UNIVERSIDADE FEDERAL DE MINAS GERAIS
Instituto de Ciências Exatas
Programa de Pós-Graduação em Ciência da Computação

Túlio Marcondes Moreira

Automatic Flexible Control of Tidal Range Structures

Belo Horizonte
2022

Túlio Marcondes Moreira

Automatic Flexible Control of Tidal Range Structures

Final Version

Dissertation presented to the Graduate Program in Computer Science of the Federal University of Minas Gerais in partial fulfillment of the requirements for the degree of Doctor in Computer Science.

Advisor: Gilberto Medeiros Ribeiro
Co-Advisor: Pedro Olmo Stancioli Vaz De Melo

Belo Horizonte
2022

Moreira,Túlio Marcondes.

M838A Automatic flexible control of tidal range structures [manuscrito] / Túlio Marcondes Moreira. – 2022.
185 f. il.

Orientador: Gilberto Medeiros Ribeiro.
Coorientador: Pedro Olmo Stancioli Vaz De Melo.
Tese (Doutorado) - Universidade Federal de Minas Gerais, Instituto de Ciências Exatas, Departamento de Ciências da Computação.
Referências: f. 167-179.

1. Computação – Teses. 2. Aprendizado profundo – Teses. 3. energia maremotriz – Teses. 4. Barragens maremotriz – Teses. I. Ribeiro, Gilberto Medeiros. II. Melo, Pedro Olmo Stancioli Vaz de. III. Universidade Federal de Minas Gerais, Instituto de Ciências Exatas, Departamento de Computação. IV.Título.

CDU 519.6*82.10(043)

Ficha Ficha catalográfica elaborada pela bibliotecária Belkiz Inez Rezende Costa CRB 6/1510 Universidade Federal de Minas Gerais - ICEx

UNIVERSIDADE FEDERAL DE MINAS GERAIS
INSTITUTO DE CIÊNCIAS EXATAS
DEPARTAMENTO DE CIÊNCIA DA COMPUTAÇÃO
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIA DA COMPUTAÇÃO

FOLHA DE APROVAÇÃO

AUTOMATIC FLEXIBLE CONTROL OF TIDAL RANGE STRUCTURES

TÚLIO MARCONDES MOREIRA

Tese defendida e aprovada pela banca examinadora constituída pelos Senhores:

Prof. Gilberto Medeiros Ribeiro - Orientador
Departamento de Ciência da Computação - UFMG

Prof. Pedro Olmo Stancioli Vaz de Melo - Coorientador
Departamento de Ciência da Computação - UFMG

Prof. Rodney Rezende Saldanha
Departamento de Engenharia Elétrica - UFMG

Prof. Adriano Alonso Veloso
Departamento de Ciência da Computação - UFMG

Prof. Luiz Chaimowicz
Departamento de Ciência da Computação - UFMG

Prof. Divanilson Rodrigo de Sousa Campelo
Centro de Informática - UFPE

Prof. Segen Farid Estefen
COPPE - UFRJ

Belo Horizonte, 28 de junho de 2022.

Documento assinado eletronicamente por **Gilberto Medeiros Ribeiro**, **Membro de comissão**, em 31/08/2022, às 14:50, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

Documento assinado eletronicamente por **Pedro Olmo Stancioli Vaz de Melo**, **Professor do Magistério Superior**, em 31/08/2022, às 18:28, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

Documento assinado eletronicamente por **Adriano Alonso Veloso**, **Professor do Magistério Superior**, em 01/09/2022, às 09:32, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

Documento assinado eletronicamente por **Luiz Chaimowicz**, **Professor do Magistério Superior**, em 01/09/2022, às 09:38, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

Documento assinado eletronicamente por **Rodney Rezende Saldanha**, **Membro de comissão**, em 05/09/2022, às 15:45, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

Documento assinado eletronicamente por **Divanilson Rodrigo de Sousa Campelo**, **Usuário Externo**, em 09/09/2022, às 10:10, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

Documento assinado eletronicamente por **Segen Farid Estefen**, **Usuário Externo**, em 12/09/2022, às 10:21, conforme horário oficial de Brasília, com fundamento no art. 5º do Decreto nº 10.543, de 13 de novembro de 2020.

A autenticidade deste documento pode ser conferida no site https://sei.ufmg.br/sei/controlador_externo.php?acao=documento_conferir&id_orgao_acesso_externo=0, informando o código verificador **1723430** e o código CRC **A828F7A6**.

Dedico este trabalho à minha família, por desde cedo fomentar minha curiosidade e me ensinar a valorizar o trabalho acadêmico. Aos meus amigos, pelo apoio (mesmo que à distância) durante a pandemia e doutorado. À minha amada Beta, pelo companheirismo e suporte em meus momentos mais difíceis.

Acknowledgments

To my advisors, Professors Gilberto Medeiros Ribeiro and Pedro Olmo, for the invaluable assistance and guidance during my PhD.
To my family, specially my father, for the countless hours of discussion and insights.
To my beloved Beta, for her companionship and love.
To my friends from Oxford University, specially Ahmad Firdaus and Vance Tan Zong Hao, for their support even across the globe.
To all my Brazilian friends, whose presence I missed dearly during the Pandemic.
To the PPGCC's Professors and Staff, for their knowledge, professionalism and assistance.
To Centro Nacional de Processamento de Alto Desempenho (CENAPAD – UFMG) and Advanced Research Computing (ArC – Oxford University), for providing super-computing facilities.
To the funding agency Comissão de Aperfeiçoamento do Pessoal de Ensino Superior (CAPES), for funding this research.

“*What we want is a machine*
that can learn from experience.”
(Alain Turing, 1947)

Resumo

Nos últimos anos, a necessidade de aumentar a capacidade instalada total de fontes renováveis de energia (eólica, solar, bioenergia e hídrica) tem recebido reconhecimento mundial como forma de mitigar os efeitos danosos causados pelas mudanças climáticas resultantes do uso extensivo de combustíveis fósseis. Embora a utilização de fontes renováveis como a energia solar e eólica tenha expandido, a energia das marés permanece praticamente inexplorada. Até hoje, apenas dois grandes projetos de barragens maremotrizes foram construídos, a saber, La Rance (França) e Lake Sihwa (Coreia do Sul), com 240 MW e 254 MW de capacidade instalada, respectivamente.

Um problema recorrente entre as várias propostas de barragens maremotrizes (e.g., Swansea Bay Tidal Lagoon, Severn Barrage, Cardiff Tidal Lagoon) tem sido os altos custos de eletricidade, que estão fortemente relacionados à estratégia operacional utilizada ao estimar a produção de energia desses sistemas. Portanto, esta tese pretende preencher lacunas na literatura, propondo métodos generalistas que possibilitem a primeira operação automática (em tempo real) e otimizada de barragens maremotrizes. Isso permitiria reduzir os custos de energia elétrica por meio (i) da maximização da energia elétrica produzida e (ii) da redução do custo operacional com a operação automática desses sistemas, permitindo que essa tecnologia seja mais atrativa economicamente.

Dada uma simulação precisa das estruturas hidráulicas que compõem uma barragem maremotriz, conseguimos alcançar os objetivos apresentados com (a) modificações nos métodos estado da arte, permitindo a operação em tempo real e com controle independente de vertedouros e (b) uma nova abordagem com Aprendizado por Reforço Profundo (DRL) projetada para operar barragens maremotrizes. Como estudo de caso para o nosso trabalho, focamos na modelagem 0D de dois projetos: La Rance Tidal Barrage (já contruída) e o projeto da Lagoa de Swansea no Reino Unido (SBL).

Por meio de técnicas paramétricas e de leis de afinidade, uma representação 0D para La Rance é criada e validada em relação a dados medidos pela Electricité de France (EDF) – companhia responsável pelo operação de La Rance, desde sua construção. O modelo 0D parametrizado é então operado de forma ótima com nosso método DRL, gerando resultados quantitativos e uma estratégia de operação comparáveis às observadas em La Rance. Ao melhor de nosso conhecimento, esta é a primeira validação de um modelo 0D e da estratégia de operação de TRS em comparação a uma TRS construída. Além disso, a parametrização das estruturas de La Rance permitiu uma descrição generalizável do modo operacional de bombeamento das turbinas, que é implementado no modelo 0D de SBL.

Mostramos que, utilizando o modelo 0D de SBL com bombeamento, ambos os métodos (a) e (b) desenvolvidos alcançam uma geração de energia superior ao upper bound do estado da arte (cenários onde as previsões de maré são iguais às medições) em $2,68\%$ e $3,14\%$, respectivamente (levando a um aumento do lucro anual de até 967.267£, para SBL), enquanto operam a TRS em tempo real. Adicionalmente, o método DRL (b) alcança esse objetivo sem necessitar de previsões de maré, enquanto performa uma operação dinâmica das turbinas em modo de bombeamento, incluindo cenários com queda positiva (auxiliada pela gravidade) – essas duas últimas características são observadas em medições de La Rance, e até então não nenhum método de controle de TRS havia sido capaz de realizar essa tarefa.

Palavras-chave: energia maremotriz, barragens maremotrizes, aprendizado por reforço profundo, “proximal policy optimisation”.

Abstract

In recent years, the need to increase the total installed capacity of renewable energy sources (wind, solar, bioenergy and hydro) has received worldwide recognition as a means to mitigate the damaging effects of climate change that result from the extensive use of fossil fuels. While significant progress has been made in expanding solar and wind resources, tidal energy remains practically untapped. As of today, only two successful large Tidal Range Structure (TRS) projects have been built, namely, La Rance (France) and Lake Sihwa (South Korea), with 240 MW and 254 MW of installed capacity, respectively.

A recurrent issue among the several proposals for TRS (e.g. Swansea Bay Tidal Lagoon, Severn Barrage, Cardiff Tidal Lagoon) has been the high electricity costs, which are strongly related to the operational strategy utilised when estimating TRS energy yield. Therefore, this thesis's project intents to fill gaps in the literature by proposing general methods that enable the first real-time optimised operation of TRS. This would allow for reducing the costs of electricity through (i) energy maximisation of TRS and (ii) operational cost reduction (automatic TRS), allowing this technology to be more economically attractive.

Given an accurate simulation of the hydraulic structures that compose a TRS, objectives (i) and (ii) are solved either through: (a) improvements to the state-of-art approach from the literature, enabling real-time control of TRS with independent sluice operation, or (b) utilising a Deep Reinforcement Learning approach (DRL) designed to operate TRS. As a case study for our work we focused on the 0D modelling of two TRS projects: the constructed La Rance Tidal Barrage and the Swansea Bay Tidal Lagoon (SBL) pathfinder project.

Through parametric and affinity laws techniques, a 0D La Rance model representation is created using measured data from Electricité de France (EDF) – the utility company responsible for operating La Rance since its construction. The parametrised 0D model is then optimally operated with our DRL method, yielding comparable quantitative results and strategy to observed measurements in La Rance. To the best of our knowledge, this is the first validation of a 0D model and TRS operational strategy against a constructed TRS. Furthermore, the development of the 0D La Rance model enabled a generalizable description of pump operational modes, which is implemented into the 0D model of the SBL.

