Dynamic Simulation and Design Optimisation of Power Conversion Systems for a Pulsed Fusion Power Plant

Oliver M. G. Ward

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Chemical Engineering

University College London

30th July, 2025

I, Oliver M. G. Ward, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

The UK Atomic Energy Authority is leading the Spherical Tokamak for Energy Production (STEP) project to design and build a power plant that demonstrates net electricity generation from nuclear fusion. Due to challenges with tokamak operation, prototype fusion power plants may be operated in a pulsed regime. Additionally, thermal heat will be distributed to multiple tokamak components at different temperatures. Both of these complicate the design of the power conversion system. Components, like turbines, are sensitive to stresses from thermal transients. Heat sources must be integrated to utilise all available heat efficiently for net electrical power generation to be feasible. In this thesis, an optimisation methodology based on dynamic process models is developed to address this design challenge.

Dynamic models of possible power conversion systems for pulsed fusion power plants are developed using novel dynamic heat exchanger models made with computational complexity and ease of automated simulation in mind. Thermal energy storage is considered as a solution to the pulsed tokamak operation. These models are then used within the objective functions of multiple optimisation problems.

The optimal tuning of a load following proportional-integral-derivative controller is considered for a system with a constant heat source. The optimised controller shows large error reductions at the cost of increased low amplitude oscillations due to over-tuned integral action. The trade-off between thermal energy storage system size and fraction of nominal power output during a dwell is considered as a bi-objective optimisation problem for a tokamak with one or three pulsed heat sources. Both case studies show the potential for significant size reductions, but these designs lead to issues with low molten salt temperatures. The methodology is shown

to generate good designs within the constraints of a fixed process configuration, and could be extended to the design of other transient processes.

Impact Statement

Novel optimisation methodologies that use dynamic process models of thermal power conversion systems for process design are presented. To the best of the author's knowledge, such methodologies have not been developed previously for thermal power system, although similar coupling of optimisation and dynamic process models can be seen in dynamic optimisation studies for optimal operation [218, 248] and the simultaneous thermodynamic and control optimisation of solar thermal power plants with dynamic solar field models [14, 175]. Similar simulation-optimisation design methodologies have been applied in the literature to other processes [37]. They are demonstrated in the context of pulsed fusion power plants, but are general for any model developed in OpenModelica, which can model systems representable as systems of differential-algebraic equations. Novel lumped-parameter heat exchanger models are also developed. While they are developed for use in design optimisation, they could be useful in applications such as real-time optimisation or model predictive control with constrained simulation times [248].

Aspects of this work have been disseminated at multiple conferences, including a poster presentation at ChemEngDayUK 2024 and oral presentations at ICHEAP16 and ESCAPE33. Conference papers for ICHEAP16 [274], ESCAPE33 [273], ESCAPE35 [276] and a full journal paper [275] have also been published. These make the contributions of the work more accessible to other researchers.

Presentations based on this work were also given to multiple London schools at two separate outreach events. These presentations are hoped to have raised public awareness of fusion energy as a distinct technology from fission energy and shown the students the importance of STEM subjects for addressing global issues.

Acknowledgements

I would like to express my deepest thanks to Prof. Eric S. Fraga for welcoming me into his group at UCL, offering his time to help me whenever it was needed and encouraging my growth as a researcher. I would also like to thank Prof. Federico Galvanin and Dr Nelia Jurado for their supervision and guidance of my development and this project. They always found time in their schedules to give me guidance when I needed it. Thank you to all the other staff at UCL who have helped me over the last 4 years.

This project exists thanks to the funding of the United Kingdom Atomic Energy Authority, whom I want to thank for allowing me to take part in this work. From the guidance, discussions and support of Chris Clements, Mohamad Abdallah, Daniel Blackburn, Robert J. Warren and Jack Acres, this project became what it is today.

Thank you to all of my friends, both inside and out of UCL, who have been there for me during the highs and lows of the PhD. Thank you to my group, Andrés, Maryam and Raymond, for their friendship. Special thanks to Lucy and CheERS, without whom I would only know a small fraction of the friends I have made at UCL.

My family has changed a lot over the course of this PhD. I would like to dedicate this to my grandparents, Kenneth, Winifred, Peter and Jean, who always encouraged my curiosity. To my brother Chris and his partner Chantelle, and their beautiful children Alfie and Eleanor. To my uncle Mark, my dad Richard, Bev and my mum Alison. Without their constant love, encouragement and support, I would not be where I am today.

Finally, I would like to dedicate this to my partner Bea. The PhD has not been an easy journey, but her love, positivity and kindness (and humouring me by playing over-complicated boardgames) has kept me going through rain or shine.

Illustration by Beatrix Hatcher

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Global energy demands are rising due to population growth and industrialisation, requiring an equivalent growth in energy production. Global electricity generation specifically has grown by an average of 2.5% per annum between 2013 and 2023. Despite electricity generation by renewables growing by an average of 5.5% per annum over the same period, fossil fuels still generated 60% of the world's electricity in 2023 [122]. To meet global emission targets, electricity generation must shift towards cleaner technologies.

Renewables have been estimated to have lower levelised costs of electricity than conventional power plants at large scale [99], and are predicted to form a large fraction of electricity generation in the near-term future under many scenarios [39]. The high variability of solar and wind availability require either the integration of energy storage technologies or the flexible dispatch of other power generation systems to maintain stable grid operation. They are also constrained by requiring suitable geographic conditions.

Firm low-carbon energy sources, which can reliably meet demands in all seasons over long time periods, could reduce the cost of the decarbonisation of electrical grids, even if these technologies have much higher levelised costs than renewables. Short-term energy storage by batteries, although important, cannot fulfil the same role as firm sources in these scenarios [223]. Possible technologies include flexible fission plants, fossil fuel plants with carbon capture and storage, geothermal plants and fusion plants.

Fusion is an attractive prospect for future power generation, with the potential for no carbon emissions, abundant fuels, low land use and little long-lived radioactive waste. It also faces many novel engineering challenges before it can become an economically feasible energy technology. Among these are control of the plasma, extraction of waste products and heat, development of durable materials, self-sustaining fuel breeding and efficient conversion of heat into electricity by the power conversion system [45]. Fusion programmes, such as STEP and EU Demonstration Power Plant (EU-DEMO), must address all of these to demonstrate that fusion can be a commercial energy source. This work addresses aspects of the last challenge.

1.1 Nuclear Fusion

Nuclear fusion is the process by which atomic nuclei fuse together to form larger nuclei. From figure 1.1, the fusion of light nuclei is an exothermic process, with the excess energy being carried as kinetic energy with the reaction products. It is the exothermic fusion of hydrogen nuclei into helium that generates energy to sustain the Sun's core at a temperature of 15 million degrees Celsius. It is natural to consider whether this energy source can be replicated on Earth for power generation.

For fusion on Earth, deuterium-tritium fusion has been identified as one of the most promising candidates due to its higher reaction rates at lower temperatures than other fusion reactions. Deuterium and tritium are heavy isotopes of hydrogen with 1 and 2 neutrons respectively. Deuterium ^{2}H is a naturally occurring stable isotope of hydrogen that can be extracted from water. Tritium ^{3}H is an unstable isotope of hydrogen, with a half-life of 12.3 years, that can be produced by the neutron activation of lithium-6 or lithium-7. For long-term operation of fusion power plants, it is necessary for reactors to generate a self-sufficient supply of tritium fuel in breeder blankets of lithium [156]. Fusing a deuterium and a tritium nuclei together, a neutron and a helium-4 nucleus are produced [46]:

$$^{2}\mathrm{H} + {}^{3}\mathrm{H} \rightarrow \mathrm{n}\ (14.07\mathrm{MeV}) + {}^{4}\mathrm{He}\ (3.52\mathrm{MeV}) \tag{1.1}$$

The majority of the released energy is carried by the neutron. This can be intercepted

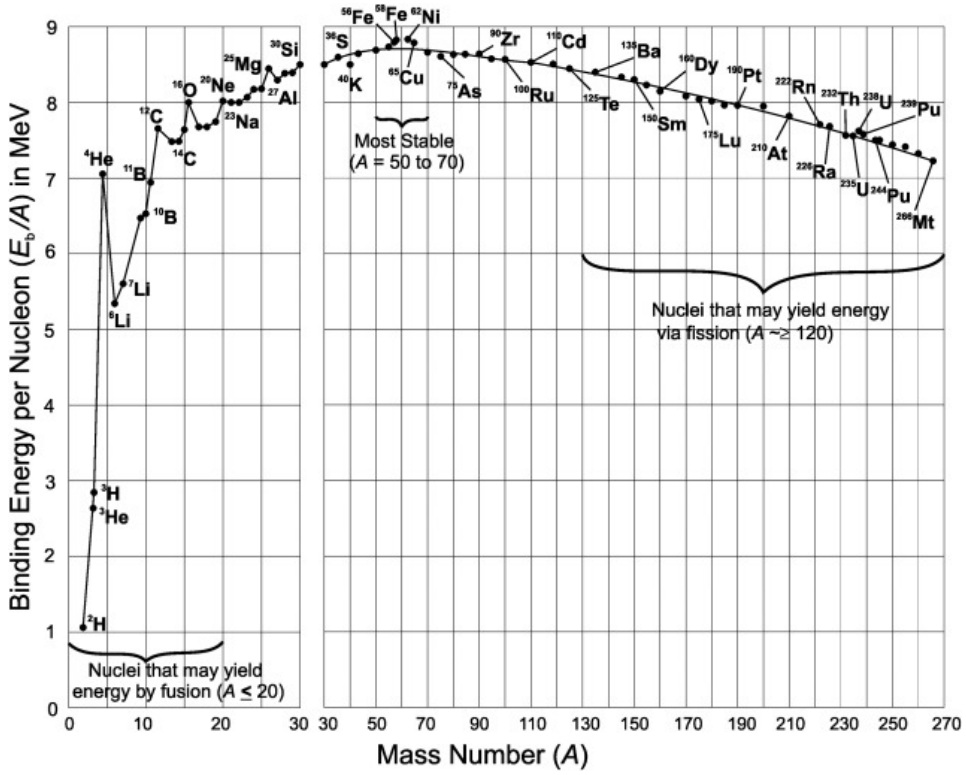

Figure 1.1: The average binding energy per nucleon of different nuclei shows that the fusion of lighter nuclei or the fission of heavier nuclei can release energy [140]. Image reproduced with permission of the rights holder, Elsevier.

by materials to convert the kinetic energy into thermal energy. While the reaction products are themselves stable, the neutron interception can lead to the formation of radioactive isotopes. This neutron activation of materials not only leads to radioactive waste that must be handled upon decommissioning, but also degrades the properties of materials. Research into neutron-resilient materials is a priority when moving from test devices to long-term operating power plants [203].

Fusion reactions require high temperatures and pressures, at which point the nuclei and electrons of the fuel dissociate into a charged plasma. These conditions must be maintained for long enough that sufficient fuel can undergo fusion to generate net energy. This requirement is described by the triple product inequality, which for deuterium-tritium fusion in the 116-232 MK range is:

$$nT\tau_E \geq 3 \times 10^{21} \text{keV s m}^{-3} \tag{1.2}$$

where n is the particle density, T is the plasma temperature and τ_E is the confinement time [172]. Even deuterium-tritium fusion requires temperatures >100,000,000°C for suitable reaction rates.

Looking at the slope of the binding energy per nucleon in figure 1.1, fusion can release more energy per unit mass of reactants than fission can. The energy density of deuterium-tritium fusion fuel, uranium-235 fission fuel and bituminous coal are of order 3.4×10^{14}J/kg, 2.1×10^{12}J/kg and 2.8×10^{7}J/kg [316, 49]. If tritium self-sufficiency is achieved, it is estimated that there is sufficient fuel to last tens of thousands of years due to the extremely high energy density [91].

A particular focus of fusion energy research is the reactor design. In a tokamak, the fuel is heated to become a charged plasma, at which point it can be controlled and confined, using magnets, into a ring within a larger vacuum vessel. Energy is released primarily as fast neutrons by the fusion reactions in the plasma. These neutrons are intercepted by the tokamak walls, transferring their kinetic energy to thermal energy. Additional energy is also transferred to the tokamak device by radiation, the exhaust of waste gases and breeding reactions. Energy will therefore be distributed to many tokamak components, which may need separate cooling systems.

Conventional power plants can control the thermal energy input to the system to maintain steady state conditions. It is desirable therefore for fusion reactors to operate the same, allowing for consistent power generation and standard power conversion system configurations. The steady state control of plasmas in tokamaks has proven to be challenging however. Hot spots, dust generation and wall flaking can lead to fast plasma disruptions. Additionally, technological advancements are required for steady state magnets, heating and current drive systems. It is therefore possible that, particularly for early demonstration power plants, operation of tokamaks may be pulsed [102]. Auxiliary heat sources, such as thermal energy storage or gas heaters, would be necessary to avoid thermal cycling of components and disconnection from the grid. This would then require alternative design methodologies to conventional power conversion systems.

If the associated engineering challenges can be overcome, fusion energy is

likely to be a critical part of future decarbonised power grids. A further discussion of the state of the art in fusion energy is presented in section 2.2.

1.2 Research Gap

As highlighted in section 1.1, two design issues for fusion energy are that heat will be available from multiple tokamak components, likely at different temperatures and powers, and that the operation of the tokamaks may follow a pulsed schedule. If the latter issue is assumed true, the process is inherently transient between when fusion reactions are occurring in the tokamak, termed a pulse, and when they are not, termed a dwell. Design procedure for conventional power conversion systems does not consider process dynamics until later in the design process, and are not suited for integrating multiple heat sources [168].

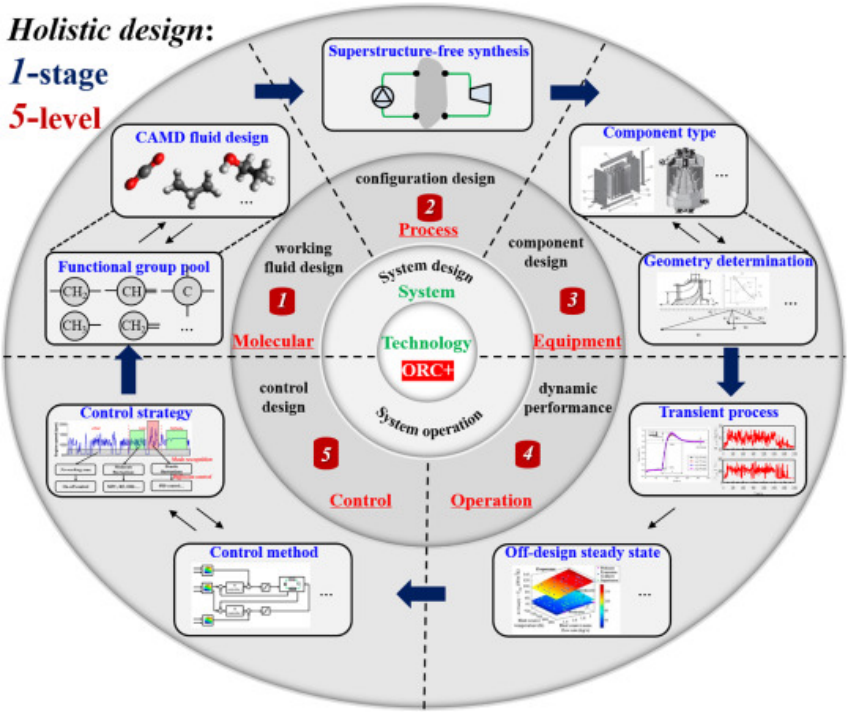

Figure 1.2: A proposed methodology for the design of organic Rankine cycles using fluctuating waste heat. Process dynamics should be explicitly considered in the design, as they are an inherent part of the system's operation [147]. Image reproduced with permission of the rights holder, Elsevier.

Similar design issues have been raised by researchers for organic Rankine cycles utilising fluctuating waste heat [147]. The need for an integrated design methodology

for such systems was highlighted, as off-design performance and process inertia will be part of the normal operation of the cycle. The researchers proposed a possible method for integrated design, shown in figure 1.2. While highlighted for organic Rankine cycles specifically, such a methodology could be expanded to the design of general transient processes. It was noted that no studies had yet considered time-dependent characteristics of organic Rankine cycles in the design. To the best of the author's knowledge, no study since then has considered the dynamics of the power cycle during design optimisation, for organic Rankine cycles or otherwise. Dynamic optimisation studies consider process dynamics in the optimisation, but are limited to a fixed process design where the control variables in time are chosen as the decision variables [218]. Other studies have used dynamic models of a solar field for optimisation, but neglect inertia in the power cycle models [14, 175]. Simulation-optimisation methodologies have been used for the design of other dynamic processes [37].

1.3 Thesis Outline

To address the highlighted research gap, this work considers the use of dynamic process models within an optimisation framework such that the transient performance of the system can be optimised. This methodology is then applied to the design of power conversion systems of pulsed fusion tokamaks, a review of which is provided in chapter 2.

Chapter 3 introduces the requirements and scope of the process modelling in this work. In chapter 4, the modelling of typical process components and fluids is presented. Heat exchangers are noted to be a significant source of inertia in thermal power systems, and as such should be represented by dynamic process models [43, 82, 121]. Novel heat exchanger models are developed specifically for use in optimisation.

Optimisation is widely used within the design of thermal power conversion systems, as presented in chapter 5. For this work, the plant propagation algorithm of Fresa.jl [92] is chosen as it is well suited for simulation-based black-box objective

functions. For multi-objective optimisation, a methodology using the multi-agent optimisation system Cocoa.jl [96] is also implemented.

The optimisation of a proportional-integral-derivative controller for load following is considered for a steam Rankine power conversion system with an ideal constant heat source in chapter 6. A molten salt thermal energy storage system allows for load following.

For a pulsed fusion reactor, the operation of the power conversion system shifts between pulse and dwell modes. Reduced power output during a dwell could be considered instead to allow for savings on the auxiliary heat system. In chapter 7, steam Rankine power conversion systems are presented for either a single pulsed heat source or three pulsed heat sources. These designs utilise molten salt thermal energy storage systems to provide thermal power for electrical power generation during dwells. Bi-objective optimisations are performed to assess the trade-off between the variability in turbine power output and the size of the thermal energy storage systems.

Chapter 2

State of the Art in Thermal Power Generation and Fusion Energy

The efficient use of thermal power from fusion devices by a power conversion system has been identified as a key design integration issue for programmes aiming to demonstrate fusion energy's viability [15, 1]. The state of the art in thermal power generation must therefore be adapted to the particularities of fusion devices as a heat source. A review is presented first of thermal power plants, and the current research into improvements of thermal efficiency and flexibility under variable operating conditions.

A fusion power plant will consist of a fusion reactor and a power conversion system. Current approaches to fusion reactor designs are outlined first, and then a review of current studies into power conversion system design for fusion energy is presented. This review is limited to tokamak devices, as this is the technology selected for current programmes, like STEP and EU-DEMO, to develop prototype fusion power plants.

2.1 Thermal Power Plants

Thermal power plants convert thermal energy into electrical energy. A useful metric of thermal power plant performance is the thermal efficiency η_{th}, which is the ratio of the net work produced per unit of thermal energy supplied. From the theory of thermodynamics, a fundamental limit on the maximum possible thermal efficiency

by a thermodynamic cycle can be defined. The Carnot cycle is a fully reversible cycle consisting of isothermal heat addition, isentropic expansion, isothermal heat rejection and isentropic compression. For a heat source at temperature T_{high} and a heat sink at temperature T_{low}, the thermal efficiency of a Carnot cycle is:

$$\eta_{\text{th,carnot}} = 1 - \frac{T_{\text{low}}}{T_{\text{high}}} \tag{2.1}$$

The Carnot cycle cannot be implemented practically due to irreversible losses, such as friction, heat losses and heat transfer over finite temperature differences, and the difficulty of achieving isothermal processes. Real power cycles compromise on thermal efficiency for feasible implementation in reality. The Carnot cycle is still a useful reference when looking for efficiency improvements in these power cycles however.

The most common power cycles used for thermal power generation are based on the Rankine and Brayton cycles. Both of these cycles consist of the isobaric heat addition, isentropic expansion, isobaric heat rejection and isentropic compression of the working fluid. If the working fluid undergoes phase change during the cycle, it is called a Rankine cycle. If the working fluid remains as a gas or supercritical fluid throughout the cycle, it is instead a Brayton cycle [34].

The steam Rankine cycle is a technology dating back to the 1900's, but is still a critical technology for contemporary power generation, with a global capacity of 1,500GW from coal-fired steam cycles alone in 2019 and thermal efficiencies as high as 46% [168]. Components for steam Rankine cycles are commercially available and high efficiencies can be achieved at high temperatures and power rates. They are also widely used for solar thermal power plants due to this availability, although the trade-off between decreasing solar field efficiency and increasing power cycle thermal efficiency with increasing temperature has led to interest in organic Rankine cycles [179].

Organic Rankine cycles use organic working fluids with lower boiling points than water. This allows them to efficiently utilise lower temperature heat sources than steam Rankine cycles. The working fluid can be selected to have a dry saturation

curve such that expansion through a turbine does not result in condensation, which allows for simpler cycle configurations [179]. They are also well-suited for waste heat recovery systems [111, 130, 112], and also for renewable energy sources like geothermal [118].

Brayton cycles are typically implemented as open cycles for combustion gas turbines, where ambient air is used as the working fluid. Fuel is mixed with compressed air and internally combusted, before the resulting flue gases are passed through a gas turbine. The flue gas is then expelled into the atmosphere, as it cannot be reused for combustion. This limits the exhaust pressure and the amount of power that can be extracted from the flue gas. The high temperature of the flue gas presents an opportunity for waste heat recovery. Combined cycle power plants utilise an open Brayton cycle with a steam Rankine cycle recovering waste heat to achieve thermal efficiencies of $> 60\%$ [168].

Closed Brayton cycles using supercritical fluids, such as carbon dioxide (CO_2) or helium, are a research area of interest. For supercritical CO_2 specifically, its thermodynamic properties could allow for more compact process components, higher thermal efficiencies for high temperature heat sources and suitability of air as a heat sink relative to steam Rankine cycles, among other benefits [146]. There is great interest in supercritical CO_2 Brayton cycles for nuclear applications, both fission and fusion [280]. The large parasitic loads of fusion plants make high thermal efficiency a necessity for net electricity generation [1]. Supercritical CO_2 Brayton cycles are still in the experimental demonstration phase however.

Traditional thermal power plant design is based on a single heat source. Novel plants are being proposed to integrate multiple heat sources, such as waste heat recovery from multiple processes. Structural synthesis of novel cycles has been considered for such systems, using optimisation methodologies, and is discussed more in section 5.1.2.

For all thermal power cycles, research can be seen to focus on improvements in efficiency and operational flexibility [168, 103], progress in which is presented in the following sections.

2.1.1 Efficiency of Thermal Power Plants

Efficiency improvements are dependent on the specifications of each given plant, as variations in the heat source, heat sink, ambient environment and more will impact the efficiency. There are some standard techniques that are often used to improve efficiency.

Common modifications to improve thermal efficiency in steam Rankine cycles are reheating, superheating and regeneration. Reheating and superheating address a particular problem of steam turbines that excess condensation of steam within the low-pressure sections of turbines results in erosion of the blades and a decrease in turbine efficiency. This limits the maximum and minimum pressure of the cycle, the increasing or decreasing of which will improve thermal efficiency respectively [34].

Reheating is the cycle modification by which the temperature of the steam is increased by the heat source after partial expansion through a turbine. This increases the mean temperature of heat addition, increasing efficiency, and allows the boiler pressure to be safely increased while maintaining safe moisture levels in the low-pressure turbine stages. This is seen in figure 2.1 by comparing the turbine outlet states of 5 with reheat and $5'$ without. This higher pressure improves the thermal efficiency more significantly.

Superheating the steam also reduces the moisture content in the low-pressure turbine stages, and also increases thermal efficiency by increasing the mean temperature of heat addition to the cycle. Referring to figure 2.1, it can be seen that if the steam could be heated to state $2'$, the reheat would be unnecessary to protect the turbine. However, the superheating temperature is limited by the heat source and material properties of process components

Regeneration uses extractions of hot steam to preheat the cold feedwater, increasing the mean temperature of heat addition and hence the cycle efficiency. This can be done by mixing streams in open feedwater heaters or via closed feedwater heaters where the streams are kept separate. The former has better heat transfer characteristics and are cheaper, but require additional pumps relative to closed heaters [34].

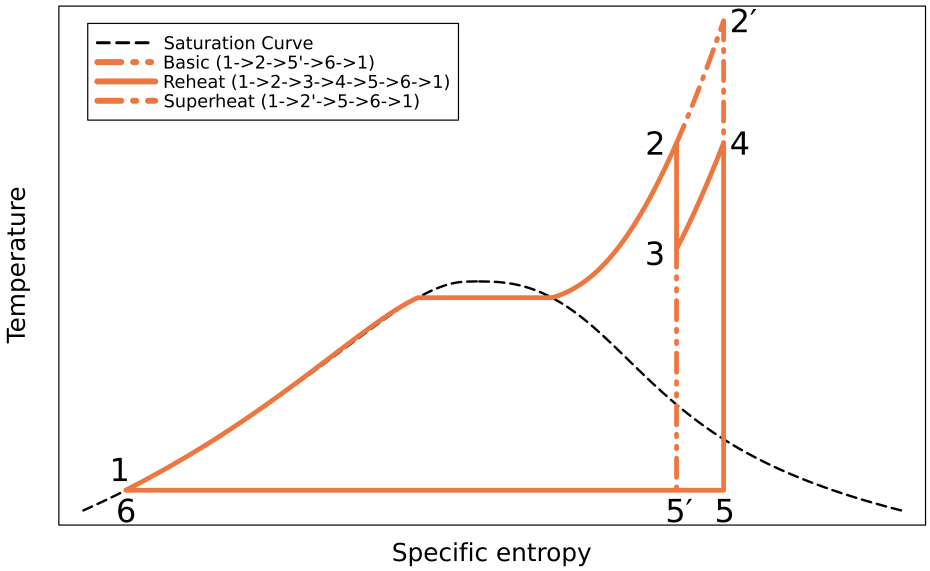

Figure 2.1: T-s diagrams of steam Rankine cycles with no modifications, with reheating or with superheating. Both modifications can protect the turbines from excess condensation at low pressures.

Valid cycle modifications are dependent on the specific working fluid. For example, cycle modifications like turbine split flow, inter-cooling, recompression and more have been proposed to improve the thermal efficiency of supercritical CO_2 Brayton cycles [154, 264].

The state of the art coal-fired steam Rankine cycles, so called ultra supercritical plants, can reach temperatures and pressures of $\geq$ 600°C, 270bar, far above the critical point of water. These cycles use reheat, or even double reheat, superheating and multiple feedwater heaters to achieve efficiencies of $\sim$ 50% [103]. Steam cycles in combined cycle power plants forgo regeneration as it harms heat recovery and the plant's net efficiency. Multiple evaporation pressure levels are used instead to match the flue gas temperature profile and increase heat recovery [67, 168].

For these conventional cycles, some researchers believe that the remaining possible efficiency gains are negligible with the exception of improved materials [103, 27]. Higher temperatures and pressures would raise the efficiency "ceiling" imposed by thermodynamics.

The structural design of conventional steam cycles is typically guided by previous industrial experience of the thermodynamics and economics [168, 103]. The

design then goes through a series of design stages, typically starting with steady state modelling and optimisation before moving on to detailed component design and finally dynamic modelling of key transients [168].

For non-conventional power cycles, like solar thermal and waste heat recovery plants, optimisation can supplement the lack of industrial experience, allowing effective novel cycles to be designed. The use of optimisation for design of thermal power conversion systems is discussed more in section 5.1.

2.1.2 Operational Flexibility of Thermal Power Plants

Power plants are increasingly required to operate under transient conditions. Increased electricity generation by renewables like solar and wind require that other plants are dispatchable to make up any shortcomings in generation. Solar thermal plants experience large fluctuations in thermal power availability, even during the day, that could strain the process components. Improved operability can be achieved via design modifications and upgraded process control. Dynamic modelling plays an essential role in flexibility studies, allowing alternative strategies to be assessed in a cost- and time- effective manner. Operability is often assessed by metrics such as start-up/shut-down time, maximum rate of load change and minimum load [168].

Alternative turbine control strategies can be considered to control the turbine power output in response to electrical load changes or frequency changes in the electrical grid. Admission governing of steam into the turbine allows for fast responses at the cost of throttling losses. Sliding pressure control avoids these losses at the cost of slower response times. These methods can be combined to achieve greater performance for plants that often operate at partial load, like combined heat and power plants [126].

Steam cycles are limited in their load following, start-up and shut-down capabilities due to sensitivity of components like turbines and boilers to stresses. Modifications such as electric heater blankets and increased temperature of sealing steam allow turbines to be kept warm in solar thermal plants for faster start-up flexibility [246], while alternative materials can withstand larger thermal transients [103]. Gas turbines are able to withstand much larger transients, but are typically

constrained in combined cycle plants by the steam Rankine heat recovery cycle. Steam bypasses and supplementary firing have been shown to improve the start-up flexibility of combined cycle plants [7].

Thermal energy storage is a widely used technology for solar thermal power plants, allowing them to operate without daily start-up and shut-downs and providing means to flexibly match generation to grid loads. As of 2017, 47% of thermal solar power plants operated with thermal energy storage, with an increase to $>$70% in future plants. The most established technology is sensible heat storage, where the temperature of a medium is increased to store thermal energy. Heat can then be exchanged from the medium to another fluid to discharge the stored energy. The medium can be a solid, like concrete or rocks, liquid, like molten salts or thermal oil, or gas, such as compressed air or steam. Latent and thermochemical heat storage technologies are of interest due to the possibility of higher energy densities, but are not yet viable for commercial plants [192].

Integration of thermal energy storage into conventional power plants has also been considered. Using extracted steam to charge a molten salt thermal energy storage system allows the minimum load of the plant to be reduced and improves the load increase response speed [143, 305]. Compressed air energy storage can take excess grid or plant power to be stored for later discharge [285].

The aforementioned modifications must be paired with an effective control system to function properly. Modifications of the control schemes alone can also improve flexibility.

Coordinated control of the turbine and boiler in coal-fired power plants can utilise the natural thermal energy stored in the boiler for fast load control. Extending this coordinated strategy by throttling turbine steam extractions, the steam can instead be used to generate power for increased load ramp rate. Thermal inertia of the feedwater heaters can meet high frequency load variations while lower frequency changes can be met by the boiler control, which has a slower response [311, 155]. Feedwater bypass throttling similarly controls the steam extraction, and can be used in conjunction with steam extraction throttling for improved control [270, 262].

For solar thermal power plants, control is particularly challenging due to intermittent and uncertain ambient conditions. Control of flow through the solar field and from the thermal energy storage system to the power cycle are areas of particular research [248]. These plants need to include long-term forecasting in their control scheme for robust operation.

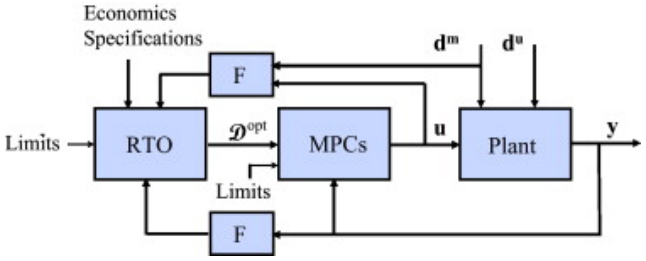

Figure 2.2: Implementation of real-time optimisation and model predictive control with a real plant. $\mathbf{d}^{m}$ and $\mathbf{d}^{u}$ are measured and unmeasured disturbances to the plant. RTO passes long-term set points $\mathscr{D}^{opt}$ to the model predictive controller, which in turn passes shorter-term set points $\mathbf{u}$ to the plant, which includes the distributed control system which directly controls the plant. The plant outputs $\mathbf{y}$ are fed back to the control system [66]. Image reproduced with permission of the rights holder, Elsevier.

Decomposing the control system into a hierarchy of time scales, such as in figure 2.2, can make the long-term control tractable. Real-time optimisation establishes optimal process set points over the time-frame of hours. These are passed to the model predictive control, which calculates optimal set point trajectories over a moving time horizon on the order of minutes based on a dynamic process model. These trajectories are passed to the distributed control system that directly regulates the plant. Typically, the distributed control is handled by proportional-integral-derivative controllers [66, 248].

Model predictive control has been shown to outperform stand-alone proportional-integral derivative controllers for thermal power plants, as they can handle multiple variables simultaneously as well as process constraints [121, 248].

Nevertheless, proportional-integral-derivative controllers are still considered in control studies of thermal power plants due to their simplicity [50, 49, 261, 153], even for highly transient pulsed fusion power plants [20]. Alternative control schemes based on them have also been developed to address some of their weaknesses, such

as the decoupling of two controlled variables [101].

It's been seen that the control scheme is driven largely by the heat source dynamics. Development of robust control schemes for fusion power plants will be critical for their operation, particularly prototype plants with uncertain operating regimes and risks of tokamak trips where the fusion reactor unexpectedly stops [1].

2.2 Fusion Energy

Research into nuclear fusion as an energy source began in the 1950's, with the global unveiling of previously classified research at the Second United Nations Conference on the Peaceful Uses of Atomic Energy. Various concepts for controlling plasmas were proposed, but the tokamak design quickly became the dominant design of interest when record performance was reported by a Soviet tokamak in 1968. Various national and international tokamak programmes began over the following decades, such as the Joint European Torus, the US Tokamak Fusion Test Reactor and the Japan Torus-60 Upgrade. These projects led to the start of the International Thermonuclear Experimental Reactor programme in 1988, a multinational collaboration to demonstrate the feasibility of fusion power [18]. It is currently being built in France, with an estimated start of experiments in 2034.

Advances in the area of fusion reactors are presented, followed by current studies on power conversion system design for fusion power plants.

2.2.1 Reactor Designs

Tokamaks are the most researched and common type of fusion reactor. They are an example of magnetic confinement fusion. Other categories are inertial confinement fusion and magneto-inertial fusion [190].

Magnetic confinement fusion uses magnetic fields to confine and control fusion in a charged plasma, and are the most common approach of public fusion programmes. The plasma density is typically low, so the plasma must be confined for longer and at higher temperatures to achieve suitable reaction rates for net energy output from the plasma relative to other reactor designs. In a tokamak, the plasma is confined to a ring, as shown in figure 2.3. Tokamaks can be further distinguished as

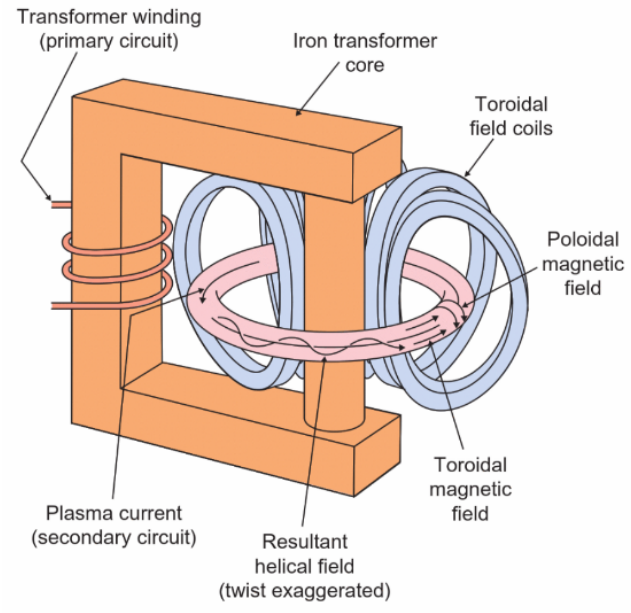

Figure 2.3: Magnetic confinement of ring-shaped plasma in a tokamak fusion reactor [190]. Image reproduced with permission of the rights holder, IOP Publishing Limited.

conventional or spherical, depending on the aspect ratio of the plasma. The latter is cited as having the potential for better plasma performance within smaller reactor volumes, with attractive prospects for smaller, cheaper and lower power spherical tokamaks for pilot fusion plants [63, 102]. A spherical tokamak has been selected as the reactor design for the STEP programme, while EU-DEMO has selected a conventional tokamak design more in line with the International Thermonuclear Experimental Reactor.

Inertial confinement fusion rapidly compresses fusion fuel to achieve the high plasma temperatures and pressures required for fusion reactions. The typical approach for compression is through the use of high-powered lasers, either targeted directly at a pellet of fuel or at a specially designed targetting capsule containing a fuel pellet. Proper targetting of the laser leads to a symmetrical implosion that can reach the necessary temperatures and pressures in a small hot spot in the pellet. An alternative approach to compression is to use shock-waves from a physical projectile to implode the fuel.

Magneto-inertial fusion blends the other two approaches, typically forming a plasma that is contained using magnets before being compressed and heated to

ignition by inertial techniques. This sort of approach is common in the private fusion sector, but is lacking in experimental testing and theoretical understanding that has been established for the other approaches.

Due to the long plasma confinement time required for magnetic confinement fusion, plasma control and current drive is a research focus. For inertial confinement fusion, the main challenges are the reliable targetting and implosion of the fuel, as well as making the waste disposal and replacement of the pellets fast enough that the devices can be pulsed at a high frequency. Magneto-inertial aims to avoid these issues, requiring short enough confinement times that plasma control is less critical, but requires reliable plasma production, injection and compression systems. These systems will likely be concept specific.

Steady state operation of fusion devices is desirable for the avoidance of thermal cycling of components both inside and out of the reactor, as well as for power grid integration. Inertial confinement fusion is inherently pulsed, and will require methods of rapidly repeating fuel injection and compression to be scaled to a viable power plant. For tokamaks, it has become apparent that steady state operation relies on improvements in key technologies like magnets, heating and current drive systems for viability. Plasma disruptions and dust generation due to hot spots could also limit the maximum pulse duration. Pulsed tokamaks, both conventional and spherical, have been proposed as a competitive alternative to steady state tokamaks for pilot fusion plants [220, 102].

2.2.2 Power Conversion Systems

The challenges of fusion power do not end once thermal energy has been generated in the reactor; this energy must then be extracted from the tokamak and converted into useful electrical energy. Considering tokamaks only, as they seem the most mature technology for demonstrating fusion power, thermal energy is available from multiple “sources”: fast neutrons from fusion reactions are absorbed by tokamak components; radiation from the plasma is absorbed by plasma-facing components; charged plasma coming into direct contact with certain components; exothermic breeding reactions in the blanket due to absorption of neutrons.

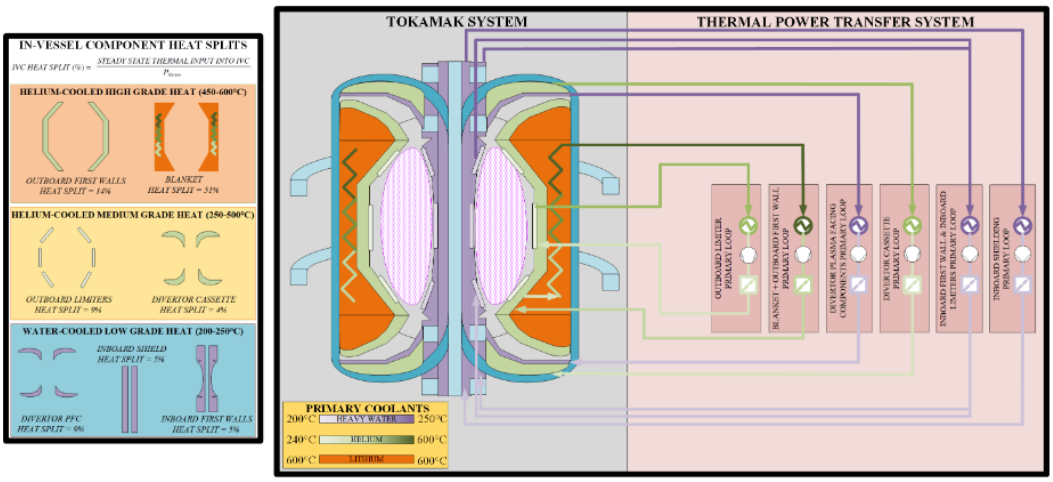

Figure 2.4: Expected heat splits and coolant streams of the tokamak in the STEP Prototype Plant [1]. Image reproduced under the Creative Commons Attribution License.

Regardless of the exact quantity of thermal power from each, the net effect is that heat will be distributed to multiple tokamak components in different amounts and at different temperatures. Multiple coolant streams are required to cool these components, and these coolant streams must be integrated into a power conversion system efficiently to enable net power generation. A possible distribution of heat in a tokamak, expected for the STEP Prototype plant, is shown in figure 2.4.

Integration of multiple coolant streams, and the possibility of highly transient thermal power output due to pulsed operation or tokamak trips, make the design of the power conversion system a challenge. Programmes such as the International Thermonuclear Experimental Reactor are not aiming to generate electrical power, focussing on the tokamak operation, with generated heat being discharged to a heat sink. For the next generation of fusion projects, such as the STEP, EU-DEMO and Affordable Robust and Compact programmes, conversion of heat to electrical power is a main objective, and power conversion system design has been identified as a key design integration issue [15].

EU-DEMO is planned to operate on a pulsed schedule with a pulse duration of 2h, a dwell of 10min and ramps of 100s between them. Two variants of the breeding blanket, the in-vessel component with the largest available thermal power, were selected during the pre-concept design phase, namely the water-cooled lithium lead

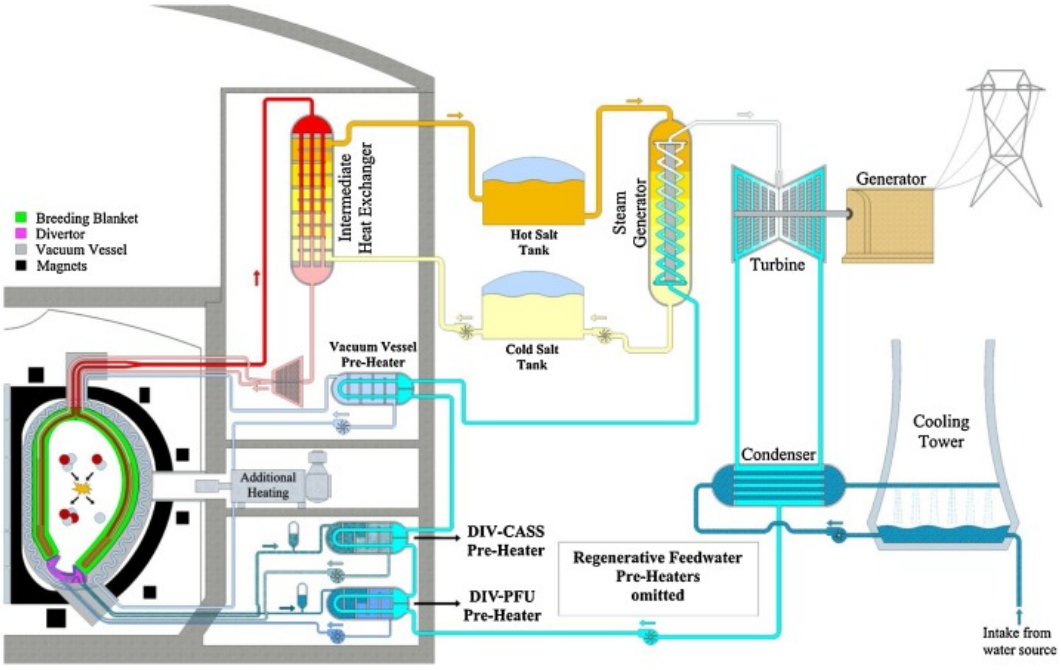

Figure 2.5: An indirect coupling design of the power conversion system for the pulsed tokamak of EU-DEMO. All heat from the breeding blanket is passed through an intermediate thermal energy storage loop to help decouple the tokamak and power cycle operation [19]. Image reproduced with permission of the rights holder, Elsevier.

and helium-cooled pebble bed technologies. Each has different design parameters, requiring alternative power conversion system designs, with a total of nine highlighted variants. To maintain grid synchronisation and minimise thermal stresses during dwells, the variants included combinations of molten salt or concrete thermal energy storage, electric heaters and gas heaters to provide thermal power during dwells at different load values. Molten salt thermal energy storage is an attractive option as it draws from established technology from solar thermal power plants. All variants used a steam Rankine cycle as the power cycle.

A distinction was made between direct and indirect coupling designs. Direct coupling designs are those where the power cycle exchanges heat directly with the breeding blanket coolant. Indirect designs, such as in figure 2.5, use an intermediate heat transfer loop equipped with thermal energy storage to further decouple the power cycle's operation from the tokamak's. The direct designs rely on auxiliary heaters to provide heat during a dwell, typically at a low fraction of nominal turbine

load to reduce the size of the auxiliary system. The indirect designs, if the thermal energy storage is suitably large, can maintain a constant turbine load across pulse and dwell, which is attractive for control and operational purposes [19, 21]. A design that integrates both breeding blanket technologies and coupling approaches has also been proposed to test both technologies and achieve better stability during a dwell than pure direct coupling designs [296].

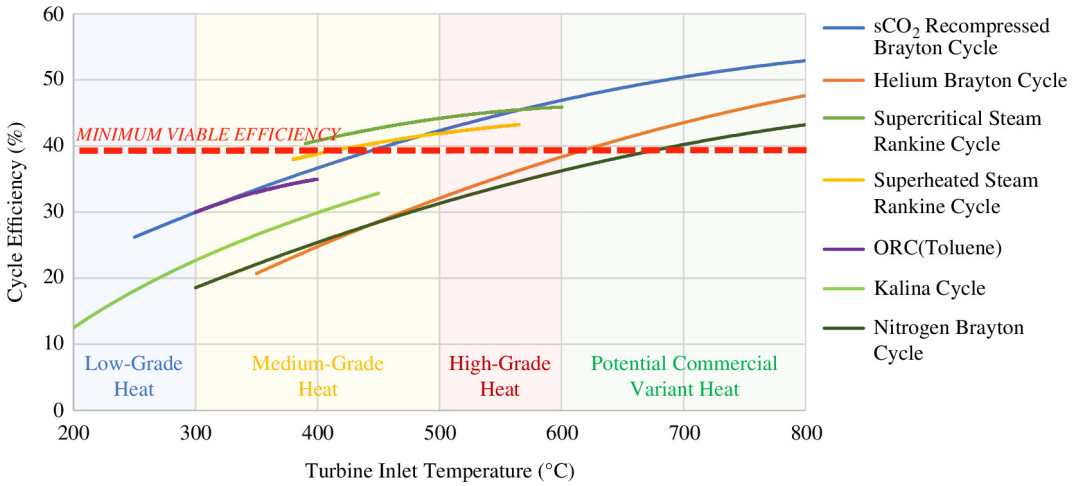

Figure 2.6: Possible efficiencies of different power cycle options for different turbine inlet temperatures. A minimum cycle efficiency for the STEP Prototype Plant suggests the supercritical CO_2 Brayton cycle and the supercritical steam Rankine cycle are the most viable options [1]. Image reproduced under the Creative Commons Attribution License.

The STEP Prototype Plant power conversion system design is instead planning for steady state tokamak operation, although the design should be able to handle the uncertainty of prototypic operation and resulting transients via auxiliary heat sources. Supercritical CO_2 Brayton cycles are being considered as the first option for the design, due to high efficiencies and flexibility, although steam Rankine cycle remains a fallback to consider uncertainties in development of the relevant technologies. The possibility of matching the different tokamak heat sources with multiple power cycles is also proposed [1]. The state of the art efficiency of the different power cycle options, and the efficiency required for net power generation, in the STEP Prototype Plant is presented in figure 2.6. Similar conclusions have been made for the power cycle options of reactors like the Affordable Robust and Compact reactor [58].

2.3 Summary

A summary of power cycles employed in thermal power conversion systems, and their typical applications, is presented. Steam Rankine cycles are the most established technology and are widely used, but there is great interest in novel technologies like organic Rankine and supercritical CO_2 Brayton cycles.

The current research into the efficiency of thermal power plants shows that steam Rankine cycles in conventional power plants are near the maximum efficiency that can be feasibly achieved by process optimisation. Non-conventional power plants lack the same level of industrial experience, and optimisation methodologies are being used to improve the efficiency of such systems.

Increased flexibility is desirable for thermal power plants due to the penetration of variable power generation from renewable sources in the grid. Dynamic simulations allow alternative control strategies and process designs to be tested under variable operating conditions in a cost-effective manner.

The alternative approaches to fusion reactor design are magnetic confinement, inertial confinement and magneto-inertial. Magnetic confinement designs, specifically tokamaks, are the most common approach in publicly funded programmes, and are considered to be a promising approach to fusion energy. Current research into the design of the power conversion system for prototype fusion power plants is discussed. Auxiliary heat, either by thermal energy storage or additional fossil fuel heaters, is identified as a necessary design feature to account for uncertain operation of tokamaks in prototype plants.

Chapter 3

Process Modelling of Fusion Power Plants

From the state of the art review in chapter 2, it is seen that a fusion power plant can be expected to consist of a fusion reactor, where fusion reactions generate thermal power, and a power conversion system, which converts this thermal power into useful electrical power.

The dynamic modelling of a fusion reactor is a daunting task, requiring the simulation of plasma dynamics, some aspects of which are still poorly understood [102]. Fortunately, the coupling of the fusion reactor and power conversion system dynamics can be considered to be largely uni-directional, with the reactor dynamics driving the power conversion system dynamics rather than vice versa.

There will be some feedback the other direction. Considering tokamaks, as they are the most mature fusion technology, the temperature of the coolant streams returning to the tokamak must be regulated to ensure that sufficient heat can be removed from the reactor. A significant fraction of the power being generated by the power conversion system will also be used to power the auxiliary subsystems of the reactor, such as magnet cooling and coolant pumping.

These feedbacks are neglected in this work. For the former, control systems are implemented to regulate coolant return temperatures. While the control is not perfect, it is assumed that any remaining temperature differences can be supplemented by auxiliary heating or cooling systems. For the latter, it is necessary that the fusion

reactor can be initiated without any power from the power conversion system for cold starts. Hence, it is reasonable to assume that the auxiliary systems are equipped to draw power from either the grid or the power conversion system, and that the amount of generated power does not impact the tokamak's operation.

Based on these assumptions, the scope of the modelling can be reduced to just the power conversion system, only requiring as inputs the flow profiles of the coolants exiting the tokamak. The specifics of the inputs to the power conversion system from the tokamak are specified in each case study presented.

To simulate the dynamics of a power conversion system, and hence guide the design of the process, the behaviour of the internal thermofluids in each component must be modelled. A brief introduction to the modelling tools and philosophy is given in the following sections, while detailed descriptions of the models used are given in chapter 4.

3.1 Modelica

Models are an abstraction of a system that can be used to answer questions about that system, and are an essential concept in science. For example, miniature scale-models of buildings and vehicles may be constructed to estimate aerodynamic properties of the real objects. Mathematical models describe relationships between variables of a system as mathematical expressions, and form our best understanding of the laws of nature. Even non-fundamental mathematical models are powerful tools in science. For example, Ohm's law allows us to make useful predictions of electrical circuits without needing to consider, or even be aware of, the dynamics and interactions of countless charged particles within the circuit [98].

The advent of computers opened the possibility of studying mathematical models by computational simulation that previously were too complex to analyse. Computational modelling has become common across all scientific disciplines [120], prompting the need for robust and multi-disciplinary software to represent and solve these models.

The Modelica language was released in 1997 for general-purpose modelling of

physical systems across domains, unlike other modelling tools available at the time [171]. Properties of the Modelica language include:

- Equation-based: Models are written in terms of equations instead of assigning values to variables
- Object-oriented: Models can be instantiated and extended in other models for ease of reuse
- Acausal: Inputs and outputs are not specified, and the solution direction can change depending on the available input data

These properties are well suited to physical modelling. Systems are often understood by mathematical equations, such as Ohm's law or Newton's second law, which can easily be translated into the Modelica language. A model can be defined for common objects, such as resistors or springs, which can then be instantiated multiple times with different parameters in another model instead of repeating equations for each instance. By not defining the input and output variables, the same model can be reused with no extra effort in different contexts, such as a circuit where either the current or voltage is fixed.

Modelica supports models represented as systems of differential-algebraic-equations, ordinary differential equations or algebraic equations, with support for discrete events and algorithmic assignments. Differential-algebraic-equations are common in physical systems due to conservation laws and constraints, and can be used to approximate systems described by partial differential equations. Modelica also includes a standard library of models across multiple domains, including thermofluids, heat transfer, electronics and mechanics [13].

Modelica is used as the language of many simulation environments, which implement the numerical solvers. Of particular note are Dymola and OpenModelica [239, 61]. Dymola is a commercial tool for modelling and simulation with state of the art solvers and specialised model libraries, and is widely used for thermal system modelling [42, 32, 72, 31, 213, 33, 90, 255, 305]. OpenModelica is a free, open-source modelling and simulation environment aiming to create an accessible

environment for research, teaching and industry. It has also been used for research on thermal systems [240, 301], although it has been noted that the solvers are less powerful than Dymola's in the past [64].

OpenModelica is chosen for this work to support open science and promote ease of reproducibility of the work. The chosen solver for the models presented is the default DASSL solver, a method that replaces derivatives with k^{th} order backward difference approximations for k between 1 and 5 and then solves the resulting equations using Newton's method. The order k and time step at each step is chosen adaptively based on the behaviour of the solution [193].

3.1.1 Initialisation

A system of differential-algebraic equations can be described in an explicit form as:

$$\boldsymbol{f}(\dot{\boldsymbol{x}}, \boldsymbol{x}, \boldsymbol{a}, \boldsymbol{p}, t) = \boldsymbol{0} \tag{3.1}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is the vector of state variables, $\dot{\boldsymbol{x}} \in \mathbb{R}^n$ is the vector of the state variables' time derivatives, $\boldsymbol{a} \in \mathbb{R}^m$ is the vector of algebraic variables and $\boldsymbol{p} \in \mathbb{R}^l$ is the vector of constant parameter values.

For a unique solution to be defined, n additional equations must be given for $t = 0$ to define initial values for $\dot{\boldsymbol{x}}$, $\boldsymbol{x}$ and $\boldsymbol{a}$. This is called initialisation, or the initial value problem, and is a challenging step for solvers, often requiring iterative solvers and simplified model transformations [228, 121].

When developing dynamic models, special care must be taken with initialisation to aid the solver's efforts [82]. In the Modelica language, this can consist of proper guess values for state variables, suitable initial equations and the use of the homotopy operator, which can be the difference between simulations failing at initialisation and completing successfully. In each case study, the initialisation equations of the dynamic process model are presented. In section 6.3.1, the behaviour of the process model after initialisation is discussed.

3.1.2 Modelica Fluid Modelling Framework

The Modelica Standard Library [13] includes the *Fluid* package, which implements a set of interfaces and models for the modelling of zero- or one- dimensional thermofluid flows. In this work, all models are developed using the *FluidPort* interface model provided in the package, designed for numerically robust simulation of bi-directional flows.

The Modelica language is designed around object-oriented programming, where individual component models can be connected together easily via connectors or inherit from other models. Connectors relate variables between the different models. The variables in a connector can be assigned three types: potential, flow or stream. Modelica can then automatically generate a set of connection equations between models.

Potential variables, such as pressure or electric potential, are scalar values that are equalised between connected interfaces. Flow variables, such as mass flow rate or current, have an associated direction and must sum to zero between all connected interfaces.

Stream variables were added to the language to improve numerical robustness of bi-directional flows [97]. These are convective transport quantities carried by the flow variable in a connector, of which there can only be one when using stream variables. For thermofluid flows, this includes specific enthalpy and flow composition. Stream variables are specified in both directions, irrespective of the actual flow direction, to avoid non-linear equations with Boolean variables for the flow direction. Modelica also automatically generates mixing equations for stream variables when more than two connectors are connected.

While flow reversal does not occur in any of the process models of this work, the robust *FluidPort* interface prevents Modelica from generating these non-linear equations regardless, and hence improves the overall robustness of the simulations.

3.2 Philosophy of Model Design

When developing computational models, a trade-off needs to be kept in mind between model accuracy and computational effort. Highly sophisticated models, such as computational fluid dynamics, can give accurate predictions at high resolution [208]. However, to achieve this level of accuracy, they require detailed information about the system geometry, fluid properties and physical processes. They also are extremely complex and time consuming to develop and solve. Together, these make them unsuitable for many applications, such as early process design and simulation.

Simplifications of models are necessary to make them tractable, while still capturing the key characteristics of the system. In the context of heat exchangers, even simplified models like the finite volume method and the moving boundary models encounter issues such as chattering [30, 204] and complex switching conditions [121] for two-phase flows.

In this work, dynamic models are used within an optimisation algorithm. Each design evaluation involves a full dynamic simulation, so models should be as computationally efficient and robust against numerical instabilities as possible for automated simulation. Additionally, at the early design stage, component parameters such as detailed pipe geometries and pipe materials are not available and would require excessive and wasted effort to estimate from literature if they are not fixed in the design.

Based on this, the models presented in chapter 4 emphasise the lumped-parameter modelling approach. Pressure dynamics are not considered as they lead to small time steps that increase simulation time [121], and typically has a shorter time-constant than the mass and energy dynamics [72].

3.3 Summary

By making two key assumptions about the feedbacks between the fusion tokamak and the power conversion system, the scope of the required process model is reduced to the power conversion system alone. Modelling such systems primarily involves modelling the flow of thermofluids.

Modelica is a multi-domain modelling language that is well-established for the modelling of thermal systems. The Modelica Standard Library includes a framework for thermofulid modelling that can automatically and computationally efficiently generate connection equations between different process components. An example of such a power conversion system model, as seen in the graphical view of OpenModelica, is presented in figure 3.1.

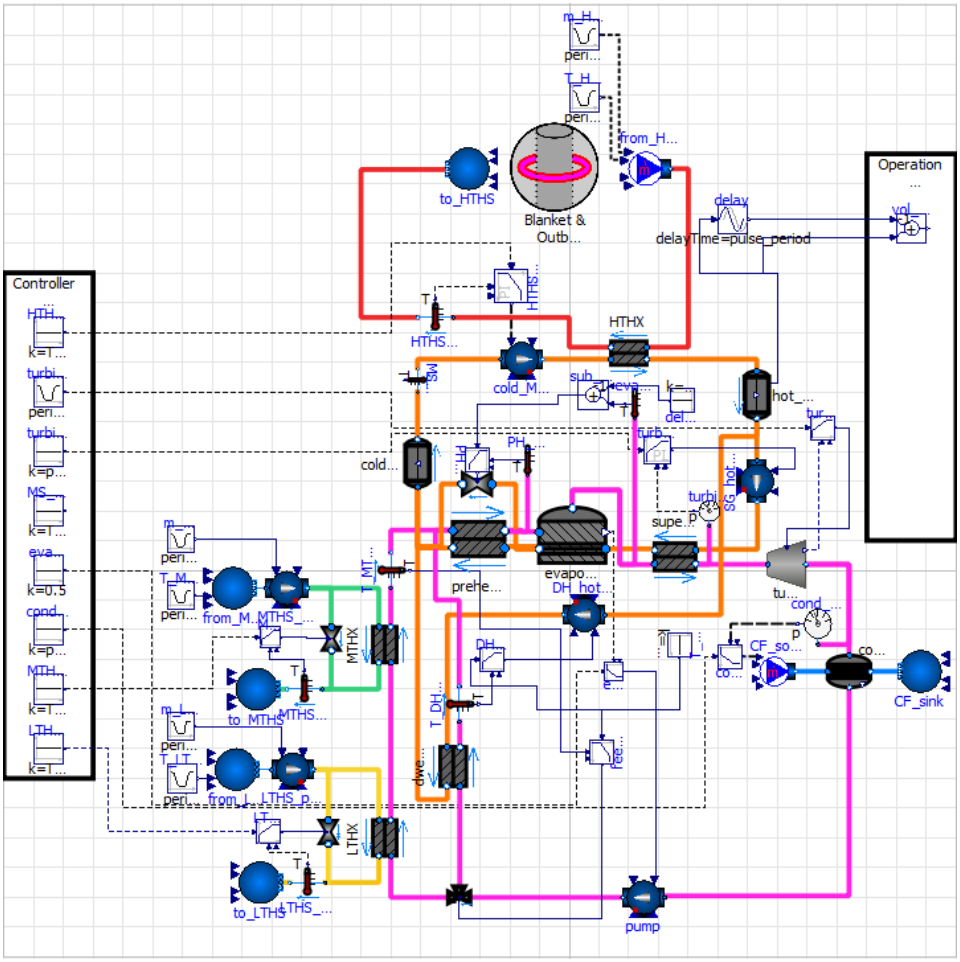

Figure 3.1: The power conversion system model of Case Study III, as seen in the OpenModelica graphical view. Individual component models can be reused and easily connected using the Modelica *Fluid* library connectors.

