing quantitative design targets for ARC that achieve a low doubling time  $t_d$  and a low start-up inventory  $I_{\text{startup}}$ , strengthening the safety features of ARC-class tokamaks.

A time-dependent, system-level models of the ARC-class fuel cycle is developed in MATLAB Simulink ®. The systems-level model provides a suitable tradeoff between model accuracy and computational cost. It enables to assess the effect of individual design parameters on tritium self-sufficiency and start-up inventory. The results show that tritium self-sufficiency can be achieved in ARC for a narrow but achievable window of operating parameters: tritium burn efficiency  $\text{TBE} > 0.5\%$ <sup>1</sup>, availability factor  $\text{AF} > 70\%$ , direct internal recycling fraction  $f_{\text{DIR}} > 0.3$ , and tritium processing time  $t_p < 4$  h. A major advantage of ARC versus large FPPs

<sup>1</sup>From a purely mathematics point of view,  $\text{TBE} = \eta_f f_b$ , where  $\eta_f$  is the fueling efficiency and  $f_b$  is the tritium burn fraction. However, Whyte et al. point out that the TBE results in a more coherent figure of merit when describing continuous fueling processes in magnetic confinement fusion [138], so the TBE is used in this work.

is found in the small size of ARC ( $P_{\text{fus}} = 525 \text{ MW}_{\text{th}}$ ), which translates to a lower start-up inventory.

The Chapter is organised as follows: Section 4.2 presents the mathematical and computational models for the fuel cycle. Section 4.3 shows the results from the sensitivity analysis on selected design and operational parameters for the ARC-class tokamak. Section 4.4 compares ARC results to those from DEMO, discusses key topics like the optimal TBE, tritium trapping, and doubling time, and proposes an R&D path to reduce uncertainties associated with this analysis methodology.

## 4.2 Methodology

### 4.2.1 ARC fuel cycle layout

ARC is in a conceptual design phase, and the specific parameters of its fuel cycle technologies and components are not finalized. A high-level layout of the expected fuel cycle is shown in Figure 4.1. The fuel cycle is divided into two sub-cycles: the outer fuel cycle (OFC) and the inner fuel cycle (IFC). The OFC encompasses most of the components and subsystems that provide tritium to the IFC. In the ARC-class FPP, the OFC includes the FLiBe LIB and FLiBe coolant channels in the vacuum vessel (VV)<sup>2</sup>, the associated tritium extraction systems (TES), the heat exchanger(s) (HX), and the divertor. The IFC processes the tritium from the OFC and from the gas stream exhausted by the divertor vacuum pumps, which are considered within the IFC. The IFC supplies the required quantities of deuterium and tritium to the fueling system. The following subsections provide a detailed description of fuel cycle components and subsystems.

#### Outer fuel cycle

The outer fuel cycle includes the divertor, FW/VV channels, LIB, TES, and HX(s). It is important to stress that in an ARC-class tokamak, the FLiBe acts as both coolant and tritium carrier. This results in a key distinction between the ARC fuel cycle

---

<sup>2</sup>The first wall (FW) is placed on the external surface of the inner VV wall in ARC [3]. In this analysis we consider the first wall (FW), the VV and the VV channels as a single component called FW/VV channels.

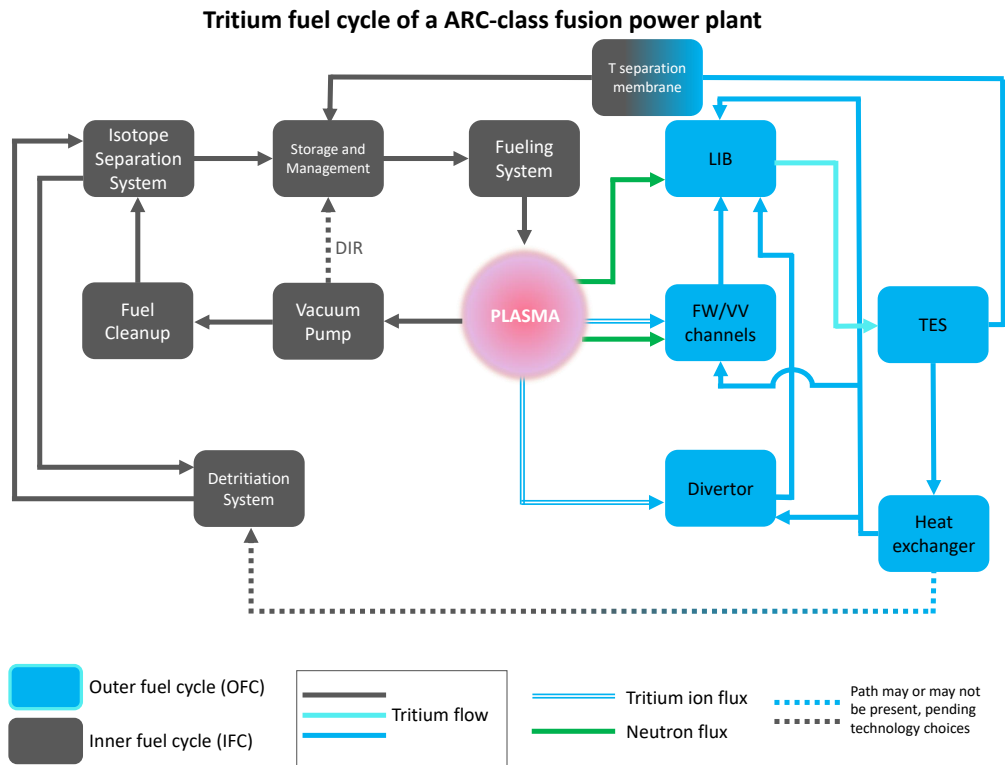


Fig. 4.1 The block diagram shows the high-level fuel cycle of an ARC-class FPP. The outer fuel cycle (OFC) breeds tritium via  $\text{Li}(n,T)$  reactions in the blanket and coolant channels. This tritium is passed to the inner fuel cycle (IFC), where it is processed into fuel. Only the components that have been modeled in this analysis are shown in the figure. The DIR system is represented by a dotted path to indicate that it may or may not be present in the IFC. The dotted path from the heat exchanger to the detritiation system is present if the chosen heat exchanger has an active tritium removal system (e.g. a double-walled heat exchanger with purge gas flow). The inclusion of DIR and the decision to include T extraction in the heat exchanger (HX) are examples of ARC design decisions which have yet to be finalized.

model and the DEMO fuel cycle model used in [30]: in ARC, the coolant line and tritium carrier line are part of a single OFC system. There is no coolant purification system (CPS) in the ARC model because the tritium extraction system (TES) fulfills this role in ARC. The layout of ARC fuel cycle is shown in Figure 4.1.

Tritium is bred from the  ${}^6\text{Li}$  and  ${}^7\text{Li}$  in the LIB and in any FLiBe coolant channels located in the VV. The VV channels receive cold ( $\sim 800$  K) FLiBe directly from the HX [3], where it is heated by the plasma and discharged into the blanket tank. Divertor channels are also fed with cold FLiBe from the HX that is discharged into the blanket tank. FLiBe in the blanket tank (which carries tritium) exits via outlet(s) at the top of the tank that is connected to the TES, which extracts tritium from the FLiBe and sends it to the storage and management system. Additional components, such as a tritium permeation membrane, are used on the pathway between the TES and tritium storage to separate tritium from other gases in the FLiBe.

The TES has not a unitary efficiency, so some fraction of tritium is inevitably carried to the HX by the FLiBe. Generally speaking, the HX is part of the power cycle, not the fuel cycle. However, since a small fraction of tritium passes through the HX, the HX must be considered in this fuel cycle model. Here it is assumed that the HX has an active tritium removal system that actively impacts tritium mass transport through the fuel cycle. An example is a double-walled HX with sweep gas flowing between the primary and the secondary coolant channels [139]). Tritium that permeates the wall of the primary coolant channel and is recovered by the sweep gas and sent to the detritiation system. Any remaining tritium in FLiBe flows back to the LIB, VV channels and divertor channels with the cold FLiBe.

The fuel cycle model should also account for unburnt tritium ions from the plasma, which act as an additional tritium source in the OFC. Unburnt tritium ions can impinge on the plasma-facing surfaces of the divertor and the FW. From there, the tritium ions might diffuse through the structural materials and permeate through the coolant, be trapped in the material, or cycle back into the plasma (note that also the opposite is true: tritium can diffuse from the coolant to the structural materials, depending on its concentration in the two media).

Certain elements of the OFC that are likely to be present in future ARC designs are not modeled here. These include a redox control unit (to monitor and control FLiBe chemistry) and a CPS for the secondary coolant loop (which will inevitably uptake a small amount of tritium from the HX). From a systems-level view, leaving

these components out of this iteration of the model should not significantly change results. The redox control unit has no direct fuel cycle functions and does not change tritium mass transfer through the fuel cycle. The secondary CPS collects tritium that reaches the secondary coolant in the HX and sends it to the detritiation system; this is already accounted for by modeling tritium losses in the HX component and the tritium flow from the HX to the detritiation system. To summarize, tritium transport in the OFC is mostly dictated by FLiBe flow through the system.

### Inner fuel cycle

The inner fuel cycle comprises the vacuum pumps and the systems for fuel clean-up, isotope separation, detritiation, tritium storage and management, and fueling. Vacuum pumps pump out tritium from the vacuum chamber, the boundaries of which are defined by the FW and the VV structures (the blanket lies outside the vacuum boundary in an ARC-class tokamak). Unburnt tritium is removed from the vacuum chamber by the vacuum pumps. Note that some fraction of the unburnt tritium will implant in the plasma facing components (PFC); in this model, this fraction is assumed as  $2 \cdot 10^{-4}$  (equally distributed between FW and divertor), per the value used in [30]<sup>3</sup>. The vast majority of tritium fuel passes through the IFC, so its dynamics have a critical impact on the performance of the overall fuel cycle.

The exhaust gas from the vacuum pumps is processed in the fuel clean-up system that separates hydrogenic species from plasma enhancement gases (Ar, Ne, and other inert gases) and He ashes. If the Direct Internal Recycling (DIR) line is implemented [145] (see Table C.3 in Appendix C for the mathematical implementation in the model), a Metal Foil Pump (MFP) [146] or a multi-stage cryopump will be located upstream of the fuel clean-up to provide a preliminary separation of hydrogenic species from the exhaust gas. Separation of hydrogenic species from the exhaust gas by the DIR reduces processing burdens on the fuel clean-up system and ISS, speeding up the IFC dynamics and reducing the tritium inventories in the IFC components.

<sup>3</sup>Tritium implantation in PFCs is just one of the many phenomena involved in plasma material interactions (PMI). Modeling of PMI requires a tremendous effort due to the complex physical and chemical mechanisms involved and the multiscale nature of the problem [140]. The fraction of unburnt tritium that implants in the PFC is usually quantified by surface analysis of small material samples from ion irradiation setups [141] or experimental devices (e.g., JET-ILW [142], JT-60U [143], ASDEX Upgrade [144], etc.). No experimental data is available for the ion fluxes and energies expected in fusion devices like ARC-class tokamaks. Therefore, we assumed a tritium implantation fraction of  $2 \cdot 10^{-4}$  for consistency with the analysis from [30].

The DIR concept was originally proposed by Day et al. [145] and Giegerich et al. [147] for the DEMO fuel cycle. It features a MFP (or a multi-stage cryopump) that can separate a large fraction of hydrogenic species from the exhaust gas (up to 90% [145]), and two identical pumping trains exploiting a linear diffusion pump and a liquid ring pump. The first train pumps the pure fuel permeated through the MFP directly to the storage and management system, whereas the second train pumps the remaining exhaust gas to the fuel clean-up system. Downstream of the fuel clean-up, the Isotope Separation System (ISS) separates hydrogen from deuterium and tritium. The separated deuterium and tritium are sent to the storage and management system. The ISS also processes tritium coming from the detritiation system. We refer to a general “black box” detritiation system without considering specific engineering details in this analysis.

The storage system fulfills many functions. First, it contains the start-up tritium inventory required to start FPP operations. The start-up inventory will be on the order of kilograms, so an effective storage system is critical to plant safety. Second, it works as a short-term buffer for fresh fuel. Lastly, it provides long term storage for the tritium that will be eventually transferred to new fusion power plants for start-up. The last major IFC component is the fueling system, which injects fresh fuel in the vacuum chamber.

## 4.2.2 Fuel Cycle Model

### Mathematical Formulation

The fuel cycle model in this work is based on the resident time method (RTM) proposed in [148] and exploited in the recent analysis of the DEMO fuel cycle [30]. In the RTM, a lumped model is used to describe each component of the fuel cycle. The tritium balance for each component is described with a time-dependent, ordinary differential equation:

$$\frac{dI_i}{dt} = \sum_{j \neq i} \left( f_{j \rightarrow i} \frac{I_j}{\tau_j} \right) - (1 + \varepsilon_i) \left( \frac{I_i}{\tau_i} \right) - \lambda I_i + S_i \quad (4.1)$$

with initial conditions described by Eq. (4.2) and Eq. (4.3):

$$I_i(t = 0) = 0, \forall i \neq \text{storage} \quad (4.2)$$

$$I_{\text{storage}}(t = 0) = I_{\text{startup}} \quad (4.3)$$

Subscripts  $i$  and  $j$  refer to the generic  $i$ -th and  $j$ -th component,  $I$  is tritium inventory,  $\tau$  is tritium residence time,  $\lambda$  is the tritium decay constant, and  $\varepsilon$  is the fraction of tritium lost due to non-radioactive phenomena (e.g., leakages). The evolution of the tritium inventory in the  $i$ -th component,  $I_i$ , depends on:

- Tritium flows entering component  $i$  from the  $j$ -th components connected to  $i$ ,  $\sum_{j \neq i} \left( f_{j \rightarrow i} \frac{I_j}{\tau_j} \right)$
- Outflow from component  $i$ ,  $\frac{I_i}{\tau_i}$
- Non-radioactive and radioactive tritium losses,  $\varepsilon_i \frac{I_i}{\tau_i} - \lambda I_i$
- Tritium generation in component  $i$ ,  $S_i$

Specifically,  $\tau_i$  is the tritium residence time in component  $i$ ,  $\varepsilon_i$  is the non-radioactive tritium loss fraction in component  $i$ , and  $f_{j \rightarrow i}$  is a multiplying factor representing the fraction of component  $j$ 's outflow that is directed towards component  $i$ . If component  $j$  has only one output, to component  $i$ , then  $f_{j \rightarrow i} = 1$ .  $I_{\text{storage}}$  is the tritium inventory in the storage system, and  $I_{\text{startup}}$  is the start-up inventory (see Section 4.2.2 for additional details on these inventories).

The key advantages of the RTM are its computational speed, design flexibility, and the ability to model tritium flow through a component using simple parameters without needing engineering details. For a target condition, such as tritium self-sufficiency, the RTM can be used to understand the required performance from each component, such as blanket TBR, providing researchers with meaningful design targets.

The physical and chemical processes underlying tritium transport in FLiBe are not well understood, and there is a wide spread in available experimental data [23]. As of 2022, there are multiple new experimental efforts underway in the United States to understand FLiBe chemistry, motivated by fusion applications like the ARC LIB and fission applications like the molten salt reactor under development by Kairos

Power. However, until the data from these projects is available, we must rely on coarse high-level models like the RTM. Furthermore, as pointed out in [30], the time and length scales that describe the FPP fuel cycle are not suitable for an atomic-scale simulation of tritium transport. Instead, each component is a black box characterized by tritium residence time within the component and component-specific parameters (e.g., tritium extraction efficiency for the TES, or DIR fraction for DIR). Eventually, new experimental data will be used to reduce uncertainty on the value of these parameters [39].

FPP operations are expected to be intermittent, and are modeled by the Availability Factor (AF)<sup>4</sup>. The AF accounts for both pulsed operations and possible outages due to system failure. Equation (4.4) defines AF,

$$AF = \frac{t_{\text{on}}}{t_{\text{on}} + t_{\text{off}}} \quad (4.4)$$

where  $t_{\text{on}}$  is the pulse duration and  $t_{\text{off}}$  is the downtime between two consecutive pulses. The AF is an effective way to consider FPP downtime in the RTM. Increasing AF by improving the reliability of FPP systems and the efficiency of maintenance procedures is an important goal for fusion engineers.

The RTM is not the only option for studying the FPP fuel cycle, and it is worthwhile to briefly consider the advantages and drawbacks of other approaches. More advanced models have been developed that include detailed component simulations within a dynamic model [149, 150]. However, a detailed design of fuel cycle components is required to implement these models, and they are computationally expensive, as a full component-level simulation is run at every time step. Another notable fuel cycle modeling effort is found in [151], which details a lumped model of the tritium fuel cycle for the EU-DEMO FPP and leverages the preliminary design parameters available for the EU-DEMO fuel cycle. Their model computes the tritium mass flow rates exchanged between the components without explicitly accounting for the residence time. FPP operations in [151] are described by sampling failure and outage times from a log-normal distribution instead of using a mean AF. This accounts for the wide range of timescales that will characterize FPP

---

<sup>4</sup>We use the term “availability factor” for consistency with the analysis carried out in [30]. However, AF has a slightly different meaning here (fraction of time during which the FPP can effectively breed tritium) than it does in its reliability engineering definition (fraction of time during which the FPP is operating).

outage times, from minutes (e.g. fixing a stuck valve) to days (e.g. divertor cassette failure) and captures a more realistic system evolution, albeit at the expense of higher computational costs. A mean AF is implemented in this model because design uncertainties for the ARC fuel cycle outweigh the benefits of the more detailed AF approach in [151]. The lumped model from [151] also uses the expected achievable TBR in EU-DEMO as a model input that can vary narrowly around the nominal value. Doubling time is then a model output in [151]. This is a reasonable approach when the achievable TBR is known and few major design changes to the FPP are expected. However, ARC is in an early design stage, and it is more useful to set a desired doubling time as an input parameter for the RTM model and then calculate the TBR that will be required to achieve it.

### Start-up inventory and reserve inventory

The overall tritium inventory within the FPP site boundary is relevant to safety and regulation. The RTM model is most concerned with the reserve inventory ( $I_{\text{res}}$ ) and the start-up inventory ( $I_{\text{startup}}$ ). Both are key to plant self-sufficiency, and both are major contributors to the overall tritium inventory at the plant. An optimized fuel cycle will minimize  $I_{\text{res}}$  and  $I_{\text{startup}}$  (for the purposes of better safety and smaller regulatory burdens) without sacrificing tritium self-sufficiency.

Equation (4.5) gives  $I_{\text{res}}$ , the amount of tritium required to keep the reactor fueled and operating for a prescribed reserve time ( $t_{\text{res}}$ ) if a fraction  $q$  of the fuel cycle fails:

$$I_{\text{res}} = \frac{\dot{N}_{\text{T,burn}}}{\text{TBE}} q t_{\text{res}} \quad (4.5)$$

The reserve time can be chosen according to economic, safety, and reliability considerations. In theory, it can be equal to zero. Robust and reliable FPPs will not require large  $I_{\text{res}}$ , which in turn eases safety and regulatory burdens by decreasing overall site inventory.

$I_{\text{startup}}$  is the amount of tritium required in fuel storage at the beginning of operations to start and initially run the plant. At the beginning of operations, the tritium bred in the blanket is not readily available as fuel: tritium will build up in components to a certain equilibrium level, plus it takes time for tritium to pass through the fuel cycle systems. Hence, the tritium for FPP fueling is initially provided

by the storage system. Furthermore, the start-up inventory must also compensate for the non-radioactive and radioactive tritium losses from the fuel cycle components. Initially, the storage inventory decreases as tritium is used for fuel, is recirculated, or is lost. As bred tritium becomes available, the storage inventory begins to increase. The minimum storage inventory is reached at a time  $t_{\min}$ , which is on the order of the total residence time in the OFC plus the residence time in the ISS (i.e., the time needed to make newly bred tritium available as fuel).

In a ideal optimized fuel cycle model, the storage inventory reaches a minimum value of zero (temporarily) at the minimum point, indicating that the start-up inventory at  $t = 0$  was minimized ( $I_{\text{startup,opt}}$ ) without sacrificing tritium self-sufficiency. For  $I_{\text{startup}} < I_{\text{startup,opt}}$ , tritium self-sufficiency cannot be achieved, and the FPP will run out of tritium before it can be sustained on bred tritium alone. However, due to uncertainties and possible failures in the fuel cycle systems, a real FPP will have a reserve inventory for the first days of operation. In this case, the storage inventory reaches a minimum value equal to the reserve inventory. The start-up inventory necessarily increases by an amount equal to  $I_{\text{res}}$ . Figure 4.2 shows  $I_{\text{storage}}$  as a function of time for an ideal and realistic FPP.

### Dynamics of a basic fuel cycle model

This subsection describes the dynamics of a basic fuel cycle to provide a common background to the readers. An analytical solution can be indeed found for very simple fuel cycle models, while this task becomes overwhelming for more complex models such as the ones developed for ARC-class tokamaks in the next sections.

Figure 4.3 shows a simple fuel cycle model where all IFC components except the storage system are collapsed in a single “IFC” block, and all the OFC components are collapsed in a single “OFC” block. The storage system is represented by its own block so to explicitly track inventory buildup. The fueling system is neglected for simplicity. We assume a set of reasonable parameters:  $\dot{N}_{\text{T,burn}} = 9 \cdot 10^{-7}$  kg/s (from a total thermal power of 525 MW<sub>th</sub>), TBE = 1%, TBR = 1.08,  $I_{\text{startup}} = 1.5$  kg,  $\tau_{\text{IFC}} = 4$  h and  $\tau_{\text{OFC}} = 24$  h. Radioactive and non-radioactive tritium losses are neglected.

The storage system provides a constant tritium flow for fueling equal to  $\frac{\dot{N}_{\text{T,burn}}}{\text{TBE}}$ . The tritium production in the OFC is equal to  $\dot{N}_{\text{T,burn}} \cdot \text{TBR}$ , while the unburned




Fig. 4.2 Evolution of the storage inventory during plant operations. (a): Ideal scenario ( $AF = 1$ ) with no reserve inventory. The minimum storage inventory is zero at the minimum point. (b): Realistic scenario ( $AF < 1$ ) with a reserve inventory of 1 kg. The minimum storage inventory equals the reserve inventory at the minimum point.  $I_{\text{storage}}$  decreases from  $t = 0$  to  $t = t_{\text{min}}$  because the bred tritium is not immediately available as fuel. The minimum point is reached when the tritium bred at  $t = 0$  has passed through the OFC, the ISS, and reached the storage system. For  $t > t_{\text{min}}$  the storage inventory increases linearly due to the constant inflow of bred tritium (assuming no interruption to plant operation). The values reported in the plots do not refer to any particular design, and were chosen to emphasize the contribution of the reserve inventory to the storage inventory.

tritium is exhausted from the vacuum chamber at a rate equal to  $\dot{N}_{\text{T,burn}} \frac{1-\text{TBE}}{\text{TBE}}$ . The characteristic timescales of the IFC and OFC are defined by their tritium residence times,  $\tau_{\text{IFC}}$  and  $\tau_{\text{OFC}}$ . Equations (4.6) - (4.10) describe the evolution of the tritium inventories in the basic fuel cycle.

$$\frac{dI_{\text{OFC}}(t)}{dt} = \dot{N}_{\text{T,burn}} \cdot \text{TBR} - \frac{I_{\text{OFC}}(t)}{\tau_{\text{OFC}}} \quad (4.6)$$

$$\frac{dI_{\text{IFC}}(t)}{dt} = \frac{1 - \text{TBE}}{\text{TBE}} \dot{N}_{\text{T,burn}} + \frac{I_{\text{OFC}}(t)}{\tau_{\text{OFC}}} - \frac{I_{\text{IFC}}(t)}{\tau_{\text{IFC}}} \quad (4.7)$$

$$\frac{dI_{\text{storage}}(t)}{dt} = \frac{I_{\text{IFC}}(t)}{\tau_{\text{IFC}}} - \frac{\dot{N}_{\text{T,burn}}}{\text{TBE}} \quad (4.8)$$

with initial conditions:

$$I_{\text{IFC}}(0) = I_{\text{OFC}}(0) = 0 \quad (4.9)$$




Fig. 4.3 A basic fuel cycle layout. All IFC components except the storage system are collapsed in a single block with residence time  $\tau_{IFC}$ . Similarly, all the components in the OFC are collapsed in a single block with residence time  $\tau_{OFC}$ . The tritium flows exchanged by the systems are depicted by the arrows.

$$I_{\text{storage}}(0) = I_{\text{startup}} \quad (4.10)$$

where  $I_{IFC}$ ,  $I_{OFC}$ , and  $I_{\text{storage}}$  are the tritium inventories in the IFC, OFC and storage system, respectively. The analytical solution of the system of equations defined by Eqs. (4.6) - (4.10) is reported in Eqs. (4.11) - (4.13).

$$I_{OFC}(t) = \dot{N}_{T,burn} \tau_{OFC} TBR (1 - e^{-t/\tau_{OFC}}) \quad (4.11)$$

Eqs. (4.12) - (4.13) highlight a crucial feature of the fuel cycle dynamics for  $TBE < 1$ . The FPP burns only a fraction of the tritium injected in the vacuum chamber, and the net tritium flow through the storage system is approximately

$$I_{\text{IFC}}(t) = \dot{N}_{\text{T,burn}} \tau_{\text{IFC}} \frac{1 - \text{TBE}}{\text{TBE}} (1 - e^{-t/\tau_{\text{IFC}}}) + \dot{N}_{\text{T,burn}} \text{TBR} \tau_{\text{IFC}} (1 - e^{-t/\tau_{\text{IFC}}}) + \\ + \dot{N}_{\text{T,burn}} \text{TBR} \frac{\tau_{\text{IFC}} \tau_{\text{OFC}}}{\tau_{\text{OFC}} - \tau_{\text{IFC}}} (e^{-t/\tau_{\text{IFC}}} - e^{-t/\tau_{\text{OFC}}}) \quad (4.12)$$

$$I_{\text{storage}}(t) = I_{\text{startup}} + \dot{N}_{\text{T,burn}} (\text{TBR} - 1)t + \dot{N}_{\text{T,burn}} \text{TBR} \frac{\tau_{\text{IFC}}^2}{\tau_{\text{OFC}} - \tau_{\text{IFC}}} (1 - e^{-t/\tau_{\text{IFC}}}) + \\ + \dot{N}_{\text{T,burn}} \text{TBR} \frac{\tau_{\text{OFC}}^2}{\tau_{\text{OFC}} - \tau_{\text{IFC}}} (1 - e^{-t/\tau_{\text{OFC}}}) - \dot{N}_{\text{T,burn}} \tau_{\text{IFC}} \frac{1 - \text{TBE}}{\text{TBE}} (1 - e^{-t/\tau_{\text{IFC}}}) \quad (4.13)$$

$-\frac{\dot{N}_{\text{T,burn}}}{\text{TBE}}$  during the initial transient. It should be noted that the unburnt tritium flow ( $\dot{N}_{\text{T,burn}} \frac{1 - \text{TBE}}{\text{TBE}}$ ) takes  $\sim \tau_{\text{IFC}}$  to be processed and to flow back to the storage system. Furthermore, the contribution of the unburnt tritium flow to the total tritium flowing in the storage system is dominant (at steady state) if  $\text{TBE} \ll 1$ . Since  $I_{\text{IFC}}(t)$  doesn't reach its steady-state value until  $t \sim 3\tau_{\text{IFC}}$ <sup>5</sup>, a large fraction of  $I_{\text{startup}}$  is used to keep the FPP operating during the initial transient.  $I_{\text{storage}}$  reaches  $t_{\text{min}}$  at  $t \sim 3\tau_{\text{OFC}}$  (i.e. when the tritium bred in the OFC reaches the storage system) and thereafter begins to increase. The exact minimum time can be computed analytically by setting the time derivative of  $I_{\text{storage}}$  equal to zero and solving for  $t$ :

$$\left. \frac{dI_{\text{storage}}}{dt} \right|_{t=t_{\text{min}}} = 0 \quad (4.14)$$

that gives  $t_{\text{min}} = 2.7$  days, which is approximately  $3\tau_{\text{OFC}}$ .