We show that, by using the SBL 0D model with pumping, both our developed (a) and (b) methods achieve energy generation superior to the state-of-art upper bound (scenarios where tidal predictions equal tidal measurements) by 2.68% and 3.14%, respectively (yielding a revenue gain up to 967, 267£ per year, for the SBL), while operating the TRS in real time. The DRL (b) method, however, manages this feat without requiring tidal predictions, performing fine-tuned operation of turbines in pump mode and pumping even with positive (gravity-assisted) head differences – these last two characteristics are observed in La Rance's measurements, but so far no TRS control method had been able to accomplish this task.

Palavras-chave: tidal energy, tidal range structures, deep reinforcement learning, proximal policy optimisation.

Related Publications

- https://doi.org/10.1115/OMAE2019-95767 SCHNABL, ANDREA M.; MOREIRA, TULIO MARCONDES; WOOD, DYLAN; KUBATKO, ETHAN J.; HOULSBY, GUY T.; MCADAM, ROSS A.; ADCOCK, THOMAS A. A.. Implementation of Tidal Stream Turbines and Tidal Barrage Structures in DG-SWEM. In: ASME 2019 38th International Conference on Ocean, Offshore and Arctic Engineering, 2019, Glasgow. Volume 10: Ocean Renewable Energy, 2019. v. 1, 10 pages.

- https://doi.org/10.1115/OMAE2019-95075 MA, QIAN; MOREIRA, TULIO MARCONDES; ADCOCK, THOMAS A. A.. Impact of the Swansea Bay Lagoon on Storm Surges in the Bristol Channel. In: ASME 2019 38th International Conference on Ocean, Offshore and Arctic Engineering, 2019, Glasgow. Volume 10: Ocean Renewable Energy, 2019. v. 1, 10 pages.

- https://doi.org/10.1007/s40722-019-00143-w MA, QIAN; MOREIRA, TULIO MARCONDES; ADCOCK, THOMAS A. A.. The impact of a tidal barrage on coastal flooding due to storm surge in the Severn Estuary. JOURNAL OF OCEAN ENGINEERING AND MARINE ENERGY (PRINT), v. 5, p. 217-226, 2019.

- https://doi.org/10.1016/j.oceaneng.2022.110657 MOREIRA, TULIO MARCONDES; FARIA JR, JACKSON GERALDO; VAZ DE MELO, PEDRO O.S.; CHAIMOVICZ, L; MEDEIROS-RIBERO G.. Prediction-Free, Real-Time Flexible Control of Tidal Lagoons through Proximal Policy Optimisation – A Case Study for the Swansea Lagoon, Ocean Engineering, v. 247, p. 110657, 2022.

- MOREIRA, TULIO MARCONDES; FARIA JR, JACKSON GERALDO; VAZ DE MELO, PEDRO O.S.; MEDEIROS-RIBERO G.. Development and Validation of an AI-Driven Model for the La Rance Tidal Barrage: A Generalisable Case Study, arXiv preprint arXiv:2202.05347 (In Review).

- https://doi.org/10.1016/j.simpa.2022.100356 MOREIRA, TULIO MARCONDES; VAZ DE MELO, PEDRO O.S.; MEDEIROS-RIBERO G.. Control Optimisation Baselines for Tidal Range Structures - CoBaseTRS, Software Impacts, 14, p.100356, 2022.

List of Figures

List of Tables

Acronyms

CH Constant head operational method.

CHV Constant head variant operational method, with independent sluice operation.

DC Direct Current.

$E.G$ Ebb generation stage.

$E.P$ Ebb pumping stage.

EHN Every half-tidal cycle and next operational method.

$EHNV$ Every half-tidal cycle and next variant operational method, with independent sluice operation.

EHT Every half-tidal cycle operational method.

$EHT-P$ Every half-tidal cycle classic operational method with pumping.

$EHTV$ Every half-tidal cycle variant operational method, with independent sluice operation.

$EHTV-\ P$ Every half-tidal cycle variant operational method with pumping and independent sluice operation.

ET Every tidal cycle operational method.

ETN Every tidal cycle and next operational method.

E_oG Ebb-only generation scheme.

$F.G$ Flood generation stage.

$F.P$ Flood pumping stage.

F_oG Flood-only generation scheme.

Ho Holding stage.

LR Inductor-Resistor circuit.

$NSSR$ Time-normalised sum of Squared Residuals.

SSR Sum of Squared Residuals.

Sl Sluicing stage.

T Turbine ON/OFF operation.

TW Two-way generation scheme.

TWP Two-way with pumping generation scheme.

0D 0-Dimensional tidal range structure simulation model.

1D 1-Dimensional tidal range structures simulation model.

2-D 2-Dimensional simulation game environment.

2D 2-Dimensional tidal range structures simulation model.

3-D 3-Dimensional simulation game environment.

3D 3-Dimensional tidal range structures simulation model.

ADCIRC ADvanced Multi-Dimensional CIRCulation Model for Shelves, Coasts and Estuaries.

AI Artificial Intelligence.

BODC British Oceanographic Data Centre.

DG-SWEM Discontinuous Galerkin Shallow Water Equations Model.

DRL Deep Reinforcement Learning.

EDF Electricité de France.

GWh Gigawatt-hour.

HC Head-Controlled.

LSTM Long Short-Term Memory.

MDP Markov Decision Process.

ML Machine Learning.

MLP Multilayer Perceptron.

MW Megawatt.

O&M Operation and Maintenance.

PDEs Partial differential equations.

PPO Proximal Policy Optimisation.

RL Reinforcement Learning.

RMS Root Mean Square differences.

SAC Soft Actor-Critic.

SBL Swansea Bay Tidal Lagoon Project.

SWE Shallow Water Equations.

TC Time-Controlled.

TRPO Trust Region Policy optimisation.

TRS Tidal Range Structures.

TWh Terawatt-hour.

List of Symbols

A_S Sluice/orifice area.

A_T Turbine area.

A_m Wave amplitude.

A_t Random variable with possible actions for the reinforcement learning agent.

Al_0 Lagoon impounded area at $z = 0$.

Al Lagoon impounded area.

C_E Other efficiency considerations for TRS turbines.

C_d Discharge coefficient.

C_f Correction factor.

C_{dt} Turbine discharge coefficient.

C Wave speed (celerity).

D Turbine or pump impeller diameter.

$E[G_t]$ Expected return of rewards.

E_{Pmax} Maximum energy that can theoretically be extracted by instantaneously emptying an enclosed body of water (from high tide to low tide).

E_f Turbine iso-efficiency curves.

E_{TW} Analytical upper bound for the available energy in a general TRS, under classic TW scheme operation.

$Eff_{E.G}$ Ebb generation turbine efficiency.

$Eff_{F.G}$ Flood generation turbine efficiency.

F_x Additional SWE forces in the x direction.

F_y Additional SWE forces in the y direction.

G_p Number of generating poles for a bulb turbine unit.

G_t Total discounted return.

$G_{(\tau)}^{future}$ Discounted future return of following trajectory τ.

$G_{(\tau)}$ Total undiscounted return, following a trajectory τ.

HS_{start} Operational water head for starting sluice operation.

H_p Pump head for HC methods.

H_t Total water depth.

H_{min} Operational water head for ending power generation.

H_{mt} Minimum required turbine operational water head.

H_{start} Operational water head for starting power generation.

Ho_{Tr} Holding trigger for HC methods.

H Time horizon.

I_{ss} Steady state electrical current.

$L_t^{CLIP}(\theta)$ Clipped surrogate loss.

$L_t^{CVS}(\theta)$ Final PPO objective loss function to be maximised.

L_t^{VF} Clipped, squared-error loss.

$L_{\theta old}^{IS}(\theta)$ Surrogate Loss.

L_{in} Hydraulic inertance.

L Lagoon water level, or electrical inductance.

M_2 Main semi-diurnal lunar tide.

N Pump rotation.

O_i Observations of the agent-environment interaction.

O Ocean water level.

P_6 Fixed pump power input $P_{in} = 6MW$.

P_C Turbine capacity.

P_o Pumping orientation.

P_{EG} Parametrised ebb generation power chart.

P_{FG} Parametrised flood generation power chart.

P_{Tr} Pumping trigger for HC methods.

P_{inMax} Maximum pump power input allowed.

P_{int} Interpolated (parametric) power output curves.

P_{in} Electrical power input, available to the pump or turbine in pump mode.

P_{out} Rate of work exerted from the pump to the fluid.

P_{ss} Steady-state power output estimate.

P Turbine power output.

Q_M Maximum pump flow rates, when $h_p = 0$.

Q_T Total directional flow rate from both turbines and sluices.

Q_g Gravitational flow rate estimate.

Q_p Pump flow rate, or turbine as pump flow rate.

Q_s Sluice flow rate.

Q_t Turbine flow rate.

Q_{11} Turbine specific unity discharge.

$Q_{\pi^*}(s_t, a_t)$ Action value function following an optimal policy π^*.

Q_{ss} Steady state flow rate.

Q Conduit flow rate.

$R(s_t, a_t)$ Reward from taking action a_t from state s_t.

R_h Hydraulic resistance.

R_t Random variable with possible rewards provided by the environment.

R Electrical resistance or pump power ratio P_{inMax}/P_{in}.

S_2 Semi-diurnal solar tide.

S_p Turbine rotational speed.

S_t Random variable with possible environmental states.

S_{op} Sluice operation.

Sl_{Tr} Sluice trigger for HC methods.

S Entropy bonus magnitude.

T_o Turbine orientation.

T_{Tr} Turbine power generation trigger for HC methods.

T_{op} Turbine operation.

T Wave period.

$U(\theta)$ Expected return of following a parameterised policy π_θ.

U Depth averaged flow velocity in the x direction.

V_0 Impounded water volume when $z = 0$.

$V_{\pi^*}(s_t)$ State value function following an optimal policy π^*.

V_π State value function following a policy π.

V Voltage, water volume or depth averaged flow velocity in the y direction.

ΔV Impounded lake water volume.

Δp Pressure drop observed in a conduit.

Δt Discretised time step.

Σ Co-variance matrix.

$\vec{\mu}$ Mean vector.

$\vec{q}(\vec{s}_t, a_i)$ Vector representation of action value function.

$\vec{s}_t$ Vector representation of input states at time t.

α Learning rate or turbine wicket gate opening.

β Hyperparameter controlling the strength of entropy regularisation (S) or turbine runner blades opening.

ϵ Hyperparameter for limiting large policy updates.

η_p Idealised pump efficiency.

η_{eff} Total turbine efficiency losses.

η Free surface elevation relative to the mean water level.

γ Discount factor.

$\hat{A}_t$ Advantage estimate.

$\hat{L}_P$ Predicted lagoon water levels, during pumping stage.

$\hat{L}_S$ Predicted lagoon water levels, during sluicing stage.

$\hat{L}$ Predicted lagoon water levels.

$\hat{P}$ Parametrised power chart ($P_{EG}(h)$, $P_{FG}(h)$) augmented with the momentum ramp function.

$\hat{g}^{IS}$ Expected return gradient with importance sampling.

λ_w Wavelength.

λ Regularisation hyperparameter.

∇_θ Gradient operator over $theta$.

ω Angular frequency.

ϕ Phase angle.