For the purpose of automated design via optimisation, where a large number of simulations will be required, a modelling philosophy is outlined to emphasise lumped-parameter models. This minimises the computational cost of simulations, allowing for a wider search of the design space, while hopefully still capturing the key process dynamics in the design process.

Chapter 4

Modelling of Thermal Power Conversion Systems

Publications based on this chapter:
"*Lumped-Parameter Heat Exchanger Models for the Robust Dynamic Modelling of Power Generation Cycles*" [273]
"*A dynamic model of a power conversion system with indirect thermal energy storage for a pulsed fusion tokamak for use in design optimisation*" [275]

Heat exchangers transfer thermal energy between fluid streams, allowing it to be transported and utilised in closed loop power cycles. Turbomachinery converts between mechanical, kinetic and potential energy, while valves and controllers are necessary to regulate the operating conditions of the plant. Accurate models of the fluids themselves are also necessary to predict how the thermodynamic properties vary around the cycle. Any auxiliary heat sources must also be modelled.

In this chapter, current approaches to modelling the different components of thermal power conversion systems are reviewed. The review is limited to components that are relevant to steam Rankine cycles and molten salt thermal energy storage systems as the scope of the project could not include alternative technologies. Based on the literature, models for these components are then presented for use in the three case studies of this work. For dynamic simulation, the models must be able to handle off-design conditions.

4.1 State of the Art in Modelling of Thermal Power Conversion Systems

4.1.1 Thermodynamic Properties of Fluids

Thermodynamic properties are measurable quantities that describe the state of a system, such as specific heat capacity and density, and transport properties [80]. When modelling thermal processes, such as thermal power plants, it is necessary to have suitably accurate models of the media properties for the model to be of any use. Developing such models is very labour intensive, requiring extensive experimental data covering a wide range of conditions and phases for each medium, further complicated when considering mixtures of these media.

For the purpose of this review, the scope will be limited to the thermodynamic property modelling of pure fluids in the context of thermal power plants, where they undergo large changes in temperature and pressure around the cycle. The thermodynamic properties of metals is discussed in section 4.2.7. All the fluids considered in this work are pure except for the molten salt, for which a detailed thermophysical model was not available. Details on thermophysical modelling of mixtures can be found in [230, 231, 195, 116]. Modelling of transport properties, such as viscosity and thermal conductivity, is also not considered here. Following the modelling philosophy outlined in section 3.2, detailed pressure loss modelling and axial heat conduction, which would require transport properties to model, are not considered to avoid the requirement of detailed parameterisation of component geometries, as well as to reduce computational simulation times. The modelling of transport properties is discussed more in [24, 116].

Comprehensive thermodynamic property modelling of fluids can be traced back to the 1840s, when many groups made efforts to measure the properties of steam at different conditions and develop steam tables that related temperature, pressure, heat content, volume and other thermodynamic variables. Different sources disagreed on steam properties, particularly at high temperatures and pressures, eventually requiring international collaboration for standardisation and the formation of the

International Association for the Properties of Water and Steam [107]. Beyond water, there was a need in the 1980s to model the properties of refrigerants for the development of new fluids to replace chlorofluorocarbon, which had been damaging the ozone layer [116].

The state of the art in thermodynamic property modelling emphasises the equation of state approach, by which all of the thermodynamic properties are calculated from a single equation. This ensures thermodynamic consistency between the properties. Early refrigerant models based on equations of state include the modified Carnaham-Starling-DeSantis equation [177], the modified Benedict-Webb-Rubin equation [117] and the extended corresponding states model [115]. The former and latter of these were appealing as they are based in theory, cover liquid, gas and supercritical regions and required minimal experimental data to parameterise. The modified Benedict-Webb-Rubin equation required more experimental data, having 32 parameters to fit, but could be applied to polar fluids and predict properties close to experimental uncertainty. The equation is explicit in pressure, of the form $p = f(T, \rho)$, and requires an additional equation for the specific heat capacity of an ideal gas as a reference state to calculate other thermodynamic properties by integration [116].

The requirement of integrals for calculating thermodynamic properties limited the mathematical expressions that could be used. Fundamental equations of state are those that require only derivatives to calculate all the single-phase thermodynamic properties, allowing more flexibility in mathematical expressions that can be included in the equation of state [142, 116]. These can be formulated from expressions for the specific Helmholtz energy $a = f(T, \rho)$, specific Gibbs energy $g = f(T, p)$, specific enthalpy $h = f(s, p)$ or specific internal energy $u = f(s, \rho)$. Experimentally measuring temperature T, pressure p and density ρ is much easier relative to specific entropy s. T and p can not uniquely define a thermodynamic state on phase boundaries or two-phase regions, unlike T and ρ. For these reasons, formulating equations of state in terms of the Helmholtz energy is preferred.

Generally, the reduced Helmholtz energy $\alpha = a/RT$ is expressed as the sum

of an ideal gas term α_0 and a residual term for real gas behaviour α_r in terms of reduced quantities $\tau = T_c/T$, $\delta = \rho/\rho_c$:

$$\alpha(\tau,\delta) = \alpha_0(\tau,\delta) + \alpha_r(\tau,\delta) \tag{4.1}$$

where R is the ideal gas constant and T_c, ρ_c are generally the critical temperature and density of the medium [231, 116]. α_0 can be fitted relatively easily from experimental measurements in the gas phase or spectroscopic data. The residual term is more complex, and consists of a sum of various types of expressions, including polynomial-like, exponential and Gaussian bell-shaped terms. These are empirical terms, and cannot be directly correlated to theory.

Fitting α_r to experimental data is a computationally intensive task, with the current state of the art equation of state for water, IAPWS-95, having 56 terms, each having multiple parameters to fit [259]. Fitting techniques use both linear regression and non-linear methods on a subset of the available data. The former can be used first to optimise the structure of the equation, adding or removing terms to α^r from a bank of possible expressions. The latter can then iterate on this equation, minimising the error in multiple measured properties simultaneously by adjusting parameters or removing terms entirely [142]. The final equation can be tested against the data not used for fitting.

Once an equation for $\alpha(\tau,\delta)$ has been fitted, other thermodynamic properties can be calculated from thermodynamic theory [142]. For example, pressure p and specific enthalpy h are given by:

$$p = \rho RT\left[1 + \delta\left(\frac{\partial \alpha_r}{\partial \delta}\right)_\tau\right] \tag{4.2}$$

$$\frac{h}{RT} = \tau\left[\left(\frac{\partial \alpha_0}{\partial \tau}\right)_\delta + \left(\frac{\partial \alpha_r}{\partial \tau}\right)_\delta\right] + \delta\left(\frac{\partial \alpha_r}{\partial \delta}\right)_\tau + 1 \tag{4.3}$$

For a pure fluid, the Maxwell criterion states that the temperature, pressure and Gibbs free energy are equal at the equilibrium between the liquid and vapour phases. This must be exploited to calculate the thermodynamic properties at the saturation

curves and in the two-phase region. This is an iterative process for the values of the saturation pressure and the saturation densities of the liquid and vapour phases. The successive substitution method is simple to implement, iteratively guessing the saturation pressure, but is sensitive to the quality of the initial values [230]. Other methods based on the Newton-Raphson method [3, 166] and Chebyshev expansions [22] have been proposed. Iterative solvers can struggle to converge near the critical point, requiring method modifications such as a spline extensions of the saturation curves [24] or extended precision arithmetic [22].

Taking the IAPWS-95 as the state of the art in thermodynamic modelling of any pure fluid, having the most extensive set of experimental data, the possible accuracy of these fundamental equations of state is illustrated in figure 4.1.

Models for the thermodynamic properties of a collection of fluids are available in various computational libraries. These are often designed with computational efficiency in mind for applications such as process simulation. Using the scientific standard formulation IAPWS-95 for water and steam, the current industry standard IAPWS-97 was developed to balance computational speed and accuracy [258]. The largest available libraries are REFPROP [116], with models of 147 pure fluids and 5 pseudo-pure fluids, and the open-source CoolProp [24], with models of 122 pure and pseudo-pure fluids. Both libraries have wrappers for use in various programming languages, and favour the Helmholtz formulation of equations of state where possible for a particular fluid. CoolProp implements interpolation methods, using cached property calculations to reduce computation time by $> 98\%$. REFPROP is undergoing a complete rewrite to use new thread-safe algorithms and techniques, such as automatic differentiation, to achieve greater computational speed and ease of implementing new fluid models.

4.1.2 Turbomachinery

Turbomachines are responsible for transferring energy between a rotor and a fluid [8]. For thermal power conversion systems, the notable turbomachines are pumps, compressors and turbines. Pumps and compressors transform mechanical energy of a rotating shaft to the potential and kinetic energy of a incompressible or compressible

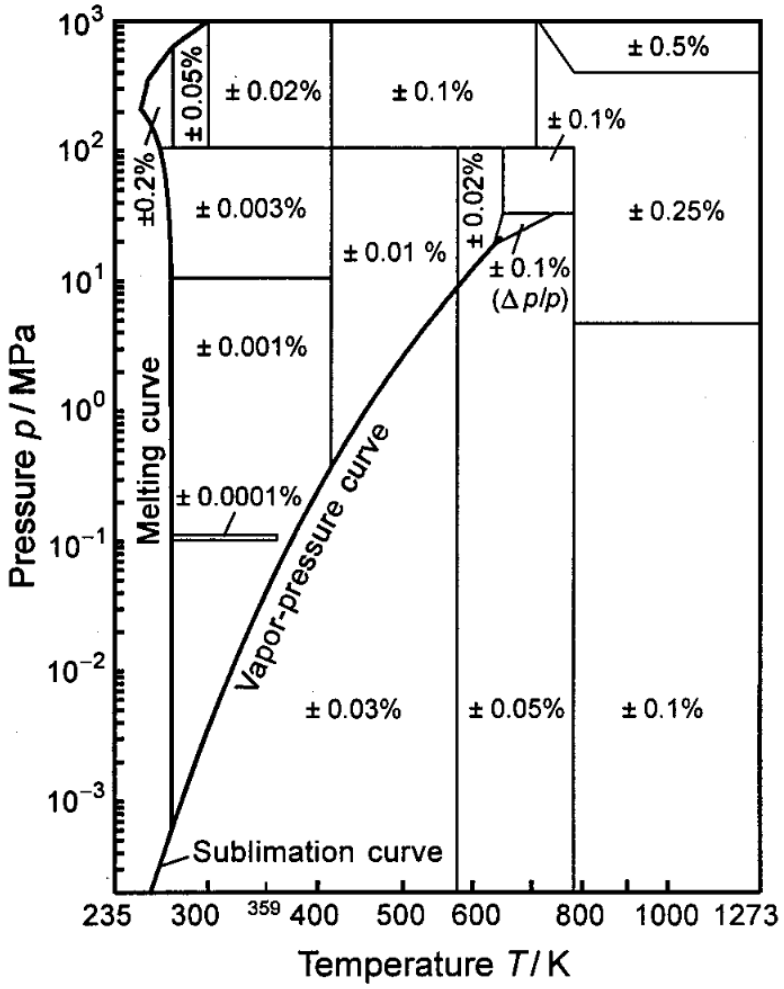

Figure 4.1: Percentage uncertainties of the density predicted by IAPWS-95 equation of state across its valid temperature and pressure range [259]. Image reproduced with permission of the rights holder, Elsevier.

fluid respectively, generating flow and increasing the fluid pressure. Turbines are the reverse, converting the energy within a fluid into mechanical energy by having the fluid expand to a lower pressure, converting potential energy into kinetic energy that can be used to drive the rotation of a shaft. An emphasis is made here on pumps and steam turbines, as they are the turbomachines present in steam Rankine cycles.

Due to the complex geometry and mechanics of turbomachinery, mechanistic modelling of their performance often relies on complex computational fluid dynamics simulations using the Navier-Stokes equations [103]. These models are used with optimisation algorithms to optimise the shape of blades and nozzles within the units, particularly for turbines where small improvements in efficiency can translate to

significant amounts of power. Objectives for these optimisations can include the formation of water droplets in the steam, which are associated with efficiency losses and blade erosion [182], and stresses in the blade [183].

For process simulation of turbomachines integrated with other process units, computational fluid dynamics models are not feasible due to the computational costs. Lumped-parameter methods are employed to reduce the model complexity, and these typically rely on some combination of similarity relations and empirical data. Performance at design conditions will be provided by manufacturers or units will be selected to fit the design operating conditions. The challenge is to model the behaviour of a turbomachine at off-design conditions, which will be necessary for any comprehensive design that considers transient operating conditions, such as load following in the context of thermal power plants.

To model the off-design performance of pumps, similarity relations can be used. These relate the behaviour of two systems based on three similarity conditions: geometric similarity, where all characteristic lengths differ by the same scale; kinematic similarity, where all characteristic velocities differ by the same scale; and dynamic similarity, where all forces differ by the same scale [44]. Geometric similarity is automatically satisfied when considering the same pump, but kinetic and dynamic similarity are weaker assumptions.

Under the three similarity conditions, the following relations between volumetric flow rate $\dot{V}$, rotational speed N and pump head H can be derived for a pump at two different operating points, denoted by subscripts 1 and 2:

$$\frac{\dot{V}_1}{\dot{V}_2} = \frac{N_1}{N_2} \tag{4.4}$$

$$\frac{H_1}{H_2} = \left(\frac{N_1}{N_2}\right)^2 \tag{4.5}$$

A single performance curve, relating two of the three variables, must also be provided to fully define the relationship between $\dot{V}$, N and H [250, 8].

Experimental analysis of volumetric flow rates for three pumping stations showed relative deviations of up to -7.9% between measurements and theoretical

similarity values for changing blade angles and rotational speeds [312]. This was attributed to the hydraulic losses due to residual velocity circulation being non-similar. An experimental comparison of a 15kW five-stage pump with different amounts of static head showed an error of 10.69% in the predicted head from the similarity laws at 120% of the rated speed, although this reduced to -2.96% for 80% of the rated speed under the same static head conditions [74]. The conclusion was that the similarity laws are more suitable at lower-than-rated speeds.

If a specific pump has been selected, empirical performance data may be available for a wide range of operating conditions. This data typically gives the values of H and the hydraulic torque τ_h at different values of $\dot{V}$ and N, covering all operating modes, such as turbine mode, where flow direction or shaft rotation can be reversed. These can be provided as tables of data, requiring two-dimensional interpolation for implementation in models.

For numerical modelling, it is advantageous to transform performance data into homologous curves based on the similarity relations, allowing the data on H or τ to be represented by a single, bounded curve, each composed of eight partial curves. In this form, only one-dimensional interpolation is used [82, 189]. Performing a further transformation of these homologous curves into a parametric form of two closed curves has been demonstrated to be even more advantageous for numerical applications, avoiding issues of zero-valued denominators and sign changes [253].

Lumped-parameter modelling for the off-design performance of steam turbines typically uses some form of the empirical relation known as Stodola's Ellipse law, which relates the mass flow rate $\dot{m}$ through a turbine and the inlet and outlet pressures P_{in}, P_{out}:

$$\dot{m} = K\sqrt{\rho_{in} P_{in}\left(1 - \frac{{P_{out}}^2}{{P_{in}}^2}\right)} \tag{4.6}$$

K is the Stodola constant and ρ_{in} is the inlet density. K can be calculated from a design operating point. The equation is derived from the "nozzle analogy", which assumes the expansion through a turbine group is equivalent to expansion through a single nozzle, and the empirical observation that the mass flow coefficient $\dot{m}/\sqrt{P\rho}$ is approximately constant at all points of the expansion downstream from a nozzle

throat. While the derivation assumes infinite expansion stages in the turbine, it has been empirically validated for as few as eight stages by Stodola [62].

The simplicity of equation 4.6, combined with the minimal design data required to parameterise it, lends itself to power plant optimisation studies, where many evaluations are required [232, 49, 105, 255]. It is also used for simulation studies on flexibility of power plants for load following, start-up or shut-down [43, 49, 307, 160, 263]. A more generalised flow model can also be considered [150]. The steady state flow equations can be justified by the short response times of turbines relative to other components in a power cycle [43, 160].

Two methods for turbine control are full arc (throttle) and partial arc (nozzle) governing. In full arc governing, the control valves at the inlet (typically four valves) to the high pressure turbine are controlled simultaneously. At part-load operation, pressure losses from the valves can cause significant reduction in the turbine's efficiency. Partial arc governing reduces this efficiency loss at partial loads by having the control valves operate sequentially, allowing some valves to be left fully open and therefore reducing the pressure losses. The disadvantage of partial arc governing relative to full arc is the larger stresses on the first stage blades. Full arc governing is common for combined cycle power plants, where the gas turbine can be controlled for faster load following, while partial arc governing is more usual for single cycle power plants [150, 241].

Modelling full arc governing can be done by considering a single control valve for the steam flow entering the turbine [25], the modelling of which are discussed more in section 4.1.3. Partial arc governing is more complicated due to the non-uniform valve operation. Most realistically, each individual control valve can be modelled in parallel with their own efficiency and flow calculation, with the valves being closed or opened sequentially. In a full process simulation, these parallel flows would likely increase the computational time required. Instead, partial arc governing can be assumed to be perfect, with no impact on turbine efficiency, and the partial arc opening fraction can be included as a scaling factor in the mass flow rate equation [41]. This will overestimate the turbine's efficiency however.

The efficiency of turbomachines will change with their operating conditions. The off-design efficiency can be measured and provided by manufacturers for specific units. Alternatively, efficiency curves can be estimated from other works [233, 232, 90]. If only a nominal efficiency is known, the efficiency can be modified based on empirical relations to account for condensation in the turbine [82] or the velocity of steam exiting the turbine [255].

Under highly-transient conditions, such as start up, shut down or power supply failure, the dynamics of turbomachinery can be important. It is during these periods that pumps encounter the largest pressures and turbines undergo significant thermal transients, which can put excess stress on components and reduce their lifetime [42, 21]. Rotational inertia, mass storage and thermal inertia may need to be considered for accurate dynamics, depending on the context.

Rotational inertia can be modelled using an angular momentum balance for the turbomachine shaft with rotational inertia I and angular speed ω:

$$\frac{\mathrm{d}}{\mathrm{d}t}\omega = \frac{\dot{W}}{I\omega} \tag{4.7}$$

where $\dot{W}$ is the net power applied to the shaft [89, 43, 82].

Thermal inertia and mass storage can be modelled by considering the volume inside a turbomachine using lumped-parameter mass and energy balance equations, either thermally coupled to the walls of the machine [250, 43] or thermally insulated from them [133].

4.1.3 Pressure Drops

Pressure changes occur along a flow due to frictional forces, changes in geometry of the flow vessel, elevation changes and fluid density variations. For thermal power conversion systems, the main areas of interest are the distributed frictional pressure drops, such as fluids flow through pipes or heat exchangers, and local pressure drops, such as over valves. Pressure losses impact the thermal efficiency of cycles, with significant economic gains possible by reducing their magnitude [148].

Modelling frictional pressure drops is a complex task, as the forces are de-

pendent on fluid properties, flow geometry and material roughness. Along a pipe segment of length dL, the frictional pressure drop dp_f is generally modelled as:

$$\mathrm{d}p_\mathrm{f} = \frac{f\dot{m}^2}{2D_\mathrm{h}\rho A^2}\mathrm{d}L \tag{4.8}$$

where f is the friction factor, D_h is the wetted perimeter of the pipe and A is the cross-sectional area of the flow [82]. The calculation of f is calculated by a staggering variety of empirical or semi-empirical correlations, the selection of which depends on the flow conditions, such as whether the flow is laminar or turbulent, the flow pattern, for two-phase flow, and the relative roughness of the pipe walls. A review of the correlations is omitted here, but please refer to recent reviews [12, 299, 88, 194]. Equation 4.8 can be used in distributed-parameter models, such as discussed in section 4.1.6.2, or integrated over the length of a pipe for a lumped pressure drop model, assuming constant values for the spatially varying quantities [82].

Valves are typically modelled in terms of a flow coefficient C_V:

$$C_\mathrm{V}^2\Delta p = \frac{\dot{m}^2}{\rho_\mathrm{in}} \tag{4.9}$$

where ρ_in is the inlet density of the fluid. This flow coefficient is a function of the valve opening position $0 \leq \Omega \leq 1$, with $\Omega = 0$ being fully closed and $\Omega = 1$ being fully open.

Different valve characteristics $C_\mathrm{V}(\Omega)/C_\mathrm{V,max}$ can be selected depending on the specific valve, with commonly chosen characteristics including linear, parabolic and equal percentage [8, 82, 314]. With a reference point for the valve at full opening, $C_\mathrm{V,max}$ can then be calculated.

4.1.4 Sensible Thermal Energy Storage

Sensible thermal energy storage stores thermal energy by raising the temperature of some insulated storage medium. For liquid media, either single- or two- tank configurations can be used, as shown in figure 4.2. The most relevant use of such storage systems is in solar thermal power plants, where excess thermal energy generated during sunny conditions can be used to maintain power generation during

cloudy conditions or at night. As discussed in section 2.2.2, there is great interest in sensible thermal energy storage for fusion power too. Common storage media used in such systems are thermal oils and molten salts, with molten salts generally being favoured for higher temperature stability and lower costs than thermal oils [79, 298, 213, 127].

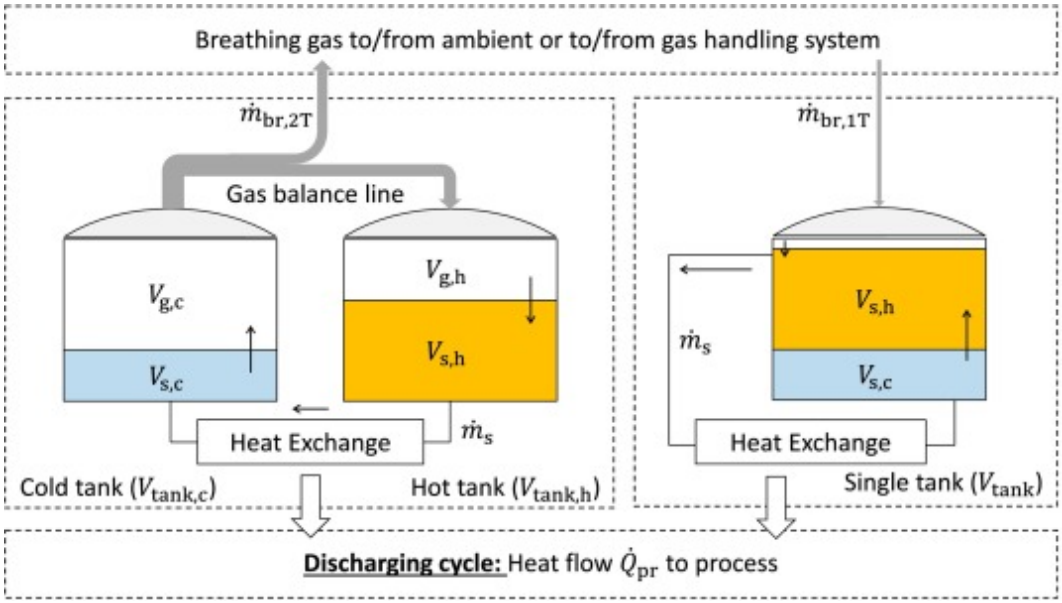

Figure 4.2: Discharge of thermal energy from a single- or two- tank sensible thermal energy storage system. During charging, the flow directions are reversed, either via the same pipes or via a separate channel. Gas exchange maintains a near constant pressure in the inert gas atmosphere [135]. Image reproduced under the Creative Commons Attribution License.

Dynamic modelling of such systems is an research area of interest, with studies investigating startup operation [127, 307], where molten salts may be prone to freezing, the impact of storage system dynamics on the power cycle [4, 213, 146] and control system design [33]. Depending on the application, the fidelity of the storage tank models will vary.

There are two key aspects of a storage tank model that need to be considered: how the temperature distribution of the storage medium within the tank is modelled; and how heat losses to the environment are modelled.

For the temperature distribution of the medium, the simplest model is to assume a perfectly mixed tank with a single homogeneous temperature T. This mixing can

be justified either through the use of mixers within the tank or by natural convection within the tanks providing adequate mixing [298, 124]. In this case, the dynamics of the mass M and specific enthalpy h of the stored medium is described by lumped-parameter mass and energy balances:

$$\frac{\mathrm{d}}{\mathrm{d}t}M = \dot{m}_{\mathrm{in}} - \dot{m}_{\mathrm{out}} \tag{4.10}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(Mh) = \dot{m}_{\mathrm{in}}h_{\mathrm{in}} - \dot{m}_{\mathrm{out}}h + \dot{Q}_{\mathrm{heating}} - \dot{Q}_{\mathrm{loss}} \tag{4.11}$$

where $\dot{m}_{\mathrm{in}}$, $\dot{m}_{\mathrm{out}}$ are the mass flow rates of medium into or out of the tank, h_{in} is the specific enthalpy of the inlet flow, $\dot{Q}_{\mathrm{heating}}$ is the thermal power provided to maintain the tank temperatures, such as by electrical heaters, and $\dot{Q}_{\mathrm{loss}}$ is the heat flow rate from the stored medium to the environment. For modelling that cannot ignore temperature variation within the tank, such as the investigation of tank cool down and mixing effectiveness, computational fluid dynamics can be used [217, 208, 124].

While storage tanks are insulated, the large surface areas and temperature differences mean that there can still be significant heat losses, with the 150MW commercial parabolic-trough plant Andasol-1 reporting temperature drops of up to 5-6°C per day at minimum tank levels [298]. Models can consider multiple heat loss mechanisms, such as heat loss to the gas environment in the tank, heat loss through the bottom of the tank, radiative heat loss and convection within the gas or storage medium at temperature resolutions of 0D lumped-parameter values [298, 33, 240, 198], 1D distributions [240, 48], 2D distributions [217, 208, 237, 260, 48] or 3D distributions [217, 210, 48].

Complex heat loss models generally require data on material emissivities, dynamic and kinematic viscosities and tank geometry, which is not always available. Often, a simplified heat loss model is used based on a single value for an effective heat transfer coefficient that tries to capture all of the processes mentioned above to a reasonable degree of accuracy [127, 146]. Heat loss can even be neglected entirely under the assumption of adequate insulation [307].

4.1.5 Process Controllers

When present, process controllers are a key driver of process dynamics, and must be included in the dynamic modelling of the system. The available dynamic model of the process generally determines how the control system is modelled. Models of industrial processes generally consist of differential equations from the fundamental mass, energy and momentum balances that govern them. These can either be solved numerically, via integration methods, or solved analytically. Analytical solutions are desirable, as they give a clear description of how changing model parameters impact the process, but not all problems can be solved analytically.

The Laplace transform is a powerful tool for finding analytical solutions to linear differential equations. It transforms a piecewise continuous function $f(t)$ from the time domain t to the form $F(s)$, where s is a complex independent variable:

$$F(s) = \mathscr{L}[f(t)] = \int_0^\infty f(t)e^{-st}dt \tag{4.12}$$

As $\mathscr{L}$ and $\mathscr{L}^-1$ are linear operators, finding solutions to differential equations can be reduced to an algebraic rearrangement. Controllers represented in the s domain are known as transfer function models, characterising the dynamics between a independent input variable and a dependent output variable [219]. Despite the limitations of transfer function models to linear systems with only one input and one output variable, their application has been extended to non-linear and multiple-input-multiple-output systems [283, 300].

The transfer function model of a proportional-integral-derivative controller with output $y(t)$ for a measured error $e(t)$ is then:

$$\frac{\mathscr{L}[y(t)-\bar{y}]}{\mathscr{L}[e(t)]} = \frac{Y'(s)}{E(s)} = K_\mathrm{c}\left[1 + \frac{1}{\tau_\mathrm{I}s} + \tau_\mathrm{D}s\right] \tag{4.13}$$

where K_c is the controller gain, τ_I is the integral time and τ_D is the derivative time [219]. The bias value $\bar{y}$ is determined by the initial conditions, such as steady state or fixed initial output.

With advancements in computational power and packages of numerical solution

methods for differential equations [6, 205, 317], complex process models can often be solved directly in the time domain. In this case, controllers will also be represented as functions of time, which for a proportional-integral-derivative controller is:

$$y(t) = \bar{y} + K_{\mathrm{c}} \left(e(t) + \frac{1}{\tau_{\mathrm{I}}} \int_0^t e(t^*)\mathrm{d}t^* + \tau_{\mathrm{D}} \frac{\mathrm{d}e(t)}{\mathrm{d}t} \right) \tag{4.14}$$

Model predictive control is handled externally to the dynamic model in conjunction with real plant measurements, as shown in figure 2.2. At each sampling step, the control actions across the whole control horizon are calculated as a dynamic optimisation problem using the dynamic model. For simulation studies, the plant model can also be used to represent the real plant [251].

4.1.6 Heat Exchangers

A typical steam Rankine cycle would include heat exchangers such as preheaters, evaporators, superheaters, regenerators and condensers. Heat exchangers come in a large variety of designs, such as shell-and-tube, plate and printed circuit. The most important aspect of a particular design is the path through which the fluids flow relative to each other. This flow path configuration will determine the limits of the heat transfer, as shown in figure 4.3. Other notable configurations are single- and multi-pass crossflow [82].

Heat exchangers are governed by the fundamental balances of mass, energy and momentum. These are complex partial differential equations with derivatives in the three spatial dimensions and time [243]. Methods such as computational fluid dynamics can be used to attempt to approximate numerical solutions of the full three-dimensional equations, but are limited by high computational costs and the need for highly detailed geometry information. It can be used for significant cost and time savings in the design of individual heat exchanger components [11], but is not suitable for process simulations.

Lower complexity models can be employed by making assumptions to simplify the governing equations. Generally, there's a trade-off between the computational complexity and required degree of parameterisation, such as geometry information,

against the model accuracy.

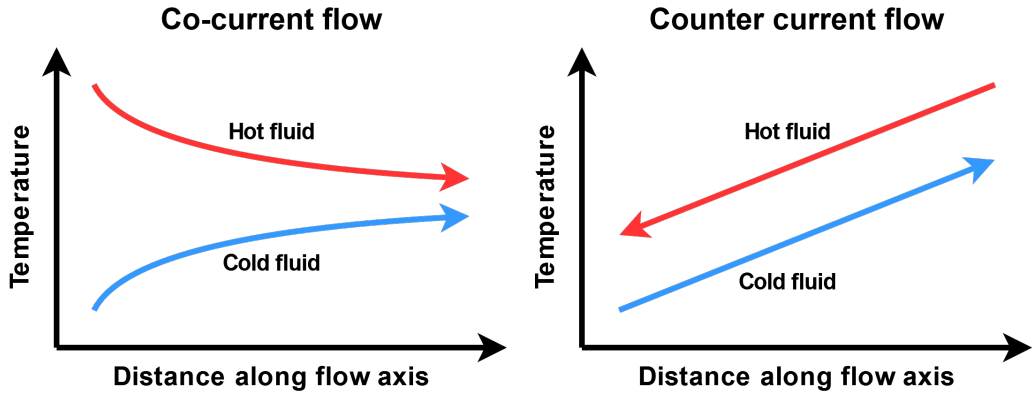

Figure 4.3: Possible flow configurations of a heat exchanger. The counter current configuration allows for a temperature gradient to be maintained along the flow path, with the outlet temperature of at least one fluid approaching the inlet temperature of the other fluid in the limit of infinite heat transfer area. The outlet temperatures of the co-current configuration will approach some intermediate temperature in the same limit.

A common assumption in heat exchanger modelling is that spatial variations in thermofluid properties perpendicular to the flow direction are negligible. This reduces the three-dimensional equations of the mass, energy and momentum balances to more tractable one-dimensional forms:

$$\frac{\partial}{\partial t}(\rho_x) + \frac{\partial}{\partial x}(\rho_x v_x) = 0 \tag{4.15}$$

$$\frac{\partial}{\partial t}(\rho_x h_x - p_x) + \frac{\partial}{\partial x}(h_x \rho_x v_x) = \dot{q}_x \tag{4.16}$$

$$\frac{\partial}{\partial t}(\rho_x v_x) + \frac{\partial}{\partial x}\left(\rho_x v_x^2 + p_x\right) = F_x \tag{4.17}$$

where ρ is the density, v is the velocity, h is the specific enthalpy, p is the pressure, $\dot{q}$ is the heat flow rate into the fluid per unit volume and F is the force per unit volume exerted on the fluid by forces such as friction and gravity. Subscript x denotes that the quantity has been averaged across the area perpendicular to the flow axis x [8]. Kinetic and potential energy are neglected from the energy balance as they are usually a small percentage of the total energy for thermofluid flows in power plants [30]. Further simplifications to the one-dimensional equations can be made

by integrating the thermofluid properties along the flow axis into zero-dimensional control volumes.

Heat transfer in heat exchangers is driven by a temperature difference between each fluid and the adjacent wall separating them. Newton's law of cooling can be used to model the convective heat flux Φ_{q} from a wall w to a fluid f at any point along the flow axis x:

$$\Phi_{\mathrm{q,w}\to\mathrm{f}} = K(T_{\mathrm{w}} - T_{\mathrm{f}}) \tag{4.18}$$

where K is the convective heat transfer coefficient and T is temperature. K can be calculated from complex non-linear correlations that depend on flow geometry, fluid properties, metal properties and flow conditions [82].

An alternative to modelling the temperature of the wall directly is to treat the wall as a thermal resistance R_w in series with two convective heat transfer coefficients K_{f1}, K_{f2} for each fluid with the wall. An overall heat transfer coefficient U between the two fluids f1, f2 can then be calculated as:

$$\frac{1}{U} = \frac{1}{K_{\mathrm{f1}}} + R_{\mathrm{w}} + \frac{1}{K_{\mathrm{f2}}} \tag{4.19}$$

with the heat flux between the two fluids then being given by:

$$\Phi_{\mathrm{q,f1}\to\mathrm{f2}} = U(T_{\mathrm{f1}} - T_{\mathrm{f2}}) \tag{4.20}$$

As heat exchange occurs along the flow direction, the temperature of each fluid will change, in general leading to a non-constant driving temperature difference. This is greatly influenced by the flow configuration, as shown in figure 4.3, and it is necessary for models to capture this non-linearity to accurately model heat transfer.

As presented above, heat transfer within heat exchangers is, in general, governed by complex systems of partial differential algebraic equations with no analytic solution methods. Various models have been developed to simplify these equations into forms that can be solved. Of interest to this work are models suitable for process simulation of thermal power conversion systems based upon the one-dimensional equations 4.15-4.17. This eliminates computationally intensive methods such as

computational fluid dynamics from the discussion. In particular, models are categorised as those using either lumped-parameter methods or distributed-parameter methods.

4.1.6.1 Lumped-Parameter Heat Exchanger Models

Lumped-parameter models do not explicitly model variations in the spatial dimension, allowing for a significant reduction in model complexity. The one-dimensional governing equations are integrated along the length of the exchanger to give lumped-parameter balances.

Typical non-linear temperature profiles within the heat exchanger must still be considered for accurate predictions of heat transfer. Two established methods for modelling heat transfer in lumped-parameter models are the logarithmic mean temperature difference (LMTD) and the effectiveness-number of transfer units (ε-NTU) methods [211].

The logarithmic mean temperature difference $\Delta T_{\mathrm{LMTD}} =$

$$\begin{cases} \dfrac{\Delta T_1 - \Delta T_2}{\ln \Delta T_1 - \ln \Delta T_2}, & \text{if } \Delta T_1 \neq \Delta T_2 \quad (4.21a) \\ \dfrac{\Delta T_1 + \Delta T_2}{2}, & \text{if } \Delta T_1 = \Delta T_2 \quad (4.21b) \end{cases}$$

approximates a mean driving temperature difference, taking into account the non-linear temperature difference along the flow axis in its derivation. ΔT_1, ΔT_2 are the temperature differences between the fluid at each end of the heat exchanger. This is valid for both co-current and counter current heat exchangers [297]. The total heat duty is then $\dot{Q} = UA\Delta T_{\mathrm{LMTD}}$.

ΔT_{LMTD} is derived under the assumptions of adiabatic conditions, negligible axial conduction, non-zero flow rates, non-zero temperature differences, steady-state conditions and constant fluid specific heat capacities and overall heat transfer coefficient along the heat exchangers length.

The LMTD method has been widely applied in the design [139, 181], off-design simulation [9, 139, 86, 162] and control analysis [86, 162, 227] of power conversion systems, as well as a wider scope of processes such as heat exchanger networks

[211, 128]. Some of these are dynamic models [9, 162]. Comparisons against more detailed models [9, 129, 214] or experimental data [162] show the LMTD can be sufficiently accurate for process simulation when computational efficiency is a requirement, despite the assumption of steady state in its derivation.

It has been shown to become less accurate if fluid properties change considerably within the heat exchanger, such as for supercritical CO_2 near the critical point [129]. Corrective factors can be added to address deviations from data [82, 139]. Extensions have also been made to the LMTD to extend it to systems with radiative heat transfer [279], phase change [82] and pressure drop with temperature glide [199]. Studies have noted the difficulty of using the LMTD for dynamic simulations due to indeterminate values of ΔT_{LMTD} for some inputs [297, 9, 278], and modifications have been proposed to improve its reliability in such applications [9, 128, 278].

To compute outlet temperatures and total heat duty from inlet temperatures, the logarithmic mean temperature requires iterative procedures. This can cause difficulty for numerical solvers, such as guesses of negative temperature differences that cause ΔT_{LMTD} to be undefined. Modifications have been made to equations 4.21a-4.21b to aid numerical solvers [9, 128].

The ε-NTU method can be used without solving implicit equations and derives from the same assumptions as the logarithmic mean temperature difference method. It considers the thermal capacity of the cold stream c and hot stream h, identifying the limiting one, and defines the number of transfer units NTU as:

$$\mathrm{NTU} = \frac{UA}{\min(\dot{m}_{\mathrm{c}} c_{\mathrm{p,c}}, \dot{m}_{\mathrm{h}} c_{\mathrm{p,h}})} \tag{4.22}$$

where $\dot{m}$ is the mass flow rate and c_{p} is the specific heat capacity. All the variables are assumed constant along the flow axis.

The method also defines the effectiveness of a heat exchanger ε as the ratio of the actual heat duty $\dot{Q}$ to the theoretical maximum heat duty $\dot{Q}_{\max}$ of an infinitely long counter current heat exchanger, where the outlet temperature of the lowest capacity stream is brought to the inlet temperature of the other stream.

The capacity ratio of the streams C_R is:

$$C_R = \frac{\min(\dot{m}_c c_{p,c}, \dot{m}_h c_{p,h})}{\max(\dot{m}_c c_{p,c}, \dot{m}_h c_{p,h})} \tag{4.23}$$

ε can then be derived for different flow configurations, such as for counter current heat exchangers [82]:

$$\varepsilon_{counter} = \frac{1 - e^{-NTU(1-C_R)}}{1 - C_R e^{-NTU(1-C_R)}} \tag{4.24}$$

The ε-NTU method is used similarly to the LMTD method for design [288, 196, 51, 57, 242], off-design simulation [206, 57] and control analysis [288, 227]. It can show good agreement with experimental data [196, 51] and be used in place of distributed models [242].

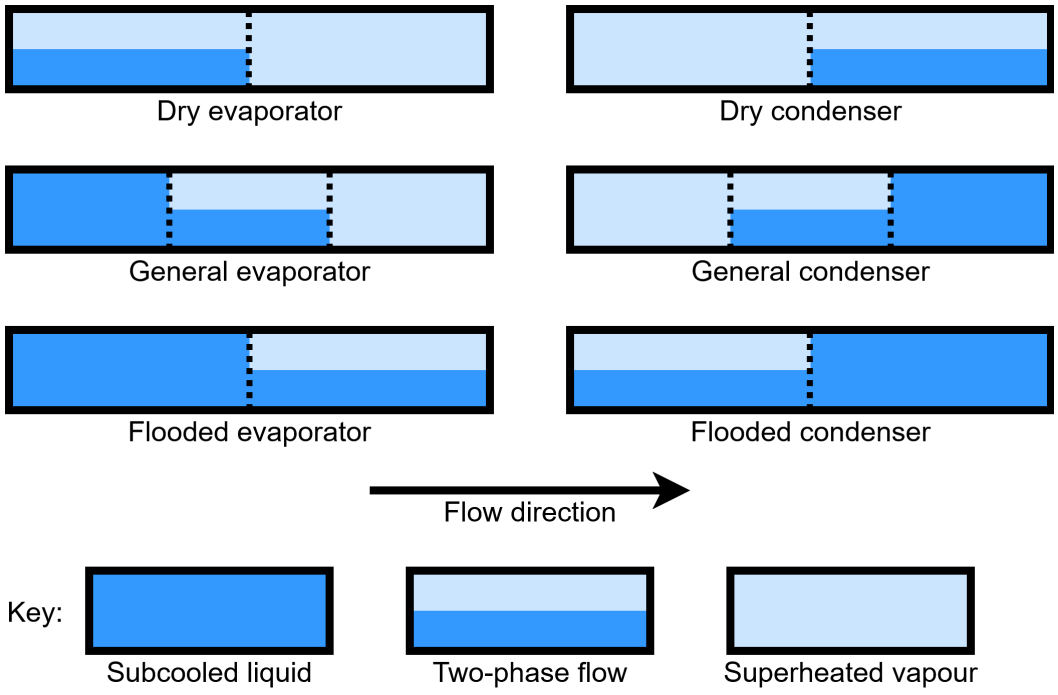

Figure 4.4: Possible phase configurations for evaporators and condensers in moving boundary models [32].

There are heat exchangers, such as once-through steam generators, where a fluid can be either a subcooled liquid, a superheated vapour or a two-phase mixture at different points within the same exchanger. The LMTD and ε-NTU methods will not model these exchangers accurately as properties such as specific heat capacity will vary significantly along the heat exchangers length. Moving boundary models

have been developed to model these exchangers with complex phase configurations. Instead of integrating the whole exchanger into a single control volume, it is divided up into cells with boundaries aligning with phase transitions, as shown in figure 4.4. Each cell is governed by its own lumped-parameter balance equations. As suggested by the name, the boundaries of these cells can move depending on the operating conditions.

Each cell can be modelled using Newton's law of cooling [32, 72, 252], the LMTD method [23, 54] or the ε-NTU method [72, 131]. Moving boundary models have been extended to include pressure drops [201], binary mixtures [131], changing phase configurations [32, 201, 252], and can accurately match results from distributed models [32, 72, 121, 281] and experiments [201, 72] while maintaining the computational benefits of lumped-parameter models. They have been used successfully for model predictive control [85, 292].

The complex derivation of the moving boundary model equations make them hard to extend to new heat exchanger types, and the implementation of switching must be handled carefully or rapid chattering between configurations can occur [200, 188, 209, 121].

4.1.6.2 Distributed-Parameter Heat Exchanger Models

Distributed-parameter models explicitly model the state of the heat exchanger along the flow path of the fluids. This allows the models to handle situations such as highly variable fluid properties, phase change boundaries and, for dynamic models, flow reversal that lumped-parameter models may struggle to handle without special consideration.

The three classical methods for approximating partial differential equations, such as equations 4.15-4.17, into ordinary differential equations are the finite volume, finite difference and finite element methods. The finite volume method is the standard choice for heat exchanger modelling, as it is well suited for conservation equations [191, 30, 204]. For one-dimensional flows, the finite volume method discretises a heat exchanger's length into an arbitrary number of control volumes, or cells, of equal size. Each of these cells is governed by its own set of zero-dimensional balance

equations derived by integrating the partial differential equations along the flow axis.

The staggered grid method is generally used where the momentum balances are integrated half a cell out of sync with the cells of the mass and energy balances. This allows the mass fluxes calculated by the momentum balance to align with the boundaries of the energy and mass cells, as shown in figure 4.5. As for moving boundary models, each cell can use Newton's cooling law [31, 307, 306, 153], the LMTD method [125, 162, 78, 55, 144] or the ε-NTU method [161, 291].

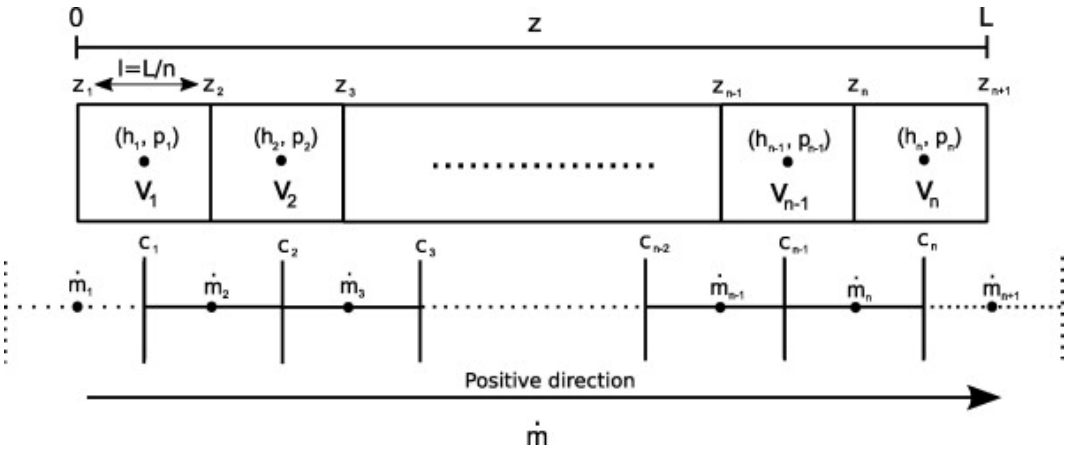

Figure 4.5: The staggered grid method for finite volume models of a one-dimensional heat exchanger. Mass and energy balances are computed in the top cells. The bottom cells, which are offset by half a cell width from the top cells, compute the momentum balances, allowing for mass flow rates to be calculated at the boundaries of the top cells [30]. Image reproduced with permission of the rights holder, Elsevier.

The specific enthalpy of fluid flow between cells must also be defined. The upwind scheme defines the flow's specific enthalpy as the enthalpy of the upstream cell. The central differences scheme defines the specific enthalpy of a cell as being the arithmetic mean of the specific enthalpy of the flow at each of its boundaries. The central differences scheme has been identified as the more accurate scheme of the two [204].

Distributed-parameter models can be used for process design [125, 214, 291, 161, 176, 78, 55, 144], off-design simulation [161, 176, 165, 144] and control analysis [214, 153, 78, 165]. As finite volume models can handle more complex flow conditions, they are also suitable for dynamic simulations of highly transient op-

erations, like start-up and shut-down scenarios [42, 119, 307, 306], where moving boundary models can exhibit problems [252]. The explicit spatial distribution modelling also allows for more accurate modelling of more complex flow configurations, such as cross-flow with baffles [31, 128].

While finite volume models are more general, and can accurately match experimental data [72, 125], they are more computationally intensive to solve. A particular issue of note is chattering, where discontinuities in the models lead to slow simulation or simulation failure. For example, in heat exchangers with phase change of water, discontinuities in the density partial derivatives on the saturated liquid curve in the IAPWS-IF97 formulation can cause flow reversals in some cells. Various methods to improve the robustness of finite volume models against chattering have been proposed [29, 30, 204].

4.2 Process Component Models

This section presents the models used in the power conversion system models presented in chapters 6 and 7. An emphasis is placed on model simplicity for the reasons outlined in section 3.2.

Except for the molten salt storage tanks and heat exchangers, the models presented here are all steady state. This is justified by the fact that the most significant source of inertia in thermal power plants is attributed to heat exchangers [43, 82].

4.2.1 Thermodynamic Properties of Media

In this work, all power conversion system designs are based on steam Rankine cycles. Water undergoes significant temperature and pressure changes in Rankine cycles, existing in the liquid, two-phase and vapour states at different stages of the cycle. The *StandardWater* model from the *Media* package of the Modelica Standard Library [13] is chosen for this work. The model is based on the industry standard IAPWS-97 formulation and is designed specifically for dynamic simulation in Modelica.

The molten salt considered in this work is OMNISTORE MS-600, sold by Globaltherm [178], chosen for its large operating temperature range of 290-600°C that is well-suited for the maximum tokamak coolant temperatures in this work.

The thermodynamic properties of the molten salt are calculated based on constant specific heat capacity c_p =1,495J/(kg K) and density ρ =1,988kg/m^3 given by the manufacturer at 300°C. It has been noted in literature that the temperature dependence of density in molten salts impacts the fill levels of equivalent molten salt masses by $\sim$ 5% between cold and hot tanks [28]. This could mean that the molten salt tank sizes estimated in this work are undersized due to the assumption of higher density in the hot tanks.

Helium is used for the higher temperature tokamak coolant streams, and exists solely as a supercritical fluid. Helium is modelled as an ideal gas for model simplicity, as it is only present at the model boundaries and does not form any closed loops. It has been shown that the ideal gas model can closely approximate the properties of a supercritical fluid at high temperatures relative to its critical temperature, as is the case here [16].

To verify this, the ideal gas helium model can be compared to the Helmholtz free energy formulation of CoolProp [24] across the range of temperatures and pressures that could be expected in the power conversion system. The relative differences in specific enthalpy and density between the models are shown in figure 4.6. The errors are deemed to be low enough to justify the use of the ideal gas model for computational savings.

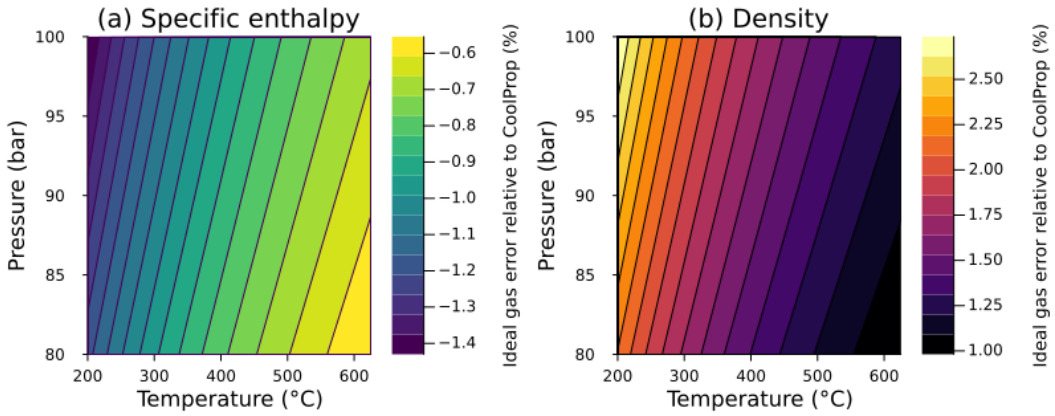

Figure 4.6: Error in (a) specific enthalpy (b) density calculated by the ideal gas helium model relative to the Helmholtz free energy helium model in CoolProp. Temperature and pressures cover the ranges that could be expected from a fusion tokamak.

4.2.2 Steam Turbines

Steam turbines are modelled here as steady-state components. Compared to components like heat exchangers, their response times are relatively short [43, 82], and this assumption should not significantly impact the dynamics of the whole process as start up and shut down are not being considered.

Stodola's Ellipse law, equation 4.6, is used to model the relationship between the off-design mass flow rate $\dot{m}$ and pressure difference across the turbine $\Delta p = p_{\text{in}} - p_{\text{out}}$ based on a design operating point. Control of the steam entering a turbine is modelled as being via partial arc admission with no loss of efficiency [41]:

$$\dot{m} = \theta_{\text{arc}} K \sqrt{\rho_{\text{in}} P_{\text{in}} \left(1 - \frac{{P_{\text{out}}}^2}{{P_{\text{in}}}^2}\right)} \tag{4.25}$$

θ_{arc} is the partial arc admission fraction, and is an optional input to the model from other process components such as controllers.

To uniquely relate the input and output variables, the efficiency of the expansion process must be used. The isentropic efficiency η_{is} measures the deviation of the real expansion from an idealised isentropic expansion. For an expansion from pressure P_{in} to P_{out} with inlet specific enthalpy h_{in}, the specific enthalpy change of the ideal isentropic expansion Δh_{is} is:

$$\Delta h_{\text{is}} = h\left(P_{\text{out}}, s\left(P_{\text{in}}, h_{\text{in}}\right)\right) - h_{\text{in}} \tag{4.26}$$

This uses the fact that two thermodynamic state variables, such as pressure, specific enthalpy and specific entropy, uniquely define any other thermodynamic variable for a pure, single-phase substance at equilibrium.

The specific enthalpy change of the real expansion Δh and the total power output of the turbine $\dot{W}$ are then calculated as:

$$\Delta h = \eta_{\text{is}} \Delta h_{\text{is}} \tag{4.27}$$

$$\dot{W} = -\eta_{\text{mech}} \dot{m} \Delta h \tag{4.28}$$

where η_{mech} is the mechanical efficiency. Turbine manufacturers can provide maps for the turbine efficiencies under different operating conditions. The isentropic and mechanical efficiency are taken to have constant values here. The model will overestimate the turbine efficiency at partial loads due to this assumption and the simplified partial arc admission modelling. This will impact the action of any controllers that act to control the turbine power output, as overestimating the turbine efficiency may mean that less steam flow is predicted than a real turbine would require to meet the control set points. Hence, the optimised designs presented in Case Study II and III may overestimate the design turbine power.

4.2.3 Pumps

Pumps are modelled using the *ControlledPump* model from the *Fluid* package of the Modelica Standard Library [13]. As with the steam turbine in section 4.2.2, it is modelled as a steady-state component with constant isentropic efficiency and mechanical efficiency with the same reasoning. This will lead to overestimates of the thermal efficiency,

The pump model is developed from the assumption of the three similarity conditions for a given pump at different operating conditions, discussed in section 4.1.2.

Lacking specific unit data, the model assumes a quadratic curve relating H and $\dot{V}$ for a specific design speed $\tilde{N}$. The curve is defined by three points in H-$\dot{V}$-N space: the design point $(\tilde{H},\tilde{\dot{V}},\tilde{N})$, the zero-flow point $(2\tilde{H},0,\tilde{N})$, and the zero-head point $(0,1.5\tilde{\dot{V}},\tilde{N})$. Applying equations 4.4 and 4.5 then allows the curve to be extrapolated to different speeds, as illustrated in figure 4.7.

The pump head H can be used to calculate the pressure change across the pump Δp from its definition:

$$H = \frac{\Delta p}{\rho g} \tag{4.29}$$

where ρ is the fluid density, calculated at the pump outlet in the ControlledPump model, and g is the gravitational acceleration, taken to be 9.80665m/s^2.

To calculate the rate of work done by the pump on the fluid $\dot{W}$, the assumption is

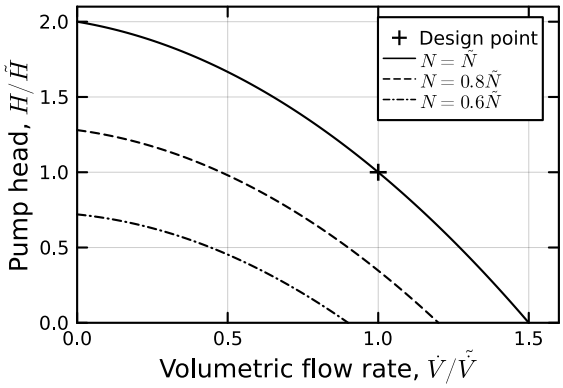

Figure 4.7: Default quadratic pump curve (solid line) used by the Modelica Standard Library *ControlledPump* model. It is based on a design operating point $(\tilde{H}, \tilde{\dot{V}}, \tilde{N})$. Using similarity relations, the curve can be extrapolated to non-design pump speeds (non-solid lines).

made that the fluid is incompressible, which is justifiable for most liquids, including water [73]. To account for inefficiencies, such as energy loss to friction, an efficiency factor η is included:

$$\dot{W} = \frac{\dot{V}\Delta p}{\eta} \tag{4.30}$$

As with the turbine model in section 4.2.2, the efficiency is assumed to have a constant value here.

Assuming the pump's speed can be controlled instantaneously to achieve a desired mass flow rate, the mass flow rate is then an input to the model, either as a fixed parameter or from a controller. This perfect control may neglect an additional source of inertia in that the pumps can only respond so fast to perturbations. This may then mean that the control system designs presented are not feasible for real systems.

4.2.4 Valves

Valves are modelled as steady-state components, used to regulate mass flow rates as part of a control system. As with the steam turbine model, the key relation that needs to be defined is the relationship between the mass flow rate of fluid through the valve $\dot{m}$ and change in fluid pressure across the valve Δp. Equation 4.9 is used here with a linear valve characteristic $C_{\mathrm{V}} = C_{\mathrm{V,max}}\Omega$ for simplicity. Different valve

characteristics would require alternative controller tuning, but should not impact the wider process dynamics if they are tuned to a similar degree of effectiveness.

In the proposed power conversion system models presented in chapters 6 and 7, valves are used either for bypassing units or splitting flows. For the later, a three-way valve configuration is used by combining two regular valves, as shown in figure 4.8.

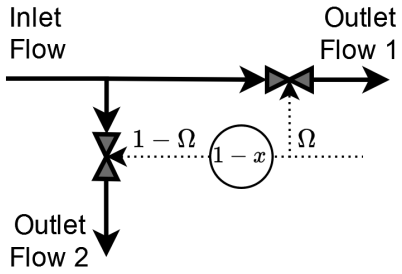

Figure 4.8: The combination of two valves to form a three-way valve. The valve opening input signal Ω is inverted between the two valves, such that when one valve is fully open, the other is fully closed and vice versa.

4.2.5 Molten Salt Storage Tanks

Molten salt tanks are modelled using lumped-parameter mass and energy balances [298, 127]:

$$\frac{\mathrm{d}}{\mathrm{d}t}M_{\mathrm{ms}} = m_{\mathrm{ms,in}} - m_{\mathrm{ms,out}} \tag{4.31}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(M_{\mathrm{ms}}h_{\mathrm{ms}}\right) = m_{\mathrm{ms,in}}h_{\mathrm{ms,in}} - m_{\mathrm{ms,out}}h_{\mathrm{ms}} \tag{4.32}$$

$$- U_{\mathrm{ms,amb}}A_{\mathrm{ms,amb}}\left(T_{\mathrm{ms}} - T_{\mathrm{amb}}\right) \tag{4.33}$$

M_{ms} is the total mass of molten salt stored in the tank, $\dot{m}_{\mathrm{ms,in}}$, $\dot{m}_{\mathrm{ms,out}}$ are the inlet and outlet mass flow rates, h_{ms}, $h_{\mathrm{ms,in}}$ are the specific enthalpies of the stored molten salt and inlet flow, $U_{\mathrm{ms,amb}}$ is the overall heat transfer coefficient between the stored molten salt and the ambient, $A_{\mathrm{ms,amb}}$ is the wetted internal surface area of the tank and T_{ms}, T_{amb} are the temperature of the stored molten salt and the ambient. The wetted internal surface area assumes a vertical cylindrical tank and includes the circular base of the tank.