The steady-state value of the tritium mass flow rate into the storage system is equal to:

$$\dot{m}_{\text{in,storage,ss}} = \dot{N}_{\text{T,burn}} \text{TBR} + \dot{N}_{\text{T,burn}} \frac{(1 - \text{TBE})}{\text{TBE}} \quad (4.15)$$

which is larger than the tritium flow exiting the storage system:

$$\dot{m}_{\text{out,storage}} = -\frac{\dot{N}_{\text{T,burn}}}{\text{TBE}} \quad (4.16)$$

---

<sup>5</sup>We recall that solutions to ODE in the form  $(1 - e^{-t/\tau})$  reach the steady state value in  $t \sim 3\tau$ , being  $(1 - e^{-3t/\tau}) \sim 0.95$ .

as expected from the fact that  $TBR > 1$ .

### Computational model

The MATLAB Simulink<sup>®</sup> model of the ARC fuel cycle is based on the layout shown in Figure 4.1. The differential equations used to model each component are given in Appendix C, and so only a brief overview of the model is provided here. Eq. (4.1) is the key equation used to describe each component. The only component that was not modeled is the fueling system. Instead, an outflux equal to the tritium fueling rate is added to the equation describing the storage and management system (see Eq. (C.11) in Appendix C for additional details). Pulsed operation was modeled using a pulse source<sup>6</sup> for the tritium generation in the blanket and channels, for the tritium implanted into the FW and the divertor, and for the tritium exhausted by the vacuum pumps. The pulse width is equal to AF.

The computation of a suitable tritium breeding ratio and start-up inventory requires an iterative approach. Both the TBR and  $I_{\text{startup}}$  are unknown at the beginning of the simulation. We stress that the TBR value returned by the simulation is the  $TBR_r$ , where  $TBR_r$  is the minimum TBR required for the plant to achieve tritium self-sufficiency and double  $I_{\text{startup}}$  in the prescribed doubling time  $t_d$ . While  $I_{\text{startup}}$  is mostly insensitive to  $TBR_r$ , the opposite is not true. If  $I_{\text{startup}}$  increases and the prescribed doubling time remains the same, the TBR must increase. At each iteration,  $(I_{\text{storage}}(t) - I_{\text{res}})$  is computed and the start-up inventory is updated accordingly. The TBR is gradually increased, and simulations are run until two constraints are satisfied:

$$I_{\text{storage}}(t) - I_{\text{res}} > 0, \forall t \geq 0 \quad (4.17)$$

$$I_{\text{storage}}(t = t_d) = 2I_{\text{startup}} \quad (4.18)$$

The first constraint, Eq. (4.17), requires that at any time during operation, the desired reserve inventory  $I_{\text{res}}$  is available in tritium storage. The second constraint, Eq. (4.18), requires that the desired  $t_d$  is achieved by the FPP.

<sup>6</sup>A pulse source (or pulse generator) in Matlab Simulink<sup>®</sup> generates a square wave signal of any real variable. The pulse source is characterized by four parameters: amplitude ( $A$ ), pulse width (% of period) ( $w$ ), period ( $T$ ) and phase delay ( $\phi$ ). In the present model  $A = TBR \cdot \dot{N}_{\text{T,burn}}$ ,  $w = \text{AF}$ ,  $T = (t_{\text{on}})$ , and  $\phi = 0$ .

There is sometimes minor confusion about the precise definition of “doubling time.” In this analysis,  $t_d$  is defined as the time needed to store an amount of tritium equivalent to  $2I_{\text{startup}}$ . However,  $t_d$  is sometimes defined as the time required to breed the minimum amount of tritium needed to start a second, identical FPP without sacrificing the tritium self-sufficiency of the first FPP:

$$I_{\text{storage}}(t = t_d) = I_{\text{startup}} + I_{\text{res}} \quad (4.19)$$

In this analysis, the more straightforward definition given in Eq. (4.18) is used. However, it is useful to remember that  $I_{\text{storage}}(t_d)$  can be very different for each definition.

An advantage of MATLAB Simulink® is that it automatically selects the (variable) stepsize and the numerical solver (that implements the adaptive step size) for the system of equations at each point in the model. This is important because the timescales involved in the fusion fuel cycle (FPP lifetimes are given in decades) are much larger than the required stepsize at the beginning of the simulation (10-1000 s). Once the tritium inventory of all components (other than tritium storage) have reached quasi-steady state, the timestep is increased for the sake of computational efficiency.

### 4.3 Tritium Self-Sufficiency in an ARC-class fusion power plant

In this section, the methods described in Section 4.2 are used to analyze the fuel cycle of an ARC-class FPP and determine the conditions needed to achieve tritium self-sufficiency. Section 4.3.1 introduces the model parameters. Section 4.3.2 reports the one-at-a-time sensitivity analysis [152]. Results in Sections 4.3.1 and 4.3.2 are normalized with respect to the mean value to better show how  $TBR_r$  and  $I_{\text{startup}}$  depend on the design parameters. Sections 4.3.3-4.3.9 present the quantitative dependencies of  $TBR_r$  and  $I_{\text{startup}}$  on individual design parameters. Section 4.3.10 shows how the choice of the non-radioactive loss fraction affects tritium self-sufficiency. Section 4.3.11 describes three example design and operational scenarios for tritium self-sufficient ARC-class FPPs.

### 4.3.1 ARC-class tokamak model parameters

Here, we choose a reference configuration for ARC and vary the design parameters to understand the sensitivity of  $TBR_r$  and  $I_{\text{startup}}$  to different design choices. This is a simple but effective way to perform sensitivity analysis, also known as one-at-a-time sensitivity analysis [152]. Table 4.1 lists the design parameters for the baseline ARC case and the range of values considered in the sensitivity analysis. Results presented in Section 4.3 reference these values unless otherwise specified. The baseline case considers a 525 MW<sub>th</sub> ARC-class tokamak with an availability factor  $> 50\%$ , a low DIR fraction ( $f_{\text{DIR}} = 0.3$ ), a 4 hour tritium processing time (time for tritium to be processed in the fuel clean-up system and ISS), the ability to operate on tritium fuel reserves for up to 24 h, and an aggressive target for  $t_d$  (2 y, in contrast to the 5-10 y expected for DEMO-class FPPs [151]).

Parameter	Symbol	Value	Range	Units
Fusion power	$P_{\text{fus}}$	525	-	MW <sub>th</sub>
Tritium burn rate	$\dot{N}_{\text{T,burn}}$	$8.99 \cdot 10^{-7}$	-	kg/s
Pulse duration	$t_{\text{pulse}}$	1800	1800 - 3600	s
Time between pulses	$t_{\text{off}}$	60	60-120	s
Tritium burn efficiency	TBE	0.02	0.005 – 0.1	-
Non-radioactive loss fraction	$\epsilon_i$	$10^{-4}$	-	-
Availability factor	AF	0.7	0.5 - 1	-
Tritium processing time	$t_p$	4	1 – 12	h
Doubling time	$t_d$	2	-	y
Fraction of the system failing	q	0.25	-	-
Reserve time	$t_{\text{res}}$	24	0 - 48	h
Direct Internal Recycling fraction	$f_{\text{DIR}}$	0.3	0.1 – 0.9	-

Table 4.1 Design parameters for the baseline ARC-class FPP fuel cycle analyzed in Section 4.3.

A preliminary analysis suggests that TBE=1-3% for ARC [153]. The non-radioactive loss fraction is set to  $10^{-4}$  as in [30]. However, the reader should be aware that the non-radioactive loss fraction significantly affects the results of the fuel cycle analysis and  $TBR_r$ . In this analysis it is considered a meta-parameter, i.e. a parameter that is set *a priori* without a direct link to physics phenomena occurring in the fuel cycle. An availability factor lower than 50% is not considered, as commercial

electricity generating plant designs with  $AF < 50\%$  will not be commercially viable<sup>7</sup>. The tritium processing time  $t_p$  ranges from 1 h to 12 h. Abdou et al. [156] assumed a tritium processing time in the range 1-24 h based on the experience from the TSTA (Tritium Systems Test Assembly) experiment [157], which reported  $t_p = 24$  h. The same author restricted the range to 1-12 h in the more recent work on tritium fuel cycle analysis [30]. A recent analysis for DEMO fuel cycle assumed  $t_p = 4800$  s [145]. Coleman et al. [151] performed fuel cycle analyses considering a  $t_p = 5 \pm 2$  h. The range 1 h to 12 h is therefore considered suitable to investigate a wide range of possible operating scenarios. The reader should be aware that  $t_p$  here is considered to be the sum of the fuel clean-up time  $\tau_{fc}$  and the ISS residence time  $\tau_{ISS}$  only. We kept this definition of  $t_p$  for consistency with the work from [30]. An alternative definition of  $t_p$  for an FPP with DIR technology can be easily derived from a tritium mass balance in the DIR loop, fuel cleanup system and ISS:

$$I_{tot} = I_{DIR} + I_{fc} + I_{ISS} \quad (4.20)$$

where  $I_{tot}$  is the total mass of tritium in the control volume encompassing the DIR loop, fuel cleanup system and ISS. Recalling that  $\dot{m}_i = \frac{I_i}{\tau_i}$  (Section 4.2.2), Eq. (4.20) can be written as:

$$\dot{m}_{tot} t_{p,eff} = \dot{m}_{DIR} \tau_{DIR} + \dot{m}_{fc} \tau_{fc} + \dot{m}_{ISS} \tau_{ISS} \quad (4.21)$$

where  $\dot{m}_{tot} = \dot{m}_{DIR} + \dot{m}_{fc} + \dot{m}_{ISS}$ , and the effective (or average) tritium processing time ( $t_{p,eff}$ ) is defined as the average time required by the tritium to flow from the vacuum pumps to the storage system (passing through the DIR or through the fuel cleanup and isotope separation systems). Since  $f_{DIR} \equiv \frac{\dot{m}_{DIR}}{\dot{m}_{tot}}$ , Eq. (4.21) can be rewritten as:

<sup>7</sup>A simple calculation can show why  $AF < 50\%$  is not commercially viable. The payback time (PT) of a FPP can be computed to a first approximation as  $PT = C_{tot} / (8760 P_{el} AF c_{el})$ , where PT [y] is the payback time,  $C_{tot}$ [\$] is the total FPP cost (here assumed equal to the overnight cost for sake of simplicity), and  $c_{el}$  [\$/kWh] is the electricity cost. For FPPs,  $C_{tot}$  is still unknown, but a reasonable value can be extrapolated from innovative nuclear reactor designs, such as GEN-IV molten salt reactors (MSR). The overnight cost of an MSR ranges from \$1B to \$10B [154]. Assuming  $C_{tot} = \$5B$  and  $c_{el} = 0.129$  \$/kWh (the average price of U.S. electricity in June 2022 [155]), for  $P_{el} = 200$  MW<sub>e</sub> (the ARC-class tokamak power output) and  $AF = 0.5$ , the resulting payback time is  $PT = 44$  y, which is too long to make the FPP economically attractive.

Component	Symbol	Value	Range	Units
Breeding zone	$\tau_{BZ}$	1.25	1 – 240	h
First wall	$\tau_{FW}$	1000	-	s
Divertor	$\tau_{div}$	1000	-	s
Tritium extraction system	$\tau_{TES}$	24	1 - 240	h
Heat exchanger	$\tau_{HX}$	1000	-	s
Vacuum pump	$\tau_{VP}$	600	-	s
Fuel clean-up	$\tau_{fc}$	0.3	0.1 - 1	h
Isotope separation system	$\tau_{ISS}$	3.7	0.9 - 11	h
Detritiation system	$\tau_{det}$	1	-	h

Table 4.2 Tritium residence time in the fuel cycle components and systems. Nominal value for the baseline case and the range used in the parametric analysis are reported.

$$t_{p,eff} = \tau_{DIR}f_{DIR} + (1 - f_{DIR})(\tau_{ISS} + \tau_{fc}) \quad (4.22)$$

This definition is useful when one wants to investigate the effect of DIR at a very high level, without entering into system-level details. Since this model is system-level and allows for a more detailed treatment, this definition is not used. Hence, the value of  $t_p$  is independent of  $f_{DIR}$  in the following analysis.  $t_{res}$  ranges from 1-48 h. Larger values of  $t_{res}$  leads to impractical start-up inventories, as shown in Section 4.3.8.

The second set of relevant parameters for the model are the tritium residence times in individual fuel cycle components and systems. These values are reported in Table 4.2. Note that we refer to a general “breeding zone” and not to the LIB, because tritium will be produced in FLiBe-containing VV and divertor channels. The residence time in the breeding zone was quantified from the steady-state tritium inventory computed by Ferrero et al. [27]. Fuel cycle simulations with the present model were run for different  $\tau_{BZ}$  until the breeding zone tritium inventory reached 3 g, which was the value found in [27]. This resulted in a  $\tau_{BZ} = 1.25$  h. The fuel clean-up and ISS residence time sum to 4 h for the baseline case, 1 h for a short processing time case, and 12 h for a long processing time case [30]. The storage residence time does not appear in Table 4.2 because the tritium outflow from the tritium storage system is defined by the required fueling rate.

Parameter	Symbol	Mean value	Range	Units
Availability factor	AF	0.75	$\pm 0.25$	-
Tritium burn efficiency	TBE	2.5	$\pm 2.5$	%
DIR fraction	$f_{\text{DIR}}$	0.5	$\pm 0.5$	-
Breeding zone residence time	$\tau_{\text{BZ}}$	13	$\pm 11$	h
TES efficiency	$\eta_{\text{TES}}$	0.7	$\pm 0.3$	-
Tritium processing time	$t_p$	6.5	$\pm 5.5$	h
Reserve time	$t_{\text{res}}$	24	$\pm 24$	h
Doubling time	$t_d$	6	$\pm 5$	y

Table 4.3 Mean value  $\mu$  and range  $\sigma$  of the parameters used in the Sec. 4.3.2 sensitivity analysis. Each parameter was varied across the specified range, centered on the mean value, in increments of (range)/6.

### 4.3.2 Sensitivity analysis

The following eight parameters have been varied for the sensitivity analysis: availability factor (AF), tritium burn efficiency (TBE), DIR fraction ( $f_{\text{DIR}}$ ), breeding zone residence time ( $\tau_{\text{FW}}$ ), TES efficiency ( $\eta_{\text{TES}}$ ), tritium processing time ( $t_p$ ), reserve time ( $t_{\text{res}}$ ) and doubling time ( $t_d$ ). During each simulation, all but one of the parameters are fixed. The parameter being studied is varied through the range specified in Table 4.3. The mean values and the range specified in Table 4.3 are slightly different from the nominal values specified in Table 4.1 and Table 4.2 to better emphasize the sensitivities of  $\text{TBR}_r$  and  $I_{\text{startup}}$ .

Figure 4.4 shows the results from the sensitivity analysis for  $\text{TBR}_r$ . The sensitivity index coefficient (IC) was computed as:

$$\text{IC} = \frac{\Delta Y}{\Delta X} \frac{X_{\text{mean}}}{Y_{\text{mean}}} \quad (4.23)$$

where  $\Delta Y$  is the change in  $\text{TBR}_r$  or  $I_{\text{startup}}$  due to a change of  $\Delta X$  in one of the design parameters.  $\Delta X$  is computed as:

$$\Delta X = x_{\text{max}} - x_{\text{min}} \quad (4.24)$$

where  $x_{\text{max}} = \mu + \sigma$  and  $x_{\text{min}} = \mu - \sigma$ ,  $\mu$  is the mean value and  $\sigma$  is half of the range of variation. The relative sensitivity index coefficient (RIC) makes it easier

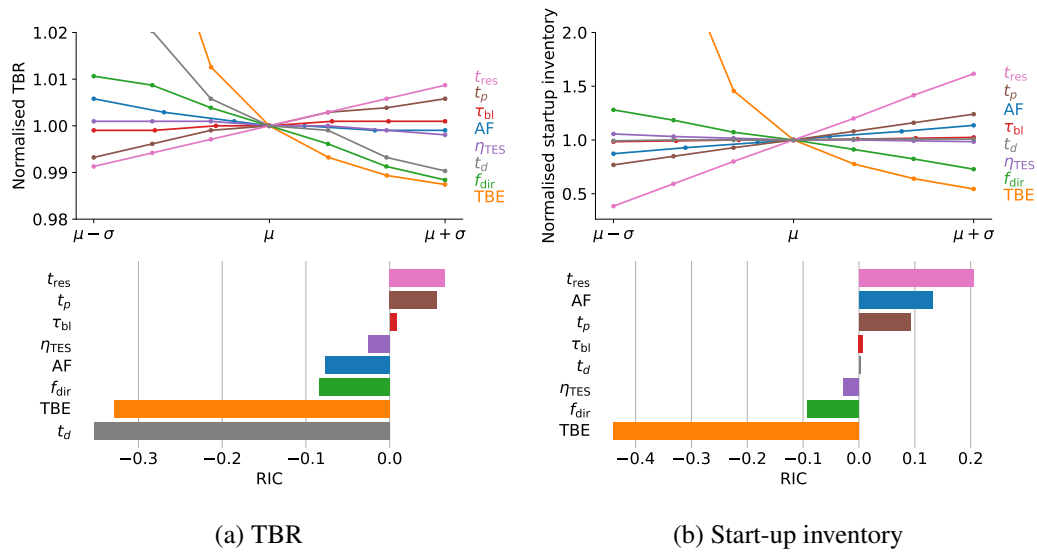


Fig. 4.4 (a) Normalized TBR and relative sensitivity index coefficient (RIC) show the sensitivity of  $TBR_r$  to the eight parameters listed in Table 4.3. (b) Normalized  $I_{startup}$  and RIC show the sensitivity of  $I_{startup}$  to the same eight parameters. To normalize  $TBR_r$  and  $I_{startup}$ , all design parameters were set to the mean value reported in Table 4.3. Each design parameter has been varied in  $[\mu - \sigma, \mu + \sigma]$  range, where  $\sigma$  is the range reported in Table 4.3, choosing 7 equally-spaced points. Both  $TBR_r$  and  $I_{startup}$  sharply increase as TBE decreases. TBE shows the second largest (by absolute value) RIC for  $TBR_r$  and the largest RIC for  $I_{startup}$ . With the exception of AF, the RIC for each design parameter has the same sign for  $TBR_r$  and  $I_{startup}$ , meaning that the minimization of  $TBR_r$  and  $I_{startup}$  can be achieved without tradeoffs. AF considerations are explained in detail in Section 4.3.7.

to understand the dependence of  $TBR_r$  and  $I_{\text{startup}}$  on the design parameters. It is plotted in Figure 4.4 and is given by:

$$\text{RIC} = \frac{\text{IC}_i}{\sum_n^{N_{\text{params}}} \text{IC}_n} \quad (4.25)$$

where  $\text{IC}_i$  is the sensitivity index coefficient of the  $i$ -th parameter and the summation runs over all parameters.

The most impactful parameters for  $TBR_r$  are TBE and the doubling time  $t_d$ .  $TBR_r \rightarrow \infty$  as  $\text{TBE} \rightarrow 0$  because the fuel utilization is so low that the amount of bred tritium is negligible relative to the  $I_{\text{startup}}$  that must be doubled. This limit also follows from the definition of the fuel injection rate  $\dot{T}_{\text{fuel}}$ :

$$\dot{T}_{\text{fuel}} = \frac{\dot{N}_{\text{T,burn}}}{\text{TBE}} \quad (4.26)$$

As  $\text{TBE} \rightarrow 0$ ,  $\dot{T}_{\text{fuel}} \rightarrow \infty$ , and so the required  $I_{\text{startup}} \rightarrow \infty$  as well (Figure 4.4). It is therefore impossible to meet a given target  $t_d$  unless  $TBR \rightarrow \infty$ . Similarly,  $TBR_r \rightarrow \infty$  as  $t_d \rightarrow 0$ . In order to achieve both tritium self-sufficiency and shorter  $t_d$ , it is necessary to achieve high TBE. Other factors that impact  $TBR_r$  - although not to the extent that  $t_d$  and TBE do - include:

- DIR fraction  $f_{\text{DIR}}$ : affects IFC dynamics. Higher DIR fractions reduce the time tritium spends in the IFC, decreasing  $I_{\text{startup}}$  and  $TBR_r$ ;
- Reserve time  $t_{\text{res}}$ : a larger  $t_{\text{res}}$  requires a larger  $I_{\text{startup}}$  and thus a higher  $TBR_r$  (given a fixed  $t_d$ );
- Tritium processing time  $t_p$ : affects IFC dynamics in a similar way to  $f_{\text{DIR}}$ ;
- Availability factor (AF): a lower AF indicates less operating time and less tritium production, so  $TBR_r$  must increase to meet the  $t_d$  target if AF decreases.

$\tau_{\text{BZ}}$  and  $\eta_{\text{TES}}$  have a weak effect on  $TBR_r$ . While they affect the tritium inventory in the OFC, they do not significantly impact system dynamics and tritium production. This is explained in greater detail in Sections 4.3.3-4.3.4.

### 4.3.3 Tritium burn efficiency

$TBR_r$  and  $I_{\text{startup}}$  are shown in Figure 4.5 as a function of TBE for different values of  $t_p$ . It is highly unlikely that we can attain  $TBR > 1.3$ ; extremely challenging but possible to attain  $TBR = 1.2-1.3$ , challenging but reasonable to expect that we can attain  $TBR = 1.1-1.2$ , and likely to achieve  $TBR < 1.1$ . The lowest value of TBE considered, 0.5%, might be a realistic value for the first generation of ARC-class tokamaks [153], but will require fast processing times and high TBR to achieve tritium self-sufficiency. For low values of TBE,  $t_p$  has a large impact on  $TBR_r$  and  $I_{\text{startup}}$  because most of the injected tritium is exhausted from the vacuum boundary and processed in the IFC. Figure 4.5 highlights the importance of achieving high TBE and fast  $t_p$  in order to relax TBR requirements on the blanket design.

As TBE increases, less tritium is flowing through the IFC and the importance of  $t_p$  diminishes. Slow IFC dynamics ( $t_p > 4$  h) are acceptable for  $TBE > 4\%$ , and  $TBR_r < 1.1$  and  $I_{\text{startup}} < 2\text{kg}$  are still within reach. Figure 4.5 shows that targeting a longer  $t_d$  is an effective way to reduce  $TBR_r$ . However,  $I_{\text{startup}}$  is weakly affected by the choice of  $t_d$ , and other strategies are needed to decrease it.