$\pi(a_t|s_t)$ Reinforcement learning agent's policy as a conditional probability distribution.

π^* Optimal reinforcement learning agent's policy.

π Reinforcement learning agent's policy.

ρ Seawater density ($= 1024kg/m^3$).

τ Sequence of sampled state-action pairs.

θ Neural network weights.

ξ Bottom drag friction coefficient.

ζ_s ζ for sluices.

ζ_{AE} ζ value during $E.G$, when turbine is accelerating.

ζ_{AF} ζ value during $F.G$, when turbine is accelerating.

ζ_{DE} ζ value during $E.G$, when turbine is decelerating.

ζ_{DF} ζ value during $F.G$, when turbine is decelerating.

ζ Ratio of hydraulic inertance L_{in} over resistance R_h.

a_i Sampled instance from A_i.

b Bathymetric depth.

c_1 Hyperparameter responsible for modulating L_t^{VF}, during $L_t^{CVS}(\theta)$ calculation.

f Grid frequency.

g Gravity constant.

h_p Negative water head surpassed during pumping.

h_s Pump shutoff head.

h Water head difference.

i Electrical current.

m Number of sampled trajectories from the “actor” neural network.

n_o Policy network node outputs.

n_t Number of turbines.

n_{11} Turbine unity speed.

$r(t)$ Heuristic sinusoidal ramp function.

$r_t(\theta)$ Probability ratio.

r_{i+1} Sampled instance from R_i.

s_i Sampled instance from S_i.

t_m Time when the current operation TRS operational mode was triggered.

t_r Transition ramp time for opening/closing hydraulic structures.

ub Maximum flow rate upper bound.

z Tidal datum, defined from the lowest tide level.

Contents

Chapter 1

Introduction

In recent years, concerns about the damaging effects of climate change combined with political and social pressures have pushed the world to increase the installed capacity of renewable energy sources (wind, solar, bioenergy and hydro) allowing renewable energy to account for 28% of global electricity energy generation in 2020 [1]. While significant progress has been made in expanding solar and wind resources, tidal energy remains practically untapped. Given that the worldwide potential for tidal range resource (considering reasonable thresholds for energy output and water depth) is around $25,880$ TWh [2], and that the world electric energy consumption in 2018 was $\approx 23,398$ TWh [3], the potential for tidal energy to become a significant contributor to the installed capacity of renewable energy sources in the world cannot be ignored.

In the current literature there are three main renewable energy schemes under study aimed at harnessing the energy from the tides: (i) dynamic tidal power, (ii) tidal stream turbines and (iii) Tidal Range Structures (TRS) [4; 5]. While the first two technologies are reasonably "new", TRS have been studied for nearly 90 years [6], with successful implementations in France, South Korea, Canada, Russia and China [4]. Beyond the advantage of construction experience and operation know-how, TRS can be designed to enable large-scale electricity generation on pair with hydroelectrics ($>$ 100 MW). As an example, the two most successful TRS projects built, namely, La Rance (France) and Lake Sihwa (South Korea), have 240 MW and 254 MW of installed capacity, respectively [2]. The theory behind TRS is to create an artificial impoundment between an enclosed body of water and the ocean. When a sufficient head is available, water is allowed to pass through turbines, generating energy in a similar fashion to hydroelectrics [7; 6]. The artificial enclosed body of water can be obtained either by impounding entire basins or by creating artificial lagoons in coastal regions, distinguishing TRS design into two types of projects: tidal barrages (e.g. La

Rance, Severn Barrage) and tidal lagoons (e.g. Swansea Bay Tidal Lagoon, Cardiff Lagoon). In contrast to tidal barrages, tidal lagoons permit the construction of smaller scale developments, with lower energy outputs but less environmental impacts [2].

A 2016 review by UK's ex-minister of energy [8] has helped in drawing attention to TRS as a competitive choice among renewables. In his report, the construction of "small-scale" tidal lagoons, such as the Swansea Bay Lagoon (SBL), is suggested as a pathfinder project before moving to larger-scale lagoons [4; 9; 10]. The report also emphasises that TRS have proposed lifetimes of operation of 120 years – far surpassing any other renewable energy type, allowing for very low electricity cost for years. As an example, La Rance, which is in operation for 55 years, took 20 years to amortise the initial investment, generating energy at competitive cost of nuclear or offshore wind sources [11; 8].

While La Rance is a significant example of how TRS can be economically viable given enough time of operation, novel projects, such as the SBL, have been questioned if they represent good value for money. This concern is exacerbated by the high initial investment costs that TRS require for construction ($\approx$ £1.3 billion for SBL) and the recent reduction of solar and wind payback periods to $\approx 2 - 20$ (year range) and ≈ 13 years, respectively (considering tax incentives and a maximum lifespan of 30 years) [12]. These current issues led the UK Government to halt support for the proposed SBL project in 2018 with the argument of high cost of electricity [13]. Nevertheless, in a recent turn of events, the Welsh Government launched a tidal lagoon challenge in Wales to ascertain market engagement in the construction of several TRS projects. A total of 55 companies registered for the challenge, with the winner possibly receiving Welsh Government financial support for designing a pathfinder TRS project [14; 15]. These initiatives demonstrate that TRS development in the near future is still very likely.

In order to estimate the potential energy of such systems, analytical or numerical models (0D through 3D), can be considered. While analytical solutions are useful in estimating a theoretical upper bound for the available power [7; 16], numerical models can aid in optimising the operational time-sequence of the hydraulic structures present in a TRS (e.g. turbine and sluices), while accounting effects not considered in analytical models (e.g. variable bathymetry, pump operational modes). Each of these approaches have different strengths, weaknesses and unique physical simplifications. For instance, while 0D models can be computationally inexpensive, they are solely based on the simple statement of mass conservation for the impounded lagoon. In contrast, 1D through 3D models utilise the shallow water equations to obtain (in an increasing order of detail and computational time), the water profile and velocity components in and out of the lagoon [17; 18; 19]. Shallow water models are helpful in assessing environmental impacts of TRS [20; 21], while also detecting the sudden high flow-rates developed when turbines start operating – a situation that can abruptly perturb water velocity and elevation (therefore power generation) in the vicinity of hydraulic structures [22]. Due to the characteristic features of the described approaches, 0D models are usually implemented into optimisation routines of TRS, since these often require a huge number of iterations for convergence [23; 24; 25]. As a next step, 1D to 3D models are employed to perform verification steps of the operational strategies devised from utilising the 0D model. 2D models are typically picked, among 1D and 3D options, given their satisfactory trade-off between accounting for detailed bathymetry, coastal geography and accurately predicting complex shallow water effects (e.g. funnelling, resonance) [26; 18], while not being computationally expensive as 3D models.

After choosing an appropriate location with high tidal range and low water depth, optimal TRS dispatch occurs in two fronts: operational and design optimisation. With a fixed design, operational optimisation focuses on finding optimal strategies for the sequential operation of the hydraulic structures [23; 25], with the goal of maximising energy extraction or meeting energy demands. In the design optimisation, the number and scale of turbines and sluices is optimised, considering a fixed impoundment perimeter [27; 2]. The design and operation optimisation can also be performed simultaneously, although with increased computational cost [28]. Both these approaches help increasing the utilisation factor (ratio of actual energy generated to installed capacity) and consequently the competitiveness of TRS against wind and solar alternatives.

Although the current literature has advanced in increasing theoretical power generation capabilities of TRS [2; 4; 29], there is room for improvement, considering that state-of-art 0D operational optimisation methods (i) cannot explore fine-tuned power input control of turbines in pump mode, (ii) can be computationally expensive, (iii) do not perform real-time control, only assuming an upper bound estimate for perfect forecast scenarios. In order to fill these gaps in the literature, this thesis proposes two general approaches that enable real-time control of 0D TRS models, surpassing current energy and revenue extraction capabilities of state-of-art methods. With approach (a), we improve the energy extraction capabilities of the current state-of-art operational method for operating TRS, by enabling independent sluice operation (named "variant" TRS operation) and propose modifications for allowing real-time TRS control (although requiring tidal predictions). With approach (b) we utilise Deep Reinforcement Learning (DRL) methods, more specifically Proximal Policy Optimisation (PPO) through the Unity Machine Learning (ML)-Agents package [30]. In a SBL (with pumping) case study, both (a) and (b) approaches manage to surpassed the upper bound energy estimate of state-of-art methods (scenarios where tidal predictions equal tidal measurements) by 2.61% and 3.14%, respectively (yielding a revenue gain up to 967, 267£ per year), while operating the TRS in real-time. However, the DRL (b) approach was shown to be superior from (a), since (i) it does not require tidal predictions, (ii) can perform fine tuned operation of turbines in pump mode and (iii) pump with positive head differences (aided by gravity). DRL was chosen among machine learning techniques due to the nature of TRS operational problem, which involves sequential decision-making of a reactive environment (lagoon water levels vary depending on the operation of hydraulic structures) with the goal of maximising expected return (energy), and also because a target optimal operation of the tidal lagoon is not known "a priori" – a requirement for supervised learning techniques.

The real-time TRS control approaches introduced in this thesis have been tested with 0D model representations. Therefore, revised 0D numerical models for TRS were also developed for this work. For the La Rance tidal barrage, a 0D model was created with novel representations for turbines (power generation and pump modes), wetted area and transition ramp functions (for switching between operational modes). These representations were obtained utilising parametrisation and affinity laws techniques applied to measured data from Electricité de France (EDF) – the utility company responsible for operating La Rance since its construction. To the best of our knowledge, this is the first 0D validated model of a constructed TRS made available to academics. A 0D model for the SBL was also created. While turbines (in power generation mode), sluices and wetted area follow techniques available from the literature, novel derivations for transition ramp functions and turbines in pump mode (obtained from La Rance 0D parametrisation) characterise our revised 0D model. This augmentation was required in order to showcase our (a) and (b) approaches ability to surpass the upper bound of state-of-art methods, in energy extraction.

A 2D barrage model was also developed in this work, by utilising the finite element tool DG-SWEM (Discontinuous Galerkin Shallow Water Equations Model), a discontinuous Galerkin version of ADCIRC (ADvanced Multi-Dimensional CIRCulation Model for Shelves, Coasts and Estuaries) [31; 32; 33]. For including the effects of sluices and turbines inside DG-SWEM, modifications were made to an existing and validated culvert boundary model [34; 35]. Agreement of results against 0D TRS models and applications of the developed 2D model for coastal protection (from storm surge events) can be seen in [20; 21; 36; 37; 38].

For reproducibility, all developed code implementations for this work are available in GitHub: `https://github.com/TmoreiraBR/PhD-Thesis-codes`.

1.1 Thesis Statement

The statement of this thesis is that: *Deep Reinforcement Learning and enhanced state-of-art methods for control optimisation, coupled with revised 0D models, can enable TRS to be more economically competitive.* We support this statement by showing that, in contrast to state-of-art methods, our techniques can (i) further increase energy extraction capabilities of TRS, (ii) perform real-time control and, exclusive to our DRL approach, perform (iii) prediction-free, automatic operation of TRS, reducing Operation and Maintenance (O&M) costs, (iv) fine tuned operation of turbines in pump mode and (v) pump with positive head differences (aided by gravity). Statements "iv" and "v" allow for control optimisation of TRS that is more realistic, i.e. more closely resembling observed data from constructed TRS, as explained in our La Rance case study (Section 5.2).

