UK Atomic Energy Authority

UKAEA-STEP-CP(23)10

A.N. Petrov, S. Wray, K. Zarebski, T. Stroud, J. Acres, T. Owoeye, M. Cannon, F. Christie, T. Jackson, V. Laksharam, S. Stewart, R. Griffiths, N. Brewer, K. Rochford, S. Gribben

An open-source power balance model for the estimation of tokamak net electrical power output

This document is intended for publication in the open literature. It is made available on the understanding that it may not be further circulated and extracts or references may not be published prior to publication of the original when applicable, or without the consent of the UKAEA Publications Officer, Culham Science Centre, Building K1/0/83, Abingdon, Oxfordshire, OX14 3DB, UK.

Enquiries about copyright and reproduction should in the first instance be addressed to the UKAEA Publications Officer, Culham Science Centre, Building K1/0/83 Abingdon, Oxfordshire, OX14 3DB, UK. The United Kingdom Atomic Energy Authority is the copyright holder.

The contents of this document and all other UKAEA Preprints, Reports and Conference Papers are available to view online free at scientific-publications.ukaea.uk/

An open-source power balance model for the estimation of tokamak net electrical power output

A.N. Petrov, S. Wray, K. Zarebski, T. Stroud, J. Acres, T. Owoeye, M. Cannon, F. Christie, T. Jackson, V. Laksharam, S. Stewart, R. Griffiths, N. Brewer, K. Rochford, S. Gribben

This paper has been submitted to
32nd Symposium on Fusion Technology (SOFT), Dubrovnik, Croatia, 18–23 September 2022

An open-source power balance model for the estimation of tokamak net electrical power output

Alexander Petrov*, Tom Stroud, Daniel Blackburn, Taiwo Owoeye, Steven Wray*, Kristian Zarebski, Jack Acres, Mohamad Abdallah, Chris Clements, Marius Cannon, Finlay Christie, Timothy Jackson, Vignesh Laksharam, Samuel Stewart, Rhys Griffiths, Nicholas Brewer, Katherine Rochford, Sophie Gribben

United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, United Kingdom

- Multi-system tokamak power balance
- Python API with OpenModelica back-end
- Open-source software

Abstract: One of the largest hurdles for commercialisation of magnetic confinement fusion has historically been achieving net power – existing experiments require more electricity to keep the tokamak running than it could theoretically generate, and none have been equipped with thermal to electric conversion equipment. When designing a machine that intends to overcome this, there must be a cheap and robust way for the designer to estimate what the net power will be, preferably with the ability to perform parametric sweeps, without having to know the detailed design of each system.

The work presented in this paper is an integrated time-dependent model, describing the power demands of the major tokamak components (magnets, cryogenics, heating and current drive, etc.), as well as the power generated. The physics are implemented in OpenModelica and make use of a Python API (Application Programming Interface) to collect inputs, run studies and record outputs. The model cannot be validated against real world data, since there is no operational tokamak in the world designed for electrical power generation. Therefore, the correctness of each submodule (i.e., the magnet model, the cryogenics model) has been validated either from first principles or via validation against data from JET (Joint European Torus) where possible. The model has been used extensively as part of the work on the UK's Spherical Tokamak for Energy Production (STEP) and has informed decisions on the STEP concept. It is going to be publicly available on GitHub.

Keywords: power balance, net power, tokamak, python modelling

1 Introduction

As is well known, the primary reason to develop nuclear fusion is the production of safe, clean and abundant electricity. One of the major challenges that lies before achieving this is that the supporting systems, such as magnets and heating and current drive, demand more power than what the plasma can produce, or the plant can extract. Significant work has been done in tokamak research facilities around the world over the last few decades to solve this [1, 2, 3], culminating in the International Thermonuclear Experimental Reactor (ITER). For example, ITER is going to use more advanced plasma-facing components [4] and lower-loss superconducting magnets that can support higher magnetic field at the plasma centre for longer operation [5] than most older machines, as well as a multitude of

other systems aimed at increasing the fusion power and decreasing the consumed power. However, the fact remains that no fusion experiment has exported electrical power to the grid, and while ITER will generate far more fusion energy than existing tokamak experiments (the record holder for tokamaks is JET from 2022 [3]), it will not be able to export electricity either.

This poses an anticipated challenge for STEP – the UK's Spherical Tokamak for Energy Production prototype power plant. Can a given tokamak design produce net power during steady-state operation? That question must be answered early on via a power balance model (PBM) relying only on a low fidelity design, meaning any sub-model used will be basic, dependent on high-level equations and simplified engineering. At the same time, the model must take into account all major systems and non-negligible interactions. This is not a trivial task, especially considering that no real system exists that can validate the operation of all models working together. Any power balance model is likely to be only good enough to provide a 'ballpark' value for the power consumption and generation – this may be considered acceptable when most of the plant is in an early concept design stage.

One power balance model including the major parasitic loads – magnets, heating and current drive, cryogenics, waste heat, fuel cycle, vacuum pumping – and power generation, is presented in this paper. It has been used extensively in the early stages of the STEP design. The method of modelling for each system is presented and the validation is discussed briefly. It is assumed the tokamak in the model is in a deuterium-tritium plasma fuelling mode.