### 4.3.4 Outer fuel cycle residence time

The relevant timescales for tritium transport through the OFC range from a few minutes in the divertor, FW, and heat exchanger to days in the breeding zone and TES. OFC dynamics are therefore dominated by  $\tau_{\text{BZ}}$  and  $\tau_{\text{TES}}$ . The baseline case assumes  $\tau_{\text{BZ}} = 1.25$  h and  $\tau_{\text{TES}} = 24$  h (Table 4.2). The components with a short tritium residence time do not significantly affect nominal system operations, even if the residence time is increased by a factor of 10 (i.e., from 20 min to 3 h). Figure 4.6 shows how different values of  $\tau_{\text{BZ}}$  affect  $TBR_r$  and  $I_{\text{startup}}$ . Decreasing  $\tau_{\text{BZ}}$  does not have a significant impact on  $TBR_r$  or  $I_{\text{startup}}$  for  $\tau_{\text{BZ}} < 24$  h. As  $\tau_{\text{BZ}}$  goes above 1 d, there is a small corresponding increase in  $TBR_r$  and  $I_{\text{startup}}$ . Generally, as long as  $\tau_{\text{BZ}}$  is on the order of 1 d, engineering efforts directed at decreasing  $\tau_{\text{BZ}}$  are unlikely to meaningfully reduce  $TBR_r$  or  $I_{\text{startup}}$ .

Figure 4.7 shows the evolution of tritium inventories in the ISS, tritium storage, breeding zone, and TES systems for  $\tau_{\text{BZ}} = 1$  h and  $\tau_{\text{BZ}} = 10$  d. Longer values of  $\tau_{\text{BZ}}$  are associated with a larger breeding zone inventory and a delayed minimum point.

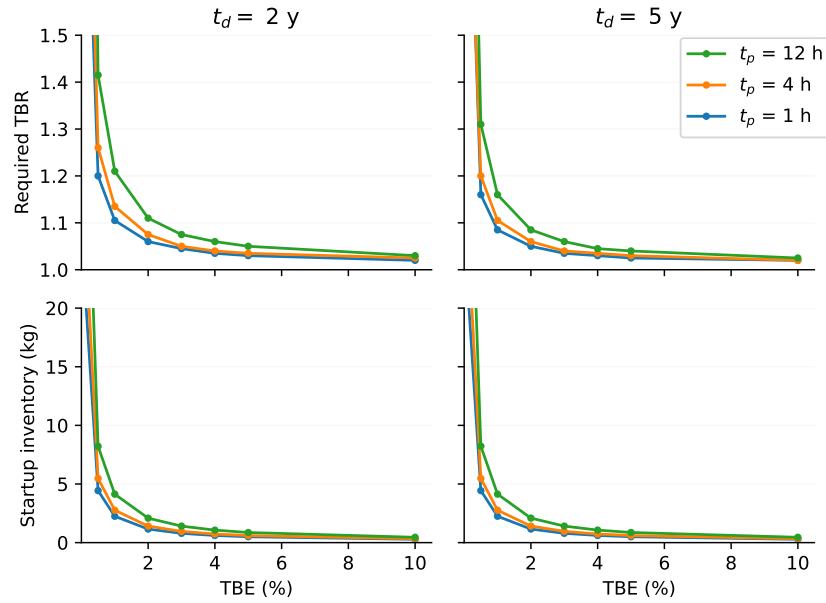


Fig. 4.5  $TBR_r$  and  $I_{\text{startup}}$  as a function of TBE, for different processing times and doubling times. For  $t_d = 2$  y (left column),  $t_p$  has a small impact on  $TBR_r$  when TBE is high ( $> 4\%$ ). For low TBE, tritium self-sufficiency becomes difficult to practically achieve unless  $t_p$  is quite low. For the longer  $t_d$  of 5 y (right column), engineering and design constraints on the blanket are relaxed: tritium production rates can be lower and longer  $t_p$  is acceptable. The plots also show that for very low TBE  $< 0.5\%$ ,  $I_{\text{startup}}$  and  $TBR_r$  sharply increase to impractical values.

In Figure 4.7, the increase in  $\tau_{\text{BZ}}$  shifts the minimum point time from  $\sim$ hours after start-up to  $\sim$ days. Nevertheless, this shift is negligible compared to the doubling time, and increasing  $\tau_{\text{BZ}}$  results in very small increases in  $TBR_r$ . The observed increase in  $TBR_r$  compensates for the larger tritium losses in the breeding zone that inevitably occur (e.g., due to permeation or radioactive decay) when tritium resides there for longer.

Figure 4.8 shows the steady-state breeding zone inventory for different residence times. This relationship is linear. A  $\tau_{\text{BZ}}$  of 10 d results in a steady-state breeding zone inventory of  $\sim 600$  g; for  $\tau_{\text{BZ}} = 1$  h, this drops to  $\sim 3$  g. Here we see the advantage of efficient T extraction from the blanket: while  $\tau_{\text{BZ}}$  may not have a large impact on  $TBR_r$  or  $I_{\text{startup}}$ , it determines T inventory in the blanket. Large quantities of tritium in any part of the plant present maintenance challenges and may be unacceptable from a regulatory standpoint: we want to attain tritium self-sufficiency while still minimizing overall on-site inventory.




Fig. 4.6  $TBR_r$  and  $I_{startup}$  for different residence times in the breeding zone.  $\tau_{BZ}$  has effectively no impact on  $TBR_r$  and  $I_{startup}$  for  $\tau_{BZ} < 1$  day. At longer  $\tau_{BZ}$ ,  $TBR_r$  and  $I_{startup}$  increase slightly.

$\tau_{TES}$  and  $\tau_{BZ}$  have the same effects on the OFC dynamics and tritium self-sufficiency. Varying  $\tau_{TES}$  results in equivalent inventory evolution trends as those shown in Figs. 4.6 and 4.7, and the plotted results are omitted here.




Fig. 4.7 Evolution of relevant system inventories for (a)  $\tau_{BZ} = 1.25$  h and (b)  $\tau_{BZ} = 10$  d. The slower dynamics and the shift of the minimum time for  $\tau_{BZ} = 10$  d is depicted.




Fig. 4.8 Steady-state tritium inventory in the breeding zone for different residence times. Longer  $\tau_{BZ}$  is associated with larger T inventories in the blanket, which presents maintenance and regulatory challenges. Higher values of TBE are associated with smaller tritium inventories in the breeding zone, as the fuel is being used more efficiently and  $TBR_r$  is lower.

### 4.3.5 Tritium extraction efficiency

The OFC is further characterized by the TES tritium extraction efficiency  $\eta_{\text{TES}}$ . It is difficult to assume a baseline value for  $\eta_{\text{TES}}$  for ARC because tritium extraction from FLiBe has been poorly investigated [158, 159], extrapolation from small scale experiments is not straightforward, and full-scale breeding systems have not been tested yet. Experimental data on tritium solubility and diffusivity in FLiBe is limited in scope and varies by multiple orders of magnitude [23]. In this analysis,  $\eta_{\text{TES}}$  is varied over a wide range to explore the consequences of multiple operating points for an ARC-class plant. Data from the upcoming LIBRA experiment will narrow the range of possible values of  $\eta_{\text{TES}}$  for an ARC-class FPP [39].

Counterintuitively,  $\eta_{\text{TES}}$  does not have a strong impact on tritium self-sufficiency. If  $\eta_{\text{TES}}$  is lowered, the concentration of tritium in the FLiBe of the ARC breeding zones increases, but the rate of tritium extraction is maintained at the prior level. This affects the tritium inventories in the OFC (higher tritium concentrations in FLiBe lead to higher concentrations in the other OFC components), but has minimal impact on tritium self-sufficiency. The OFC components take more time to reach their steady-state tritium inventory values, but  $\text{TBR}_r$  remains almost the same. We have explained in Section 4.3.2 how the  $\text{TBR}_r$  is strictly related to  $t_d$ . The slowing of OFC dynamics is negligible with respect to the doubling time (days vs years). Furthermore, the additional tritium lost due to buildup in the OFC components (10-100 g, as shown in Figure 4.9) is 1-2 orders of magnitude smaller than the initial inventory that must be doubled (1-10 kg). The TES tritium inventory for different values of  $\eta_{\text{TES}}$  is shown in Figure 4.9. Even in the worst-case scenario ( $\text{TBE} = 0.5\%$ ) the difference between a high efficiency TES ( $\eta_{\text{TES}} = 95\%$ ) and a low efficiency TES ( $\eta_{\text{TES}} = 40\%$ ) is a  $\sim 100$  g increase in the TES inventory.  $\eta_{\text{TES}}$  has a marginal effect on  $\text{TBR}_r$ . For  $\text{TBE} = 1\%$ , if we reduce  $\eta_{\text{TES}}$  from 95% to 40%,  $\text{TBR}_r$  increases from 1.15 to just 1.16 - less than a 1% increase.




Fig. 4.9 Steady-state tritium inventory in the TES for different values of  $\eta_{TES}$  and TBE. Higher efficiency in the TES and in the fueling system corresponds to lower TES tritium inventory.

### 4.3.6 Direct internal recycling fraction

DIR impacts IFC dynamics only. As shown for the DEMO fuel cycle in [30], the inclusion of a DIR line can have a dramatic impact on  $TBR_r$  and  $I_{\text{startup}}$  in the ARC fuel cycle if operating at low TBE values. DIR technology is therefore an important tool for achieving tritium self-sufficiency, especially if fueling technologies are not significantly improved by the time the first FPPs are built. A DIR fraction ( $f_{\text{DIR}}$ ) of 0.9 (indicating that 90% of exhausted fuel is directly recycled back to tritium storage and fueling) reduces  $TBR_r$  by 10% for TBE=0.5%. The  $TBR_r$  reduction is less dramatic when operating with high-efficiency fueling: at higher values of TBE, the tritium flow rate in the IFC is decreased, thus reducing the impact of DIR on  $TBR_r$ . For example, at TBE=5%, implementing DIR with  $f_{\text{DIR}}=0.9$  reduces  $TBR_r$  by just 1%. This relationship between  $TBR_r$  and  $f_{\text{DIR}}$  is shown in Figure 4.10(a).  $I_{\text{startup}}$  follows a similar trend: increasing  $f_{\text{DIR}}$  decreases  $I_{\text{startup}}$ , but the trend is most pronounced at low TBE values.  $I_{\text{startup}}$  is strongly affected by IFC dynamics because the majority of tritium in the plant will flow through the IFC (95% for TBE=5%). The DIR line is therefore also a powerful tool for reducing  $I_{\text{startup}}$  and easing regulatory burdens for ARC. However, including DIR leads to a more complex IFC design. If ARC is able to achieve high TBE, the increase in system complexity may not be worth the small reduction in  $I_{\text{startup}}$  and  $TBR_r$  (especially for low  $f_{\text{DIR}}$ ).




Fig. 4.10  $TBR_r$  and  $I_{startup}$  as a function of the direct internal recycling fraction. DIR is more advantageous when operating at low tritium burn efficiency.

### 4.3.7 Availability factor

FPPs must reliably output power to be economically feasible. It is explained in [30] that the mean time between failures (MTBF) in DEMO and the mean time to repair (MTTR) will be respectively short (on the order of days) and long (on the order of months). This is because the DEMO blanket, FW, and divertor are located inside the vacuum vessel, which makes reliability, accessibility, maintainability, and inspectability (RAMI) quite complex for these components. Here, we see a key advantage of the ARC-class FPP. The vacuum vessel is located inside the LIB, which uses a minimum of structural materials. This has important consequences for RAMI. The liquid blanket is not damaged by neutron radiation, cannot fail due to thermo-mechanical stresses, and provides a large thermal sink in case of severe accidents. Furthermore, demountable magnet coils and a drainable blanket allow easy access to the VV for maintenance and replacement<sup>8</sup>. In-vessel components can be removed and replaced without dismantling a great number of solid structural components, reducing MTTR. Ideally, these design choices will increase the AF of ARC-class FPPs relative to other designs, as AF is inversely proportional to MTTR [30]):

$$AF = \frac{MTBF}{MTBF + MTTR} \quad (4.27)$$

This is an alternative definition of AF (compare to Eq. (4.4) in Section 4.2.2) that explicitly shows the inverse dependence of AF on MTTR. The effect of plant AF on  $TBR_r$  is shown in Figure 4.11. Improving AF has the most significant impact on reducing  $TBR_r$  when TBE is low. The two dashed lines identify the best-case AF operating points for two possible combinations of  $(t_{\text{pulse}}, t_{\text{off}})$ . For  $(t_{\text{pulse}} = 1800 \text{ s}, t_{\text{off}} = 60 \text{ s})$  the best possible AF (no system failures or need for maintenance) is 97%. For  $(3600 \text{ s}, 120 \text{ s})$  the best possible AF is 94%.

At  $TBE = 0.5\%$ , tritium self-sufficiency is almost impossible if  $AF < 90\%$ : it is difficult to see, based on neutronics studies, how we could achieve a  $TBR > 1.25$ . In this scenario, a large DIR fraction would be required to bring down the  $TBR_r$ , and even then it is not obvious that tritium self-sufficiency would be possible. On the

<sup>8</sup>In ARC-class FPPs, the VV is not a lifetime component, but is replaced regularly (i.e., every few years). This eases the challenge of robust VV design, because the structure needs to retain its integrity for a shorter time.

other hand, for  $TBE > 1\%$ , tritium self-sufficiency could theoretically be achieved even with low availability factors, e.g.  $AF = 50\%$  (although  $TBR_r > 1.15$  is still not necessarily “easy” to achieve). For  $TBE = 5\%$ ,  $TBR_r$  is almost insensitive to improvement in  $AF > 70\%$ . Note that if  $AF$  drops below  $50\%$ ,  $TBR_r$  increases significantly, as described in [30]. This is because at low  $AF$ , radioactive decay loss becomes significant (inventory will decrease at a rate of  $\sim 5.5\%$  per year if the plant is not producing tritium). Furthermore, if  $t_d$  is fixed, more tritons must be produced for every neutron to compensate for the shorter operating window associated with low  $AF$ . This concept is shown in Figure 4.12.

$I_{startup}$  is not affected by  $AF$  for any value of  $TBE$ . A lower  $AF$  shifts  $t_{min}$  further in time. However, even for  $AF = 50\%$ , when  $t_{min}$  is almost doubled, the timescales involved are of the order of days. The additional tritium required to compensate for the associated decay losses is negligible, and  $I_{startup}$  does not increase when the  $AF$  decreases. So while improving  $AF$  will not reduce  $I_{startup}$ , high  $AF$  is necessary for faster  $t_d$  and, of course, reliable production of power.




Fig. 4.11  $TBR_r$  as a function of availability factor ( $AF$ ). The dashed lines identify the  $AF$  for two different combinations of ( $t_{pulse}$  and  $t_{off}$ ).




Fig. 4.12 Evolution of the tritium inventories in the tritium extraction system, isotope separation system, and tritium storage system. These are the most relevant tritium inventories in the plant. At fixed TBR, a FPP with AF = 50% takes almost twice the time to double the initial inventory with respect to a FPP with AF = 100%.

### 4.3.8 Doubling time

The doubling time  $t_d$  is the most important parameter affecting  $TBR_r$ , and it is assumed that the FPP is intentionally designed to be capable of achieving a target  $t_d$ . A shorter  $t_d$  requires a larger tritium production rate and a higher TBR. Conversely, tritium self-sufficiency requirements can be relaxed if a longer  $t_d$  is acceptable to the FPP designers. Figure 4.13 shows  $TBR_r$  for the ARC system for different doubling times. For  $TBE = 0.5\%$ , FPP designers would likely need to settle for long doubling times (on the order of years) in order to get to a reasonably achievable  $TBR_r$ . If  $t_d = 5$  y is acceptable, then tritium self-sufficiency should not be a leading issue, especially if  $TBE > 1\%$  (i.e.,  $TBR_r < 1.1$ ). It should be noted that short  $t_d$  (e.g.,  $t_d < 2$  y) may not even be desirable in early plants if FPPs are slow to penetrate the electricity market. Rapidly doubling the tritium inventory every year without a corresponding demand could result in unnecessary regulatory challenges due to the need to safely store large amounts of tritium. For now, a target  $t_d = 2$  y seems reasonable in terms of design requirements and possible benefits for a first generation of ARC-class FPPs.




Fig. 4.13  $TBR_r$  for different doubling times  $t_d$ . A short  $t_d$  requires high TBR. For  $t_d = 1$  y, tritium self-sufficiency is highly unlikely if  $TBE < 1\%$ . If a longer  $t_d$  is acceptable, tritium self-sufficiency is much easier to achieve. For  $t_d = 5$  y, tritium self-sufficiency is possible even if  $TBE = 0.5\%$ . The bottom plot zooms in on  $TBR_r$  in the region of interest ( $TBR_r < 1.3$ ).

### 4.3.9 Reserve time

The reserve time  $t_{\text{res}}$  has a fundamental impact on FPP operations in the event that some part of the fuel cycle fails. Longer reserve times make it possible to keep operating when failure requires a lengthy repair, but leads to higher  $I_{\text{startup}}$  and thus higher  $\text{TBR}_r$  for a fixed  $t_d$  (Figure 4.14). At  $\text{TBE} = 0.5\%$ ,  $t_{\text{res}} > 24$  h is infeasible as this leads to unachievable  $\text{TBR}_r$  ( $> 1.3$ ). Improving TBE enables a longer  $t_{\text{res}}$ , but it is important to balance the advantages of a longer  $t_{\text{res}}$  against the disadvantages of a larger  $I_{\text{startup}}$ . At low values of TBE, small increases in  $t_{\text{res}}$  correspond to large increases in  $I_{\text{startup}}$ , as shown in Figure 4.14.




Fig. 4.14  $\text{TBR}_r$  and  $I_{\text{startup}}$  for an ARC-class FPP as a function of tritium reserve time.  $I_{\text{startup}}$  increases linearly with  $t_{\text{res}}$ . The  $\text{TBR}_r$  increases accordingly for fixed  $t_d$  as a higher inventory must now be doubled.

### 4.3.10 Non-radioactive loss fraction

The non-radioactive loss fraction is a parameters of great interest because it models tritium losses and, indirectly, the efficiency of the tritium permeation barriers.  $\epsilon = 10^{-4}$  is a necessary target to achieve tritium self-sufficiency. For  $TBE > 2\%$  a lower  $\epsilon$  does not strongly affect tritium self-sufficiency, whereas at lower TBE reducing  $\epsilon$  below  $10^{-5}$  could be beneficial (Figure 4.15). Tritium self-sufficiency is impossible for  $\epsilon > 10^{-3}$  unless  $TBE \gg 5\%$ , and even in that case it would be extremely challenging to achieve tritium self-sufficiency ( $TBR_r > 1.2$ ).

The start-up inventory is not affected by  $\epsilon$ . Section 4.2 showed how the start-up inventory is mostly affected by the initial transient occurring before the minimum time. The tritium lost during this transient due to large  $\epsilon$  is negligible, hence  $I_{\text{startup}}$  does not change even at high  $\epsilon$ .

Tritium self-sufficiency might not be the leading goal in minimizing  $\epsilon$ .  $\epsilon \sim 10^{-4}$  seems acceptable from the tritium economy standpoint, but more strict requirement might be imposed by safety regulation.




Fig. 4.15  $TBR_r$  for different non-radioactive loss fractions  $\epsilon$ . For  $\epsilon > 10^{-3}$  the  $TBR_r$  increases exponentially for any value of TBE. Tritium self-sufficiency is unpossible unless  $TBE > 5\%$ . For  $\epsilon < 10^{-3}$  tritium self-sufficiency becomes achievable, especially for  $TBE > 2\%$ . For  $\epsilon < 10^{-5}$  no significant changes are observed if  $TBE > 1\%$ .

### 4.3.11 Implications for ARC-class FPP design and operations

Continued advances in both fuel cycle technology and plasma performance should make tritium self-sufficiency in ARC very feasible. However, the goal of this section is to provide design guidelines to ensure tritium self-sufficiency in the first generation of ARC-class FPPs even if minimal progress is made in fuel cycle or plasma operations technologies in the coming years. Table 4.5 provides the  $TBR_r$  and  $I_{\text{startup}}$  for an ARC-class FPP operating with an “ambitious” fuel cycle (targeting  $t_d = 1$  y) and a “moderate” fuel cycle (targeting  $t_d = 2$  y). For each, three cases are considered: (1) fuel cycle (FC) technology has improved significantly relative to the present day (but plasma operation technologies have not), (2) plasma operations (PO) technologies have improved significantly relative to the present day (but FC technologies have not), and (3) both the FC advances from (1) and the PO advances from (2) are implemented in a “best case” scenario. In the model presented here, improvements in the FC correspond to improvements in  $AF$ ,  $f_{\text{DIR}}$ , and  $t_p$ . Improvements in PO correspond to improvements in TBE. Results in Table 4.5 are based on a simplified view of a very complex engineering system, but they provides useful guidelines for organizing near-term R&D of ARC-class tokamaks. As was the case in [30] for DEMO, improving the TBE lowers the required  $AF$ ,  $f_{\text{DIR}}$ , and  $t_p$  to achieve a given  $t_d$  target. For both the ambitious and moderate  $t_d$  targets,  $TBR_r$  is well within the expected achievable range ( $<1.15$ ) if FC and PO technologies improve with respect to the baseline case. The neutronics analysis carried out by Bae et al. [160] reported a  $TBR_a = 1.053$  for an Inconel-718 VV in ARC. A similar result is reported in [161], which used a simplified neutronics model of an ARC-class plant and calculated a  $TBR_a = 1.07$ . We note that while  $TBR_r$  can be lowered by improving PO and FC technology, we can also increase  $TBR_a$  by optimizing plant design and materials choices (for example, if V-Cr-Ti alloys can be used for the VV,  $TBR_a$  may be as high as 1.2 [161]).

A second consideration involves the DIR, which should lower  $TBR_r$  and  $I_{\text{startup}}$  for a given FPP design at the cost of increased complexity. We note that if the TBE is sufficiently high, advantages gained from DIR are small and may not be worth the associated costs. Furthermore, the successful implementation of DIR in ARC depends on the development of metal foil pumps that process the exhaust flow. If DIR is deemed necessary in order to achieve tritium self-sufficiency, then this is an obvious argument in favor of prioritizing R&D of these pumps.

Component	T inventory	Units
Divertors	$1.4 \cdot 10^{-4}$	g
FW	$1.4 \cdot 10^{-4}$	g
Breeding zone	3	g
TES	60	g
HX	$3.5 \cdot 10^{-4}$	g
Vacuum pumps	9	g
Fuel clean-up	77	g
Detritiation system	8	g
ISS	260	g
Storage system	1.14	kg

Table 4.4 Tritium inventories in FC components at steady-state, for the baseline design parameters. The tritium inventory for the separation membrane downstream of the TES is not reported because it is negligible. The storage system inventory does not reach a steady-state, and the start-up inventory is reported as maximum value (assuming that tritium is removed at each  $t_d$ ).

Lastly, the tritium inventories in the FC components (Table 4.4) and the start-up inventories reported in Table 4.5 are reasonably achievable for the first ARC-class FPP. The initial tritium inventory required to begin operations can be provided by present reserves [133].

## 4.4 Discussion

The results suggest that, based on the current conceptual design of ARC-class FPP, tritium self-sufficiency is feasible assuming the baseline parameters of Table 4.1 are achieved. Section 4.4.1 compares the ARC fuel cycle results to those calculated (in [4]) for STEP-EBB (Encapsulated Breeding Blanket [162]), STEP-Li (liquid lithium blanket) and EU-DEMO (in [30]). Section 4.4.2 highlights the link between TBE and power production and shows why the TBE cannot be increased indefinitely. Section 4.4.3 briefly discusses the potential impact of tritium trapping, which has not been considered in this model. Section 4.4.4 highlights the next steps to increase the accuracy of fuel cycle modeling. Section 4.4.5 analyses the relationship between the tritium doubling time and FPP construction time.

Parameters	Baseline case	Ambitious ( $t_d=1$ y)			Moderate ( $t_d=2$ y)		
	No FC or PC advances	FC advances only	PO advances only	Both	FC advances only	PO advances only	Both
AF (%)	70	90	70	90	90	70	90
TBE (-)	0.02	0.02	0.1	0.1	0.02	0.1	0.1
$f_{\text{DIR}}$ (-)	0.3	0.7	0.3	0.7	0.7	0.3	0.7
$t_p$ [h]	4	1	4	1	1	4	1
$t_d$ [y]	1 or 2	1	1	1	2	2	2
$\text{TBR}_r$ (-)	1.113 ( $t_d = 1$ ) 1.067 ( $t_d = 2$ )	1.064	1.029	1.022	1.039	1.017	1.012
$I_{\text{startup}}$ [kg]	1.14 ( $t_d = 1$ ) 1.14 ( $t_d = 2$ )	1.12	0.33	0.30	1.12	0.33	0.30

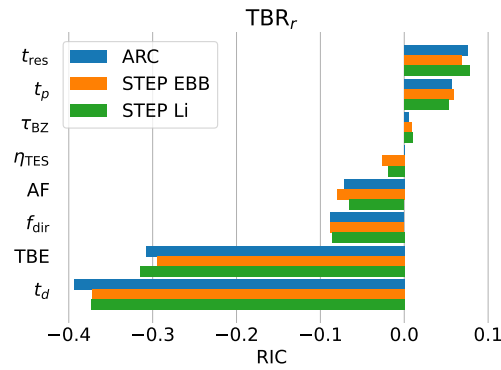
Table 4.5  $\text{TBR}_r$  and  $I_{\text{startup}}$  for an ARC-class FPP, assuming baseline design parameters, an ambitious operating scenario ( $t_d=1$  y), and a moderately ambitious operating scenario ( $t_d=2$  y). This allows to consider how parameters change if we can implement significant improvements in fuel cycle (FC) technologies, in plasma operations (PO), or in both. The advanced-FC scenario considers AF = 90%,  $f_{\text{DIR}} = 0.7$ ,  $t_p = 1$  h, and TBE = baseline case (0.02, based on the current state-of-the-art). The advanced-PO scenario considers TBE = 0.1 and all other parameters unchanged from the baseline scenario. The “both” scenario assumes that both FC and PO advances have been implemented.