1.2 Problem and Motivation

The main problem investigated in this thesis is how to make Tidal Range Structures projects more economically viable. Our motivations can be summarised as follows:

- Help revivify a renewable energy type that has the potential of providing hundreds of gigawatts of clean energy worldwide, while offering coastal protection for storm surges and sea level rise.
- Help kick-start pathfinder projects, such as the Swansea Bay Tidal Lagoon, so that companies around the world have their interest renewed in funding TRS schemes.
- Develop novel approaches and improvements to the state-of-the-art algorithms for TRS control, increasing energy and revenue extraction capabilities of these systems.

1.3 Research Objective

Our objective with this research is to create general operational optimisation approaches for TRS that are automatic (real-time) and can surpass energy and revenue extraction capabilities of current state-of-art operational methods. With this final goal in mind, several gaps in the literature needed to be filled. Our achieved objectives can be summarised as:

- Propose improvements to the state-of-art control of TRS, further increasing energy extraction capabilities and allowing for real-time control.
- With DRL techniques, create the first automatic operation of TRS that does not rely on external models for tidal predictions and does not require re-runs, in contrast to state-of-art methods.
- Create and validate a 0D parametric model of a constructed tidal barrage (La Rance).
- From the parametric model of La Rance, create new generalizable models for pumps, so that TRS projects can further benefit from our DRL and enhanced state-of-art optimisation approaches.
- Enable energy extraction results superior to the upper bound of state-of-art methods through optimal control strategies derived from our (i) enhanced state-of-art and (ii) DRL approaches.
- Investigate state-of-art and enhanced state-of-art upper bound (and real-time) estimates, switching between energy and revenue maximisation targets.
- Provide comparison analysis between our developed real-time control optimisation methods and the current state-of-art.

1.4 Contributions

- To showcase the first automatic (real-time) method for optimising TRS operation, through either a (i) enhanced state-of-art approach or (ii) a novel Deep Reinforcement Learning (DRL) implementation.
- The novel DRL approach does not require future tidal predictions (concurrent with measured data) as inputs to optimally operate TRS in real-time and needs training a single time only, in contrast to state-of-art and enhanced state-of-art methods.
- Similarly to state-of-art methods, the DRL approach performs flexible operation of TRS, optimally adjusting its operational strategy according to the experienced tidal range.
- For the SBL case study, both the enhanced state-of-art and novel DRL approaches surpassed the upper bound energy estimate of state-of-art methods (idealised

perfect tidal forecast scenarios) by 2.61% and 3.14%, respectively, while operating the TRS in real-time, with pumping capabilities.

- For the SBL case study, both the enhanced state-of-art and novel DRL approaches surpassed the real-time yearly revenue estimate of state-of-art methods by $967,267£$ and $861,301£$, respectively, while operating the TRS in real-time, with pumping capabilities.
- The development and validation of a methodology for parametrising constructed TRS, utilising La Rance as a case study.
- The developed methodology for La Rance enabled novel and generalisable representations for ramp functions, turbines in pump mode and equivalent wetted area that are applicable to any TRS.
- The operational strategy developed by our DRL agent for the La Rance case study enabled fine-tuned pump power input and pumping with positive head differences. These characteristics are (i) unique to the DRL agent and (ii) more realistic than state-of-art (or enhanced state-of-art) approaches, i.e. comparable to the observed operation of real TRS that, due to intellectual property rights, have their details on control optimisation concealed from academia.

1.5 Thesis Outline

Given the multidisciplinary nature of this work, the next Chapter of this thesis (2) concerns itself on the fundamentals of: tides, tidal range structures and reinforcement learning subjects (along with a deeper description of Deep Reinforcement Learning, Proximal Policy Optimisation method and Unity ML-agents). Next, in Chapter 3, a literature review on the state-of-art operational schemes, components parametrisation and techniques for operational optimisation of TRS is presented, criticised and discussed. In Chapter 4 we propose improvements to the state-of-art operational optimisation of TRS, enabling increased power generation and real-time control. Also, novel methodologies for parametrisation of a real TRS are developed, with the tidal barrage of La Rance (France) being utilised as a generalisable case study. The 0D and 2D TRS models developed in this thesis are then presented, together with respective pseudocode implementations. Finally, Chapter 4 ends with a description of the Markov Decision Process (MDP) utilised for our DRL method, with a thorough description of rewards, actions, input states and developed environment components. A visual framework of the environment is also developed in Unity3D, within the Unity ML-Agents package. In Chapter 5, we apply our developed operational optimisation methods (enhanced state of art and DRL approach) to a tidal lagoon and tidal barrage TRS project types, namely: the Swansea Bay Tidal Lagoon pathfinder project and La Rance tidal barrage. Where applicable, our results are compared against the current state-of-art for operational optimisation of TRS. Finally, in Chapter 6, we draw Conclusions and Future Research Directions of this work.

Chapter 2

Fundamentals

2.1 Characteristics of Tides

The phenomenon of oceanic tides is a direct result from the difference between the centrifugal forces due to Earth's rotation and the combined gravitational attraction exerted by celestial bodies, mainly the Moon and the Sun. The major precursor of the study concerning tides is available in the work of Newton [39], where the attractive forces pulling the tides derive from Newton's law of universal gravitation.

2.1.1 Equilibrium Theory

From the Equilibrium Theory of Tides, Newton explained the $12h : 25min$ period of the main semi-diurnal lunar tide (M_2). His theory suggested that the opposing bulges of water, formed due to the attraction of the Moon, would have to cross the same location twice in a lunar day (moon-rise period of $24h : 50min$ for a same location on Earth). Similarly, the "bulges" concept is also used to describe the exact $12h$ principal semi-diurnal solar tide (S_2), developed as the Earth revolves around its axis while facing the Sun. The constructive interference (spring tides) happens when both Sun and Moon are aligned towards Earth, leading to the summation of both M_2 and S_2 tidal amplitudes. Intuitively, the destructive interference peaks in an opposite scenario (neap tides), when a 90° phase shift, relative to the Earth, occurs between the celestial bodies.

2.1.2 Dynamic Theory

In order to model the combined effect of different tidal constituents, while accounting for other physical interactions such as wave reflections and bed friction in estuaries, two mathematical tools helped advance the understanding of tides: Harmonic Analysis and Shallow Water Theory. Harmonic Analysis is a mathematical tool that captures the time evolution of combined tidal patterns. The analysis is based on the assumption that the periodic motion of each tidal constituent can separately be interpreted as a sine wave. For instance, by considering amplitude and frequencies of the two major tide constituents (S_2 and M_2), we can obtain an adequate description, in a first approximation, of the tidal phenomenon.

The harmonic decomposition of tides not only allows the main constituents responsible for explaining tidal patterns to be determined [40]; it also enables the analytical and numerical simulation of tides, since it provides a forcing boundary condition in the form of a time series approximation.

While Harmonic Analysis provides good approximations to describe changes in the oceanic surface at fixed locations, it does not explain well how the water waves dissipate, reflect and resonate. In 1775, Laplace [41] proposed a basic form of the Shallow Water Equations (SWE), which would become the foundation for the Dynamic/Real Theory of Tides. The coupled equations of hyperbolic partial differential equations (PDEs), derived from the Navier-Stokes mass and momentum conservation equations, can take friction and variable bathymetry as inputs, capturing coastline reactions to the forcing tides, including resonant behaviour. The SWE are cast in the "conservative form" by depth averaging velocities in "x" and "y" directions [20]. This process eliminates one dimension, thus providing a two-dimensional set of three equations: the first being the continuity equations Eq. 2.1, while Eq. 2.2 and Eq. 2.3 are the momentum equations in the "x" and "y" directions, respectively [32].

$$\frac{\partial \eta}{\partial t} + \frac{\partial (H_t U)}{\partial x} + \frac{\partial (H_t V)}{\partial y} = 0, \tag{2.1}$$

$$\frac{\partial (H_t U)}{\partial t} + \frac{\partial (H_t U^2 + g({H_t}^2 - b^2)/2)}{\partial x} + \frac{\partial (H_t U V)}{\partial y} = g\eta \frac{\partial b}{\partial x} - \xi U H_t + F_x, \tag{2.2}$$

$$\frac{\partial (H_t V)}{\partial t} + \frac{\partial (H_t U V)}{\partial x} + \frac{\partial (H_t V^2 + g({H_t}^2 - b^2)/2)}{\partial y} = g\eta \frac{\partial b}{\partial y} - \xi V H_t + F_y. \tag{2.3}$$

Here, η represents the free surface elevation relative to the mean water level, b the bathymetric depth and H_t the total water depth ($H_t = \eta + b$). From the depth averaging process, the plane velocities U and V (along x and y, respectively) are simplified to be constant throughout H_t, while g is the gravity constant, ξ the bottom drag friction (or bed friction) coefficient and F_x and F_y additional terms that can be incorporated into the equation (e.g. Coriolis force, air pressure variations). The conservative form (flux-form) is widely used in 2D numerical methods, since it allows for conservation of mass even when discontinuities are present in the problem (e.g. hydraulic jumps).

As its name implies, SWE equations have a theoretical basis that relies on the wavelength λ_w being several times larger than the bathymetric depth b [42]. A useful guide is that when

$$\frac{b}{\lambda_w} < 0.05 \tag{2.4}$$

is satisfied, then flow velocities can be described as purely horizontal (i.e. the use of depth averaged velocities in the model is reasonable).

2.1.3 Tidal Resonance

With the condition of shallow water waves being satisfied Eq. 2.4, analytical solutions show that the wave speed, or celerity C, depends solely on depth (i.e. waves are non-dispersive):

$$C = \sqrt{gb}. \tag{2.5}$$

This result is useful in explaining a very complex phenomenon observed in natural channels and estuaries on Earth: tidal resonance. In a simplified model (Fig. 2.1), an incoming wave from deeper water level (Wave I) is reflected by a vertical wall inside a constant depth shallow water channel.

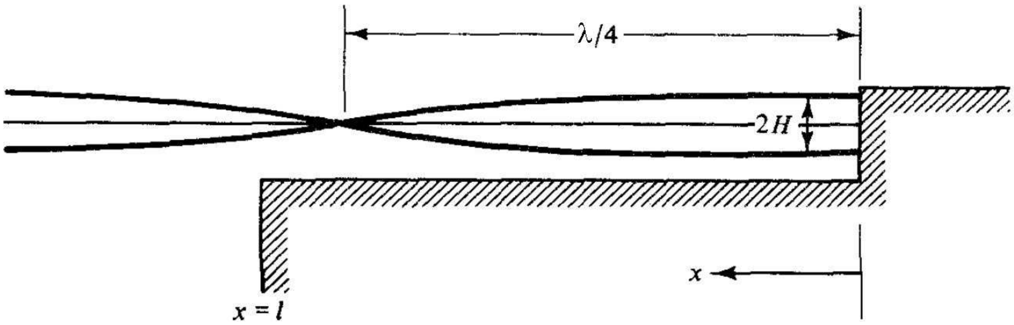

Figure 2.1. Superposition of incident and reflective wave yields a standing wave inside an idealised channel where $l \approx \lambda_w/4$. Edited from [42].