2 Software Engineering

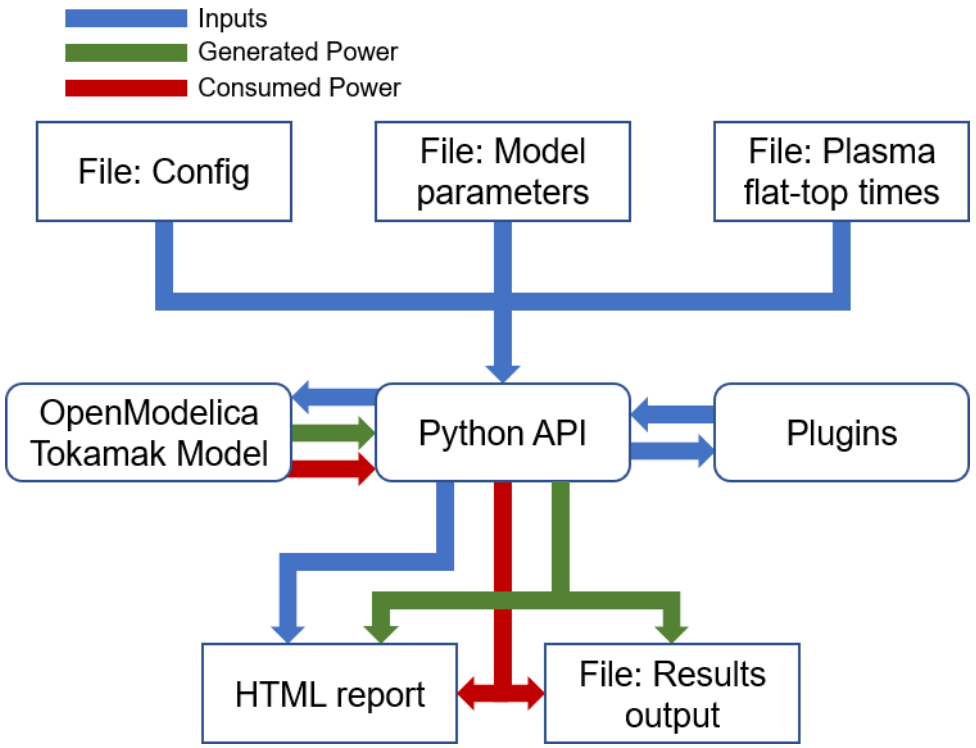

Figure 1: Diagram displaying the flow of data in the power balance model from a software point of view.

All models are built using OpenModelica – a free and open-source equation-based modelling language. Unlike standard programming languages such as C and Python, the execution of OpenModelica code is not sequential. This is because the language collects all lines at once and builds a system of equations which it then solves for each time step, rather than executing one line at a time.

The benefit of using such a language is that it allows the user to easily define its own components or extend the ones already available using equations directly linking two or more variables without having to write functions to describe each physical relationship, thus simplifying the modelling process. Implementing multiphysics is also straightforward. OpenModelica uses object-oriented programming in the sense that components and system models can be defined then reused in a similar way that classes and class instances work in C and Python. It also comes at no cost, pre-packaged with a wide array of components and libraries (electrical, thermal, etc.).

While OpenModelica is good for implementing and simulating the relationships between engineering phenomena, on its own it lacks the ability to handle automated inputs and data handling. Thankfully, it presents to the user a command line interface that can be interacted with via external code. Python was chosen as it is a highly capable, user-friendly language with a large amount of built-in and third-party functionality; it is also the language that the authors are most experienced in. A Python API was developed to communicate with the OpenModelica process and relay information to and from the user in a reliable manner.

3 Physical Models

The PBM system model is built as a tree. At the top is the overall tokamak model, which acts as a point of entry for simulation, instantiates all engineering models, and defines the interactions between them. Each distinct engineering model also instantiates its own sub-models in a similar fashion.