#### 4.4.1 Comparison of D-T FPP designs

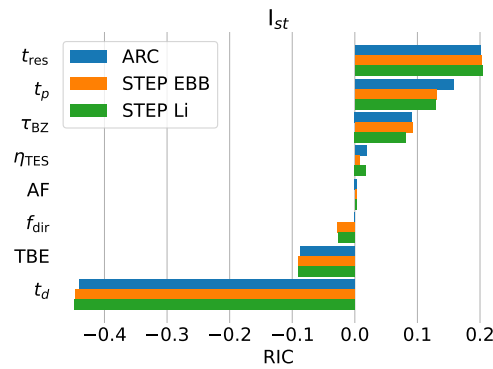
The fuel cycle analysis for DEMO presented in [30] highlighted serious vulnerabilities in the design of the 3 GW<sub>th</sub> FPP with regards to tritium self-sufficiency. For low tritium burn fractions ( $f_b = 0.36\%$ ) and fueling efficiencies ( $\eta_{\text{TES}} = 25\%$ ) that currently characterize the technologies used in the design,  $\text{TBR}_r$  is far larger than the achievable TBR deemed possible by the inherent neutronics of the plant. If such a low TBE characterizes ARC-class FPP, tritium self-sufficiency will be similarly unlikely.

Section 4.3.2 showed how  $\text{TBR}_r$  is mostly determined by the  $t_d$  target, which is independent of the fusion power. Figure 4.16 shows the relative sensitivity index coefficient (RIC) for  $\text{TBR}_r$  and  $I_{\text{startup}}$  for the three FPP designs. Practically no differences exist between ARC and other designs. Since the layouts of their IFCs are the same, and their OFC differs only for a few parameters, this is to be expected. Figure 4.16 suggests that the results presented in this analysis are applicable to other tokamak designs with different OFC components, provided that  $\tau_{\text{OFC}} \sim$  days and  $\eta_{\text{TES}} > 10\%$ .

Compact size, modular structure and the liquid immersion blanket concept bring many advantages to the ARC-class design. A very low availability factor is expected



(a)



(b)

Fig. 4.16 Relative sensitivity index coefficient (RIC) of  $TBR_r$  (a) and  $I_{startup}$ (b) of ARC, STEP-EBB and STEP-Li [4]. No significant differences are found between ARC, STEP-EBB and STEP-Li as they have similar fuel cycle designs.

for DEMO and near-term FPP [30]. Therefore, a high TBE will be required to compensate for the low availability in order to achieve tritium self sufficiency. This will require major technology advances in a short time frame with regards to plasma operations and fueling technologies. ARC-class tokamaks may be able to achieve much better availability factors from their first generation as a simple result of their characteristic design. Demountable magnets and a replaceable VV aim to minimize mean-time-to-repair and thus improve AF. The LIB can be considered a FPP lifetime component, which is undamaged by radiation damage or thermomechanical stresses. The maintenance required by the LIB is primarily chemistry control: FLiBe purity must be continuously monitored to prevent the formation aggressive radiochemical byproducts of the tritium breeding reaction (e.g, TF) and remove impurities, as this is expected to cause severe corrosion issues in FLiBe-facing components. Of

		<b>Baseline</b>	<b>FC</b>	<b>PO</b>	<b>Both</b>
ARC	$TBR_r$	1.067	1.039	1.017	1.012
	$I_{\text{startup}}$ [kg]	1.14	1.12	0.33	0.30
STEP-EBB	$TBR_r$	1.068	1.040	1.018	1.012
	$I_{\text{startup}}$ [kg]	8.94	7.50	2.36	2.20
STEP-Li	$TBR_r$	1.070	1.041	1.020	1.015
	$I_{\text{startup}}$ [kg]	9.15	7.20	2.36	2.20

Table 4.6 Comparison of D-T FPP designs for the baseline case and the three possible advanced scenarios. The points refer to the “moderate” scenarios ( $t_d = 2$  y).

course, the successful operation of the ARC-class tokamak depends on the rapid development of demountable magnet joints and LIB technology. If this happens, it becomes realistic to target an availability factor larger than 50% in ARC, which will relax requirements on TBE.

The lower fusion power of ARC does not directly lower  $TBR_r$ , but  $I_{\text{startup}}$  is significantly lowered (Table 4.6). We showed that  $I_{\text{startup}}$  in ARC depends on the IFC dynamics and the desired reserve inventory (Sections 4.3.6 and 4.3.9). A 525 MW<sub>th</sub> FPP burns much less tritium than a 3000 MW<sub>th</sub> one. Hence, less tritium is required to sustain the initial operations of the FPP, and a lower reserve inventory is needed to account for possible fuel cycle failures. This translates to start-up inventories much lower than those required for STEP or DEMO (i.e.,  $\sim 1$ - kg [4] vs  $\sim 10$ -100 kg [30]). Table 4.7 lists the design parameters,  $TBR_r$ , and  $I_{\text{startup}}$  for a “worst-case” scenario for ARC, so-called because it assumes a low AF (below what would be targeted for commercial operation), no technological advances in burn efficiency, and no implementation of DIR. In this case,  $TBR_r = 1.35$ , above  $TBR_q$  for ARC, so improvements in the fuel cycle or in plasma operations would be required to achieve at least the baseline design parameters in Table 4.1 and to achieve tritium self-sufficiency. As far as the start-up inventory is concerned,  $I_{\text{startup}} < 10$  kg even at TBE = 0.5 %,  $f_{\text{DIR}} = 0$ , and AF = 50%. Considering that ARC is expected to start operations before ITER D-T operations, the initial tritium inventory for ARC can be covered by the available tritium supply [133].

Parameter	Value	Units
AF	50	%
TBE	0.005	-
$f_{\text{DIR}}$	0	-
$t_p$	4	h
$t_d$	2	y
$\text{TBR}_r$	1.35	-
$I_{\text{startup}}$	5.4	kg

Table 4.7 ARC design parameters, required TBR and start-up inventory for a “worst-case” scenario.

Large  $I_{\text{startup}}$  does not need to be an unavoidable, intrinsic property of commercial FPPs. It can be minimized by relaxing constraints on the required tritium reserve inventory  $I_{\text{res}}$ , as shown in Figure 4.17. This is evident for the advanced fuel cycle (“FC”) scenario.  $I_{\text{res}}$  accounts for  $\approx 85\%$  of  $I_{\text{startup}}$  in the STEP-class tokamak [4]. The ARC start-up (and reserve) inventory is about  $6\times$  lower than STEP’s because  $P_{\text{fus,STEP}} \sim 6P_{\text{fus,ARC}}$ . Without a reserve inventory,  $I_{\text{startup}} < 200$  g. Decreasing or eliminating  $I_{\text{res}}$  is a key strategy for minimizing  $I_{\text{startup}}$ , and may be required for regulatory and safety purposes. This will be an important risk analysis for plant designers to consider:  $I_{\text{res}}$  allows the plant to keep running even if part of the fuel cycle breaks temporarily. However, we will have to consider whether the fuel cycle components are the most likely points of failure in the FPP. If the plant is shut down for reasons other than a fuel cycle failure, having a large reserve tritium inventory won’t turn the plant back on, but it does add potential safety liability.




Fig. 4.17  $I_{\text{startup}}$  with and without a reserve inventory at the beginning of the operations (FC scenario) for ARC, STEP-EBB and STEP-Li [4].  $I_{\text{startup}}$  of high-power-output FPPs is significantly increased if the reserve inventory is included at the beginning of the operations.

#### 4.4.2 Limitations to the maximum TBE in a D-T FPP

The impact of a high TBE on  $TBR_r$  and on tritium inventory has been highlighted multiple times in this Chapter. It might seem natural to target extremely high TBE in a FPP. However, a TBE > 10-20% would be deleterious for the plant economics. TBE depends on  $f_{\text{He,core}}$  (the helium ash fraction in the core), and, via the dilution factor,  $P_{\text{fus}}$  (the fusion power density) depends on  $f_{\text{He,core}}$  [138]. This imposes a tradeoff between efficient tritium burn and fusion power density. Figure 4.18 shows the required TBR and the fusion power density ratio as a function of TBE. The optimal TBE range is between 10% and 15%, depending on the reduction in  $P_{\text{fus}}$  that is considered acceptable for a given FPP design. A FPP with a very efficient fuel cycle (e.g. short  $t_p$  and large  $f_{\text{DIR}}$ ) might target relatively low TBE (e.g. TBE = 5%) to maximize  $P_{\text{fus}}$  and electricity production. Conversely, a FPP with a less efficient fuel cycle might choose to maximize TBE, and thus account for smaller revenues from electricity production with a cheaper fuel cycle design. The game-changing variable here is the pump selectivity  $\Sigma$ , which allows for a very high TBE with no drawbacks on the fusion power density if  $\Sigma > 3$  [138].




Fig. 4.18  $TBR_r$  and fusion power density ratio as a function of TBE for an ARC-class FPP with baseline parameters ( $\eta_{He} = 0.5$ ,  $t_d = 2y$ ). The fusion power reduction is approximately 30% at TBE = 10% for  $\Sigma = 1$ , and 10% for  $\Sigma = 3$ . TBE < 10% is therefore necessary to avoid excessive power reduction at  $\Sigma < 3$ , while 10% < TBE < 15% might be acceptable for higher  $\Sigma$ .

### 4.4.3 Potential effect of T trapping in solid components

In this analysis, hydrogen trapping has been neglected. Hydrogen can however be trapped in defects of solid materials (vacancies, grain boundaries, impurities, voids, etc.) [163, 164]. The concentration of trapped hydrogen depends on the concentration of mobile hydrogen, the material (and its inherent trapping properties), and the temperature. Dark et al. showed that neglecting trapping in WCLL breeding blankets would lead to an underestimation of the tritium inventory (between 14% and 35% in the structural materials and by a factor of 4.5 in the W first wall) [165]. Including trapping will impact (although hard to quantify) the dynamics of the fuel cycle models. Components may build up an inventory of trapped hydrogen, delaying tritium retrieval and affecting  $I_{\text{startup}}$ ,  $TBR_r$ , and/or  $t_d$ . Conducting a thorough investigation on hydrogen trapping and its potential influence on the results presented here is outside of the scope of this study. However, this point should be addressed in future work.

### 4.4.4 Reducing uncertainty in fuel cycle modeling and design

The parametric analysis presented in this work explored ARC operations across a wide range of possible operating parameters. This range should be narrowed to better focus technology development moving forward. This will require design teams to (1) specify the technologies for the IFC, TES, and HX; (2) obtain better data describing tritium transport phenomena in FLiBe and lithium; (3) quantify the attainable tritium burn efficiency through detailed plasma physics simulations and/or experiments; and (4) refine this model according to the results from (1)-(3).

The first step is a system engineering exercise similar to the one carried out by Day et al. [166] for the conceptual design of DEMO fuel cycle. Existing technologies will probably be implemented in the fuel cycle of a FOAK FPP because of the short timeframe available for development. Residence times at each step of the model can then be better-specified. This step will move the model from technology-agnostic to technology-informed.

Better understanding of tritium transport in FLiBe will be pursued by experiments like LIBRA [39]. Tritium transport properties in FLiBe are affected by large uncertainties [23]. Tritium transport through the IFC and OFC is modeled here

by Eq. (4.1), but this certainly is too simplistic. Improved datasets describing tritium solubility and permeability through FLiBe and different materials are needed to better predict residence times, improve technologies like the TES, and predict (and mitigate) tritium loss due to permeation of structural materials. For example, LIBRA will investigate tritium transport and radiochemistry in FLiBe in ARC relevant conditions, providing more precise input data to tritium transport models. Tritium extraction methods (gas sparging through salt and sweep gas flowing in tank interstitials) will be tested as well, providing a rough estimate of achievable TES efficiency. Steps like this move the fuel cycle model from physics-agnostic to physics-informed.

The third goal will be pursued by devices like SPARC, which will investigate the high-field path to practical fusion energy [167]. Preliminary results for ARC tritium burn efficiency may come even earlier through detailed plasma physics simulations. We know from [30] and the present analysis that the FPP's utilization of fuel has a profound effect on the tritium self-sufficiency of the system. These results will improve confidence in the fuel cycle analysis and system design for ARC.

The last step is a modeling exercise. The fuel cycle model is scalable and flexible thanks to its implementation in MATLAB Simulink®. A component level description can be added in a similar way to the one described by Riva et al. [150] once the uncertainties related to tritium transport properties are reduced. Alternatively, standalone component simulations might be carried out to quantify the tritium residence time. Then, the precise value for  $\tau_i$  can be implemented in the MATLAB Simulink® model, improving calculation speed.

#### **4.4.5 Doubling time, construction time, and FPP penetration into the electricity market**

In this analysis, the doubling time  $t_d$  is defined as the time required to double the FPP's start-up inventory  $I_{\text{startup}}$ . This definition assumes implicitly that the limiting factor on the construction and operation of new FPPs is the availability of tritium fuel for plant start-up. The broader definition of "doubling time" is the time required to double the number of a given FPP. We refer to this broader definition hereunder as the construction time ( $t_c$ ) for grid-connected FPP. For the first generation of FPP,  $t_d$  and  $t_c$  will hardly match. The  $t_d$  targets described in this model are affected by large

uncertainties. A minimum achievable  $t_c$  will remain an unknown data point until the first FPPs are built and operated. However, it is reasonable to expect construction times much longer than the doubling times targeted in this analysis (i.e., 1 or 2 years). For instance, the median construction time for fission power plants in 2021 was 88 months ( $\sim 7.3$  y), despite eight decades of experience building fission power plants [168]. In this paper, short doubling times are considered: 1 or 2 years. If we optimistically assume a  $t_c$  of 6 years for the construction of early FPPs, we see the obvious issue with the aggressive  $t_d$  scenarios. A doubling time  $t_d = \sim 0.3t_c$  sets unnecessarily demanding requirements on the first FPPs without a corresponding advantage in tritium economy. A large amount of tritium would need to be safely stored for years before it could be used. It seems reasonable, then, to target longer  $t_d$  for the first generation of FPPs.

However,  $t_d$  can be a misleading figure of merit in this context. If the first generation of FPPs demonstrates the technical and economical feasibility of fusion energy production, the construction of many FPP for the second generation might start in parallel, potentially regardless of whether there is immediately tritium available to start the finished plant. Also note that this analysis treats the number of FPPs as a straightforward mathematical variable. In reality, commercial FPPs will be affected by a far more complex political, social, economic, and commercial landscape than can be accounted for in a simple fuel cycle analysis. In this Chapter we are considering tritium self-sufficiency for individual FPPs; eventually, we will need to consider what it means for a fusion power sector as a whole to be tritium self-sufficient. As the fusion landscape evolves, we must incorporate more nuanced discussions of fusion economics and future energy scenarios into our understanding of tritium self-sufficiency.

# Chapter 5

## Conclusions and future work

### 5.1 Summary

This work presented a preliminary design, model, and safety assessment of two safety-critical systems (primary cooling system and fuel cycle) in an ARC-class FPP. Both the PCS and the FC layouts and models are the first attempt (in literature) to design those systems in an ARC-class FPP. The thermal-hydraulics model of the PCS is also one of the few models described in literature to integrate a deterministic and a probabilistic approach towards big data applications. In fact, the study also investigated novel approaches based on the integration of probabilistic and deterministic analyses, with a focus on machine learning algorithms for nuclear safety. DNNs were developed to showcase the potentiality of ML for FPP safety (end state prediction and accident progression), and a framework to foster ML in nuclear safety has been proposed based on the experience gained during ARC PCS safety analysis.

Chapter 2 delves deeply into the PCS, where a preliminary layout of the system is developed through a rigorous functional analysis. An FMEA is built on top of this layout to determine the relevant failure modes and failure probabilities of PCS components, and to identify the postulated initiating events. The FMEA sets the basis for a probabilistic safety assessment from a classical perspective. The most relevant PIEs involve pump and instrumentation failures, as expected from the review of past experience on molten salt and blanket systems. Nevertheless, the ultimate goal of Chapter 2 was to develop suitable tools for an integrated deterministic and

probabilistic safety assessment. The qualitative and quantitative results from the FMEA were used to inform a thermal-hydraulics model of the PCS. The TH model was developed from scratch in OpenModelica to allow integration with probabilistic routines, and was verified against available data in literature. A Monte Carlo failure injection engine was coupled with the TH model to analyze unanticipated transients started from the failures identified by the FMEA. This approach culminated in a fully automated simulator equipped with a user-friendly python API (OMPpython) for pre-processing, and a suitable output for big data applications (i.e., machine learning applications). The dynamics of the PCS were deeply analyzed, and the conditions leading to hot or cold failures of the system were identified by means of data science techniques.

The approach presented in Chapter 2 offers many advantages over a classical safety analysis approach. Once automated tools are developed, the iterations between design and safety analysis are quite straightforward from a practical perspective. The integrated approach also allows for the removal, or at least the strong limitation, of analyst bias and the transition from *postulated initiating events* to *randomly generated initiating events*. This last feature is essential because it captures the real dynamics of the system, which is defined by the failure probability and failure modes of the system components only, without any assumptions from the analyst. The simulated transients involved both safe and fault states because no system failure was defined *a priori*. Clearly, such an approach is not exempt by criticality, and it is worth mentioning the most critical issues and opportunities for future research.

First and foremost, it is important to note that the model presented in this work was not validated, as no experimental data in relevant conditions for the ARC PCS is available. However, this is not a major issue, as the model was not intended to replace any established and validated codes, but rather to investigate the workflow and requirements for a tool that leverages both deterministic and probabilistic subroutines and can be coupled with machine learning algorithms. The decision not to use a validated code was based on the goals of this work. Established codes like MELCOR or RELAP are not suitable for the applications proposed in Chapter 2 and Chapter 3. While pre-processing and post-processing of input files can be managed by developing ad-hoc routines, this requires a significant effort. Moreover, the computational cost for a single simulation is much higher, making it challenging to build a database of transients without significant computational resources. Finally, the closed-group nature of these codes is not suitable for big

data applications, as sharing models and results is impractical. This highlights the opportunity to develop modern computational tools that integrate deterministic and probabilistic approaches [58], and/or different physics phenomena (e.g., FERMI [169]).

Second, this approach may lack of conservatism. On the one hand, the classic deterministic approach lacks domain coverage as it analyzes only a few severe accidents, while on the other hand, an integrated probabilistic and deterministic approach may not pinpoint to these very critical accidents because transients are randomly sampled. A natural solution is to exploit these novel tools to support the classical safety analysis, providing a more complete view on the whole spectrum of possible accidents, and improving both the FPP safety and capacity factor.

Lastly, the use of data science tools to extract results from a large database is more susceptible to analyst interpretation and choices. While the proposed approach limits analyst bias in input, it still leaves some freedom of interpretation in the post-processing phase. The different results obtained by clustering algorithms simply by varying the number of clusters (Section 2.7) has demonstrated the need for a rigorous methodology to apply data science tools to safety analysis. This topic has been discussed in more detail in Chapter 3.

Chapter 3 presented a methodological discussion and practical applications of machine learning (ML) to nuclear safety, building upon the results obtained in Chapter 2. The key finding is that ML algorithms are capable of handling safety assessment tasks if appropriately integrated within a suitable framework. Simple ML models based on logistic regression and fully-connected neural networks were successfully employed to predict the final state of unanticipated transients simulated by the TH model of ARC PCS. Key issues such as handling false positives (negatives) were addressed by analysing different performance metrics, showing that the area under the curve of the receiver operator characteristic performs better than the widely-used accuracy, and allows to minimize the misclassification of accidental scenarios. The task was shown to be easy from the ML standpoint (i.e., the capabilities of simple ML models are sufficient to perform safety-critical tasks), demonstrating that a fully-connected neural network can aid in the decision-making process of operators during accident scenarios.

More advanced tasks were accomplished by developing deep neural networks capable of handling time series. An LSTM-based deep neural network was trained

on a large dataset (15000 samples) of transients to predict the evolution of a safety variable (FLiBe outlet blanket enthalpy) and the time-to-failure (TTF) of the system. Predicting safety variables' evolution may become critical during an accidental scenario if direct measurements are not available, while the TTF provides a powerful figure of merit, informing the operators in real-time of the predicted time before the system reaches complete failure.

The analysis carried out in Chapter 3 also highlighted criticalities in ML for safety applications. The nuclear community has little experience in ML, and the literature contains few applications, which are strongly case-dependent. Furthermore, a common framework for developing ML models for nuclear safety is lacking. Since data is required to develop ML models, the absence of a common framework limits the availability of large databases for training and testing models. Therefore, a possible framework was presented, consisting of the development of user-friendly, lightweight tools to automate simulations, standardization of the output to build shared databases, sharing codes and models, and applying best practices inherited from other safety-critical fields to build robust ML models. These steps were followed in Chapter 2 and Chapter 3, demonstrating the effectiveness of such a framework.

Chapter 4 was devoted to the quantification of the primary radiological source term in an ARC-class FPP, namely tritium. Tritium inventories in an FPP are determined by the nominal operating conditions, indicating that safety is closely tied to tritium self-sufficiency in an FPP. Therefore, a fuel cycle layout was proposed, and a dynamic, system-level model was developed to assess tritium flows and inventories. Three alternative scenarios (advanced fuel cycle, advanced plasma operations, and both) were examined to identify potential R&D directions and design targets. The fuel cycle design is constrained by two factors: the required TBR and the start-up inventory. The optimal design point is the set of design parameters (tritium processing time, availability factor, DIR fraction, reserve time, and tritium burn efficiency) that minimizes the required TBR and start-up inventory, while still doubling the initial tritium inventory within the prescribed doubling time. It has been shown that various combinations of design parameters result in low required TBR and start-up inventory, thereby allowing for the possibility of trading off some design parameters (e.g., low TBE but high DIR fraction, or vice versa) to reduce the technological risk associated with individual components (e.g., metal foil pumps). Nevertheless, the results also

indicate that improvements in some of the design parameters are necessary in the coming years to achieve tritium self-sufficiency.

The results show that tritium self-sufficiency is theoretically achievable in an ARC-class FPP thanks to unique design features, such as demountable coils, replaceable VV and liquid immersion blanket. Low tritium burn efficiencies ( $TBE < 1\%$ ) would require high TBR in the breeding zone ( $TBR_r > 1.10$ ), but the implementation of a direct internal recycling line in this case can lower the  $TBR_r$  to easier-to-attain values. For  $TBE < 1\%$ , lower DIR fractions (e.g.,  $f_{DIR} = 0.3$ ) and FPP availability (e.g.,  $AF = 70\%$ ) are acceptable from the tritium self-sufficiency standpoint, with  $TBR_r < 1.1$ . Neutronics analyses of the ARC concept indicate that  $TBR_a > TBR_r$  in these scenarios. Research on ARC-relevant plasma physics performance should prioritize achieving  $TBE \simeq 1\%$ , at a minimum.

The analysis showed that doubling time is an effective design target when carrying out FPP tritium self-sufficiency analyses. The choice of  $t_d$  has a dramatic impact on tritium self-sufficiency. Very aggressive doubling times, such as  $t_d = 1$  y, will prevent tritium self-sufficiency unless FPPs can achieve relatively high TBE ( $> 1\%$ ), most obviously because the  $TBR_r$  rapidly becomes physically unobtainable. Relaxing the target doubling time is an effective way to decrease  $TBR_r$ . However, doubling time does not exist in a vacuum: it must make sense given projected construction times for new grid-connected FPPs. Knowledge of other power plant industries (and of fusion device construction) indicates that FPP construction time is likely to be much longer than 1-2 y, especially for FOAK plants. FOAK FPPs may target slightly longer doubling times (3-4 y), making tritium self-sufficiency much easier to achieve, and decreasing the complexity and cost of the plant.

Tritium inventories in an ARC-class FPP are much lower than large-sized reactors. The total tritium inventory ranges from  $\sim 100$  g to few kilograms. A significant, but still reasonable, effort is needed to achieve a less-than-one-kilogram FPP. To do so, a large TBE ( $TBE \approx 5\%$ ) and short tritium residence times in the IFC and OFC are required. The largest tritium inventories are found in the storage system (up to kilograms, depending on the doubling time and the start-up inventory), the isotope separation system ( $\sim 250$  g), and the tritium extraction system ( $\sim 50$  g). Safety analyses must prioritize these systems, and multiple safety barriers (e.g., glovebox) and a detritiation system seem necessary.