The wave reflection resembles a reflected wave from a rope with free end: in both cases a node is formed at a quarter wave-length distance from the vertical wall and the amplitude is increased at the wall (antinode) by the constructive interference between incident and reflected waves i.e. a standing wave is created. The difference, when dealing with the shallow water model, is that at the abrupt transition of bathymetries, a portion of the incoming wave will also be reflected (Wave II). This reflected wave is shifted 180° and moves to the "left", away from when $x = l$. For a constant depth, celerity is constant throughout the channel and the wavelength can be calculated as

$$\lambda_w = CT, \tag{2.6}$$

where T is the period. If the channel itself has a length which is a quarter of the wavelength from Wave I, it is possible to see that Wave I will take half the period to travel into and out of the channel. As it encounters Wave II, both waves will be in phase, and resonance will occur. In other words: the natural period of the channel matches the period of the forcing wave. The analytical solution to this problem gives an insight to understanding the great tidal ranges observed in natural resonating channels (e.g. Bay of Fundy, Canada and Bristol Channel, UK). According to the literature [43; 44], both of these channels are said to resonate primarily with the S_2 and M_2 tides. Using the M_2 as an example, if we consider a simplified model of a channel with constant depth of $40m$ [26], and knowing that the M_2 period is $12h : 25min$, one can obtain from Eq. (2.5, 2.6) a quarter resonant length of $221km$ for the M_2 tide. This value is roughly close to the approximate distance to the edge of the continental shelf for the Bristol Channel ($\approx 160km$) and Bay of Fundy ($\approx 290km$), indicating that the natural periods of these channels is close to the forcing period of the M_2 tide. As a matter of fact, according to numerical simulations from [26; 45], the natural period of the Bristol Channel and Bay of Fundy are close to $11h$ and $13h$, respectively, which reinforces the theory. It is important to note that the equations presented in this section are an extreme simplification of the actual phenomenon, given that no friction effects are considered, and that the variable bathymetry present in natural channels should interfere with the wave celerity in the direction of the shore. In order to account for these physical effects, 2D numerical models utilising the SWE should be considered.

2.2 Tidal Range Structures

Tidal barrages (and lagoons) are large-scale systems that utilise the same principle as hydroelectrics to generate power: hydraulic head. By allowing water to flow through sluice gates and turbines into an artificial impoundment, the high tide is confined until a operational head is established between the ocean and basin [7]. When this condition is achieved, water is allowed to pass through the turbines, thus generating power. The main difference between these types of impoundments is that tidal barrages are designed to enclose entire estuaries, while tidal lagoons impound significant smaller sub-areas inside the main portion of water [8]. The oscillatory motion of tides, combined with the minimum starting head required for an efficient operation of the turbines, characterises the predictable, but intermittent, power production of TRS.

From a design perspective, TRS are composed of four main components [6]:

- Turbines: hydrodynamic structures embedded into submerged water passages. Turbines are responsible for converting the potential energy from the water head difference across the barrage into kinetic, then rotational, and finally electric energy. For TRS, the most common type of turbine utilised are low-head bulb turbines, that can operate either to generate energy or as pumps.

- Sluices: designed to (i) control water levels inside the impounded lagoon during the construction phase and (ii) help increase the available head differences (i.e. potential energy) by reducing/increasing the impounded lagoon water levels at the end of turbine operation.

- Locks or similar apparatuses: designed to allow for safe passage of ships and boats across the impoundment.

- Embankments or concrete caissons: compose the majority of the perimeter of the TRS. These have to be as impermeable as possible, in order to ensure water head differences are as high as possible.

These main components can be seen in constructed TRS. However, since the debut of La Rance Tidal Barrage in 1966 (with $240 MW$ of installed capacity), few other TRS with power production on the order of MW have been built (Table 2.1).

Table 2.1. Characteristics of constructed TRS (edited from [2]).

TRS	Year	Capacity (MW)	Max. Basin Area (km^2)	Operational Scheme[b]
La Rance, France	1966	240	22	TWP
Kislaya Guba, Russia	1968	1.7	2	TW
ARGS[a], Canada	1984	20	6	E_oG
Jiangxia, China	1980	3.9	2	TW
Lake Sihwa, Korea	2011	254	30	F_oG

[a] Annapolis Royal Generating Station; [b] Defined in Section 3.1.

This contrasts with the annual energy potential of feasible sites around the world (Table 2.2), demonstrating how TRS technology is being underutilised.

Table 2.2. Annual tidal range potential energy per country (from [2]).

Country	Annual Potential Energy (TWh)
Canada (Hudson) (extensive sea ice)	20110
Australia	1760
Canada (Fundy)	1357
UK	734
France	732
US (Alaska) (partial sea ice)	619
Brazil	298
South Korea	107
Argentina	62
Russia	75
India	19
China	12

Since this work focuses on increasing the energy extraction capabilities of TRS projects through the smart operation of turbines and sluices, the construction steps and strategies for finding best suited locations for TRS deployment will not be discussed. In Section 3.1, state-of-art operational schemes and techniques for turbine and sluice parametrisation and simulation from the literature will be presented.

2.3 Reinforcement Learning

As shown in the work of Sutton and Barto [46], a reinforcement learning (RL) problem can be mathematically formalised as a Markov Decision Process (MDP). In an MDP, an agent interacts with an environment through actions (A_t), and these actions lead to new environmental states (S_{t+1}) and possible rewards (R_{t+1}) for the agent. The quantities A_t, S_t and R_t are random variables, with well-defined probability distributions. A general agent-environment MDP framework can be visualised in Fig. 2.2.

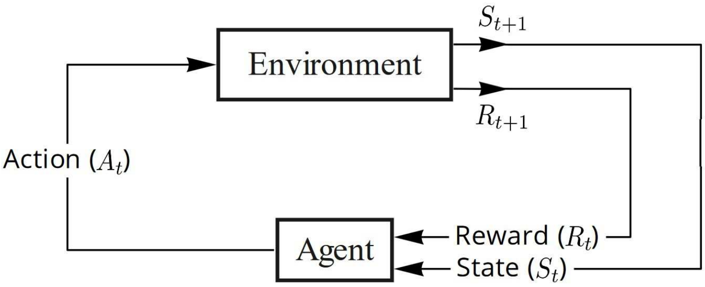

Figure 2.2. Agent-environment interaction in an MDP, illustrating the state-action, next reward triples sequence. Adapted from [46].

By sampling multiple time-steps $t = 0, 1, 2, 3...$, observations (O_i) of the agent-environment interaction are organised as a sequence of state-action, next reward triples:

$$O_i =< s_i, a_i, r_{i+1} >, \tag{2.7}$$

where s_i, a_i and r_{i+1} are sampled instances of the random variables (S_t, A_t and R_t). The sequence of state-action pairs defines a trajectory τ:

$$\tau = s_0, a_0, s_1, a_1, s_2, a_2... \tag{2.8}$$

Also, in an MDP, we can say that the probabilities of R_{t+1} and S_{t+1} are completely conditioned only on the preceding state and action (S_t and A_t), that is:

$$p(s_{t+1}, r_{t+1}|s_t, a_t). \tag{2.9}$$

The probability distribution Eq. 2.9 defines the dynamics of the MDP. It can also be manipulated to yield the state-action-transition probability distribution (which is just the sum of probabilities over all possible future rewards):

$$p(s_{t+1}|s_t, a_t) = \sum_{r \in R_{t+1}} p(s_{t+1}, r_{t+1}|s_t, a_t). \tag{2.10}$$

For estimating Eq. 2.10 for a given state, we also need to condition an action. In non-deterministic scenarios, the selection of possible actions by the agent is a stochastic process, defined by a conditional probability distribution (known as policy) of the form:

$$\pi(a_t|s_t). \tag{2.11}$$

Using Eq. 2.10,2.11, the probability distribution of starting in a state s_t and ending in s_{t+1}, given a policy, can be estimated as:

$$p_\pi(s_{t+1}|s_t) = \sum_{a \in A_t} \pi(a_t|s_t)p(s_{t+1}|s_t, a_t). \tag{2.12}$$

For a single path i:

$$p_\pi(s_{t+1}^{(i)}|s_t^{(i)}) = \pi(a_t^{(i)}|s_t^{(i)})p(s_{t+1}^{(i)}|s_t^{(i)}, a_t^{(i)}). \tag{2.13}$$

Also, with some defined policy, we can sum the observed rewards for each state-action pair (as shown in Eq. 2.7 and calculate a total discounted return G_t for a given policy, at time-step t:

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} ... = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}, \tag{2.14}$$

where γ is a discount factor between 0 and 1, that regulates the relative importance of immediate rewards to the learning agent.

The objective of reinforcement learning problems is to find an optimal policy π^*, that maximises the expected return of rewards $E[G_t]$ conditioned on any initial state s_t, i.e.

$$\pi^* = \arg\max_{\pi} E_{\pi}[G_t | S_t = s_t], \forall s_t. \tag{2.15}$$

Finally, the expected reward, following the optimal policy, can be conditioned on a state or state-action pair, yielding state value and action value functions, respectively:

$$V_{\pi^*}(s_t) = E_{\pi^*}[G_t | s_t]. \tag{2.16}$$

$$Q_{\pi^*}(s_t, a_t) = E_{\pi^*}[G_t | s_t, a_t]. \tag{2.17}$$

There are several techniques for modelling the the environment and agent components of an MDP. The environment component can be represented either by a probability dynamics model (Eq. 2.10) or a virtual simulator (as is the case with this work), that return next state s_{t+1} and reward r_{t+1} for some previous state and action s_t and a_t. When the probability dynamics model is available, the reinforcement learning problem is framed as model-based. Conversely, when using a simulator, the problem is framed as model-free. In both cases, and considering sampled actions from a given policy $\pi(a_t|s_t)$, the environment is sampled multiple times, so that multiple observations (Eq. 2.7) are used during training.

Similarly to the environment definition, there are also many approaches that can be used for modelling the agent in an MDP. These approaches are related to the chosen algorithm (value based vs policy based) used for finding the optimal policy π^* (i.e. the agent). In value based reinforcement learning, the goal is to obtain a representation for the optimal action (or state) value function (Eq. 2.16 and 2.17). Then, the optimal policy is obtained indirectly, i.e., given a state s_i, the optimal action (a_j) to take is the one that maximises the expected return. As an example for value-based methods, if the problem consists of a discrete number of states and actions (e.g s_i and a_j), a Q-table with i rows and j columns can be build. Each element in a Q-table corresponds to the expected return of choosing the j action in state i. The expected return estimate for each cell can be calculated through several pseudo-code algorithmic implementations, such as:

- First-visit Monte Carlo Prediction.
- Every-visit Monte Carlo Prediction.
- Monte Carlo Control (ϵ - greedy).

Given enough episodes, any of the methods above are guaranteed to converge to an optimal Q-table, which is the matrix form of the action value function $Q_{\pi^*}(s_t, a_t)$.