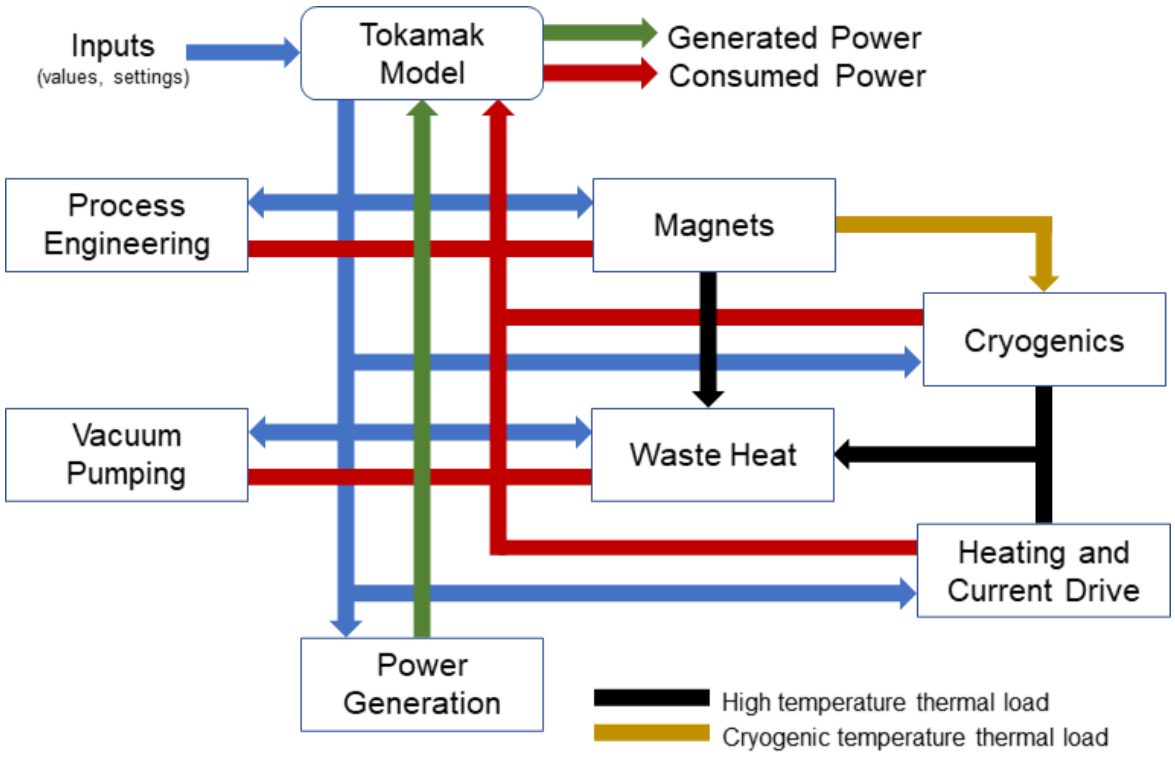

Figure 2: A diagram displaying the flow of information in the OpenModelica part of the PBM. All inputs are grouped in one stream for simplicity.

Using the generated power and each model's consumed power, it can be determined how much electricity the plant may be able to produce. Any abnormal or problematic behaviour can be spotted by an experienced engineer by analysing the generated and consumed powers alongside the input

parameters. See Figure 2 for an illustration of the flow of information in the OpenModelica part of the power balance model.

A PBM simulation is transient and must be given a start and finish time. The plasma current timings can be roughly considered via specifying the start and end time of the flat-top operation. The PBM is primarily interested in the "steady-state" power balance of the plant during this flat-top operation.

3.1 Magnets

In the PBM, an electrical circuit is used to estimate the power use of a single magnet (Figure 3). An individual magnet can be represented as a combination of resistive, inductive and capacitive components. The capacitance may be ignored because one, its contribution is negligible compared to the inductance, and two, the calculation of capacitance is only possible in a more advanced design stage; it is also well-known that fusion coils are far more inductive than capacitive. In terms of resistance, if the coil is superconducting then it may be assumed that it has zero resistance [6] (implemented by not using a resistor, or using a near-zero resistance value for numerical stability); other parts of the magnet system, such as joints between parts of the magnet or between the magnet and feeder, may still contribute with a resistive load. Only the self-inductance is taken into account, although it is possible to extend the model with a full mutual inductance matrix if one is available. Hence, different magnets run in isolation from each other, but their power consumptions are put together to determine the total power consumption of the magnet system. A pre-determined current waveform per coil is the input, and absolute power in MVA is the output.

Additional systems are considered – power supplies, feeders, joints. The power supply efficiency can be represented either as a fixed voltage drop, or more accurately – as a fixed efficiency value α such at the total supplied power is equal to P/α where P is the useful power provided by the power supply. It may be possible to extend the model to use a variable efficiency value based on percentage load. Feeders are represented using a resistance component only, which is acceptable for short length cables for similar reason as the coils – the capacitive and inductive impedance are expected to be far lower than the resistance and are also not readily available at the early design stage. They may be easily used should the need arise. Joints are simply represented as a constant resistance in series with the coil. The losses from the power supplies and feeders occur at room temperature, thus they are passed to the environment and dealt with via the waste heat model; the losses from the joints and the coil itself occur at cryogenic temperature and are thus passed to the cryogenic system.

During transient operation, superconductors experience AC loss due to the change of current (magnetic field) inside them [7]. This loss manifests itself in the form of heat, which can be considerable, but modelling its generation is not trivial. The PBM attempted to calculate rudimentary hysteresis losses in HTS tapes via non-numerical means, but it has not been possible to validate to a high standard; hence AC losses in superconductors are not calculated and will not be discussed in this paper. However, the expected neutronic radiation heating for each coil can be provided as an input.