In conclusion, conducting preliminary (conceptual) design and safety analysis in parallel proved to be effective in identifying criticalities at an early stage of the design process, resulting in a robust safety-oriented conceptual design of the primary cooling system and fuel cycle. The novelty of fusion systems presents an opportunity to develop modern computational tools and models that can incorporate machine learning in the nuclear safety field. An ARC-class fusion power plant is very well positioned with regards to tritium self-sufficiency and tritium safety. It should require a relatively low start-up tritium inventory, be able to maintain relatively low plant inventory, breed/process sufficient tritium to fuel itself, and provide tritium fuel for future plant startup.

## **5.2 Recommended future work**

### **5.2.1 Primary cooling system and thermal-hydraulics**

The primary cooling system layout includes all the components required to provide heat removal from the VV and blanket, but many other components may be present (e.g., a chemical control unit for FLiBe, a FLiBe furnace, etc). While not strictly relevant from the TH standpoint, the safety analysis would be deepened by including these components in the design. More detailed tests on fusion-specific components (e.g., first wall and divertor tiles) and ARC-class FPP components (e.g., molten salt pumps) may provide accurate failure rates to improve the probabilistic safety assessment.

From the modeling standpoint, the main criticality is the calculation of the heat transfer coefficient in the tank. Developing a simple, yet accurate model for the tank looks hard due to the complex geometry and the inherent impact of 3D flows on heat transfer. Therefore, the heat transfer coefficient as a function of the FLiBe inlet velocity may be quantified by component-level simulations and the function can be implemented in the TH model to account for heat transfer changes during transients. Also, a validation of the model should be performed as soon as relevant TH experiments provide results for the validation.

### **5.2.2 Machine learning and nuclear safety**

Future works on machine learning for nuclear safety involve very different level of applications.

At a very low level, the most suitable metrics for nuclear safety applications should be identified. For classification tasks, penalizing misclassification looks as a leading goal when choosing the metric, while the impact of different metrics on regression applications has not been assessed in this work. Also, real-time applications require very short computational time for ML models, thus requiring a careful choice of the model architecture and of the input parameters.

From the safety perspective, ML models must demonstrate their reliability and robustness. This requires a tremendous effort on topics that are quite far from the nuclear engineering background, such as ML safety, UQ on ML models, and model interpretability. Collaboration between different fields (nuclear engineering and data science) looks crucial for this task.

Connected to the previous point, and from a boarder perspective, the nuclear community should engage discussion with practitioners from other safety-relevant fields, inheriting best practices and avoiding mistakes that might jeopardise the adoption of ML in nuclear safety. The proposed framework presented in this work is just a conceptual flow diagram that must be detailed and expanded to be accessible and fruible.

### **5.2.3 Fuel cycle and tritium transport**

The analysis of the fuel cycle highlighted many R&D steps needed to precisely assess tritium inventories and tritium self-sufficiency in FPP. Tritium transport properties (diffusivity, solubility, recombination and dissociation coefficients) in fusion-relevant materials are highly uncertain. Dedicated experiments on tungsten and FLiBe systems are required to produce more accurate correlations, and tritium transport models for ARC-class FPP components must be developed to quantify the uncertainty propagation from tritium transport at component-level to system-level quantities (e.g., tritium inventories). Tritium extraction from FLiBe has to be modeled and investigated experimentally to quantify the tritium extraction efficiency as a function of the tritium concentration, since the TES efficiency strongly affects all the tritium

inventories in the OFC. More detailed modeling of fuel cycle components is also required to capture transients that can affect the fuel cycle dynamics (e.g., disruptions, plasma ramp-up and ramp-down, etc.). Also, a coupled CFD and tritium transport analysis of the heat exchanger is deemed necessary to set a tentative design that achieves an acceptable heat transfer coefficient and limits tritium losses. Lastly, the impact of tritium trapping on tritium retention in ARC FW must be assessed by a dedicated tritium transport model that accounts for all the uncertainties described above, and the production of additional trapping sites due to neutron irradiation damages needs to be quantified.

# References

- [1] G. Ferrero, S. Meschini, R. Testoni, M. Zucchetti, Exploration of ARC-class reactor vessel and divertor cooling system, *Fusion Engineering and Design* 192 (2023) 113818.
- [2] S. Segantin, S. Meschini, R. Testoni, M. Zucchetti, Preliminary investigation of neutron shielding compounds for ARC-class tokamaks, *Fusion Engineering and Design* 185 (2022) 113335.
- [3] A. Kuang, N. Cao, A. J. Creely, C. A. Dennett, J. Hecla, B. LaBombard, R. A. Tinguely, E. A. Tolman, H. Hoffman, M. Major, et al., Conceptual design study for heat exhaust management in the ARC fusion pilot plant, *Fusion Engineering and Design* 137 (2018) 221–242.
- [4] S. Meschini, S. Ferry, R. Delaporte-Mathurin, D. Whyte, Modeling and analysis of the tritium fuel cycle for ARC- and step-class d-t fpps (2023).
- [5] B. Sorbom, J. Ball, T. Palmer, F. Mangiarotti, J. Sierchio, P. Bonoli, C. Kasten, D. Sutherland, H. Barnard, C. Haakonsen, et al., ARC: A compact, high-field, disassemblable fusion nuclear science facility and demonstration power plant, in: *APS Division of Plasma Physics Meeting Abstracts*, volume 2014, 2014, pp. TP8–005.
- [6] R. d. E. Atkinson, F. G. Houtermans, Zur frage der aufbaumöglichkeit der elemente in sternern, *Zeitschrift für Physik* 54 (1929) 656–665.
- [7] M. L. E. Oliphant, P. Harteck, E. Rutherford, Transmutation effects observed with heavy hydrogen, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 144 (1934) 692–703.
- [8] H. A. Bethe, Energy production in stars, *Physical Review* 55 (1939) 434.
- [9] J. Luxon, T. Simonen, R. Stambaugh, D.-D. Team, Overview of the diiii-d fusion science program, *Fusion Science and Technology* 48 (2005) 807–827.
- [10] J. Pamela, E. R. Solano, J. E. Contributors, Overview of jet results, *Nuclear fusion* 43 (2003) 1540.

- [11] Y. Kamada, P. Barabaschi, S. Ishida, et al., Progress of the jt-60sa project, *Nuclear Fusion* 53 (2013) 104010.
- [12] E. Fermi, The development of the first chain reacting pile, *Proceedings of the American Philosophical Society* 90 (1946) 20–24.
- [13] F. D’auria, B.-S. ANIS, A. Petruzzi, A. D. Nevo, State of the art in using best estimate calculation tools in nuclear technology, *Nuclear engineering and technology* 38 (2006) 11–32.
- [14] U. N. R. Commission, et al., Reactor safety study: An assessment of accident risks in us commercial nuclear power plants, wash-1400 (nureg-75/014), October. Available from National Technical Information Service, Springfield, VA 22161 (1975).
- [15] U. N. R. Commission, et al., An approach for using probabilistic risk assessment in risk-informed decisions on plant-specific changes to the licensing basis, *Regulatory Guide*, 1.174, Rev. 1 (2002).
- [16] Integrated Safety Assessment of Nuclear Installations by the Regulatory Body, number 1940 in TECDOC Series, INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2021. URL: <https://www.iaea.org/publications/14775/integrated-safety-assessment-of-nuclear-installations-by-the-regulatory-body>.
- [17] E. Zio, Integrated deterministic and probabilistic safety assessment: concepts, challenges, research directions, *Nuclear Engineering and Design* 280 (2014) 413–419.
- [18] S. Segantin, R. Testoni, M. Zucchetti, ARC reactor–neutron irradiation analysis, *Fusion Engineering and Design* 159 (2020) 111792.
- [19] J. P. Sharpe, D. Petti, H.-W. Bartels, A review of dust in fusion devices: Implications for safety and operational performance, *Fusion Engineering and Design* 63 (2002) 153–163.
- [20] J. Martin, M. Bacharis, M. Coppins, G. Counsell, J. Allen, Modelling dust transport in tokamaks, *Europhysics Letters* 83 (2008) 65001.
- [21] A. Y. Pigarov, S. Krashenninikov, T. Soboleva, T. Rognlien, Dust-particle transport in tokamak edge plasmas, *Physics of plasmas* 12 (2005) 122508.
- [22] S. Zinkle, Summary of physical properties for lithium, pb-17li, and (lif) n• bef2 coolants, in: APEX study meeting, Sandia National Laboratories, 1998, pp. 1–8.
- [23] P. W. Humrickhouse, T. F. Fuerst, Tritium Transport Phenomena in Molten-Salt Reactors, Technical Report, Idaho National Lab.(INL), Idaho Falls, ID (United States), 2020.

- [24] B. Sorbom, J. Ball, T. Palmer, F. Mangiarotti, J. Sierchio, P. Bonoli, C. Kasten, D. Sutherland, H. Barnard, C. Haakonsen, et al., ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Engineering and Design* 100 (2015) 378–405.
- [25] S. Segantin, R. Testoni, M. Zucchetti, Neutronic comparison of liquid breeders for ARC-like reactor blankets, *Fusion Engineering and Design* 160 (2020) 112013.
- [26] J. W. Bae, E. Peterson, J. Shimwell, Ar reactor neutronics multi-code validation, *Nuclear Fusion* 62 (2022) 066016.
- [27] G. Ferrero, S. Meschini, R. Testoni, A preliminary CFD and tritium transport analysis for ARC blanket, *Fusion Science and Technology* 78 (2022) 617–630.
- [28] S. Meschini, R. Testoni, S. Segantin, M. Zucchetti, ARC reactor: A preliminary tritium environmental impact study, *Fusion Engineering and Design* 167 (2021) 112340.
- [29] S. Segantin, A. Bersano, N. Falcone, R. Testoni, Exploration of power conversion thermodynamic cycles for ARC fusion reactor, *Fusion Engineering and Design* 155 (2020) 111645.
- [30] M. Abdou, M. Riva, A. Ying, C. Day, A. Loarte, L. Baylor, P. Humrickhouse, T. F. Fuerst, S. Cho, Physics and technology considerations for the deuterium–tritium fuel cycle and conditions for tritium fuel self sufficiency, *Nuclear fusion* 61 (2020) 013001.
- [31] D. Maisonnier, Rami: The main challenge of fusion nuclear technologies, *Fusion Engineering and Design* 136 (2018) 1202–1208.
- [32] Deterministic Safety Analysis for Nuclear Power Plants, number SSG-2 (Rev.1) in Specific Safety Guides, INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2019. URL: <https://www.iaea.org/publications/12335/deterministic-safety-analysis-for-nuclear-power-plants>.
- [33] P. Fritzson, V. Engelson, Modelica—a unified object-oriented language for system modeling and simulation, in: *European Conference on Object-Oriented Programming*, 1998, p. 67.
- [34] D. E. Holcomb, A. J. Huning, A. Guler Yigitoglu, M. D. Muhlheim, W. P. Poore III, G. F. Flanagan, Molten Salt Reactor Initiating Event and Licensing Basis Event Workshop Summary, Technical Report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2019.
- [35] S. Beall, P. Haubenreich, R. Lindauer, J. Tallackson, Msre design and operations report. part v. reactor safety analysis report, Technical Report, Oak Ridge National Lab., Tenn., 1964.
- [36] P. N. Haubenreich, J. Engel, Experience with the molten-salt reactor experiment, *Nuclear Applications and technology* 8 (1970) 118–136.

- [37] R. Guymon, MSRE systems and components performance, Technical Report, ed. and comp.; Oak Ridge National Lab., Tenn.(USA), 1973.
- [38] L. C. Cadwallader, Preliminary failure modes and effects analysis of the us dcll test blanket module, 2007. doi:[doi : 10.2172/918191](https://doi.org/10.2172/918191).
- [39] S. E. Ferry, K. B. Woller, E. E. Peterson, C. Sorensen, D. G. Whyte, The libra experiment: Investigating robust tritium accountancy in molten flibe exposed to a dt fusion neutron spectrum, *Fusion Science and Technology* (2022) 1–23.
- [40] T. Allen, S. Ball, E. Blandford, T. Downar, G. Flanagan, C. Forsberg, E. Greenspan, D. Holcomb, L.-W. Hu, R. Matzie, et al., Fluoride-salt-cooled, high-temperature reactor (fhr) subsystems definition, functional requirement definition, and licensing basis event (lbe) identification white paper, Nuclear Energy University Programs, US Department of Energy, 2013.
- [41] C. Foletti, P. Batistoni, L. Boccaccini, C. Fazio, F. Gabrielli, A. Li Puma, M. Mariani, F. Moretti, V. Peluso, M. Polidori, et al., Enea experience in lbe technology, *Journal of Nuclear Materials* 356 (2006) 264. doi:[10.1016/j.jnucmat.2006.05.020](https://doi.org/10.1016/j.jnucmat.2006.05.020).
- [42] H. Feuerstein, K. Helm, R. Scheuerer, F. Serr, Tritex, a forced convection loop with pb-17li, *Journal of Nuclear Materials* 155 (1988) 520. doi:[10.1016/0022-3115\(88\)90302-9](https://doi.org/10.1016/0022-3115(88)90302-9).
- [43] H. Feuerstein, R. Werner, F. Serr, Tritex a ferritic steel loop with pb-158li facility and operation, 1999. URL: [http://inis.iaea.org/search/search.aspx?orig\\_q=RN:30042852](http://inis.iaea.org/search/search.aspx?orig_q=RN:30042852).
- [44] A. V. Drobyshev, B. M. Terent'ev, Some problems of operation and design of liquid-metal loops, *Liq. Met.*, NASA TT F-522, National Aeronautics and Space Administration Technical Translation series (1969).
- [45] B. F. Gromov, Y. A. Kazanskii, V. G. Muratov, E. V. Osipov, The analysis of operating experience of reactor installations using lead-bismuth coolant and accidents happened, in: *Proc. Heavy Liq. Met. Cool. Nucl. Technol*, volume 1, 1998.
- [46] J. G. Walford, Five years experience of the operation of the sodium circuits of the downreay prototype fast reactor (inis-mf-4857), 1978. URL: [http://inis.iaea.org/search/search.aspx?orig\\_q=RN:10449707](http://inis.iaea.org/search/search.aspx?orig_q=RN:10449707).
- [47] L. C. Cadwallader, Liquid metal, gas, molten salt, and organic cooling system operating experience review for fusion applications, 1999.
- [48] V. Kodur, Properties of concrete at elevated temperatures, *ISRN Civ. Eng.* 2014 (2014) 468510. doi:[10.1155/2014/468510](https://doi.org/10.1155/2014/468510).
- [49] D. Smith, B. Loomis, D. Diercks, Vanadium-base alloys for fusion reactor applications—a review, *Journal of nuclear materials* 135 (1985) 125–139.

- [50] International Atomic Energy Agency, Component Reliability Data for Use in Probabilistic Safety Assessment, IAEA-TECDOC-478, Technical Report, Vienna, 1988.
- [51] L. C. Cadwallader, Selected component failure rate values from fusion safety assessment tasks, Technical Report, 1998. URL: [https://inis.iaea.org/search/search.aspx?orig\\_q=RN:31031699](https://inis.iaea.org/search/search.aspx?orig_q=RN:31031699).
- [52] S. A. Eide, et al., GENERIC, COMPONENT FAILURE DATA BASE FOR LIGHT WATER AND LIQUID SODIUM REACTOR PRAs, Technical Report, 1990. doi:10.2172/975488.
- [53] D. of Defense Standard Practice, System safety, mil-std-882e, 2012.
- [54] B. M. Chisholm, et al., A unique molten salt reactor feature – the freeze valve system: Design, operating experience, and reliability, Nucl. Eng. Des. 368 (2020) 110803. doi:10.1016/j.nucengdes.2020.110803.
- [55] A. Mosleh, Common cause failures: an analysis methodology and examples, Reliability Engineering & System Safety 34 (1991) 249–292.
- [56] R. A. Berry, et al., RELAP-7 Theory Manual, Technical Report, Idaho National Lab.(INL), Idaho Falls, ID (United States), 2018.
- [57] R. O. Gauntt, et al., MELCOR Computer Code Manuals, Technical Report, Sandia Natl. Lab. NUREG/CR, 6119, 2000.
- [58] V. Destino, et al., Identification of LOFA precursors in iter superconducting magnet cryogenic cooling circuit, Reliability Engineering & System Safety 209 (2021) 107426. doi:10.1016/j.ress.2020.107426.
- [59] E. Zio, F. Di Maio, Processing dynamic scenarios from a reliability analysis of a nuclear power plant digital instrumentation and control system, Annals of Nuclear Energy 36 (2009) 1386–1395.
- [60] A. Cammi, et al., Modelling and control strategy of the italian Ibe-xads, Progress in Nuclear Energy 48 (2006) 578–594.
- [61] N. Siu, Risk assessment for dynamic systems: an overview, Reliability Engineering & System Safety 43 (1994) 43–52.
- [62] T. Aldemir, et al., A benchmark implementation of two dynamic methodologies for the reliability modeling of digital instrumentation and control systems, Technical Report, Nureg/Cr-6985. Washington, DC US Nucl. Regul. Comm., 2009.
- [63] P.-E. Labeau, et al., Dynamic reliability: towards an integrated platform for probabilistic risk assessment, Reliability Engineering & System Safety 68 (2000) 219–232.

- [64] P. Fritzson, P. Aronsson, H. Lundvall, K. Nyström, A. Pop, L. Saldamli, D. Broman, The openmodelica modeling, simulation, and development environment, in: 46th Conference on Simulation and Modelling of the Scandinavian Simulation Society (SIMS2005), Trondheim, Norway, October 13-14, 2005, 2005.
- [65] F. Casella, A. Leva, Modelica open library for power plant simulation: design and experimental validation, Proceedings of the 2003 American Control Conference, 2003. (2003).
- [66] F. Casella, A. Leva, Modelling of thermo-hydraulic power generation processes using Modelica, Mathematical and Computer Modelling of Dynamical Systems 12 (2006) 19–33.
- [67] K. Wang, C. Xiao, J. Wang, D. Zhu, Y. Ma, D. Yang, Performance improvement of a liquid molten salt pump: Geometry optimization and experimental verification, Symmetry 11 (2019) 423.
- [68] G. Ferrero, R. Testoni, M. Zucchetti, Impact assessment of radiative heat transport in ARC-class reactor flibe liquid immersion blanket, Nuclear Science and Engineering (Submitted).
- [69] P. K. Romano, B. Forget, The openmc monte carlo particle transport code, Annals of Nuclear Energy 51 (2013) 274–281.
- [70] R. H. Winterton, Where did the dittus and boelter equation come from?, International journal of heat and mass transfer 41 (1998) 809–810.
- [71] B. R. Munson, D. F. Young, T. H. Okiishi, W. Huebsch, Fluid Mechanics, Wiley Singapore, 2013.
- [72] T. Bergman, A. Lavine, F. Incropera, D. DeWitt, Fundamentals of Heat and Mass Transfer, 7th Edn, Wiley, 2011.
- [73] V. Gnielinski, New equations for heat and mass transfer in turbulent pipe and channel flow, International Chemical Engineering 16 (1976) 359–368.
- [74] T. Pinna, R. Caporali, G. Cambi, L. Burgazzi, A. Poucet, M. Porfiri, Failure mode and effect analysis on iter heat transfer systems, Fusion Engineering and Design 42 (1998) 431–436.
- [75] T. Pinna, D. Carloni, A. Carpignano, S. Ciattaglia, J. Johnston, M. Porfiri, L. Savoldi, N. Taylor, G. Sobrero, A. Ugenti, et al., Identification of accident sequences for the demo plant, Fusion Engineering and Design 124 (2017) 1277–1280.
- [76] W. E. Vesely, F. F. Goldberg, N. H. Roberts, D. F. Haasl, Fault tree handbook, Technical Report, Nuclear Regulatory Commission Washington DC, 1981.
- [77] M. H. Kalos, P. A. Whitlock, Monte carlo methods, John Wiley & Sons, 2009.

- [78] S. Aranganayagi, K. Thangavel, Clustering categorical data using silhouette coefficient as a relocating measure, in: International Conference on Computational Intelligence and Multimedia Applications (ICCIMA 2007), volume 2, 2007, p. 13.
- [79] S. Ren, K. He, R. Girshick, J. Sun, Faster R-CNN: Towards real-time object detection with region proposal networks, in: Advances in neural information processing systems, volume 28, 2015.
- [80] J. Lemley, S. Bazrafkan, J. J. Corso, Deep learning for consumer devices and services: Pushing the limits for machine learning, artificial intelligence, and computer vision., IEEE Consumer Electronics Magazine 6 (2017) 48–56. doi:10.1109/MCE.2016.2640698.
- [81] W. Schwarting, J. Alonso-Mora, M. F"arber, et al., Planning and decision-making for autonomous vehicles, Annual Review of Control, Robotics, and Autonomous Systems 1 (2018) 187–210.
- [82] J. Janai, F. Güney, A. Behl, et al., Computer vision for autonomous vehicles: Problems, datasets and state of the art, Foundations and Trends® in Computer Graphics and Vision 12 (2020) 1–308.
- [83] S. Aradi, Survey of deep reinforcement learning for motion planning of autonomous vehicles, IEEE Transactions on Intelligent Transportation Systems (2020).
- [84] T. Fischer, C. Krauss, Deep learning with long short-term memory networks for financial market predictions, European Journal of Operational Research 270 (2018) 654–669.
- [85] O. B. Sezer, A. M. Ozbayoglu, Algorithmic financial trading with deep convolutional neural networks: Time series to image conversion approach, Applied Soft Computing 70 (2018) 525–534.
- [86] J. Ma, J. Jiang, Applications of fault detection and diagnosis methods in nuclear power plants: A review, Progress in Nuclear Energy 53 (2011) 255–266.
- [87] Z. Xu, J. H. Saleh, Machine learning for reliability engineering and safety applications: Review of current status and future opportunities, Reliability Engineering System Safety 211 (2021) 107530.
- [88] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (2015) 436.
- [89] R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction, MIT press, 2018.
- [90] A. Krizhevsky, I. Sutskever, G. E. Hinton, ImageNet classification with deep convolutional neural networks, Advances in neural information processing systems 25 (2012).

- [91] S. Ferdowsi, Z. Navabi, N. Meskin, Robust deep reinforcement learning for security and safety in autonomous vehicle systems, in: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 2018, p. 307.
- [92] C. P. Andriotis, K. G. Papakonstantinou, Deep reinforcement learning driven inspection and maintenance planning under incomplete information and constraints, *Reliability Engineering & System Safety* 212 (2021) 107551.
- [93] T. M. Mitchell, *Machine Learning*, McGraw Hill, 1997.
- [94] R. E. Wright, *Logistic regression* (1995).
- [95] D. W. Hosmer Jr, S. Lemeshow, R. X. Sturdivant, *Applied logistic regression*, John Wiley & Sons, 2013.
- [96] J. A. Suykens, J. Vandewalle, Least squares support vector machine classifiers, *Neural processing letters* 9 (1999) 293–300.
- [97] W. S. Noble, What is a support vector machine?, *Nature biotechnology* 24 (2006) 1565–1567.
- [98] P. H. Swain, H. Hauska, The decision tree classifier: Design and potential, *IEEE Transactions on Geoscience Electronics* 15 (1977) 142–147.
- [99] J. R. Quinlan, Learning decision tree classifiers, *ACM Computing Surveys (CSUR)* 28 (1996) 71–73.
- [100] S. Dudani, The distance-weighted k-nearest-neighbor rule, *IEEE Transactions on Systems, Man, and Cybernetics SMC-6* (1976) 325–327.
- [101] L. E. Peterson, K-nearest neighbor, *Scholarpedia* 4 (2009) 1883.
- [102] L. Breiman, Random forests, *Machine Learning* 45 (2001) 5–32.
- [103] I. Goodfellow, Y. Bengio, A. Courville, *Deep learning*, MIT Press, 2016.
- [104] S. B. Kotsiantis, I. Zaharakis, P. Pintelas, Supervised machine learning: A review of classification techniques, *Emerging artificial intelligence applications in computer engineering* 160 (2007) 3–24.
- [105] D. C. Montgomery, E. A. Peck, G. G. Vining, *Introduction to linear regression analysis*, John Wiley & Sons, 2021.
- [106] A. J. Smola, B. Schölkopf, A tutorial on support vector regression, *Statistics and Computing* 14 (2004) 199–222.
- [107] C. Williams, C. Rasmussen, Gaussian processes for regression, in: *Advances in Neural Information Processing Systems*, 1995, pp. 514–520.
- [108] J. Schmidt-Hieber, Nonparametric regression using deep neural networks with relu activation function, *Annals of Statistics* 48 (2020) 1875–1903.