Differently from value-based methods, policy-based methods parametrise the policy π^* directly during training, with state-of-the-art approaches utilising neural network implementations. A full derivation of a state-of-the-art policy-based method is shown in Section 2.4.1, when discussing the Proximal Policy Optimisation (PPO) algorithm. In this work, PPO was chosen for solving all of our reinforcement learning problems.

2.4 Deep Reinforcement Learning

Solutions to reinforcement learning problems can be augmented with the use of neural network representations. For Deep Q-Network methods [47], for instance, the neural network is a parametric representation of the Q-table discussed in Section 2.3, receiving a vector $\vec{s}_t$ with the state representation and outputting a vector with the expected return of taking one of each possible action $\vec{q}(\vec{s}_t, a_i)$. The idea of utilising neural network representations for reinforcement learning problems enable solutions to much more complex state and action spaces, due to the fact that neural networks are universal function approximators and can parametrise policies, state-value and action value functions.

2.4.1 Proximal Policy Optimisation

Once the reinforcement learning problem is formalised as an MDP, several algorithms can be used for finding an optimal control policy π^*. In this work, the process of finding π^* has been achieved through Proximal Policy Optimisation (PPO) [48], built in the Unity ML-Agents package. PPO was shown to outperform several other “on-policy” gradient methods [48] and is one of the preferred methods for control optimisation when the cost of acquiring new data is low [49]. Furthermore, a revised and up-to date PPO algorithm is already implemented in Unity ML-Agents. Nevertheless, for sake of completeness, we present in this section the mathematical derivation of the PPO method. The designing of the TRS operation as an MDP is shown in Section 4.4.

Differently from approaches that try to infer the policy through state-value or action-value functions (e.g. Deep Q-Network) [47], PPO uses an “on-policy” approach that maximises the expected sum of rewards by improving its current policy – smoothly shifting the probability density function estimate of the policy towards π^*.

The PPO algorithm is an updated form of Policy Gradients. As TRPO (Trust Region Policy optimisation) [50], it tries to increase sample efficiency (re-using data from previous policies), while constraining gradient steps to a trust region. It is also actor-critic, since it utilises an estimate of the state-value function for its baseline [51]. An overview of Policy Gradients and PPO is presented below.

2.4.1.1 Policy Gradients

Policy gradient methods rely on the fact that a stochastic policy can be parameterised by an “actor” neural network with weights $\vec{\theta}$ (simplified as θ going forward). As represented in Fig. 2.3, this Multilayer Perceptron (MLP) neural network receives a vector state representation of s_t. For the case of discrete actions, the neural network outputs the probabilities of each possible action in that state using a softmax layer. For continuous actions, each node in the last layer outputs the moments of a multivariate Gaussian distribution of the form [52; 53]:

$$\pi_\theta(a|s) = N(\vec{\mu}(s;\theta), \Sigma(s;\theta)), \tag{2.18}$$

where $\vec{\mu}$ and Σ are the parameterised mean vector (for each possible action) and covariance matrices, respectively. While training, actions are randomly sampled from the distribution. In the beginning, the standard deviation of all possible actions (covariance matrix's diagonal) is large, favouring exploration. As training progresses, the neural network weights are updated to favour optimal actions, i.e. the standard deviation for each action reduces, favouring exploitation. During testing the performed actions can either stochastic (directly sampled from the trained distribution) or deterministic (as is the case with this work), where $\vec{\mu}$ is taken as the joint optimum action for every input state s_t.

Considering a trajectory τ, the expected return of following a parameterised policy π_θ is $U(\theta) = E[G_{(\tau)}; \pi_\theta]$, where

$$G_{(\tau)} = \sum_{t=0}^{H} R(s_t, a_t), \tag{2.19}$$

and $R(s_t, a_t)$ is the reward from taking action a_t from state s_t. We also note that $G_{(\tau)}$

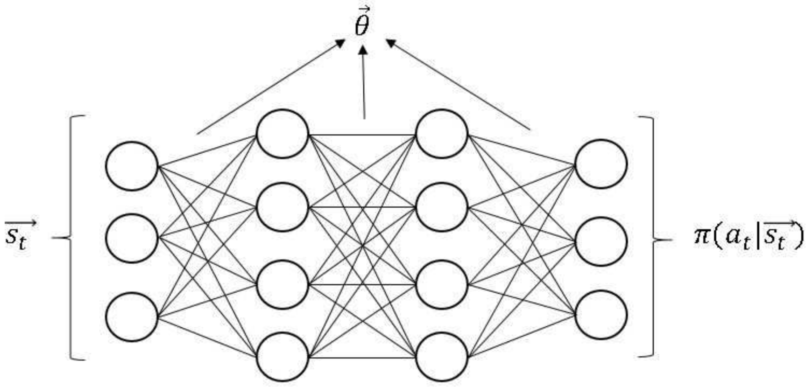

Figure 2.3. Input-output representation of policy (actor) neural network.

represents the undiscounted return following a sampled trajectory τ for a time horizon H. With these considerations, finding an optimal policy can be viewed as tuning θ to maximise $U(\theta)$, i.e. to perform gradient ascent of θ:

$$\theta \leftarrow \theta + \alpha \nabla_\theta U(\theta), \tag{2.20}$$

where α is the learning rate hyperparameter.

A sample based estimate for $\nabla_\theta U(\theta)$ assumes the form:

$$\nabla_\theta U(\theta) \approx \frac{1}{m} \sum_{i=1}^{m} \sum_{t=0}^{H} \nabla_\theta log(\pi_\theta(a_t^{(i)}|s_t^{(i)}))\hat{A}_t, \tag{2.21}$$

where m is the number of sampled trajectories from the "actor" neural network. For vanilla policy gradient methods, following the i^{th} trajectory $\tau^{(i)}$, we get

$$\hat{A}_t = G_{(\tau)}^{future} - V_\pi(s_t^{(i)}), \quad \text{and} \tag{2.22}$$

$$G_{(\tau)}^{future} = \sum_{k=t}^{H} \gamma^{(k-t)} R(s_k^{(i)}, a_k^{(i)}). \tag{2.23}$$

$V_\pi(s_t^{(i)})$, also parameterised by a "critic" MLP neural network, is the estimate for the value function of being in state s_t (starting point of the i^{th} trajectory) and following policy π thereafter; $G_{(\tau)}^{future}$ is the discounted future return of following the i^{th} trajectory, from time $k = t$; and $\hat{A}_t$ is the advantage estimate of taking this trajectory in respect to the current estimate of $V_\pi(s_t^{(i)})$. A complete derivation of $\nabla_\theta U(\theta)$ can be seen in [54].

2.4.1.2 Clipped Surrogate Loss derivation for PPO

In order to increase sampling efficiency [55], importance sampling can be used to rewrite the gradient term in Eq. 2.21 as:

$$\hat{g}^{IS} = \hat{E}_t \left[\nabla_\theta \frac{\pi_\theta(a|s)}{\pi_{\theta old}(a|s)} \hat{A}_t \right]. \tag{2.24}$$

Eq. 2.24 importance sampling form allows for re-utilising samples from an older policy ($\pi_{\theta old}$) to perform gradient ascent steps, when refining a new policy. It is obtained when differentiating the Surrogate Loss:

$$L^{IS}_{\theta old}(\theta) = \hat{E}_t[r_t(\theta)\hat{A}_t], \tag{2.25}$$

where $r_t(\theta) = \pi_\theta(a|s)/\pi_{\theta old}(a|s)$ is a probability ratio.

While in TRPO [50] the maximisation of the surrogate loss from Eq. 2.25 is subjected to a Kullback–Leibler divergence constraint, in PPO [48] the surrogate loss is constrained through a clipping procedure, yielding the clipped surrogate loss objective:

$$L^{CLIP}_t(\theta) = \hat{E}_t[min(r_t(\theta)\hat{A}_t, clip(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t)], \tag{2.26}$$

where ϵ is a hyperparameter that limits large policy updates, clipping the surrogate loss when large probability ratios are observed, i.e., situations where $|r_t(\theta)| > 1 + \epsilon$.

To further reduce variance when estimating the advantage, [48; 51] utilise a truncated version of generalised advantage estimation [56], where $\hat{A}_t$ is estimated as

$$\hat{A}_t = \delta_t + (\gamma\lambda)\delta_{t+1} + ... + (\gamma\lambda)^{H-t+1}\delta_{H-1}, \tag{2.27}$$

t is a time index within the sampled trajectory time horizon $[0, H]$, λ is a regularisation hyperparameter that performs the exponential weighted average of k-step estimators of the returns [56], and $\delta_t = r_t + \gamma V_\pi(s_{t+1}) - V_\pi(s_t)$.

When utilising shared parameters for the "actor" and "critic" neural networks (as is the case with this work), the loss function needs to be augmented with a value function error term [48]. To ensure exploration, an entropy term "S" is also added. Finally, the loss function to be maximised at each iteration becomes:

$$L^{CVS}_t(\theta) = \hat{E}_t[L^{CLIP}_t(\theta) - c_1 L^{VF}_t(\theta) + \beta S\pi_\theta(s_t)]. \tag{2.28}$$

For this study, Unity ML-Agents package fixes $c_1 = 0.5$ [57], β is a hyperparameter controlling the entropy bonus magnitude S, and L^{VF}_t is a clipped, squared-error loss between the estimate of the state-value function $V_\pi(s_t;\theta)$ and the actual return value obtained when following a trajectory τ [48]. The implementation of L_t^{VF} in Unity ML-Agents is seen in [57].

Additionally, parallel training can also be implemented as a way of substituting experience replay by running the same policy neural network on multiple instances of the environment. By guaranteeing that each environment starts in a random initial state during training, this parallelism helps decorrelate the sampled data, stabilising learning [58].

2.4.2 Unity ML-Agents

The Unity3D graphics engine is a popular game developing environment that has been used to create games and simulations in 2-D and 3-D since its debut in 2005. It has received widespread adoption in other areas as well, such as architecture, engineering and construction [30].

Unity ML-Agents is an open-source project that allows for designing 2-D and 3-D MDP environments where a DRL agent can learn through interactions [30; 57]. As a platform for creating learning environments for smart agents, Unity ML-Agents has grown significantly over the years, competing with other simulation environments utilised for training Artificial Intelligence (AI) agents [59; 60; 61]. The learning environments can be designed in Unity with various scales of complexity, considering complex physical simulation if required (rigid body, soft body, particle, fluid dynamics and ragdoll physics [30]). For solving the developed environments, single-agent, multi-agent cooperative, and multi-agent competitive scenarios via several Deep Reinforcement Learning algorithms are available: Soft Actor-Critic (SAC) [62], Proximal Policy Optimisation (PPO) [48], MultiAgent POsthumous Credit Assignment (MA-POCA) [63] and self play [64]. Examples of learning environments developed in Unity ML-Agents are shown in Fig. 2.4, taken from [65].

Fig. 2.4a, showcases an environment where the agent has to learn to push a block in order to jump over a wall, while Fig. 2.4b presents multiple agents learning how to play soccer in a cooperative/competitive setting. Finally Fig. 2.4c showcases the environment example of an agent learning to walk.