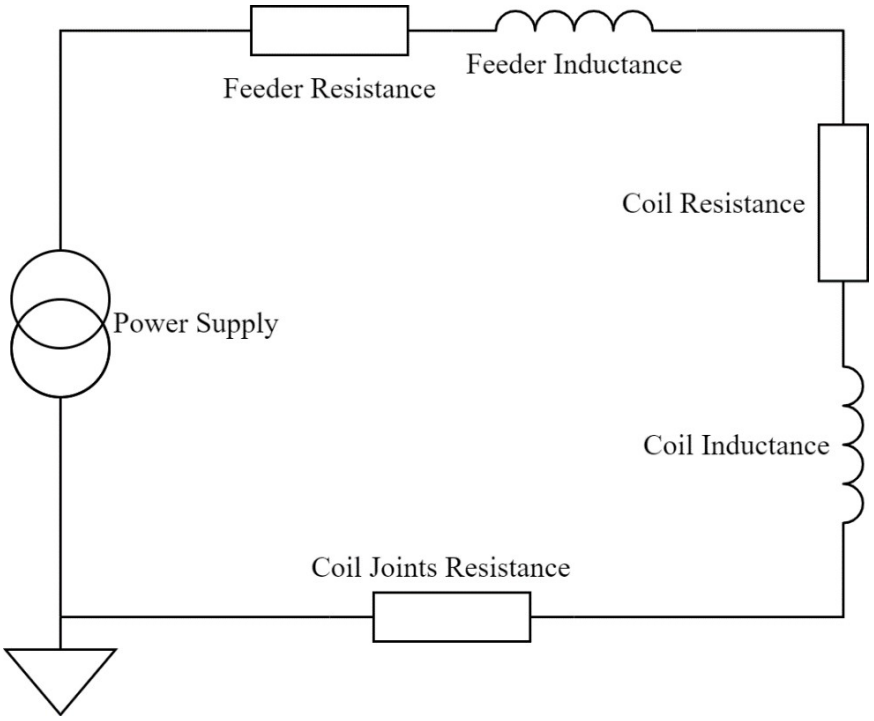

Figure 3: Circuit diagram showing the setup of the magnet model for a single magnet.

3.2 Heating and Current Drive

The heating and current drive (HCD) system calculates the electrical power drawn from the grid and the thermal power delivered to the plasma based on the various design parameters of the HCD system. The design parameters are crucial to accurately evaluate the efficiency of subsystems and more importantly the overall system efficiency, otherwise the model would run the default values. The HCD system library comprises of the Neutral Beam Injector (NBI) such as the Negative Ion Neutral Beam Injector (NINI); and the Radio Frequency (RF) plasma heating such as the Electron Cyclotron Resonance Heating (ECRH). It is possible to have a mix of different HCD systems in the PBM such as a mix of NBI, electron cyclotron resonant heating and electron Bernstein wave.

A simple wall-plug efficiency value for the HCD may be used if the individual parameters are not known.

3.3 Power Generation

Power generation is a simplified model, approximating the performance of either a CO_2 Brayton or Steam Rankine thermodynamic power cycle, in conjunction with a primary blanket coolant. The user then also sets the pressure ratio (for Brayton), system pressure, and finally outlet temperature T, which points the model towards the appropriate 4th order polynomial, that outputs an efficiency ε based on the outlet temperature provided. This is then multiplied by the reactor thermal power and thus an estimate on the power generation is made.

$$\varepsilon = aT^4 + bT^3 - cT^2 + dT - e \quad (1)$$

The coefficients a, b, c and d are selected based on the type of power cycle and the values of the aforementioned system variables. Creation of the polynomials involves a lot of background modelling of the power cycles (external to the PBM), taking into consideration losses and pumping head, making assumptions and simplifications on the coolant loops and their interactions. The modelling was done from fundamentals, and the best fit curves generated for the selected criteria in the PBM. This approach made implementation and running of the PBM simpler.

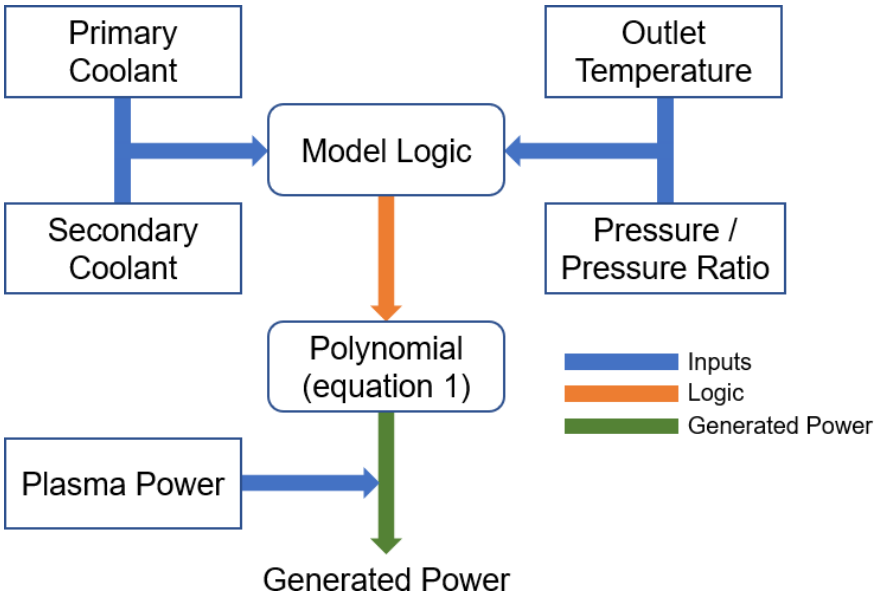

Figure 4 – A flowchart representing the Power Generation model.

3.4 Detritiation

This has been reduced to 4 key streams identified that require detritiation – blanket, power conversion cycle coolant, air-gas, and water. Assumptions have been made up front on applicable technologies and is not representative of the final design or technologies selected for the STEP programme. Each area scales based on the reactor thermal power.

Further work could be performed to include various technologies and combinations, but the model is intended as a high-level overview and the technologies selected are deemed reasonable and potentially the ones that may represent the highest power loads, to be conservative.