- [109] J. A. Hartigan, M. A. Wong, Algorithm as 136: A K-means clustering algorithm, *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28 (1979) 100–108.
- [110] F. Murtagh, P. Contreras, Algorithms for hierarchical clustering: an overview, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2 (2012) 86–97.
- [111] H.-P. Kriegel, P. Kröger, J. Sander, A. Zimek, Density-based clustering, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 1 (2011) 231–240.
- [112] D. A. Reynolds, Gaussian mixture models, in: *Encyclopedia of Biometrics*, Springer, 2009, pp. 659–663.
- [113] M. Caron, P. Bojanowski, A. Joulin, M. Douze, Deep clustering for unsupervised learning of visual features, in: *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 132–149.
- [114] A. K. Jain, M. N. Murty, P. J. Flynn, Data clustering: a review, *ACM Computing Surveys (CSUR)* 31 (1999) 264–323.
- [115] A. Saxena, M. Prasad, A. Gupta, N. Bharill, O. P. Patel, A. Tiwari, M. J. Er, W. Ding, C.-T. Lin, A review of clustering techniques and developments, *Neurocomputing* 267 (2017) 664–681.
- [116] T. Kohonen, The self-organizing map, *Proceedings of the IEEE* 78 (1990) 1464–1480.
- [117] I. Syarif, M. Ngadiman, A. Abdullah, et al., Unsupervised clustering approach for network anomaly detection, in: *International Conference on Networked Digital Technologies*, Springer, 2012, pp. 135–145.
- [118] R. Zhang, F. Jiang, One class support vector machine for anomaly detection in the communication network performance data, in: *Proceedings of the 5th conference on Applied electromagnetics, wireless and optical communications*, ACM, 2007, pp. 31–34.
- [119] A. Nassif, E. Mezghani, M. Chikh, Y. El Hassani, A. Ait Ouahman, M. Ait Ouahman, Machine learning for anomaly detection: A systematic review, *IEEE Access* 9 (2021) 51661–51677.
- [120] H. Daumé III, D. Marcu, Domain adaptation for statistical classifiers, *Journal of Artificial Intelligence Research* 26 (2006) 101–126.
- [121] D. J. Hand, C. Anagnostopoulos, When is the area under the receiver operating characteristic curve an appropriate measure of classifier performance?, *Pattern Recognition Letters* 34 (2013) 492–495.
- [122] A. P. Bradley, The use of the area under the roc curve in the evaluation of machine learning algorithms, *Pattern recognition* 30 (1997) 1145–1159.

- [123] T. Saito, M. Rehmsmeier, The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets, *PloS one* 10 (2015) e0118432.
- [124] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural computation* 9 (1997) 1735–1780.
- [125] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., Scikit-learn: Machine learning in python, the *Journal of machine Learning research* 12 (2011) 2825–2830.
- [126] J. Li, M. Lin, Research on robustness of five typical data-driven fault diagnosis models for nuclear power plants, *Annals of Nuclear Energy* 165 (2022) 108639.
- [127] R. Liu, B. Yang, E. Zio, X. Chen, Artificial intelligence for fault diagnosis of rotating machinery: A review, *Mechanical Systems and Signal Processing* 108 (2018) 33–47.
- [128] L. Guo, Y. Lei, S. Xing, T. Yan, N. Li, Deep convolutional transfer learning network: A new method for intelligent fault diagnosis of machines with unlabeled data, *IEEE Transactions on Industrial Electronics* 66 (2018) 7316–7325.
- [129] Y. Lei, F. Jia, J. Lin, S. Xing, S. X. Ding, An intelligent fault diagnosis method using unsupervised feature learning towards mechanical big data, *IEEE Transactions on Industrial Electronics* 63 (2016) 3137–3147.
- [130] T. Ince, S. Kiranyaz, L. Eren, M. Askar, M. Gabbouj, Real-time motor fault detection by 1-d convolutional neural networks, *IEEE Transactions on Industrial Electronics* 63 (2016) 7067–7075.
- [131] B. Boyack, I. Catton, R. Duffey, K. Katsma, G. Lellouche, S. Levy, G. Wilson, N. Zuber, Quantifying reactor safety margins part 1: an overview of the code scaling, applicability, and uncertainty evaluation methodology, *Nuclear Engineering and Design* 119 (1990) 1–15.
- [132] L. Briggs, Uncertainty quantification approaches for advanced reactor analyses. (2009).
- [133] M. Kovari, M. Coleman, I. Cristescu, R. Smith, Tritium resources available for fusion reactors, *Nuclear Fusion* 58 (2017) 026010.
- [134] L. Baylor, T. Jernigan, P. Parks, G. Antar, N. Brooks, S. Combs, D. Fehling, C. Foust, W. Houlberg, G. Schmidt, Comparison of deuterium pellet injection from different locations on the diii-d tokamak, *Nuclear fusion* 47 (2007) 1598.
- [135] S. K. Combs, L. R. Baylor, S. J. Meitner, J. Caughman, D. A. Rasmussen, S. Maruyama, Overview of recent developments in pellet injection for iter, *Fusion Engineering and Design* 87 (2012) 634–640.

- [136] U. Fischer, L. Boccaccini, F. Cismondi, M. Coleman, C. Day, Y. Hörstensmeyer, F. Moro, P. Pereslavtsev, Required, achievable and target tbr for the european demo, *Fusion Engineering and Design* 155 (2020) 111553.
- [137] R. J. Pearson, A. B. Antoniazzi, W. J. Nuttall, Tritium supply and use: a key issue for the development of nuclear fusion energy, *Fusion Engineering and Design* 136 (2018) 1140–1148.
- [138] D. Whyte, S. Ferry, S. Meschini, R. Delaporte-Mathurin, Limitations on tritium burn efficiency in magnetic fusion (2023).
- [139] C. W. Forsberg, S. Lam, D. M. Carpenter, D. G. Whyte, R. Scarlat, C. Contescu, L. Wei, J. Stempien, E. Blandford, Tritium control and capture in salt-cooled fission and fusion reactors: status, challenges, and path forward, *Nuclear Technology* 197 (2017) 119–139.
- [140] K. Nordlund, C. Björkas, T. Ahlgren, A. Lasa, A. E. Sand, Multiscale modelling of plasma–wall interactions in fusion reactor conditions, *Journal of Physics D: Applied Physics* 47 (2014) 224018.
- [141] R. A. Causey, Hydrogen isotope retention and recycling in fusion reactor plasma-facing components, *Journal of Nuclear Materials* 300 (2002) 91–117.
- [142] A. Widdowson, E. Alves, A. Baron-Wiechec, N. Barradas, N. Catarino, J. Coad, V. Corregidor, A. Garcia-Carrasco, K. Heinola, S. Koivuranta, et al., Overview of the jet iter-like wall divertor, *Nuclear materials and Energy* 12 (2017) 499–505.
- [143] Y. Hirohata, Y. Oya, H. Yoshida, Y. Morimoto, T. Arai, K. Kizu, J. Yagyu, K. Masaki, Y. Gotoh, K. Okuno, et al., Depth profile and retention of hydrogen isotopes in graphite tiles used in the w-shaped divertor of jt-60u, *Journal of nuclear materials* 329 (2004) 785–789.
- [144] V. Rohde, M. Mayer, V. Mertens, R. Neu, K. Sugiyama, et al., Dynamic and static deuterium inventory in asdex upgrade with tungsten first wall, *Nuclear fusion* 49 (2009) 085031.
- [145] C. Day, T. Giegerich, The direct internal recycling concept to simplify the fuel cycle of a fusion power plant, *Fusion Engineering and Design* 88 (2013) 616–620.
- [146] B. Peters, S. Hanke, C. Day, Metal foil pump performance aspects in view of the implementation of direct internal recycling for future fusion fuel cycles, *Fusion Engineering and Design* 136 (2018) 1467–1471.
- [147] T. Giegerich, C. Day, The kalpurex-process—a new vacuum pumping process for exhaust gases in fusion power plants, *Fusion Engineering and Design* 89 (2014) 1476–1481.
- [148] M. A. Abdou, E. Vold, C. Gung, M. Youssef, K. Shin, Deuterium-tritium fuel self-sufficiency in fusion reactors, *Fusion Technology* 9 (1986) 250–285.

- [149] A. Ying, H. Zhang, B. Merrill, M.-Y. Ahn, S. Cho, Breeding blanket system design implications on tritium transport and permeation with high tritium ion implantation: a matlab/simulink, comsol integrated dynamic tritium transport model for hccr tbs, *Fusion Engineering and Design* 136 (2018) 1153–1160.
- [150] M. Riva, A. Ying, M. Abdou, M.-Y. Ahn, S. Cho, Impact of outer fuel cycle tritium transport on initial start-up inventory for next fusion devices, *Fusion Science and Technology* 75 (2019) 1037–1045.
- [151] M. Coleman, Y. Hörstensmeyer, F. Cismondi, Demo tritium fuel cycle: performance, parameter explorations, and design space constraints, *Fusion Engineering and Design* 141 (2019) 79–90.
- [152] D. M. Hamby, A review of techniques for parameter sensitivity analysis of environmental models, *Environmental monitoring and assessment* 32 (1994) 135–154.
- [153] M. Zucchetti, Z. Hartwig, S. Meschini, S. Segantin, R. Testoni, D. Whyte, ARC reactor: Radioactivity safety assessment and preliminary environmental impact study, *Fusion Engineering and Design* 162 (2021) 112132.
- [154] B. Mignacca, G. Locatelli, Economics and finance of molten salt reactors, *Progress in Nuclear Energy* 129 (2020) 103503.
- [155] U. E. I. A. (EIA), Electric power monthly, average price of electricity to ultimate customers by end-use sector, 2022. [https://www.eia.gov/electricity/monthly/epm\\_table\\_grapher.php?t=epmt\\_5\\_6\\_a](https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_6_a), Last accessed on 21/09/2022.
- [156] M. Abdou, N. B. Morley, S. Smolentsev, A. Ying, S. Malang, A. Rowcliffe, M. Ulrickson, Blanket/first wall challenges and required r&d on the pathway to demo, *Fusion Engineering and Design* 100 (2015) 2–43.
- [157] J. L. Anderson, J. R. Bartlit, R. V. Carlson, D. O. Coffin, F. A. Damiano, R. H. Sherman, R. S. Willms, H. Yoshida, T. Yamanishi, T. Naito, et al., Experience of TSTA milestone runs with 100 grams-level of tritium, *Fusion Technology* 14 (1988) 438–443.
- [158] F. Satoshi, N. Masabumi, S. Akio, Calculation of recovery rates of tritium from flibe blanket, *Fusion technology* 39 (2001) 1073–1077.
- [159] S. Fukada, M. Nishikawa, A. Sagara, T. Terai, Mass-transport properties to estimate rates of tritium recovery from flibe blanket, *Fusion science and technology* 41 (2002) 1054–1058.
- [160] J. W. Bae, E. Peterson, J. Shimwell, ARC reactor neutronics multi-code validation, *Nuclear Fusion* (2022).
- [161] S. Segantin, R. Testoni, Z. Hartwig, D. Whyte, M. Zucchetti, Optimization of tritium breeding ratio in ARC reactor, *Fusion Engineering and Design* 154 (2020) 111531.

- [162] J. Fradera, S. Sádaba, F. Calvo, S. Ha, S. Merriman, P. Gordillo, J. Connell, A. Elfaraskoury, B. Echeveste, Pre-conceptual design of an encapsulated breeder commercial blanket for the step fusion reactor, *Fusion Engineering and Design* 172 (2021) 112909.
- [163] R. Delaporte-Mathurin, H. Yang, J. Denis, J. Dark, E. A. Hodille, G. D. Temmerman, X. Bonnin, J. Mougenot, Y. Charles, H. Bufferand, G. Ciruolo, C. Grisolia, Fuel retention in WEST and ITER divertors based on FESTIM monoblock simulations, *Nuclear Fusion* 61 (2021) 126001. URL: <https://doi.org/10.1088/1741-4326/ac2bbd>. doi:10.1088/1741-4326/ac2bbd, publisher: IOP Publishing.
- [164] E. Hodille, Study and modeling of the deuterium trapping in ITER relevant materials, Ph.D. thesis, Aix-Marseille, 2016. URL: <https://tel.archives-ouvertes.fr/tel-01477426>.
- [165] J. Dark, R. Delaporte-Mathurin, Y. Charles, E. A. Hodille, C. Grisolia, J. Mougenot, Influence of hydrogen trapping on WCLL breeding blanket performances, *Nuclear Fusion* 61 (2021) 116076. URL: <https://doi.org/10.1088/1741-4326/ac28b0>. doi:10.1088/1741-4326/ac28b0, publisher: IOP Publishing.
- [166] C. Day, B. Butler, T. Giegerich, P. Lang, R. Lawless, B. Meszaros, Consequences of the technology survey and gap analysis on the EU-DEMO R&D programme in tritium, matter injection and vacuum, *Fusion Engineering and Design* 109 (2016) 299–308.
- [167] A. Creely, M. Greenwald, S. Ballinger, D. Brunner, J. Canik, J. Doody, T. Fülöp, D. Garnier, R. Granetz, T. Gray, et al., Overview of the SPARC tokamak, *Journal of Plasma Physics* 86 (2020).
- [168] W. N. Association, World Nuclear Performance Report 2022, 2022. <https://www.world-nuclear.org/getmedia/9dafaf70-20c2-4c3f-ab80-f5024883d9da/World-Nuclear-Performance-Report-2022.pdf.aspx>, Last accessed on 24/02/2023.
- [169] V. Badalassi, A. Sircar, J. Solberg, J. Bae, K. Borowiec, P. Huang, S. Smolentsev, E. Peterson, Fermi: Fusion Energy Reactor Models Integrator, *Fusion Science and Technology* (2023) 1–35.
- [170] U. Chandra, Managing modernization of nuclear power plant instrumentation and control systems, IAEA-TECDOC 1389 (2004) 35–48.
- [171] H. Xu, Y. Gao, F. Yu, T. Darrell, End-to-end learning of driving models from large-scale video datasets, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2174–2182.
- [172] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al., End to end learning for self-driving cars, *arXiv preprint arXiv:1604.07316* (2016).

- [173] C. Richter, N. Roy, Safe visual navigation via deep learning and novelty detection, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 3–11.
- [174] S. Fuertes, E. Carmona, I. Arruego, J. Hervas, J. Rodriguez, Improving spacecraft health monitoring with automatic anomaly detection techniques, in: *14th International Conference on Space Operations*, 2016, p. 2430.
- [175] K. K. Reddy, B. R. Divakar, B. K. Panigrahi, A. Prasad, A. Agrawal, et al., Anomaly detection and fault disambiguation in large flight data: A multi-modal deep auto-encoder approach, in: *Annual Conference of the PHM Society*, volume 8, 2016.
- [176] S. Wold, K. Esbensen, P. Geladi, Principal component analysis, *Chemometrics and intelligent laboratory systems 2* (1987) 37–52.
- [177] Q. Meng, W. An, L. Deng, Relational autoencoder for feature extraction, in: *2017 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2017, pp. 364–371.
- [178] A. Y. Appiah, W. Lin, J. Lv, Q. Li, H. Li, Long short-term memory networks based automatic feature extraction for photovoltaic array fault diagnosis, *IEEE Access* 7 (2019) 30089.
- [179] A. A. Jiménez, G. Robles, J. López, A. Arnaiz, F. J. Molina, R. Gutiérrez, Dirt and mud detection and diagnosis on a wind turbine blade employing guided waves and supervised learning classifiers, *Reliability Engineering & System Safety* 184 (2019) 2.
- [180] M. Gomez-Fernandez, et al., Status of research and development of learning-based approaches in nuclear science and engineering: A review, *Nuclear Engineering and Design* 359 (2020) 110479.
- [181] International Atomic Energy Agency, Control room systems design for nuclear power plants. technical report, [http://www.iaea.org/inis/collection/NCLCollectionStore/\\_Public/27/002/27002051.pdf](http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/27/002/27002051.pdf), 1995.
- [182] R. Souza, J. Moreira, Neural network correlation for power peak factor estimation, *Annals of Nuclear Energy* 33 (2006) 594–601.
- [183] K. Moshkbakshayesh, M. Ghofrani, Transient identification in nuclear power plants: A review, *Progress in Nuclear Energy* 67 (2013) 23–38.
- [184] Z. Yang, F. Dong, J. Ma, Y. Zhang, Y. Wu, J. Gu, Z. Zhang, Application of convolution neural network to flow pattern identification of gas-liquid two-phase flow in small-size pipe, in: *2017 Chinese Automation Congress (CAC)*, IEEE, 2017, pp. 1389–1393. doi:10.1109/CAC.2017.8242984.
- [185] J. Park, S. Yang, M. Lee, J. Bae, W. Jung, Providing support to operators for monitoring safety functions using reinforcement learning, *Progress in Nuclear Energy* 118 (2020) 103123.

- [186] G. Lee, B. Kim, J. H. Park, P. H. Seong, A convolutional neural network model for abnormality diagnosis in a nuclear power plant, *Applied Soft Computing* 99 (2021) 106874.
- [187] J. Degraeve, G. Van Oost, D. De Baere, P. Maget, I. Novikau, A. Loarte, C. Reux, F. Ryter, D. Ryan, F. Taccogna, et al., Magnetic control of tokamak plasmas through deep reinforcement learning, *Nature* 602 (2022) 414–419.
- [188] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, D. Mané, Concrete problems in ai safety, *arXiv preprint arXiv:1606.06565* (2016).
- [189] Z. Li, D. Hoiem, Improving confidence estimates for unfamiliar examples, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2686–2695.
- [190] I. J. Goodfellow, J. Shlens, C. Szegedy, Explaining and harnessing adversarial examples, *arXiv preprint arXiv:1412.6572* (2014).
- [191] S. Kokalj-Filipovic, N. Bezmalinovic, I. Maric, I. Knezevic, Adversarial examples in rf deep learning: Detection and physical robustness, in: *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, IEEE, 2019, pp. 1–5.
- [192] S. Mohseni, A. Ahmadi, H. Mao, M. Al-Shedivat, Practical machine learning safety: A survey and primer, *arXiv preprint arXiv:2106.04823* (2021).
- [193] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al., Explaining how a deep neural network trained with end-to-end learning steers a car, *arXiv preprint arXiv:1704.07911* (2017).
- [194] C. Shorten, T. Khoshgoftaar, A survey on image data augmentation for deep learning, *Journal of big data* 6 (2019) 1–48.
- [195] I. Cristescu, F. Priester, D. Rapisarda, A. Santucci, M. Utili, Overview of the tritium technologies for the eu demo breeding blanket, *Fusion Science and Technology* 76 (2020) 446–457.
- [196] S. Brezinsek, J.-E. contributors, et al., Plasma-surface interaction in the be/w environment: conclusions drawn from the jet-ilw for iter, *Journal of nuclear materials* 463 (2015) 11–21.
- [197] R. Maurizio, S. Elmore, N. Fedorczak, A. Gallo, H. Reimerdes, B. Labit, C. Theiler, C. Tsui, W. Vijvers, T. Team, et al., Divertor power load studies for attached l-mode single-null plasmas in tcv, *Nuclear Fusion* 58 (2017) 016052.
- [198] T. Petrie, B. Grierson, T. Osborne, C. Petty, F. Turco, S. Allen, M. Fenstermacher, J. Ferron, H. Guo, E. Hinson, et al., High performance double-null plasmas under radiating divertor and mantle scenarios on diiii-d, *Nuclear Fusion* 59 (2019) 086053.

- [199] H. Reimerdes, G. Canal, B. Duval, B. Labit, T. Lunt, W. Vijvers, S. Coda, G. De Temmerman, T. Morgan, F. Nespoli, et al., Power distribution in the snowflake divertor in tcv, *Plasma Physics and Controlled Fusion* 55 (2013) 124027.
- [200] E. Martelli, F. Giannetti, G. Caruso, A. Tarallo, M. Polidori, L. Barucca, A. Del Nevo, Study of eu demo wcll breeding blanket and primary heat transfer system integration, *Fusion Engineering and Design* 136 (2018) 828–833.
- [201] D. Q. Kern, D. Q. Kern, *Process heat transfer*, volume 5, McGraw-Hill New York, 1950.
- [202] B. Zohuri, *Heat pipe design and technology*, Springer, 2011.
- [203] F. Papa, M. Utili, A. Venturini, G. Caruso, L. Savoldi, R. Bonifetto, D. Valerio, A. Allio, A. Collaku, M. Tarantino, Engineering design of a permeator against vacuum mock-up with niobium membrane, *Fusion Engineering and Design* 166 (2021) 112313.
- [204] B. Garcinuno, D. Rapisarda, I. Fernández, C. Moreno, I. Palermo, Á. Ibarra, Design of a permeator against vacuum for tritium extraction from eutectic lithium-lead in a dcll demo, *Fusion Engineering and Design* 117 (2017) 226–231.
- [205] P. W. Humrickhouse, B. J. Merrill, Vacuum permeator analysis for extraction of tritium from dcll blankets, *Fusion Science and Technology* 68 (2015) 295–302.
- [206] V. A. Maroni, R. D. Wolson, G. E. Staahl, Some preliminary considerations of a molten-salt extraction process to remove tritium from liquid lithium fusion reactor blankets, *Nuclear Technology* 25 (1975) 83–91.
- [207] F. Okino, K. Noborio, Y. Yamamoto, S. Konishi, Vacuum sieve tray for tritium extraction from liquid pb–17li, *Fusion Engineering and Design* 87 (2012) 1014–1018.
- [208] C. Day, A. Antipenkov, M. Dremel, H. Haas, V. Hauer, A. Mack, D. Murdoch, M. Wykes, R&d and design for the cryogenic and mechanical vacuum pumping systems of iter, *Vacuum* 81 (2007) 738–747.
- [209] T. Giegerich, C. Day, Conceptuation of a continuously working vacuum pump train for fusion power plants, *Fusion Engineering and Design* 88 (2013) 2206–2209.
- [210] T. Giegerich, C. Day, M. Jäger, Mercury ring pump proof-of-principle testing in the theseus facility, *Fusion Engineering and Design* 124 (2017) 809–813.
- [211] B. Bornschein, M. Glugla, K. Günther, R. Lässer, T. Le, K. Simon, S. Welte, Tritium tests with a technical permcath for final clean-up of iter exhaust gases, *Fusion Engineering and Design* 69 (2003) 51–56.

- 
- [212] K. Kotoh, M. Tanaka, S. Takashima, T. Tsuge, Y. Asakura, T. Uda, T. Sugiyama, Verification of hydrogen isotope separation/enrichment by pressure swing adsorption process: Successive enrichment of deuterium using sz-5a column, *Fusion Engineering and Design* 85 (2010) 1992–1998.
- [213] X. Xiao, H. T. Sessions, R. Rabun, Advanced isotope separation technology for fusion fuel, *Fusion Science and Technology* 78 (2022) 253–257.
- [214] A. Perevezentsev, B. Andreev, M. Rozenkevich, Y. S. Pak, A. Ovcharov, S. Marunich, Wet scrubber technology for tritium confinement at iter, *Fusion engineering and design* 85 (2010) 1206–1210.
- [215] W. Shmayda, P. Mayer, Uranium beds for temporary tritium storage, *Journal of the Less Common Metals* 104 (1984) 239–250.
- [216] W. Shmayda, A. Heics, N. Kherani, Comparison of uranium and zirconium cobalt for tritium storage, *Journal of the Less Common Metals* 162 (1990) 117–127.
- [217] S. Combs, Pellet injection technology, *Review of scientific instruments* 64 (1993) 1679–1698.