The significant differences between the learning tasks presented in Fig. 2.4 showcase the general capabilities of Unity ML-Agents of enabling solutions for various reinforcement learning tasks. Moreover, developing the agent-environment MDP in Unity3D allows for having access to a visual framework representation of the reinforcement learning problem, allowing for easy debugging and assessing the behaviour

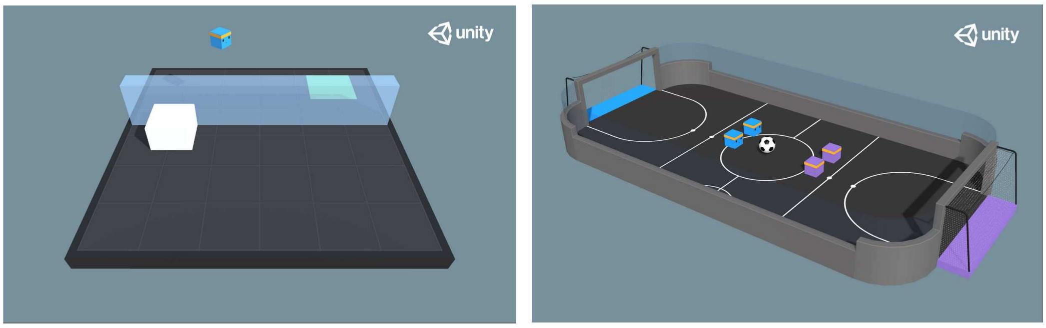

(a) Jumper agent. **(b)** Soccer player agents.

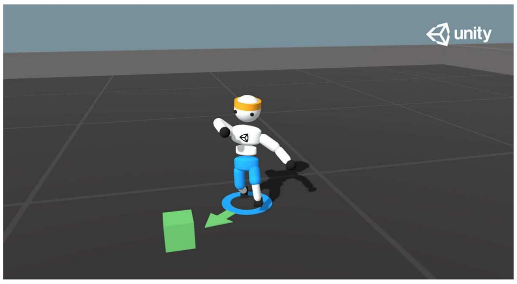

(c) Walker Agent.

Figure 2.4. Visual representations from learning environments (from [65]).

(policy) of the agent, during and after training.

Chapter 3

State-of-Art Literature Review

3.1 TRS Operation Schemes

3.1.1 Without Pumping

TRS operation can follow three main schemes: ebb-only generation (E_oG), flood-only generation (F_oG) or a two-way generation (TW). While for E_oG or F_oG, power generation only occurs during the ebb or flood tides, respectively, TW schemes generate energy for both flood and ebb tides.

As reported by Baker [6], when comparing one way E_oG or F_oG schemes, the former has usually been set as the preferred option for TRS for its ability to generate more energy. For justifying this claim, ocean and lagoon water levels for E_oG and F_oG schemes, together with associated periods of power production, are represented in Figs. 3.1a and 3.1b. From these images, we note how lagoon (basin) water levels are lower for F_oG than for E_oG. Considering how the wetted surface area of basins tend to naturally decrease with water depth (e.g. Swansea Bay Lagoon wetted area in Fig. 3.2), the lower mean water level during F_oG means that the barrage will operate with a smaller flooded area. Therefore, smaller volumes of water will flow through the turbines, lowering energy production per tide.

In Fig. 3.1, the operational stages: holding, power generation (E_oG or F_oG) and sluicing (or filling), are highlighted. During sluicing stage, both sluices and idling turbines allow water to flow into the impounded area so that the water level inside the lagoon gets as close as possible to the original tidal range. During holding time, turbine wicket gates and sluices are closed, allowing head difference to increase until a sufficient operational head H_{start} is established. Then, during power generation, flow is allowed to pass through turbines only.

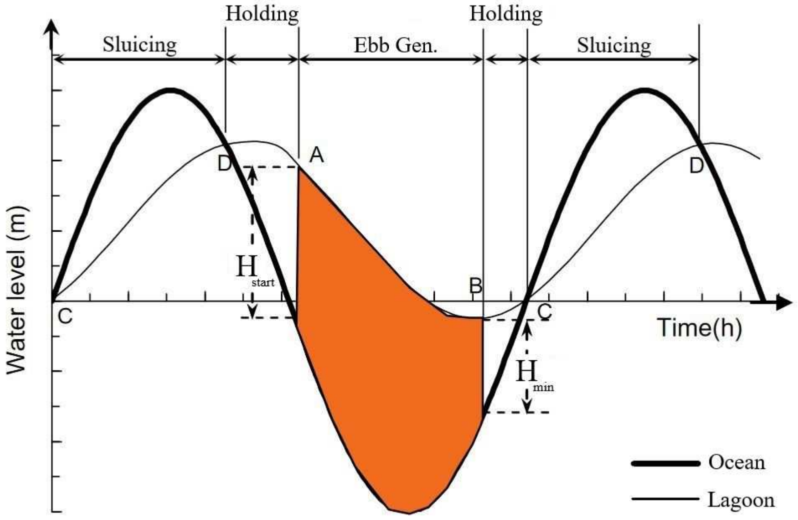

(a) E_oG scheme operation.

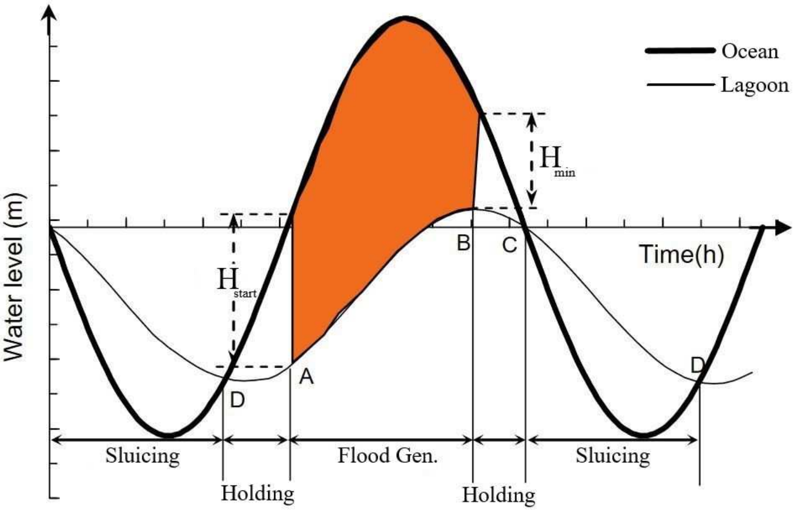

(b) F_oG scheme operation.

Figure 3.1. Representation of E_oG and F_oG schemes with starting (H_{start}) and finishing heads (H_{min}) (edited from [66]).

Although TW schemes are theoretically able to generate more energy than E_oG and F_oG [7], this option has been debated in the literature due to inefficiencies in reverse flow and high costs of more complex turbines [6]. In spite of this, successful TW operation in La Rance [67] and new advances in turbine technology [4] have allowed TW schemes to become the primary choice for most modern TRS proposals.

Considering a scenario starting at high tide, TW schemes can be described sequentially:

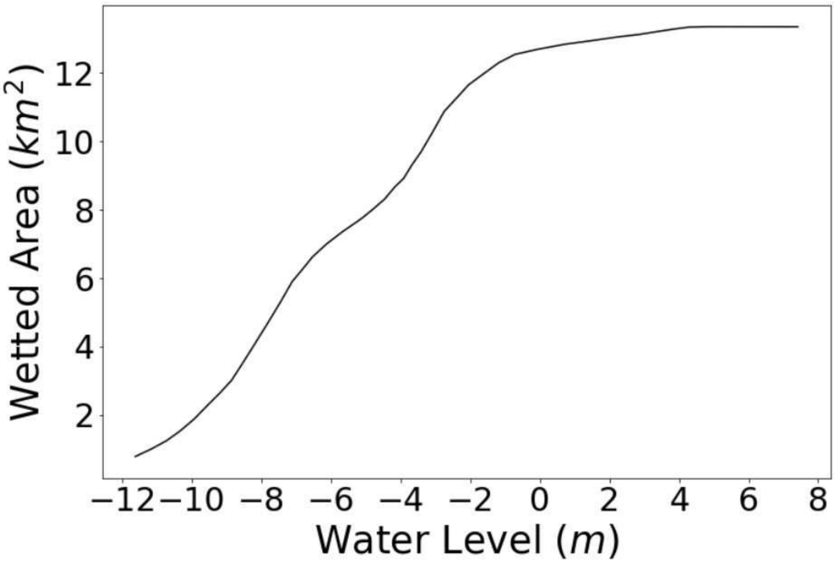

Figure 3.2. Swansea Bay TRS approximation for lagoon wetted area as a function of water level (digitised from [25]).

- 1) The incoming tide (flood tide) is confined within the lagoon at high level (holding stage).
- 2) Then, during the receding tide (ebb tide), power generation begins when a high operational head (H_{start}) is established between the basin and ocean [7].
- 3) Power generation stops when a minimum operational head (H_{min}) is achieved.
- 4) A sluicing sequence immediately follows, where idling turbines and sluices allow water to flow in order to increase lagoon tidal range for the next operation.
- 5) Following the same procedure, generating energy is also possible during the flood tide, although with reduced efficiency due to turbines usually being ebb-oriented [23].

The described operational modes for a classic TW operation are shown in Fig. 3.3 and detailed in Table 3.1.

Considering a TW scheme scenario, TRS turbines can be operated to either generate energy or to increase flow rates through the barrage during sluicing stage (idle operation of turbines) at either incoming or receding tides. Also, a minimum head H_{mt}, usually in the $[1m - 2m]$ range [68], is required for the turbine to generate energy. Considering that the holding stage begins automatically when the difference between ocean and lagoon is negligible, and that power generation is not possible with

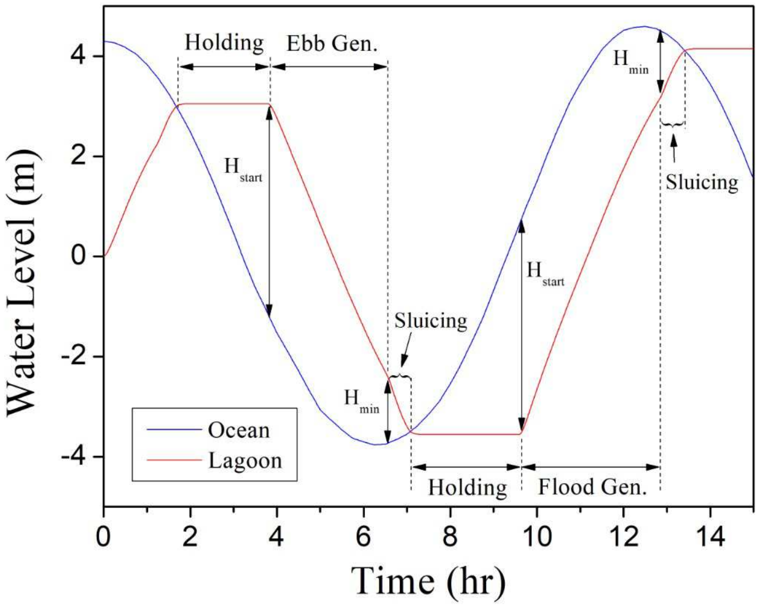

Figure 3.3. Classic “two-way scheme” operation. Ocean level is represented by the blue line, while the lagoon level is shown in red.