The majority of power requirements for the fuel cycle come from heating and pumping operations. These are calculated using the following fundamental calculations:

- **Heating:** All heating is assumed electric, with no heat recovery or integration. This is calculated using the specific heat formula.
- **Gas Compression:** Assumes isentropic compression.

Given the large uncertainty in technologies, simplifications, flowrates and scaling applicability, a contingency factor is applied to all the detritiation power outputs. This is implemented via a constant that can be set by the user.

3.4.1 Coolant Detritiation

Due to the sheer number of coolants and configurations that may exist, this section only focuses on the blanket coolant. Furthermore, it assumes a solid breeder blanket configuration, except in the case of a molten coolant. Depending on the user input for the primary blanket coolant, the following detritiation options will be applied.

- **Water Coolant:** Based on a Combined Electrolysis and Catalytic Exchange (CECE) model *(see 3.4.4 Water Detritiation)*
- **Gas Coolant:** Choice of either He or CO_2. Technology modelled is a catalytic recombination and drier (*see 3.4.2 Air Detritiation)*. Furthermore, the choice can be further subdivided, which affects the gas flowrate through the system:

    - **Carrier:** Coolant is in direct contact with the breeder material and thus has a higher requirement for detritiation.
    - **Non-carrier:** Coolant is not in direct contact with the breeder material and thus requires detritiation based on an estimated quantity that has permeated (*see 3.4.3 Blanket Detritiation* for additional power requirement).
- **Molten Salts Coolant:** Choice of either FLiBe or LiPb. The technology modelled is that of a helium bubble tower, with the tritium removed using a getter bed (*see 3.4.3 Blanket Detritiation*).

The amount of tritium bred is estimated from the tritium breeding ratio (TBR), tritium burn rate, and tritium injection rate into the plasma (for a given power output).

Some of tritium this will then permeate into the non-carrier coolants (i.e. water and gas non-carriers) which is roughly estimated using the correlation proposed by le Claire *et. al.* [8] and assuming a surface area for the blanket and coolant which scales on the thermal duty.

The flowrate of coolant through the detritiation processes is then found based off an allowable limit on activity in the coolant, which by default has been set using CANDU water reactor limits [9] and converting based on molecular weight of the appropriate coolant selected.

The detritiation technologies and thus the power requirements for each coolant are described in other sections, with the except on the molten coolants. This estimates the helium flowrate requirement, scaling off a method employed by Fukada *et. al.* [10] for FLiBe breeder. The helium is then detritiated using a getter as described in 3.4.3 Blanket Detritiation, with heating and compression calculations.

3.4.2 Air Detritiation

Uses a catalytic recombination and drier technology as default to estimate the power associated with detritiation of the tokamak hall. The size of the hall and therefore volume of air passing through the system scales based off the thermal power and estimates of the tokamak hall volume. The power associated with this system is driven by:

- Compression of HVAC air prior to the recombiner reaction
- Heating of the HVAC air prior to the recombiner reaction at 200 °C [11]
- Heating of the regeneration gas (air) for adsorber bed regeneration
- Compression of the regeneration gas (air) for adsorber regeneration

The outstanding parameter left to find is the regeneration gas flowrate, which will then determine the heating and compression requirements. As an overview, this is found by first estimating the bed volume using the gas hourly space velocity (GHSV), which by default is set to 6000 hr^{-1} as per ITER estimates [11]. Using the known HVAC air volume flowrate, the bed size can be found and will scale based off this value. This is a conservative estimate, assuming all the air will be routed through the system to ensure no tritium is present.

Inputting the conditions of the HVAC air, psychrometric correlations can be used to quantify the amount of water in the air, which will be removed by the adsorbent material, regardless of tritiation.

Then by selecting the adsorbent properties (by default these are based off Sigma Aldrich Molecular Sieve 3A [12]), operating conditions (25 and 300 °C for normal and regeneration respectively), and regeneration time, the air flowrate requirement for regeneration can be found.

3.4.3 Blanket Detritiation

This section considers the event of a solid breeder, non-carrier coolant being chosen, which requires a separate stream for tritium extraction. This is assumed to be a helium purge utilising a getter bed material, for which the power requirements arrive from compression and heating of the helium purge and the regeneration of the getter and regeneration gas; which is similar to the catalytic recombination and drier set-up described in 3.4.1 Coolant Detritiation.

In this instance, however, the getter bed temperatures are higher (300 and 600 °C for operation and regeneration respectively), with a hydrogen saturation (0.02 wt/wt%) and lower heat capacity (0.4 kJ/(kg K)). The default values are based off ST-707 getter material. [13]

The final difference is the amount of tritium in the helium purge is purely the amount being bred with no permeation allowance.

3.4.4 Water Detritiation

Modelled on a Combined Electrolysis and Catalytic Exchange (CECE) technology, whereby the power users are water heating, electrolysis, and hydrogen gas compression. The electrolysis requirement has been set as a constant of 26895 kJ/kg_{Water} [14], but the water heating requirement must account for the latent heat, as well as the sensible heat, as the water is being vaporised.