# Appendix A

## Summary of the FMEA of ARC PCS

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
Tank wall	External leakage	Material flaw, penetration welding failure, corrosion	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	1.7E-9	E	4	Design a corrosion resistant coating for the wall. Check welding and penetration during construction/installation	Assumed to be made of the same material of the VV wall.	INEEL/EXT-98-00892
	External rupture	Material failure	Medium or large loss of coolant	Depressurization, loss of heat removal function	Reactor shutdown	2.2E-10	E	1	Pre-service inspection. Non destructive testing.	Assumed to be made of the same material of the VV wall.	INEEL/EXT-98-00892
Vacuum vessel joint	Rupture	Thermomechanical stresses, mechanical overstress during assembly, fatigue	Missing connection between vacuum vessel and supporting structure	Stress increase on functioning joints	Depends on the number of functioning joints	2.0E-7	D	3	Assess fatigue threshold. Minimize fatigue cycles	Reference reports a 2.0E-6 upper bound failure rate (90% loosen, 10% rupture).	Bickford 1995
First wall (including the structural material)	Internal leakage	Material flaw, penetration welding failure, corrosion (from the structural material side)	Gas or FLiBe penetration inside the vacuum vessel. Disruption is likely to occur	Pressurization of the vacuum vessel, possible mobilization of radioactive dust	Reactor shutdown, possible damages to the first containment	1.7E-9	E	3	Design a corrosion resistant coating for the wall. Check welding and penetration during construction		INEEL/EXT-98-00892

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
	Internal rupture	Material failure, creep, fatigue, disruptions	Medium or large gas or FLiBe penetration inside the vacuum vessel	Pressurization of the vacuum vessel, loss of blanket heat removal function, mobilization of radioactive dust	Reactor shutdown, damage to the first containment, radiological release	2.2E-10	E	1	Assess fatigue threshold. Minimize fatigue cycles. Set operation conditions to avoid creep	A single disruption should not lead to first wall failure.	INEEL/EXT-98-00892
Temperature sensor	Fails to operate	Fouling, calibration error, no power	No information on FLiBe temperature	Possible temperature drift from nominal value	Need operator action	1.0E-6	C	4	System should be designed with at least two sensors. Include sensor test during scheduled maintenance		INEEL/EXT-98-00892
	Improper reading	Calibration error	Wrong temperature information	Wrong temperature reading to the operators	No consequences if other sensors available. Otherwise need shutdown	7.0E-7	C	3	System should be designed with at least two sensors. Voting logic and three sensors can filter the spurious signal. Select remote calibrating sensors to limit workers exposure. Include sensor test during scheduled maintenance.	Improper reading may lead to more severe consequences that fails to operate. Overheating or undercooling may result from wrong reading.	INEEL/EXT-98-00892
	External leakage	Instrumentation crack, welding failure	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	3.0E-5	B	3	Check welding sealing during scheduled maintenance. Isolate the leak if possible.	Coolant loss is probably negligible. Reactor can operate until scheduled shutdown	EGG-SSRE-8875
Heaters (pipe heat tracing)	Fail to start	Electric fault, control failure, no power	No FLiBe temperature increase	Possible loss of flow due to FLiBe solidification	Isolation of the failed heaters section	1.6E-6	C	3	Periodic testing, QA on control system	Higher bound chosen for failure rate. Failure mode is not specified (fail to function)	IAEA-TECDOC-478

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
	Improper starting	Control failure	FLiBe temperature increase	Degraded heat transfer in blanket, pressure increase	Isolation of the failed heaters section	1.6E-6	C	3	QA on control system. Identify operator procedures.	Higher bound chosen for failure rate. Failure mode is not specified (fail to function). FLiBe boiling improbable	IAEA-TECDOC-478
Pressure sensor	Fails to operate	Fouling, calibration error, no power	No information on FLiBe pressure	Possible pressure drift from nominal value	Need operator action	1.0E-6	C	4	System should be designed with at least two sensors. Include sensor test during scheduled maintenance.		INEEL/EXT-98-00892
	Improper reading	Calibration error	Wrong pressure information	Wrong pressure reading to the operators	No consequences if other sensors available. Otherwise need shutdown	8.7E-7	C	3	System should be designed with at least two sensors. Voting logic and three sensors can filter the spurious signal. Select remote calibrating sensors to limit workers exposure. Include sensor test during scheduled maintenance.		INEEL/EXT-98-00892
	External leakage	Instrumentation crack, welding failure	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	3.0E-5	B	3	Check welding sealing during scheduled maintenance. Isolate the leak if possible.	Frequency assumed the same as Temperature sensor	INEEL/EXT-98-00892
Pressurizer	External leakage	Material flaw, penetration welding failure	Small gas or coolant loss, depending on the leakage position	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	1.0E-8	E	3	Materials selection in design, pre-inspection.		INEEL/EXT-98-00892

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
	External rupture	Material failure	Medium or large loss of gas and/or coolant	Depressurization of the system	Isolation of the pressurizer. Reactor shutdown	1.0E-10	E	2	Materials selection in design, pre-inspection. Isolate the pressurizer if rupture occurs		INEEL/EXT-98-00892
Pressurizer heater	Fails to start	Electric fault, control failure, no power	No heat to the FLiBe in the pressurizer	Pressurizer cannot increase system pressure	Cannot reach the required pressure in the FLiBe circuit	2.2E-6	C	4	Periodic testing. QA on control system	Reported failure mode is "won't heat" and "others"	IAEA-TECDOC-478
	Improper starting	Control failure	Heat to the FLiBe in the pressurizer when not required	FLiBe temperature increases in the pressurizer	Pressure and temperature in the FLiBe circuit increase	2.2E-6	C	3	QA on control system. Identify operator procedures. Safety relief valve will open when the maximum pressure is reached.	Reported failure mode is "won't heat" and "others". FLiBe boiling improbable	IAEA-TECDOC-478
Rupture disk	Improper opening	Pressure pulse, improper installation, corrosion, construction fault	Pressurizer gas leakage	Depressurization of the system	Isolation of the pressurizer. Reactor shutdown	3.3E-6	C	3	Replace the disk at regular interval.		IAEA-TECDOC-478
Temperature safety relief valve	Improper opening	Wrong signal, mechanical malfunctioning	Coolant flows out	Depressurization of the system	Isolation of the pressurizer. Reactor shutdown	3.4E-6	C	4	Periodic inspection and testing		IAEA-TECDOC-478. VSPDV
	Internal leakage	Valve wear, corrosion	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	1.5E-7	D	4	Periodic inspection		IAEA-TECDOC-478. VR68F
	External leakage	Welding flaw, construction fault	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	1.5E-7	D	3	Pre-service inspection and periodic inspection	No data are found for external leakage. Same probability of Internal leakage is assumed	IAEA-TECDOC-478

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
	Internal rupture	Material failure	Medium or large loss of coolant from primary circuit to safety tank	Depressurization of the system	Isolation of the pressurizer. Reactor shutdown	1.0E-8	E	4	Inspection, non-destructive testing	No data are found for internal rupture. A frequency value is assumed according to similar valve failure	IAEA-TECDOC-478
	External rupture	Material failure, welding failure	Medium or large loss of coolant	Depressurization of the system	Isolation of the pressurizer. Reactor shutdown	1.0E-8	E	3	Pre-service inspection, non-destructive testing	No data are found for external rupture. A frequency value is assumed according to similar valve failure	IAEA-TECDOC-478
Salt melter	Improper starting	Wrong signal	Heat production in the storage tank	Solid FLiBe starts melting	No direct effect unless FLiBe reaches boiling temperature	1.6E-6	C	4	QA on system control. Design the system to accommodate the volume increase due to phase change.	The technology used to melt the salt greatly affects failure modes and failure rates. A common electrical heater system is assumed	IAEA-TECDOC-478
FLiBe storage tank	External leakage	Material flaw, welding failure, corrosion	Small FLiBe loss	FLiBe storage inventory decrease	Continuous loss of FLiBe from storage tank	1.0E-8	E	3	Check welding and penetration during construction/installation	Same failure rate as pressurizer leak	INEEL/EXT-98-00892
	External rupture	Material failure	Medium or large FLiBe loss	FLiBe storage unavailable	Cannot provide FLiBe to the coolant circuit if required	2.7E-8	E	2	Materials choice in design. Pre-service inspection. Non-destructive testing.		IAEA-TECDOC-478

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
Pipes (from blanket to tritium extractor)	Rupture	Material cracks, corrosion, overpressure	Large loss of FLiBe	Depressurization of the system	LOCA. FLiBe has maximum tritium concentration in this section	7.0E-11	E	2	Pre-service inspection. Non-destructive testing. Limit corrosion by neutralizing corrosive agents through a dedicated chemical unit.	Data includes contribution from all pressure bound components (nozzle, fittings, valve bodies). Pipes are nuclear grade. FLiBe has high tritium concentration	IAEA-TECDOC-478
	Plugging	Corrosion product buildup, impurity buildup, external materials buildup, foreign material intrusion	Flow reduced or completely stopped	FLiBe flow reduced, higher outlet temperature from blanket, structural materials temperature increase	LOFA	2.5E-9	E	1	Limit corrosion and impurity in FLiBe flow	Average failure rate for plugging is 2.5E-11-m. 100m of piping is assumed	INEEL/EXT-98-00892
Pipes (from tritium extractor to blanket)	Rupture	Material cracks, corrosion, overpressure	Large loss of FLiBe	Depressurization of the system	LOCA. FLiBe has minimum tritium concentration in this section	7.0E-11	E	2	Pre-service inspection. Non-destructive testing. Limit corrosion by neutralizing corrosive agents through a dedicated chemical unit.	Data includes contribution from all pressure bound components (nozzle, fittings, valve bodies). Pipes are nuclear grade. FLiBe has low tritium concentration	IAEA-TECDOC-478
	Plugging	Corrosion product buildup, impurity buildup, external materials buildup, foreign material intrusion	Flow reduced or completely stopped	FLiBe flow reduced, higher outlet temperature from blanket, structural materials temperature increase	LOFA	2.5E-9	E	1	Limit corrosion and impurity in FLiBe flow	Average failure rate for plugging is 2.5E-11-m. 100m of piping is assumed	INEEL/EXT-98-00892

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
Blanket inlet valve (MOV)	Fails to remain in position (open)	Wrong signal, motor fault	FLiBe flow blocked by the valve	No inlet flow to the blanket	Reactor shutdown. Need actions to the secondary cycle as well.	2.0E-7	D	3	Periodic testing. Re-open the valve or shutdown the reactor		IAEA-TECDOC-478
	External leakage	Welding flaw, construction fault	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	2.0E-6	C	4	Check welding sealing during scheduled maintenance.		IAEA-TECDOC-478. VMATU
	External rupture	Material failure, welding failure	Medium or large loss of coolant	Depressurization of the system	Reactor shutdown	1.0E-8	E	3	Periodic inspection and testing. QA on valve construction		IAEA-TECDOC-478. VMATW
	Plugging	Impurities buildup, corrosion products buildup, foreign material intrusion	FLiBe flow reduced or stopped	No flow to the blanket	LOFA	5.0E-9	E	2	Limit corrosion and impurity in FLiBe flow.		EGG-SSRE-8875
Blanket outlet valve (MOV)	Fail to remain in position (open)	Wrong signal, motor fault	FLiBe flow blocked by the valve	No outlet flow from the blanket	Reactor shutdown. Need actions to the secondary cycle as well.	2.0E-7	D	4	Periodic testing. Re-open the valve or shutdown the reactor		IAEA-TECDOC-478
	External leakage	Welding flaw, construction fault	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	2.0E-6	C	3	Check welding sealing during scheduled maintenance.		IAEA-TECDOC-478
	External rupture	Material failure, welding failure	Medium or large loss of coolant	Depressurization of the system	Reactor shutdown	1.0E-8	E	3	Periodic inspection and testing. QA on valve construction		IAEA-TECDOC-478
	Plugging	Impurities buildup, corrosion products buildup, foreign material intrusion	FLiBe flow reduced or stopped	No flow from the blanket	LOFA	5.0E-9	E	2	Limit corrosion and impurity in FLiBe flow.		EGG-SSRE-8875
Flow sensor	Spurious function	Calibration error, no power	Wrong flow information	Wrong information to the operator	No significant plant effect	4.3E-6	D	4	Other sensors should provide enough information to keep the reactor running		IAEA-TECDOC-478

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
Pump (motor driven)	Fail to run, given start	Motor fault, motor overheated, shaft damage	FLiBe flow stopped	No flow in blanket	LOFA	3.0E-5	B	2	System should be designed with two pumps. Each pump should have twice the required pumping power.	Include pump and motor. Data have been chosen according to the ultimate source reported	IAEA-TECDOC-478. PMYRW
Heat exchanger	Degraded heat transfer	Fouling	FLiBe temperature increase in the HX	Higher coolant outlet temperature from HX	Higher inlet coolant temperature in blanket. Lower secondary coolant temperature	1.0E-6	C	4	Periodic inspection and maintenance. If possible, choose a heat exchanger that can be easily inspected and clean		EGG-SSRE-8875
	Tube leakage	Welding tube flaw	Small loss of primary fluid	Possible contamination between primary and secondary fluid	Negligible loss of primary fluid. Possible contamination between primary and secondary fluid	3.0E-9	E	3	Periodic inspection and testing	Failure rate is given per tube	IAEA-TECDOC-478
	Tube rupture	Welding failure, tube failure, overheating, overpressure	Loss of primary flow	Loss of primary flow. Mixing between primary and secondary fluid and contamination	LOCA. Contamination of the secondary fluid	1.0E-6	C	2	Periodic inspection and testing. Isolate the heat exchanger to avoid contamination of the secondary fluid	Failure rate per tube bundle. Consequences depend on secondary fluid pressure	EGG-SSRE-8875
	Shell leakage	Material welding flaw	Small loss of secondary fluid	Possible depressurization of the secondary circuit	Continuous loss of secondary fluid	3.0E-6	C	3	Periodic inspection and testing		IAEA-TECDOC-478
	Shell rupture	Material welding failure	Loss of secondary flow	Depressurization of the secondary circuit. No heat extraction from coolant	Loss of heat sink. Coolant temperature rises	1.0E-7	D	1	Periodic inspection and testing. Isolate the heat exchanger and shut down the reactor	Consequences depend on the secondary fluid choice	EGG-SSRE-8875

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
	Plugging	Corrosion product buildup, impurity buildup, external materials buildup, foreign material intrusion	Flow stopped in the tube bundle	Higher pressure drop in the circuit. Plugged heat exchanger unavailable	Small LOFA	3.0E-7	C	2	Limit corrosion and impurity in FLiBe flow. Clean heat exchanger tubes during maintenance. Isolate failed heat exchanger and redirect FLiBe flow to the other heat exchanger	Severity lower than pipe plugging due to the possibility to redirect the flow to other heat exchangers	EGG-SSRE-8875
HX inlet valve (MOV) - Primary circuit	Fail to remain in position (open)	Wrong signal, motor fault	FLiBe flow blocked by the valve	No inlet flow to the blanket	Reactor shutdown.	2.0E-7	D	4	Periodic testing. Re-open the valve or shutdown the reactor		IAEA-TECDOC-478
	External leakage	Welding flaw, construction fault	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	2.0E-6	C	4	Check welding sealing during scheduled maintenance.		IAEA-TECDOC-478. VMATU
	External rupture	Material failure, welding failure	Medium or large loss of coolant	Depressurization of the system	Reactor shutdown	1.0E-8	E	3	Periodic inspection and testing. QA on valve construction		IAEA-TECDOC-478. VMATW
	Plugging	Impurities buildup, corrosion products buildup	FLiBe flow reduced or stopped	No flow to the heat exchanger	Small LOFA	5.0E-9	E	3	Limit corrosion and impurity in FLiBe flow. Redirect FLiBe flow to the other heat exchangers		EGG-SSRE-8875
HX outlet valve (MOV) - Primary circuit	Fail to remain in position (open)	Wrong signal, motor fault	FLiBe flow blocked by the valve	No inlet flow to the blanket	Reactor shutdown	2.0E-7	D	4	Periodic testing. Re-open the valve or shutdown the reactor		IAEA-TECDOC-478
	External leakage	Welding flaw, construction fault	Small coolant loss	Possible depressurization depending on leakage magnitude	Continuous loss of coolant if reactor is not shut down	2.0E-6	C	4	Check welding sealing during scheduled maintenance.		IAEA-TECDOC-478. VMATU
	External rupture	Material failure, welding failure	Medium or large loss of coolant	Depressurization of the system	Reactor shutdown	1.0E-8	E	3	Periodic inspection and testing. QA on valve construction		IAEA-TECDOC-478. VMATW

Component	Failure mode	Possible causes	Local effects	System effects	Plant effects	Frequency [h <sup>-1</sup> ]	Prob. Cat.	Severity	Preventive and corrective measures	Comments	Source
	Plugging	Impurities buildup, corrosion products buildup, foreign material intrusion	FLiBe flow reduced or stopped	No flow from the heat exchanger to the blanket	Small LOFA	5.0E-9	E	3	Limit corrosion and impurity in FLiBe flow. Redirect FLiBe flow to the other heat exchangers		EGG-SSRE-8875
Freeze valve	Spurious Thaw	Mechanical failure, control failure	FLiBe starts flowing through the freeze valve	Depressurization of the system	Small LOCA	2.54E-05	B	3	Periodic inspection and testing		Chisholm, 2020

# **Appendix B**

## **Machine learning for safety applications**

This Appendix frames the analysis shown in Chapter 3 in the broader context of ML for safety applications. Readers who are not familiar with ML may benefit from this Appendix to build a basic knowledge on ML models for safety applications. A more detailed discussion on the main criticalities of ML algorithm in safety critical industries is presented.

### **Overview of ML models in safety-critical industries**

Machine learning is not an exact science. There is no model suitable for any kind of application: the most effective ML model for a given application is usually identified among potential models after training and test phases. Such an approach may be extremely time consuming. The nuclear sector may benefit from the studies and applications carried out in different fields, leveraging on past experiences and best practices to modernize the nuclear industry [170]. This Appendix presents real-world applications of ML in safety critical industries. Basics tasks such as classification, regression, clustering and anomaly detections are placed in context. The fields considered are autonomous driving, aeronautics and spacecraft, rotating machinery, and energy production. For additional information, a comprehensive overview of ML applications in industrial systems can be found in [87].

---

## Non-nuclear fields

Among safety critical industries, autonomous driving is probably the one that mostly benefited from ML. Autonomous driving is based on computer vision, but it should not be reduced just to it. Autonomous vehicles must first “understand” the surrounding environment, which is done through computer vision, and then select the action to undertake. A complicating factor is the presence of dynamic objects in the environment whose behaviour may be hard to predict, such as other drivers or pedestrians. A very detailed review on the topic can be found in [81]. The most widely ML models exploited in autonomous driving are Convolutional Neural Networks (CNN) [88]. CNN are multi-layered neural networks that take advantage of the convolution operation and pooling layers to extract high-level features (e.g., shapes) from low-level features (e.g, pixels in an image). This is at the basis of object recognition, whose goal is to correctly identify different regions in an image and output a label (the object name) for them. The output can then be passed to another DNN or an agent to undertake an action. Xu et al. [171] coupled a CNN with a Long Short-Term Memory (LSTM) [124] network in an end-to-end system capable to undertake discrete actions (e.g., stop, start, turn left) and to predict continuous behaviours (e.g., steering wheel angle control). Similarly, NVIDIA PilotNet predicts the steering angle based on raw images from a front-facing camera exploiting a CNN only [172]. The input images for PilotNet were also rotated and distorted to increase the robustness of the model. More importantly, the researchers managed to visualize salient objects from the layers of the CNN [56]. The CNN learned salient features such as road borders or the shape of other vehicles, providing an explanation of how the trained network worked. An important step forward required for these safety relevant systems is the capability of the model to understand when a new situation is encountered. This issue is treated in detail in Section 3 when discussing the generalization capabilities of ML models. As for the autonomous driving is concerned, Richter [173] attempted to develop a method to detect anomalies (in this case, novelties) and revert the operating approach of the system from a learning-based to a (safer) non-learning-based. Differently from computer vision, the input to the ML models comes as time series (i.e., physical or logical variables as a function of time) in the other fields described in that paper. Aeronautics and spacecrafts face a different kind of challenges. Inputs are usually provided by multiple, multimodal sensors as a stream of raw, noisy measurements.

The signals are analysed to predict some target variables of the system (regression) or to detect anomalies. A feature extraction process is usually required to select only those input signals that are relevant for the task. Features extraction may be handcrafted [174] or embedded in the model [175]. Fuertes et al. [174] developed a software for automatic anomaly detection in spacecraft by starting from manual selection of the features, reducing the dimensionality of the input with PCA [176] and exploiting a one-class support vector machine to detect the anomalies. Reddy et al. [175] proposed a deep auto-encoder to automate the feature extraction process and detect faults exploiting data from the NASA DASHlink open database. Both the approaches reported good accuracy. However, ML models that achieve autonomous feature extraction may unveil latent features that are instead overlooked by the expert. There is indeed a rising trend to let the model discovering the relevant features from the input data by exploiting DNN architectures [177]. A more transversal field is rotating machinery and motors. The most critical task in this field is again fault diagnosis (i.e., anomaly detection). Ince et al. [130] proposed a 1-D CNN to perform real-time motor fault detection. The use of a CNN is justified by the need to extract features prior to the fault diagnostics. Even though Recurrent Neural Networks (RNN) (e.g., LSTM) may look more suitable for analysis of time series (e.g., the raw sensor measurements), a 1-D CNN is light enough to perform real-time monitoring and to correctly classify the system state. A similar approach is followed by Lei et al. [129], who developed a DNN where the first two layers are devoted to features extraction and the deeper layers compute the health condition of the system. Additionally, Lei et al. [129] shew that the weights of the layers performing feature extractions can be physically interpreted as a Gabon-like filter. This is again an extremely important contribution to the understanding and visualization of DNN. Another critical issue is addressed by Guo et al [128], who developed a deep convolutional transfer learning network. As a matter of fact, real-world data may not belong to the same probability distribution of the training set. This usually translates in bad predicting capabilities of the model. Guo et al. [128] exploited a 1-D CNN to assess system conditions, coupled with a domain adaptation module that allows to learn domain invariant features. In such a way the robustness of the model can be greatly increased. For further applications in this field the reader is referred to [127]. Even in the energy production field, fault diagnosis is of paramount importance. Appiah et al. [178] developed a LSTM that processes input from simulated photovoltaic (PV) arrays to predict PV failure and

---

possible fire hazard. The input does not come from real-world measurement but from computer simulations, and the data are not processed in real time but offline. Nevertheless, the model seems suitable to predict failures in PV arrays, and highlights the importance of a tight coupling between system simulators and ML. Five different supervised models are instead investigated by Jiménez et al. [179] for mud level detection in wind turbines. The data is acquired by non-destructive techniques on an experimental setup. Feature extraction is carried out separately from the classifier by autoregressive models and PCA. The reduced input is then fed into the five ML models to train and test the accuracy. A systematic review of machine learning models for energy system can be found in [68].

## **Nuclear science and engineering**

Nuclear industry is implementing machine learning techniques more gradually than other fields. Strict regulation, high safety standards and intellectual properties may slow the penetration of these methods for nuclear applications [180]. Nevertheless, the International Atomic Energy Agency (IAEA) addressed this topic in the TECDOC-812 [181], considering the introduction of AI tools to support decision. Many applications of ML are related to fault detection and diagnosis, which is a transversal topic touching different fields, as shown in Section B. A detailed reference on this topic can be found in [86]. PCA, Multivariate State Estimation Technique (MSET) and ANN (Artificial Neural Networks) are the most exploited methods for Nuclear Power Plants (NPP). PCA has been used to monitor instruments and equipment [86]. Its capability to reduce the dimension of the feature space is exploited to select the most relevant signals coming from the sensors. ANN are particularly suitable for this task, thanks to their capability to process large inputs in short time. For instance, Souza and Moreira [182] demonstrated that the reactor power peak factor can be computed with a low relative error (approximately 1%) by an ANN. ANN have been used also to identify transients [183]. More recent applications of ML methods include DNN for flow regime identification [184], Modified Self-Organizing Maps (MSOM) and SVM for transient identification [183], RL for operation and decision support [185], and CNN for abnormality diagnosis [186]. The reader should be aware that differently from the other safety critical industries presented previously, most of the ML applications in the nuclear field have been applied to real-world data or simulators data, but the implementation

on real systems is usually missing. A broad and recent overview of learning-based methods for nuclear science and engineering can be found in [180]. Lastly, a notable effort to integrate a ML model on the hardware of a real reactor is found in the nuclear fusion field. A reinforcement learning approach has been successfully exploited for plasma control by acting on the control coils of the TCV reactor [187].

## Criticalities and challenges in ML for safety applications

Machine learning applications have undoubtedly a different level of maturity in different fields. Some of the applications are closer to real-world applications (e.g., autonomous driving), while others are more restricted to academic studies, such as for nuclear. Nevertheless, ML algorithms have a tremendous potential for both real-world applications and academic research in the nuclear safety field. A deep understanding of the issues will foster ML diffusion by carrying out more robust analysis. Hence, this section provides a discussion on high-level and low-level criticalities for ML applications in safety relevant fields. High-level criticalities refer to those issues that are case-independent: ML safety, model interpretability, robustness, and verification, to cite few. Low-level criticalities are more technical related, such as availability of small dataset only, curse of dimensionality, model complexity and loss of salient information, effective feature extraction, and so on. Stated in another way, high-level criticalities affect the applicability of ML models in a safety framework, whereas low-level criticalities affect the efficiency and effectiveness of a given model to fulfil its task.