The assumption of lumped-parameter balances will result in the model having

less inertia than real systems as the molten salt outlet temperature will instantly begin changing due to a perturbation in the inlet temperature. Even with mixing modifications, the time for complete mixing can be of the order of hours [124]. For the processes presented in the case studies, the most significant changes in molten salt temperature is expected at the hot tank inlet during ramp up and ramp down. The control systems should quickly respond to reduce the molten salt inlet mass flow rate, such that the lumped tank temperature does not change significantly and the instantaneous response of the outlet temperature to inlet perturbations will have negligible impact on downstream dynamics.

A MEng project was done in parallel to this work to investigate alternative molten salt tank models. Stratified tank models divide the fluid into multiple layers with their own energy balances and capture the thermal inertia due to mixing better than the lumped-parameter model, but could not account for changing tank levels. A new model was developed based on a stratified tank with a single tank-wide mass balance equation and equal but non-constant layer thickness. The time delay between inlet perturbations and the outlet temperature was significantly impacted by the number of layers, with an order of an hour difference in the response time between a 10 layer model and a 100 layer model. Due to this sensitivity of the model, and the increased computational cost of discretised models, this model was not implemented in this work.

Estimates of the dwell period of a prototype fusion power plant have been given as $\sim$600s, far shorter than the overnight period of solar thermal power plants [21]. Literature suggests that insulation is effective at maintaining molten salt temperatures [298, 197, 240]. A simple heat loss model based on a constant value of $U_{\text{ms,amb}}$ is therefore used here. It is thought that the additional model complexity of detailed heat losses will not significantly impact the dynamics of the tanks, and hence the generated designs.

The assumption of instantaneous mixing of the inlet molten salt with the molten salt in the tank means the outlet specific enthalpy is equal to the lumped-parameter specific enthalpy.

The cover gas environment in a tank is assumed to be regulated such that it is at a fixed pressure. This can be achieved in tanks by having an exchange of gas between the two tanks, as shown in figure 4.2. Hydraulic pressure from the stored molten salt is not considered. As the molten salt thermodynamic property model presented in section 4.2.1 assumes the properties are constant and the pumps control mass flow rate instantaneously, this assumption will only impact the power required for pumping molten salts. None of the case studies presented consider thermal efficiency for optimisation, so this assumption will not impact the generated designs.

The level of stored molten salt must also be constrained. In a real tank, molten salt cannot be pumped out if the level goes below the inlet of the pump, and the there should be a maximum level to avoid overflowing. Hence, an assertion statement is added to the model:

$$\chi_{\text{min}} \leq \frac{V_{\text{ms}}}{V_{\text{tank}}} \leq \chi_{\text{max}} \tag{4.34}$$

If this assertion is false, the simulation will be terminated prematurely.

4.2.6 Proportional-Integral-Derivative Controllers

The only type of controllers used in this work are proportional-integral and proportional-integral-derivative controllers. These are modelled using either the *PI* or *LimPID* model from the *Block* package of the Modelica Standard Library [13]. The *LimPID* model allows the user to select which combination of actions are used, with available combinations of proportional, proportional-integral, proportional-derivative or proportional-integral-derivative. The *PI* model has only the proportional-integral option.

The output of the controller $y(t)$, under non-saturated conditions, is given by the time domain model, equation 4.14. Controllers are initialised from an initial output value in this work, although other initialisation options are available.

To ensure that controller outputs stay within a reasonable range, a controller's output can be assigned upper and lower bounds. If a controller is acting on a valve for example, the valve opening must be between 0% and 100%. However, if there is a sustained error while the controller output is saturated at one of these bounds,

a phenomenon called integral windup occurs. The build-up of the integral term while the controller is unable to exert further action is undesirable, and can lead to overshoot or oscillations in the controlled variable.

The *LimPID* model implements the back-calculation method to counter integral windup [318]. A new error is introduced, e_s, which is the difference between the controller output and the same output when limited to some upper and lower bounds. This error is then fed back into the integral with gain $(N_I \tau_I)^{-1}$. If the output does not saturate at the limits, this feedback does not change the controller behaviour as $e_s = 0$. If the output is saturated and $e_s \neq 0$, this modification will drive the integral output towards a value for which the integral input is zero. By default, the value of N_I is 0.9, and this is not changed in any of the models in this work.

4.2.7 Dynamic Modelling of Heat Exchangers

As heat exchangers connect different process streams, they are the components through which transients in one part of the process can propagate to the rest of the process. In the context of pulsed fusion power plants, the unavoidable source of transients is the tokamak itself. Dynamic models of the heat exchangers are then necessary to investigate how the tokamak dynamics influence the downstream power conversion system and electrical output.

Different modelling methods were attempted to be implemented in OpenModelica during the project. Only models that are able to be used for complete simulations are presented, namely a one-phase finite volume model and two lumped-parameter models. Distinct models for one- and two-phase heat exchangers are developed using the LMTD method. As counter current heat exchangers have the best heat transfer performance and are relatively easy to model with a single flow axis, the models presented in this section assume this configuration. Printed circuit heat exchangers are a promising technology for the development of compact, high performance heat exchangers for novel applications, including counter current configurations [125, 55].

A moving boundary model is unable to be successfully implemented in OpenModelica by the author. Additionally, the finite volume model is not able to be used for heat exchangers with phase change due to chattering creating non-physical outlet

conditions. The mean densities method outlined in Quoilin et al. (2014) [204] is unable to be implemented successfully.

With two dynamic models for one-phase heat exchangers, a comparison can then be made between them in terms of computational efficiency and predicted dynamics. There is no experimental data for this work to verify the models or for comparison of the two-phase lumped parameter model.

The models are based on the following assumptions:

1. Constant coefficient of heat transfer between the fluids and the wall
2. No conduction in the fluids or wall in the flow direction
3. Infinite thermal conductivity in the walls in the radial direction (uniform temperature of a wall cross-section)
4. Constant specific heat capacity of the wall
5. Thermodynamic variables of the fluids vary only in the flow direction
6. Static momentum balance
7. The kinetic and potential energy of the fluids can be neglected
8. No heat exchange with the ambient environment

The Rankine cycles presented in the Case Studies generally operate between two steady state regimes, with short transition periods between if the control system performs well. The use of flow-dependent heat transfer correlations would change the heat fluxes, and hence change the required heat exchanger areas to meet the design operating conditions. Scaling area would impact the thermal inertia of a heat exchanger in Case Studies II and III, but combined with the inertia of other components like the molten salt storage tanks, the overall process dynamics should not differ significantly. The missing dynamics from fluctuating flow conditions during the transitions are expected to be smaller than the driving dynamics of large changes to inlet conditions. Combined with the short duration of the transition

periods relative to the steady state periods, the impact of assumption 1 on the optimisation results is expected to be small.

Assumption 2 is justified by the dominant form of heat transfer being advection along the flow direction. The poor thermal conductivity of fluids means axial conduction should be negligible in the fluids [174]. For macro-channels, the ratio of axial conduction in the walls to convective heat transfer in the flow is small [167].

For sufficiently thin pipes, the wall temperature differences in the radial axis are not expected to be large. When considering high pressure systems like thermal power plants, some components will require thicker walls however. Assumption 3 will lead to an overestimation of the heat flux between the fluids and the wall relative to a finite radial thermal conductivity. This will have an impact on the required heat exchanger area, similar to assumption 1 as outlined above.

The specific heat capacity of steel has an approximately linear temperature dependence of 0.505J/(kg K^2) across the temperature range of 0-600°C relevant to the Case Studies [87]. The heat exchangers for the tokamak coolants are expected to experience the largest temperature changes in the processes, so this assumption may underestimate their thermal inertia. The high temperature heat exchanger in particular has the largest wall mass. The impact of its dynamics on the Rankine cycle is buffered by mixing the molten salt into the hot tank, so the additional inertia of this heat exchanger would likely not impact the rest of the cycle significantly. Hence, assumption 4 is made for simplicity.

From assumption 5, thermodynamic variables in the fluids are represented by cross-sectional averages. Spatial variations above or below these means would be expected to approximately cancel out in more detailed simulations. Additionally, as phase change is limited only to occur in the large internal volumes of evaporators and condensers, complex two-phase flow configurations should not occur in the process, which may require more complex models to accurately represent [8].

Assumption 6 is justified by the low time-constant for pressure propagation relative to mass and energy dynamics [72]. Kinetic and potential energy are generally dominated by enthalpy energy in turbulent flows, which heat exchangers design for

to improve heat transfer [30, 106], so the impact of assumption 7 should be small. Perfect insulation is impossible, so assumption 8 will mean that the model overestimates the thermal efficiency of each process. However, it is reasonable to assume that adequate investment will be made to insulate high temperature components as efficiency is a key consideration for fusion energy to generate net electricity.

4.2.7.1 Finite Volume Model

In addition to the assumptions presented in section 4.2.7, the finite volume heat exchanger model is based on the additional assumptions:

1. Thermodynamic variables are constant within cells
2. Flow reversal does not occur
3. Linear pressure drop along the heat exchanger's length under steady flow conditions

The one-dimensional conservation equations 4.15-4.16 are discretised into an arbitrary number of cells N that each has its own lumped-parameter conservation equations for both fluids and the wall, as shown in figure 4.5. The impact of assumption 1 is determined by N, with larger values being assumed to provide more accurate results. The mass and energy conservation equations for either the hot or cold fluid $\mathrm{f} \in \{\mathrm{hf}, \mathrm{cf}\}$ in cell i are:

$$\frac{V_\mathrm{f}}{N}\frac{\mathrm{d}}{\mathrm{d}t}\rho_{i,\mathrm{f}} = \dot{m}_{i,\mathrm{f,in}} - \dot{m}_{i,\mathrm{f,out}} \tag{4.35}$$

$$\frac{V_\mathrm{f}}{N}\frac{\mathrm{d}}{\mathrm{d}t}\left(u_{i,\mathrm{f}}\rho_{i,\mathrm{f}}\right) = \dot{m}_{i,\mathrm{f,in}}h_{i,\mathrm{f,in}} - \dot{m}_{i,\mathrm{f,out}}h_{i,\mathrm{f,out}} - \dot{Q}_{i,\mathrm{f}} \tag{4.36}$$

where V_f is the internal volume of the heat exchanger for fluid f, $\rho_{i,\mathrm{f}}$, $u_{i,\mathrm{f}}$ are the averaged density and specific internal energy of fluid f in cell i, $\dot{m}_{i,\mathrm{f,in}}$, $\dot{m}_{i,\mathrm{f,out}}$ are the inlet and outlet mass flow rates of fluid f for cell i, $h_{i,\mathrm{f,in}}$, $h_{i,\mathrm{f,out}}$ are the specific enthalpies of the inlet and outlet flows of fluid f for cell i and $\dot{Q}_{i,\mathrm{f}}$ is the heat flow rate from fluid f to the walls in cell i. Using the upwind scheme, the specific enthalpy and pressure of fluid flows between cells is equal to the value in the upstream cell [30, 82].

To avoid the additional model complexity and parameterisation required for detailed friction factor correlations, a simple pressure drop model is considered based on design values for the pressure drop $\Delta p_{\mathrm{des,f}}$, mass flow rate $\dot{m}_{\mathrm{des,f}}$ and density $\rho_{\mathrm{des,f}}$ [82]. From the assumption of linear pressure drop, the pressure drop in fluid f across cell i is then:

$$p_{i,\mathrm{f,in}} - p_{i,\mathrm{f,out}} = \frac{\Delta p_{\mathrm{des,f}} \rho_{\mathrm{des,f}}}{N \dot{m}_{\mathrm{des,f}}^2} \frac{\dot{m}_{i,\mathrm{f,in}}^2}{\rho_{i,\mathrm{f,in}}} \tag{4.37}$$

where $p_{i,\mathrm{f,in}}$, $p_{i,\mathrm{f,out}}$ are the pressures of fluid f flowing into and out of cell i, and $\rho_{i,\mathrm{f,in}}$ is the density of fluid f flowing into cell i. This assumption will impact the accuracy of the model's prediction of thermal efficiency, but is not expected to influence the dynamics significantly. Pressure drops are discussed further in section 4.1.3.

As the heat exchanger is counter current, the cold fluid is defined to be flowing from cell 1 to N, and the hot fluid from cell N to 1. The inlet and outlet flow variables for the cold fluid in cell i will be from cell $i-1$ to i and from i to $i+1$ respectively. The inlet and outlet flow variables for the hot fluid in cell i will be from cell $i+1$ to i and from i to $i-1$ respectively.

The energy balance of the wall section in cell i is:

$$c_{\mathrm{w}} \frac{M_{\mathrm{w}}}{N} \frac{\mathrm{d}}{\mathrm{d}t} T_{i,\mathrm{w}} = \dot{Q}_{i,\mathrm{cf}} + \dot{Q}_{i,\mathrm{hf}} \tag{4.38}$$

where c_{w} is the constant specific heat capacity of the wall, M_{w} is the total mass of the heat exchanger walls and $T_{i,\mathrm{w}}$ is the averaged temperature of the wall in cell i. Heat flows are calculated based on Newton's cooling law [82] for a cell's averaged wall and fluid temperatures:

$$\dot{Q}_{i,\mathrm{f}} = \frac{A_{\mathrm{f}}}{N} U_{\mathrm{f}} \left(T_{i,\mathrm{f}} - T_{i,\mathrm{w}} \right) \tag{4.39}$$

where A_{f}, U_{f} are the total heat exchange area and heat transfer coefficient between the wall and fluid f, and $T_{i,\mathrm{f}}$ is the averaged temperature of fluid f in cell i.

4.2.7.2 One-Phase Lumped-Parameter Model

The one-phase lumped-parameter heat exchanger model derives from the $N = 1$ finite volume model governing equations 4.35-4.37, making the following additional assumptions:

1. Both fluids are single phase
2. The density time derivatives are negligible for single phase fluids
3. The pressure time derivatives are negligible
4. The temperature profile of the heat exchanger walls is linear along the flow axis
5. The mass flow rates of both fluids are non-zero
6. Density variation can be neglected for the pressure drop

The model diagram of a single-phase heat exchanger is shown in figure 4.9. Discussion of the impact of these additional assumptions, relative to the finite volume model, is made in section 4.2.7.3.

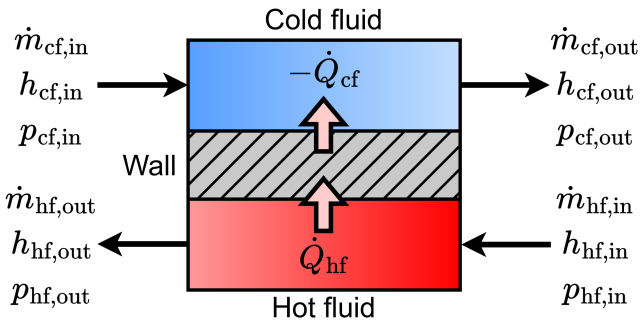

Figure 4.9: Structure of the one-phase lumped-parameter heat exchanger model. Heat exchange between the fluids occurs via the heat exchanger walls.

The conservation equations of the single control volume for each fluid then become [146, 127]:

$$\dot{m}_{\mathrm{f,in}} = \dot{m}_{\mathrm{f,out}} \tag{4.40}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{f}} = \frac{\dot{m}_{\mathrm{f,in}} \left(h_{\mathrm{f,in}} - h_{\mathrm{f,out}} \right) - \dot{Q}_{\mathrm{f}}}{V_{\mathrm{f}} \rho_{\mathrm{f}}} \tag{4.41}$$

$$p_{\mathrm{f,in}} - p_{\mathrm{f,out}} = \frac{\Delta p_{\mathrm{f,des}}}{\dot{m}^2_{\mathrm{f,des}}} \dot{m}^2_{\mathrm{f,in}} \tag{4.42}$$

Using the upwind scheme, the specific enthalpy at the outlet $h_{\mathrm{f,out}}$ is equal to the specific enthalpy of the control volume h_{f}. Unlike the finite volume model however, the pressure drop is considered to be lumped after the heat exchange. Hence, calculation of state variables such as ρ_{f} and h_{f} are performed using the inlet pressure. A central difference method, where the control volume pressure is the arithmetic mean of the inlet and outlet pressures, would be more physically accurate. However, under the assumption of single phase fluids, the thermodynamic properties should not change significantly for typical pressure drops. This is supported by good agreement of a moving boundary model using a lumped pressure drop at the outlet to experimental data [200] and the comparison in section 4.2.7.3.

The modified LMTD method proposed by Altés Buch et al. (2015) [9] is adapted here, with the thermal masses of the fluids being kept separate from the wall and modelled using equation 4.41. Crossings are a particular issue during initialisation and the transition between tokamak pulses and dwells where temperatures can be highly transient. The unmodified logarithmic mean temperature difference is therefore unsuitable here, as simulations would likely fail shortly after initialisation.

The assumption of a constant wall temperature from equation 4.38 can no longer be justified here. Instead, a linear wall temperature profile with axial distance is assumed. The profile is then defined by the mean wall temperature $\bar{T}_{\mathrm{w}} = (T_{\mathrm{w},1} + T_{\mathrm{w},2})/2$ and the temperature difference between each end of the heat exchanger $\Delta T_{\mathrm{w}} = T_{\mathrm{w},2} - T_{\mathrm{w},1}$. These are governed by the following equations:

$$c_{\mathrm{w}} M_{\mathrm{w}} \frac{\mathrm{d}}{\mathrm{d}t} \bar{T}_{\mathrm{w}} = \dot{Q}_{\mathrm{hf}} + \dot{Q}_{\mathrm{cf}} \tag{4.43}$$

$$\begin{aligned} c_{\mathrm{w}} M_{\mathrm{w}} \frac{\mathrm{d}}{\mathrm{d}t} \Delta T_{\mathrm{w}} = A_{\mathrm{hf}} U_{\mathrm{hf}} \left(T_{\mathrm{hf,in}} - T_{\mathrm{hf,out}} - \Delta T_{\mathrm{w}} \right) + \\ A_{\mathrm{cf}} U_{\mathrm{cf}} \left(T_{\mathrm{cf,out}} - T_{\mathrm{cf,in}} - \Delta T_{\mathrm{w}} \right) \end{aligned} \tag{4.44}$$

where c_{w}, M_{w} are the specific heat capacity and total mass of the walls.

Heat flow rates are calculated for each fluid f with the walls only as $\dot{Q}_\mathrm{f} = U_\mathrm{f} A_\mathrm{f} \Delta T_{\mathrm{RLMTD,f}}$. The robust logarithmic mean temperature difference $\Delta T_\mathrm{RLMTD} =$

$$\begin{cases} \dfrac{\Delta T_1 - \Delta T_2}{\ln \Delta T_1 - \ln \Delta T_2}, & \text{if } \Delta T_1 \geq \varepsilon,\ \Delta T_2 \geq \varepsilon \text{ and } \Delta T_1 \neq \Delta T_2 \quad (4.45a) \\ \dfrac{\Delta T_1 + \Delta T_2}{2}, & \text{if } \Delta T_1 \geq \varepsilon, \Delta T_2 \geq \varepsilon \text{ and } \Delta T_1 = \Delta T_2 \quad (4.45b) \\ \dfrac{\Delta T_1 - \varepsilon}{\ln \frac{\Delta T_1}{\varepsilon} \left(1 - \xi \left(\Delta T_2 - \varepsilon\right)\right)}, & \text{if } \Delta T_1 > \varepsilon \text{ and } \Delta T_2 < \varepsilon \quad (4.45c) \\ \dfrac{\Delta T_2 - \varepsilon}{\ln \frac{\Delta T_2}{\varepsilon} \left(1 - \xi \left(\Delta T_1 - \varepsilon\right)\right)}, & \text{if } \Delta T_1 < \varepsilon \text{ and } \Delta T_2 > \varepsilon \quad (4.45d) \\ \dfrac{\varepsilon}{\left(1 - \xi \left(\Delta T_1 - \varepsilon\right)\right)\left(1 - \xi \left(\Delta T_2 - \varepsilon\right)\right)}, & \text{if } \Delta T_1 \leq \varepsilon \text{ and } \Delta T_2 \leq \varepsilon \quad (4.45e) \end{cases}$$

is adopted here [9]. This equation is C^0-continuous for all possible inlet and outlet temperature differences ΔT_1, ΔT_2, preventing simulation crashes if the iterative solver tries a negative temperature difference and aiding solver convergence.

The robust logarithmic mean temperature difference uses two parameters, ε and ξ. ε is the threshold value for ΔT_1, ΔT_2 that switches from the logarithmic mean temperature difference in equations 4.45a-4.45b to decreasing, non-physical functions in equations 4.45c-4.45e. ξ determines the rate at which equations 4.45c-4.45e decrease towards zero. Note that $\Delta T_\mathrm{RLMTD}(\Delta T_1, \Delta T_2) \geq 0$ for all $\Delta T_1 \in \mathbb{R}$, $\Delta T_2 \in \mathbb{R}$. Hence, there will always be a small leakage heat flow from the "hot" fluid to the "cold" fluid even if the "cold" fluid is actually hotter.

As only equations 4.45a-4.45b have a physical derivation, ε should be chosen such that the dynamic models are operating in their region for the bulk of a simulation. Low values of ε are reported by Altés Buch et al. (2015) [9] to result in "slow and non-robust simulation". Similarly, higher values of ξ results in smaller leakage heat flows in the non-physical equations, but steep variations in the function can cause simulation failures.

The performance of the robust logarithmic mean temperature difference is discussed more in section 7.3.1 of Case Study II, where a pulsed heat source is modelled.

4.2.7.3 Comparison of the Finite Volume and Lumped-Parameter Models

Parameter	Lumped-parameter model	Finite volume model
Wall mass, M_w	100kg	
Wall specific heat capacity, c_w	466J/(kg K)	
Heat transfer area for fluid f, A_f	15m^2	
Heat transfer coefficient for fluid f, U_f	1,000W/(m^2 K)	
Internal volume for fluid f, V_f	0.037m^3	
Design mass flow rate for hot fluid, $\dot{m}_{hf}$	1kg/s	
Design mass flow rate for cold fluid, $\dot{m}_{cf}$	2kg/s	
Design pressure drop for hot fluid, $\Delta p_{des,hf}$	0.2bar	
Design pressure drop for cold fluid, $\Delta p_{des,cf}$	0.1bar	
Number of cells, N	-	100

Table 4.1: Parameters for the two comparisons of the lumped-parameter and finite volume dynamic heat exchanger models, based on [9].

To assess the suitability of the one-phase lumped parameter model, it is compared to the finite volume model in two simulations using water for both fluids. In one simulation, both fluids are liquid phase water; in the other, both fluids are vapour phase water. For both simulations, the heat exchanger parameters are given in table 4.1.

For the liquid phase comparison, the cold water has an inlet temperature of 10°C at 2kg/s and an outlet pressure of 9.9bar. The hot water has an outlet pressure of 19.8bar, and its inlet undergoes two step changes of 40°C to 90°C at $t = 500$s and 1kg/s to 0.5kg/s at $t = 1{,}000$s. The design densities $\rho_{des,f}$ of the hot and cold fluid are 993kg/m^3 and 1,000kg/m^3 respectively. The outlet temperatures of both fluids predicted by each model are shown in figure 4.10.

The cold fluid outlet temperature has excellent agreement between the models,

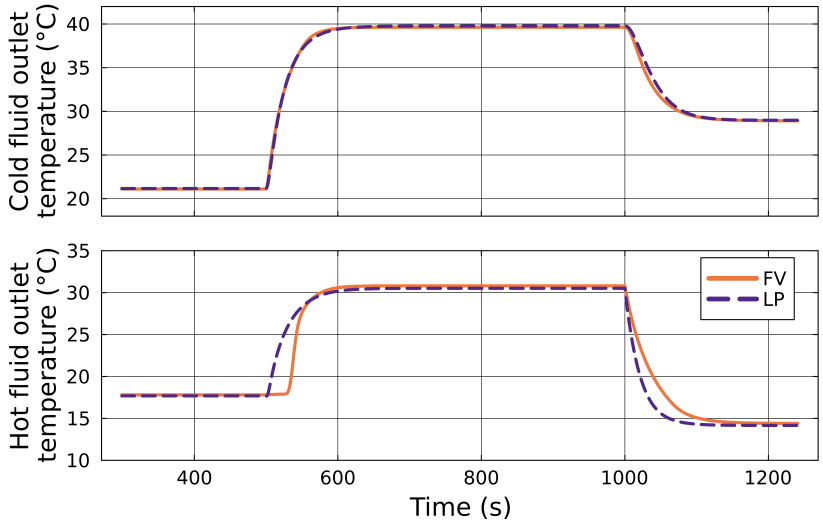

Figure 4.10: Comparison of the outlet temperatures predicted by the finite volume (FV) and lumped-parameter (LP) dynamic heat exchanger models with both fluids being liquid water.

with a maximum difference of < 1.5K during the transients and < 0.15K at steady state. There is a greater discrepancy with the hot fluid outlet temperature, notably that the lumped-parameter model does not capture the delayed response of the outlet relative to the inlet. This is expected from equation 4.41 as the derivative of the outlet enthalpy is a function of the inlet enthalpy and mass flow rate. Referring to table 4.2 for $N = 100$, the finite volume simulation takes $\sim 133\times$ longer than for the lumped-parameter, but this is reduced to $\sim 42\times$ longer for $N = 50$ with similar simulation results.

For the vapour phase comparison, the cold steam has an inlet temperature of 250°C at 2kg/s and an outlet pressure of 9.9bar. The hot steam has an outlet pressure of 19.8bar, and its inlet undergoes two step changes of 280°C to 330°C at $t = 500$s and 1kg/s to 0.5kg/s at $t = 1{,}000$s. The design densities $\rho_{\text{des,f}}$ of the hot and cold fluid are 8.33kg/m^3 and 4.30kg/m^3 respectively. The outlet temperatures of both fluids predicted by each model are shown in figure 4.11.

While it is expected that the vapour phase may stress the assumption of the model that the density time derivative is zero, the agreement between the models is largely better than for the liquid phase. The lower density of vapour results in less inertia on the hot fluid side, so the models are in closer agreement there, although the

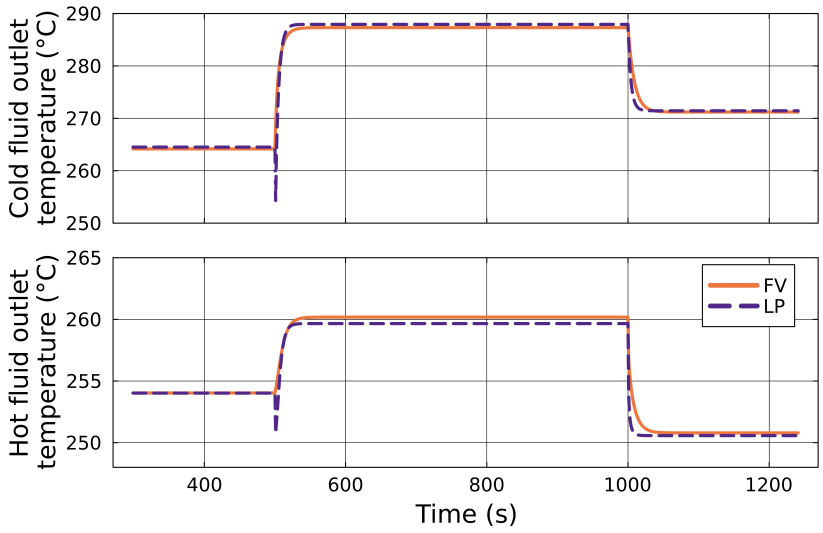

Figure 4.11: Comparison of the outlet temperatures predicted by the finite volume (FV) and lumped-parameter (LP) dynamic heat exchanger models with both fluids being steam.

finite volume model still shows greater inertia after the mass flow rate step change. The maximum difference in outlet temperatures at steady state is < 0.65K. The instantaneous response of the lumped-parameter model does result in sharp drops in outlet temperature immediately after the temperature step change. This sharp change is exasperated by the perturbation being a step, which wouldn't be expected in real systems. Referring again to table 4.2 for $N = 100$, the finite volume simulation takes $\sim 172\times$ longer than the lumped parameter simulation, or $\sim 45\times$ longer for $N = 50$ with similar simulation results.

From these comparisons, it is justified that the lumped-parameter model be employed in place of the finite volume model for one-phase heat exchangers. The models agree well on transient behaviour and better at steady state, even for the vapour phase where the assumptions are less justifiable. The lumped-parameter model produces these simulation results for a greatly reduced computational cost.

Simulation time	Lumped-parameter model	Finite volume model, $N = 50$	Finite volume model, $N = 100$
Liquid phase	0.205s	8.58s	27.2s
Vapour phase	0.233s	10.4s	40.0s

Table 4.2: Simulation times of the heat exchanger model tests. N is the number of cells in the finite volume model.

4.2.7.4 Two-Phase Lumped-Parameter Model

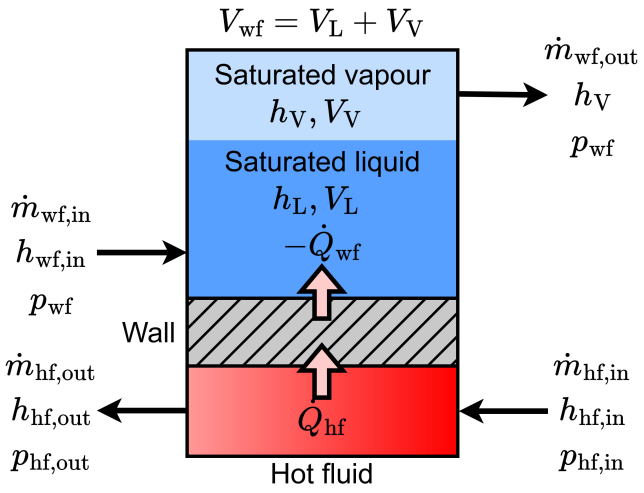

Figure 4.12: Structure of the two-phase lumped-parameter heat exchanger model in the configuration of an evaporator. Saturated liquid and vapour phases coexist in the same volume of the heat exchanger.

When considering heat exchangers where phase change occurs, such as evaporators and condensers, the first assumption in section 4.2.7.2 can no longer be justified. For water at standard pressure, the density of saturated liquid is 1,624 times greater than that of the saturated vapour. Depending on the heat transfer within a heat exchanger, the mass of the liquid and vapour phases in the heat exchanger may not be constant, leading to mass hold-up or discharge.

To model this, we take the one-phase lumped-parameter model and replace the first assumption “Both fluids are single phase” with the following additional assumptions:

1. Only one fluid (the working fluid) is undergoing phase change
2. The working fluid in the heat exchanger is at saturated conditions
3. Both phases of the working fluid are at the same temperature and pressure
4. Pressure drop is neglected on the working fluid side
5. For an evaporator, the outlet is saturated vapour

6. For a condenser, the outlet is saturated liquid

The fluid not undergoing phase change, the heat exchanger walls and the heat transfer are modelled as in section 4.2.7.2, with the working fluid temperatures used to calculate heat transfer being the saturated temperature. The model diagram of an evaporator with coexisting saturated phases is shown in figure 4.12.

Assumption 3 means the model will not capture shrink-swell evaporator dynamics, which can complicate level control [82]. A real boiler would have to control for this however, so the impacts of such phenomena on the process dynamics of a real system should not be significant.

Pressure losses in a real system would likely lead to partial condensation of steam leaving the evaporator. As this steam immediately then passes through a superheater, neglecting this partial condensation should only result in a overestimation of the cycle efficiency by the model, with minimal impact on the wider process dynamics.

The conservation equations for the working fluid wf are:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(V_{\mathrm{L}}\rho_{\mathrm{L}} + V_{\mathrm{V}}\rho_{\mathrm{V}}\right) = \dot{m}_{\mathrm{wf,in}} - \dot{m}_{\mathrm{wf,out}} \tag{4.46}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(V_{\mathrm{L}}\rho_{\mathrm{L}}h_{\mathrm{L}} + V_{\mathrm{V}}\rho_{\mathrm{V}}h_{\mathrm{V}}\right) = \dot{m}_{\mathrm{wf,in}}h_{\mathrm{wf,in}} - \dot{m}_{\mathrm{wf,out}}h_{\mathrm{wf,out}} - \dot{Q}_{\mathrm{wf}} \tag{4.47}$$

where L denotes properties of the saturated liquid phase and V denotes properties of the saturated vapour phase [146, 127, 121]. For evaporators, $h_{\mathrm{wf,out}} = h_{\mathrm{V}}$, and for condensers, $h_{\mathrm{wf,out}} = h_{\mathrm{L}}$.

As the two saturated phases occupy the same volume in the heat exchanger, they are constrained by the heat exchanger's internal volume on the working fluid side V_{wf} such that $V_{\mathrm{wf}} = V_{\mathrm{L}} + V_{\mathrm{V}}$. This model is most suitable to represent kettle boilers and shell-and-tube condensers with a hotwell [121]. For these heat exchangers, the secondary fluid tubes should only be in contact with one of the saturated phases. This avoids issues of simultaneous heat transfer between the wall and both phases, which in a real system would have significantly different coefficients of heat transfer. The model cannot capture these effects due to assumption 1 in section 4.2.7.

4.2.8 Steady State Modelling of Heat Exchangers

Steady state models are developed for model parameterisations in chapter 7. These are based on the lumped-parameter dynamic models presented in sections 4.2.7.2 and 4.2.7.4. Simplifications are made so that the models can be solved sequentially using the turbine inlet as a tear stream, where the conditions are most often specified in a design [168].

The steady state models neglect the modelling of wall temperatures. Instead, an overall heat transfer coefficient U between the two fluids is used. This can be related to the heat transfer coefficient of each fluid $\mathrm{f} \in \{\mathrm{hf}, \mathrm{cf}\}$ with the wall U_f from the dynamic models using equation 4.19, neglecting thermal resistance from the wall:

$$\frac{1}{U} = \frac{1}{U_\mathrm{cf}} + \frac{1}{U_\mathrm{hf}} \tag{4.48}$$

All time derivatives are considered to be zero, giving the follow mass and energy balances:

$$\dot{m}_\mathrm{f,in} = \dot{m}_\mathrm{f,out} \tag{4.49}$$

$$\dot{m}_\mathrm{f,in}\left(h_\mathrm{f,out} - h_\mathrm{f,in}\right) = \begin{cases} \dot{Q}, & \text{if f} = \text{cf} \\ -\dot{Q}, & \text{if f} = \text{hf} \end{cases} \tag{4.50}$$

where Q is calculated using the logarithmic mean temperature difference:

$$\dot{Q} = \begin{cases} UA\frac{\Delta T_1 - \Delta T_2}{\ln \Delta T_1 - \ln \Delta T_2}, & \text{if } \Delta T_1 \neq \Delta T_2 \\ UA\Delta T_1, & \text{otherwise} \end{cases} \tag{4.51}$$

with $\Delta T_1 = T_\mathrm{hf,out} - T_\mathrm{cf,in}$ and $\Delta T_2 = T_\mathrm{hf,in} - T_\mathrm{cf,out}$. For evaporators and condensers, heat transfer is calculated based on the saturated temperature of the working fluid. A fixed-percentage pressure drop is assumed across the one-phase sides of heat exchangers [304].

4.3 Summary

Modelling approaches for the off-design modelling of various power conversion system components are discussed. Thermodynamic models are generally developed

by fitting semi-empirical equations to experimental data. Turbomachinery and pressure drops require empirical relations to model, barring complex computational fluid dynamics simulations. Sensible thermal energy storage can be modelled at different resolutions depending on the context; computational fluid dynamics may be used when studying mixing, while lumped-parameter models can be used to study heat losses to the environment. Process controllers can either be modelled in the frequency- or time- domain, depending on the process model.

As a fluid of great scientific interest, water has a highly accurate property model suited for computational use. Helium and molten salt are modelled more simply as an ideal gas and as having constant properties respectively. The former is verified to be suitably accurate against more accurate equations of state. Lumped-parameter models are presented for the process components. Empirical models are parametrised based on a single design point as more detailed data is unavailable at the concept design stage.

A wide spectrum of heat exchanger models have been presented in literature at different levels of accuracy and computational complexity. For process simulation, either lumped-parameter or distributed-parameter models are typical.

Dynamic models are developed based on the finite volume method, for one phase flows only, and the logarithmic mean temperature difference method, either for one-phase flows or phase change in one fluid only. Comparison of the finite volume model with the one-phase lumped-parameter model shows good agreement in both the liquid and vapour water phases, justifying the use of the lumped-parameter models in the remainder of this work. A simplified steady-state model of the lumped-parameter heat exchangers is also presented for the parameterisation of the dynamic process models in chapter 7.

Chapter 5

Optimisation for Thermal Power Conversion System Design

Publications based on this chapter:
"*Optimization of the Power Conversion System for a Pulsed Fusion Power Plant with Multiple Heat Sources using a Dynamic Process Model*" [276]
"*A dynamic model of a power conversion system with indirect thermal energy storage for a pulsed fusion tokamak for use in design optimisation*" [275]

Optimisation is the process of searching for a set of variable values that maximise or minimise the value of some objective functions, possibly subject to constraints on the values that these variables can take. It is a powerful tool across the engineering sciences, with applications to process design, operation, model calibration and more [229].

The general form of an optimisation problem is as follows:

$$\begin{aligned} &\min_{\boldsymbol{x}\in X} \boldsymbol{f}(\boldsymbol{x}) \\ &\text{s.t. } \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0} \\ &\quad \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \end{aligned} \tag{5.1}$$

$\boldsymbol{x}$ is a vector of D design (or decision) variables, $X \subseteq \mathbb{R}^D$ is the design (or decision) space, $\boldsymbol{f} : X \to \mathbb{R}^k$ is the objective function, D is the dimension of the optimisation

problem, k is the number of objectives and $\boldsymbol{g}(\boldsymbol{x})$, $\boldsymbol{h}(\boldsymbol{x})$ are the inequality and equality constraints, respectively.

Optimisation problems can be classified by various features, as outlined in figure 5.1, although the list is by no means exhaustive. Of note to this work is the distinction of whether derivatives of $\boldsymbol{f}(\boldsymbol{x})$ are available. It may be the case that the evaluation of $\boldsymbol{f}(\boldsymbol{x})$ involves simulations where no explicit form of $\boldsymbol{f}(\boldsymbol{x})$ can be found, and estimating the gradient via finite difference methods is computationally infeasible. These are known as black-box optimisation problems.

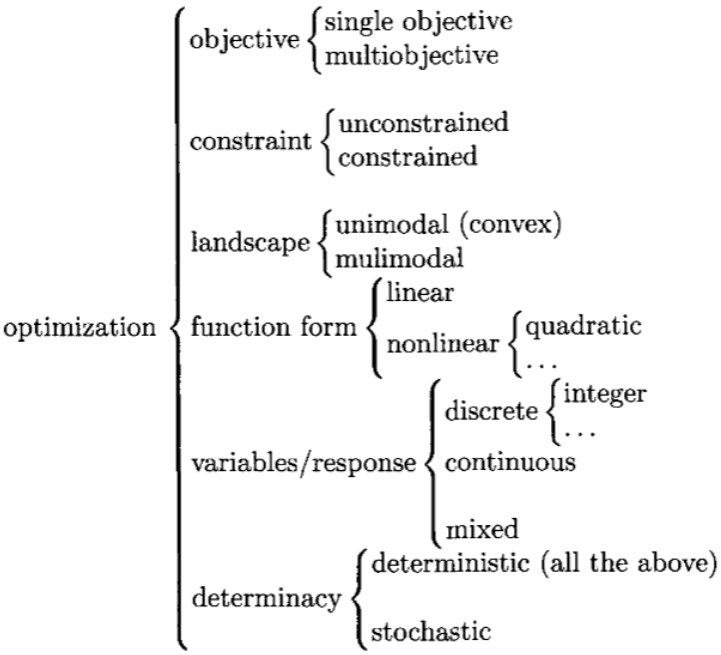

Figure 5.1: Possible classifications of optimisation problems [289]. Image reproduced with permission of the rights holder, Wiley.

Referring to figure 5.2, the range of available optimisation algorithms is large, with meta-heuristic algorithms alone having such diversity to motivate a taxonomy for its sub-category of biologically-inspired algorithms [222]. The selection of an algorithm is informed by the classification of the optimisation problem, and the proper pairing of algorithm and problem can improve the results of the optimisation. For example, the gradient descent algorithm can only be used for unconstrained problems, and may fail to find a global optimum of multimodal problems, depending on its starting point.

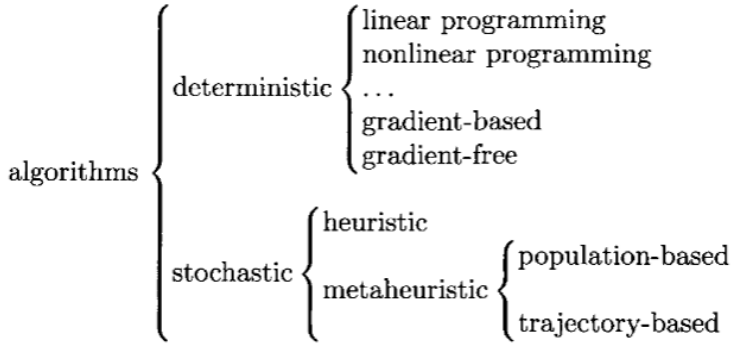

Figure 5.2: Possible classifications of optimisation algorithms [289]. Image reproduced with permission of the rights holder, Wiley.

Define S as the feasible set of X considering the constraints $\boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}$ and $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$. For single-objective optimisation problems where $k = 1$, there will be at least one globally optimal design vector $\boldsymbol{x}^* \in S$ for which:

$$\boldsymbol{f}(\boldsymbol{x}^*) \leq \boldsymbol{f}(\boldsymbol{x}), \forall \boldsymbol{x} \in S \tag{5.2}$$

Single-objective optimisation aims to find design(s) as close as possible to a global optimum $\boldsymbol{x}^*$.

For multi-objective optimisation problems with $k > 1$, there is no guarantee for the existence of a design vector that satisfies 5.2. Instead, the concept of dominance between design vectors is used. A design vector $\boldsymbol{x}_1 \in S$ is said to strictly dominate a design vector $\boldsymbol{x}_2 \in S$ if:

$$f_i(\boldsymbol{x}_1) < f_i(\boldsymbol{x}_2), \forall i \in \{1, ..., k\} \tag{5.3}$$

A design vector $\boldsymbol{x}_1 \in S$ is said to instead Pareto-dominate a design vector $\boldsymbol{x}_2 \in S$ if:

$$\begin{gathered} f_i(\boldsymbol{x}_1) \leq f_i(\boldsymbol{x}_2), \forall i \in \{1, ..., k\} \\ \text{and } f_i(\boldsymbol{x}_1) < f_i(\boldsymbol{x}_2), \text{ for at least one of } i \in \{1, ..., k\} \end{gathered} \tag{5.4}$$

For a given set of design vectors, a set of them can be said to be non-dominated if no other design vector strictly or Pareto- dominates them. Considering instead the whole set of feasible design vectors S, the Pareto front can then be defined

as its non-dominated set, as illustrated in figure 5.3. The goal of multi-objective optimisation is to approximate the Pareto front as closely as possible.

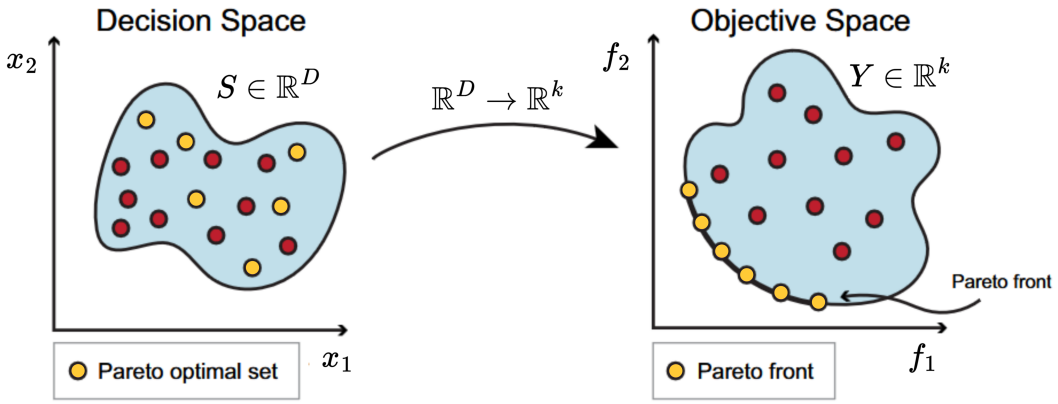

Figure 5.3: The Pareto optimal set and associated Pareto front of a bi-objective minimisation problem where $k = 2$. Adapted from [212]. Image reproduced under the Creative Commons Attribution License.

In this chapter, the application of optimisation to thermal power conversion system design in literature will first be reviewed. The optimisation methodology of this work will then be outlined, including the integration of the dynamic process models presented previously into the optimisation.

5.1 State of the Art in Optimisation for Thermal Power Conversion System Design

For the design optimisation of thermal power conversion systems, problems are typically constrained and non-linear [269, 168]. Power conversion system performance is dependent on the thermophysical properties of fluids, discussed more in section 4.1.1, which are generally complex and non-linear equations of state. The range of valid temperatures and pressures is constrained by the available cold and hot sources, as well as by material limits. Formal application of mathematical optimisation to thermal power plants began in the 1940's to analytically optimise feedwater heaters [269], and is now a driver of performance improvements in modern designs [168]. Innovation in novel cycles, such as supercritical CO_2 Brayton cycles, and integration with other technologies, as seen in combined cooling, heating and power or

waste heat recovery systems, also motivates the use of optimisation where industry experience may be less developed. Two categories are identified for the use of optimisation for thermal power conversion systems design in modern literature: parametric optimisation and structural synthesis. These are discussed in the following sections.

Dynamic optimisation of thermal power plants has also been seen, where the control of the process is optimised [218, 248]. As the plant designs are fixed, dynamic optimisation is not discussed further here.

While the choice of algorithm is dependent on the problem, algorithms should not be problem-specific. As such, the development of novel algorithms is not directly relevant to the design of power conversion systems, and a review of the state of the art is not included here. Please refer to the following for details on meta-heuristic algorithms [35, 94, 222], derivative-free algorithms [207] and general optimisation methodologies in engineering [289].

5.1.1 Parametric Optimisation

Parametric optimisation is the optimisation of a design for a fixed process structure. The design variables typically include key temperatures or pressures identified for the specific process, and can include other variables such as the split ratio of flows between two streams [159, 202, 294, 282, 265, 77, 310, 163, 302, 303, 290, 173] and the effectiveness of regenerators [202, 158, 77]. Often, parametric analysis is performed before optimisation to select the most suitable design variables. Objective variables can generally be categorised as thermodynamic, economic or environmental.

Thermodynamic objectives capture the performance of a power conversion system. Thermal efficiency, the effectiveness with which the thermal energy input to the system is converted into useful work, is a natural choice for optimisation [104, 290, 77, 161, 202, 185, 159, 173, 302, 282, 301, 78], as higher efficiency can be correlated to more power output, less fuel consumption or a combination. Some studies consider the thermal input and power output directly as objectives [236, 304, 303, 310, 272, 173, 282, 51], as they can trade-off against efficiency. Exergy efficiency is also a useful objective for power conversion systems [265, 112,

163, 158, 294], as it can give more information of energy losses that are avoidable within a system, although component-based exergy analysis is necessary to identify which components are responsible for the most exergy destruction and candidates for efficiency improvements [269].

The trade-off with increasing the efficiency of a power conversion system is often increasing capital costs. For example, increasing the turbine inlet temperature, which increases the thermal efficiency [34], requires larger heat exchangers to bring the working fluid closer to the temperature of the heat source. It is common therefore for optimisation to include economic objectives in conjunction with thermodynamic objectives. Total system cost can be used [104, 186], but thermoeconomic objectives are more common, such as payback period, levelised cost of energy or equivalents for systems with other outputs like cooling power [294, 304, 112, 272, 271, 77, 163, 302, 290, 173, 181, 51, 78, 114]. All of these objectives require cost models for the system components.

Environmental objectives have also become of interest with the pressures of climate change. The simplest environmental objectives consider the emissions associated with the fuel consumption [186, 315, 114] or a single conversion factor for the whole system [271]. More detailed analyses quantify environmental impact of individual components, using life cycle assessment to include construction, operation, maintenance and disposal, as in the Eco-indicator 99 method [272].

It is noted that the majority of recent parametric optimisation studies cited here focus on novel power conversion systems and applications. This could be due to global pressure to shift away from conventional fossil fuel thermal power plants, the minimal possible improvements to conventional power plants [103] or the advancements of technologies such as organic Rankine cycles and supercritical CO_2 Brayton cycles making these novel applications feasible. Waste heat recovery systems [304, 303, 272, 186, 282, 294] and co-generation systems [290, 158, 186, 272, 271, 114] offer large efficiency gains for systems by using energy that would previously have been lost as waste heat. Extreme applications such as submarine and space power systems are also considered, where the size and mass of the systems are

constrained [152, 77, 173].

The majority of parametric optimisation studies here used meta-heuristic population-based algorithms, most notably the genetic algorithm, due to their applicability to multi-modal, non-linear, non-continuous and black-box objectives. This is particularly useful when external simulation programs are being used for objective evaluations. Alternative algorithms were sequential quadratic programming [159], Powell's method [158] and a hybrid of the genetic and gradient descent algorithms [310].

For certain applications, such as solar thermal power systems, incorporating uncertainty into the optimisation process will help to generate robust process designs. Uncertain parameters can be incorporated into deterministic process models using Monte Carlo simulations, randomly sampling uncertain variables and evaluating the model repeatedly to estimate the probability distribution of the objectives [249, 224, 157]. To propagate uncertainty from the model inputs to the objectives this way requires a large number of function evaluations for each design, of the order of 1,000 evaluations for 14 uncertain parameters. To make this computationally tractable, surrogate modelling of the process can be used [224, 157]. Multi-parametric programming can also be used to generate optimal solutions as a function of uncertain parameters [47]. Design robustness to uncertainties can also be quantitatively integrated into the optimisation by using properties of the simulated probability distributions as objective functions [249, 224, 157].

The parametric optimisation studies cited above all assumed steady state processes. After performing steady state optimisation, some studies then analysed the dynamics of a non-dominated design [186, 158]. Others considered variations in operating conditions, either by optimising on a single-case basis for each set of conditions [255, 294] or a multi-case basis that considers all operating conditions simultaneously [249, 255]. These approaches neglect the dynamics of the system between operating points however, which could be important to a design's performance and feasibility for systems with fast transients such as solar thermal plants.

Studies have been performed on the simultaneous optimisation of thermody-

namic objectives with the dynamic settling time of a controlled solar field [14, 175]. These studies consider a steady state steam Rankine cycle model coupled to dynamic models of the solar field, so the optimisation does not consider inertia in the steam Rankine cycle.

While many studies do consider and compare multiple system layouts [202, 290, 302, 310, 185, 159, 282], the layouts are not chosen systematically, meaning optimal designs may be excluded from the design space. This limitation is what structural synthesis methodologies aim to address.

5.1.2 Structural Synthesis

Optimising the topology of a process adds another layer of complexity to parametric optimisation, as the design parameters considered in such studies should still be optimised either as a sub-level problem or simultaneously with the topology. Generally, a synthesis problem will consist of both integer and continuous design variables. Two approaches to optimisation-based process synthesis of thermal energy systems have been identified in literature: superstructure and superstructure-free.

The superstructure approach requires the definition of a structural design space within which are embedded all possible designs of interest. Within such a superstructure, the existence of different process units and streams can be toggled via binary design variables. For a specified set of binary design variables, a process model should then be generated that represents the governing equations of the design and can be used to predict its performance. An alternative to binary variables is to use zero mass flow rates to model non-existence of units in the model [234], reducing the optimisation problem from mixed integer non-linear to only non-linear. Problems can also be decomposed into a bi-level problem, with all integer variables being optimised by the upper level problem [169, 83, 84].

Some superstructure studies for waste heat recovery plants expand on the methods for heat exchanger networks. Taking a set of hot and cold streams, with predefined flow rates, inlet and outlet temperatures, heat exchanger network synthesis generates a system of heat exchangers, splits and mixers that exchanges heat between these streams and optional utility streams to achieve the desired temperatures for a

minimal cost [211]. One method for heat exchanger network synthesis is based on superstructures [113, 132, 2]. In figure 5.4, the integration of two superstructures of a heat exchanger network and a Rankine cycle with multiple optional pressure levels is shown. Integrating the Rankine cycle streams into the network as process streams allows the optimisation of both superstructures sequentially [293] or simultaneously [111, 169, 83, 84, 277], with the latter approach claiming improved designs. More simply, the heat exchanger network of a Rankine cycle can be considered for synthesis alone [149].

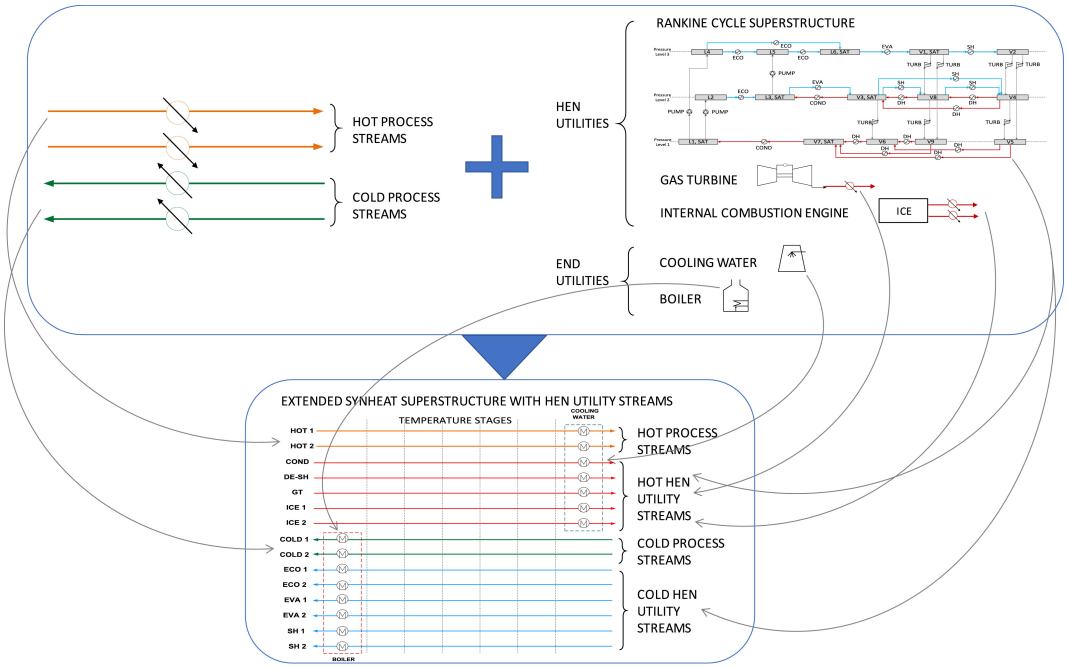

Figure 5.4: Integrating superstructures of heat exchanger networks and Rankine cycles allows for simultaneous optimisation of both systems for waste heat recovery systems [83]. Image reproduced with permission of the rights holder, Elsevier.

Other studies consider superstructures of stand-alone cycles. These can be generated by defining a number of pressure levels, any pairs of which can be connected by optional pumps, turbines or valves [268, 313, 130]. Alternatively, a structure can be defined with switches that alternate between different stream configurations [136, 17, 164, 286, 309].

Superstructure methods have also been extended to include uncertain or variable parameters. The robustness of a design can be assessed post-optimisation [10] or

incorporated directly into the optimisation [238, 221].

A weakness of superstructure approaches is that biases can be built into the problem formulation. As the superstructure is necessarily a subset of all possible designs, good or optimal designs may be excluded from the design space. The chance of losing optimum designs in the design space is noted to be related to the level of prerequisite knowledge applied to define the superstructure, with a trade-off against the volume of the design space and hence the computational difficulty of finding optimal designs [100]. Superstructure-free approaches have been developed to avoid this issue.

Superstructure-free approaches typically develop designs from the ground up. Starting with a simple configuration, a set of mutation rules can be defined that allow the structure to evolve into any feasible structure. Evolutionary algorithms are typically used due to the natural implementation of a method's mutation rules with the algorithms. Two approaches featured in literature are the SYNTHSEP [141] and energy conversion hierarchy [267] methods.

The SYNTHSEP method decomposes an energy conversion system into a set of elementary thermodynamic cycles, each of which consists of a heating, expansion, cooling and compression process, that share at least one process with another elementary cycle. Mixers and splitters mark where these elementary cycles begin or end shared processes respectively. Configurations can be mutated by altering the binary design variables stating whether two particular elementary cycles are sharing a particular process. The complexity of designs can be limited by imposing a fixed number of elementary cycles. The optimisation is decomposed into a bi-level problem. From the HEATSEP method, all heat exchange processes are replaced with "thermal cuts", which are all treated as one single black-box heat exchanger network, as shown in figure 5.5. The temperatures at both ends of these cuts are free design variables in the upper level optimisation, in addition to the configuration and pressures of the elementary cycles. The lower level optimisation then adjusts the mass flow rates in each elementary cycle, with one constraint from the Pinch Analysis theory to ensure the heat exchange process is feasible. This lower level is a

linear problem when maximising net power output. The black-box heat exchanger network can then be synthesised separately once the optimum design has been found. This method has been applied to Rankine cycles [141, 244, 151] and then extended for trans-critical or supercritical cycles [40, 52, 287, 53, 65] and semi-closed cycles [284]. These methods have generated designs that outperform any others found in literature by the authors at the time [40, 65].

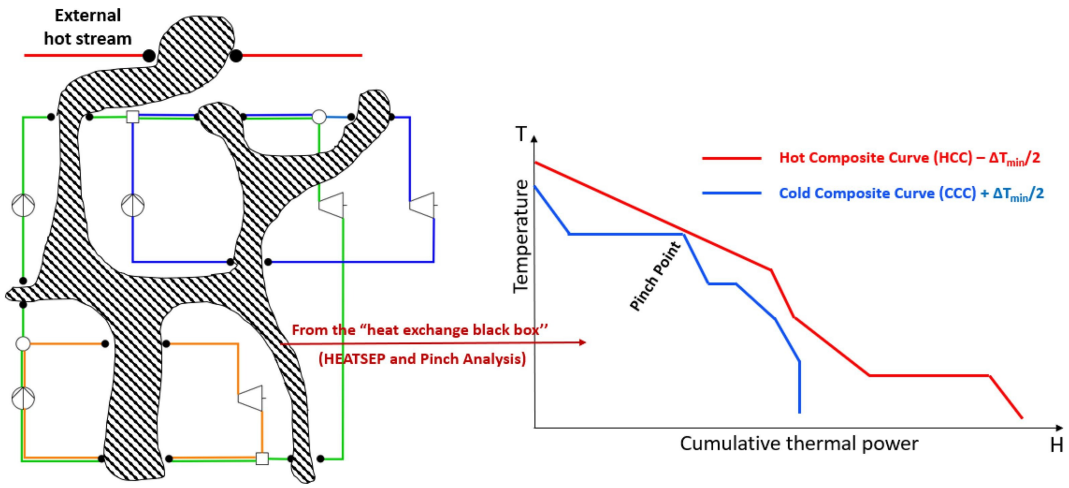

Figure 5.5: A process composed of three elementary cycles, with all heat exchange processes lumped into a single black-box per the HEATSEP method. The composite curve from Pinch Analysis theory ensures there is a feasible heat exchange network for the black-box [65]. Image reproduced with permission of the rights holder, Elsevier.

The energy conversion hierarchy method assigns various functions to each process unit, such as the "water heater" function of both steam to water and gas to water heat exchanger units. A set of six mutation rules allow units to be deleted, replaced or added based on these functions. The optimisation is also decomposed into a bi-level problem, with the upper level determining the process structure and the lower level determining the unit sizing and operation. Few uses of this method for thermal power plants were seen in literature, possibly due to the higher level of expert knowledge required to define the hierarchy [65].