Further parameters are required as the water is flowing through a column in order to detritiate the water. This requires setting a number of parameters such as the reflux ratio and vapour fraction, which are best described by Boniface *et. al.* [15]

The water flowrate into the unit scales linearly on the thermal power, assuming an amount is produced by the fuel cycle and from the various catalytic recombination and drier units as discussed previously.

3.5 Vacuum Pumping

Vacuum pumping requires selecting either a cryogenic or turbomolecular pump, with a series of roughing pumps further downstream. The flowrate through the pumps and out of the reactor is scaled on the thermal power and matter injection.

Turbomolecular pumping power is found quite simply by taking a given power, speed and throughput for a single pump and scaling based off a given flowrate. The same is then applied for the roughing pumps and summed to give a total duty for the turbomolecular pump system.

The cryogenic pump has 2 loads: cryogen and backing pumps. The cryogen requirement is found by estimating the heat of fusion, vaporisation and change in specific heat of the hydrogen isotopes (other impurities have been ignored to simplify the calculation). The quantity of these entering is an input from the matter injection and burn rate, scaled for a given thermal power. The cryogen usage is then converted to an electrical power as described in 3.6 Cryogenic. The backing pump power is found using the same approach as for a turbopump, using known power requirements and scaling appropriately.

3.6 Cryogenic Plant

For the cryoplant, 7 main potential loads have been identified: Toroidal Field (TF) Coils, Poloidal Field (PF) Coils, Current Leads, Fuel Matter Injection (FMI), Cryo-distillation (CD) and Cryopumps (if appropriate). All models except FMI and CD have been described elsewhere in this paper, but they all calculate either a cryogenic (cooling) power or a cryogen flowrate requirement, which is converted to an electrical load as described below.

Similar to the Detritiation models, a contingency factor is applied to the power output to allow for the uncertainties and over-simplification of the models. This can be set by the user.

3.6.1 Converting Cryogenic to an Electrical Load

The cryogenic power is converted to an electrical load using a figure of merit (FOM) assigned to the cryogen temperature ranges that are used in STEP, taking the cryogen temperature as an input and assigning the relevant FOM to it. The temperature is also used to calculate a Carnot efficiency based off a reference temperature (default is room temperature at 298 K) and multiplying together to give an equivalent electrical load.

Cryogenic flow is converted to an electrical load by finding the enthalpy change of the cryogen from reference to the cryogenic temperature and multiplying by the flowrate and the FOM as described above.

3.6.2 Fuel Matter Injection

This model is very simply scaled based off JET data requirements for cryogenic flowrate, which is then scaled on the reactor power output. Alternatively, it can be input as a fixed value by the user if known. The cryogenic flowrate is then fed into the conversion calculation as described above (3.6.1 Converting Cryogenic to an Electrical Load).

3.6.3 Cryo-distillation

Inspiration for this has been sought by looking at the ITER isotope separation system (ISS) modelling approach as described by Noh *et. al.* [16] As the modelling of distillation is a very complex and involved process, the reflux ratios have been fixed based on the paper and the top product flowrates scale based on the reactor thermal output and matter injection. Using the latent and sensible heats for the hydrogen isotopes, the cryogenic requirement is estimated and then converted to an electrical load using the function described above (3.6.1 Converting Cryogenic to an Electrical Load).

3.7 Waste Heat

The aim of this model is to capture a value for the overall low grade waste heat produced by the tokamak systems and calculate the added parasitic load of a waste heat rejection system. The model captures the waste heat figures from the major plant items and separates them into either being cooled by water or by the heating, ventilation, and air conditioning (HVAC) system. In general, water-cooling is less energy intensive than HVAC cooling [17]. Which of the two methods is used depends on which part of the plant the waste heat is coming from, and whether it is possible to water-cool it. For example, the electrical generator can be water-cooled, whereas it would not be practical to water-cool heat loss from a pipe. Some components must be HVAC-cooled because of the tritium contamination challenges involved and the need for detritiation (removal of tritium from a volume of material). The model also scales the waste heat loads linearly with the thermal output from the reactor changes.

The added parasitic load for the water-cooling system is calculated as follows:

- Use the water-cooled waste heat value to calculate the mass flow rate of cooling water needed using the specific heat equation [18].
- Length/elevation/fitting constant figures were estimated once and then used to calculate the pumping power.
- Finally, a scalable cooling tower load is added to the total extra parasitic load assuming an ‘induced draft’ system [19].

The HVAC loads in the model are cooled using a chiller rather than solely a flow of air. This is because in most of the tokamak contamination zones, an air detritiation is needed to stop tritium being released into the atmosphere. This approach will result in a more conservative value. The chiller power consumption is estimated using a scalable rule of thumb based on the amount of waste heat – every megawatt of waste heat results in 0.3 MW chiller power consumption (found using previous expertise based on JET). The chilled water pumping power is calculated using the same method shown above for the cooling water pumping power.

The overall added parasitic load of the waste heat rejection system is the sum of the pumping power required for water-cooled components, the cooling tower load, the HVAC chiller power consumption and the HVAC chilled water pumping power.

4 Methods of Validation

It is not possible to validate the power balance model as a whole, because there is no machine in existence that contains all systems required. Hence, each model is validated on its own – against limited experimental/simulation data or via first principles.