### High-level criticalities

An open field of research is ML safety, which can be considered a subset of AI safety [188]. Machine learning safety deals with unexpected and unintended behaviour of ML algorithms that may harm people or damage the surrounding environment. Machine learning models have a black-box nature, especially when considering DNN and RL agents (which usually embeds a DNN). The output of a layer is computed as  $f(\vec{w} \cdot \vec{x} + \vec{b} \cdot \vec{z})$ , where  $\vec{w}$  and  $\vec{b}$  are the learnable model's parameters

(weights and biases),  $\vec{z}$  is the output from the previous layer and  $f(\cdot)$  is a non-linear function. The information about a trained model is therefore embedded in the values of  $w$  and  $b$ , which have no physical meaning. While shallow networks may be somehow interpretable due to the limited number of parameters, DNN cannot be easily interpreted. Model interpretability is indeed an issue when ML is applied to safety-relevant cases. The requirement for high performances coupled with the high dimensionality of the domain inevitably leads to the choice of deep, complex neural networks that lack of transparency. Since RL is taking the lead in executing series of task in complex environment, we will use the term agent instead of ML model or ML algorithm in the discussion hereunder. The considerations made throughout this paragraph applies also to simpler ML models. The actions undertaken by the agent in a given environment are the main source of risk for ML safety. The effectiveness (and safety) of the agent policy largely depends on its reward function and the possibility to explore the environment. Generally speaking, the environment is in a good state (e.g., a reactor during nominal operations), and any agent actions may worsen this state. This is likely to occur when the reward function is not carefully defined. As an example, imagine an agent that must keep the reactor safe, where the reactor safety is defined by some metric involving reactor criticality. The easiest way to fulfil this task is to have a criticality equal to zero. The agent may try to shut down the reactor to maximize its reward function, with a detrimental effect on the reactor availability (and a corresponding economic loss). Such a behaviour is a typical reward hacking: the agent does not follow the expected behaviour because it found a shortcut to maximize its reward. A related issue is the possibility of producing negative side effects. If the reward function accounts only for a small part of the environment, the agent may take actions that have a negative effect on the part of the environment that is neglected by the reward function. If the unconsidered environment also includes operators, this behaviour may be extremely risky. From a practical viewpoint, defining an objective function with a high level of abstraction may solve these issues. A poor agent training may lead to safety issues as well. If the exploration of the environment is strongly constrained, the agent may face new situations for which it is not properly trained. Conversely, if the agent is free to explore the whole environment, it may undertake dangerous actions, which is clearly not acceptable in a safety critical environment. A safe exploration may be achieved if the risk profile in a given environment is known. This goal highlights the importance of deterministic models for high-fidelity simulations of the environment,

which is treated in Section B. The capability of the model to deal with new situations is instead discussed below. A major assumption in ML is that the training data belongs to the same probability distribution of operational data (i.e., the real-world data that the model will process when deployed). However, this assumption may not be satisfied in real-world scenarios. For instance, the training dataset may contain data representing the nominal flow regime of the reactor coolant. In accidental conditions the flow regime may drastically change, transitioning from turbulent flow to natural circulation for a Loss Of Flow Accident (LOFA). A model trained only with nominal condition data would fail to provide the desired output during an accidental scenario. Out-of-distribution data may appear even if care is taken in considering the accidental scenario domain. In the nuclear industry, accidents are rare events. Thus, few data is available. Furthermore, this data may not be representative of the whole spectrum of possible accidental scenarios that may occur. Using a ML model for a safety-relevant task during an accidental scenario (e.g., predicting the plant end-state) may lead to incorrect predictions if the accidental scenario differs from the ones used for model training. Such a behaviour may be experienced even if the training data and the operational data look similar, at least superficially, as demonstrated [120]. The generalization capability of the model is therefore a leading issue for safety applications. An additional complication is the lack of self-awareness of the ML model. In fact, DNN are prone to be overconfident on their predictions, especially on unseen data [189]. Generalization capabilities and self-awareness should be therefore tackled together when developing a ML model. An issue related to generalization and distributional error is the adversarial error. Despite the out-of-distribution samples, which occur “naturally” in a given environment, adversarial samples are crafted with the explicit goal of misleading the model without fooling the human [190]. Adversarial examples have been deeply investigated for computer vision (i.e., crafting adversarial images), while little interest has been shown by other fields that are not directly related to image processing. It is legit to ask whether this is a topic that may affect those other fields. Adversarial examples may be crafted following the same approach of computer vision, but targeting different domains, such as signal monitoring. For instance, Kokalj-Filipovic et al. [191] reported possible adversarial attacks to radio frequency signals. Few works address this problem outside computer vision, and more detailed discussion for the nuclear field are premature. Nevertheless, it should be noted that adversarial errors and adversarial attacks rise concerns not only for safety, but also

---

for security. Even if the model is developed following best-practices and design guidelines, the initial question remains valid: how can we be sure that the ML model behaves as expected during real-world operations? To address this problem in a rigorous way a verification procedure is needed. Verification of a ML model is a two-level issue. From the formal viewpoint, verification of ML algorithms cannot be achieved following the guidelines for software safety [192]. A formal verification requires exact specification of the algorithm properties to ensure that it exhibits the expected properties when operating in a given environment. It is evident that a formal verification of a black-box model is a hard task. Furthermore, the given environment should be correctly specified, but machine learning algorithms rely only on limited data samples from the environment. The environment might be not properly described, as described previously for the generalization issue. Therefore, from a practical point of view the verification of the model is affected also by the available data, which may be missing, incorrect or noisy.

### **Low-level criticalities**

The high-level criticalities described in the previous section usually find their counterpart in technical, low-level criticalities. Interpretability, for instance, may be lost when the feature extraction procedure is automated, or when the initial feature space is reduced by means of dimensionality reduction techniques (e.g., PCA). Nevertheless, manual feature extraction is time consuming, it requires expert knowledge and may lead to loss of generality. This last event may occur whenever the expert judges some features not relevant for the problem. However, many applications in Section 2 reported increased model capabilities when the feature extraction was performed by the algorithm itself, because the downsampling of the input may involve combination of features that a human may ignore. Thus, automatic feature extraction should be set as a target, increasing the effort in explaining the results and the model, as done in [193]. Real-world data from safety relevant fields comes mostly from nominal operations. Data from accidental scenarios are much more limited. Hence, the available training sets are usually small and unbalanced. Small datasets may be a leading issue for the generalization capability of the model. Unbalanced datasets (e.g., datasets when most of the samples come from nominal operations) have a more subtle effect on the model. Models trained with an unbalanced training set may report high accuracy, but their

accuracy drops if tested against relevant input data (e.g., only accidental scenarios). Data augmentation is a common solution to increase number of training samples from the relevant domain [194]. However, data augmentation is mostly exploited in computer vision, where rotation, distortion and similar operations do not alter the nature of an image (i.e., a cow remains a cow even if the image is rotated by  $90^\circ$ ). Stated in another way, the operations applied for data augmentation do not modify the invariants of the sample. If the model is properly trained, it can still identify the invariants of the input and fulfil its task. This approach becomes much harder when dealing with signals from an industrial plant, that is, when dealing with temporal series instead of images. Modifying the input signal during a transient may alter the end-state of the system. This happens because the signal is related to a physical variable of a system in evolution – if we change the physics of the process, the process may evolve in a different way. The only way to check if the end-state (which may be the target variable) remains the same is to simulate that specific physical condition. Therefore, deterministic codes become of paramount importance to increase the available data for training the model. The reader should be aware that the previous discussion refers to data augmentation. Many applications in literature report increases in the robustness of the model by adding noise and perturbations to the input signal. This is a valid and effective approach to improve the robustness of the model against noisy signals but does not tackle the issues related to small or unbalanced datasets. Lastly, results from ML models inevitably include false positive and false negative. While this is a minor issue for some ML applications (e.g., wrong advertisement on a user homepage), it becomes extremely relevant for safety application. Labelling an incipient accidental transient as a safe transient may have harsh consequences if an operator, or a higher-level software, does not undertake the required actions.

# **Appendix C**

## **ARC fuel cycle model**

A detailed layout of the IFC is shown in Figure C.1. The FC components and the tritium flows, with the corresponding nomenclature, are shown in Figure C.2. A brief description of the components of the OFC and IFC (and examples of technologies that can be used) is provided in Table C.1 and Table C.2.

All the quantities used in the residence-time-model are defined in Tab. C.3.

Component	Description	Examples
Blanket (LIB)	The blanket employs a lithium-containing material (FLiBe in the LIB), which interacts with neutrons from the D-T reactions in the plasma to produce tritium.	LIB (ARC) [5], Helium-cooled lithium lead (DEMO) [195], Water-cooled lithium lead (DEMO) [195], Helium-cooled pebble bed (DEMO), Dual-coolant lithium lead (DEMO) [195], liquid lithium (STEP).
First wall and vacuum vessel (FW/VV)	The FW surrounds the plasma. It must withstand high heat loads ( $\sim$ MW/m <sup>2</sup> ) and damaging PMI. It is usually made of high-Z materials. In ARC, the FW is immediately adjacent to the VV, which is inside the blanket.	Tungsten FW (ARC, expected), Be FW (JET-ILW [196]) VV: TBD, but low-activation material likely necessary.
Divertor	The divertor provides ash removal and heat extraction. The divertor design depends on the plasma magnetic configuration. It must withstand very high heat loads ( $\sim$ 10 MW/m <sup>2</sup> ) and strong particle bombardment.	Tungsten divertor (ARC, expected [3]), ITER tungsten cassette assemblies. Single null [197], double null [198], snowflake [199].
Heat exchanger (HX)	The HX transfers thermal energy from the primary coolant to the secondary coolant.	Shell and tube heat exchanger [200], plate heat exchanger, concentric pipe heat exchanger [201], heat pipes [202].
Tritium extraction system (TES)	The TES removes tritium from the tritium carrier (which may be solid, liquid, or gas depending on design). The extracted tritium is sent to the inner fuel cycle.	Permeator against vacuum [203–205], liquid-liquid contractor [206], gas-bubble columns [158], vacuum sieve tray [207].

Table C.1 Key components in the outer fuel cycle of a DT-burning fusion power plant, along with a non-comprehensive list of examples.

<b>Component</b>	<b>Description</b>	<b>Examples</b>
Vacuum pump	Vacuum pumps provide exhaust pumping from the vacuum chamber. The exhaust contains unburnt fuel, He ashes, and plasma enhancement gases.	Cryopumps [208], vapor diffusion pumps [209], liquid ring pumps [210].
Fuel clean-up	Fuel clean-up separates hydrogenic species from the other exhaust gases (plasma enhancement gases, helium ash).	Pd-Ag alloy permeator [211].
Isotope Separation System (ISS)	ISS separates hydrogen, deuterium, and tritium.	Pressure swing absorption [212], thermal diffusion [213], cryogenic distillation.
Detritiation system	The detritiation system removes tritium in trace amount from water (water detritiation system) and exhaust (exhaust detritiation system).	Wet scrubber [214], getter beds [215],[216].
Storage and management system	The storage and management system is used to store the initial tritium inventory and the tritium that eventually will be transferred to new FPP.	Uranium beds [215], ZrCo beds [216].
Fueling system	The fueling system delivers the required fuel mixture to the plasma.	Gas gun pellet injection [217], centrifuge pellet injection [217].
DIR loop	The DIR loop separates hydrogenic species from the gas exhaust by exploiting hydrogen superpermeation through metal foils (for MFP) or hydrogen cryosorption in different pump chambers (for cryopumps)	Possible DIR technologies include MFP or a multi-stage cryopump, linear diffusion pumping, and liquid ring pumping [147]

Table C.2 Key IFC components and a non-exhaustive list of associated technologies.

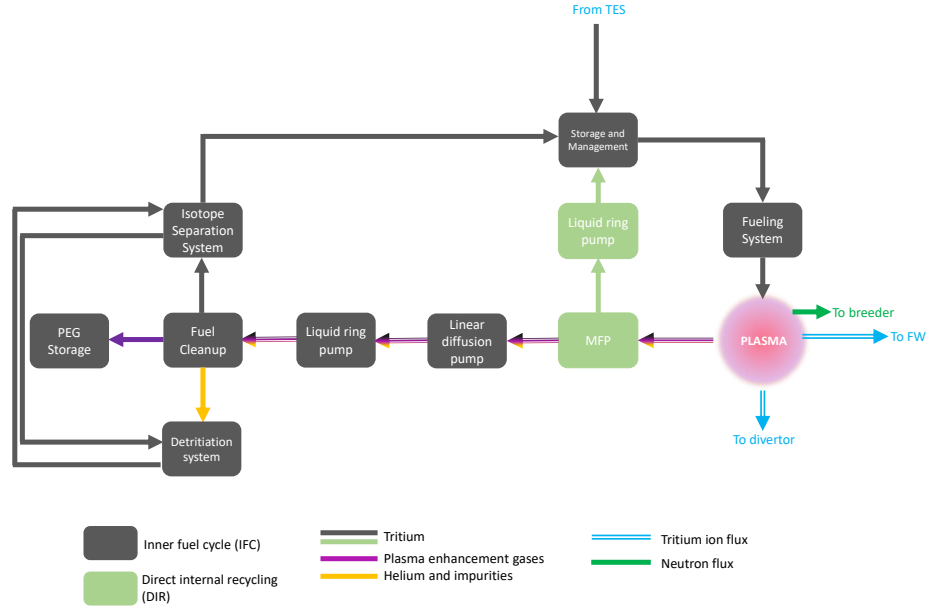


Fig. C.1 Detailed layout of ARC IFC considering the DIR loop. Tritium, PEG, helium and impurity flows are shown in different colors.

The general form for tritium balance in the  $i$ -th component is adapted from [30]:

$$\frac{dI_i}{dt} = \sum_{j \neq i} \left( f_{j \rightarrow i} \frac{I_j}{\tau_j} \right) - (1 + \varepsilon_i) \left( \frac{I_i}{\tau_i} \right) - \lambda I_i + S_i \quad (\text{C.1})$$

with  $\sum_j f_{i \rightarrow j} = 1$ . The system of ODE that describes the model is:

$$\frac{dI_1}{dt} = \text{TBR} \dot{N}_{\text{T,burn}} + \frac{I_3}{\tau_3} + \frac{I_4}{\tau_4} + f_{5 \rightarrow 1} \frac{I_5}{\tau_5} - \frac{I_1}{T_1} \quad (\text{C.2})$$

$$\frac{dI_2}{dt} = \frac{I_1}{\tau_1} - \frac{I_2}{T_2} \quad (\text{C.3})$$

$$\frac{dI_3}{dt} = f_{p-3} \frac{\dot{N}_{\text{T,burn}}}{\text{TBE}} + f_{5 \rightarrow 3} \frac{I_5}{\tau_5} - \frac{I_3}{T_3} \quad (\text{C.4})$$

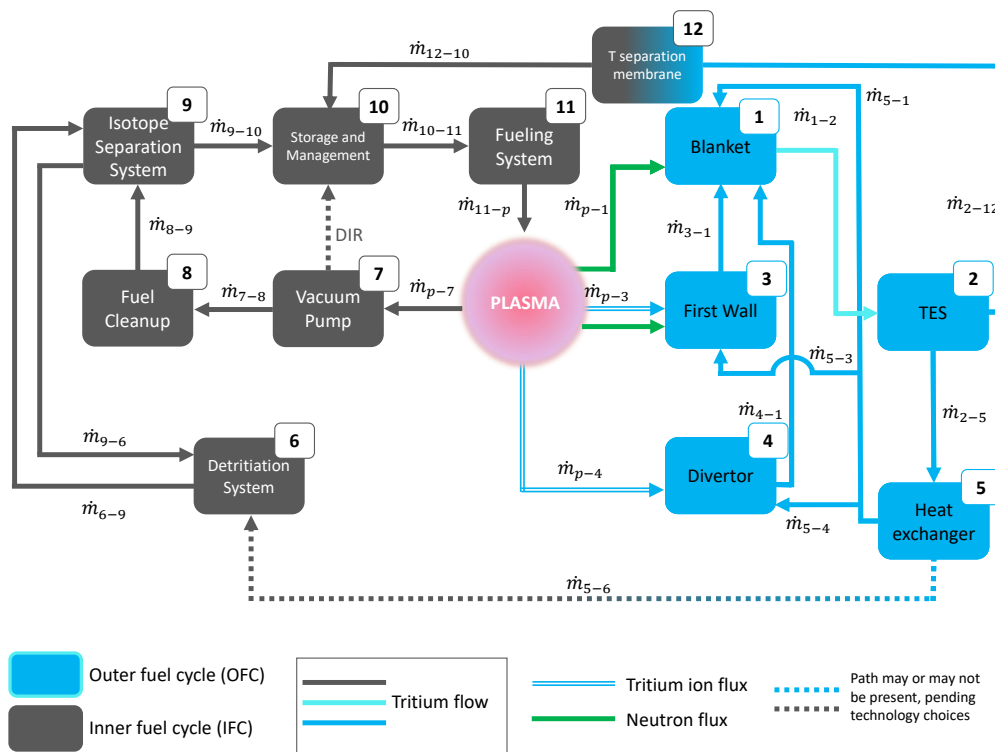


Fig. C.2 ARC fuel cycle layout. Only the components that have been modeled are shown in the figure. The tritium flows between components are depicted. For the definition of each tritium flow see Tab. C.3.

$$\frac{dI_4}{dt} = f_{p-4} \frac{\dot{N}_{T,\text{burn}}}{\text{TBE}} + f_{5-4} \frac{I_5}{\tau_5} - \frac{I_4}{T_4} \quad (\text{C.5})$$

$$\frac{dI_5}{dt} = (1 - \eta_2) \frac{I_2}{\tau_2} - \frac{I_5}{T_5} \quad (\text{C.6})$$

$$\frac{dI_6}{dt} = f_{5-6} \frac{I_5}{\tau_5} + f_{9-6} \frac{I_9}{\tau_9} - \frac{I_6}{T_6} \quad (\text{C.7})$$

$$\frac{dI_7}{dt} = (1 - \text{TBE} - f_{p-3} - f_{p-4}) \frac{\dot{N}_{T,\text{burn}}}{\text{TBE}} - \frac{I_7}{T_7} \quad (\text{C.8})$$

$$\frac{dI_8}{dt} = (1 - f_{\text{DIR}}) \frac{I_7}{\tau_7} - \frac{I_8}{T_8} \quad (\text{C.9})$$

$$\frac{dI_9}{dt} = \frac{I_6}{\tau_6} + \frac{I_8}{\tau_8} - \frac{I_9}{T_9} \quad (\text{C.10})$$

$$\frac{dI_{10}}{dt} = (1 - f_{9-6}) \frac{I_9}{\tau_9} + f_{\text{DIR}} \frac{I_7}{\tau_7} + \frac{I_{12}}{\tau_{12}} - \frac{\dot{N}_{T,\text{burn}}}{\text{TBE}} - \lambda I_{10} \quad (\text{C.11})$$

$$\frac{dI_{12}}{dt} = (1 - \eta_2) \frac{I_2}{\tau_2} - \frac{I_{12}}{T_{12}} \quad (\text{C.12})$$

where:

$$\frac{1}{T_i} = \frac{1 + \varepsilon_i}{\tau_i} + \lambda \quad (\text{C.13})$$

$$\varepsilon_i = 10^{-4}, \forall i \neq 3, 4, 10 \quad (\text{C.14})$$

$$\varepsilon_3 = \varepsilon_4 = \varepsilon_{10} = 0 \quad (\text{C.15})$$

and the initial conditions are given by:

$$I_i(t = 0) = 0, \forall i \neq 10 \quad (\text{C.16})$$

$$I_{10}(t = 0) = I_{\text{startup}} \quad (\text{C.17})$$

The non-radioactive loss fraction has been assumed as  $\varepsilon_i = 10^{-4}$  in this analysis. The values of the  $f_{i \rightarrow j}$  terms are reported in Table C.4.

Component	Flow rate, source or sink
1	$S_1 = \text{TBR} \dot{N}_{\text{T, burn}}$ $\dot{m}_{5-1} = f_{5-1} \left( \frac{I_5}{\tau_5} \right)$ $\dot{m}_{3-1} = \frac{I_3}{\tau_3}$ $\dot{m}_{4-1} = \frac{I_4}{\tau_4}$ $\dot{m}_{1-2} = -\frac{I_1}{\tau_1}$
2	$\dot{m}_{1-2} = \frac{I_1}{\tau_1}$ $\dot{m}_{2-12} = -\eta_2 \frac{I_2}{\tau_2}$ $\dot{m}_{2-5} = -(1 - \eta_2) \left( \frac{I_2}{\tau_2} \right)$
3	$\dot{m}_{p-3} = f_{p-3} \frac{\dot{N}_{\text{T, burn}}}{\text{TBE}}$ $\dot{m}_{5-3} = f_{5-3} \left( \frac{I_5}{\tau_5} \right)$ $\dot{m}_{3-1} = -\frac{I_3}{\tau_3}$
4	$\dot{m}_{p-4} = f_{p-4} \frac{\dot{N}_{\text{T, burn}}}{\text{TBE}}$ $\dot{m}_{4-1} = -\frac{I_4}{\tau_4}$ $\dot{m}_{5-4} = f_{5-4} \frac{I_5}{\tau_5}$
5	$\dot{m}_{2-5} = (1 - \eta_2) \left( \frac{I_2}{\tau_2} \right)$ $\dot{m}_{5-1} = -f_{5-1} \left( \frac{I_5}{\tau_5} \right)$ $\dot{m}_{5-3} = -f_{5-3} \left( \frac{I_5}{\tau_5} \right)$ $\dot{m}_{5-4} = -f_{5-4} \frac{I_5}{\tau_5}$ $\dot{m}_{5-6} = -f_{5-6} \left( \frac{I_5}{\tau_5} \right)$
6	$\dot{m}_{5-6} = f_{5-6} \left( \frac{I_5}{\tau_5} \right)$ $\dot{m}_{9-6} = f_{9-6} \left( \frac{I_9}{\tau_9} \right)$ $\dot{m}_{6-9} = -\frac{I_6}{\tau_6}$
7	$\dot{m}_{p-7} = (1 - \text{TBE} - f_{p-3} - f_{p-4}) \frac{\dot{N}_{\text{T, burn}}}{\text{TBE}}$ $\dot{m}_{7-8} = -(1 - f_{\text{DIR}}) \left( \frac{I_7}{\tau_7} \right)$ $\dot{m}_{7-10} = -f_{\text{DIR}} \left( \frac{I_7}{\tau_7} \right)$
8	$\dot{m}_{7-8} = (1 - f_{\text{DIR}}) \left( \frac{I_7}{\tau_7} \right)$ $\dot{m}_{8-9} = -\frac{I_8}{\tau_8}$
9	$\dot{m}_{6-9} = \frac{I_6}{\tau_6}$ $\dot{m}_{2-9} = \eta_2 \left( \frac{I_2}{\tau_2} \right)$ $\dot{m}_{8-9} = \frac{I_8}{\tau_8}$ $\dot{m}_{9-6} = -f_{9-6} \left( \frac{I_9}{\tau_9} \right)$ $\dot{m}_{9-10} = -(1 - f_{9-6}) \left( \frac{I_9}{\tau_9} \right)$
10	$\dot{m}_{9-10} = (1 - f_{9-6}) \left( \frac{I_9}{\tau_9} \right)$ $\dot{m}_{7-10} = f_{\text{DIR}} \left( \frac{I_7}{\tau_7} \right)$ $\dot{m}_{10-11} = -\frac{\dot{N}_{\text{T, burn}}}{\text{TBE}}$ $\dot{m}_{12-10} = \frac{I_{12}}{\tau_{12}}$
11	$\dot{m}_{10-11} = \frac{\dot{N}_{\text{T, burn}}}{\text{TBE}}$ $\dot{m}_{11-p} = -\frac{\dot{N}_{\text{T, burn}}}{\text{TBE}}$
12	$\dot{m}_{2-12} = \eta_2 \frac{I_2}{\tau_2}$ $\dot{m}_{12-10} = -\frac{I_{12}}{\tau_{12}}$
All components	Losses = $-\varepsilon_i \left( \frac{I_i}{\tau_i} \right) - \lambda I_i$

Table C.3 Mathematical definition of the terms appearing in ARC fuel cycle model (Fig. C.2).

---

<b>Flow rate fraction</b>	<b>Value</b>
$f_{5-1}$	0.33
$f_{5-3}$	0.33
$f_{5-6}$	$10^{-4}$
$f_{9-6}$	0.1
$f_{p-3}$	$10^{-4}$
$f_{p-4}$	$10^{-4}$

Table C.4 Tritium flow rate fractions for the fuel cycle model. The flow rate fractions from the FLiBe tank to the other components have been estimated according to the FLiBe flow rates specified in [3] ( $\sim 66\%$  of coolant flow rate to the main chamber,  $\sim 33\%$  to the divertors). An equal distribution from the main chamber to the FW/VV channels and the FLiBe tank was assumed (i.e.,  $f_{5-3} = f_{5-1}$ )

# Appendix D

## List of Acronyms

AI: Artificial Intelligence  
AUC: Area Under the Curve  
ARC: Affordable, Robust, Compact  
BDBA: Beyond Design Basis Accidents  
CANDU: CANadian Deuterium Uranium  
CDF: Cumulative Density Function  
CPS: Coolant purification system  
DBA: Design Basis Accidents  
DEMO: DEMONstration nuclear FPP  
DL: Deep Learning  
DIR: Direct internal recycling  
DNN: Deep Neural Network  
DSA: Deterministic Safety Assessment  
EBB: Encapsulated Breeding blanket  
FC: Fuel Cycle  
FCNN: Fully Connected Neural Network  
FMEA: Failure Modes and Effects Analysis  
FN: False Negative  
FNR: False Negative Ratio  
FP: False Positive  
FPP: Fusion power plant  
FPR: False Positive Ratio  
FPY: Full power year

FW: First wall  
HX: Heat Exchanger  
JET-ILW : Joint European Torus - ITER like wall  
IE: Initiating Event  
IFC: Inner fuel cycle  
ISS: Isotope separation system  
LBFGS: Limited-memory Broyden–Fletcher–Goldfarb–Shanno  
LIB: Liquid immersion blanket  
LOCA: Loss of Coolant Accident  
LOFA: Loss of Flow Accident  
LSTM: Long Short-Term Memory  
LWR: Light Water Reactors  
MC: Monte Carlo  
MFP: Metal foil pump  
ML: Machine Learning  
MLP: Multi-Layer Perceptron  
MSE: Mean squared error  
MSR: Molten Salt Reactor  
MSRE: Molten Salt Reactor Experiment  
ODE: Ordinary differential equation  
OFC: Outer fuel cycle  
OOM: Object-Oriented Modeling  
PCS: Primary Cooling System  
PDF: Probability Density Function  
PI: Proportional Integral  
PFC: Plasma facing component  
PSA: Probabilistic Safety Assessment  
RAMI: Reliability, Availability, Maintainability, Inspectability  
ROC: Receiver Operating Characteristic  
RL: Reinforcement Learning  
RIC: Relative sensitivity Index Coefficient  
RTM: Residence time method  
STEP: Spherical Tokamak for Energy Production  
TBE: Tritium burn efficiency  
TBR: Tritium breeding ratio

TES: Tritium extraction system

TH: Thermal-Hydraulics

TN: True Negative

TP: True Positive

TSS: Tritium self-sufficiency

TTF: Time-To-Failure

UQ: Uncertainty Quantification

VV: Vacuum vessel