Thermal efficiency or net power output are often considered as an objective function for synthesis. Used alone, they can encourage designs towards high complexity. Economic or thermoeconomic objectives can be included in synthesis optimisation to favour reasonable complexity [266, 118]. A maximum limit on the number of

each process unit can alternatively be imposed [266].

As synthesis brings in integer design variables to the problems, some researchers utilise this to simultaneously consider the selection of working fluids and the structure with minimal complexity added to the solution process [295, 17, 151]. These studies demonstrate how optimal structure depends on the fluid properties. Other studies consider the different fluids as separate cases [149].

It has been demonstrated that the optimal structure can change depending on many process parameters [286, 65]. Even more so than for parametric optimisation, synthesis methods are most promising for novel plants, such as those integrating multiple different heat sources [149], where there will be a lack of research literature and industry experience to guide the design process.

5.2 Optimisation Methodology

As highlighted in the state of the art review in section 5.1, no studies are known to directly consider process dynamics in the power cycle for the optimal design of thermal power conversion systems. With the growing interest in novel cycles for applications such as nuclear fusion, waste heat recovery and more dispatchable power generation, the transient operation should be considered from the earliest stages of design [147]. This work begins to address this research gap by demonstrating the design optimisation of thermal power conversion systems using dynamic process models.

The process models are developed by connecting individual component models from chapter 4, implemented in the Modelica language, and solved using the OpenModelica program. While OpenModelica supports some optimisation functionalities via OMOptim, it does not seem mature enough for research use yet. Hence, OpenModelica is used as an external simulator for design evaluations, as seen in other works [134, 301]. This allows the use of languages like Julia [26] for the optimisation, where a wide variety of optimisation packages are available and the user has greater control over the optimisation process.

When optimising a system, a simulation model is required to evaluate the

objective function and constraints for each design. A simulation may involve evaluating a single function, or it may require solving a large system of equations using specialised software. Once the simulation has been performed, the state variables in the model will be required to evaluate the design. In this work, the use of dynamic models in OpenModelica mean the simulation must be treated as an external black-box. This also means derivatives of the objective function with respect to the design variables will not be directly available. Derivatives can be estimated via finite difference approximations, but this would increase the computational effort of optimisation due to the larger number of simulations required per design [59, 5].

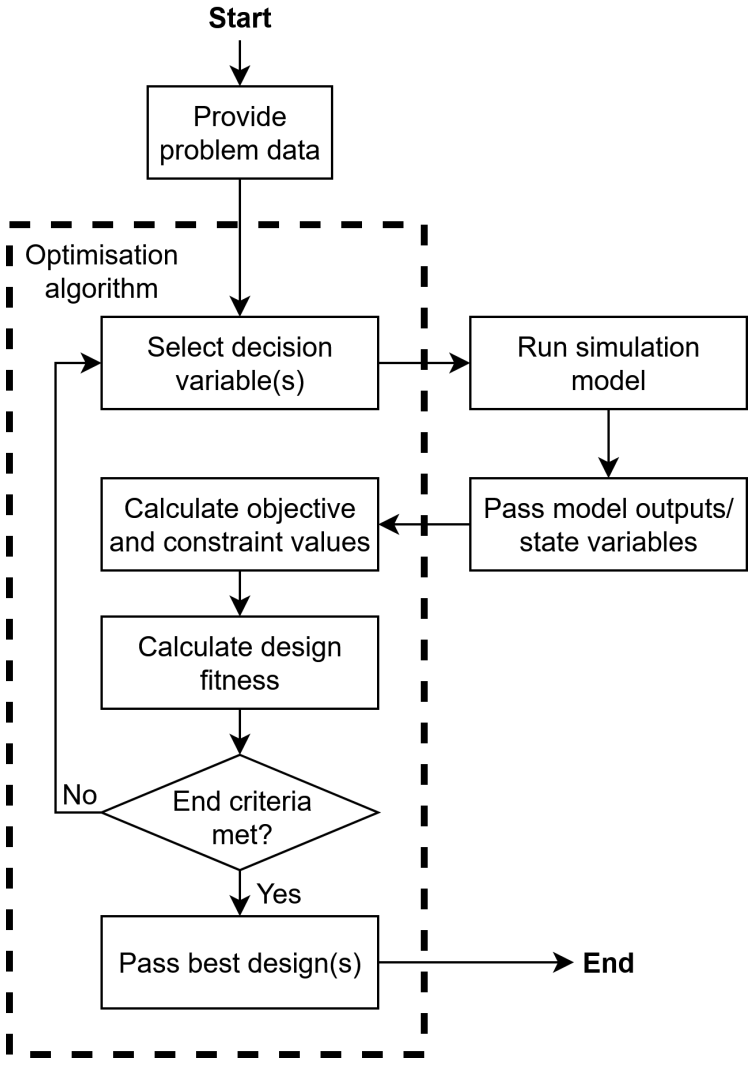

Figure 5.6: Flow of a general meta-heuristic optimisation algorithm. The simulation model is external to the algorithm, and only the simulation state variables are required to evaluate the objective functions and constraints [35].

Meta-heuristic algorithms have been widely used for optimisation in engineering domains, including thermal power conversion system design as seen in section 5.1. They are based on targeted sampling of the design space, where each algorithm differs in how it explores the space and generates new designs using information from previously evaluated designs. The algorithms are independent of the simulation models, and only require the state variables from the simulation to evaluate the objective functions and constraints, as shown in figure 5.6. This means they can be applied to any problem [245]. The drawback of this generality and the often stochastic nature of the algorithms is that there is no mathematical guarantee of convergence towards global optima. The algorithms often have a large number of parameters, with a poor understanding of how these impact the performance across different problems.

When using expensive simulation models, the independence of the algorithm and simulation model allow another benefit in the form of population-based meta-heuristic methods. These methods maintain a population of multiple designs simultaneously which are iterated upon in "generations". This feature allows the algorithm to exploit parallel processing of computers to evaluate multiple designs in parallel, reducing the real-time cost of optimisation. These methods are also well-suited for multi-objective optimisation problems where there may be a continuum of non-dominated designs, allowing a set of designs to be generated and returned as the optimisation outcome [123].

5.2.1 Plant Propagation Algorithm

Following from the previous discussion, the plant propagation algorithm implemented in the Fresa.jl package is chosen for this work [216, 92]. It is a population-based, meta-heuristic algorithm inspired by the propagation strategy of strawberry plants. A parent strawberry plant will send out runners that can root and grow into daughter plants. It has been observed that well-established plants growing in good soil and ambient conditions will produce a large number of runners that root in close proximity to the parent plant. Weaker plants growing in poorer conditions will send a smaller number of runners, due to the more limited resources available for

propagation, that travel comparatively far from the parent plant in the hope of the young plants finding better conditions.

Mapping this to an optimisation algorithm, plants are instead individual designs in a population, and each generation, some are selected to persist or propagate new designs to the next generation based on their fitness. Fitter designs generate a larger number of runners that are relatively close in the design space, while less fit designs send fewer runners a further distance. This behaviour aims to balance the exploration and exploitation of the design space such that the algorithm should be capable of finding global minima. An outline of the implementation of the plant propagation algorithm in Fresa is given in algorithm 1.

Algorithm 1 Psuedocode of the plant propagation algorithm as implemented in the Fresa.jl package [92, 93].

Require: Objective $\boldsymbol{f}(\boldsymbol{x})$, $\boldsymbol{x} \in X$, Parameters $n_{\text{gen,max}}$, n_{p}, $n_{\text{r,max}}$
  $P \leftarrow$ initial population ▷ Pre-defined or randomly generated
  for $n_{\text{gen,max}}$ generations **do**
    $N \leftarrow$ fitness(P) ▷ Evaluate fitness of each design in population
    if elitism **then** ▷ Keep some of the best solutions in the population
      $\tilde{P} \leftarrow \{$up to $\lceil n_{\text{p}}/2 \rceil$ of non-dominated designs in $P\}$
    else
      $\tilde{P} \leftarrow \varnothing$
    end if
    for $i \leftarrow 1, ..., \min\{n_{\text{p}}, |P|\}$ **do** ▷ Propagate up to n_{p} designs
      $j \leftarrow$ select(P,N) ▷ Fitness-based selection
      $\tilde{P} \leftarrow \tilde{P} \cup \{\boldsymbol{x}_j\} \cup \{\boldsymbol{x}_k | \boldsymbol{x}_k = \text{runner}(\boldsymbol{x}_j, N_j); k = 1, ..., n_{\text{r},j} \propto N_j n_{\text{r,max}}\}$
    end for
    if pruning **then**
      prune $\tilde{P}$ ▷ Remove similar designs from new population
    end if
    $P \leftarrow \tilde{P}$
  end for
  return non-dominated set of P

In Fresa, each design point is associated with a design vector $\boldsymbol{x}$, an objective vector $\boldsymbol{z} = \boldsymbol{f}(\boldsymbol{x})$, an infeasibility value g and an ancestral point. g represents the violation of the constraints in equation 5.1 by a design, with designs being feasible if $g \leq 0$ and infeasible otherwise. Fresa computes the fitness N of a population based on a combination of $\boldsymbol{z}$ and g, with the calculation differing for single- and multi-

objective problems and for feasible or infeasible designs, as outlined in table 5.1. If both feasible and infeasible designs are present, their fitness values are mapped to the open intervals $(0.5, 1)$ and $(0, 0.5)$ respectively. Otherwise, the only present set will be mapped to the open interval $(0, 1)$. The fitness values of a generation will then always lie in the open interval $N \in (0, 1)$, with the best designs close to 1 and the worst close to 0.

Fitness N	Single-objective	Multi-objective
Feasible	Ordered by $\boldsymbol{z}$	Hadamard: Ordered by product of rank for each objective z_j [95] Borda: Ordered by sum of rank for each objective z_j Non-dominated: Ordered by recursively removing the non-dominated set from generation [70]
Infeasible	Ordered by g	Ordered by g

Table 5.1: Fresa's calculation of a design's fitness depends on the optimisation problem and its feasibility. Three different ranking methods are available for feasible designs in multi-objective problems.

For multi-objective problems, the fitness ranking method for feasible designs can have a large impact on the search behaviour. The Hadamard method ranks designs near the extreme objective values in the non-dominated set as fitter than those compromising multiple objectives. This tends to lead to a wider range of objective values at the cost of fewer designs at intermediate values. Conversely, the non-dominated method ranks methods by recursively removing the non-dominated set from the population, with the fitness of a design decreasing with increasing number of recursive removals. As the original non-dominated set are ranked equally, this tends to emphasise more evenly dispersed populations at the expense of a smaller range of objective values.

Variations of the original plant propagation algorithm have been presented for discrete optimisation [138], image generation [187], modified runner generation [235], dynamic fitness functions [257, 256] and modified selection [36]. Fresa implements the tournament selection procedure and supports the dynamic fitness calculation of the latter two, both shown to improve the algorithm's performance on

benchmark functions. The former two can be implemented with Fresa using Julia's multiple-dispatch capability to define functions for arbitrary design representations, or even multiple representations simultaneously [93].

A common criticism of meta-heuristic algorithms is the large number of algorithmic parameters that have to be tuned relatively arbitrarily for different problems. Compared to popular methods, such as genetic algorithms and particle swarm optimisation, the plant propagation algorithm has relatively few parameters to tune, depending on the exact implementation. For Fresa, the main parameters to be considered are the number of designs to propagate n_{p}, the maximum number of runners per propagated design $n_{\mathrm{r,max}}$ and either the maximum number of generations $n_{\mathrm{gen,max}}$ or the maximum number of function evaluations $n_{\mathrm{f,max}}$. For multi-objective problems, the choice between the three ranking methods presented in table 5.1 is also an important parameter.

Parameter sensitivity studies have looked at the parameters n_{pop} and $n_{\mathrm{r,max}}$, with fixed $n_{\mathrm{f,max}}$. n_{pop} is the fixed size of the population, and is most analogous to n_{p} in Fresa, but is not equivalent. For the application to multidimensional benchmark functions, it was shown that the algorithm is largely insensitive to these parameters [69, 68]. For the discrete Travelling Salesman Problem, a log-normal pattern was observed in the performance parameter sensitivity, although an underlying cause of this was not identified. On the other hand, the algorithm seemed quite insensitive to different instances of the same problem [137].

The conclusions of these sensitivity studies cannot in good confidence be extended to this work. The benchmark functions considered were all single objective, whereas the case studies in chapter 7 are multi-objective and include an additional parameter for the fitness ranking method. Fresa's implementation of the algorithm differs from the presented works, such as the features of tournament selection and elitism, with undocumented impact on the parameter sensitivity. As a result, the algorithm parameters were based on trial-and-error, lacking the computational resources for a rigorous sensitivity analysis for each case study.

A comparison was made between the performance with the Hadamard and

non-dominated fitness ranking methods in section 7.3.2, with the aid of a multi-agent system that allowed for both methods to be used simultaneously.

5.2.2 Multi-Agent Optimisation System

Multi-agent systems are programs that allow multiple independent programs, called agents, to interact with each other. The Cocoa.jl package implements a multi-agent based framework that allows multiple optimisation agents, which could be the same algorithm with different parameters, to solve a problem simultaneously [96]. The agents and the interactions between them are shown in figure 5.7. The scheduler is the core agent of the system that shares information between the different optimisation algorithms, labelled as the solver agents. The scheduler controls the solvers' access to the model, which actually evaluates the objective function and constraints for each design.

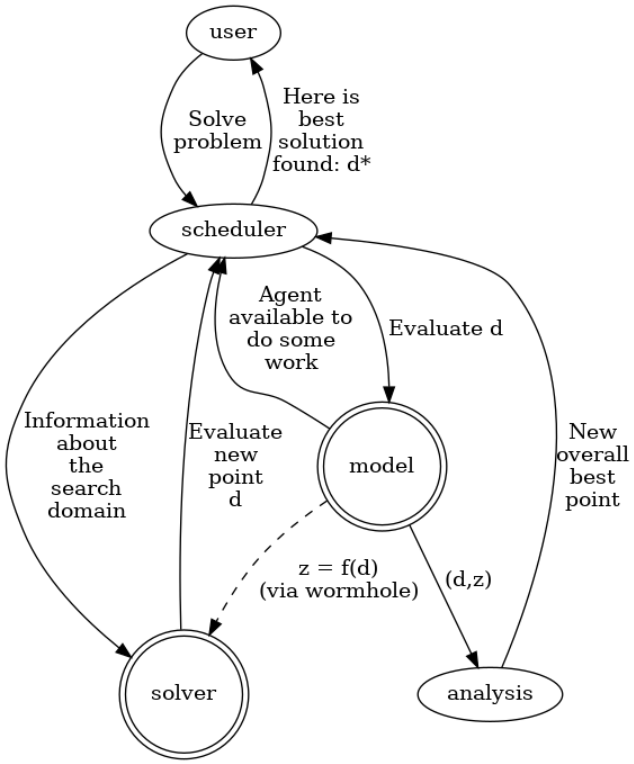

Figure 5.7: The agents, and the information flows between them, in the Cocoa.jl multi-agent optimisation system. Double circles indicate that multiple instances of the agent may be present [96]. Image reproduced with permission of the rights holder, Elsevier.

The separation of the optimisation algorithm and model evaluation agents is

motivated by the aim to allow agents to compete for computational resources based on their performance [247], which isn't currently implemented. The system does currently allow the scheduler to share any new non-dominated solutions found with all the solver agents however, with the prospect of exploiting the search behaviour of different algorithms simultaneously to improve the optimisation.

This cooperative multi-agent system is tested in Case Study II, and its performance is discussed in section 7.3.2.

5.2.3 Integration Between Julia & OpenModelica

To use the dynamic models developed in OpenModelica within the Fresa optimisation algorithm, the scripting interface OMJulia.jl is used [60]. OMJulia allows OpenModelica instances to be created in Julia, within which a dynamic process model can be simulated. A similar coupling between OpenModelica and Python has been seen for the steady-state design optimisation of a Brayton cycle [301].

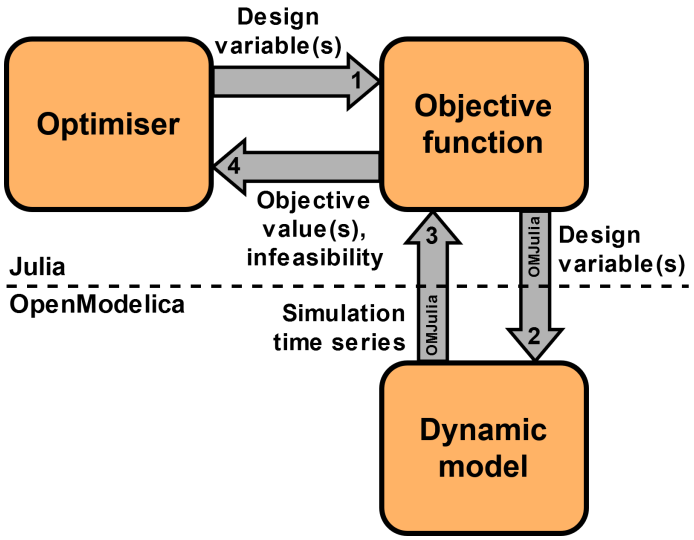

Figure 5.8: Flow chart for the evaluation of a design. OMJulia is used to create an OpenModelica instance for simulating the dynamic model with the provided design variables. The results of the simulation can then be passed back to Julia to evaluate $\boldsymbol{z}$ and g for the design.

The basic flow of information in the program for a single design evaluation is shown in figure 5.8. The optimiser calls the objective function with a design vector it wants to evaluate. The objective function then creates an OpenModelica instance,

using OMJulia, loads the dynamic process model and sets model parameters per the provided design vector. Once the simulation is complete, the objective function can then access the time series of all model variables to calculate the values of $\boldsymbol{z}$ and g, which are then passed back to the optimiser.

For Case Study I, presented in chapter 6, some designs fail to simulate successfully. Manual inspection of these designs shows they can be simulated successfully in the OpenModelica client, and no behaviour in the simulation results can be identified that would lead to crashes. It is concluded that this is an issue solely with the coupling between Julia and OpenModelica, not the quality of the designs. To circumvent this, the objective function is modified to allow the simulation of each design to be attempted up to 10 times. A tally of the repeats required across the optimisation run in Case Study I is presented in figure 5.9. This modification allows all 229 design evaluations to be performed successfully.

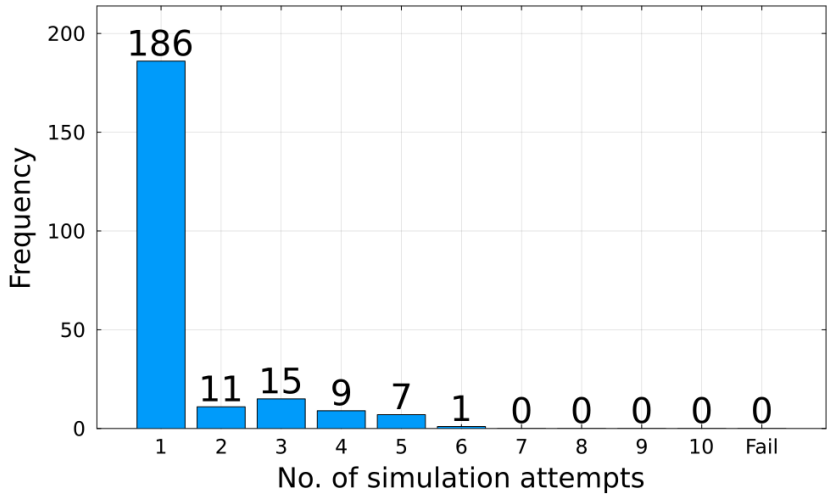

Figure 5.9: Tally of the number of attempted simulations for each design evaluation in Case Study I. All 229 evaluations were eventually completed successfully.

The repeated simulation method is not applied to the other case studies. The models are complex enough that each simulation had an evaluation times $>$1min, and the inclusion of pulsed heat sources meant that designs could fail due to the tank level assertion in equation 4.34. It is decided the time consumed by re-evaluating every failed design would harm the optimisation results more than occasionally losing a feasible design.

Another issue encountered when using dynamic process models is the variable time required for simulations. Poor designs could lead to oscillations in the process, which requires the OpenModelica solver to use smaller time steps and hence increases the evaluation time of the design. This resulted in poor designs drastically slowing down the optimisation, holding up the algorithm from moving on to the next generation. For the case studies in chapter 7, a time limit is imposed on the evaluation of each design to mitigate this issue.

Further testing of the coupling found that some instances of OpenModelica created using OMJulia would not be correctly terminated, leading to a gradual increase in used memory until the program crashed for longer optimisation runs. To circumvent this, a higher level shell script is written that can repeatedly call the Julia optimisation script. This allows any hanging OpenModelica sessions to be purged, clearing the memory, after which the optimisation can be reinitialised with the final population of the previous Julia script. The Serialization package of the Julia Standard Library is used to save and read population files [26]. The package allows arbitrary object types to be saved, allowing the population to be read directly as *Fresa.Point* objects and avoiding additional computational cost of re-evaluating designs. The flow chart for the full optimisation methodology is summarised in figure 5.10.

Reinitialisation is not necessary for Case Study I where the number of function evaluations was low enough that the program did not crash prematurely. The flow diagram for that case study can be considered by removing the reinitialisation script level and its associated loop from figure 5.10.

The state of the Fresa algorithm is solely described by the current population. Fresa returns the entire population upon reaching its termination criterion. Hence, the entire population will be restored upon reinitialisation and there should be no impact on the optimisation outcomes.

When using Cocoa however, the system state is described by the current population of each solver agent. Past the initialisation of the system, the populations of the solver agents will not be identical for stochastic or non-identical solvers,

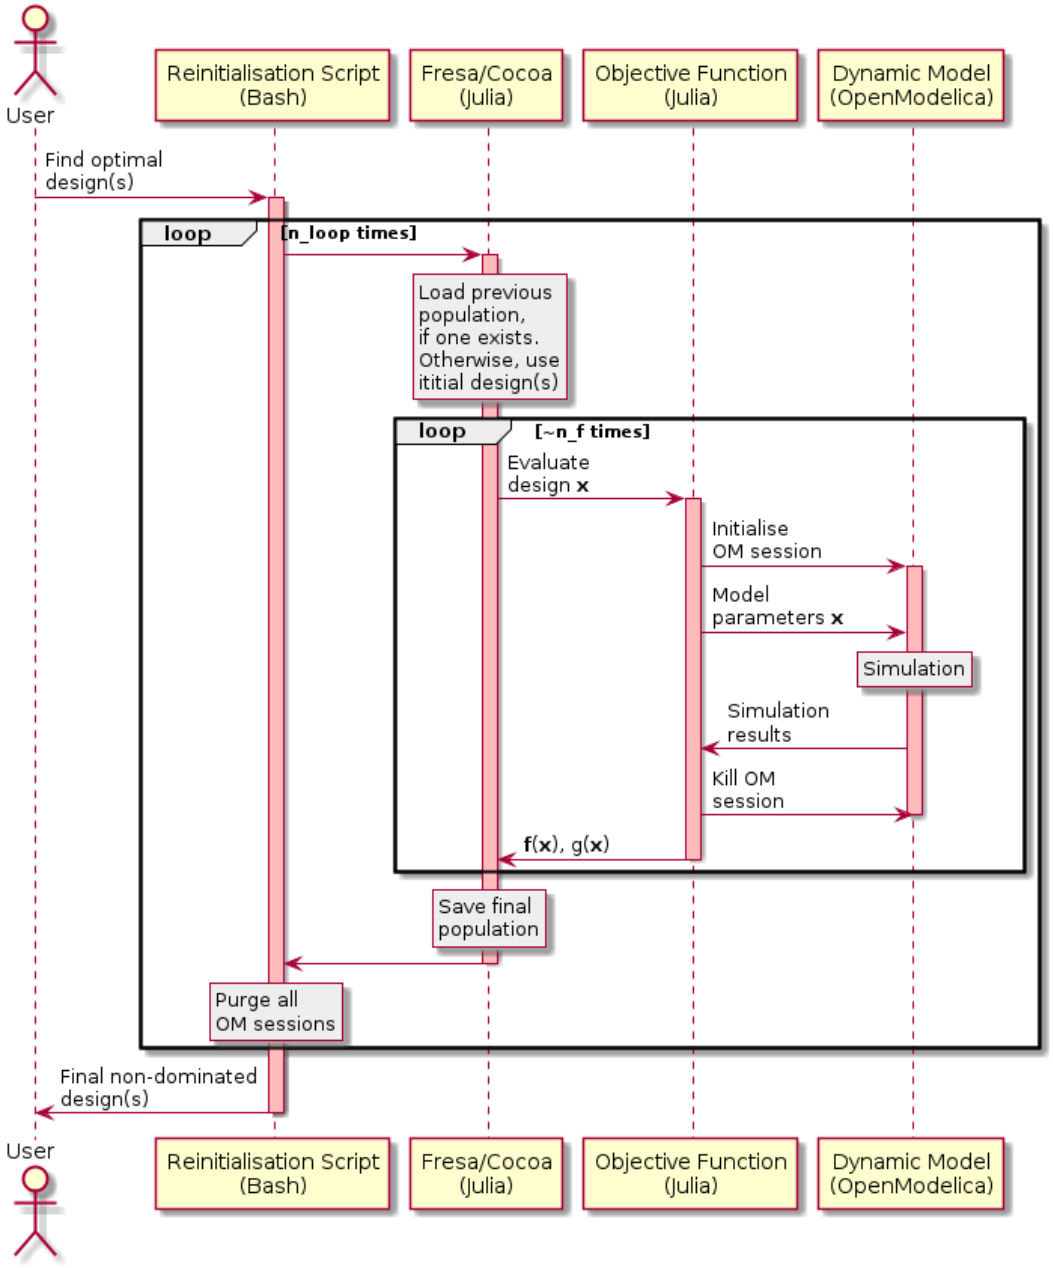

Figure 5.10: Flow chart of the optimisation methodology. At the highest level, a Bash script repeatedly calls a Julia script using Fresa (or Fresa and Cocoa) as the optimiser. This allows hanging OpenModelica sessions to be purged by the Bash script to clear memory periodically. The optimiser calls an objective function, also in Julia, that can initialise an OpenModelica session and pass some design parameters $\boldsymbol{x}$ to the dynamic model. After a simulation, the objective function can then access the results as a time series to evaluate $\boldsymbol{f}(\boldsymbol{x})$ and $g(\boldsymbol{x})$ before attempting to kill the OpenModelica session.

and only new non-dominated designs will be shared between solver populations due to the scheduler sharing them. Cocoa only returns the non-dominated designs upon termination. Upon initialisation, all solver agents will be passed the same population, which in this case will be the non-dominated designs of the previous loop. A consequence of these two details is that when the Julia script is reinitialised, the populations will be homogenised between all solver agents and any diversity added by dominated designs will be lost.

This will impact the optimisation results, but it is not possible to quantify this for the case studies while the memory issues persist. It is suspected that the loss of diversity will harm the gains offered by combining different search behaviours simultaneously. With large enough gaps between reinitialisations however, it is hoped the individual agents would have enough time to demonstrate their particular behaviours and find some non-dominated designs that can persist to the next loop. A comparison between the optimisation framework's performance with and without Cocoa is presented in section 7.3.2.

5.3 Summary

For the design of thermal power plants, two main areas of use for optimisation are identified. Parametric optimisation concerns the optimisation of process parameters, such as temperatures, pressures and units sizes, for a fixed process structure. Objectives can be thermodynamic, like thermal efficiency, economic, like levelised cost of electricity, environmental, like total emissions, context specific, like total system mass, or a combination of these. The fixed structure can limit the scope of possible improvement for designs. Structural synthesis allows for the optimisation of the process structure and process variables, either by defining a superstructure or defining a set of mutation rules.

For the optimisation of dynamic process models, the plant propagation optimisation algorithm from Fresa.jl is used. Being a black-box, population-based algorithm, it is well suited for objective functions involving costly simulations. Compared to other meta-heuristic algorithms, it has relatively few parameters and has been

shown to be quite insensitive to these parameters on benchmark functions. For multi-objective problems, multiple fitness ranking methods are available that change the search behaviour.

To benefit from the strengths of different fitness ranking methods' simultaneously, the multi-agent optimisation system Cocoa.jl is used. This allows for two instances of Fresa, with either the Hadamard or non-dominated ranking methods, to act in parallel while sharing the best solutions found.

The coupling between OpenModelica and Julia is noted to cause some implementation issues. Some OpenModelica sessions are not cleared after simulation, resulting in a gradual increase in memory usage that can lead to the program crashing. To circumvent this, a higher level Bash script is used to periodically kill hanging processes and restart the optimisation with the final population of the previous loop.

Chapter 6

Case Study I: Optimisation of a Controller for Load Following with a Constant Heat Source

Publications based on this chapter:
"*Optimization of a PID Controller within a Dynamic Model of a Steam Rankine Cycle with Coupled Energy Storage*" [274]

Proportional-integral-derivative controllers, discussed more in section 4.2.6, are widespread in industrial processes due to their ease of implementation. While proportional-integral-derivative controllers are single-input-single-output, each with one controlled and one manipulated variable, they are generally used in processes with multiple controlled variables.

Traditional tuning methods are based on linearised models about a specific operating point and are performed one controller at time [170]. Drawbacks of such methods are interactions between control loops impacting the performance of controllers and the possibility of poor control for processes with a wide range of operating conditions.

Optimisation algorithms have been used as an alternative tuning method for proportional-integral-derivative controllers in power generation systems, such as for load frequency control [81, 180, 215, 56, 75, 226] and boiler control [225]. Optimisa-

tion methods allow for the full range of operating conditions to be considered during tuning. The choice of objective functions also lets different controller behaviours be prioritised depending on context, such as by weighted sums of maximum overshoot, rise time, settling time or steady state error [215]. Optimisation methods have also been extended to simultaneously tune multiple controllers [180, 225, 226].

In this chapter, the dynamic process models and optimisation methodology presented in chapters 4 and 5 are demonstrated for the optimisation of a proportional-integral-derivative controller in a power conversion system for a constant heat source.

6.1 Power Conversion System Design for a Constant Heat Source

The design of the power conversion system considered for utilising a constant heat source is shown in figure 6.1. The process is a steam Rankine cycle coupled to a sensible thermal energy storage system using molten salt. This thermal energy storage system is of the two-tank design, storing cold and hot molten salt in separate insulated tanks at 1bar with design temperatures of 300°C and 450°C respectively.

Molten salt from the cold tank is pumped at a constant mass flow rate of 1,500kg/s through a heat exchanger that delivers 335MWth of heat before entering the hot tank. This is an idealised heat exchanger with a constant heat duty that will not vary with changing inlet temperature of the cold molten salt. This was chosen as estimates for the flow rate and temperature of coolant leaving a tokamak were not yet available. The molten salt in the hot tank is then pumped through the superheater, evaporator and then the preheater in a counter current configuration with the flow of water in the steam Rankine cycle.

Water is pumped to the preheater before entering the evaporator, exiting as saturated steam that is then heated in the superheater. A spray valve bypass goes from after the preheater to before the turbine inlet, where the water mixes with superheated steam from the superheater before entering the turbine. The steam is expanded in the turbine, generating work, before being condensed back to liquid water in the condenser for pumping again. The condenser is cooled by a source of

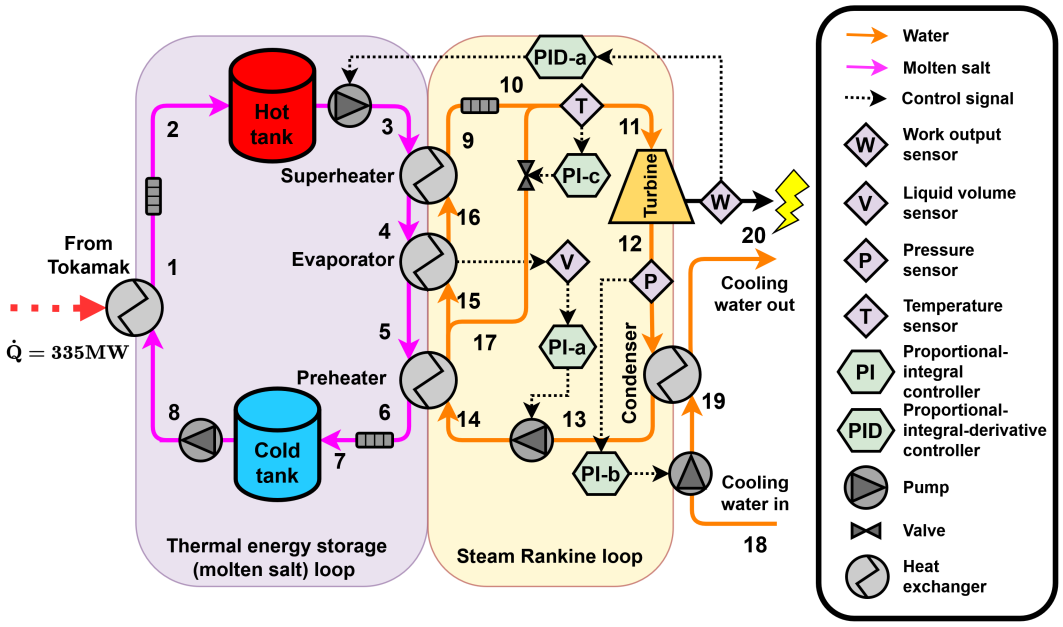

Figure 6.1: Power conversion system for an ideal, constant heat source in Case Study I. Thermal energy is delivered to a molten salt thermal energy storage loop. Hot molten salt can then be pumped on demand to drive a steam Rankine cycle to meet load demands.

20°C water that exits the condenser at 1bar.

Pressure drops, based on equation 4.42, are included after the superheater and before each molten salt tank. The former is necessary so that there is a non-zero pressure difference across the valve, which is necessary for non-zero flows per equation 4.9. The molten salt pressure drops are unnecessary, but will not impact the optimisation results, as discussed in section 4.2.5.

6.1.1 Control System Design

From figure 6.1, there are four controllers in this process: three proportional-integral controllers and one proportional-integral-derivative controller.

A proportional-integral controller (PI-a) controls the liquid level in the evaporator by manipulating the feedwater mass flow rate. This is necessary to ensure the heat being delivered to the evaporator is balanced by the amount of saturated steam being generated. If they are imbalanced, the evaporator could flood or dry out. The set point is chosen to be 0.5, where equal volumes of saturated liquid and vapour occupy the evaporator.

Another proportional-integral controller (PI-b) controls the condenser pressure by manipulating the flow of cooling water. Maintaining a low pressure in the condenser allows the turbine to extract more energy from the steam, improving the thermal efficiency. The set point is 0.05bar.

The final proportional-integral controller (PI-c) controls the temperature of steam entering the turbine by manipulating the valve opening position of the spray valve. Opening this valve allows liquid subcooled feedwater to mix with superheated steam, bringing down its temperature. Hence, the spray valve can only bring down the steam temperature from the superheater outlet. The purpose of this is to prevent thermal cycling if the turbine load changes significantly. The set point is 420°C.

The proportional-integral-derivative controller (PID-a) manipulates the flow of hot molten salt to the Rankine cycle heat exchangers to control the power produced by the turbine. The hot molten salt mass flow rate determines the thermal power being delivered to the cycle, and hence the cycle's power output. Reducing the thermal power will reduce the amount of superheated steam that can be generated, lowering the turbine inlet pressure and the mass flow rate of water around the cycle. This is comparable to sliding pressure control, where the turbine inlet pressure is controlled to regulate turbine power. A consequence of this control method is that the changing mass flow through the superheater may lead to temperature fluctuations at the superheater outlet with changing turbine load. The spray valve and its controller is included to counter this. The set point of PID-a is discussed in section 6.2. Derivative action is included as the heat exchangers are the main source of inertia in thermal power plants, and derivative action may be able to stabilise against any oscillations that arise.

The controllers were all hand-tuned by repeated simulations until the controller action was deemed acceptable. Detailed controller parameters are given in appendix A. Hand-tuning is a time and labour intensive process that does not guarantee any level of controller performance. These downsides only become more pronounced as the number of controllers and complexity of the model increases. Optimisation is investigated as an alternative tuning method for the proportional-integral-derivative

controller only.

The controllers are activated at $t = 100$s to reduce transients immediately after initialisation.

6.1.2 Model Parameters

The parameter values used for the dynamic process model are presented in appendix A. Many of the values are not based on literature, and are rather selected by repeated simulations to achieve desired process performance, such as being able to meet the requested power demand and having suitable molten salt storage to last the full simulation. The other case studies presented introduce a more rigorous parameterisation based on steady state process models and literature values to address this limitation.

Following the discussion of initialisation challenges in section 3.1.1, each dynamic model is also given a set of initial equations. In this process, the dynamic models are the two molten salt tanks, the four heat exchangers and the four controllers. The initialisation equations for each dynamic model are also presented in appendix A. Further discussion of the initialisation process and potential impacts on the optimisation results is given in section 6.3.1.

The initial conditions were partially chosen by trial and error until the model could be initialised. OpenModelica was unable to initialise the model with all of the heat exchangers at steady state, and hence fixed initial temperatures and pressures are given to the superheater and condenser.

6.2 Case Study I: Optimisation Problem Statement

To optimise the action of the proportional-integral-derivative controller regulating the turbine power output, the mean absolute error is chosen as the objective to minimise, with the design vector $\boldsymbol{x} = (K_\mathrm{c}, \tau_\mathrm{I}, \tau_\mathrm{D})^\mathrm{T}$ being the three controller tuning parameters. This objective is equal to the integral absolute error, often seen in literature, divided by the integration time, and is chosen to be more a more interpretable quantity for discussion. The problem statement is then given by:

$$\begin{aligned} &\min_{\boldsymbol{x}\in X} \left(\frac{1}{t_\mathrm{end} - t_\mathrm{start}} \int_{t_\mathrm{start}}^{t_\mathrm{end}} \left| \dot{W}_\mathrm{demand}(t) - \dot{W}_\mathrm{turbine}(t) \right| \mathrm{d}t \right) \\ &\text{s.t. } \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0} \\ &\quad \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \end{aligned} \tag{6.1}$$

$\dot{W}_\mathrm{demand}$ is the set point of the proportional-integral-derivative controller, $t_\mathrm{start} =$ 3,000s is the start time of the power demand perturbation and $t_\mathrm{end} = 4{,}000$s is the time at the end of the simulation. $\boldsymbol{g}(\boldsymbol{x})$, $\boldsymbol{h}(\boldsymbol{x})$ encapsulate the equations of the dynamic process model that cannot be written in a closed form. A smooth change from 100MW to 90MW is applied to $\dot{W}_\mathrm{demand}$ over a period of 150s, shown in figure 6.2, to test the controller's response to load changes. The domain of the search space X is given in table 6.1, with the lower limits of the time constants chosen based on typical values [219] and the other limits to give a reasonable range relative to the initial design.

Design variable	Minimum value	Maximum value
K_c	-0.01kg/(s W)	-0.0001 kg/(s W)
τ_I	2s	100s
τ_D	2s	100s

Table 6.1: Domains of the design variables for equation 6.1.

To reduce any impact of the somewhat arbitrary initial conditions discussed in section 6.1.2 on the optimisation results, 3,000s of simulated time is allowed to pass before the set point perturbation and error integral begin. By this time, any fluctuations introduced by initialisation will have faded, and the controllers are hoped

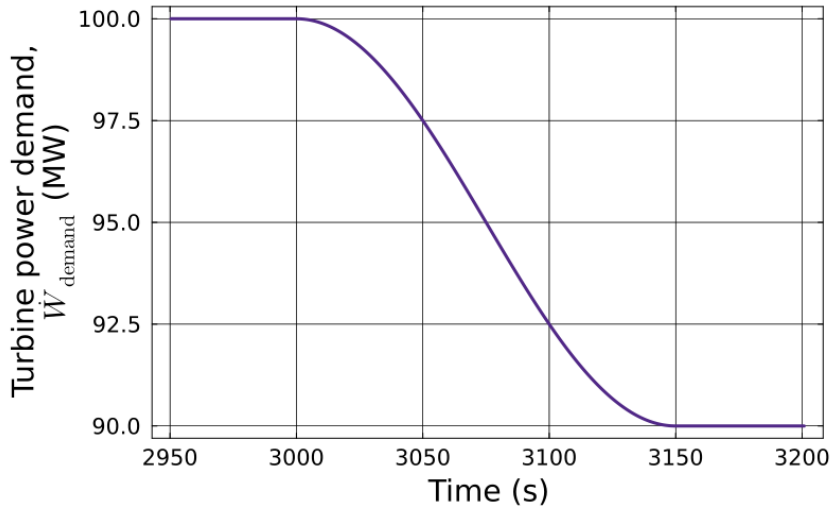

Figure 6.2: Turbine power set point perturbation to test the performance of the proportional-integral-derivative controller in Case Study I.

to have brought the system to operating conditions that are independent of the initial state. This method is assessed in section 6.3.1.

The integral in equation 6.1 is calculated within the OpenModelica model and is consistent with the accuracy of the model itself.

Fresa is used for the optimisation. All designs are considered feasible with $g = 0$. As discussed in section 5.2.3, up to 10 simulation attempts are allowed for each design, and reinitialisation is not necessary due to the lower number of function evaluations. The parameters used for Fresa are given in table 6.2. A single design point $\boldsymbol{x} = (-0.001\text{kg/(s W)}, 30\text{s}, 10\text{s})^{\text{T}}$ is given as the initial population.

Parameter	Value	Description
$n_{\text{gen,max}}$	10	Maximum number of generations for the algorithm
n_{p}	10	Maximum number of designs to propagate each generation
elite	*true*	Non-dominated solutions are always kept in the population
multithreading	*true*	Evaluate designs in parallel

Table 6.2: Parameters of Fresa used for solving equation 6.1 in Case Study I. All other parameters are left as default. Please refer to `https://www.ucl.ac.uk/~ucecesf/fresa.html` for default values and more detailed explanations.

6.3 Case Study I: Optimisation Results

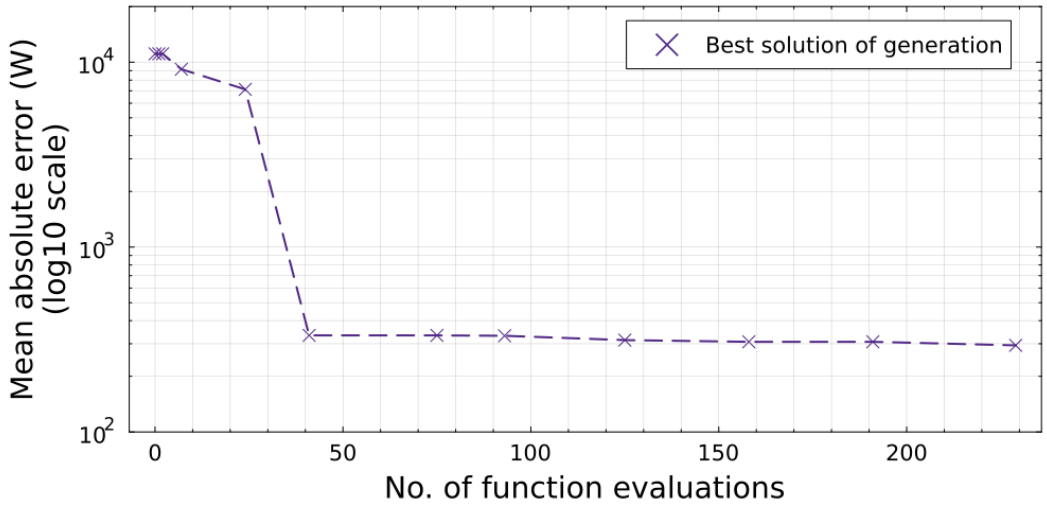

Figure 6.3: Evolution of the objective function value of the best design found by Fresa at the end of each generation in Case Study I. Note the logarithmic scale.

The evolution of the best design found throughout the optimisation is presented in figure 6.3. The best design found is $\boldsymbol{x} = (-0.00318\text{kg/(s W)}, 2.00\text{s}, 8, 41\text{s})^{\text{T}}$ after 229 function evaluations.

From figure 6.3, the objective consistenly decreases each generation. While the rate of reduction is small relative to the first 5 generations, it suggests that improvements are still possible for the design. The selection of $n_{\text{gen,max}} = 10$ is therefore too low for this problem.

To assess the improvement achieved by the optimisation, the power difference $\Delta\dot{W} = \dot{W}_{\text{demand}}(t) - \dot{W}_{\text{turbine}}(t)$ during the perturbation is shown in figure 6.4 for the initial and optimal controller designs. $\Delta\dot{W}$ is also the error of the controller. The optimisation greatly reduces the magnitude of the controller error during and after the power demand perturbation. The optimised design shows a 97% reduction in the mean absolute error and a 95% reduction in the maximum absolute error and a faster return towards the set point of $\Delta\dot{W} = 0\text{W}$ over the time period $3{,}000\text{s} \leq t \leq 4{,}000\text{s}$.

The optimised design demonstrates more oscillatory behaviour than the initial design. It is seen that the optimal design is at the lower bound of τ_{I}, corresponding to a strong integral action and suggesting the objective could be minimised further

by going past this bound. Oscillations are a known downside of integral action, but small oscillations are acceptable for the improvement in response speed and reduced offset errors. Oscillations can either be offset by derivative action or proper tuning [219].

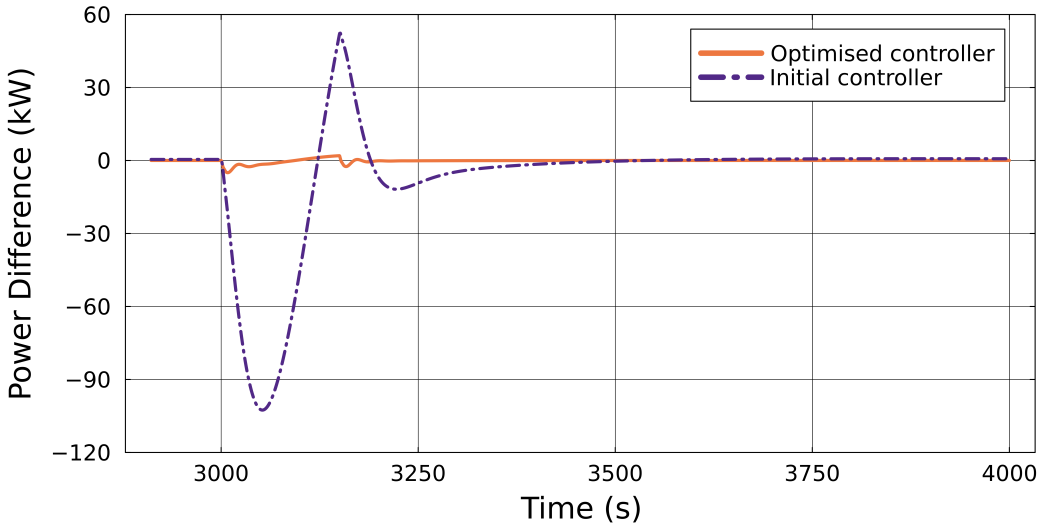

Figure 6.4: Difference between the turbine power output and the power demand due to a power demand perturbation for the initial and optimised proportional-integral-derivative controller designs in Case Study I.

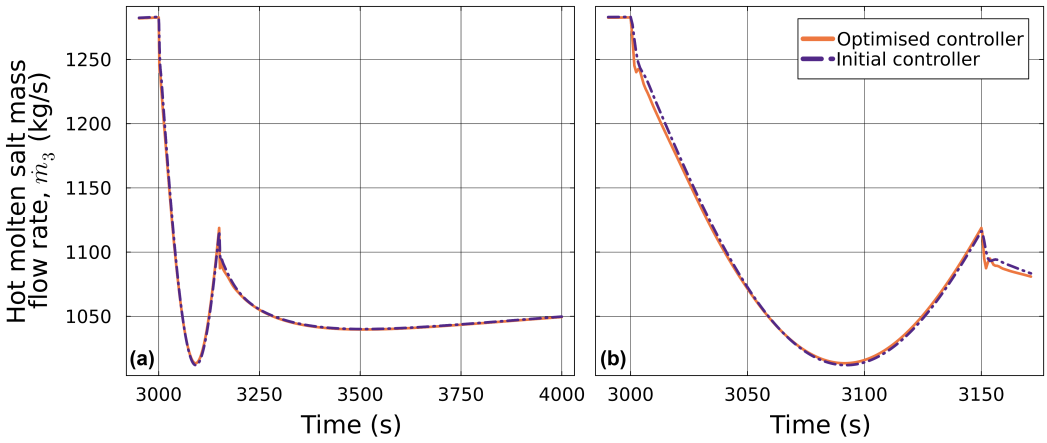

Figure 6.5: Action of the initial and optimised proportional-integral-derivative controller designs across (a) the period of the objective integral (b) the period of the perturbation in Case Study I. The controllers act on the mass flow rate of hot molten salt to the superheater

An offset error will increase the mean absolute error more than an oscillation of the same magnitude would, so it is unsurprising that oscillations are present in the

optimised design with the chosen objective. Despite derivative action being included, there is no metric in the optimisation that directly penalises oscillations to allow it to be properly tuned.

A comparison of the controllers' actions is shown in figure 6.5. The most noticeable difference between the controllers is that the optimised design reacts faster to the start of the power demand perturbation. It may be the case that the more aggressive change in mass flow rate is not feasible for the pump, but the pumps considered in this work are idealised to control mass flow rate perfectly. Note that the mass flow rate is changing long after the perturbation. This is due to the changing temperature of molten salt in the hot tank, which is not controlled in this process.

To improve the controller optimisation, other objectives should be considered. A simple time-domain metric that could quantify oscillations is to evaluate the number of times the derivative of $\Delta\dot{W}$ goes through zero. This could be treated as a constraint or as an objective function. The integral of time multiplied by the absolute error or the integral of time multiplied by the squared error could be considered in the objective. Including time in the integral will promote faster settling times, which may eliminate persistent oscillations. A parallel MEng project investigated some of these alternative objectives, and their analysis showed that the integral time absolute error and integral time squared error could lead to reductions in the maximum error observed. The overall conclusions were that the four integral objectives generated very similar controller designs, but this could have been due to the limited computational and time budget for the project.

Another limitation of this controller optimisation approach is that only a single perturbation is considered. This could lead to over-fitting of the controller, which may then perform poorly for other feasible perturbations. This could be avoided by testing each design on a number of different perturbation scenarios, similar to reliability modelling of renewable systems [10].

6.3.1 Initialisation Behaviour

To assess the effectiveness of the initial stabilisation period at negating any impacts of initialisation on the optimisation results, the turbine inlet conditions before the perturbation at $t = 3{,}000$s are compared for the initial and optimised designs in figure 6.6.

After the controllers are activated at $t = 100$s, they bring the system to a fixed operating point before the perturbation begins, despite different controller designs and the model being initialised far from the fixed operating point. This suggests the method of allowing a stabilisation period is sufficient to decouple the optimisation results from the initialisation of the model.

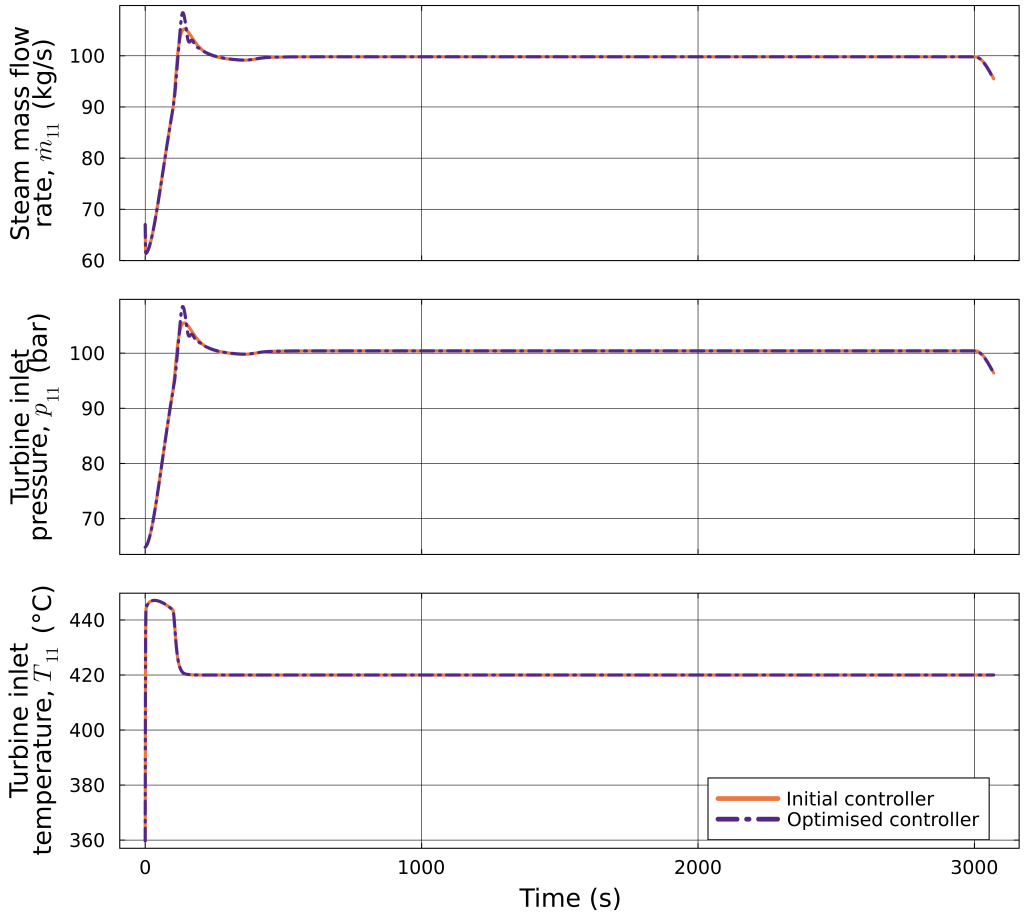

Figure 6.6: Process dynamics at the turbine inlet between initialisation and the perturbation for the initial and optimised proportional-integral-derivative controller designs in Case Study I.

6.4 Summary

In this chapter, a dynamic model of a power conversion system with a constant, ideal heat source is presented. The thermal power is delivered to a molten salt thermal energy storage system, which can then deliver heat on demand to a steam Rankine cycle for variable turbine loads. Proportional-integral and proportional-integral-derivative controllers are used to regulate the thermal energy storage system and the steam Rankine cycle.

The developed optimisation methodology for dynamic process models is applied to optimise the three tuning parameters of a proportional-integral-derivative controller that controls the turbine power output using a sliding pressure control scheme. For a simulation of a -10% change in power demand, the mean absolute error is used to quantify the performance of the controller. A 97% reduction in the objective is achieved by optimisation relative to the original hand-tuned controller. Non-steady state initialisation is deemed to not impact the optimisation by allowing for a suitable settling time before the perturbation.

The optimised controller does introduce higher frequency oscillations, at a smaller amplitude, than the original controller. The objective function does not penalise oscillations, leading to strong integral action that causes these oscillations. This could be improved by alternative objectives that penalise oscillations increasingly with time or explicitly count oscillations. This methodology is also susceptible to over-fitting to the single perturbation considered.

Chapter 7

Case Study II & III: Optimisation of the Thermal Energy Storage System Size with Pulsed Heat Sources

Publications based on this chapter:
"*Optimization of the Power Conversion System for a Pulsed Fusion Power Plant with Multiple Heat Sources using a Dynamic Process Model*" [276]
"*A dynamic model of a power conversion system with indirect thermal energy storage for a pulsed fusion tokamak for use in design optimisation*" [275]

As discussed in section 2.2, there is uncertainty in how prototype tokamaks will be operated. Steady state operation is desirable, allowing for power conversion system designs akin to traditional fossil fuel thermal power plants. More conservative designs may instead plan for pulsed operation, which will reduce the technology requirements for the tokamak control [102] at the cost of additional challenges for the power conversion system. As highlighted by the EU-DEMO pre-conceptual design of the power conversion system, the main design objectives are avoiding disconnections from the grid during dwells and minimising component damage from thermal transients.

The decay thermal power available during dwell is estimated to be 1% of the maximum pulse power, so auxiliary energy sources are required to maintain enough

electric power generation to avoid grid disconnections. Proposed dwell energy sources include fossil fuel heaters, electric heaters and thermal energy storage [109]. Sensible thermal energy storage, based on molten salts, was identified as the most promising option for the EU-DEMO reference design, citing the industrial experience with the technology from concentrated solar thermal power plants [21].

With a pulsed heat supply, the charging and discharging of the thermal energy storage system must be balanced between the pulse and dwell periods for long-term operability. The thermal energy system must also be sized to match the thermal energy requirements. To reduce the size, and hence cost, of the storage, reduced electric power generation can be considered during a dwell. The EU-DEMO reference design considers a dwell electrical power output of $\sim 10\%$ of the maximum output, chosen as the minimum to maintain grid synchronicity and maintain acceptable thermal stresses [21]. Earlier concept designs considered the other extreme of generating more power during dwells than during pulses [109]. The dwell power is clearly a key decision for the power conversion system design.

The heat generated in a fusion tokamak is distributed to multiple components, as shown in figure 2.4, each of which will require a dedicated coolant loop. The pumping power of these coolant loops, together with the load of other tokamak sub-systems, like cryogenic cooling, plasma heating and current drive systems, result in significant parasitic loads [1]. Fully utilising all available thermal power from the tokamak efficiently is therefore a requirement for net electricity generation.

The temperatures and thermal powers available from each coolant stream will differ significantly and must be integrated into the power conversion system sequentially with temperature for maximum heat utilisation. Some streams may have low enough temperatures that they are incompatible with some energy storage technologies, such as molten salts with an operating range of >290°C [178], requiring either direct integration into the power cycle or alternative energy storage systems. The integration is further complicated when all of these sources are considered to be pulsed.

In this chapter, the dynamic process models and optimisation methodology

presented in chapters 4 and 5 are used to investigate the trade-off between the thermal energy storage system size and the fraction of electrical power generated during a dwell for a fusion power plant with either a single pulsed heat source or three pulsed heat sources.

7.1 Power Conversion System Design for a Pulsed Heat Source

Based on the reference EU-DEMO design [21], the design of the power conversion system developed to utilise a single pulsed heat source is shown in figure 7.1. A sensible thermal energy storage system using molten salt delivers thermal energy to a steam Rankine cycle. A two-tank design is considered, as in chapter 6, as it enables the operation of the Rankine cycle to be decoupled from the heat source. The tanks have design temperatures of 580°C and 300°C and are kept at 1bar.