4.1 Magnets

The validation of the magnet systems is aimed at verifying that an RL circuit can accurately capture the losses, i.e., if there isn't any loss contributor that hasn't been identified. Current and voltage waveforms for one coil were taken from one pulse in the Joint European Torus (JET) alongside the coil's resistance and inductance. The current waveform was applied to the magnet model in OpenModelica and yielded the same voltage, thus power, within 6% of the JET data. Hence, an RL-circuit is sufficient. Note that only the coil itself has been validated – no validation has been made for the power loss inside the feeders, joints and power supply because their representation as circuit elements is self-evident from first principles.

4.2 Heating and Current Drive

The validation of the HCD system is aimed at accurately estimating the overall electrical efficiency of the system and the thermal power reaching the plasma. The empirical and mathematical models of the HCD system were first developed, breaking down the system into sub-models to estimate their efficiencies. These were then validated using real experimental data from JET (NINI model) or results published in journal papers [20, 21, 22, 23, 24, 25].

4.3 Power Generation

As described above, the model is based off a polynomial curve that was generated by modelling various power cycle arrangements. These were generated from first principles and represent the best guess arrangements at the time of creation. As such the validation is performed simply by ensuring the background calculations were implemented correctly and the assumptions made are reasonable.

Furthermore, it should be noted that process simulation packages (e.g. Aspen HYSYS) has been used to generate more accurate power outputs and validate the original calculations made. In order simplify the PBM, the decision to stick with the original equations was made.

4.4 Detritiation

Similar to the Power Generation section, calculations are based off first principles and input data taken from sources were possible. For example, reaction conditions are broadly known, and vessel scaling can be made using GHSV from known operating units or designs for ITER. This in itself acts as the validation and is kept customisable to the user as better/different information is made available.

Furthermore, a contingency factor is set to allow for the broad range of uncertainty in the modelling and captures the confidence that is had in the calculations.

4.5 Vacuum Pumping

Again, validation is performed on first principles calculation basis and using known power values for pumps and scaling. Input was sought internally from vacuum pumping experts and verified (where possible) against more detailed modelling.

4.6 Cryogenics

First principal calculations and figure-of-merit application is the easiest conversion to make for electrical conversion, with the values for each cryogenic temperature being set by the internal cryogenics expert. Again, it is verified against real units or designs such as ITER were possible.

4.7 Waste Heat

Firstly, the overall amount of low-grade waste heat was checked against values for similar balance of plant systems and individual components to ensure a realistic result. Two of the main contributors to the overall waste heat were the cryogenic plant and the heating and current drive system. The figures for both were based on a combination of scaling from ITER and the work previously done by the STEP programme in these areas. Models for mass flow rate and pumping power are based on standardised calculations and the inputs and assumptions were sanity checked.

5 Conclusions

This paper presents a power balance model (PBM) used in the early design of the Spherical Tokamak for Energy Production (STEP). Each sub-model's modelling and validation methodology and underlying principles are discussed without going in detail.

The power balance model is far from complete. It is hoped that with the eventual operation of STEP, it can be corrected, expanded and validated against a real-world machine. Of most interest is the model of power generation and any non-negligible but missed sources of power consumption. The authors hope to engage the wider fusion community in the development and maintenance of the PBM. Any interested party may benefit from the established API – thus skipping the need to develop from scratch – and to contribute to it for the betterment of fusion design.

The software is open sourced under a GNU LGPL 2.1 license and can be found with a DOI: 10.5281/zenodo.6341685, and the code itself can be found on the following web address:

https://github.com/ukaea/powerbalance

6 Acknowledgements

This work is funded by the United Kingdom Atomic Energy Authority (UKAEA).

References

[1] Tokamak Energy, "Tokamak Energy moves closer to commercial fusion: 100M degree plasma a world record for a spherical tokamak," 2022 March 2022. [Online]. Available: https://www.tokamakenergy.co.uk/tokamak-energy-moves-closer-to-commercial-fusion/. [Accessed 14 July 2022].

[2] ITER Organization, “60 Years of Progress,” [Online]. Available: https://www.iter.org/sci/beyonditer. [Accessed 14 July 2022].

[3] EUROfusion, “European researchers achieve fusion energy record,” 9 February 2022. [Online]. Available: https://www.euro-fusion.org/news/2022/european-researchers-achieve-fusion-energy-record/. [Accessed 14 July 2022].

[4] M. Merola, D. Loesser, A. Martin, P. Chappuis, R. Mitteau, V. Komarov, R. A. Pitts, S. Gicquel, V. Barabash, L. Giancarli, J. Palmer, M. Nakahira, A. Loarte, D. Campbell, R. Eaton, A. Kukushkin, M. Sugihara, F. Zhang, C. S. Kim, R. Raffray, L. Ferrand, D., “ITER plasma-facing components,” *Fusion Engineering and Design,* vol. 85, no. 10, pp. 2312-2322, 201.

[5] N. Mitchell, A. Devred, P. Libeyre, B. Lim and F. Savary, “The ITER Magnets: Design and Construction Status,” *IEEE Transactions on Applied Superconductivity,* vol. 22, no. 3, pp. 4200809-4200809, 2012.

[6] K. Fossheim and A. Sudboe, Superconductivity: Physics and Applications, Wiley, 2005.