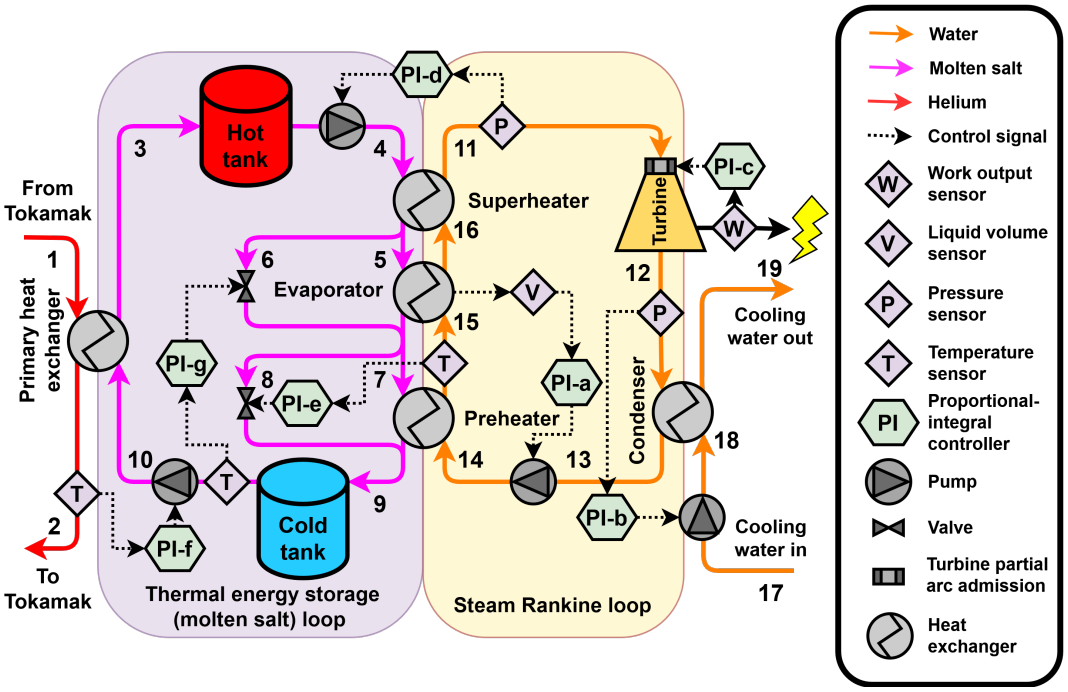

Figure 7.1: Power conversion system for a pulsed heat source from a fusion tokamak in Case Study II. Thermal energy from the tokamak is delivered to a molten salt thermal energy storage loop. Hot molten salt can then be pumped on demand to drive a steam Rankine cycle to meet load demands, even during a dwell when minimal thermal power is available from the tokamak.

The main heat source of a fusion tokamak is typically the breeding blanket, with

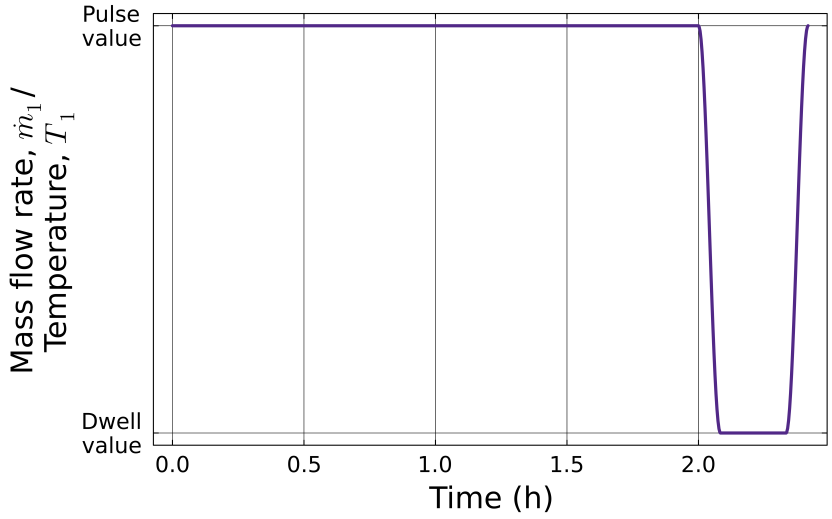

Figure 7.2: A full cycle of the mass flow rate and temperature of helium coolant from the tokamak with respect to their pulse and dwell values in Case Study II.

the largest heat split and temperatures [1]. During a pulse, 1,350MWth is available from the 1,732kg/s of 600°C helium coolant exiting the tokamak at the target return conditions of 450°C, 81bar. The return temperature is not fixed in the simulation, so the actual thermal power delivered to the molten salt is also not fixed. The return temperature is regulated by a controller to meet the target temperature, as discussed in chapter 3. During a dwell, the mass flow rate of helium from the tokamak drops to 1% of its pulse value, and the temperature drops to 450°C. The mass flow rate does not go to zero as the heat exchanger model presented in section 4.2.7.2 cannot handle zero flow, but this flow can be justified as the flow required to remove residual and decay heat from the tokamak, estimated as $\sim 1\%$ of the thermal power during a pulse. These values are estimated from early design values for STEP and EU-DEMO [110, 1].

The power profile of the tokamak is estimated from EU-DEMO pre-concept values [21]: a pulse of 2h, a ramp-down of 5min, a dwell of 15min and a ramp-up of 5min. Both the mass flow rate and temperature of helium coolant follow this profile between their pulse and dwell values, with the ramp-up and ramp-down being a C^1-continuous sigmoid curve using a cosine function. A full tokamak operating cycle is shown in figure 7.2.

Cold molten salt is pumped through the primary heat exchanger, where it is

heated by helium from the tokamak and then stored in the hot molten salt tank. This is then pumped to the superheater, evaporator and preheater in series, counter-current to the water flow in the steam Rankine cycle. There are two bypasses for the molten salt around either the evaporator or the preheater, each regulated by a valve. The molten salt is then stored in the cold molten salt tank.

Water is pumped through the preheater, evaporator and superheater in series, exiting as superheated steam that drives the turbine to produce work. The spray valve bypass considered in chapter 6 is omitted here to reduce model complexity. A suitably sized superheater is found to maintain a consistent temperature difference between the molten salt in the hot tank and the steam entering the turbine. The flow exiting the turbine is condensed to saturated water by a condenser before being pumped around the cycle again. The condenser is cooled by a source of 15°C water that exits the condenser at 2bar. An additional assertion is included in the model that the outlet water of the preheater must not enter the two-phase region after the first 1,000s of the simulation, allowing for the possibility of large transients after initialisation that need time to settle.

7.1.1 Control System Design

All seven controllers shown in figure 7.1 are proportional-integral controllers. Derivative action is not included to reduce the tuning complexity in the main optimisation problem, but is considered for a singled controller in section 7.3.3.

Two controllers (PI-a, PI-b) are implemented as described in section 6.1.1, with set-points of 0.5 and 0.0508bar respectively.

A controller (PI-c) manipulates the partial arc admission valve opening to control the turbine power output. This is a faster control method compared to the sliding pressure control method used in chapter 6, and improves thermal efficiency of the cycle by maintaining a higher mean temperature of heat supply [150, 126]. The set-point of this controller is discussed in section 7.2. Another controller (PI-d) can then control the pressure at the turbine inlet to a set-point of 165.5bar by manipulating the flow of hot molten salt to the superheater, and hence the thermal power delivered to the steam Rankine cycle. In section 7.3.3, PI-d is replaced by a

proportional-integral-derivative controller (PID-a).

A controller (PI-e) manipulates the molten salt bypass valve opening around the preheater to control the temperature of water leaving the preheater. The purpose of this is to prevent the water undergoing phase change before the evaporator, which is undesirable for the real operation of power plants and would violate the modelling assumptions presented in section 4.2.7.2. The set-point is 341°C.

A controller (PI-f) manipulates the mass flow rate of cold molten salt to the primary heat exchanger to control the return temperature of helium coolant to the tokamak at a set-point of 450°C. During a dwell, this minimises cold molten salt being pumped into the hot tank. The flow does not go to zero due to the heat exchanger model's limitations.

The final controller (PI-g) manipulates the opening of the other molten salt bypass around the evaporator to control the temperature of the molten salt in the cold tank. Without this controller, the temperature of the cold tank would drift throughout operation due to heat losses or temperature differences with the molten salt entering the tank. This would in turn cause the temperature of the hot tank to drift. This could eventually lead to issues such as molten salt freezing or temperatures beyond the molten salt's stable range in a real system, although these processes aren't modelled here. This is seen and discussed further in section 7.6.1. The set-point is the cold tank's design operating temperature, 300°C.

The controllers are all hand-tuned by repeated simulations until the controller action is deemed acceptable. To reduce the dimension of the search space and the number of objective functions, optimisation of controllers is not considered initially. Preliminary results from the simultaneous optimisation of the process design and controller design are presented in section 7.3.3. This gives an insight into the potential impact of the hand-tuning on the optimisation results. Detailed controller parameters are given in appendix B.

7.1.2 Model Parameters

The parameters used for the dynamic process model are presented in appendix B. To more rigorously set the required model parameters relative to chapter 6, a steady

state model is used to calculate parameters such as heat exchange areas and design mass flow rates based on design values for key temperatures and pressures. The model is presented in algorithm 2. It designs for a constant turbine power output. The turbine design mass flow rate is scaled by a factor of 1.2 to allow controller PI-c to bring the power output to its set point.

The heat exchangers are modelled using the steady state models presented in section 4.2.8, and the molten salt tanks are considered to have a fixed temperature of stored molten salt. The turbine inlet is chosen as the tear stream where the calculations can begin for the Rankine loop as the temperature and pressure there are typically specified as design variables.

The ramp-up and ramp-down periods are considered as step changes between pulse and dwell at the middle of their time period. During a dwell, no thermal power is delivered to the molten salts.

Note the fluid enthalpy change across the pump is calculated based on the isentropic efficiency $\eta_{\text{is,pump}}$ instead of the efficiency η used in equation 4.30. Equation 4.30 predicts pump powers $\sim 20\%$ greater than the steady state model if $\eta = \eta_{\text{is,pump}}$. However, of the parameters that are passed to the dynamic model, such as heat exchanger areas, the differences are $< 0.1\%$, so this discrepancy should not impact the use of the steady state model for parameterisation.

The parameter values used for both the steady state and dynamic process models are presented in appendix B. The initialisation equations, chosen by trial and error until the model could be initialised reliably, are also given in appendix B.

7.2 Case Study II: Optimisation Problem Statement

To investigate the design implications of reduced dwell power, the set-point of controller PI-c is chosen to be pulsed in sync with the tokamak, using the same interpolation curve between its pulse and dwell values. This is established as a valid loading strategy for EU-DEMO for rapid changes in steam flow rate [20]. During a pulse, the set-point value is the design pulse power $\dot{W}_{\text{pulse}}$, and during a dwell, the set-point is $\varphi\dot{W}_{\text{pulse}}$, where φ is the dwell power fraction. Changing the power

Algorithm 2 Steady state model of the power conversion system in Case Study II. The results are used to parameterise the dynamic process model. PHX, SH, EVP, PH, CND are the primary heat exchanger, superheater, evaporator, preheater and condenser. Medium thermodynamic properties are calculated using CoolProp if the equation is not specified [24].

Require: U_{PH}, U_{EVP}, U_{SH}, U_{CND}, U_{PHX}, $\dot{Q}_{\mathrm{helium}}$, $p_{\mathrm{turbine,in}}$, p_{CND}, $T_{\mathrm{ms,hot}}$, $T_{\mathrm{ms,cold}}$, $c_{\mathrm{p,ms}}$, ρ_{ms}, τ_{pulse}, τ_{dwell}, $T_{\mathrm{helium,in}}$, $T_{\mathrm{helium,out}}$, $c_{\mathrm{p,helium}}$, $T_{\mathrm{cw,in}}$, $T_{\mathrm{cw,out}}$, $p_{\mathrm{cw,out}}$, $\Delta T_{\mathrm{pinch}}$, $\Delta p_{\%}$, $\eta_{\mathrm{is,turbine}}$, $\eta_{\mathrm{is,pump}}$

Primary heat exchanger (PHX) calculations during pulse:

$\Delta T_{\mathrm{LMTD,PHX}} \leftarrow \frac{(T_{\mathrm{helium,out}}-T_{\mathrm{ms,cold}})-(T_{\mathrm{helium,in}}-(T_{\mathrm{ms,hot}}+\Delta T_{\mathrm{pinch}}))}{\ln(T_{\mathrm{helium,out}}-T_{\mathrm{ms,cold}})-\ln(T_{\mathrm{helium,in}}-(T_{\mathrm{ms,hot}}+\Delta T_{\mathrm{pinch}}))}$ ▷ Eq. 4.21a

$A_{\mathrm{PHX}} \leftarrow \frac{\dot{Q}_{\mathrm{helium}}}{U_{\mathrm{PHX}}\Delta T_{\mathrm{LMTD,PHX}}}$ ▷ Eq. 4.50

$\dot{m}_{\mathrm{helium}} \leftarrow \frac{\dot{Q}_{\mathrm{helium}}}{c_{\mathrm{p,helium}}(T_{\mathrm{helium,in}}-T_{\mathrm{helium,out}})}$

$\dot{m}_{\mathrm{ms,PHX}} \leftarrow \frac{\dot{Q}_{\mathrm{helium}}}{c_{\mathrm{p,ms}}((T_{\mathrm{ms,hot}}+\Delta T_{\mathrm{pinch}})-T_{\mathrm{ms,cold}})}$

Iterate over Rankine loop to estimate balance between pulse and dwell:

$\dot{W}_{\mathrm{turbine}} \leftarrow 1\mathrm{MW}$

while *true* **do**

  $T_{\mathrm{turbine,in}} \leftarrow T_{\mathrm{ms,hot}} - \Delta T_{\mathrm{pinch}}$

  $h_{\mathrm{turbine,in}} \leftarrow h(p_{\mathrm{turbine,in}}, T_{\mathrm{turbine,in}})$

  $\Delta h_{\mathrm{is,turbine}} \leftarrow h(p_{\mathrm{CND}}, s(p_{\mathrm{turbine,in}}, T_{\mathrm{turbine,in}})) - h_{\mathrm{turbine,in}}$ ▷ Eq. 4.26

  $h_{\mathrm{turbine,out}} \leftarrow h_{\mathrm{turbine,in}} + \eta_{\mathrm{is,turbine}}\Delta h_{\mathrm{is,turbine}}$ ▷ Eq. 4.27

  $\dot{m}_{\mathrm{rankine}} \leftarrow \frac{\dot{W}_{\mathrm{turbine}}}{h_{\mathrm{turbine,in}}-h_{\mathrm{turbine,out}}}$ ▷ Eq. 4.28

  $h_{\mathrm{CND,out}} \leftarrow h_{\mathrm{saturated,L}}(p_{\mathrm{CND}})$

  $T_{\mathrm{CND,out}} \leftarrow T_{\mathrm{saturated}}(p_{\mathrm{CND}})$

  $\dot{Q}_{\mathrm{CND}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{CND,out}} - h_{\mathrm{turbine,out}})$

  $p_{\mathrm{pump,out}} \leftarrow \frac{p_{\mathrm{turbine,in}}}{(1-\Delta p_{\%})^2}$

  $\Delta h_{\mathrm{is,pump}} \leftarrow h(p_{\mathrm{pump,out}}, s(p_{\mathrm{CND}}, h_{\mathrm{pump,in}})) - h_{\mathrm{pump,in}}$ ▷ Eq. 4.26

  $h_{\mathrm{pump,out}} \leftarrow h_{\mathrm{CND,out}} + \eta_{\mathrm{is,pump}}\Delta h_{\mathrm{is,pump}}$

  $T_{\mathrm{pump,out}} \leftarrow T(p_{\mathrm{pump,out}}, h_{\mathrm{pump,out}})$

  $T_{\mathrm{PH,out}} \leftarrow T_{\mathrm{saturated,L}}(p_{\mathrm{pump,out}}(1-\Delta p_{\%})) - \Delta T_{\mathrm{pinch}}$

  $h_{\mathrm{PH,out}} \leftarrow h(p_{\mathrm{pump,out}}(1-\Delta p_{\%}), T_{\mathrm{PH,out}})$

  $\dot{Q}_{\mathrm{PH}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{PH,out}} - h_{\mathrm{pump,out}})$

  $h_{\mathrm{EVP,out}} \leftarrow h_{\mathrm{saturated,V}}(p_{\mathrm{pump,out}}(1-\Delta p_{\%}))$

  $T_{\mathrm{EVP,out}} \leftarrow T_{\mathrm{saturated}}(p_{\mathrm{pump,out}}(1-\Delta p_{\%}))$

  $Q_{\mathrm{EVP}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{EVP,out}} - h_{\mathrm{PH,out}})$

  $Q_{\mathrm{SH}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{turbine,in}} - h_{\mathrm{EVP,out}})$

  $\dot{m}_{\mathrm{ms,SH}} \leftarrow \frac{\dot{Q}_{\mathrm{PH}}+\dot{Q}_{\mathrm{EVP}}+\dot{Q}_{\mathrm{SH}}}{c_{\mathrm{p,ms}}(T_{\mathrm{ms,hot}}-T_{\mathrm{ms,cold}})}$

  if $\tau_{\mathrm{pulse}}(\dot{m}_{\mathrm{ms,PHX}} - \dot{m}_{\mathrm{ms,SH}}) \leq \tau_{\mathrm{dwell}}\dot{m}_{\mathrm{ms,SH}}$ **then**

    break

  end if

  $\dot{W}_{\mathrm{turbine}} \leftarrow \dot{W}_{\mathrm{turbine}} + 1\mathrm{MW}$

end while

Calculate heat exchanger areas and cooling water flow:
$\Delta T_{\text{LMTD,CND}} \leftarrow \frac{(T_{\text{CND,out}}-T_{\text{cw,out}})-(T_{\text{CND,out}}-T_{\text{cw,in}})}{\ln(T_{\text{CND,out}}-T_{\text{cw,out}})-\ln(T_{\text{CND,out}}-T_{\text{cw,in}})}$ ▷ Eq. 4.21a
$A_{\text{CND}} \leftarrow \frac{-\dot{Q}_{\text{CND}}}{U_{\text{CND}}\Delta T_{\text{LMTD,CND}}}$
$\dot{m}_{\text{cw}} \leftarrow \frac{-\dot{Q}_{\text{CND}}}{h(p_{\text{cw,out}},T_{\text{cw,out}})-h(p_{\text{cw,out}},T_{\text{cw,in}})}$
$T_{\text{ms,SH,out}} \leftarrow T_{\text{ms,hot}} - \frac{\dot{Q}_{\text{SH}}}{c_{\text{p,ms}}\dot{m}_{\text{ms,SH}}}$
$\Delta T_{\text{LMTD,SH}} \leftarrow \frac{(T_{\text{ms,hot}}-T_{\text{turbine,in}})-(T_{\text{ms,SH,out}}-T_{\text{EVP,out}})}{\ln(T_{\text{ms,hot}}-T_{\text{turbine,in}})-\ln(T_{\text{ms,SH,out}}-T_{\text{EVP,out}})}$ ▷ Eq. 4.21a
$A_{\text{SH}} \leftarrow \frac{\dot{Q}_{\text{SH}}}{U_{\text{SH}}\Delta T_{\text{LMTD,SH}}}$
$T_{\text{ms,EVP,out}} \leftarrow T_{\text{ms,SH,out}} - \frac{\dot{Q}_{\text{EVP}}}{c_{\text{p,ms}}\dot{m}_{\text{ms,SH}}}$
$\Delta T_{\text{LMTD,EVP}} \leftarrow \frac{(T_{\text{ms,SH,out}}-T_{\text{EVP,out}})-(T_{\text{ms,EVP,out}}-T_{\text{PH,out}})}{\ln(T_{\text{ms,SH,out}}-T_{\text{EVP,out}})-\ln(T_{\text{ms,EVP,out}}-T_{\text{PH,out}})}$ ▷ Eq. 4.21a
$A_{\text{EVP}} \leftarrow \frac{\dot{Q}_{\text{EVP}}}{U_{\text{EVP}}\Delta T_{\text{LMTD,EVP}}}$
$\Delta T_{\text{LMTD,PH}} \leftarrow \frac{(T_{\text{ms,EVP,out}}-T_{\text{PH,out}})-(T_{\text{ms,cold}}-T_{\text{pump,out}})}{\ln(T_{\text{ms,EVP,out}}-T_{\text{PH,out}})-\ln(T_{\text{ms,cold}}-T_{\text{pump,out}})}$ ▷ Eq. 4.21a
$A_{\text{PH}} \leftarrow \frac{\dot{Q}_{\text{PH}}}{U_{\text{PH}}\Delta T_{\text{LMTD,PH}}}$
Estimate size of molten salt tanks, allowing for fill constraints
$V_{\text{tank}} \leftarrow 1.25\frac{\tau_{\text{dwell}}\dot{m}_{\text{ms,SH}}}{\rho_{\text{ms}}}$
return $\dot{m}_{\text{cw}}$, $\dot{m}_{\text{helium}}$, $\dot{m}_{\text{ms,PHX}}$, $\dot{m}_{\text{ms,SH}}$, $\dot{m}_{\text{rankine}}$, A_{PHX}, A_{SH}, A_{EVP}, A_{PH}, A_{CND}, V_{tank}

generation during a dwell will change the amount of hot molten salt that must be stored during a pulse. The design vector is chosen to be $\boldsymbol{x} = \left(V_{\text{tank}}, \dot{W}_{\text{pulse}}, \varphi\right)^{\text{T}}$, where V_{tank} is the volume of each molten salt tank. The bi-objective problem statement is given by:

$$\begin{aligned} &\min_{\boldsymbol{x}\in X} \begin{bmatrix} V_{\text{tank}} \\ -\phi \end{bmatrix} \\ &\text{s.t. } \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0} \\ &\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \end{aligned} \tag{7.1}$$

As before, $\boldsymbol{g}(\boldsymbol{x})$, $\boldsymbol{h}(\boldsymbol{x})$ encapsulate the equations of the dynamic process model. The domain of the search space X is given in table 7.1. The radius and height of the tanks are calculated from V_{tank} such that the surface area is minimised.

While the objective functions do not contain any time-dependent variables, a dynamic simulation of each design must still be performed to verify if the design is feasible and suitably balances charging and discharging of molten salt between pulse and dwell. This is assessed by running the simulation for $t_{\text{sim}} = 68{,}100\text{s}$, covering an

Design variable	Minimum value	Maximum value
V_{tank}	1,000m^3	6,000m^3
$\dot{W}_{\text{pulse}}$	400MW	550MW
φ	0.25	1

Table 7.1: Domains of the design variables for equation 7.1.

Simulation Outcome	Infeasibility, g	Reasoning
Completed successfully	0	Molten salt usage is suitably balanced
Simulation failed after initialisation	$1 - \frac{t_{\text{fail}}}{t_{\text{sim}}}$	Molten salt is being over-/under-utilised - later simulation failures imply that designs are closer to balancing the molten salt usage and becoming feasible
Simulation failed before initialisation	10^3	Error may not reflect on the design if due to coupling between OpenModelica and Julia
Simulation evaluation time exceeded 300s	10^6	Undesirable behaviour, such as oscillations, are likely slowing the simulation

Table 7.2: Infeasibility assignments for the different possible simulation outcomes in Case Study II. t_{fail} is the internal simulation time at failure. It is assumed that all mid-simulation failures are caused by the molten salt tank model assertion in equation 4.34.

initial pulse period, allowing for transients from initialisation to settle, followed by 7 full cycles. This approach has been used for the operability assessment of a DEMO power conversion system [28]. If a design completes the simulation successfully without failing due to equation 4.34, it suggests the design balances charge and discharge well or has oversized thermal energy storage. Designs of the latter type should be selected out by the optimiser.

The infeasibility g of a design is calculated according to table 7.2. Only the ranked order of infeasible designs with respect to g impact selection in Fresa, not the magnitudes of g. Unknown errors, such as those mentioned in section 5.2.3, are considered as less infeasible than time-outs as the former cannot be found to reflect any attribute of the design, while the latter often represents undesirable oscillatory behaviour in the simulation.

Recall that the steady state model in section 7.1.2 designs for a constant turbine

power output. In the dynamic model, the turbine power set-points and tank volumes are entirely determined by the optimiser. Hence, the sizing of the heat exchangers by the steady state model will not necessarily be well suited for a particular design. This could be addressed by integrating the steady state and dynamic models into a two-step objective function. It could run the steady state model for a fixed turbine power, provided by the optimiser, and then pass its outputs to the dynamic model as parameters. The steady state model could also be extended to consider reduced power output during a dwell.

Three optimisers are used to solve equation 7.1. The first two use Fresa as the optimiser, each with identical parameters except one uses the non-dominated fitness ranking method (optimiser 1) [70] and the other uses the Hadamard fitness ranking method (optimiser 2) [95]. The third approach uses the Cocoa multi-agent optimisation system to utilise two Fresa instances with both ranking methods simultaneously (optimiser 3) [96]. The parameters used for Fresa, or the instances of Fresa, in each optimiser are presented in table 7.3. The initial design provided to all the optimisers is $\boldsymbol{x} = \left(2{,}125\text{m}^3, 455\text{MW}, 1\right)^{\text{T}}$, calculated by the steady state model. The simulation of this design does not complete successfully due to the hot molten salt tank overflowing, highlighting the need for the dynamic model to assess design feasibility.

The number of loops of the reinitialisation script, discussed in section 5.2.3, for optimiser 1 and 2 is 100. Cocoa tracks the number of communication ticks, which is related to function evaluations by slightly more than 2 ticks per evaluation due to other messages between agents. These ticks are shared across all agents, but not necessarily equally. For optimiser 3, the number of reinitialisation loops is 200, and each loop has 200 ticks.

A $27 \times 27 \times 27$ point grid search is also performed for comparison.

7.3 Case Study II: Optimisation Results

The non-dominated set found by optimiser 3 after 18,813 objective function evaluations is presented in figure 7.3 with a comparison to the grid search. A good

Parameter	Optimiser 1	Optimiser 2	Optimiser 3	Description
$n_{\text{f,max}}$	200	200	Controlled by Cocoa	Maximum number of evaluations for the algorithm
n_{p}	(10,30)			Maximum number of designs to propagate each generation
fitness	*:nondominated*	*:hadamard*	*:nondominated* or *:hadamard*	Fitness ranking method
elite	*true*	*true*	*false*	Non-dominated solutions are always kept in the population
archive-elite	*true*	*true*	*false*	Save pruned non-dominated solutions
issimilar	*Fresa.isimilarx*			Similarity function
ε	0.01			Similarity threshold
multi-threading	*true*			Evaluate designs in parallel

Table 7.3: Parameters of Fresa used by each optimiser to solve equation 7.1 in Case Study II. All other parameters are left as default. Please refer to `https://www.ucl.ac.uk/~ucecesf/fresa.html` for default values and more detailed explanations.

distribution of designs is seen in the objective space, with the sparsest region being around $\varphi = 0.9$. The grid search is too coarse to suitably search the design space with the requirement of balanced molten salt utilisation, with the design at $\varphi = 1$ having a value of $V_{\text{tank}} > 50\%$ higher than the design found by Cocoa. It is able to find two designs that are non-dominated by the final population of the Cocoa optimiser, but it is clearly a poor approximation of the Pareto front by comparison.

As a benchmark for the bi-objective optimisation problem, a separate single-objective optimisation is performed, using Fresa only, equivalent to equation 7.1 with fixed $\varphi = 1$. The best design found after 25,093 objective evaluations is $D_{\text{SO}} = \left(2{,}003\text{m}^3, 468.0\text{MW}, 1\right)^{\text{T}}$, which is in good agreement with design D_3.

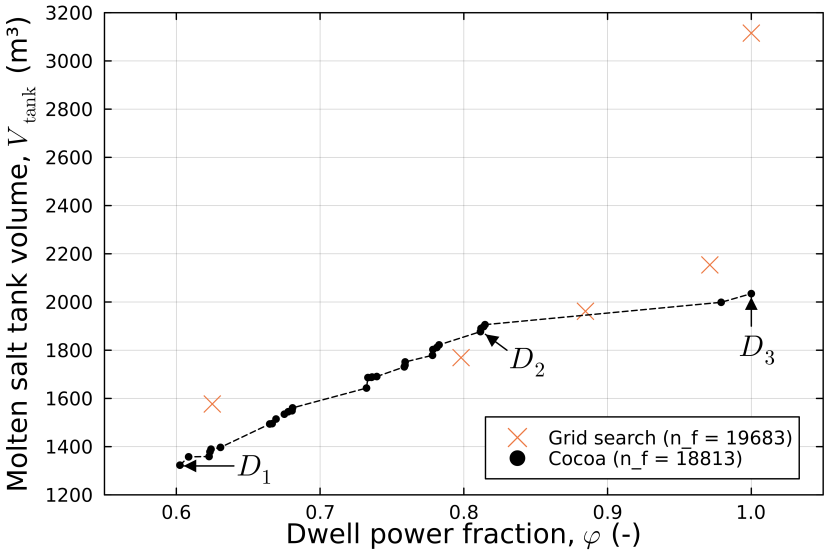

Figure 7.3: Non-dominated set at the end of the optimisation by optimiser 3 in Case Study II. Three designs $D_1 = \left(1{,}323\text{m}^3, 494.4\text{MW}, 0.6025\right)^{\text{T}}$, $D_2 = \left(1{,}877\text{m}^3, 480.1\text{MW}, 0.8116\right)^{\text{T}}$ and $D_3 = \left(2{,}035\text{m}^3, 467.9\text{MW}, 1\right)^{\text{T}}$ are highlighted for further discussion in the text. Results from a grid search with comparable function evaluations is included for comparison.

The dynamic behaviour of the hot molten salt storage and the turbine power output for the highlighted designs D_1, D_2, D_3 are shown in figures 7.4, 7.5 and 7.6 respectively. As the simulations cover 7 full tokamak cycles, the volume of stored hot molten salt can be seen to be quite stable for each design after each cycle. As the total supply of hot helium coolant from the tokamak is independent of the designs, each design must balance its time-averaged energy change to be near zero, or the simulation will fail due to equation 4.34. Hence, as φ decreases, $\dot{W}_{\text{pulse}}$ must increase to generate more power during pulses for balanced designs. It can also be seen that the turbine power controller (PI-c) performs qualitatively well, maintaining steady power during pulse and dwell and following the sigmoid curve during ramp-up and ramp-down closely.

Designs D_1, D_3 utilise the available storage well, with the molten salt levels approaching the constraints imposed by 4.34. Note that as the model is initialised at $V_{\text{ms,hot}}/V_{\text{tank}} = 0.1$, while the constraint is at $V_{\text{ms,hot}}/V_{\text{tank}} \geq 0.05$, the optimisation could not reduce the tank size to its theoretical minimum between the constraints.

Design D_2 shows a larger gap between the maximum volume of stored

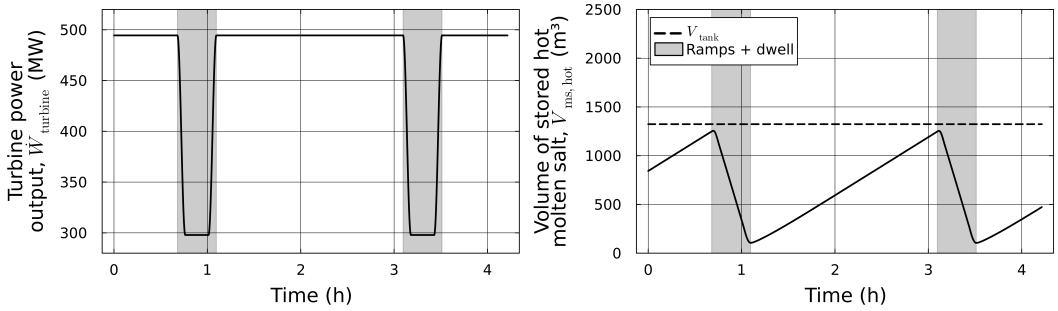

Figure 7.4: The dynamics of the turbine power output and the volume of stored hot molten salt in the simulation of design D_1 in Case Study II. A 40% reduction in the dwell power output allows for a 35% reduction in the molten salt tank sizes.

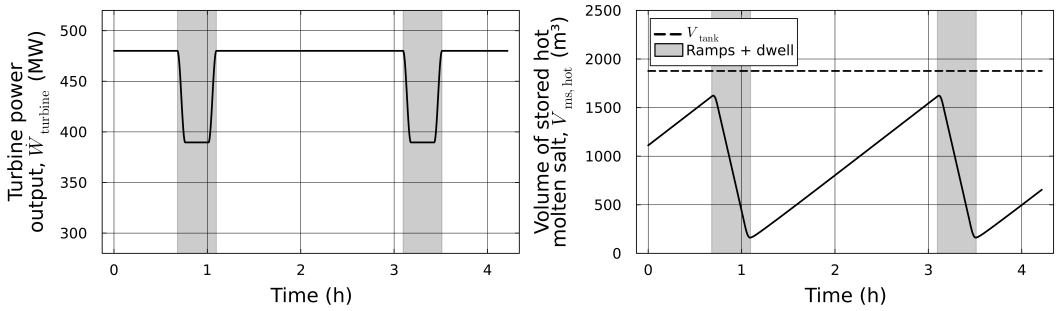

Figure 7.5: The dynamics of the turbine power output and the volume of stored hot molten salt in the simulation of design D_2 in Case Study II. This design compromises between the two objectives, although the difference between the maximum volume of stored hot molten salt and the total tank volume suggests the tank volume can be reduced further.

hot molten salt and the tank volume. Manual simulations confirm that $\boldsymbol{x} = \left(1{,}777\text{m}^3, 480.1\text{MW}, 0.8116\right)^{\text{T}}$ is a feasible design that Pareto-dominates D_2. Similarly, improved designs are found by the grid search method. These suggest that the concave shape of the non-dominated set in figure 7.3 is a result of insufficient time for the optimiser to converge towards the true Pareto front.

The minimum value of φ seen in the designs is 0.6025, while the lower bound of the design variable is 0.25. Inspecting the dynamics of design D_1, it is believed that this limit is due to the model assertion of no two-phase flow before the evaporator, discussed in section 7.1. During a dwell, the temperature of feedwater leaving the preheater is within $\sim$ 1K of the temperature of saturated water in the evaporator, as opposed to a 10K difference during a pulse. When power is reduced during a dwell, the combined action of the control system is to reduce the flow of steam to the

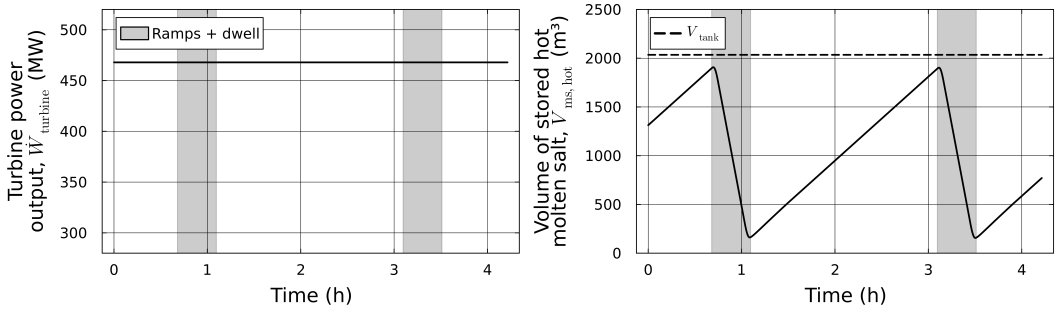

Figure 7.6: The dynamics of the turbine power output and the volume of stored hot molten salt in the simulation of design D_3 in Case Study II. This design achieves a steady-state power output at the expense of the maximum thermal energy storage system size.

turbine while maintaining the same inlet pressure. As the flows of water and molten salt through the heat exchangers drops, the heat exchangers, which are sized for larger flows, bring the fluids closer to the temperature of the other. During a dwell, the molten salt preheater bypass valve is fully open for D_1, allowing half the molten salt flow to bypass the preheater. This then leads to low molten salt temperatures of 170°C leaving the preheater, below the freezing point of the molten salt. Hence, the proposed control system is not suitable and should be redesigned to allow for reductions in the dwell power fraction further.

7.3.1 Performance of the Robust Logarithmic Mean Temperature Difference Method

The robust logarithmic mean temperature difference [9], introduced in section 4.2.7.2, extends the standard logarithmic mean temperature difference with additional equations such that it is defined and C^0-continuous for all possible inlet and outlet temperature combinations. This is better suited for dynamic simulations where transients could result in temporary temperature crossings.

The heat exchanger that undergoes the largest transients is the primary heat exchanger due to the 150°C change in the helium inlet temperature and the 99% changes in mass flow rates between pulses and dwells. Hence, the analysis of the robust logarithmic mean temperature difference method is limited to this heat exchanger.

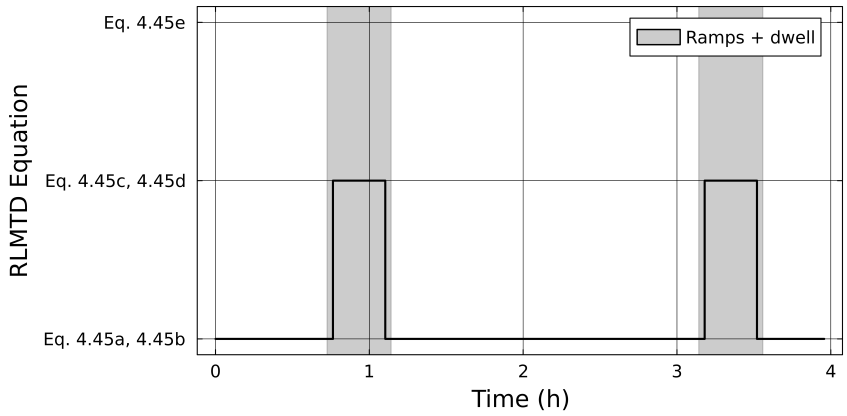

Figure 7.7: Usage of the different robust logarithmic mean temperature difference equations 4.45a-4.45e during the simulation of design D_3 on the helium side of the primary heat exchanger in Case Study II. Non-physical equations are only utilised during dwells.

Inspecting dynamic simulations of design D_3, the specific equation of equations 4.45a-4.45e used to calculate the heat exchange on the helium side of the primary heat exchanger is presented in figure 7.7. The plots are similar for designs D_1, D_2 and for the molten salt side of all designs. The non-physical equations 4.45c-4.45e are only utilised during dwells.

During a dwell, the mass flow rate of helium goes to a minimum of 17.3kg/s, with an inlet temperature of 450°C, and controller PI-f brings down the flow of cold molten salt to its minimum of 31.1kg/s. The cold tank temperature is $\sim$ 300°C due to controller PI-g. The non-physical equations result in the molten salt outlet temperature exceeding the helium inlet temperature, exiting at $\sim$ 461°C. The total heat flow rate from the helium to the molten salt is $\sim$ 7.46MW during a dwell, resulting in a minimum helium outlet temperature of 367°C.

For comparison, the ε-NTU method introduced in section 4.1.6.1 is used to calculate the steady state heat transfer under the inlet conditions specified above. From the counter-current effectiveness in equation 4.24 and using the steady state heat transfer coefficient between fluids from equation 4.48, it is found that $\varepsilon = 1$. This is due to the large heat exchange area relative to the flow rates. As the molten salt is the lowest capacity stream, its outlet temperature is brought up to 450°C and the helium's outlet temperature is 372°C. The total heat flow rate from the helium to

the molten salt is ~ 6.98MW.

Despite the non-physical equations, the robust logarithmic mean temperature difference predicts outlet temperatures within a reasonable degree of the ε-NTU method. Relative to the pulse power of 1,350MW, the 0.5MW difference between the models should have a negligible impact on the dynamics of the rest of the process. The outlet molten salt at an impossible temperature of ~ 461°C is instantly mixed into the hot tank, which is at a higher temperature of ~ 590°C and has a large stored mass relative to the incoming flow rate. Hence, the non-physical temperature will not lead on to create other non-physical temperatures around the process.

During a dwell, it is seen that the return temperature of helium to the tokamak is not satisfied. Despite this, the flow rate is low enough that it is assumed an alternative heating systems could manage this difference due to the low flow rate, so it is not considered as a critical design failure.

From this assessment, it can be concluded that the robust logarithmic mean temperature difference is justified in its use here, and does not results in non-physical behaviour in the wider process.

7.3.2 Performance of the Multi-Agent Optimisation System

To assess the performance of the Cocoa multi-agent optimisation system, the non-dominated sets found by each optimiser from table 7.3 is presented in figure 7.8. The optimisers are compared at an approximately equal number of function evaluations for fairness.

As discussed in section 5.2.1, the Hadamard and non-dominated fitness ranking methods tend to emphasise a wider or more evenly distributed non-dominated set, respectively. This statement is supported by figure 7.8, where the range of V_{tank} and φ covered by the non-dominated set of the Hadamard optimiser is much greater than the set of the non-dominated optimiser, but there is a large gap in the non-dominated set of the former in the range $0.77 < \varphi < 0.93$.

Comparatively, the Cocoa optimiser is able to find a wider non-dominated set than the other solvers, and shows a more even distribution of points than the Hadamard optimiser around $0.6 < \varphi < 0.82$. There is still a sizeable gap in its

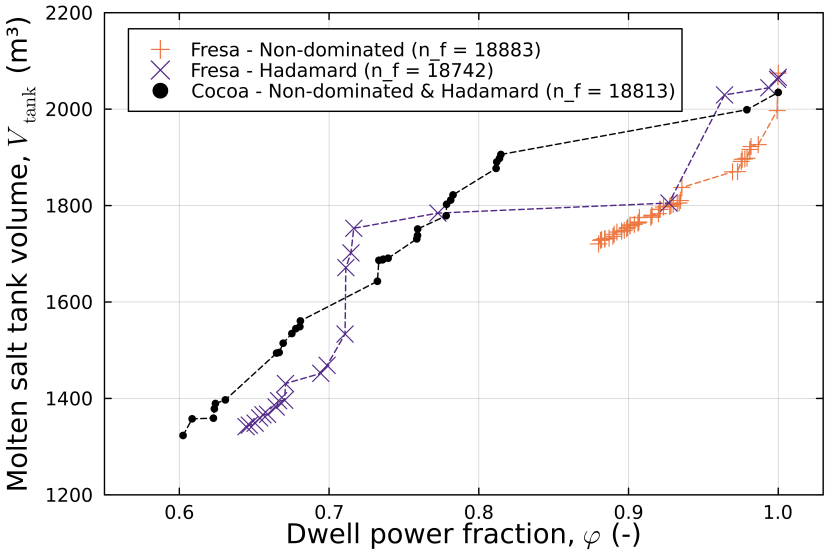

Figure 7.8: The final non-dominated sets generated by the three optimisers presented in table 7.3 in Case Study II.

non-dominated set comparable to that of the Hadamard optimiser. Considering the combined final populations of the three solvers, each solver contributes non-dominated designs in different ranges of φ. The choice of the solver for longer optimisation runs could therefore be selected based on the designer's weighting of the objectives' importance.

Besides needing more function evaluations, the performance of the Cocoa optimiser is likely being harmed by the loss of population diversity after each reinitialisation, as discussed in section 5.2.3. The impact of this cannot be quantified in this comparison unfortunately. Regardless, Cocoa is a promising system for optimisation, especially considering the combination of completely different algorithms.

7.3.3 Simultaneous Optimisation of Design & Control Variables

As highlighted in section 1.2, there is a need to integrate process and control system design for inherently transient processes. As a preliminary investigation of this, an extension of equation 7.1 is considered based on the work of chapter 6:

$$\begin{aligned} \min_{\boldsymbol{x}\in X} \quad & \begin{bmatrix} V_{\text{tank}} \\ -\phi \\ \frac{1}{t_{\text{end}}-t_{\text{start}}} \int_{t_{\text{start}}}^{t_{\text{end}}} \left| p_{11,\text{des}} - p_{11}(t) \right| \mathrm{d}t \end{bmatrix} \\ \text{s.t. } & \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0} \\ & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \end{aligned} \tag{7.2}$$

where $\boldsymbol{x} = \left(V_{\text{tank}}, \dot{W}_{\text{pulse}}, \varphi, K_{\text{c}}, \tau_{\text{I}}, \tau_{\text{D}}\right)^{\text{T}}$. In this problem, the controller PI-d has been replaced by a proportional-integral-derivative controller (PID-a) and its tuning parameters are new decision variables. This controller was chosen as there is likely to be significant inertia between the controller action and changes to the measured variable due to the large thermal inertia of the evaporator and stored water. The bounds of the decision variables are given in table 7.4. The initial design is $\boldsymbol{x} = \left(2{,}125\text{m}^3, 455\text{MW}, 1, 0.001\text{kg/(s Pa)}, 600\text{s}, 1\text{s}\right)^{\text{T}}$, which fails due to the hot tank overflowing.

Design variable	Minimum value	Maximum value
V_{tank}	500m^3	5,000m^3
$\dot{W}_{\text{pulse}}$	355MW	555MW
φ	0.25	1
K_{c}	0.0001kg/(s Pa)	0.1kg/(s Pa)
τ_{I}	1s	1000s
τ_{D}	0.1s	100s

Table 7.4: Domains of the design variables for equation 7.2.

Fresa is used to solve equation 7.2 with parameters $n_{\text{f,max}} = 100$, $n_{\text{p}} = (20, 40)$ and all other parameters as presented in table 7.3 for optimiser 1. To account for the increased dimension of the search space and the additional objective, the number of reinitialisation loops is set to 1,000. To reduce the computational time required for the optimisation, infeasibility is calculated as in table 7.7, with a maximum

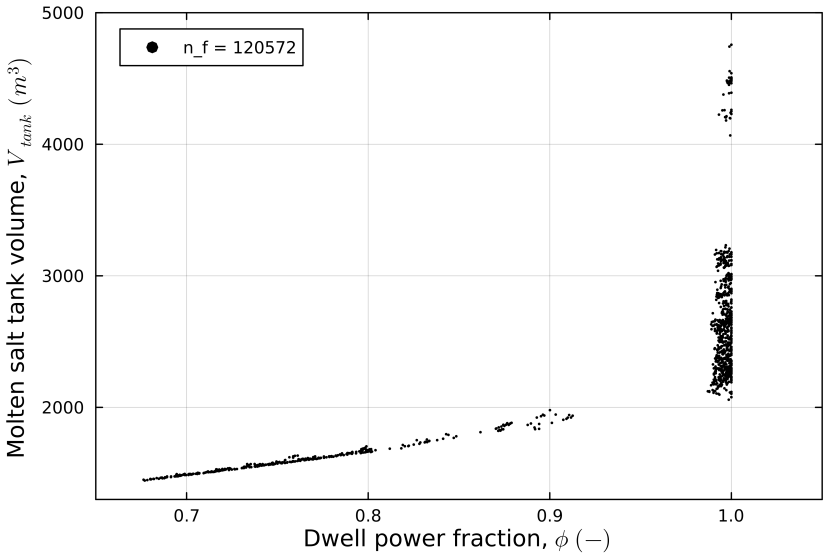

Figure 7.9: Non-dominated set for equation 7.2 found by Fresa. The set is projected onto the objective space of the first two of its three objectives.

simulation evaluation time of 120s, and the internal simulation time t_{sim} is 24,600s, encompassing three pulse periods and two dwell periods.

The final non-dominated set, projected into the objective space of V_{tank} and φ, is presented in figure 7.9 after 120,572 objective function evaluations. There is a large clustering of designs near to $\varphi = 1$, likely due to the initial design being at this value. The optimisation finds a good distribution of non-dominated designs in the range of $0.67 < \varphi < 0.92$, although there is still a gap in designs between $0.92 < \varphi < 0.98$ comparable to optimiser 1 in the main case study.

To simplify this preliminary analysis of the final population, the non-dominated set is reduced to the designs that are non-dominated in V_{tank} and φ only. This weighting of the objectives is justified by the control not needing to be the optimum, but rather good enough, which should be ensured by having the new control objective in the search. The reduced non-dominated sets at different stages of the optimisation, including for similar number of function evaluations to the main optimisation problem, is presented in figure 7.10.

At a similar number of function evaluations to the optimisers in the main case study, it can be seen that large improvements are still possible for designs with higher values of φ. The designs at lower values of φ are already comparable to the final

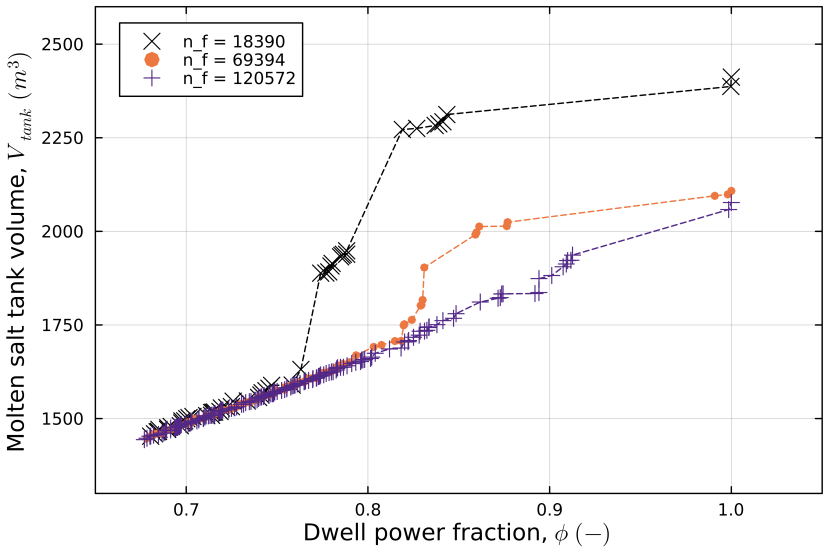

Figure 7.10: Reduced non-dominated set for equation 7.2 found by Fresa at different number of function evaluations.

non-dominated set however. By the final set, there is a good distribution of designs except the gap mentioned above.

Comparing the results to figure 7.8, designs at the extreme values of φ have V_{tank} values $< 2.1\%$ above the combined sets of the three solvers. Some of the designs at intermediate φ values are non-dominated relative to the populations of the three solvers. The range of φ in the non-dominated set is not as wide however, and does not seem to expand significantly over 100,000 additional evaluations. The additional degrees of freedom with the controller likely make it harder to find designs at lower values of φ where the model is being pushed to its limits, as discussed in section 7.3.

From this comparison, it is unclear whether the optimised control could allow the optimiser to find improved designs relative to the two shared objectives that couldn't be found by the original optimisers given more function evaluations, as noted earlier. This is not unexpected, as the overall molten salt usage would be the same for an oscillating flow rate and a constant flow rate at the time-averaged mean of the oscillations.

The dynamics of the turbine inlet pressure of the non-dominated design found by optimiser 3 at $\varphi = 1$ is compared with the non-dominated design $\boldsymbol{x} = \left(2{,}077\text{m}^3, 467.8\text{MW}, 1, 0.01347\text{kg/(s Pa)}, 606.9\text{s}, 3.038\text{s}\right)^{\text{T}}$ found in this sec-

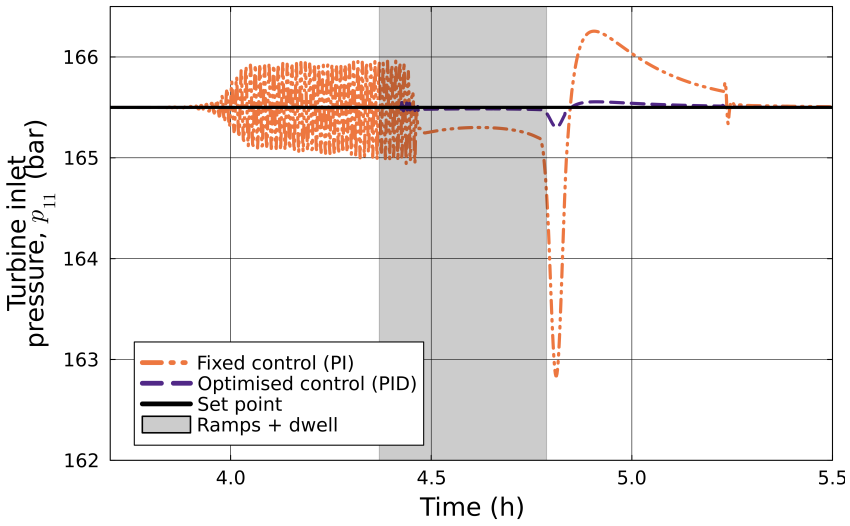

Figure 7.11: Dynamics of the turbine inlet pressure for steady state turbine power designs found by optimisation with a fixed proportional-integral controller and an optimised proportional-integral-derivative controller.

tion in figure 7.11. High-frequency oscillations seen in the model with the fixed proportional-integral controller are thought to be due to interactions between this controller and the two controllers regulating the molten salt bypasses. The optimised controller is far more stable, eliminating the high frequency oscillations and reducing the pressure drop that occurs immediately after the ramp up and the return to set point time.

This suggests the treatment of the control objective as secondary to the remaining objectives could be an effective method to incorporate simultaneous control design into process design without further complicating the final design selection of the end user. The additional function evaluations required for the additional design variables and objectives could be prohibitive when extending larger design optimisation problems however. An alternative option could be to incorporate the control objective as an inequality constraint.

7.4 Power Conversion System Design for Three Pulsed Heat Sources

Three pulsed heat sources are considered from the tokamak, representing thermal power from the blanket and outboard first wall, outboard limiter and divertor tokamak

components. The specifications of these heat sources is given in table 7.5. As in section 7.1, the temperature and mass flow rates of each source vary as a sigmoid curve between their pulse and dwell values.

Source	Coolant	Outlet temperature (°C)	Return temperature* (°C)	Return pressure (bar)	Mass flow rate (kg/s)	Thermal power available* (MW)
Blanket + outboard first wall	Helium	600 (450)	450 (450)	81 (81)	1,735 (17.35)	1,350 (0)
Outboard limiter	Helium	450 (225)	225 (225)	99 (99)	184 (1.84)	215 (0)
Divertor	Water	300 (275)	275 (275)	149 (149)	2,067 (20.67)	270 (0)

Table 7.5: Parameters of the available heat sources from the tokamak during a pulse, with values during a dwell given in brackets, in Case Study III. *return temperatures to the tokamak and thermal power are not fixed in the simulation

As in section 7.1, the thermal energy of the high temperature heat source is delivered to a molten salt thermal energy storage system, with design tank temperatures of 350°C and 550°C, that can be pumped to drive a steam Rankine cycle.

The temperatures of the medium and low temperature sources are incompatible with the operating range of the molten salt 290-600°C in the thermal energy storage system. Direct integration of these sources with the Rankine cycle was considered to avoid the complexity of additional thermal energy storage systems. Their temperatures make them suitable for feedwater preheating, as seen in figure 2.5 for an EU-DEMO concept design [21, 20]. During a dwell however, an alternative preheating method must be employed to avoid large transients in the Rankine cycle.

A parallel preheating stream is considered with a dwell heater fed by molten salt directly from the hot tank. A split valve then controls whether feedwater flows to the medium and low temperature heat exchangers during a pulse or to the dwell heater during a dwell. A preheater is still included after the parallel streams mix, allowing the design temperature at the outlet of the medium temperature heat exchanger and the dwell heater to be far from the saturation temperature. This avoids issues of two-phase flow before the evaporator during the transition between pulse and dwell

operation. The design configuration is shown in figure 7.12.

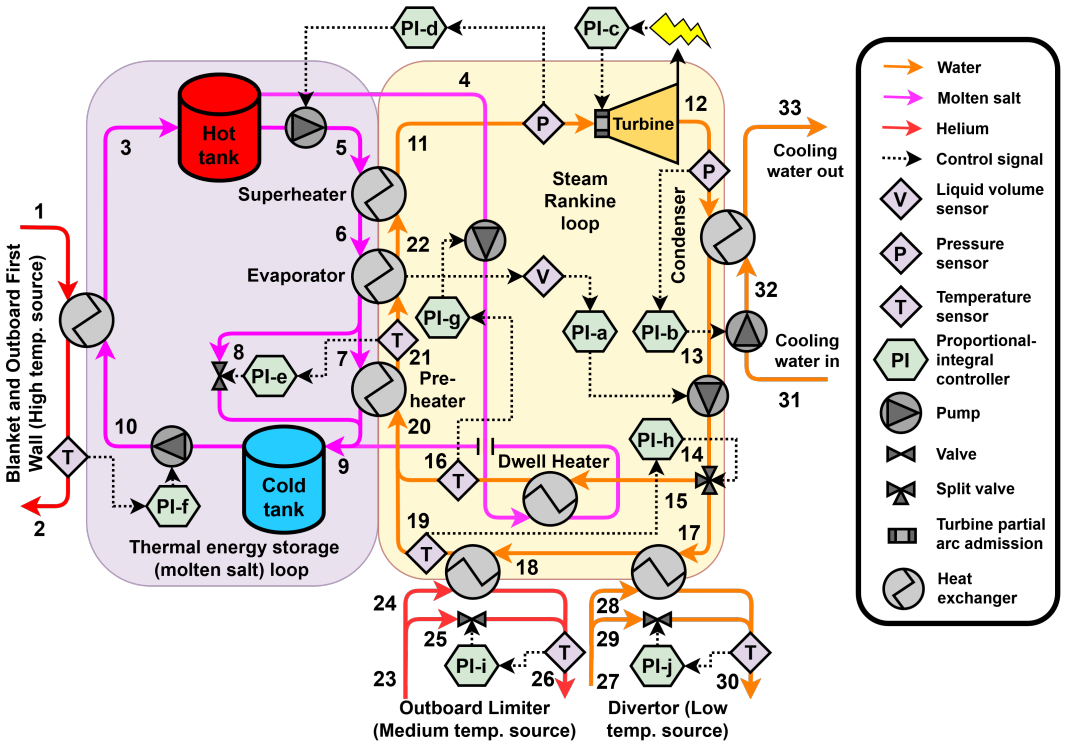

Figure 7.12: Power conversion system for three pulsed heat sources from a fusion tokamak in Case Study III. Thermal energy from the high temperature source is delivered to a molten salt thermal energy storage loop. The medium and low temperature sources deliver heat directly to the Rankine cycle. A parallel dwell heater fed by molten salt takes over their heat duty during a dwell.

Valve-regulated bypasses are included to allow the return temperature of the medium and low temperature coolants to be controlled. The molten salt bypass around the evaporator from section 7.1 is not included here for reduced model complexity. The impact of this is discussed in section 7.6.1.

7.4.1 Control System Design

As in section 7.1, all controllers are proportional-integral and are tuned by hand via multiple simulations. Due to the increased simulation time of this process relative to the process in section 7.2, the simultaneous optimisation of controllers with the process design presented in section 7.3.3 is not considered here. Detailed controller parameters are provided in appendix C.

Six of the controllers (PI-a, PI-b, PI-c, PI-d, PI-e, PI-f) are implemented as described in section 7.1.1. The set points of controllers PI-a, PI-b, PI-c, PI-d and

PI-f are the same as in section 7.1.1. The preheater outlet temperature set point of controller PI-e is set to be 10K below the temperature of steam leaving the evaporator.

A controller (PI-h) manipulates the feedwater split valve between the parallel streams to control the feedwater temperature at the outlet of the medium temperature heat exchanger. During a tokamak ramp-down, the controlled temperature will drop due to declining thermal power. The controller will respond by sending more feedwater to the dwell heater. Another controller (PI-g) detects this change in flow rate's impact on the feedwater temperature at the dwell heater output, increasing the flow of hot molten salt to the dwell heater to compensate. Both controllers have a set point of 267°C. The opposite response will occur during a ramp-up. As discussed in section 7.1, the heat exchanger models cannot handle zero mass flow rates. There will always be at least a small flow of feedwater through both parallel streams, but the flow should be small enough that the heat duty on the heat exchangers is negligible.

Two controllers (PI-i, PI-j) control the return temperatures of the medium and low temperature coolants to their set points of 225°C and 275°C respectively. They manipulate the openings of their respective bypass valves. EU-DEMO designs have considered controlling the bypass of feedwater instead to regulate heat transfer [38].

7.4.2 Model Parameters

The steady state model presented in section 7.1.2 is extended to include the additional heat sources and the dwell heater. The new model is presented in appendix C for brevity. The same assumptions used in the steady state model of Case Study II are made here. The turbine design mass flow rate is scaled by a factor of 1.05 to allow controller PI-c to bring the power output to its set point.

The parameters of the dynamic and steady state model, in addition to the initial equations of the dynamic models, are also presented in appendix C.

The temperature profiles of the Rankine cycle during a pulse and dwell, from the steady state model, are given in figures 7.13 and 7.14. The process is parametrised such that the preheater inlet temperature is constant. The lines only approximate the temperature profile; for example, in the evaporator the feedwater temperature would first increase before reaching saturation and maintaining a constant temperature.

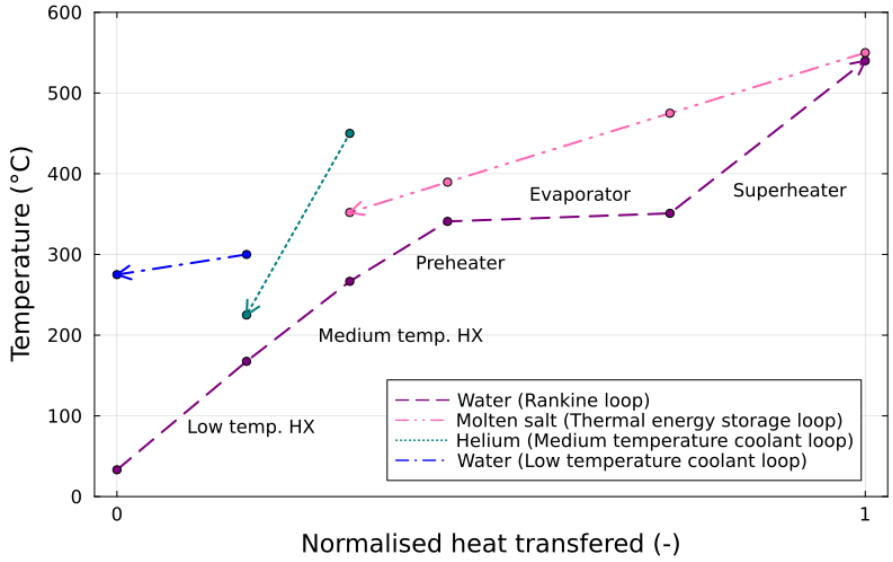

Figure 7.13: Heat sources of the steam Rankine cycle during a tokamak pulse, calculated by the steady state model in algorithm 3. Heat from the medium and low temperature heat sources is used to directly heat the feedwater.

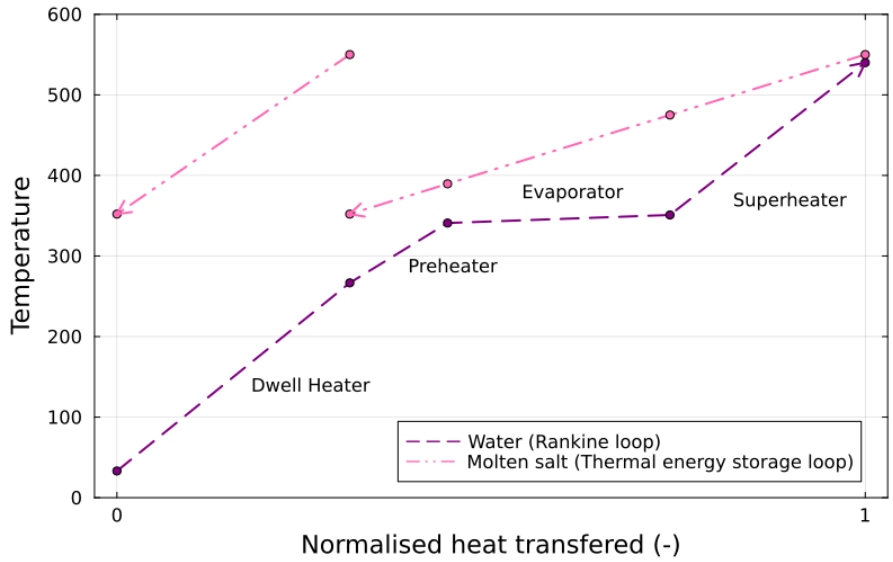

Figure 7.14: Heat sources of the steam Rankine cycle during a tokamak dwell, calculated by the steady state model in algorithm 3. Heat from the high temperature heat source that was stored during the pulse is used to replace the heat duty of the medium and low temperature sources.

7.5 Case Study III: Optimisation Problem Statement

The optimisation problem is formulated as seen in section 7.2:

$$\begin{aligned} &\min_{\boldsymbol{x}\in X} \begin{bmatrix} V_{\text{tank}} \\ -\phi \end{bmatrix} \\ &\text{s.t. } \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0} \\ &\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \end{aligned} \tag{7.3}$$

$\boldsymbol{g}(\boldsymbol{x})$, $\boldsymbol{h}(\boldsymbol{x})$ now include the equations of the dynamic process model for the new power conversion system design presented in section 7.4. The domain of the search space X is given in table 7.6, accounting for the larger amount of available thermal power.

Design variable	Minimum value	Maximum value
V_{tank}	2,000m^3	8,000m^3
$\dot{W}_{\text{pulse}}$	500MW	750MW
φ	0.25	1

Table 7.6: Domains of the design variables for equation 7.3.

With the added model complexity of the additional heat sources, parallel streams and additional controllers, simulations of the new model take $\sim 28\times$ longer to evaluate relative to the model in section 7.1 for the same t_{sim}. Most of the computational time is used to solve the period between pulses when transients are large and small time steps are required. To improve the speed of the optimisation, simulations cover an initial pulse period, a dwell and another pulse period with a ramp-down and ramp-up between them for a total time of $t_{\text{sim}} = 15{,}900$s. This will greatly reduce simulation time compared to the 7 dwells simulated in Case Study II.

To compensate for this, an alternative metric of long-term operability for balancing molten salt usage must be given. The proposed metric is to measure the change in volume of stored hot molten salt $\Delta V_{\text{ms,hot}}$ at the end of the simulation relative to one tokamak cycle period earlier, at the instant the ramp-down begins. Positive values of $\Delta V_{\text{ms,hot}}$ show excessive charging that will eventually lead to overflow, while negative values show overutilisation of the molten salt that will eventually

drain the hot tank. This can then be used in the calculation of infeasibility of designs that simulate successfully, shown in table 7.7.

Simulation Outcome	Infeasibility, g	Reasoning
Completed successfully	$-\frac{\Delta V_{\text{ms,hot}}}{V_{\text{tank}}} - 0.01$	Designs that underutilize molten salt will be feasible, but outcompeted by propagated designs that balance molten salt better and have smaller tanks – useful to keep in the population
Simulation failed after initialisation	$1 + \left(1 - \frac{t_{\text{fail}}}{t_{\text{sim}}}\right)$	Molten salt is being critically over-/under- utilised - later simulation failures imply that designs are closer to balancing the molten salt usage and becoming feasible
Simulation failed before initialisation	10^3	Error may not reflect on the design if due to OpenModelica-Julia coupling
Simulation evaluation time exceeded 1,800s	10^6	Undesirable behaviour, such as oscillations, are slowing the simulation

Table 7.7: Infeasibility assignments for the different possible simulation outcomes in Case Study III. t_{fail} is the internal simulation time at failure. $\Delta V_{\text{ms,hot}}$ is the change in stored hot molten salt at the end of the simulation relative to one cycle period earlier. It is assumed that all mid-simulation failures are caused by the molten salt tank model assertion in equation 4.34.

Designs that simulate successfully may now be considered infeasible if the volume of stored hot molten salt changes by $< -1\%$ of the tank volume V_{tank}. Designs that underutilise molten salt are considered feasible so that they are more likely to propagate with the motivation that they can produce offspring with smaller tank sizes and higher turbine pulse power set points that can outcompete them. For simulations that fail after initialisation, infeasibility is offset by 1 so that they will always be given a lower fitness than designs that simulate successfully.

Fresa alone is used for this Case Study to avoid any influence from the loss of diversity in the Cocoa implementation due to reinitialisation. Its parameters are given in table 7.8, and the reinitialisation script is looped 126 times. The initial population consists of a single design calculated from the steady state model for constant turbine

power, $\boldsymbol{x} = \left(3{,}976\text{m}^3, 523.5\text{MW}, 1\right)^{\text{T}}$. This design fails after initialisation.

Parameter	Value	Description
$n_{\text{f,max}}$	100	Maximum number of evaluations
n_{p}	(20,40)	Maximum number of designs to propagate each generation
fitness	*:nondominated*	Fitness ranking method
elite	*true*	Non-dominated solutions are always kept in the population
archiveelite	*true*	Save pruned non-dominated solutions
issimilar	*Fresa.isimilarx*	Similarity function
ε	0.01	Similarity threshold
multithreading	*true*	Evaluate designs in parallel

Table 7.8: Parameters of Fresa used for solving equation 7.3 in Case Study III. All other parameters are left as default. Please refer to `https://www.ucl.ac.uk/~ucecesf/fresa.html` for default values and more detailed explanations.

7.6 Case Study III: Optimisation Results

The non-dominated set found by the optimiser after 15,917 objective function evaluations is presented in figure 7.15. The set can be seen to have both a good range and distribution in the objective space, with an approximately linear relationship between objectives. It is believed that the improved optimisation results relative to Case Study II can be attributed to the changed calculation of g for designs that simulate successfully. Keeping designs that under-utilise molten salt in the feasible population encourages the optimiser to iteratively tweak these designs over the generations with small changes, as opposed to making larger changes to the designs as the plant propagation algorithm does for less fit designs.

To assess the shape of the non-dominated set, a steady state energy balance model is proposed to approximate the Pareto front. Assuming a constant ratio of turbine power output to the heat flow rate into the Rankine cycle $\varepsilon_{\text{turbine}}$, perfect turbine control, a 100% efficient thermal energy storage system and molten salt temperatures equal to the tank design values, the change in volume of stored hot

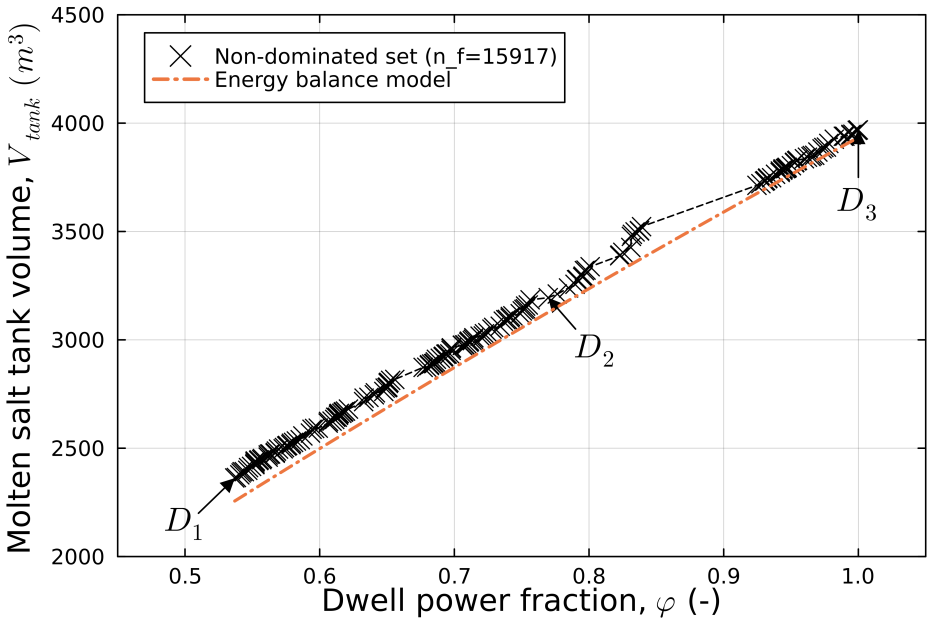

Figure 7.15: Non-dominated set at the end of the optimisation by the optimiser in Case Study III. Three designs $D_1 = \left(2{,}363\text{m}^3, 673.2\text{MW}, 0.5369\right)^\text{T}$, $D_2 = \left(3{,}195\text{m}^3, 651.4\text{MW}, 0.7697\right)^\text{T}$ and $D_3 = \left(3{,}969\text{m}^3, 630.6\text{MW}, 1\right)^\text{T}$ are highlighted for further discussion in the text. An estimate of the Pareto front is calculated based on a steady state energy balance model.

molten salt during a pulse and dwell is estimated as:

$$\Delta V_\text{pulse} = \frac{\dot{Q}_\text{total} - \frac{\dot{W}_\text{pulse}}{\varepsilon_\text{turbine}}}{(T_\text{ms,hot} - T_\text{ms,cold}) c_\text{p,ms} \rho_\text{ms}} \tau_\text{pulse} \tag{7.4}$$

$$\Delta V_\text{dwell} = \frac{-\frac{\varphi \dot{W}_\text{pulse}}{\varepsilon_\text{turbine}}}{(T_\text{ms,hot} - T_\text{ms,cold}) c_\text{p,ms} \rho_\text{ms}} \tau_\text{dwell} \tag{7.5}$$

where τ_pulse, τ_dwell are the pulse and dwell periods and $\dot{Q}_\text{total}$ is the sum of the thermal power of all the tokamak heat sources. Ramp-ups and ramp-downs are assumed to be a pulse for half their period and a dwell for the other half, as in the steady state models in sections 7.1.2 and 7.4.2.

Designs are balanced if $\Delta V_\text{pulse} + \Delta V_\text{dwell} = 0$, which then defines the turbine power during pulse as:

$$\dot{W}_\text{pulse} = \frac{\varepsilon_\text{rankine} \dot{Q}_\text{total} \tau_\text{pulse}}{\tau_\text{pulse} + \varphi \tau_\text{dwell}} \tag{7.6}$$

To estimate the volume of a molten salt tank, the initialisation of the tank at

10% volume fill and the upper constraint of 95% volume fill must be considered. The volume change during a dwell is hence scaled by 0.85^{-1}, giving the tank volume of a balanced design as:

$$V_{\text{tank}} = \frac{\varphi \dot{Q}_{\text{total}} \tau_{\text{pulse}} \tau_{\text{dwell}}}{0.85(\tau_{\text{pulse}} + \varphi \tau_{\text{dwell}})(T_{\text{ms,hot}} - T_{\text{ms,cold}}) c_{\text{p,ms}} \rho_{\text{ms}}} \tag{7.7}$$

The Pareto front predicted by equation 7.7 is shown in figure 7.15, showing very close agreement with the non-dominated set. This suggests that the optimiser has converged close to the real Pareto front of equation 7.3.

By calculating $\varepsilon_{\text{turbine}} = 0.400$ from algorithm 3 and using equations 7.6 and 7.7, three new designs are generated at the φ values of designs D_1, D_2 and D_3. Putting the design parameters of the new designs into the dynamic process model, it is found that the design $\boldsymbol{x} = \left(2, 256\text{m}^3, 675.9\text{MW}, 0.5369\right)^{\text{T}}$ fails due to the hot molten salt tank overflowing. The designs $\boldsymbol{x} = \left(3, 127\text{m}^3, 653.5\text{MW}, 0.7697\right)^{\text{T}}$ and $\boldsymbol{x} = \left(3, 934\text{m}^3, 632.8\text{MW}, 1\right)^{\text{T}}$ simulate successfully, but are infeasible due to over-utilisation of hot molten salt.

It is concluded that these designs are unbalanced due to discrepancies between the dynamic process model and the steady state model. This is supported by the difference between the non-dominated set and the estimated Pareto front increasing with decreasing φ. As φ decreases, the magnitude of transients in the process generally increase due to the changing turbine power set point forcing the other controllers to have to react in turn to maintain their set points. This statement is supported by the results presented in sections 7.6.1, 7.6.2 and 7.6.3.

It is interesting to ask why the non-dominated set has no designs for $\varphi < 0.5369$. Unlike in Case Study II, the preheater outlet temperature is well controlled by PI-e for all three designs D_1, D_2 and D_3. It is noted that for lower values of φ, the molten salt tank temperatures begin to deviate more from their design values, as shown in section 7.6.1. For design D_1, it does not appear extreme enough to potentially cause simulation failures. No reasons can be identified that would prevent further reductions in φ. It is suggested that the range of the non-dominated set in figure 7.15 is a product of the algorithm's stochasticity and the total number of function

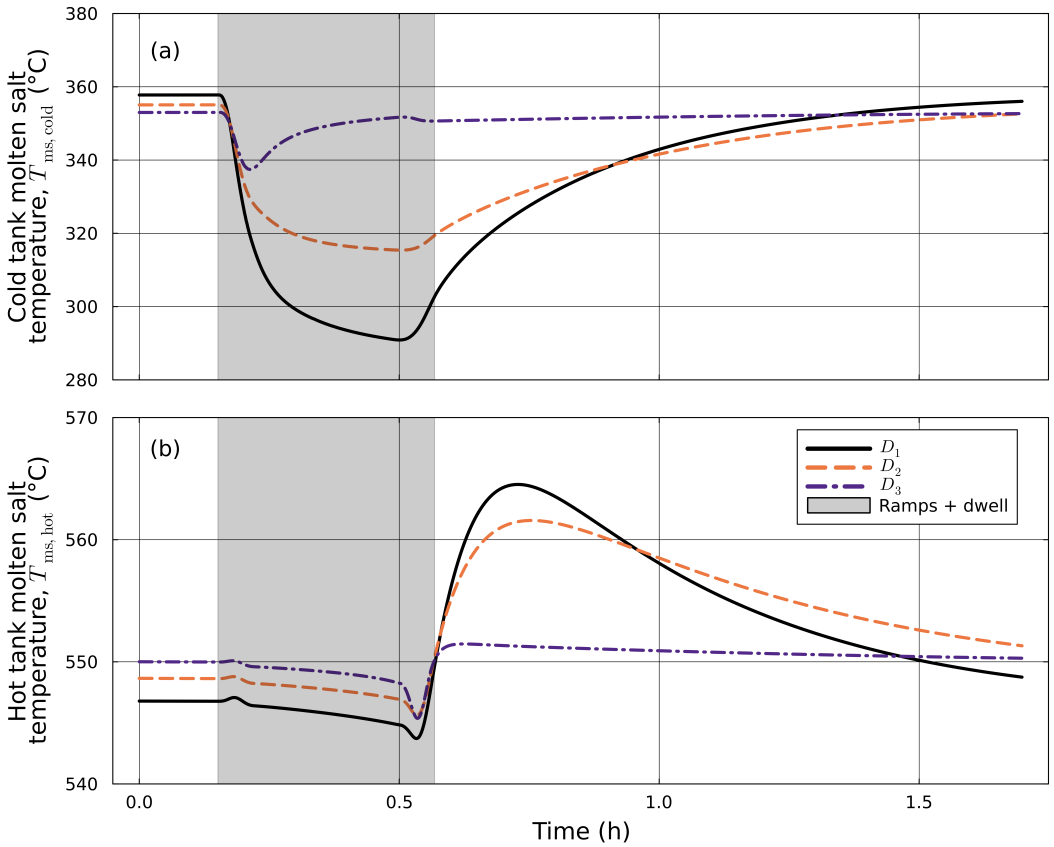

Figure 7.16: Dynamics of the temperature in the (a) cold tank (b) hot tank of designs D_1, D_2 and D_3 from the final non-dominated set in Case Study III. The control system does not regulate the temperatures, so large deviations from the design temperatures are seen, especially for smaller dwell power fractions φ.

evaluations, with a wider non-dominated set possible if the optimisation is continued for longer. This statement is supported by the existence of a manually-found feasible design at $\boldsymbol{x} = \left(2, 300\text{m}^3, 676.5\text{MW}, 0.5\right)^{\text{T}}$.

7.6.1 Impact of Absent Molten Salt Tank Temperature Control

To reduce model complexity, the evaporator molten salt bypass used to control the cold tank temperature in Case Study II is omitted in this process. To assess the impact of this simplification on the process dynamics, the dynamics of the molten salt tank temperatures for designs D_1, D_2 and D_3 are shown in figure 7.16.

Over a single dwell period, there are significant fluctuations in the tank temperatures, particularly the cold tank. The magnitude of the fluctuations increases with decreasing φ.

The cold tank temperature is driven down by the lower inlet temperature of the molten salt. The steady state model designs for steady state turbine power output, and hence sizes the Rankine cycle heat exchangers for a design flow rate of water. If the turbine power set point is reduced during a dwell, the flow rate of water is reduced by the control system. Hence, the heat exchangers in the Rankine cycle exchange too much heat between the molten salt and water, bringing the molten salt temperatures below the design values in the steady state model. The dwell heater in particular is poorly parametrised, as it always receives a reduced mass flow rate of water if $\varphi < 1$. The smaller the value of φ, the smaller the mass flow rate of feedwater during a dwell and the more pronounced the drop in the cold tank temperature is.

The dwell heater sizing could be addressed by including φ into the steady state model and parametrising each simulation independently with the steady state model. The over-sizing of the preheater, evaporator and superheater during a dwell cannot be addressed this way, and would require modifications to the process to address, such as the evaporator bypass considered in Case Study II or auxiliary heaters [251]. These methods can lead to issues of low molten salt temperatures outside the cold tank however, as discussed in section 7.3.

The varying cold tank temperature has a knock on effect on the hot tank temperature. The lower the cold tank temperature, the less molten salt mass flow rate is required to meet the helium return temperature set point of controller PI-f during a pulse. The lower flow rate then allows the molten salt to be brought to a higher temperature in the high temperature heat exchanger, leading to an increase in the hot tank temperature.

Even before the dwell, there are temperature differences between the designs. This is again due to the single steady state parameterisation used. As φ decreases, $\dot{W}_{\text{pulse}}$ should increase to balance the storage and discharge of molten salt, as seen in equation 7.6. Hence, the feedwater flow rate during a pulse also increases, leading to increased molten salt flow to the superheater to maintain the turbine inlet pressure and increased molten salt temperature at the preheater outlet. This temperature difference in the cold tank is then inverted in the hot tank by the mechanism mentioned above.

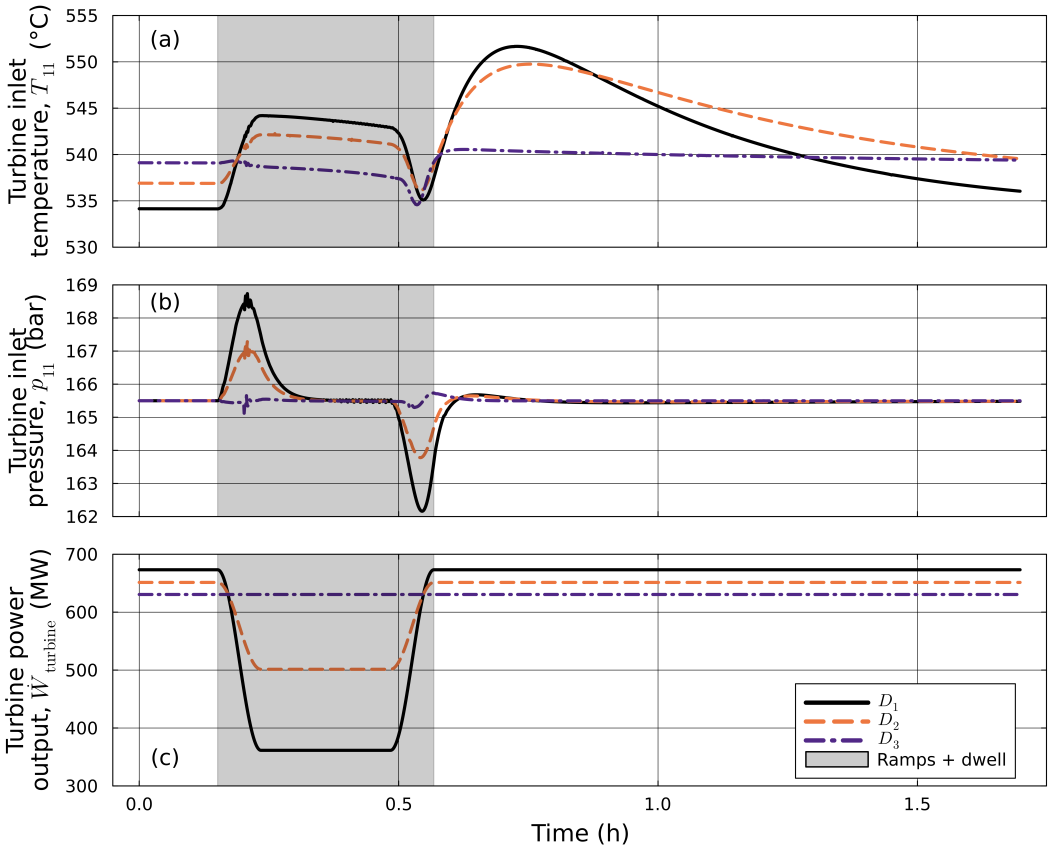

Figure 7.17: Dynamics of the (a) turbine inlet temperature (b) inlet pressure (c) power output of designs D_1, D_2 and D_3 from the final non-dominated set in Case Study III. Note that the turbine inlet temperature is not directly controlled by a controller, unlike the other two variables.

The lack of molten salt tank temperature control is a limitation of this power conversion system design. Electric heaters could be fitted in one tank to help stabilise the process, although this will impact net plant efficiency as approximately $\sim 1 - \eta_{\text{rankine}}$ of the energy will be lost. In addition, the methodology should improve the integration of the steady state model to allow for more design versatility.

7.6.2 Dynamics at the Turbine Inlet

A primary goal of the power conversion system design of pulsed fusion power plants is minimising the impact of thermal transients on process components. Steam turbines are particularly sensitive components [42, 103], so large fluctuations of their inlet conditions should be avoided. The turbine inlet dynamics of designs D_1, D_2 and D_3 are presented in figure 7.17.

The turbine inlet temperature is not controlled by any controllers in this process, and is largely governed by the temperature of hot molten salt entering the superheater. During a dwell, this is compounded with reductions in steam flow through the superheater increasing the outlet temperature if $\varphi < 1$. Hence, the profile of D_3 qualitatively matches the hot tank temperature in figure 7.16b, whereas for D_2 and D_3 the profile shape is elevated during the dwell. Across all three designs, the inlet temperature variation is < 4K/min, which is within the turbine rotor tolerance of typical materials [103].

Controller PI-d controls the turbine inlet pressure by manipulating the flow of hot molten salt to the superheater. All three designs show fluctuations during ramps, with the magnitude increasing with decreasing φ. Even for design D_1, the magnitude and time period of the peaks agree with typical responses expected for load following in conventional steam Rankine power cycles where ~ 12bar changes are seen over a period of < 800s [184]. During ramp down, a period of smaller amplitude, higher frequency oscillations in the pressure are seen for each design. This could be a result of feedback between multiple process controllers and poor controller tuning. This could be addressed by more rigorous tuning methods, such as in Case Study I or section 7.3.3, or more advanced multiple-input-multiple-output control methods like model predictive control, which are impractical to implement here.

The turbine power output appears to be well controlled, as in Case Study II, closely following the pulsed profile imposed by the set point. The maximum error seen by PI-c is 12.7kW for design D_1 during the ramp up. As predicted by equation 7.6, the design pulse power $\dot{W}_{\text{pulse}}$ increases with decreasing φ to balance the system's energy inputs and outputs.

Overall, it can be seen that the design is successful at maintaining stable operating conditions for the turbine, a key requirement of the power conversion system design. For long-term operation however, the lack of molten salt tank temperature control would begin to compromise the turbine's operation.

7.6.3 Performance of the Parallel Feedwater Heaters

The effectiveness of the two parallel preheating streams is analysed by looking at the dwell heater operation across the ramp and dwell periods. The transients in the mass flow rate and outlet temperature of feedwater in the dwell heater are shown in figure 7.18.

Large oscillations can be seen in both the feedwater outlet temperature and mass flow rate at the start and end of the dwell period. It is believed that this is due to poor tuning of controller PI-h, the action of which on the split valve is shown in figure 7.19. Oscillations can be seen to grow in magnitude during a dwell in design D_1, just notable in figure 7.18b between t=500s and t=750s. This is then replaced by periodic constant magnitude perturbations in the later half of the dwell period.

In figure 7.18a, design D_3 enters the ramp down period with a feedwater outlet temperature far above the other designs. Of the three designs, it has the lowest value of $\dot{W}_{\text{pulse}}$, and hence the lowest mass flow rate of water around the Rankine cycle during a pulse. The low temperature and medium temperature heat exchangers are sized for $\varphi = 1$, so controller PI-h saturates and opens the split valve to let the maximum flow of feedwater to them. The controller PI-g then saturates at the minimum flow rate of hot molten salt before it can bring the dwell heater feedwater outlet temperature to its set point, resulting in the higher pulse temperature.

Designs D_1 and D_2 have a higher flow rate of water around the Rankine cycle, but controller PI-h adjusts the split valve opening until approximately the same flow goes through the low temperature and medium temperature heat exchangers as in design D_3. This is because the thermal powers of the low and medium temperature heat sources are the same across designs. The excess feedwater flow goes through the dwell heater, even during a pulse. The larger flow rate means that controller PI-g can bring the dwell heater feedwater outlet temperature to its set point before it saturates. The feedwater mass flow rate through the dwell heater during a pulse in design D_3 is $< 0.2\%$ of the flow rate around the cycle, so the high outlet temperature will have negligible impact on the process dynamics after mixing with the other stream.

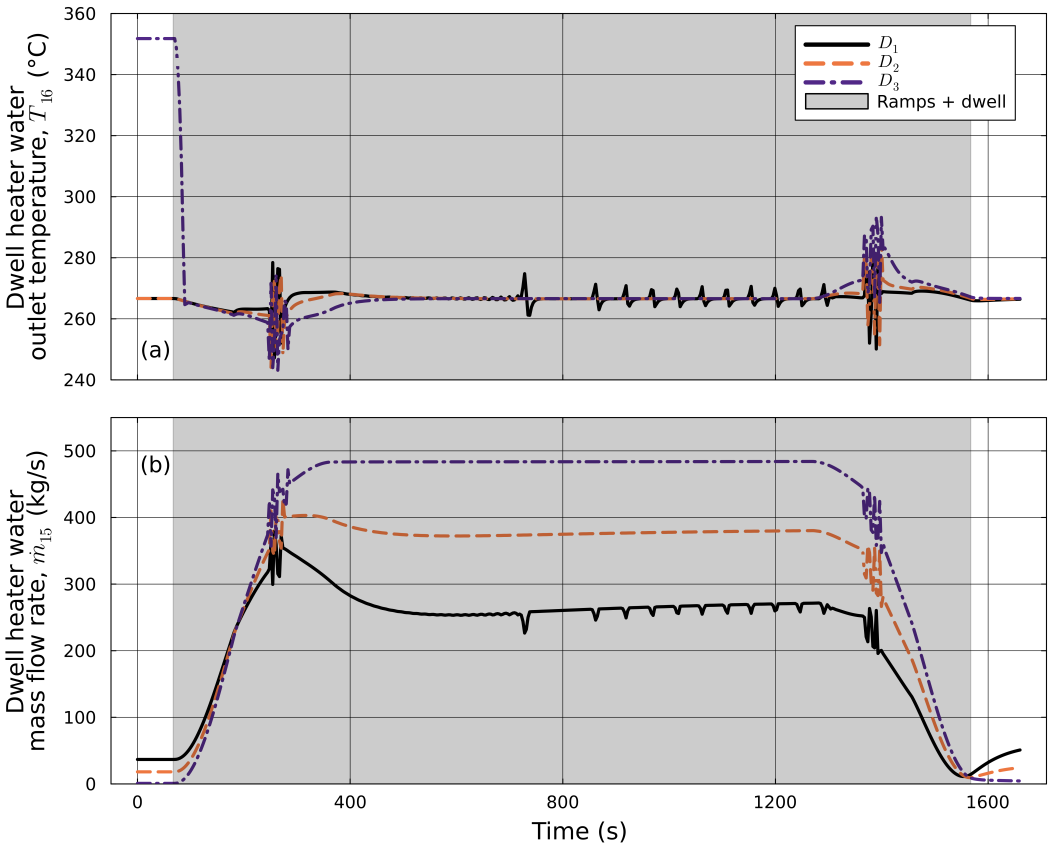

Figure 7.18: Dynamics of the dwell heater's (a) feedwater outlet temperature (b) feedwater mass flow rate of designs D_1, D_2 and D_3 from the final non-dominated set in Case Study III. Oscillations are present in all designs due to the control system.

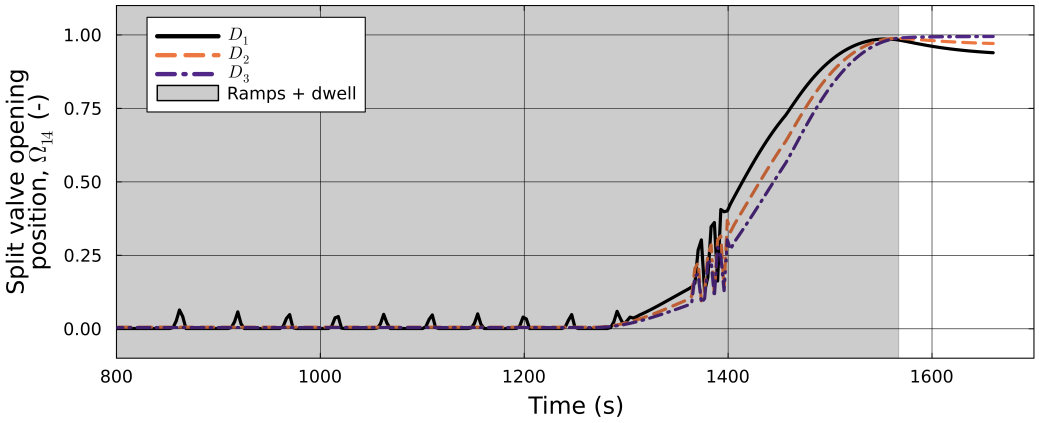

Figure 7.19: Dynamics of the split valve opening position controlled by controller PI-h in Case Study III. An opening of $\Omega = 1$ and $\Omega = 0$ corresponds to all the feedwater being sent to the low temperature heat exchanger or the dwell heater, respectively.

Multiple-input-multiple-output control would likely be far more competent at managing the manipulation of multiple process components simultaneously to handle the pulse-dwell transition than multiple single-input-single-output controllers. Regardless, it seems clear that the performance of the existing control system has room for improvement with better tuning.

7.7 Summary

In this chapter, dynamic models of power conversion systems with one or three pulsed heat source from a tokamak are presented. The necessary size of the thermal energy storage tanks scales with the amount of power produced by the turbine during a tokamak dwell. The trade-off between the size of the thermal energy storage system V_{tank} and the dwell power fraction φ is posed as a bi-objective optimisation problem for both systems.

For the single pulsed heat source system, three different optimisers are used to solve the problem: Fresa with Hadamard ranking, Fresa with non-dominated ranking and Cocoa with two Fresa instances using either Hadamard or non-dominated ranking. The Cocoa optimiser has the widest non-dominated set in φ, allowing for a 35.0% reduction in the molten salt tank volume at the cost of a 39.8% reduction in the turbine power during dwell. The dynamics of three chosen designs show good control of the turbine power, but the intermediate design is noted to have excess tank volume that could be further reduced. Reduction of φ is noted to lead to issues such as low molten salt temperatures and two-phase flow before the evaporator. The non-physical equations of the robust logarithmic mean temperature difference are only used during dwell in one heat exchanger, and are not thought to have a significant impact on the wider process dynamics.

Comparing the non-dominated sets of each optimiser, it can be seen that Cocoa does seem to share the search behaviour of the individual Fresa instances, with a wide non-dominated set and more evenly dispersed designs in objective space. The Cocoa optimiser is missing possible reductions in the molten salt tank volume achieved by the other optimisers. This could be caused by the loss of population

diversity due to the reinitialisation script.

Preliminary results of the simultaneous optimisation of process and control objectives are also presented. The increased dimension of the search space and additional objective require additional function evaluations at the expense of additional computational evaluation time. Comparing the designs for steady state turbine power, treating the control objective as secondary to the others, the optimised controller shows far more stable behaviour, completely eliminating some high frequency oscillations in the controlled variable.

For the system with three pulsed heat sources, Fresa is used with the Hadamard ranking method based on the conclusions of section 7.3.2. A good range and distribution of non-dominated designs is found, which agrees well qualitatively with a steady state model of the process. This suggests that the optimiser has found a good approximation of the true Pareto front for equation 7.3. It also shows that the methodology is able to produce designs that are comparable to an idealised process with perfect and instantaneous control. This supports the use of the methodology for other optimisation problems that cannot be well-approximated by steady state models.

The power conversion system design is seen to be capable of handing the transition between pulse and dwell with acceptable transients at the turbine inlet. Uncontrolled molten salt tank temperatures compromises the long-term process operability, and hand-tuning of controllers leads to process oscillations. This can be addressed by adjusting the control system and considering simultaneous control and process design optimisation, as shown in section 7.3.3. The evaluation methodology is also limited by the fixed design of the dwell heater. This leads to infeasible molten salt temperatures that would result in freezing in a real system for designs with a low value of φ. Individual parameterisation of designs could address this and open up the methodology to more extensive design adjustments, such as having key cycle pressures and temperatures as design variables.

Chapter 8

Conclusion & Future Recommendations

Nuclear fusion is an appealing technology for firm power generation in a future decarbonised power grid. With the possibility of abundant fuels, minimal waste or emissions and minimal geographic requirements, it could supplement intermittent renewable sources. Commercialisation of fusion energy requires many engineering challenges to be addressed first, one of which is the efficient conversion of thermal energy from multiple sources, which may or may not be pulsed, into electrical energy. The state of the art in thermal power generation and fusion devices is discussed in chapter 2.

Traditional design approaches to thermal power plants often rely on industrial experience for process configurations and optimise the system with steady state models before assessing process dynamics at a later design stage. Optimisation and dynamic modelling are useful tools for the designs of novel power plants [269, 168]. For fusion tokamaks, heat will be available from multiple coolant streams and may be highly transient due to pulsed operation or tokamak trips. Efficient integration of the available heat into the power cycle is necessary for net energy generation and the resulting process dynamics from a transient heat source will determine the feasibility of designs. These features should therefore be incorporated into the design process from the beginning.

This work develops an optimisation-based methodology to design power conver-

sion systems for pulsed fusion tokamaks, described in chapter 5. The pulsed nature of the heat sources requires that dynamic process models are developed, which are discussed in chapter 4. This is done with an object-oriented approach in the Modelica language [13], introduced in chapter 3, allowing individual component models to be easily reused to create models of different process structures. An emphasis is made on low-complexity computational models for reduced simulation times and parametrisation. By performing a dynamic simulation for each design evaluation using OpenModelica [61], the process dynamics are explicitly considered by the optimisation. The plant propagation optimisation algorithm in Fresa.jl [92] is used as it is well suited for simulation-based design evaluations. A multi-agent system Cocoa.jl [96] is also considered to combine the search behaviour of different fitness ranking methods of the algorithm, which is seen in section 7.3.

Control systems are essential for the operation of thermal power plants, enabling the operation to be robust to changing conditions. Control of the turbine power output is particularly important to maintain synchronisation with the power grid. When designing a dynamic process, the control should be considered simultaneously as the dynamic behaviour of the system will depend on both. In chapter 6, the optimisation methodology is demonstrated for the optimal tuning of a proportional-integral-derivative controller that regulates the turbine power output of a power conversion system with a constant heat source. The optimised controller shows significantly reduced amplitude of errors, but does introduce more oscillations compared to the hand-tuned controller.

If a fusion tokamak is operated on a pulsed schedule, this would drive the dynamics of the whole power conversion system. Unlike solar thermal power plants, the availability of thermal power may follow a predictable schedule. This could allow for design optimisations that are not possible for other systems, like the exact sizing of the thermal energy storage system to the thermal power profile. In chapter 7, power conversion systems for a tokamak with one and three pulsed heat sources are considered. The trade-off between the size of the thermal energy storage system and the fraction by which the turbine power output is reduced during a dwell is

presented for each system. In each case, a significant reduction in the tank volume is possible by reducing the dwell power. The control of the molten salt temperatures around the cycle is noted to be a challenge for both systems, with temperatures below the freezing point occurring in the designs with minimal dwell power fractions. Preliminary results on simultaneous optimisation of control and process variables show that greatly improved control can be achieved with minimal losses in the original design objectives, at the cost of increased computational time.

The novel contributions of this work are the development of a design optimisation methodology that uses dynamic process models from OpenModelica and the development of dynamic process models of power conversion systems for fusion power plants with pulsed heat sources. To achieve this, novel lumped-parameter heat exchanger models are described that balance computational complexity and the representation of their main dynamics. By considering the process dynamics in the early stages of design, it is hoped that better designs can be found, regardless of the specific system. The methodology could also be applied to optimise general systems that can be represented as systems of differential-algebraic equations in OpenModelica.

8.1 Future Recommendations

The dynamic process models developed are computationally tractable and capture key dynamics that allow for discussion of the final designs. The modelling side of this work could be improved however, with some possible directions given below:

- Regenerative heaters are a common feature of steam Rankine cycles to improve thermal efficiency. Extracted steam is used to preheat feedwater, partially condensing the steam. The developed lumped-parameter models do not support partial phase change and are unsuitable for this use. A moving boundary model would be a computationally efficient way to model such heat exchangers and better represent real power plant designs. Such models could also be able to be adapted for dual-phase change heat exchangers utilising phase change slurry as the thermal storage medium [145].

- Besides water, simple thermodynamic property models are used. If packages such as CoolProp could be integrated into the modelling environment, helium could be modelled more accurately. For molten salt, polynomial temperature dependence could be estimated from available property data.

- While lumped-parameter models were emphasised during modelling, simulations of the power conversion system in section 7.4 still take >5min, which limits the degree of optimisation that is feasible. The training of faster surrogate models using full dynamic models has been demonstrated as an effective method to improve optimisation of thermal power systems [308, 71, 310].

- Control of the molten salt temperature is identified as a major failing of the proposed power conversion system designs for low dwell power fractions. Process modifications, such as electric heaters or alternative flow configurations, could be considered to address these issues. The process models could be given explicit limits on the safe operating conditions of each medium. Alternative thermal energy storage technologies could also be considered, such as latent heat storage.

- While many of the individual component models have been used in works that have been successfully validated against experimental data, being able to validate the lumped-parameter heat exchanger models would increase the confidence in the optimisation results.

- Only proportional-integral and proportional-integral-derivative controllers are used in this work. Advanced control methods, such as model predictive control, have been seen to outperform purely proportional-integral-derivative control in thermal power systems and could be implemented for fusion power plant models to better handle pulsed operation [121].

While the proposed methodology has been shown to generate well-optimised designs, there are limitations. Possible improvements to the optimisation methodology are given below:

- Expanding the simultaneous optimisation of control and process design variables seen in section 7.3.3 to multiple controllers could allow for more stable dynamics throughout the whole process.

- The need for regular reinitialisation due to memory issues with the OpenModelica-Julia coupling are thought to have impacted the performance of the Cocoa multi-agent optimisation system in section 7.3. This could be addressed by using a Julia-based simulation environment, such as ModellingToolkit.jl, Modia.jl or Dyad. These may benefit from the high performance metrics of Julia to outperform OpenModelica. Another option is to export the Modelica models to Functional Mock-up Units, which can then be imported and simulated in Julia directly.

- The use of the steady state model to parametrise the dynamic models is based on a dwell power fraction of $\varphi = 1$. This leads to the poor sizing of the dwell heater in section 7.6 for $\varphi < 1$. Allowing each design to be parametrised uniquely would fix this, and also open up other design variables like key temperatures and pressures for use as design variables.

- The structural configuration of the power conversion systems are chosen and modelled by hand using the available component models. Structural synthesis methodologies, presented in section 5.1.2, would allow for simultaneous structural and parametric optimisation, which could make up for the lack of industrial experience of the user. Such methodologies could simultaneously optimise multiple process components.

- Bayesian optimisation is a method commonly used for optimising expensive objective functions with no available derivatives, such as computational fluid dynamics models [76]. By training a surrogate model of the objective function, candidate points can be selected for evaluation to maximise the information gained of the objective space. While the objective functions in this work are comparatively cheap, the optimisation could still benefit from the method.

- By including cost models for each process model presented in this work, economic or thermoeconomic optimisation objectives could be considered. These economic objectives often trade off against thermodynamic objectives, and such trade offs are important to provide to decision makers. The current work uses the thermal energy storage system size as a proxy for economic cost, but when considering simultaneous design of multiple components, such as heat exchangers, the cost of each component must be converted to a standard currency for use in a single objective.

- Metrics related to process safety could be incorporated into the design optimisation. Possible metrics include the maximum rate of change of temperatures in sensitive components, maximum rate of change in generated power and the proximity of molten salt tank levels to maximum and minimum safety thresholds. These could either be considered as objective variables or as inequality constraints.

Appendix A

Model Parameters for Case Study I

Unit	Initial equations
Preheater	$\frac{\mathrm{d}}{\mathrm{d}t}h_{\mathrm{hf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t}h_{\mathrm{cf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t}\bar{T}_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$, $\frac{\mathrm{d}}{\mathrm{d}t}\Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$
Evaporator	$\frac{\mathrm{d}}{\mathrm{d}t}h_{\mathrm{hf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t}p_{\mathrm{wf}}\|_{t=0} = 0\mathrm{Pa/s}$, $\left.\frac{V_{\mathrm{L}}}{V_{\mathrm{wf}}}\right\|_{t=0} = 0.5$, $\frac{\mathrm{d}}{\mathrm{d}t}\bar{T}_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$, $\frac{\mathrm{d}}{\mathrm{d}t}\Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$
Superheater	$T_{\mathrm{hf,out}}\|_{t=0} = 350°C$, $T_{\mathrm{cf,out}}\|_{t=0} = 340°C$, $\frac{\mathrm{d}}{\mathrm{d}t}\bar{T}_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$, $\frac{\mathrm{d}}{\mathrm{d}t}\Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$
Condenser	$T_{\mathrm{cf,out}}\|_{t=0} = 25°C$, $p_{\mathrm{wf}}\|_{t=0} = 1\mathrm{bar}$, $\left.\frac{V_{\mathrm{L}}}{V_{\mathrm{wf}}}\right\|_{t=0} = 0.5$
Hot tank	$T_{\mathrm{ms}}\|_{t=0} = 450°C$, $\left.\frac{V_{\mathrm{ms}}}{V_{\mathrm{tank}}}\right\|_{t=0} = 0.5$
Cold tank	$T_{\mathrm{ms}}\|_{t=0} = 300°C$, $\left.\frac{V_{\mathrm{ms}}}{V_{\mathrm{tank}}}\right\|_{t=0} = 0.5$
PID-a	$y\|_{t=0} = 1{,}500\mathrm{kg/s}$
PI-a	$y\|_{t=0} = 75\mathrm{kg/s}$
PI-b	$y\|_{t=0} = 5{,}000\mathrm{kg/s}$
PI-c	$y\|_{t=0} = 0$

Table A.1: Initial equations for the dynamic models in Case Study I in chapter 6. Please refer to figure 6.1 or section 6.1.1 for controller labels.

Component	Parameter	Description
Cold tank & hot tank	$Z_{\text{tank}} = 20\text{m}$	Tank height
	$r = 10\text{m}$	Tank radius
	$U_{\text{ms,amb}} = 10\text{W/(m}^2\text{ K)}$	Molten salt-ambient heat transfer coefficient
	$T_{\text{amb}} = 30°\text{C}$	Ambient temperature
	$p_{\text{tank}} = 1\text{bar}$	Tank pressure
Turbine	$\eta_{\text{is}} = 0.9$	Isentropic efficiency
	$\eta_{\text{mech}} = 0.9$	Mechanical efficiency
	$\dot{m}_{\text{des}} = 125\text{kg/s}$	Design mass flow rate
	$T_{\text{in,des}} = 400°\text{C}$	Design inlet temperature
	$p_{\text{in,des}} = 120\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 0.05\text{bar}$	Design outlet pressure
All heat exchangers*	$\varepsilon = 0.1\text{K}$	RLMTD threshold
	$\xi = 5\text{K}^{-1}$	RLMTD steepness
	$c_{\text{w}} = 500\text{J/(kg K)}$	Wall specific heat capacity
Preheater	$A_{\text{cf}} = A_{\text{hf}} = 500\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{cf}} = U_{\text{hf}} = 2{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{cf}} = 3.15\text{m}^3$	Cold fluid side volume
	$V_{\text{hf}} = 13\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,cf}} = \Delta p_{\text{des,hf}} = 0\text{Pa}$	Design pressure drop of fluid
	$M_{\text{w}} = 500\text{kg}$	Wall mass

Component	Parameter	Description
Evaporator	$A_{\text{wf}} = A_{\text{hf}} = 2{,}000\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{wf}} = U_{\text{hf}} = 3{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{wf}} = 93\text{m}^3$	Working fluid side volume
	$V_{\text{hf}} = 10\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,hf}} = 0\text{Pa}$	Design pressure drop of hot fluid
	$M_{\text{w}} = 500\text{kg}$	Wall mass
Superheater	$A_{\text{cf}} = A_{\text{hf}} = 2{,}350\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{cf}} = U_{\text{hf}} = 1{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{cf}} = 12.5\text{m}^3$	Cold fluid side volume
	$V_{\text{hf}} = 31.5\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,cf}} = \Delta p_{\text{des,hf}} = 0\text{Pa}$	Design pressure drop of fluid
	$M_{\text{w}} = 500\text{kg}$	Wall mass
Condenser	$A_{\text{wf}} = A_{\text{cf}} = 7{,}500\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{wf}} = U_{\text{cf}} = 5{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{wf}} = 500\text{m}^3$	Working fluid side volume
	$V_{\text{cf}} = 100\text{m}^3$	Cold fluid side volume
	$\Delta p_{\text{des,cf}} = 0\text{Pa}$	Design pressure drop of cold fluid
	$M_{\text{w}} = 10\text{kg}$	Wall mass

Component	Parameter	Description
Cold molten salt pump	$\dot{m}$ =1,500kg/s	Mass flow rate
	$\eta = 0.8$	Efficiency
	$p_{\text{in,des}} = 1\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1.1\text{bar}$	Design outlet pressure
Hot molten salt pump	$\dot{m}_{\text{des}}$ =1,500kg/s	Design mass flow rate
	$\eta = 0.8$	Efficiency
	$p_{\text{in,des}} = 1\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1.1\text{bar}$	Design outlet pressure
Feedwater pump	$\dot{m}_{\text{des}}$ =75kg/s	Design mass flow rate
	$\eta = 0.8$	Efficiency
	$p_{\text{in,des}} = 0.05\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 80\text{bar}$	Design outlet pressure
Hot tank pressure drop	$\dot{m}_{\text{des}}$ =1,500kg/s	Design mass flow rate
	$\Delta p_{\text{des}} = 0.1\text{bar}$	Design pressure drop
Cold tank pressure drop	$\dot{m}_{\text{des}}$ =1,500kg/s	Design mass flow rate
	$\Delta p_{\text{des}} = 0.1\text{bar}$	Design pressure drop
Superheater pressure drop	$\dot{m}_{\text{des}}$ =100kg/s	Design mass flow rate
	$p_{\text{in,des}} = 0.5\text{bar}$	Design pressure drop
Spray valve	$\dot{m}_{\text{des}}$ =100kg/s	Design mass flow rate
	$p_{\text{in,des}} = 100\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 95\text{bar}$	Design outlet pressure
	$T_{\text{in,des}} = 300°\text{C}$	Design inlet temperature

Table A.2: Parameters of individual component models in Case Study I in chapter 6. *the tokamak heat source is not modelled as a heat exchanger

Label	Model	Controlled variable	Manipulated variable	Set point	K_c	τ_I (s)	τ_D (s)	Minimum output	Maximum output
PID-a	*LimPID*	Turbine power, $\dot{W}_{\text{turbine}}$	Mass flow rate of hot molten salt, $\dot{m}_3$	Figure 6.2	-0.001kg/(s W)*	30*	10*	500kg/s	2,500kg/s
PI-a	*PI*	Evaporator liquid fraction, $\frac{V_{\text{L,EVP}}}{V_{\text{wf,EVP}}}$	Mass flow rate of feedwater, $\dot{m}_{14}$	0.5	-500kg/s	500	-	25kg/s	250kg/s
PI-b	*PI*	Condenser pressure, p_{12}	Mass flow rate of cooling water, $\dot{m}_{19}$	0.05bar	50kg/(s Pa)	50	-	2,500kg/s	25,000kg/s
PI-c	*PI*	Turbine inlet temperature, T_{11}	Valve opening position, Ω_{valve}	420°C	0.1K^{-1}	20	-	0	1

Table A.3: Outline of the process controllers for Case Study I in chapter 6. Note that only the *LimPID* model implements integral anti-windup. Please refer to figure 6.1 or section 6.1.1 for controller labels. *values for the initial design only

Appendix B

Model Parameters for Case Study II

Unit	Initial equations
Primary heat exchanger	$T_{\mathrm{hf,out}}\vert_{t=0} = 450°C$, $T_{\mathrm{cf,out}}\vert_{t=0} = 500°C$, $\bar{T}_{\mathrm{w}}\vert_{t=0} = 475°\mathrm{C}$, $\Delta T_{\mathrm{w}}\vert_{t=0} = 25\mathrm{K}$
Preheater	$\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{hf,out}}\vert_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{cf,out}}\vert_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t} \bar{T}_{\mathrm{w}}\vert_{t=0} = 0\mathrm{K/s}$, $\frac{\mathrm{d}}{\mathrm{d}t} \Delta T_{\mathrm{w}}\vert_{t=0} = 0\mathrm{K/s}$
Evaporator	$T_{\mathrm{hf,out}}\vert_{t=0} = 340°\mathrm{C}$, $p_{\mathrm{wf}}\vert_{t=0} = 167\mathrm{bar}$, $\left.\frac{V_{\mathrm{L}}}{V_{\mathrm{wf}}}\right\vert_{t=0} = 0.5$, $\bar{T}_{\mathrm{w}}\vert_{t=0} = 340°\mathrm{C}$, $\Delta T_{\mathrm{w}}\vert_{t=0} = 0\mathrm{K}$
Superheater	$T_{\mathrm{hf,out}}\vert_{t=0} = 450°C$, $T_{\mathrm{cf,out}}\vert_{t=0} = 550°C$, $\bar{T}_{\mathrm{w}}\vert_{t=0} = 500°\mathrm{C}$, $\Delta T_{\mathrm{w}}\vert_{t=0} = 50\mathrm{K}$
Condenser	$T_{\mathrm{cf,out}}\vert_{t=0} = 25°C$, $p_{\mathrm{wf}}\vert_{t=0} = 0.0508\mathrm{bar}$, $\left.\frac{V_{\mathrm{L}}}{V_{\mathrm{wf}}}\right\vert_{t=0} = 0.5$, $\bar{T}_{\mathrm{w}}\vert_{t=0} = 25°\mathrm{C}$, $\Delta T_{\mathrm{w}}\vert_{t=0} = 0\mathrm{K}$
Hot tank	$T_{\mathrm{ms}}\vert_{t=0} = 580°C$, $\left.\frac{V_{\mathrm{ms}}}{V_{\mathrm{tank}}}\right\vert_{t=0} = 0.1$
Cold tank	$T_{\mathrm{ms}}\vert_{t=0} = 300°C$, $\left.\frac{V_{\mathrm{ms}}}{V_{\mathrm{tank}}}\right\vert_{t=0} = 0.9$
PI-a	$y\vert_{t=0} = 338\mathrm{kg/s}$
PI-b	$y\vert_{t=0} = 16{,}157\mathrm{kg/s}$
PI-c	$y\vert_{t=0} = 1$
PI-d	$y\vert_{t=0} = 2{,}690\mathrm{kg/s}$
PI-e	$y\vert_{t=0} = 0.1$
PI-f	$y\vert_{t=0} = 3{,}115\mathrm{kg/s}$
PI-g	$y\vert_{t=0} = 0.1$

Table B.1: Initial equations for the dynamic models in Case Study II in section 7.1. Please refer to figure 7.1 or section 7.1.1 for controller labels.

Component	Parameter	Description
Cold tank & hot tank	$U_{\mathrm{ms,amb}} = 0.2\mathrm{W/(m^2\ K)}$	Molten salt-ambient heat transfer coefficient
	$T_{\mathrm{amb}} = 30°\mathrm{C}$	Ambient temperature
	$p_{\mathrm{tank}} = 1\mathrm{bar}$	Tank pressure
Turbine	$\eta_{\mathrm{is}} = 0.9$	Isentropic efficiency
	$\eta_{\mathrm{mech}} = 1$	Mechanical efficiency
	$\dot{m}_{\mathrm{des}} = 337.6\mathrm{kg/s}$	Design mass flow rate
	$T_{\mathrm{in,des}} = 570°\mathrm{C}$	Design inlet temperature
	$p_{\mathrm{in,des}} = 165.5\mathrm{bar}$	Design inlet pressure
	$p_{\mathrm{out,des}} = 0.0508\mathrm{bar}$	Design outlet pressure
All heat exchangers	$c_{\mathrm{w}} = 500\mathrm{J/(kg\ K)}$	Wall specific heat capacity
	$\xi = 5\mathrm{K}^{-1}$	RLMTD steepness
Preheater	$A_{\mathrm{cf}} = A_{\mathrm{hf}} = 1{,}566\mathrm{m}^2$	Fluid-wall heat transfer area
	$U_{\mathrm{cf}} = U_{\mathrm{hf}} = 4{,}000\mathrm{W/(m^2\ K)}$	Fluid-wall heat transfer coefficient
	$V_{\mathrm{cf}} = 0.008731\mathrm{m}^3$	Cold fluid side volume
	$V_{\mathrm{hf}} = 0.06957\mathrm{m}^3$	Hot fluid side volume
	$\Delta p_{\mathrm{des,cf}} = 1.689\mathrm{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\mathrm{des,hf}} = 0.01010\mathrm{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\mathrm{des,cf}} = 337.6\mathrm{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\mathrm{des,hf}} = 2{,}690\mathrm{kg/s}$	Design mass flow rate of hot fluid
	$M_{\mathrm{w}} = 9{,}397\mathrm{kg}$	Wall mass
	$\varepsilon = 0.01\mathrm{K}$	RLMTD threshold

Component	Parameter	Description
Evaporator	$A_{\text{wf}} = A_{\text{hf}} = 1{,}551\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{wf}} = U_{\text{hf}} = 4{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{wf}} = 39.11\text{m}^3$	Working fluid side volume
	$V_{\text{hf}} = 14.10\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,hf}} = 0.01020\text{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\text{des,hf}} = 2{,}690\text{kg/s}$	Design mass flow rate of hot fluid
	$M_{\text{w}} = 45{,}630\text{kg}$	Wall mass
	$\varepsilon = 0.1\text{K}$	RLMTD threshold
Superheater	$A_{\text{cf}} = A_{\text{hf}} = 3{,}030\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{cf}} = U_{\text{hf}} = 4{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{cf}} = 0.01689\text{m}^3$	Cold fluid side volume
	$V_{\text{hf}} = 0.1346\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,cf}} = 1.672\text{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\text{des,hf}} = 0.01041\text{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\text{des,cf}} = 337.6\text{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\text{des,hf}} = 2{,}690\text{kg/s}$	Design mass flow rate of hot fluid
	$M_{\text{w}} = 18{,}180\text{kg}$	Wall mass
	$\varepsilon = 0.01\text{K}$	RLMTD threshold

Component	Parameter	Description
Condenser	$A_{\mathrm{wf}} = A_{\mathrm{cf}} = 10{,}820\mathrm{m}^2$	Fluid-wall heat transfer area
	$U_{\mathrm{wf}} = U_{\mathrm{cf}} = 10{,}000\mathrm{W/(m^2\ K)}$	Fluid-wall heat transfer coefficient
	$V_{\mathrm{wf}} = 272.7\mathrm{m}^3$	Working fluid side volume
	$V_{\mathrm{cf}} = 98.35\mathrm{m}^3$	Cold fluid side volume
	$\Delta p_{\mathrm{des,cf}} = 0.02020\mathrm{bar}$	Design pressure drop of cold fluid
	$\dot{m}_{\mathrm{des,cf}} = 16{,}160\mathrm{kg/s}$	Design mass flow rate of cold fluid
	$M_{\mathrm{w}} = 318{,}200\mathrm{kg}$	Wall mass
	$\varepsilon = 0.1\mathrm{K}$	RLMTD threshold
Primary heat exchanger	$A_{\mathrm{cf}} = A_{\mathrm{hf}} = 13{,}060\mathrm{m}^2$	Fluid-wall heat transfer area
	$U_{\mathrm{cf}} = U_{\mathrm{hf}} = 4{,}000\mathrm{W/(m^2\ K)}$	Fluid-wall heat transfer coefficient
	$V_{\mathrm{cf}} = 0.4195\mathrm{m}^3$	Cold fluid side volume
	$V_{\mathrm{hf}} = 0.2333\mathrm{m}^3$	Hot fluid side volume
	$\Delta p_{\mathrm{des,cf}} = 0.01010\mathrm{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\mathrm{des,hf}} = 0.8182\mathrm{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\mathrm{des,cf}} = 3{,}115\mathrm{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\mathrm{des,hf}} = 1{,}732\mathrm{kg/s}$	Design mass flow rate of hot fluid
	$M_{\mathrm{w}} = 78{,}340\mathrm{kg}$	Wall mass
	$\varepsilon = 0.01\mathrm{K}$	RLMTD threshold

Component	Parameter	Description
Cold molten salt pump	$\dot{m}_{\text{des}}$ =3,115kg/s	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{\text{in,des}} = 1\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1.010\text{bar}$	Design outlet pressure
Hot molten salt pump	$\dot{m}_{\text{des}}$ =2,690kg/s	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{\text{in,des}} = 1\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1.031\text{bar}$	Design outlet pressure
Feedwater pump	$\dot{m}_{\text{des}}$ =333.9kg/s	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{\text{in,des}} = 0.0508\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 168.9\text{bar}$	Design outlet pressure
Evaporator bypass valve	$\dot{m}_{\text{des}}$ =2690kg/s	Design mass flow rate
	$p_{\text{in,des}} = 1.020\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1.010\text{bar}$	Design outlet pressure
	$T_{\text{in,des}} = 450°\text{C}$	Design inlet temperature
Preheater bypass valve	$\dot{m}_{\text{des}}$ =2690kg/s	Design mass flow rate
	$p_{\text{in,des}} = 1.010\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1\text{bar}$	Design outlet pressure
	$T_{\text{in,des}} = 350°\text{C}$	Design inlet temperature

Table B.3: Parameters of individual component models in the dynamic process model of Case Study II in section 7.1. Condenser and evaporator parameters extrapolated linearly with area from [254, 82]. Other heat exchange parameters extrapolated linearly with area from [108, 9]. Turbine and condenser pressure from [184]. Tank heat transfer coefficient from [197].