[7] A. P. Malozemoff, J. Yuan and C. M. Rey, “5 - High-temperature superconducting (HTS) AC cables for power grid applications,” in *Superconductors in the Power Grid*, Woodhead Publishing, 2015, pp. 145-150.

[8] A. D. Le Claire, “Permeation of Hydrogen Isotopes in Structural Alloys,” *Journal of Nuclear Materials,* vol. 122 & 123, pp. 1558-1559, 1984.

[9] S. K. Sood, C. Fong, K. M. Kalyanam and K. Woodall, “A Compact, Low Cost, Tritium Removal Plant for CANDU-6 Reactors,” in *Applied Nuclear Research and Development*, Ontario, 1997.

[10] S. Fukada, M. Nishikawa, A. Sagara and T. Terai, “Mass-transport Properties to Estimate Rates of Tritium Recovery from FLiBe Blanket,” *Fusion Science and Technology,* vol. 41, pp. 1054-1058, 2002.

[11] F. Borgognoni, C. Rizzello and S. Tosti, “Experimental Study of Detritiation Catalyst Poisoning,” *Fusion Engineering and Design,* vol. 83, pp. 1375-1379, 2008.

[12] Sigma Aldrich, “Molecular Sieves - Technical Information Bulletin,” Merck, 2021. [Online]. Available: https://www.sigmaaldrich.com/chemistry/chemical-synthesis/learning-center/technical-bulletins/al-1430/molecular-sieves.html. [Accessed 2021].

[13] A. Santucci, M. Incelli, L. Noschese, C. Moreno, F. D. Fonzo, M. Utili, S. Tosti and C. Day, “The Issue of Tritium in DEMO Coolant Mitigation Strategies,” *Fusion Engineering and Design,* vol. 158, 2020.

[14] M. Lehner, R. Tichler, H. Steinmuller and M. Koppe, Power-to-Gas: Technology and Business Models, Springer, 2014.

[15] H. Boniface, I. Castillo, A. E. Everatt and D. Ryland, “A Light-water Detritiation Project at Chalk River Laboratories,” *Fusion Science and Technology,* vol. 60, no. 4, pp. 1327-1330, 2011.

[16] J. Noh, A. M. Fulgueras, L. J. Sebastian, H. G. Lee, D. S. Kim and J. Cho, “Estimation of Thermodynamic Properties of Hydrogen Isotopes and Modeling of Hydrogen Isotope Systems

Using Aspen Plus Simulator," *Journal of Industrial and Engineering Chemistry,* vol. 46, pp. 1-8, 2017.

[17] Lubna A. Hussein, Adnan A. Ateeq, and Raad Z. Homod, "Energy Saving by Reinforcement Learning for Multi-Chillers of HVAC Systems," in *IMDC-IST 2021*, Sakarya, Turkey, 2022.

[18] R. K. Sinnott, J. M. Coulson, and J. F. Richardson, Chemical Engineering Design, Volume 6, Elsevier Butterworth-Heinemann, 1999.

[19] SPX Cooling Technologies, Cooling Tower Fundamentals, Overland Park, Kansas USA: SPX Cooling Technologies Inc., 2009.

[20] Y. Oda, R. Ikeda, K. Kajiwara, T. Kobayashi, K. Hayashi, K. Takahashi, S. Moriyama, K. Sakamoto, T. Eguchi, Y. Kawakami, Y. Mitsunaka, C. Darbos and M. Henderson, "Development of the first ITER gyrotron in QST," *Nucl. Fusion,* vol. 59, 2019.

[21] M. Tsuneoka, H. Fujita, K. Sakamoto, A. Kasugai, T. Imai, T. Nagashima, T. Asaka, N. Kamioka, M. Yasuda, T. Iiyama, T. Yoshida, H. Nara and M. Ishibashi, "Development of d.c. power supply for gyrotron with energy recovery system," *Fusion Engineering and Design,* vol. 36, no. 4, pp. 461-469, 1997.

[22] P. C. Kalaria, M. V. Kartikeyan and M. Thumm, "Design of 170GHz, 1.5-MW Conventional Cavity Gyrotron for Plasma Heating," *IEEE Transactions on Plasma Science,* vol. 42, no. 6, pp. 1522-1528, June 2014.

[23] D. Ciazynski, J. L. Duchateau, P. Decool, P. Libeyre and B. Turck, "Large Superconducting Conductors and Joints for Fusion Magnets: from Conceptual Design to Test at Full Size Scale," Euratom-CEA Association, CEA/Cadarache, France.

[24] M. A. Henderson, S. Alberti, J. Bird, J. Doane, B. Elzendoorn, C. Flemming, T. P. Goodman, F. Hoekzema, J. P. Hogge, J. C. Magnin, B. Pioscyk, L. Porte, M. Q. Tran and A. G. A. Verhoeven, "An ITER-relevant evacuated waveguide transmission system for the JET-EP ECRH project," *Nuclear Fussion,* vol. 43, no. 11, pp. 1-14, 2003.

[25] R. S. Hemsworth and D. Boilson, "Research, Design, and Development Needed to Realise a Neutral Beam Injection System for a Fusion Reactor," in *Fusion Energy*, IntechOpen, 2019.

[26] S. Stewart, "Modelling of Radio Frequency Heating and Current Drive Systems for Nuclear Fusion," 2020.