Parameter	Description	Parameter	Description
$U_{\text{PH}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of preheater	$U_{\text{EVP}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of evaporator
$U_{\text{SH}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of superheater	$U_{\text{CND}} =$ 5,000W/(m^2 K)	Heat transfer coefficient of condenser
$U_{\text{PHX}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of primary heat exchanger	$\dot{Q}_{\text{helium}} = 1{,}350\text{MW}$	Thermal power of helium coolant
$p_{\text{turbine,in}} = 165.5\text{bar}$	Turbine inlet pressure	$p_{\text{CND}} = 0.0508\text{bar}$	Condenser pressure
$T_{\text{ms,hot}} = 580°\text{C}$	Hot tank temperature	$T_{\text{ms,cold}} = 300°\text{C}$	Cold tank temperature
$c_{\text{p,ms}} =$ 1,494.6J/(kg K)	Molten salt specific heat capacity	$\rho_{\text{ms}} = 1{,}899.2\text{kg/m}^3$	Molten salt density
$\tau_{\text{pulse}} = 125\text{min}$	Pulse period	$\tau_{\text{dwell}} = 20\text{min}$	Dwell period
$T_{\text{helium,in}} = 600°\text{C}$	Helium inlet temperature	$T_{\text{helium,out}} = 450°\text{C}$	Helium outlet temperature
$c_{\text{p,helium}} =$ 5,196.5J/(kg K)	Helium specific heat capacity	$T_{\text{cw,in}} = 15°\text{C}$	Cooling water inlet temperature
$T_{\text{cw,out}} = 25°\text{C}$	Cooling water outlet temperature	$p_{\text{cw,out}} = 2\text{bar}$	Cooling water outlet pressure
$\Delta T_{\text{pinch}} = 10\text{K}$	Pinch temperature difference	$\Delta p_{\%} = 0.01$	Fractional pressure drop
$\eta_{\text{is,turbine}} = 0.9$	Turbine isentropic efficiency	$\eta_{\text{is,pump}} = 0.9$	Pump isentropic efficiency

Table B.2: Arguments of the steady state process model of Case Study II in section 7.1.2. Turbine and condenser pressure from [184].

Label	Controlled variable	Manipulated variable	Set point	K_c	τ_I (s)	τ_D (s)	Minimum output	Maximum output
PI-a	Evaporator liquid fraction, $\frac{V_{L,EVP}}{V_{wf,EVP}}$	Mass flow rate of feedwater, $\dot{m}_{14}$	0.5	500kg/s	500	-	3.38kg/s	676kg/s
PI-b	Condenser pressure, p_{12}	Mass flow rate of cooling water, $\dot{m}_{18}$	0.0508bar	-50kg/(s Pa)	50	-	162kg/s	32,314kg/s
PI-c	Turbine power, $\dot{W}_{turbine}$	Partial arc admission valve opening, θ_{arc}	See section 7.2	$0.0001W^{-1}$	60	-	0	1
PI-d	Turbine inlet pressure, p_{11}	Mass flow rate of hot molten salt, $\dot{m}_4$	165.5bar	0.001kg/(s Pa)	600	-	26.9kg/s	5,381kg/s
PI-e	Preheater outlet water temperature, T_{15}	Preheater molten salt bypass valve opening, Ω_8	300°C	$-2K^{-1}$	20	-	0	1
PI-f	Helium return temperature, T_2	Mass flow rate of cold molten salt, $\dot{m}_{10}$	450°C	-500kg/(s K)	60	-	31.2kg/s	6,230kg/s
PI-g	Cold tank temperature, $T_{ms,cold}$	Evaporator molten salt bypass valve opening, Ω_6	341°C	$2K^{-1}$	20	-	0	1
PID-a	Turbine inlet pressure, p_{11}	Mass flow rate of hot molten salt, $\dot{m}_4$	165.5bar	0.001kg/(s Pa)*	600*	1*	26.9kg/s	5,381kg/s

Table B.4: Outline of the process controllers for Case Study II in section 7.1. Please refer to figure 7.1 or section 7.1.1 for controller labels. All models are *LimPID*. *values for the initial design only

Appendix C

Model Parameters for Case Study III

Algorithm 3 Steady state model of the power conversion system in Case Study III. The results are used to parameterise the dynamic process model. HTHX, MTHX, LTHX, DH, SH, EVP, PH, CND are the high temperature heat exchanger, medium temperature heat exchanger, low temperature heat exchanger, dwell heater, super heater, evaporator, preheater and condenser. Medium thermodynamic properties are calculated using CoolProp if equation not specified [24].

Require: U_{PH}, U_{EVP}, U_{SH}, U_{CND}, U_{HTHX}, U_{MTHX}, U_{LTHX}, U_{DH}, $\dot{Q}_{\mathrm{HTC}}$, $\dot{Q}_{\mathrm{MTC}}$, $\dot{Q}_{\mathrm{LTC}}$, $p_{\mathrm{turbine,in}}$, p_{CND}, $T_{\mathrm{ms,hot}}$, $T_{\mathrm{ms,cold}}$, $c_{\mathrm{p,ms}}$, ρ_{ms}, τ_{pulse}, τ_{dwell}, $T_{\mathrm{HTC,in}}$, $T_{\mathrm{HTC,out}}$, p_{HTC}, $T_{\mathrm{MTC,in}}$, $T_{\mathrm{MTC,out}}$, p_{MTC}, $T_{\mathrm{LTC,in}}$, $T_{\mathrm{LTC,out}}$, p_{LTC}, $T_{\mathrm{cw,in}}$, $T_{\mathrm{cw,out}}$, $p_{\mathrm{cw,out}}$, $\Delta T_{\mathrm{pinch}}$, ΔT_{tank}, $\Delta p_{\%}$, $\eta_{\mathrm{is,turbine}}$, $\eta_{\mathrm{is,pump}}$

Iterate over Rankine loop to estimate balance between pulse and dwell:

$\dot{W}_{\mathrm{turbine}} \leftarrow 1000\mathrm{MW}$

while *true* **do**

  High temperature heat exchanger (HTHX) calculations during pulse:

  $\Delta T_{\mathrm{LMTD,HTHX}} \leftarrow \frac{(T_{\mathrm{HTC,out}}-T_{\mathrm{ms,cold}})-(T_{\mathrm{HTC,in}}-(T_{\mathrm{ms,hot}}+\Delta T_{\mathrm{tank}}))}{\ln(T_{\mathrm{HTC,out}}-T_{\mathrm{ms,cold}})-\ln(T_{\mathrm{HTC,in}}-(T_{\mathrm{ms,hot}}+\Delta T_{\mathrm{tank}}))}$ ▷ Eq. 4.21a

  $A_{\mathrm{HTHX}} \leftarrow \frac{\dot{Q}_{\mathrm{HTC}}}{U_{\mathrm{HTHX}}\Delta T_{\mathrm{LMTD,HTHX}}}$ ▷ Eq. 4.50

  $\dot{m}_{\mathrm{HTC}} \leftarrow \frac{\dot{Q}_{\mathrm{HTC}}}{h(p_{\mathrm{HTC}},T_{\mathrm{HTC,in}})-h(p_{\mathrm{HTC}},T_{\mathrm{HTC,out}})}$

  $\dot{m}_{\mathrm{ms,HTHX}} \leftarrow \frac{\dot{Q}_{\mathrm{HTC}}}{c_{\mathrm{p,ms}}((T_{\mathrm{ms,hot}}+\Delta T_{\mathrm{tank}})-T_{\mathrm{ms,cold}})}$

  Turbine calculations:

  $T_{\mathrm{turbine,in}} \leftarrow T_{\mathrm{ms,hot}} - \Delta T_{\mathrm{pinch}}$

  $h_{\mathrm{turbine,in}} \leftarrow h(p_{\mathrm{turbine,in}}, T_{\mathrm{turbine,in}})$

  $\Delta h_{\mathrm{is,turbine}} \leftarrow h(p_{\mathrm{CND}}, s(p_{\mathrm{turbine,in}}, T_{\mathrm{turbine,in}})) - h_{\mathrm{turbine,in}}$ ▷ Eq. 4.26

  $h_{\mathrm{turbine,out}} \leftarrow h_{\mathrm{turbine,in}} + \eta_{\mathrm{is,turbine}}\Delta h_{\mathrm{is,turbine}}$ ▷ Eq. 4.27

  $\dot{m}_{\mathrm{rankine}} \leftarrow \frac{\dot{W}_{\mathrm{turbine}}}{h_{\mathrm{turbine,in}}-h_{\mathrm{turbine,out}}}$ ▷ Eq. 4.28

15: Condenser (CND) calculations:
16: $h_{\mathrm{CND,out}} \leftarrow h_{\mathrm{saturated,L}}(p_{\mathrm{CND}})$
17: $T_{\mathrm{CND,out}} \leftarrow T_{\mathrm{saturated}}(p_{\mathrm{CND}})$
18: $\dot{Q}_{\mathrm{CND}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{CND,out}} - h_{\mathrm{turbine,out}})$
19: $\Delta T_{\mathrm{LMTD,CND}} \leftarrow \frac{(T_{\mathrm{CND,out}}-T_{\mathrm{cw,out}})-(T_{\mathrm{CND,out}}-T_{\mathrm{cw,in}})}{\ln(T_{\mathrm{CND,out}}-T_{\mathrm{cw,out}})-\ln(T_{\mathrm{CND,out}}-T_{\mathrm{cw,in}})}$ ▷ Eq. 4.21a
20: $A_{\mathrm{CND}} \leftarrow \frac{-\dot{Q}_{\mathrm{CND}}}{U_{\mathrm{CND}}\Delta T_{\mathrm{LMTD,CND}}}$
21: $\dot{m}_{\mathrm{cw}} \leftarrow \frac{-\dot{Q}_{\mathrm{CND}}}{h(p_{\mathrm{cw,out}},T_{\mathrm{cw,out}})-h(p_{\mathrm{cw,out}},T_{\mathrm{cw,in}})}$
22: Pump calculations:
23: $p_{\mathrm{pump,out}} \leftarrow \frac{p_{\mathrm{turbine,in}}}{(1-\Delta p_{\%})^4}$
24: $\Delta h_{\mathrm{is,pump}} \leftarrow h(p_{\mathrm{pump,out}}, s(p_{\mathrm{CND}}, h_{\mathrm{pump,in}})) - h_{\mathrm{pump,in}}$ ▷ Eq. 4.26
25: $h_{\mathrm{pump,out}} \leftarrow h_{\mathrm{CND,out}} + \eta_{\mathrm{is,pump}}\Delta h_{\mathrm{is,pump}}$
26: $T_{\mathrm{pump,out}} \leftarrow T(p_{\mathrm{pump,out}}, h_{\mathrm{pump,out}})$
27: Low temperature heat exchanger (LTHX) calculations during pulse:
28: $h_{\mathrm{LTHX,out}} \leftarrow h_{\mathrm{pump,out}} + \frac{\dot{Q}_{\mathrm{LTC}}}{\dot{m}_{\mathrm{rankine}}}$
29: $T_{\mathrm{LTHX,out}} \leftarrow T(p_{\mathrm{pump,out}}(1-\Delta p_{\%}), h_{\mathrm{LTHX,out}})$
30: $\dot{m}_{\mathrm{LTC}} \leftarrow \frac{\dot{Q}_{\mathrm{LTC}}}{h(p_{\mathrm{LTC}},T_{\mathrm{LTC,in}})-h(p_{\mathrm{LTC}},T_{\mathrm{LTC,out}})}$
31: $\Delta T_{\mathrm{LMTD,LTHX}} \leftarrow \frac{(T_{\mathrm{LTC,out}}-T_{\mathrm{pump,out}})-(T_{\mathrm{LTC,in}}-T_{\mathrm{LTHX,out}})}{\ln(T_{\mathrm{LTC,out}}-T_{\mathrm{pump,out}})-\ln(T_{\mathrm{LTC,in}}-T_{\mathrm{LTHX,out}})}$ ▷ Eq. 4.21a
32: $A_{\mathrm{LTHX}} \leftarrow \frac{\dot{Q}_{\mathrm{LTC}}}{U_{\mathrm{LTHX}}\Delta T_{\mathrm{LMTD,LTHX}}}$
33: Medium temperature heat exchanger (MTHX) calculations during pulse:
34: $h_{\mathrm{MTHX,out}} \leftarrow h_{\mathrm{LTHX,out}} + \frac{\dot{Q}_{\mathrm{MTC}}}{\dot{m}_{\mathrm{rankine}}}$
35: $T_{\mathrm{MTHX,out}} \leftarrow T(p_{\mathrm{pump,out}}(1-\Delta p_{\%})^2, h_{\mathrm{MTHX,out}})$
36: $\dot{m}_{\mathrm{MTC}} \leftarrow \frac{\dot{Q}_{\mathrm{MTC}}}{h(p_{\mathrm{MTC}},T_{\mathrm{MTC,in}})-h(p_{\mathrm{MTC}},T_{\mathrm{MTC,out}})}$
37: $\Delta T_{\mathrm{LMTD,MTHX}} \leftarrow \frac{(T_{\mathrm{MTC,out}}-T_{\mathrm{LTHX,out}})-(T_{\mathrm{MTC,in}}-T_{\mathrm{MTHX,out}})}{\ln(T_{\mathrm{MTC,out}}-T_{\mathrm{LTHX,out}})-\ln(T_{\mathrm{MTC,in}}-T_{\mathrm{MTHX,out}})}$ ▷ Eq. 4.21a
38: $A_{\mathrm{MTHX}} \leftarrow \frac{\dot{Q}_{\mathrm{MTC}}}{U_{\mathrm{MTHX}}\Delta T_{\mathrm{LMTD,MTHX}}}$
39: Dwell heater (DH) calculations during dwell:
40: $T_{\mathrm{DH,out}} \leftarrow T_{\mathrm{MTHX,out}}$
41: $\dot{Q}_{\mathrm{DH}} \leftarrow \dot{m}_{\mathrm{rankine}}(h(p_{\mathrm{pump,out}}(1-\Delta p_{\%})^2, T_{\mathrm{DH,out}}) - h_{\mathrm{pump,out}})$
42: $T_{\mathrm{ms,DH,out}} \leftarrow T_{\mathrm{ms,cold}} + \Delta T_{\mathrm{tank}}$
43: $\dot{m}_{\mathrm{ms,DH}} \leftarrow \frac{\dot{Q}_{\mathrm{DH}}}{c_{\mathrm{p,ms}}(T_{\mathrm{ms,hot}}-T_{\mathrm{ms,DH,out}})}$
44: $\Delta T_{\mathrm{LMTD,DH}} \leftarrow \frac{(T_{\mathrm{ms,DH,out}}-T_{\mathrm{pump,out}})-(T_{\mathrm{ms,hot}}-T_{\mathrm{DH,out}})}{\ln(T_{\mathrm{ms,DH,out}}-T_{\mathrm{pump,out}})-\ln(T_{\mathrm{ms,hot}}-T_{\mathrm{DH,out}})}$ ▷ Eq. 4.21a
45: $A_{\mathrm{DH}} \leftarrow \frac{\dot{Q}_{\mathrm{DH}}}{U_{\mathrm{DH}}\Delta T_{\mathrm{LMTD,DH}}}$

Preheater (PH) calculations:
$T_{\mathrm{PH,out}} \leftarrow T_{\mathrm{saturated,L}}(p_{\mathrm{pump,out}}(1-\Delta p_{\%})^3) - \Delta T_{\mathrm{pinch}}$
$h_{\mathrm{PH,out}} \leftarrow h(p_{\mathrm{pump,out}}(1-\Delta p_{\%})^3, T_{\mathrm{PH,out}})$
$\dot{Q}_{\mathrm{PH}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{PH,out}} - h_{\mathrm{MTHX,out}})$
$h_{\mathrm{EVP,out}} \leftarrow h_{\mathrm{saturated,V}}(p_{\mathrm{pump,out}}(1-\Delta p_{\%}))$
Evaporator (EVP) calculations:
$T_{\mathrm{EVP,out}} \leftarrow T_{\mathrm{saturated}}(p_{\mathrm{pump,out}}(1-\Delta p_{\%}))$
$Q_{\mathrm{EVP}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{EVP,out}} - h_{\mathrm{PH,out}})$
$Q_{\mathrm{SH}} \leftarrow \dot{m}_{\mathrm{rankine}}(h_{\mathrm{turbine,in}} - h_{\mathrm{EVP,out}})$
$\dot{m}_{\mathrm{ms,SH}} \leftarrow \frac{\dot{Q}_{\mathrm{PH}}+\dot{Q}_{\mathrm{EVP}}+\dot{Q}_{\mathrm{SH}}}{c_{\mathrm{p,ms}}(T_{\mathrm{ms,hot}}-T_{\mathrm{ms,cold}})}$
if $\tau_{\mathrm{pulse}}(\dot{m}_{\mathrm{ms,HTHX}} - \dot{m}_{\mathrm{ms,SH}}) > \tau_{\mathrm{dwell}}(\dot{m}_{\mathrm{ms,SH}} + \dot{m}_{\mathrm{ms,DH}})$ **then**
    break
end if
$\dot{W}_{\mathrm{turbine}} \leftarrow \dot{W}_{\mathrm{turbine}} - 0.1\mathrm{MW}$
end while
$T_{\mathrm{ms,SH,out}} \leftarrow T_{\mathrm{ms,hot}} - \frac{\dot{Q}_{\mathrm{SH}}}{c_{\mathrm{p,ms}}\dot{m}_{\mathrm{ms,SH}}}$
$\Delta T_{\mathrm{LMTD,SH}} \leftarrow \frac{(T_{\mathrm{ms,hot}}-T_{\mathrm{turbine,in}})-(T_{\mathrm{ms,SH,out}}-T_{\mathrm{EVP,out}})}{\ln(T_{\mathrm{ms,hot}}-T_{\mathrm{turbine,in}})-\ln(T_{\mathrm{ms,SH,out}}-T_{\mathrm{EVP,out}})}$ ▷ Eq. 4.21a
$A_{\mathrm{SH}} \leftarrow \frac{\dot{Q}_{\mathrm{SH}}}{U_{\mathrm{SH}}\Delta T_{\mathrm{LMTD,SH}}}$
$T_{\mathrm{ms,EVP,out}} \leftarrow T_{\mathrm{ms,SH,out}} - \frac{\dot{Q}_{\mathrm{EVP}}}{c_{\mathrm{p,ms}}\dot{m}_{\mathrm{ms,SH}}}$
$\Delta T_{\mathrm{LMTD,EVP}} \leftarrow \frac{(T_{\mathrm{ms,SH,out}}-T_{\mathrm{EVP,out}})-(T_{\mathrm{ms,EVP,out}}-T_{\mathrm{PH,out}})}{\ln(T_{\mathrm{ms,SH,out}}-T_{\mathrm{EVP,out}})-\ln(T_{\mathrm{ms,EVP,out}}-T_{\mathrm{PH,out}})}$ ▷ Eq. 4.21a
$A_{\mathrm{EVP}} \leftarrow \frac{\dot{Q}_{\mathrm{EVP}}}{U_{\mathrm{EVP}}\Delta T_{\mathrm{LMTD,EVP}}}$
$\Delta T_{\mathrm{LMTD,PH}} \leftarrow \frac{(T_{\mathrm{ms,EVP,out}}-T_{\mathrm{PH,out}})-(T_{\mathrm{ms,cold}}-T_{\mathrm{pump,out}})}{\ln(T_{\mathrm{ms,EVP,out}}-T_{\mathrm{PH,out}})-\ln(T_{\mathrm{ms,cold}}-T_{\mathrm{pump,out}})}$ ▷ Eq. 4.21a
$A_{\mathrm{PH}} \leftarrow \frac{\dot{Q}_{\mathrm{PH}}}{U_{\mathrm{PH}}\Delta T_{\mathrm{LMTD,PH}}}$
Estimate size of molten salt tanks, allowing for fill constraints
$V_{\mathrm{tank}} \leftarrow 1.25\frac{\tau_{\mathrm{pulse}}(\dot{m}_{\mathrm{ms,HTHX}}-\dot{m}_{\mathrm{ms,SH}})}{\rho_{\mathrm{ms}}}$
return $\dot{m}_{\mathrm{cw}}$, $\dot{m}_{\mathrm{HTC}}$, $\dot{m}_{\mathrm{MTC}}$, $\dot{m}_{\mathrm{LTC}}$, $\dot{m}_{\mathrm{ms,HTHX}}$, $\dot{m}_{\mathrm{ms,SH}}$, $\dot{m}_{\mathrm{ms,DH}}$,$\dot{m}_{\mathrm{rankine}}$, A_{HTHX}, A_{MTHX}, A_{LTHX}, A_{SH}, A_{EVP}, A_{PH}, A_{CND}, A_{DH}, V_{tank}

Unit	Initial equations
High temperature heat exchanger	$T_{\mathrm{hf,out}}\|_{t=0} = 450°C$, $T_{\mathrm{cf,out}}\|_{t=0} = 552°C$, $\bar{T}_{\mathrm{w}}\|_{t=0} = 501°\mathrm{C}$, $\Delta T_{\mathrm{w}}\|_{t=0} = 51\mathrm{K}$
Medium temperature heat exchanger	$T_{\mathrm{hf,out}}\|_{t=0} = 225°C$, $T_{\mathrm{cf,out}}\|_{t=0} = 267°C$, $\bar{T}_{\mathrm{w}}\|_{t=0} = 246°\mathrm{C}$, $\Delta T_{\mathrm{w}}\|_{t=0} = 21\mathrm{K}$
Low temperature heat exchanger	$T_{\mathrm{hf,out}}\|_{t=0} = 275°C$, $T_{\mathrm{cf,out}}\|_{t=0} = 290°C$, $\bar{T}_{\mathrm{w}}\|_{t=0} = 283°\mathrm{C}$, $\Delta T_{\mathrm{w}}\|_{t=0} = 7.5\mathrm{K}$
Dwell heater	$T_{\mathrm{hf,out}}\|_{t=0} = 360°C$, $T_{\mathrm{cf,out}}\|_{t=0} = 267°C$, $\bar{T}_{\mathrm{w}}\|_{t=0} = 313°\mathrm{C}$, $\Delta T_{\mathrm{w}}\|_{t=0} = -47\mathrm{K}$
Preheater	$\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{hf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{cf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t} \bar{T}_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$, $\frac{\mathrm{d}}{\mathrm{d}t} \Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$
Evaporator	$T_{\mathrm{hf,out}}\|_{t=0} = 360°\mathrm{C}$, $p_{\mathrm{wf}}\|_{t=0} = 167\mathrm{bar}$, $\left.\frac{V_{\mathrm{L}}}{V_{\mathrm{wf}}}\right\|_{t=0} = 0.5$, $\bar{T}_{\mathrm{w}}\|_{t=0} = 360°\mathrm{C}$, $\Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K}$
Superheater	$\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{hf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t} h_{\mathrm{cf,out}}\|_{t=0} = 0\mathrm{J/(kg\ s)}$, $\frac{\mathrm{d}}{\mathrm{d}t} \bar{T}_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$, $\frac{\mathrm{d}}{\mathrm{d}t} \Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K/s}$
Condenser	$T_{\mathrm{cf,out}}\|_{t=0} = 25°C$, $p_{\mathrm{wf}}\|_{t=0} = 0.0508\mathrm{bar}$, $\left.\frac{V_{\mathrm{L}}}{V_{\mathrm{wf}}}\right\|_{t=0} = 0.5$, $\bar{T}_{\mathrm{w}}\|_{t=0} = 25°\mathrm{C}$, $\Delta T_{\mathrm{w}}\|_{t=0} = 0\mathrm{K}$
Hot tank	$T_{\mathrm{ms}}\|_{t=0} = 550°C$, $\left.\frac{V_{\mathrm{ms}}}{V_{\mathrm{tank}}}\right\|_{t=0} = 0.1$
Cold tank	$T_{\mathrm{ms}}\|_{t=0} = 350°C$, $\left.\frac{V_{\mathrm{ms}}}{V_{\mathrm{tank}}}\right\|_{t=0} = 0.9$
PI-a	$y\|_{t=0} = 479\mathrm{kg/s}$
PI-b	$y\|_{t=0} = 32{,}217\mathrm{kg/s}$
PI-c	$y\|_{t=0} = 1$
PI-d	$y\|_{t=0} = 3{,}627\mathrm{kg/s}$
PI-e	$y\|_{t=0} = 0$
PI-f	$y\|_{t=0} = 4{,}470\mathrm{kg/s}$
PI-g	$y\|_{t=0} = 1.64\mathrm{kg/s}$
PI-h	$y\|_{t=0} = 0.999$
PI-i	$y\|_{t=0} = 0$
PI-j	$y\|_{t=0} = 0$

Table C.1: Initial equations for the dynamic models in Case Study III. Please refer to figure 7.12 or section 7.4.1 for controller labels.

Parameter	Description	Parameter	Description
$U_{\mathrm{PH}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of preheater	$U_{\mathrm{EVP}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of evaporator
$U_{\mathrm{SH}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of superheater	$U_{\mathrm{CND}} =$ 5,000W/(m^2 K)	Heat transfer coefficient of condenser
$U_{\mathrm{HTHX}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of high temperature heat exchanger	$U_{\mathrm{MTHX}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of medium temperature heat exchanger
$U_{\mathrm{LTHX}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of low temperature heat exchanger	$U_{\mathrm{DH}} =$ 2,000W/(m^2 K)	Heat transfer coefficient of dwell heater
$\dot{Q}_{\mathrm{HTC}} =$ 1,350MW	Thermal power of high temperature coolant	$\dot{Q}_{\mathrm{MTC}} = 215\mathrm{MW}$	Thermal power of medium temperature coolant
$\dot{Q}_{\mathrm{LTC}} = 270\mathrm{MW}$	Thermal power of low temperature coolant	$p_{\mathrm{HTC}} = 81\mathrm{bar}$	High temperature coolant pressure
$p_{\mathrm{MTC}} = 99\mathrm{bar}$	Medium temperature coolant pressure	$p_{\mathrm{LTC}} = 149\mathrm{bar}$	Low temperature coolant pressure
$T_{\mathrm{HTC,in}} = 600°\mathrm{C}$	High temperature coolant inlet temperature	$T_{\mathrm{HTC,out}} = 450°\mathrm{C}$	High temperature coolant outlet temperature

Parameter	Description	Parameter	Description
$T_{\text{MTC,in}} = 450°\text{C}$	Medium temperature coolant inlet temperature	$T_{\text{MTC,out}} = 225°\text{C}$	Medium temperature coolant outlet temperature
$T_{\text{LTC,in}} = 300°\text{C}$	Low temperature coolant inlet temperature	$T_{\text{LTC,out}} = 275°\text{C}$	Low temperature coolant outlet temperature
$p_{\text{turbine,in}} = 165.5\text{bar}$	Turbine inlet pressure	$p_{\text{CND}} = 0.0508\text{bar}$	Condenser pressure
$T_{\text{ms,hot}} = 550°\text{C}$	Hot tank temperature	$T_{\text{ms,cold}} = 350°\text{C}$	Cold tank temperature
$c_{\text{p,ms}} = 1{,}495\text{J/(kg K)}$	Molten salt specific heat capacity	$\rho_{\text{ms}} = 1{,}988\text{kg/m}^3$	Molten salt density
$\tau_{\text{pulse}} = 125\text{min}$	Pulse period	$\tau_{\text{dwell}} = 20\text{min}$	Dwell period
$T_{\text{cw,in}} = 15°\text{C}$	Cooling water inlet temperature	$T_{\text{cw,out}} = 25°\text{C}$	Cooling water outlet temperature
$p_{\text{cw,out}} = 2\text{bar}$	Cooling water outlet pressure	$\Delta T_{\text{tank}} = 2\text{K}$	Tank inlet temperature excess
$\Delta T_{\text{pinch}} = 10\text{K}$	Pinch temperature difference	$\Delta p_{\%} = 0.01$	Fractional pressure drop
$\eta_{\text{is,turbine}} = 0.9$	Turbine isentropic efficiency	$\eta_{\text{is,pump}} = 0.9$	Pump isentropic efficiency

Table C.2: Arguments of the steady state process model of Case Study III in section 7.4.2. Turbine and condenser pressure from [184].

Component	Parameter	Description
Cold tank & hot tank	$U_{\text{ms,amb}} = 0.2\text{W/(m}^2\text{ K)}$	Molten salt-ambient heat transfer coefficient
	$T_{\text{amb}} = 30°\text{C}$	Ambient temperature
	$p_{\text{tank}} = 1\text{bar}$	Tank pressure
Turbine	$\eta_{\text{is}} = 0.9$	Isentropic efficiency
	$\eta_{\text{mech}} = 1$	Mechanical efficiency
	$\dot{m}_{\text{des}} = 503.4\text{kg/s}$	Design mass flow rate
	$T_{\text{in,des}} = 539.9°\text{C}$	Design inlet temperature
	$p_{\text{in,des}} = 165.5\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 0.0508\text{bar}$	Design outlet pressure
All heat exchangers	$c_{\text{w}} = 500\text{J/(kg K)}$	Wall specific heat capacity
	$\xi = 5\text{K}^{-1}$	RLMTD steepness
Preheater	$A_{\text{cf}} = A_{\text{hf}} = 1{,}560\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{cf}} = U_{\text{hf}} = 4{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{cf}} = 0.009106\text{m}^3$	Cold fluid side volume
	$V_{\text{hf}} = 0.06890\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,cf}} = 1.706\text{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\text{des,hf}} = 0.01031\text{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\text{des,cf}} = 479.4\text{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\text{des,hf}} = 3{,}627\text{kg/s}$	Design mass flow rate of hot fluid
	$M_{\text{w}} = 9{,}361\text{kg}$	Wall mass
	$\varepsilon = 0.01\text{K}$	RLMTD threshold

Component	Parameter	Description
Evaporator	$A_{\mathrm{wf}} = A_{\mathrm{hf}} = 3{,}162\mathrm{m}^2$	Fluid-wall heat transfer area
	$U_{\mathrm{wf}} = U_{\mathrm{hf}} = 4{,}000\mathrm{W/(m^2\ K)}$	Fluid-wall heat transfer coefficient
	$V_{\mathrm{wf}} = 79.71\mathrm{m}^3$	Working fluid side volume
	$V_{\mathrm{hf}} = 28.75\mathrm{m}^3$	Hot fluid side volume
	$\Delta p_{\mathrm{des,hf}} = 0.01020\mathrm{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\mathrm{des,hf}} = 3{,}627\mathrm{kg/s}$	Design mass flow rate of hot fluid
	$M_{\mathrm{w}} = 93{,}010\mathrm{kg}$	Wall mass
	$\varepsilon = 0.1\mathrm{K}$	RLMTD threshold
Superheater	$A_{\mathrm{cf}} = A_{\mathrm{hf}} = 4{,}491\mathrm{m}^2$	Fluid-wall heat transfer area
	$U_{\mathrm{cf}} = U_{\mathrm{hf}} = 4{,}000\mathrm{W/(m^2\ K)}$	Fluid-wall heat transfer coefficient
	$V_{\mathrm{cf}} = 0.02621\mathrm{m}^3$	Cold fluid side volume
	$V_{\mathrm{hf}} = 0.1983\mathrm{m}^3$	Hot fluid side volume
	$\Delta p_{\mathrm{des,cf}} = 1.672\mathrm{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\mathrm{des,hf}} = 0.01020\mathrm{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\mathrm{des,cf}} = 479.4\mathrm{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\mathrm{des,hf}} = 3{,}627\mathrm{kg/s}$	Design mass flow rate of hot fluid
	$M_{\mathrm{w}} = 26{,}940\mathrm{kg}$	Wall mass
	$\varepsilon = 0.01\mathrm{K}$	RLMTD threshold

Component	Parameter	Description
Condenser	$A_{\text{wf}} = A_{\text{cf}} = 52{,}600\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{wf}} = U_{\text{cf}} = 10{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{wf}} = 1{,}326\text{m}^3$	Working fluid side volume
	$V_{\text{cf}} = 478.2\text{m}^3$	Cold fluid side volume
	$\Delta p_{\text{des,cf}} = 0.02020\text{bar}$	Design pressure drop of cold fluid
	$\dot{m}_{\text{des,cf}} = 32{,}220\text{kg/s}$	Design mass flow rate of cold fluid
	$M_{\text{w}} = 1{,}547{,}000\text{kg}$	Wall mass
	$\varepsilon = 0.1\text{K}$	RLMTD threshold
High temperature heat exchanger	$A_{\text{cf}} = A_{\text{hf}} = 9{,}527\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{cf}} = U_{\text{hf}} = 4{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{cf}} = 0.3432\text{m}^3$	Cold fluid side volume
	$V_{\text{hf}} = 0.1332\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,cf}} = 0.01010\text{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\text{des,hf}} = 0.8182\text{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\text{des,cf}} = 4{,}470\text{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\text{des,hf}} = 1{,}735\text{kg/s}$	Design mass flow rate of hot fluid
	$M_{\text{w}} = 57{,}160\text{kg}$	Wall mass
	$\varepsilon = 0.01\text{K}$	RLMTD threshold

Component	Parameter	Description
Medium temperature heat exchanger	$A_{cf} = A_{hf} = 990.9 m^2$	Fluid-wall heat transfer area
	$U_{cf} = U_{hf}$ =4,000W/(m^2 K)	Fluid-wall heat transfer coefficient
	$V_{cf} = 0.03579 m^3$	Cold fluid side volume
	$V_{hf} = 0.01375 m^3$	Hot fluid side volume
	$\Delta p_{des,cf} = 1.723$bar	Design pressure drop of cold fluid
	$\Delta p_{des,hf} = 1$bar	Design pressure drop of hot fluid
	$\dot{m}_{des,cf} = 479.4$kg/s	Design mass flow rate of cold fluid
	$\dot{m}_{des,hf} = 184.2$kg/s	Design mass flow rate of hot fluid
	$M_w = 5{,}945$kg	Wall mass
	$\varepsilon = 0.01$K	RLMTD threshold
Cold molten salt pump	$\dot{m}_{des}$ =3,115kg/s	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{in,des} = 1$bar	Design inlet pressure
	$p_{out,des} = 1.010$bar	Design outlet pressure
Superheater molten salt pump	$\dot{m}_{des}$ =3,627kg/s	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{in,des} = 1$bar	Design inlet pressure
	$p_{out,des} = 1.041$bar	Design outlet pressure
Split valve	$\dot{m}_{des}$ =479.4kg/s	Design mass flow rate
	$p_{in,des} = 175.8$bar	Design inlet pressure
	$p_{out,des} = 174.0$bar	Design outlet pressure
	$T_{in,des} = 33.16$°C	Design inlet temperature

Component	Parameter	Description
Low temperature heat exchanger	$A_{cf} = A_{hf} = 743.1 m^2$	Fluid-wall heat transfer area
	$U_{cf} = U_{hf}$ =4,000W/(m^2 K)	Fluid-wall heat transfer coefficient
	$V_{cf} = 0.006994 m^3$	Cold fluid side volume
	$V_{hf} = 0.03016 m^3$	Hot fluid side volume
	$\Delta p_{des,cf} = 1.740$bar	Design pressure drop of cold fluid
	$\Delta p_{des,hf} = 1.505$bar	Design pressure drop of hot fluid
	$\dot{m}_{des,cf} = 479.4$kg/s	Design mass flow rate of cold fluid
	$\dot{m}_{des,hf} = 2{,}067$kg/s	Design mass flow rate of hot fluid
	$M_w = 4{,}458$kg	Wall mass
	$\varepsilon = 0.01$K	RLMTD threshold
Preheater bypass valve	$\dot{m}_{des}$ =359,104kg/s	Design mass flow rate
	$p_{in,des} = 1.010$bar	Design inlet pressure
	$p_{out,des} = 1$bar	Design outlet pressure
	$T_{in,des} = 400°$C	Design inlet temperature
MTHX bypass valve	$\dot{m}_{des}$ =1,658kg/s	Design mass flow rate
	$p_{in,des} = 100$bar	Design inlet pressure
	$p_{out,des} = 99$bar	Design outlet pressure
	$T_{in,des} = 450°$C	Design inlet temperature
LTHX bypass valve	$\dot{m}_{des}$ =18,610kg/s	Design mass flow rate
	$p_{in,des} = 150.5$bar	Design inlet pressure
	$p_{out,des} = 149$bar	Design outlet pressure
	$T_{in,des} = 300°$C	Design inlet temperature

Component	Parameter	Description
Dwell heater	$A_{\text{cf}} = A_{\text{hf}} = 806.3\text{m}^2$	Fluid-wall heat transfer area
	$U_{\text{cf}} = U_{\text{hf}} = 4{,}000\text{W/(m}^2\text{ K)}$	Fluid-wall heat transfer coefficient
	$V_{\text{cf}} = 0.009126\text{m}^3$	Cold fluid side volume
	$V_{\text{hf}} = 0.03119\text{m}^3$	Hot fluid side volume
	$\Delta p_{\text{des,cf}} = 0.01740\text{bar}$	Design pressure drop of cold fluid
	$\Delta p_{\text{des,hf}} = 0.01010\text{bar}$	Design pressure drop of hot fluid
	$\dot{m}_{\text{des,cf}} = 479.4\text{kg/s}$	Design mass flow rate of cold fluid
	$\dot{m}_{\text{des,hf}} = 1{,}638\text{kg/s}$	Design mass flow rate of hot fluid
	$M_{\text{w}} = 4{,}838\text{kg}$	Wall mass
	$\varepsilon = 0.01\text{K}$	RLMTD threshold
Dwell heater molten salt pump	$\dot{m}_{\text{des}} = 1{,}638\text{kg/s}$	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{\text{in,des}} = 1\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 1.041\text{bar}$	Design outlet pressure
Feedwater pump	$\dot{m}_{\text{des}} = 479.4\text{kg/s}$	Design mass flow rate
	$\eta = 0.9$	Efficiency
	$p_{\text{in,des}} = 0.0508\text{bar}$	Design inlet pressure
	$p_{\text{out,des}} = 175.8\text{bar}$	Design outlet pressure

Table C.3: Parameters of individual component models in the dynamic process model of Case Study III in section 7.4. Condenser and evaporator parameters extrapolated linearly with area from [254, 82]. Other heat exchanger parameters extrapolated linearly with area from [108, 9]. Turbine and condenser pressure from [184]. Tank heat transfer coefficient from [197].

Label	Controlled variable	Manipulated variable	Set point	K_c	τ_I (s)	Minimum output	Maximum output
PI-a	Evaporator liquid fraction, $\frac{V_{L,EVP}}{V_{wf,EVP}}$	Mass flow rate of feedwater, $\dot{m}_{14}$	0.5	500kg/s	500	239.7kg/s	958.8kg/s
PI-b	Condenser pressure, p_{12}	Mass flow rate of cooling water, $\dot{m}_{32}$	0.0508bar	-50kg/(s Pa)	50	8,054kg/s	48,330kg/s
PI-c	Turbine power, $\dot{W}_{turbine}$	Partial arc admission valve opening, θ_{arc}	See section 7.5	$0.00001W^{-1}$	60	0	1
PI-d	Turbine inlet pressure, p_{11}	Mass flow rate of hot molten salt to superheater, $\dot{m}_5$	165.5bar	0.003kg/(s Pa)	100	36.27kg/s	7,255kg/s
PI-e	Preheater outlet water temperature, T_{21}	Preheater molten salt bypass valve opening, Ω_8	$T_{wf,EVP} - 10K$	$-0.1K^{-1}$	60	0	1
PI-f	Helium return temperature, T_2	Mass flow rate of cold molten salt, $\dot{m}_{10}$	450°C	-500kg/(s K)	60	44.7kg/s	6,706kg/s
PI-g	Dwell heater outlet water temperature, T_{16}	Mass flow rate of hot molten salt to dwell heater, $\dot{m}_4$	266.7°C	2kg/(s K)	2	1.638kg/s	2,458kg/s
PI-h	MTHX outlet water temperature, T_{29}	Split valve opening, Ω_{14}	266.7°C	$-0.005K^{-1}$	5	0.001	0.999
PI-i	Helium return temperature, T_{26}	MTHX bypass valve opening, Ω_{25}	225°C	$0.0001K^{-1}$	120	0	1
PI-j	Water return temperature, T_{30}	LTHX bypass valve opening, Ω_{29}	275°C	$0.75K^{-1}$	30	0	1

Table C.4: Outline of the process controllers for Case Study III in section 7.4. Please refer to figure 7.12 or section 7.4.1 for controller labels. All models are *LimPID*.

Appendix D

Colophon

This document was written using LaTeX and BibTeX in the Overleaf online editor, using the template from `https://github.com/UCL/ucl-latex-thesis-templates`. The LaTeX packages used were *algorithm*, *algpseudocode*, *amsmath*, *amssymb*, *array*, *biblatex*, *bm*, *empheq*, *emptypage*, *float*, *fontenc*, *gensymb*, *graphicx*, *hyperref*, *lscape*, *mathrsf*, *microtype*, *multirow*, *pdfpages*, *setspace*, *svg* and *textcomp*.

Bibliography

[1] Jack Acres, Ioannis Antoniou, Finlay Christie, Daniel Blackburn, and Samuel Knight. Staying positive: producing net power. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2280):20230404, 2024.

[2] Maria Claudia Aguitoni, Leandro Vitor Pavão, Paulo Henrique Siqueira, Laureano Jiménez, and Mauro Antonio da Silva Sá Ravagnani. Heat exchanger network synthesis using genetic algorithm and differential evolution. *Computers Chemical Engineering*, 117:82–96, 2018.

[3] Ryo Akasaka. A reliable and useful method to determine the saturation state from helmholtz energy equations of state. *Journal of Thermal Science and Technology*, 3(3):442–451, 2008.

[4] Wisam Abed Kattea Al-Maliki, Falah Alobaid, Vitali Kez, and Bernd Epple. Modelling and dynamic simulation of a parabolic trough power plant. *Journal of Process Control*, 39:123–138, 2016.

[5] Stéphane Alarie, Charles Audet, Aïmen E. Gheribi, Michael Kokkolaras, and Sébastien Le Digabel. Two decades of blackbox optimization applications. *EURO Journal on Computational Optimization*, 9:100011, 2021.

[6] Bhar Aliyu, Charles Osheku, Moshood Adetoro, and A. Funmilayo. Ordinary differential equations: Matlab/simulink solutions. *International Journal of Scientific and Engineering Research*, 3, 01 2012.

[7] Falah Alobaid. Start-up improvement of a supplementary-fired large combined-cycle power plant. *Journal of Process Control*, 64:71–88, 2018.

[8] Falah Alobaid, Nicolas Mertens, Ralf Starkloff, Thomas Lanz, Christian Heinze, and Bernd Epple. Progress in dynamic simulation of thermal power plants. *Progress in Energy and Combustion Science*, 59:79–162, 2017.

[9] Q. Altes-Buch, R. Dickes, A. Desideri, V. Lemort, and S. Quoilin. Dynamic modeling of thermal systems using a semi-empirical approach and the thermocycle Modelica library. In *ECOS 2015 - The 28th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems*, 6 2015.

[10] Oluwamayowa O. Amusat, Paul R. Shearing, and Eric S. Fraga. Optimal integrated energy systems design incorporating variable renewable energy sources. *Computers Chemical Engineering*, 95:21–37, 2016.

[11] Ehsan Askari Mahvelati, Mario Forcinito, Laurent Fitschy, and Arthur Maesen. Three-dimensional CFD model development and validation for once-through steam generator (OTSG): Coupling combustion, heat transfer and steam generation. *ChemEngineering*, 6(2), 2022.

[12] Mustafa Asker, Oguz Turgut, and Mustafa Coban. A review of non iterative friction factor correlations for the calculation of pressure drop in pipes. *Journal of Science and Technology*, 4:1–8, 06 2014.

[13] Modelica Association. Modelica libraries. `https://github.com/modelica/{Modelica}StandardLibrary`. Accessed: 11/02/2025.

[14] Mojtaba Babaelahi, Ehsan Mofidipour, and Ehsan Rafat. Combined energy-exergy-control (CEEC) analysis and multi-objective optimization of parabolic trough solar collector powered steam power plant. *Energy*, 201:117641, 2020.

[15] C. Bachmann, S. Ciattaglia, F. Cismondi, G. Federici, T. Franke, C. Gliss, T. Härtl, G. Keech, R. Kembleton, F. Maviglia, and M. Siccinio. Key design

integration issues addressed in the EU DEMO pre-concept design phase. *Fusion Engineering and Design*, 156:111595, 2020.

[16] Daniel Banuti, Muralikrishna Raju, Peter Ma, Matthias Ihme, and Jean-Pierre Hickey. Seven questions about supercritical fluids - towards a new fluid state diagram. In *55th AIAA Aerospace Sciences Meeting*, 01 2017.

[17] Junjiang Bao, Ruixiang Zhang, Yan Lin, Ning Zhang, Xiaopeng Zhang, and Gaohong He. Simultaneous optimization of system structure and working fluid for the three-stage condensation rankine cycle utilizing lng cold energy. *Applied Thermal Engineering*, 140:120–130, 2018.

[18] Matteo Barbarino. A brief history of nuclear fusion. *Nature Physics*, 16:890 – 893, 2020.

[19] L. Barucca, E. Bubelis, S. Ciattaglia, A. D'Alessandro, A. Del Nevo, F. Giannetti, W. Hering, P. Lorusso, E. Martelli, I. Moscato, A. Quartararo, A. Tarallo, and E. Vallone. Pre-conceptual design of EU DEMO balance of plant systems: Objectives and challenges. *Fusion Engineering and Design*, 169:112504, 2021.

[20] L. Barucca, A. D'Alessandro, Andrea Burlando, and A. Traverso. The new layout and regulation of the 'pulsating' power conversion system of EU DEMO WCLL for an effective and safe control of pulse-dwell transitions and increase in plant efficiency. *Nuclear Fusion*, 64, 07 2024.

[21] L. Barucca, W. Hering, S. Perez Martin, E. Bubelis, A. Del Nevo, M. Di Prinzio, M. Caramello, A. D'Alessandro, A Tarallo, E. Vallone, I. Moscato, A. Quartararo, S D'amico, F. Giannetti, P. Lorusso, V. Narcisi, C. Ciurluini, M.J. Montes Pita, C. Sánchez, A. Rovira, D. Santana, P. Gonzales, R. Barbero, M. Zaupa, M. Szogradi, S. Normann, M. Vaananen, J. Ylatalo, M. Lewandowska, L. Malinowski, E. Martelli, A. Froio, P. Arena, and A. Tincani. Maturation of critical technologies for the DEMO balance of plant systems. *Fusion Engineering and Design*, 179:113096, 2022.

[22] Ian Bell. Efficient and precise representation of pure fluid phase equilibria with chebyshev expansions. *International Journal of Thermophysics*, (42), 2021-03-25 04:03:00 2021.

[23] Ian H. Bell, Sylvain Quoilin, Emeline Georges, James E. Braun, Eckhard A. Groll, W. Travis Horton, and Vincent Lemort. A generalized moving-boundary algorithm to predict the heat transfer rate of counterflow heat exchangers for any phase configuration. *Applied Thermal Engineering*, 79:192–201, 2015.

[24] Ian H. Bell, Jorrit Wronski, Sylvain Quoilin, and Vincent Lemort. Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library coolprop. *Industrial & Engineering Chemistry Research*, 53(6):2498–2508, 2014.

[25] Alberto Benato, Anna Stoppato, and Stefano Bracco. Combined cycle power plants: A comparison between two different dynamic models to evaluate transient behaviour and residual life. *Energy Conversion and Management*, 87:1269–1280, 2014.

[26] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.

[27] Dheeraj Shankarrao Bhiogade. Ultra supercritical thermal power plant material advancements: A review. *Journal of Alloys and Metallurgical Systems*, 3:100024, 2023.

[28] Maria-Victoria Bologa, Evaldas Bubelis, and Wolfgang Hering. Parameter study and dynamic simulation of the DEMO intermediate heat transfer and storage system design using matlab®/simulink. *Fusion Engineering and Design*, 166:112291, 2021.

[29] J. Bonilla, L.J. Yebra, and S. Dormido. Mean densities in dynamic mathematical two-phase flow models. *CMES - Computer Modeling in Engineering and Sciences*, 67(1):13 – 37, 2010.

[30] J. Bonilla, L.J. Yebra, and S. Dormido. Chattering in dynamic mathematical two-phase flow models. *Applied Mathematical Modelling*, 36(5):2067–2081, 2012.

[31] Javier Bonilla, Alberto de la Calle, Margarita M. Rodríguez-García, Lidia Roca, and Loreto Valenzuela. Study on shell-and-tube heat exchanger models with different degree of complexity for process simulation and control design. *Applied Thermal Engineering*, 124:1425–1440, 2017.

[32] Javier Bonilla, Sebastián Dormido, and François E. Cellier. Switching moving boundary models for two-phase flow evaporators and condensers. *Communications in Nonlinear Science and Numerical Simulation*, 20(3):743–768, 2015.

[33] Javier Bonilla, Margarita M. Rodríguez-García, Lidia Roca, Alberto de la Calle, and Loreto Valenzuela. Design and experimental validation of a computational effective dynamic thermal energy storage tank model. *Energy*, 152:840–857, 2018.

[34] Claus Borgnakke and Richard Edwin. Sonntag. *Fundamentals of thermodynamics*. Wiley, Hoboken, N.J, 8th edition, 2014.

[35] Omid Bozorg-Haddad, Mohammad Solgi, and Hugo A. Loáiciga. *Meta-Heuristic and Evolutionary Algorithms for Engineering Optimization*. Wiley, 2017.

[36] Nielis Brouwer, Danny Dijkzeul, Levi Koppenhol, Iris Pijning, and Daan van den Berg. Survivor selection in a crossoverless evolutionary algorithm. In *the Genetic and Evolutionary Computation Conference Companion*, GECCO '22, page 1631–1639, New York, NY, USA, 2022. Association for Computing Machinery.

[37] Robert Brunet, Gonzalo Guillén-Gosálbez, and Laureano Jiménez. Combined simulation–optimization methodology to reduce the environmental impact

of pharmaceutical processes: application to the production of penicillin v. *Journal of Cleaner Production*, 76:55–63, 2014.

[38] Evaldas Bubelis, Wolfgang Hering, and Sara Perez-Martin. Industry supported improved design of DEMO BoP for HCPB BB concept with energy storage system. *Fusion Engineering and Design*, 146:2334–2337, 2019.

[39] H. Cabal, Y. Lechón, C. Bustreo, F. Gracceva, M. Biberacher, D. Ward, D. Dongiovanni, and P.E. Grohnheit. Fusion power in a future low carbon global electricity system. *Energy Strategy Reviews*, 15:1–8, 2017.

[40] Gianluca Carraro, Piero Danieli, Andrea Lazzaretto, and Tazio Boatto. A common thread in the evolution of the configurations of supercritical CO_2 power systems for waste heat recovery. *Energy Conversion and Management*, 237:114031, 2021.

[41] Francesco Casella and Alberto Leva. Modelling of thermo-hydraulic power generation processes using Modelica. *Mathematical and Computer Modelling of Dynamical Systems*, 12(1):19–33, 2006.

[42] Francesco Casella, Francesco Pretolani, et al. Fast start-up of a combined-cycle power plant: a simulation study with Modelica. In *Modelica Conference*, volume 4, pages 3–10. The Modelica Association Linköping, Sweeden, 2006.

[43] Cesar Celis, Gustavo R.S. Pinto, Tairo Teixeira, and Érica Xavier. A steam turbine dynamic model for full scope power plant simulators. *Applied Thermal Engineering*, 120:593–602, 2017.

[44] Hubert Chanson. 14 - physical modelling of hydraulics. In Hubert Chanson, editor, *Hydraulics of Open Channel Flow (Second Edition)*, pages 253–274. Butterworth-Heinemann, Oxford, second edition edition, 2004.

[45] I. T. Chapman, S. C. Cowley, and H. R. Wilson. The spherical tokamak for energy production: theme issue introduction. *Philosophical Transactions*

of the Royal Society A: Mathematical, Physical and Engineering Sciences, 382(2280):20230416, 2024.

[46] J. J. Chapman. Advanced fusion reactors for space propulsion and power systems. In *2011 Abstracts IEEE International Conference on Plasma Science*, pages 1–1, June 2011.

[47] Vassilis M. Charitopoulos and Vivek Dua. A unified framework for model-based multi-objective linear process and energy optimisation under uncertainty. *Applied Energy*, 186:539–548, 2017.

[48] Tawfiq Chekifi and Moustafa Boukraa. CFD applications for sensible heat storage: A comprehensive review of numerical studies. *Journal of Energy Storage*, 68:107893, 2023.

[49] Chen Chen and George M. Bollas. Dynamic optimization of a subcritical steam power plant under time-varying power load. *Processes*, 6(8), 2018.

[50] Chen Chen, Zhiquan Zhou, and George M. Bollas. Dynamic modeling, simulation and optimization of a subcritical steam power plant. part i: Plant model and regulatory control. *Energy Conversion and Management*, 145:324–334, 2017.

[51] Jiancong Chen, Si Tan, Jiaqiang E, Gaoliang Liao, and Feng Zhang. Dynamic performance and control analysis of supercritical CO_2 brayton cycle for solid oxide fuel cell waste heat recovery. *Applied Thermal Engineering*, 250:123556, 2024.

[52] Mengchao Chen, Ruikai Zhao, Li Zhao, Dongpeng Zhao, Shuai Deng, and Wei Wang. Supercritical CO_2 brayton cycle: Intelligent construction method and case study. *Energy Conversion and Management*, 246:114662, 2021.

[53] Xiaoting Chen, Xiaoya Li, Mingzhang Pan, and Zongrun Wang. Superstructure-free synthesis and multi-objective optimization of supercritical CO_2 cycles. *Energy Conversion and Management*, 284:116966, 2023.

[54] Yun-Guang Chen and Qing-Yuan Yang. A new moving-boundary algorithm to predict heat transfer rate of counter-flow water-cooled transcritical CO_2 gas cooler. *International Journal of Refrigeration*, 146:357–365, 2023.

[55] Yang Cheng, Qiyuan Ma, Huiyu Zhou, Lin Jiang, Zhao Ma, Qiuwang Wang, and Ting Ma. A mathematical method for thermal design of printed circuit heat exchanger with actual flow distributors. *Applied Thermal Engineering*, 263:125364, 2025.

[56] Zafer Civelek, Goksu Gorel, Murat Luy, Necaattin Barısci, and Ertugrul Cam. Effects on load-frequency control of a solar power system with a two-area interconnected thermal power plant and its control with a new BFA algorithm. *Elektronika ir Elektrotechnika*, 24(6):3–10, Dec. 2018.

[57] Nathan Colgan, Gregory Nellis, and Mark Anderson. Modelling and optimization of a forced convection heat exchanger for mars waste heat rejection for power cycle and cryocooling applications. *Applied Thermal Engineering*, 242:122463, 2024.

[58] Francesco Colliva, Cristiano Ciurluini, Andrea Iaboni, Giulia Valeria Centomani, Antonio Trotta, and Fabio Giannetti. Analysis of power conversion system options for ARC-like tokamak fusion reactor balance of plant. *Sustainability*, 16(17), 2024.

[59] Andrew R. Conn, Katya Scheinberg, and Luís Nunes Vicente. Introduction to derivative-free optimization. In *MPS-SIAM series on optimization*, 2009.

[60] Open Source Modelica Consortium. OMJulia. `https://github.com/{OpenModelica}/OMJulia.jl`. Accessed: 4/03/2025.

[61] Open Source Modelica Consortium. OpenModelica. `https://openmodelica.org/`. Accessed: 21/02/2025.

[62] D. H. Cooke. On Prediction of Off-Design Multistage Turbine Pressures by Stodola's Ellipse. *Journal of Engineering for Gas Turbines and Power*, 107(3):596–606, 07 1985.

[63] A E Costley and S A M McNamara. Fusion performance of spherical and conventional tokamaks: implications for compact pilot plants and reactors. *Plasma Physics and Controlled Fusion*, 63(3):035005, jan 2021.

[64] Erik Dahlquist, Elena Tomas Aparicio, Hailong Li, Peter A. Fritzson, and Per Östlund. Comparison open Modelica - dymola for power plant simulation. *Simul. Notes Eur.*, 23:139–146, 2013.

[65] Enrico Dal Cin, Andrea Lazzaretto, and Andrea Toffolo. A novel extension of the SYNTHSEP methodology for the optimal synthesis and design of supercritical CO_2 cycles in waste heat recovery applications. *Energy Conversion and Management*, 276:116535, 2023.

[66] Mark L. Darby, Michael Nikolaou, James Jones, and Doug Nicholson. RTO: An overview and assessment of current practice. *Journal of Process Control*, 21(6):874–884, 2011.

[67] S De and S.K Biswal. Performance improvement of a coal gasification and combined cogeneration plant by multi-pressure steam generation. *Applied Thermal Engineering*, 24(2):449–456, 2004.

[68] Marleen de Jonge and Daan van den Berg. Parameter sensitivity patterns in the plant propagation algorithm. In *In EvoSTAR2020*, pages 92–99, 01 2020.

[69] Marleen de Jonge and Daan van den Berg. Plant propagation parameterization: Offspring population size. In *In EvoSTAR2020*, 04 2020.

[70] Kalyanmoy Deb. An efficient constraint handling method for genetic algorithms. *Computer Methods in Applied Mechanics and Engineering*, 186(2):311–338, 2000.

[71] Q.H. Deng, D. Wang, H. Zhao, W.T. Huang, S. Shao, and Z.P. Feng. Study on performances of supercritical CO_2 recompression brayton cycles with multi-objective optimization. *Applied Thermal Engineering*, 114:1335–1342, 2017.

[72] Adriano Desideri, Bertrand Dechesne, Jorrit Wronski, Martijn Van den Broek, Sergei Gusev, Vincent Lemort, and Sylvain Quoilin. Comparison of moving boundary and finite-volume heat exchanger models in the Modelica language. *Energies*, 9(5), 2016.

[73] Michel O. Deville. *Incompressible Newtonian Fluid Mechanics*, pages 1–32. Springer International Publishing, Cham, 2022.

[74] Arumugam Dhanasekaran and Sivasailam Kumaraswamy. Experimental evaluation of affinity law of pumps by using multistage electric submersible pump at various speeds of operation. *AIP Conference Proceedings*, 2283(1):020106, 10 2020.

[75] Boopathi Dhanasekaran, Saravanan Siddhan, and Jagatheesan Kaliannan. Ant colony optimization technique tuned controller for frequency regulation of single area nuclear power generating system. *Microprocessors and Microsystems*, 73:102953, 2020.

[76] Mike Diessner, Joseph O'Connor, Andrew Wynn, Sylvain Laizet, Yu Guan, Kevin Wilson, and Richard D. Whalley. Investigating bayesian optimization for expensive-to-evaluate black box functions: Application in fluid dynamics. *Frontiers in Applied Mathematics and Statistics*, 8, 2022.

[77] Yadong Du, Ce Yang, Ben Zhao, Chenxing Hu, Hanzhi Zhang, Zhiyi Yu, Jianbing Gao, Wei Zhao, and Haimei Wang. Optimal design of a supercritical carbon dioxide recompression cycle using deep neural network and data mining techniques. *Energy*, 271:127038, 2023.

[78] Yang Du, Zhenghao Yang, Zeqi Zhang, Zhenbiao Wang, Guangyu He, Jiangfeng Wang, and Pan Zhao. Control strategy optimization exploration of a

novel hydrogen-fed high-efficiency x-type rotary engine hybrid power system by coupling with recuperative organic rankine cycle. *Energy*, 293:130677, 2024.

[79] Rebecca I. Dunn, Patrick J. Hearps, and Matthew N. Wright. Molten-salt power towers: Newly commercial concentrating solar storage. *IEEE*, 100(2):504–515, Feb 2012.

[80] Giampaolo D'Alessandro, Michele Potenza, Sandra Corasaniti, Stefano Sfarra, Paolo Coppa, Gianluigi Bovesecchi, and Filippo de Monte. Modeling and measuring thermodynamic and transport thermophysical properties: A review. *Energies*, 15(23), 2022.

[81] Mohamed Ebrahim, Jagatheesan Kaliannan, and Anand Baskaran. Stochastic particle swarm optimizaion for tuning of PID controller in load frequency control of single area reheat thermal power system. *International Journal of Electrical and Power Engineering*, 8:33–40, 10 2014.

[82] Baligh El Hefni and Daniel Bouskela. *Modeling and Simulation of Thermal Power Plants with ThermoSysPro: A Theoretical Introduction and a Practical Guide*. Springer Cham, 1 edition, 01 2019.

[83] Cristina Elsido, Emanuele Martelli, and Ignacio E. Grossmann. A bilevel decomposition method for the simultaneous heat integration and synthesis of steam/organic rankine cycles. *Computers Chemical Engineering*, 128:228–245, 2019.

[84] Cristina Elsido, Emanuele Martelli, and Ignacio E. Grossmann. Simultaneous multiperiod optimization of rankine cycles and heat exchanger networks. In Sauro Pierucci, Flavio Manenti, Giulia Luisa Bozzano, and Davide Manca, editors, *30th European Symposium on Computer Aided Process Engineering*, volume 48 of *Computer Aided Chemical Engineering*, pages 1495–1500. Elsevier, 2020.

[85] Marco Crialesi Esposito, Nicola Pompini, Agostino Gambarotta, Vetrivel Chandrasekaran, Junqiang Zhou, and Marcello Canova. Nonlinear model predictive control of an organic rankine cycle for exhaust waste heat recovery in automotive engines. *IFAC-PapersOnLine*, 48(15):411–418, 2015.

[86] Kunle Fan, Cheng Yang, Zhuli Xie, and Xiaoqian Ma. Load-regulation characteristics of gas turbine combined cycle power system controlled with compressor inlet air heating. *Applied Thermal Engineering*, 196:117285, 2021.

[87] Han Fang, Bill Wong, and Yu Bai. Use of kinetic model for thermal properties of steel at high temperatures. *Australian Journal of Structural Engineering*, 13, 01 2014.

[88] Xiande Fang, Lu Xu, Yuanyuan Chen, and Weiwei Chen. Correlations for friction factor of turbulent pipe flow under supercritical pressure: Review and a new correlation. *Progress in Nuclear Energy*, 118:103085, 2020.

[89] Kazem Farhadi, Anis Bousbia-salah, and Franscesco D'Auria. A model for the analysis of pump start-up transients in tehran research reactor. *Progress in Nuclear Energy*, 49(7):499–510, 2007.

[90] Simone Ferrero, Lluís Batet, José Ignacio Linares, Eva Arenas, Alexis Cantizano, and Laura Savoldi. A Modelica dynamic model of a supercritical CO_2 energy conversion system for EU-DEMO. *Fusion Engineering and Design*, 173:112826, 2021.

[91] Culham Centre for Fusion Energy. Fusion energy: Frequently asked questions. `https://ccfe.ukaea.uk/fusion-energy/faq/`. Accessed: 26/02/2025.

[92] Eric S. Fraga. Fresa: A plant propagation algorithm for black-box single and multiple objective optimization. *International Journal on Engineering Technologies and Informatics*, 2021.

[93] Eric S. Fraga. Multiple simultaneous solution representations in a population based evolutionary algorithm, 2021.

[94] Eric S. Fraga. *Nature inspired methods for optimization: A Julia primer for process engineering*. Leanpub, 2023.

[95] Eric S. Fraga and Oluwamayowa Amusat. *Understanding the Impact of Constraints: A Rank Based Fitness Function for Evolutionary Methods*, pages 243–254. Springer International Publishing, Cham, 2016.

[96] Eric S. Fraga, Veerawat Udomvorakulchai, and Lazaros G. Papageorgiou. A multi-agent system for hybrid optimization. In Flavio Manenti and Gintaras V. Reklaitis, editors, *34th European Symposium on Computer Aided Process Engineering / 15th International Symposium on Process Systems Engineering*, volume 53 of *Computer Aided Chemical Engineering*, pages 3331–3336. Elsevier, 2024.

[97] Rüdiger Franke, Francesco Casella, Martin Otter, Michael Sielemann, Hilding Elmqvist, Sven Mattson, and Hans Olsson. Stream connectors - an extension of Modelica for device-oriented modeling of convective transport phenomena. In *7th Modelica Conference*, pages 108–121, 10 2009.

[98] Peter Fritzson. *Principles of Object Oriented Modeling and Simulation with Modelica 3.3*. John Wiley Sons, Ltd, 2014.

[99] Fraunhofer-Institut für Solare Energiesysteme. Levelized cost of electricity renewable energy technologies. `https://www.ise.fraunhofer.de/content/dam/ise/en/documents/publications/studies/EN2024_ISE_Study_Levelized_Cost_of_Electricity_Renewable_Energy_Technologies.pdf`, 2024. Accessed: 23/02/2025.

[100] Lei Gao and Yunho Hwang. A review of configuration optimization for thermal energy systems. In *ASHRAE Winter Conference*, ASHRAE Transactions, pages 534–543. American Society of Heating Refrigerating and Air-Conditioning Engineers, 2024.

[101] Xianhua Gao, Shangshang Wei, Chunlin Xia, and Yiguo Li. Flexible operation of concentrating solar power plant with thermal energy storage based on a coordinated control strategy. *Energies*, 15(13), 2022.

[102] Mikhail Gryaznevich, Valery A. Chuyanov, and Yuichi Takase. Pulsed spherical tokamak—a new approach to fusion reactors. *Plasma*, 5(2):247–257, 2022.

[103] Seyfettin Can Gülen. Steam turbine—quo vadis? *Frontiers in Energy Research*, 8, 2021.

[104] Farzaneh Hajabdollahi, Zahra Hajabdollahi, and Hassan Hajabdollahi. Soft computing based multi-objective optimization of steam cycle power plant using nsga-ii and ann. *Applied Soft Computing*, 12(11):3648–3655, 2012.

[105] Lu Han, Chen Chen, and George M. Bollas. Dynamic optimization and control of chemical looping combustion combined cycle power plants. *IFAC-PapersOnLine*, 51(18):845–850, 2018.

[106] Tianding Han, Qifei Li, Lin Shang, Xiangyu Chen, Feng Zhou, and Wangxu Li. Study on the influence of reynolds number on heat exchange performance and nusselt number of spray coil heat exchanger. *Processes*, 13(2), 2025.

[107] Allan H. Harvey and James C. Bellows. A century of asme steam tables. `https://www.asme.org/topics-resources/content/a-century-of-asme-steam-tables`. Accessed: 19/12/2024.

[108] Heatric. Printed circuit heat exchangers. `www.heatric/heat-exchangers/features/characteristics/`. Accessed: 29/5/2024.

[109] Wolfgang Hering, Evaldas Bubelis, Sara Perez-Martin, and Maria-Victoria Bologa. Overview of thermal hydraulic optimization and verification for the EU-DEMO HCPB BOP ICD variant. *Energies*, 14(23), 2021.

[110] Francisco A. Hernández, Pavel Pereslavtsev, Guangming Zhou, Qinlan Kang, Salvatore D'Amico, Heiko Neuberger, Lorenzo V. Boccaccini, Béla Kiss,

Gábor Nádasi, Luis Maqueda, Ion Cristescu, Ivo Moscato, Italo Ricapito, and Fabio Cismondi. Consolidated design of the HCPB breeding blanket for the pre-conceptual design phase of the EU DEMO and harmonization with the ITER HCPB TBM program. *Fusion Engineering and Design*, 157:111614, 2020.

[111] Brígido J. Hipólito-Valencia, Eusiel Rubio-Castro, José M. Ponce-Ortega, Medardo Serna-González, Fabricio Nápoles-Rivera, and Mahmoud M. El-Halwagi. Optimal integration of organic rankine cycles with industrial processes. *Energy Conversion and Management*, 73:285–302, 2013.

[112] Shengya Hou, Yaodong Zhou, Lijun Yu, Fengyuan Zhang, and Sheng Cao. Optimization of the combined supercritical CO_2 cycle and organic rankine cycle using zeotropic mixtures for gas turbine waste heat recovery. *Energy Conversion and Management*, 160:313–325, 2018.

[113] Ke Feng Huang, Eid M. Al-mutairi, and I.A. Karimi. Heat exchanger network synthesis using a stagewise superstructure with non-isothermal mixing. *Chemical Engineering Science*, 73:30–43, 2012.

[114] Zihao Huang, Huailiang You, Jitian Han, Guoxiang Li, Yan Xiao, Bin Hu, Ze-Hang Chen, and Daifen Chen. Techno-economic evaluation and multi-objective optimization of a cogeneration system integrating solid oxide fuel cell with steam rankine and supercritical carbon dioxide brayton cycles. *Fuel*, 382:133675, 2025.

[115] Marcia L Huber and James F Ely. A predictive extended corresponding states model for pure and mixed refrigerants including an equation of state for R134a. *International Journal of Refrigeration*, 17(1):18–31, 1994.

[116] Marcia L. Huber, Eric W. Lemmon, Ian H. Bell, and Mark O. McLinden. The nist refprop database for highly accurate properties of industrially important fluids. *Industrial & Engineering Chemistry Research*, 61(42):15449–15472, 2022.

[117] Marcia L. Huber and Mark O. McLinden. Thermodynamic properties of R134a (1,1,1,2-tetrafluoroethane). In *International Refrigeration and Air Conditioning Conference*, 1992.

[118] Wolfgang R. Huster, Artur M. Schweidtmann, Jannik T. Lüthje, and Alexander Mitsos. Deterministic global superstructure-based optimization of an organic rankine cycle. *Computers Chemical Engineering*, 141:106996, 2020.

[119] Moritz Hübel, Sebastian Meinke, Marcus T. Andrén, Christoffer Wedding, Jürgen Nocke, Conrad Gierow, Egon Hassel, and Jonas Funkquist. Modelling and simulation of a coal-fired power plant for start-up optimisation. *Applied Energy*, 208:319–331, 2017.

[120] Cyrille Imbert. *Computer Simulations and Computational Models in Science*, pages 735–781. Springer International Publishing, Cham, 2017.

[121] Muhammad Imran, Roberto Pili, Muhammad Usman, and Fredrik Haglind. Dynamic modeling and control strategies of organic rankine cycle systems: Methods and challenges. *Applied Energy*, 276:115537, 2020.

[122] Energy Institute. Statistical review of world energy 2024. `https://www.energyinst.org/statistical-review`, 2024. Accessed: 21/02/2025.

[123] Robin C. Purshouse Ioannis Giagkiozis and Peter J. Fleming. An overview of population-based algorithms for multi-objective optimisation. *International Journal of Systems Science*, 46(9):1572–1599, 2013.

[124] Alfredo Iranzo, Christian Suárez, and José Guerra. Mixing enhancement in thermal energy storage molten salt tanks. *Energy Conversion and Management*, 168:320–328, 2018.

[125] Yuan Jiang, Eric Liese, Stephen E. Zitney, and Debangsu Bhattacharyya. Design and dynamic modeling of printed circuit heat exchangers for supercritical carbon dioxide brayton power cycles. *Applied Energy*, 231:1019–1032, 2018.

[126] K. Jonshagen and M. Genrup. Improved load control for a steam cycle combined heat and power plant. *Energy*, 35(4):1694–1700, 2010.

[127] Surender Kannaiyan, Sharad Bhartiya, and Mani Bhushan. Dynamic modeling and simulation of a hybrid solar thermal power plant. *Industrial & Engineering Chemistry Research*, 58(18):7531–7550, 2019.

[128] Saif R. Kazi, Ishanki A. De Mel, and Michael Short. A new trust-region approach for optimization of multi-period heat exchanger networks with detailed shell-and-tube heat exchanger designs. In Yoshiyuki Yamashita and Manabu Kano, editors, *14th International Symposium on Process Systems Engineering*, volume 49 of *Computer Aided Chemical Engineering*, pages 241–246. Elsevier, 2022.

[129] Hanbing Ke, Qi Xiao, Yiping Cao, Ting Ma, Yuansheng Lin, Min Zeng, and Qiuwang Wang. Simulation of the printed circuit heat exchanger for s-CO_2 by segmented methods. *Energy Procedia*, 142:4098–4103, 2017.

[130] Maziar Kermani, Anna S. Wallerand, Ivan D. Kantor, and François Maréchal. Generic superstructure synthesis of organic rankine cycles for waste heat recovery in industrial processes. *Applied Energy*, 212:1203–1225, 2018.

[131] Donghun Kim, Davide Ziviani, James E. Braun, and Eckhard A. Groll. A moving boundary modeling approach for heat exchangers with binary mixtures. *Energy Procedia*, 129:466–473, 2017.

[132] Sung Young Kim and Miguel Bagajewicz. Global optimization of heat exchanger networks using a new generalized superstructure. *Chemical Engineering Science*, 147:30–46, 2016.

[133] T.S Kim, H.J Park, and S.T Ro. Characteristics of transient operation of a dual-pressure bottoming system for the combined cycle power plant. *Energy*, 26(10):905–918, 2001.

[134] Lukas Kistner, Astrid Bensmann, and Richard Hanke-Rauschenbach. Optimal design of a distributed ship power system with solid oxide fuel cells under the consideration of component malfunctions. *Applied Energy*, 316:119052, 2022.

[135] Freerk Klasing, Marco Prenzel, and Thomas Bauer. Repurposing of supercritical coal plants into highly flexible grid storage with adapted 620 °C nitrate salt technology. *Applied Energy*, 377:124524, 2025.

[136] Christoph Koch, Frank Cziesla, and George Tsatsaronis. Optimization of combined cycle power plants using evolutionary algorithms. *Chemical Engineering and Processing: Process Intensification*, 46(11):1151–1159, 2007.

[137] Levi Koppenhol, Nielis Brouwer, Danny Dijkzeul, Iris Pijning, Joeri Sleegers, and Daan van den Berg. Exactly characterizable parameter settings in a crossoverless evolutionary algorithm. In *the Genetic and Evolutionary Computation Conference Companion*, GECCO '22, page 1640–1649, New York, NY, USA, 2022. Association for Computing Machinery.

[138] Birsen Irem Kuvvetli and Abdel Salhi. *The Plant Propagation Algorithm for Discrete Optimisation: The Case of the Travelling Salesman Problem*, pages 43–61. Springer International Publishing, 03 2016.

[139] Jin Su Kwon, Seong Jun Bae, Jin Young Heo, and Jeong Ik Lee. Development of accelerated PCHE off-design performance model for optimizing power system operation strategies in s-CO_2 brayton cycle. *Applied Thermal Engineering*, 159:113845, 2019.

[140] Michael F. L'Annunziata. *Radioactivity*. Elsevier, Boston, second edition edition, 2016.

[141] Andrea Lazzaretto, Giovanni Manente, and Andrea Toffolo. SYNTHSEP: A general methodology for the synthesis of energy system configurations beyond superstructures. *Energy*, 147:924–949, 2018.

[142] Eric W. Lemmon and Richard T Jacobsen. A new functional form and new fitting techniques for equations of state with application to pentafluoroethane (HFC-125). *Journal of Physical and Chemical Reference Data*, 34(1):69–108, 03 2005.

[143] Bo Li, Yue Cao, Tianyu He, and Fengqi Si. Thermodynamic analysis and operation strategy optimization of coupled molten salt energy storage system for coal-fired power plant. *Applied Thermal Engineering*, 236:121702, 2024.

[144] Hao Li, Ruhai Zhang, Zhen Li, Sangkyoung Lee, Yaping Ju, Chuhua Zhang, and Yuanshen Lu. Design optimization and off-design performance analysis of one-dimensional supercritical CO_2 brayton cycle. *Applied Thermal Engineering*, 258:124547, 2025.

[145] Sheng Li, Jinshuang Gao, Lizhe Zhang, Yazhou Zhao, and Xuejun Zhang. Exploration of dual-phase change coupled heat transfer in solar regenerative evaporator. *Energy*, 293:130560, 2024.

[146] Xiaolei Li, Zhifeng Wang, Ershu Xu, Linrui Ma, Li Xu, and Dongming Zhao. Dynamically coupled operation of two-tank indirect TES and steam generation system. *Energies*, 12(9), 2019.

[147] Xiaoya Li, Bin Xu, Hua Tian, and Gequn Shu. Towards a novel holistic design of organic rankine cycle (ORC) systems operating under heat source fluctuations and intermittency. *Renewable and Sustainable Energy Reviews*, 147:111207, 2021.

[148] Yanfeng Li, Jingru Liu, and Guohe Huang. Pressure drop optimization of the main steam and reheat steam system of a 1000 MW secondary reheat unit. *Energies*, 15(9), 2022.

[149] Zheng Liang, Yingzong Liang, Xianglong Luo, Jianyong Chen, Zhi Yang, Chao Wang, and Ying Chen. Superstructure-based mixed-integer nonlinear programming framework for hybrid heat sources driven organic rankine cycle optimization. *Applied Energy*, 307:118277, 2022.

[150] Eric Liese. Modeling of a steam turbine including partial arc admission for use in a process simulation software environment. *Journal of Engineering for Gas Turbines and Power*, 136(11):112605, 06 2014.

[151] Shan Lin, Li Zhao, Shuai Deng, Dongpeng Zhao, Wei Wang, and Mengchao Chen. Intelligent collaborative attainment of structure configuration and fluid selection for the organic rankine cycle. *Applied Energy*, 264:114743, 2020.

[152] Haiqing Liu, Zhongran Chi, and Shusheng Zang. Optimization of a closed brayton cycle for space power systems. *Applied Thermal Engineering*, 179:115611, 2020.

[153] Kairui Liu, Yalong Guo, Limin Wang, Gaofeng Fan, and Defu Che. Control of supercritical CO_2 brayton cycle for fast and efficient load variation processes. *Applied Thermal Engineering*, 244:122797, 2024.

[154] Yaping Liu, Ying Wang, and Diangui Huang. Supercritical CO_2 brayton cycle: A state-of-the-art review. *Energy*, 189:115900, 2019.

[155] Zefeng Liu, Chaoyang Wang, Jianlin Fan, Ming Liu, Yong Xing, and Junjie Yan. Enhancing the flexibility and stability of coal-fired power plants by optimizing control schemes of throttling high-pressure extraction steam. *Energy*, 288:129756, 2024.

[156] Michael Lord, Iryna Bennett, Chris Harrington, Adam Cooper, Dan Lee-Lane, Adam Cureton, Cameron Olde, Megan Thompson, Dinusha Jayasundara, and Toby Meatyard. Fusing together an outline design for sustained fuelling and tritium self-sufficiency. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2280):20230410, 2024.

[157] Yan Luo, Zhiyuan Wang, Jiamin Zhu, Tao Lu, Gang Xiao, Fengming Chu, and Ruixing Wang. Multi-objective robust optimization of a solar power tower plant under uncertainty. *Energy*, 238:121716, 2022.

[158] Panagiotis Lykas, Evangelos Bellos, Angeliki Kitsopoulou, and Christos Tzivanidis. Dynamic analysis of a solar-biomass-driven multigeneration system based on s-CO_2 brayton cycle. *International Journal of Hydrogen Energy*, 59:1268–1286, 2024.

[159] Jiaxin Ma, Bingtao Zhao, and Yaxin Su. Development of multilevel cascade layouts to improve performance of sCO_2 brayton power cycles: Design, simulation, and optimization. *Energy*, 308:132952, 2024.

[160] Jinliang Ma, Miguel A. Zamarripa, John C. Eslick, Quang M. Le, Debangsu Bhattacharyya, Lorenz T. Biegler, Stephen E. Zitney, Anthony P. Burgard, and David C. Miller. Dynamic simulation and optimization of a subcritical coal-fired power plant during load- ramping operations. In Yoshiyuki Yamashita and Manabu Kano, editors, *14th International Symposium on Process Systems Engineering*, volume 49 of *Computer Aided Chemical Engineering*, pages 1033–1038. Elsevier, 2022.

[161] Ning Ma, Zhengkun Bu, Yanan Fu, Wenpeng Hong, Haoran Li, and Xiaojuan Niu. An operation strategy and off-design performance for supercritical brayton cycle using CO_2-propane mixture in a direct-heated solar power tower plant. *Energy*, 278:127882, 2023.

[162] Teng Ma, Ming-Jia Li, Jin-Liang Xu, Jing-Wei Ni, Wen-Quan Tao, and Li Wang. Study of dynamic response characteristics of s-CO_2 cycle in coal-fired power plants based on real-time micro-grid load and a novel synergistic control method with variable working conditions. *Energy Conversion and Management*, 254:115264, 2022.

[163] Ya-Nan Ma, Peng Hu, Cheng-Qi Jia, Ze-Rui Wu, and Qi Chen. Thermo-economic analysis and multi-objective optimization of supercritical brayton cycles with CO_2-based mixtures. *Applied Thermal Engineering*, 219:119492, 2023.

[164] Yuegeng Ma, Tatiana Morosuk, Jing Luo, Ming Liu, and Jiping Liu. Superstructure design and optimization on supercritical carbon dioxide cycle for application in concentrated solar power plant. *Energy Conversion and Management*, 206:112290, 2020.

[165] Simone Maccarini, Swatara Tucker, Luca Mantelli, Stefano Barberis, and Alberto Traverso. Dynamics and control implementation of a supercritical CO_2 recuperated cycle. *Applied Thermal Engineering*, 263:125237, 2025.

[166] Shide Mao, Zhenhao Duan, Jiawen Hu, Zhigang Zhang, and Lanlan Shi. Extension of the IAPWS-95 formulation and an improved calculation approach for saturated properties. *Physics of the Earth and Planetary Interiors*, 185(1):53–60, 2011.

[167] Gaël Maranzana, Isabelle Perry, and Denis Maillet. Mini- and micro-channels: influence of axial conduction in the walls. *International Journal of Heat and Mass Transfer*, 47(17):3993–4004, 2004.

[168] Emanuele Martelli, Falah Alobaid, and Cristina Elsido. Design optimization and dynamic simulation of steam cycle power plants: A review. *Frontiers in Energy Research*, 9, 2021.

[169] Emanuele Martelli, Cristina Elsido, Alberto Mian, and Francois Marechal. Minlp model and two-stage algorithm for the simultaneous synthesis of heat exchanger networks, utility systems and heat recovery cycles. *Computers Chemical Engineering*, 106:663–689, 2017.

[170] Sammyak Mate, Pawankumar Pal, Anshumali Jaiswal, and Sharad Bhartiya. Simultaneous tuning of multiple PID controllers for multivariable systems using deep reinforcement learning. *Digital Chemical Engineering*, 9:100131, 2023.

[171] Sven Erik Mattsson, Hilding Elmqvist, and Martin Otter. Physical system modeling with Modelica. *Control Engineering Practice*, 6(4):501–510, 1998.

[172] Samuele Meschini, Francesco Laviano, Federico Ledda, Davide Pettinari, Raffella Testoni, Daniele Torsello, and Bruno Panella. Review of commercial nuclear fusion projects. *Frontiers in Energy Research*, 11, 2023.

[173] Xinyu Miao, Haochun Zhang, Fangwei Ma, MingHao Deng, and Ersheng You. Thermodynamic, exergoeconomic evaluation and optimization of s–N2O/t-N2O nuclear power cycle for the construction of the lunar base. *Energy*, 302:131828, 2024.

[174] Indrasis Mitra and Indranil Ghosh. Axial conduction in cross-flow heat exchangers: An analytical approach to the coupled heat transfer problem. *International Journal of Heat and Mass Transfer*, 200:123502, 2023.

[175] Ehsan Mofidipour and Mojtaba Babaelahi. New procedure in solar system dynamic simulation, thermodynamic analysis, and multi-objective optimization of a post-combustion carbon dioxide capture coal-fired power plant. *Energy Conversion and Management*, 224:113321, 2020.

[176] Ettore Morosini, Dario Alfani, Salma I. Salah, Abdelrahman Abdeldayem, Francesco Crespi, Gioele Di Marcoberardino, and Giampaolo Manzolini. Off-design of a CO_2-based mixture transcritical cycle for CSP applications: Analysis at part load and variable ambient temperature. *Applied Thermal Engineering*, 236:121735, 2024.

[177] G. Morrison and M. O. McLinden. Application of a hard sphere equation of state to refrigerants and refrigerant mixtures, 1987. Accessed: 19/12/2024.

[178] Globaltherm Omnistore MS-600. `https://globalhtf.com/temperatures/globaltherm-omnistore/`. Accessed: 08-02-2025.

[179] Marta Muñoz, Antonio Rovira, and María José Montes. Thermodynamic cycles for solar thermal power plants: A review. *WIREs Energy and Environment*, 11(2):e420, 2022.

[180] Jobrun Nandong. Heuristic-based multi-scale control procedure of simultaneous multi-loop PID tuning for multivariable processes. *Journal of Process Control*, 35:101–112, 2015.

[181] A. Naserbegi, M. Aghaie, and Z. Nourani. Optimal design of a concentrated solar power plant with a thermal energy storage system using the downhill simplex method. *Applied Thermal Engineering*, 240:122264, 2024.

[182] S. M. A. Noori Rahim Abadi, A. Ahmadpour, S. M. N. R. Abadi, and J. P. Meyer. CFD-based shape optimization of steam turbine blade cascade in transonic two phase flows. *Applied Thermal Engineering*, 112:1575–1589, 2017.

[183] G. Nowak and A. Rusin. Shape and operation optimisation of a supercritical steam turbine rotor. *Energy Conversion and Management*, 74:417–425, 2013.

[184] Eni Oko and Meihong Wang. Dynamic modelling, validation and analysis of coal-fired subcritical power plant. *Fuel*, 135:292–300, 2014.

[185] Ioana Opriș and Victor-Eduard Cenușă. Parametric and heuristic optimization of multiple schemes with double-reheat ultra-supercritical steam power plants. *Energy*, 266:126454, 2023.

[186] Tiancheng Ouyang, Xianlin Tan, Xiaoyu Tuo, Peijia Qin, and Chunlan Mo. Performance analysis and multi-objective optimization of a novel CCHP system integrated energy storage in large seagoing vessel. *Renewable Energy*, 224:120185, 2024.

[187] Misha Paauw and Daan van den Berg. Paintings, polygons and plant propagation. In Anikó Ekárt, Antonios Liapis, and María Luz Castro Pena, editors, *Computational Intelligence in Music, Sound, Art and Design*, pages 84–97, Cham, 2019. Springer International Publishing.

[188] Herschel Pangborn, Andrew G. Alleyne, and Ning Wu. A comparison between finite volume and switched moving boundary approaches for dynamic vapor

compression system modeling. *International Journal of Refrigeration*, 53:101–114, 2015.

[189] Jin Seok Park, Jong Wook Kim, and Jae Seon Lee. Complete and homologous pump characteristics for a reactor coolant pump. *Nuclear Engineering and Design*, 357:110425, 2020.

[190] Richard Pearson and Shutaro Takeda. *Review of approaches to fusion energy*. IOP Publishing, 2020.

[191] Joaquim Peiró and Spencer Sherwin. *Finite Difference, Finite Element and Finite Volume Methods for Partial Differential Equations*, pages 2415–2446. Springer Netherlands, Dordrecht, 2005.

[192] Ugo Pelay, Lingai Luo, Yilin Fan, Driss Stitou, and Mark Rood. Thermal energy storage systems for concentrated solar power plants. *Renewable and Sustainable Energy Reviews*, 79:82–100, 2017.

[193] Linda Petzold. A description of dassl: A differential/algebraic system solver. In *10th IMACS World Congress*, 01 1982.

[194] Gabriel Plascencia, Lamberto Diaz-Damacillo, and Minerva Robles-Agudo. On the estimation of the friction factor: a review of recent approaches. *SN Applied Sciences*, 2, 02 2020.

[195] Bruce E Poling, John M Prausnitz, John P O'connell, et al. *The properties of gases and liquids*, volume 5. Mcgraw-hill New York, 2001.

[196] L.A. Porto-Hernandez, J.V.C. Vargas, M.N. Munoz, J. Galeano-Cabral, J.C. Ordonez, W. Balmant, and A.B. Mariano. Fundamental optimization of steam rankine cycle power plants. *Energy Conversion and Management*, 289:117148, 2023.

[197] Cristina Prieto, Laia Miró, Gerard Peiró, Eduard Oró, Antoni Gil, and Luisa F. Cabeza. Temperature distribution and heat losses in molten salts tanks for CSP plants. *Solar Energy*, 135:518–526, 2016.

[198] Cristina Prieto, Pablo D. Tagle-Salazar, David Patiño, Julieta Schallenberg-Rodriguez, Padraig Lyons, and Luisa F. Cabeza. Use of molten salts tanks for seasonal thermal energy storage for high penetration of renewable energies in the grid. *Journal of Energy Storage*, 86:111203, 2024.

[199] Hongtao Qiao. Correction of log mean temperature difference method and effectiveness-NTU relations for two-phase heat transfer with pressure drop and temperature glide. In *International Refrigeration and Air Conditioning Conference 2018*, 2018.

[200] Hongtao Qiao, Vikrant C. Aute, and Reinhard Radermacher. An improved moving boundary heat exchanger model with pressure drop. In *International Refrigeration and Air Conditioning Conference 2014*, 2014.

[201] Hongtao Qiao, Christopher R. Laughman, Vikrant Aute, and Reinhard Radermacher. An advanced switching moving boundary heat exchanger model with pressure drop. *International Journal of Refrigeration*, 65:154–171, 2016.

[202] Yu Qiu, Erqi E, and Qing Li. Triple-objective optimization of sCO_2 brayton cycles for next-generation solar power tower. *Energies*, 16(14), 2023.

[203] Amanda Quadling, David Bowden, Chris Hardie, and Arti Vasanthakumaran. Developing power plant materials using the life cycle lens. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2280):20230409, 2024.

[204] Sylvain Quoilin, Ian Bell, Adriano Desideri, Pierre Dewallef, and Vincent Lemort. Methods to increase the robustness of finite-volume flow models in thermodynamic systems. *Energies*, 7(3):1621–1640, 2014.

[205] Christopher Rackauckas and Qing Nie. DifferentialEquations.jl – a performant and feature-rich ecosystem for solving differential equations in Julia. *Journal of Open Research Software*, May 2017.

[206] Alok K. Ray, Dibakar Rakshit, K. Ravi Kumar, and Hal Gurgenci. Transient discharge performance of high-temperature latent storage system integrated with supercritical CO_2 brayton cycle: A combined analytical and numerical study. *Journal of Energy Storage*, 72:108466, 2023.

[207] Luis Rios and Nikolaos Sahinidis. Derivative-free optimization: A review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56, 11 2009.

[208] E. Rivas, E. Rojas, R. Bayón, W. Gaggioli, L. Rinaldi, and F. Fabrizi. CFD model of a molten salt tank with integrated steam generator. *Energy Procedia*, 49:956–964, 2014.

[209] Erik Rodriguez and Bryan Rasmussen. A comparison of modeling paradigms for dynamic evaporator simulations with variable fluid phases. *Applied Thermal Engineering*, 112:1326–1342, 2017.

[210] I. Rodríguez, C. D. Pérez-Segarra, O. Lehmkuhl, and A. Oliva. Modular object-oriented methodology for the resolution of molten salt storage tanks for CSP plants. *Applied Energy*, 109:402–414, 2013.

[211] Wilfried Roetzel. *Design and Operation of Heat Exchangers and their Networks*. Academic Press, 2020.

[212] Shahin Rostami, Ferrante Neri, and Kiril Gyaurski. On algorithmic descriptions and software implementations for multi-objective optimisation: A comparative study. *SN Computer Science*, 1, 08 2020.

[213] Jens Pålsson Rubén M. Montañés, Johan Windahl and Marcus Thern. Dynamic modeling of a parabolic trough solar thermal power plant with thermal storage using Modelica. *Heat Transfer Engineering*, 39(3):277–292, 2018.

[214] Muhammad Saeed, Saboora Khatoon, and Man-Hoe Kim. Design optimization and performance analysis of a supercritical carbon dioxide recompression

brayton cycle based on the detailed models of the cycle components. *Energy Conversion and Management*, 196:242–260, 2019.

[215] Mouayad A. Sahib and Bestoun S. Ahmed. A new multiobjective performance criterion used in PID tuning optimization algorithms. *Journal of Advanced Research*, 7(1):125–134, 2016.

[216] Abdel Salhi and Eric Fraga. Nature-inspired optimisation approaches and the new plant propagation algorithm. In *Proceeding of ICeMATH2011*, 06 2011.

[217] Jan Schulte-Fischedick, Rainer Tamme, and Ulf Herrmann. CFD analysis of the cool down behaviour of molten salt thermal storage systems. In *ASME 2008 2nd International Conference on Energy Sustainability*, volume 2 of *Energy Sustainability*, pages 515–524, 08 2008.

[218] Simon Scolan, Sylvain Serra, Sabine Sochard, Pierre Delmas, and Jean-Michel Reneaume. Dynamic optimization of the operation of a solar thermal plant. *Solar Energy*, 198:643–657, 2020.

[219] Dale E. Seborg. *Process Dynamics and Control, 4th Edition*. Wiley, New York, 4 edition, 2016.

[220] D.J. Segal, A.J. Cerfon, and J.P. Freidberg. Steady state versus pulsed tokamak reactors. *Nuclear Fusion*, 61(4):045001, mar 2021.

[221] Egberto Selerio, Jr., Joerabell Lourdes Aro, Samantha Shane Evangelista, Fatima Maturan, and Lanndon Ocampo. Optimal configuration of polygeneration plants under process failure, supply chain uncertainties, and emissions policies. *Computers Industrial Engineering*, 172:108637, 2022.

[222] Shitharth Selvarajan. A comprehensive study on modern optimization techniques for engineering applications. *Artificial Intelligence Review*, 57, 07 2024.

[223] Nestor A. Sepulveda, Jesse D. Jenkins, Fernando J. de Sisternes, and Richard K. Lester. The role of firm low-carbon electricity resources in deep decarbonization of power generation. *Joule*, 2(11):2403–2420, 2018.

[224] Aldo Serafino, Benoit Obert, Léa Vergé, and Paola Cinnella. Robust optimization of an organic rankine cycle for geothermal application. *Renewable Energy*, 161:1120–1129, 2020.

[225] M.S. Murshitha Shajahan, D. Najumnissa Jamal, Joy Mathew, Ahamed Anas Ali Akbar, Ajithraj Sivakumar, and Mohamed Shaheen Shahul Hameed. Improvement in efficiency of thermal power plant using optimization and robust controller. *Case Studies in Thermal Engineering*, 33:101891, 2022.

[226] Amit Sharma and Navdeep Singh. Load frequency control of connected multi-area multi-source power systems using energy storage and lyrebird optimization algorithm tuned PID controller. *Journal of Energy Storage*, 100:113609, 2024.

[227] Ji-Hyun Shin and Young-Hum Cho. A variable water flow rate control method of hybrid geothermal heat pump systems. *Renewable Energy*, 226:120241, 2024.

[228] Michael Sielemann, Francesco Casella, Martin Otter, Christoph Clauß, Jonas Eborn, Sven Mattsson, and Hans Olsson. Robust initialization of differential-algebraic equations using homotopy. In *8th International Modelica Conference*, 06 2011.

[229] Mohammad Solgi, Omid Bozorg-Haddad, and Hugo Loaiciga. *Meta-Heuristic and Evolutionary Algorithms for Engineering Optimization*. Wiley, 11 2017.

[230] R. Span. *Multiparameter Equations of State: An Accurate Source of Thermodynamic Property Data*. Engineering online library. Springer, 2000.

[231] R. Span, W. Wagner, E.W. Lemmon, and R.T. Jacobsen. Multiparameter equations of state — recent trends and future challenges. *Fluid Phase Equilibria*, 183-184:1–20, 2001.

[232] James Spelling, Daniel Favrat, Andrew Martin, and Germain Augsburger. Thermoeconomic optimization of a combined-cycle solar tower power plant. *Energy*, 41(1):113–120, 2012.

[233] R. C. Spencer, K. C. Cotton, and C. N. Cannon. A method for predicting the performance of steam turbine-generators: 16,500 kW and larger. *Journal of Engineering for Power*, 85(4):249–298, 10 1963.

[234] Oliver Stein, Jan Oldenburg, and Wolfgang Marquardt. Continuous reformulations of discrete–continuous optimization problems. *Computers Chemical Engineering*, 28(10):1951–1966, 2004.

[235] Muhammad Sulaiman, Abdel Salhi, Eric Fraga, and Wali Mashwani. A novel plant propagation algorithm: Modifications and implementation. *Science International-Lahore*, 02 2016.

[236] M.V.J.J. Suresh, K.S. Reddy, and Ajit Kumar Kolar. Ann-ga based optimization of a high ash coal-fired supercritical power plant. *Applied Energy*, 88(12):4867–4873, 2011.

[237] Christian Suárez, Alfredo Iranzo, F. J. Pino, and J. Guerra. Transient analysis of the cooling process of molten salt thermal storage tanks due to standby heat loss. *Applied Energy*, 142:56–65, 2015.

[238] Charlle L. Sy, Kathleen B. Aviso, Aristotle T. Ubando, and Raymond R. Tan. Target-oriented robust optimization of polygeneration systems under uncertainty. *Energy*, 116:1334–1347, 2016.

[239] Dassault Systemes. Dymola. `https://www.3ds.com/products/catia/dymola`. Accessed: 21/02/2025.

[240] Pablo D. Tagle-Salazar, Cristina Prieto, Anton López-Román, and Luisa F. Cabeza. A transient heat losses model for two-tank storage systems with molten salts. *Renewable Energy*, 219:119371, 2023.

[241] Akinori Tani. 9 - steam turbine design for load-following capability and highly efficient partial operation. In Tadashi Tanuma, editor, *Advances in Steam Turbines for Modern Power Plants (Second Edition)*, Woodhead Publishing Series in Energy, pages 195–207. Woodhead Publishing, second edition edition, 2022.

[242] Jia tao Meng, Shuang qing Chen, Yong gang Cao, Ji yue Liang, Sheng hui Liu, Peng rui Qiao, and Xiao liang Zhu. Thermal-hydraulic analysis of 1–2 passes printed circuit heat exchanger with lead-bismuth and supercritical carbon dioxide as working fluids. *Progress in Nuclear Energy*, 180:105612, 2025.

[243] Wah Yen Tey, Yutaka Asako, Nor Azwadi Che Sidik, and Rui Zher Goh. Governing equations in computational fluid dynamics: Derivations and a recent review. *Progress in Energy and Environment*, 1:1–19, Jun. 2017.

[244] A. Toffolo, S. Rech, and A. Lazzaretto. Generation of complex energy systems by combination of elementary processes. *Journal of Energy Resources Technology*, 140(11):112005, 06 2018.

[245] Vinita Tomar, Mamta Bansal, and Pooja Singh. Metaheuristic algorithms for optimization: A brief review. *Engineering Proceedings*, 59(1), 2023.

[246] Monika Topel and Björn Laumert. Improving concentrating solar power plant performance by increasing steam turbine flexibility at start-up. *Solar Energy*, 165:10–18, 2018.

[247] Veerawat Udomvorakulchaia, Miguel Pineda, and Eric S. Fraga. Introducing competition in a multi-agent system for hybrid optimization. Accepted for publication in proceedings of ESCAPE35.

[248] Alix Untrau, Sabine Sochard, Frédéric Marias, Jean-Michel Reneaume, Galo A. C. Le Roux, and Sylvain Serra. Analysis and future perspectives for the application of dynamic real-time optimization to solar thermal plants: A review. *Solar Energy*, 241:275–291, 2022.

[249] Luca Urbanucci and Daniele Testi. Optimal integrated sizing and operation of a chp system with monte carlo risk analysis for long-term uncertainty in energy demands. *Energy Conversion and Management*, 157:307–316, 2018.

[250] H. van Putten and P. Colonna. Dynamic modeling of steam power cycles: Part ii – simulation of a small simple rankine cycle system. *Applied Thermal Engineering*, 27(14):2566–2582, 2007.

[251] Manuel Jesús Vasallo and José Manuel Bravo. A mpc approach for optimal generation scheduling in CSP plants. *Applied Energy*, 165:357–370, 2016.

[252] Yannic Vaupel, Wolfgang R. Huster, Flemming Holtorf, Adel Mhamdi, and Alexander Mitsos. Analysis and improvement of dynamic heat exchanger models for nominal and start-up operation. *Energy*, 169:1191–1201, 2019.

[253] Marcelo A. Veloso and Joao R. L. de Mattos. Parametric representation of centrifugal pump homologous curves. In *2015 International Nuclear Atlantic Conference*, 2015.

[254] GE Vernova. Condenser | GE Steam Power. `www.ge.com/steam-power/products/heat-exchange/condenser`. Accessed: 29/5/2024.

[255] Guilherme Vescovi, Nicolas Alpy, David Haubensack, Catherine Azzaro-Pantel, and Pascal Stouffs. Paving the way to multi-case optimization of a steam rankine cycle for cogeneration in nuclear power plants. In Ludovic Montastruc and Stephane Negny, editors, *32nd European Symposium on Computer Aided Process Engineering*, volume 51 of *Computer Aided Chemical Engineering*, pages 67–72. Elsevier, 2022.

[256] Wouter Vrielink and Daan van den Berg. A dynamic parameter for the plant propagation algorithm. In *In EvoSTAR2021*, 03 2021.

[257] Wouter Vrielink and Daan van den Berg. Parameter control for the plant propagation algorithm. In *In EvoSTAR2021*, 03 2021.

[258] W. Wagner, J. R. Cooper, A. Dittmann, J. Kijima, H.-J. Kretzschmar, A. Kruse, R. Mareš, K. Oguchi, H. Sato, I. Stöcker, O. Šifner, Y. Takaishi, I. Tanishita, J. Trübenbach, and T. Willkommen. The IAPWS industrial formulation 1997 for the thermodynamic properties of water and steam. *Journal of Engineering for Gas Turbines and Power*, 122(1):150–184, 01 2000.

[259] W. Wagner and A. Pruß. The IAPWS formulation 1995 for the thermodynamic properties of ordinary water substance for general and scientific use. *Journal of Physical and Chemical Reference Data*, 31(2):387–535, 06 2002.

[260] Zhenjie Wan, Jinjia Wei, Mumtaz A. Qaisrani, Jiabin Fang, and Nan Tu. Evaluation on thermal and mechanical performance of the hot tank in the two-tank molten salt heat storage system. *Applied Thermal Engineering*, 167:114775, 2020.

[261] Chaoyang Wang, Zefeng Liu, Mengyang Fan, Yongliang Zhao, Ming Liu, and Junjie Yan. Enhancing the flexibility and efficiency of a double-reheat coal-fired power unit by optimizing the steam temperature control: From simulation to application. *Applied Thermal Engineering*, 217:119240, 2022.

[262] Di Wang, Deying Liu, Chaonan Wang, Yunlong Zhou, Xiaoli Li, and Mei Yang. Flexibility improvement method of coal-fired thermal power plant based on the multi-scale utilization of steam turbine energy storage. *Energy*, 239:122301, 2022.

[263] Di Wang, Deying Liu, Chaonan Wang, Yunlong Zhou, Xiaoli Li, and Mei Yang. Flexibility improvement method of coal-fired thermal power plant based on the multi-scale utilization of steam turbine energy storage. *Energy*, 239:122301, 2022.

[264] Enhua Wang, Ningjian Peng, and Mengru Zhang. System design and application of supercritical and transcritical CO_2 power cycles: A review. *Frontiers in Energy Research*, 9, 2021.

[265] Kun Wang and Ya-Ling He. Thermodynamic analysis and optimization of a molten salt solar power tower integrated with a recompression supercritical CO_2 brayton cycle based on integrated modeling. *Energy Conversion and Management*, 135:336–350, 2017.

[266] Ligang Wang, Matthias Lampe, Philip Voll, Yongping Yang, and André Bardow. Multi-objective superstructure-free synthesis and optimization of thermal power plants. *Energy*, 116:1104–1116, 2016.

[267] Ligang Wang, Philip Voll, Matthias Lampe, Yongping Yang, and André Bardow. Superstructure-free synthesis and optimization of thermal power plants. *Energy*, 91:700–711, 2015.

[268] Ligang Wang, Yongping Yang, Changqing Dong, Tatiana Morosuk, and George Tsatsaronis. Parametric optimization of supercritical coal-fired power plants by minlp and differential evolution. *Energy Conversion and Management*, 85:828–838, 2014.

[269] Ligang Wang, Zhiping Yang, Shivom Sharma, Alberto Mian, Tzu-En Lin, George Tsatsaronis, François Maréchal, and Yongping Yang. A review of evaluation, optimization and synthesis of energy systems: Methodology and application to thermal power plants. *Energies*, 12(1), 2019.

[270] Wei Wang, Jizhen Liu, Deliang Zeng, Yuguang Niu, and Li Sun. Flexible electric power control for coal-fired units by incorporating feedwater bypass. *IEEE Access*, 7:91225–91233, 2019.

[271] Yiming Wang and Gongnan Xie. 4e multi-objective optimization of cold electricity co-generation system based on supercritical CO_2 brayton cycle. *Energy Conversion and Management*, 283:116952, 2023.

[272] Zhiping Wang, Xiaoyu Mo, Peijia Qin, Zhongkai Zhao, and Tiancheng Ouyang. Multi-dimensional assessment and multi-objective optimization of electricity-cooling cogeneration system driven by marine diesel engine waste heat. *Journal of Cleaner Production*, 334:130187, 2022.

[273] Oliver Ward, Federico Galvanin, Nelia Jurado, Chris Clements, Mohamad Abdallah, Daniel Blackburn, and Eric Fraga. Lumped-parameter heat exchanger models for the robust dynamic modelling of power generation cycles. In Antonios C. Kokossis, Michael C. Georgiadis, and Efstratios Pistikopoulos, editors, *33rd European Symposium on Computer Aided Process Engineering*, volume 52 of *Computer Aided Chemical Engineering*, pages 3271–3276. Elsevier, 2023.

[274] Oliver M. G. Ward, Federico Galvanin, Nelia Jurado, Chris Clements, Mohamad Abdallah, Daniel Blackburn, and Eric S. Fraga. Optimization of a PID controller within a dynamic model of a steam rankine cycle with coupled energy storage. *Chemical Engineering Transactions*, 99:583–588, May 2023.

[275] Oliver M.G. Ward, Eric S. Fraga, Federico Galvanin, Nelia Jurado, Daniel Blackburn, Robert J. Warren, and Jack Acres. A dynamic model of a power conversion system with indirect thermal energy storage for a pulsed fusion tokamak for use in design optimisation. *Fusion Engineering and Design*, 219:115289, 2025.

[276] Oliver M.G. Ward, Federico Galvanin, Nelia Jurado, Daniel Blackburn, Robert J. Warren, and Eric S. Fraga. Optimization of the power conversion system for a pulsed fusion power plant with multiple heat sources using a dynamic process model. *Systems and Control Transactions*, 4, 2025.

[277] Supaluck Watanapanich, Sung-Ta Li, and Jui-Yuan Lee. Optimal integration of organic rankine cycles into process heat exchanger networks: A simultaneous approach. *Energy Conversion and Management*, 260:115604, 2022.

[278] Niels Weber and Dirk Zimmer. Regularization of the logarithmic mean for robust numerical simulation. In *10th Vienna International Conference on Mathematical Modelling, MATHMOD 2022*, volume 55 of *IFAC Papers Online*, pages 529–533. Elsevier, July 2022.

[279] King-Leung Wong, Ming-Tsun Ke, and Shi-Shi Ku. The log mean heat transfer rate method of heat exchanger considering the influence of heat radiation. *Energy Conversion and Management*, 50(11):2693–2698, 2009.

[280] Pan Wu, Yunduo Ma, Chuntian Gao, Weihua Liu, Jianqiang Shan, Yanping Huang, Junfeng Wang, Dan Zhang, and Xu Ran. A review of research and development of supercritical carbon dioxide brayton cycle technology in nuclear engineering applications. *Nuclear Engineering and Design*, 368:110767, 2020.

[281] Yingjie Wu, Baokun Liu, Han Zhang, Jiong Guo, and Fu Li. A movable boundary model for helical coiled once-through steam generator using preconditioned jfnk method. *International Journal of Advanced Nuclear Reactor Design and Technology*, 4(1):1–8, 2022.

[282] Liangtao Xie, Jianguo Yang, Xin Yang, Zheng Qin, and Sicong Sun. Design, optimisation and evaluation of the s-CO_2 brayton cycle for marine low-speed engine flue gas. *Case Studies in Thermal Engineering*, 61:105161, 2024.

[283] Qiang Xiong, Wen-Jian Cai, and Mao-Jun He. Equivalent transfer function method for PI/PID controller design of MIMO processes. *Journal of Process Control*, 17(8):665–673, 2007.

[284] Hongyu Xu, Yu Han, Shuo Chen, Cheng Xu, and Yongping Yang. An improved free-superstructure method for optimal synthesis of semi-closed CO_2 power cycles. *Energy*, 316:134372, 2025.

[285] Xiaosheng Yan, Xiaodong Wang, Xu Han, Chunqi Sun, Peng Li, and Zhonghe Han. Study on the thermodynamic performance of a coupled compressed air

energy storage system in a coal-fired power plant. *Journal of Energy Storage*, 68:107727, 2023.

[286] Chendi Yang, Yuanyuan Deng, Ning Zhang, Xiaopeng Zhang, Gaohong He, and Junjiang Bao. Optimal structure design of supercritical CO_2 power cycle for gas turbine waste heat recovery: A superstructure method. *Applied Thermal Engineering*, 198:117515, 2021.

[287] Chendi Yang, Zhiqiang Ni, Xiaopeng Zhang, Ning Zhang, Gaohong He, and Junjiang Bao. Optimal structure design of supercritical CO_2 power cycle for waste heat recovery: An improved intelligent synthesis method. *Energy Conversion and Management*, 272:116365, 2022.

[288] He Yang, Jinduo Li, Zhihua Ge, Lijun Yang, and Xiaoze Du. Dynamic characteristics and control strategy of pumped thermal electricity storage with reversible brayton cycle. *Renewable Energy*, 198:1341–1353, 2022.

[289] Xin-She Yang. *Engineering Optimization: An Introduction with Metaheuristic Applications*. Wiley Online Library: Books. Wiley, 2010.

[290] Yiping Yang, Yulei Huang, Peixue Jiang, and Yinhai Zhu. Multi-objective optimization of combined cooling, heating, and power systems with supercritical CO_2 recompression brayton cycle. *Applied Energy*, 271:115189, 2020.

[291] Lichao Yao and Zhengping Zou. A one-dimensional design methodology for supercritical carbon dioxide brayton cycles: Integration of cycle conceptual design and components preliminary design. *Applied Energy*, 276:115354, 2020.

[292] Adamu Yebi, Bin Xu, Xiaobing Liu, John Shutty, Paul Anschel, Zoran Filipi, Simona Onori, and Mark Hoffman. Estimation and predictive control of a parallel evaporator diesel engine waste heat recovery system. *IEEE Transactions on Control Systems Technology*, 27(1):282–295, Jan 2019.

[293] Haoshui Yu, John Eason, Lorenz T. Biegler, and Xiao Feng. Process integration and superstructure optimization of organic rankine cycles (orcs) with heat exchanger network synthesis. *Computers Chemical Engineering*, 107:257–270, 2017.

[294] Mingzhe Yu, Fubin Yang, Hongguang Zhang, Yinlian Yan, Xu Ping, Yachao Pan, Chengda Xing, and Anren Yang. Thermoeconomic performance of supercritical carbon dioxide brayton cycle systems for cng engine waste heat recovery. *Energy*, 289:129972, 2024.

[295] Shichang Yun, Dalin Zhang, Xinyu Li, Xuanang He, Wenxi Tian, Suizheng Qiu, G.H. Su, and Quanbin Zhao. Superstructure design and optimization on closed brayton cycle system of fluoride-salt-cooled high-temperature advanced reactor. *Applied Energy*, 347:121404, 2023.

[296] Matteo Zaupa, Mauro Dalla Palma, Ivo Moscato, and Luciana Barucca. Balance of plant conceptual design of EU DEMO integrating different breeding blanket concepts. *Fusion Engineering and Design*, 200:114235, 2024.

[297] A. Zavala-Río, Ricardo Femat, and Raul Santiesteban. An analytical study of the logarithmic mean temperature difference. *Revista Mexicana de Ingeniería Química*, 4, 01 2005.

[298] Fritz Zaversky, Javier García-Barberena, Marcelino Sánchez, and David Astrain. Transient molten salt two-tank thermal storage modeling for CSP performance simulations. *Solar Energy*, 93:294–311, 2013.

[299] Lotfi Zeghadnia, Jean Loup Robert, and Bachir Achour. Explicit solutions for turbulent flow friction factor: A review, assessment and approaches classification. *Ain Shams Engineering Journal*, 10(1):243–252, 2019.

[300] Chi Zhai, Wei Sun, and Ahmet Palazoglu. Analysis of periodically forced bioreactors using nonlinear transfer functions. *Journal of Process Control*, 58:90–105, 2017.

[301] Ao Zhang and Xiang Wang. Development of Modelica-based one-dimensional thermodynamic cycle library and its application in simulation and multi-objective optimization of a he–xe closed-brayton-cycle system. *Progress in Nuclear Energy*, 172:105205, 2024.

[302] Dong Zhang, Haochun Zhang, Ying Luo, and Shuting Zhao. Comparative and optimal study of energy, exergy, and exergoeconomic performance in supercritical CO_2 recompression combined cycles with organic rankine, transcritical CO_2, and kalina cycle. *Energy*, 311:133347, 2024.

[303] Qiang Zhang, Zewei Luo, Yongjie Zhao, and Suniaikin Pavel. Thermodynamic analysis and multi-objective optimization of a transcritical CO_2 waste heat recovery system for cruise ship application. *Energy Conversion and Management*, 227:113612, 2021.

[304] Qiang Zhang, Ryan M. Ogren, and Song-Charng Kong. Thermo-economic analysis and multi-objective optimization of a novel waste heat recovery system with a transcritical CO_2 cycle for offshore gas turbine application. *Energy Conversion and Management*, 172:212–227, 2018.

[305] Qijun Zhang, Jianning Dong, Heng Chen, Fuyuan Feng, Gang Xu, Xiuyan Wang, and Tong Liu. Dynamic characteristics and economic analysis of a coal-fired power plant integrated with molten salt thermal energy storage for improving peaking capacity. *Energy*, 290:130132, 2024.

[306] Shunqi Zhang, Ming Liu, Yuegeng Ma, Jiping Liu, and Junjie Yan. Flexibility assessment of a modified double-reheat rankine cycle integrating a regenerative turbine during recuperative heater shutdown processes. *Energy*, 233:121068, 2021.

[307] Shunqi Zhang, Ming Liu, Yongliang Zhao, Jiping Liu, and Junjie Yan. Dynamic simulation and performance analysis of a parabolic trough concentrated solar power plant using molten salt during the start-up process. *Renewable Energy*, 179:1458–1471, 2021.

[308] Hang Zhao, Qinghua Deng, Wenting Huang, Dian Wang, and Zhenping Feng. Thermodynamic and economic analysis and multi-objective optimization of supercritical CO_2 brayton cycles. *Journal of Engineering for Gas Turbines and Power*, 138(8):081602, 2016.

[309] Qiao Zhao, Mounir Mecheri, Thibaut Neveux, Romain Privat, Jean-Noël Jaubert, and Yann Le Moullec. Search for the optimal design of a supercritical-CO_2 brayton power cycle from a superstructure-based approach implemented in a commercial simulation software. *Energies*, 16(14), 2023.

[310] Tao Zhou, Zhengxian Liu, Xiaojian Li, Ming Zhao, and Yijia Zhao. Thermodynamic design space data-mining and multi-objective optimization of sCO_2 brayton cycles. *Energy Conversion and Management*, 249:114844, 2021.

[311] Yunlong Zhou and Di Wang. An improved coordinated control technology for coal-fired boiler-turbine plant based on flexible steam extraction system. *Applied Thermal Engineering*, 125:1047–1060, 2017.

[312] Feng Zhu, Rentian Zhang, Xueqin Liu, Hao Chen, Meng Sun, and Wei Zhou. Experimental research on similarity deviation between prototype and model pumping systems. *Journal of Physics: Conference Series*, 2707(1):012028, 2 2024.

[313] Qiannan Zhu, Xianglong Luo, Bingjian Zhang, and Ying Chen. Mathematical modelling and optimization of a large-scale combined cooling, heat, and power system that incorporates unit changeover and time-of-use electricity price. *Energy Conversion and Management*, 133:385–398, 2017.

[314] Dirk Zimmer, Michael Meißner, and Niels Weber. The DLR ThermoFluid Stream library. *Electronics*, 11(22), 2022.

[315] Mohammad Zoghi, Saleh Gharaie, Nasser Hosseinzadeh, and Ali Zare. 4e analysis and optimization comparison of solar, biomass, geothermal, and wind power systems for green hydrogen-fueled sofcs. *Energy*, 313:133740, 2024.

[316] Karl Zuvela, Isabella Adlington, Syed Salim Aljunied, and John Edwards. Determining the best approach to commercial fusion power. *Energy Science and Technology*, 1:3–19, 1970.

[317] David Zwicker. py-pde: A python package for solving partial differential equations. *Journal of Open Source Software*, 5(48):2158, 2020.

[318] Karl J. Åström and Tore Hägglund. *PID Controllers - Theory, Design, and Tuning (2nd Edition)*. International Society of Automation (ISA), 1995.