Table 3.1. TRS control stages.

Operational Mode	Description
Ebb Gen ($E.G$):	Power generation during receding tide
Flood Gen ($F.G$):	Power generation during incoming tide
Sluicing (Sl):	Operate sluice gates and/or idle turbines
Holding (Ho):	Stop operation of all hydraulic structures

head differences below H_{mt}, the classic operation of tidal lagoons [7] is reduced to two variables: H_{start} and H_{min}. As seen in Fig. 3.3, pairs H_{start} and H_{min} occur every half-tide period, when ocean oscillates between its valleys and peaks.

The water level variations within the lagoon, following classic TW scheme operation of hydraulic structures, can be seen in Fig. 3.3. Furthermore, Table 3.2 show all possible combined operations of turbines and sluices, with resulting power generation, for each control stage.

A modification of the discussed classical operation, called “variant” operation, allows for opening the sluice gates at the end of “flood” and “ebb” generation stages [6; 23] independently of turbine operation, with the possibility of increasing power generation (increased lagoon tidal range when starting the next “ebb” or “flood” stages).

Table 3.2. Classic TW scheme's hydraulic structures operation.

Operational Mode	Turbines	Sluices	Power Gen.
$E.G$ (if $T_{Tr} == True$):	On	Off	Yes (if $h > H_{min}$)
$F.G$ (if $T_{Tr} == True$):	On	Off	Yes (if $h > H_{min}$)
Sl (if $Sl_{Tr} == True$):	On	On	No
Ho (if $Ho_{Tr} == True$):	Off	Off	No

h: Head difference between ocean and lagoon
T_{Tr}: Turbine power generation trigger ($h > H_{start}$)
Sl_{Tr}: Sluice trigger ($h < H_{min}$)
Ho_{Tr}: Holding trigger ($h == 0$)

So far, in the academic literature, variant TW operation has only been possible through Time-Controlled (TC) operation methods (Section 3.3.3.2). A Head-Controlled (HC) method for performing the variant operation is introduced in this work in Section 4.1.1.

Since TW schemes are the preferred options when the goal is to maximise energy capabilities of TRS, E_oG e F_oG options are not considered when developing the control optimisation algorithms of this study. Therefore, TW classic and variant approaches to operate TRS are used in the optimisation routines of our baselines in Section 5.1.1. In the next section, classic TW schemes augmented with pumping will be presented.

3.1.2 Two-Way With Pumping (TWP)

As discussed in the seminal work of Gibrat from 1955 [69] – prime chief investigator of La Rance Tidal Barrage project –, the augmentation of TRS operation with pumping capabilities can increase energy (and revenue) extraction of these systems significantly. Indeed, state-of-art research [23; 25] has supported this claim, by predicting a $20 - 40\%$ increase of energy output when implementing pumping in classic TW scheme approaches. However, as discussed in the sequence, the literature interpretation of two-way with pumping schemes (TWP) is still not aligned with observed measurements of the only case study in the world: La Rance (Fig. 3.4). Firstly, in Fig. 3.5a, predicted lagoon water level variations for state-of-art literature interpretation of TWP schemes are coloured according to the operational mode chosen for turbines and sluices, with the sequence of operations for the hydraulic structures being dictated by literature constraints [23; 25]. While the logic of operation for turbines and sluices is similar to the observed in La Rance (Fig. 3.4), literature constraints limit turbine pumping stage to occur only when negative head differences (against gravity) are observed (red regimes in Fig. 3.5). In contrast, measurements from La Rance [67; 70], showcase pump operation in both positive and negative water head scenarios (power input stages in

Fig. 3.4), with pump shutoff negative heads (null pump flow rate) up to 6 m. Secondly, as shown in Fig. 3.5b, state-of-art optimisation routines need to choose to fix either power input or pump flow-rate to reduce computational costs [71] – differently from what is observed in La Rance (Fig. 3.4), where fine tuned power input is dynamically adjusted for turbines during pump mode. Finally, HC methods for TWP schemes do not consider the independent operation of sluices (variant TRS operation), also contemplated in La Rance [67; 6].

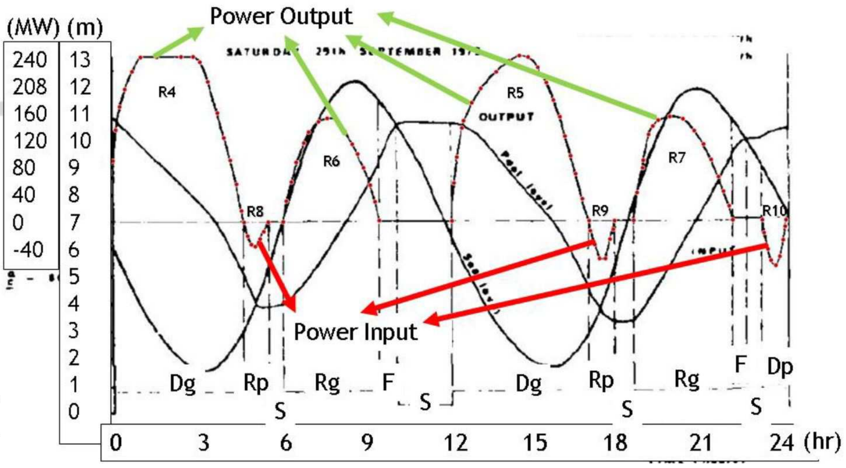

Figure 3.4. Measured water levels and power outputs of La Rance during a TWP scheme operation (edited from [67]). The labelled regions read as follow: Rp: reverse pumping (or ebb pumping); Dp: direct pumping (or flood pumping); Dg: direct generation (or ebb generation); Rg: reverse generation (or flood generation); S: standing (or holding); F: filling (or sluicing).

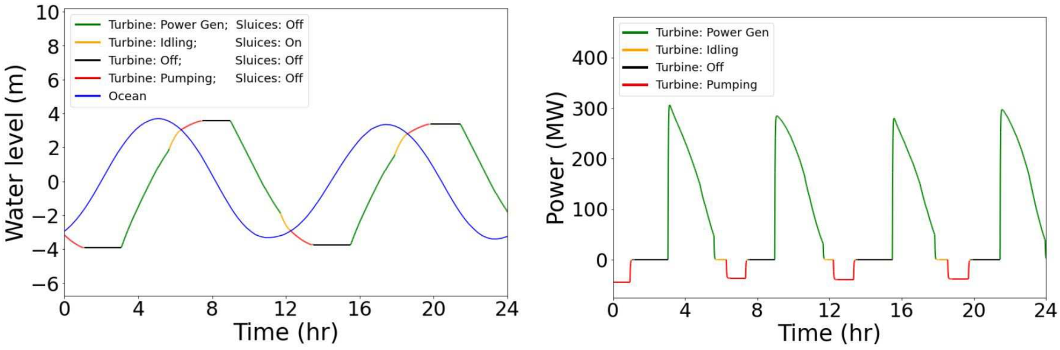

(a) Coloured lagoon water levels.

(b) Power generation.

Figure 3.5. Classic TWP HC operation of TRS, following state-of-art constraints by [25]. Ocean is coloured in blue, while lagoon water levels are shown and coloured following turbines and sluices operational modes, with green representing power generation mode, orange – idling/sluicing mode, black – offline mode and red – pumping mode.

A comparison between measured and literature interpretations for TWP scheme of operation can be see in Figs. 3.4, 3.5, where the "*Rp*" and "*Dp*" labelled regions (Fig. 3.4) correspond to reverse pumping (or ebb pumping) and direct pumping (or flood pumping) stages. The measured data from Fig. 3.4 are thoroughly detailed (and re-labelled) in Section 4.

In Section 5.2, we show that the operational strategy devised by our trained DRL agent, for a 0D parametric representation of La Rance, is capable of filling these gaps in the literature, enabling TRS operation in a more realistic way [72]. Furthermore, in Section 4.1.1 we propose improvements to the HC state-of-art approach to operate TRS, through the independent operation of sluices with pumping capabilities, leading to significant increases in power generation and revenue (Section 5.3).

3.2 TRS Components – Literature Parametrisation

For enabling the simulation of hydraulic units, 0D to 2D TRS models utilise parametric equations for predicting flow rates (and power output, when applicable) from turbines and sluices. While for sluices the orifice equation (derived from mass conservation) has usually been the common choice, turbine simulation has made use of experimental charts made available by few bulb turbine manufactures (e.g. Sulzer Escher Wyss of Zurich, Andritz) [6; 27]

Beyond turbine and sluice parametrisation, appropriate descriptions for the wetted area as a function of water depth are required in order to predict lagoon water level variations during numerical simulations. Finally, heuristic ramp functions are also usually implemented in numerical simulations when switching operational modes of TRS, simulating the time required for starting/shutting off the hydraulic structures, while also stabilising 2D numerical models by smoothing boundary flow rates.

In this section, the most modern and utilised methods available in the literature for parametrising turbines, sluices, ramp functions and lagoon wetted area will be presented. In Section 4 novel methodologies for parametrising turbines (in both power generation and pump mode), wetted area and ramp functions are developed in this work.

3.2.1 Sluices

In order to estimate flow rate through sluice gates, the literature has utilised the orifice equation [7; 6]:

$$Q_s = C_d A_S \sqrt{2g|h|}, \tag{3.1}$$

where flow rate Q_s is a function of h (the head difference between ocean and lagoon), C_d is the dimensionless discharge coefficient (greatly dependent on sluice gate design [6]), A_S the sluice/orifice area and g the gravity acceleration ($9.81 m/s^2$).

Although modern experimental results show that C_d can reach values in the range $1.3 \leq C_d \leq 1.45$ [73], which can lead to significant increases in power output for TRS [74], the literature has chosen a conservative $C_d = 1$, when utilising Eq. 3.1. This conservative approach is corroborated by [6], following experimental results of a large sluice gate prototype that measured $C_d \approx 1$ for flow rates up to $2000 m^3/s$ [75].

3.2.2 Turbines – Power, Pump & Idling Modes

Among the various hydraulic turbines available and utilised in hydroelectrics, bulb turbines have been the option of choice for TRS due to their high efficiency at low heads. Amid bulb turbine types, double-regulated bulb turbines (Fig. 3.6a) have shown to perform optimally in TRS, for their ability to maintain high efficiency when submitted to a variable water head, generate power in reverse direction (although with efficiency penalties) and of being operated as pumps if required, further increasing TRS energy extraction capabilities [68; 71; 6; 67].

As discussed by [6], the ability of double-regulated bulb turbines to sustain high efficiencies for different water head scenarios is possible due to their capacity of regulating flow quantities through distributor vanes (or wicket gates) and keeping an optimum angle between flow profile and runner blades when pitching the latter. By experimenting with small-scale turbines under different flow and head conditions, dimensionless performance turbine "hill charts" (Fig. 3.6c) can be acquired. By fixing turbine diameter D, capacity P_C and rotational speed S_p, power and flow charts can be drawn (Fig. 3.6b), with power and flow estimates (P and Q_t, respectively) following Eq. 3.2 and Eq. 3.3.

$$P = |\rho g Q_t h \eta_{eff}|, \tag{3.2}$$

$$Q_t = Q_{11} D^2 \sqrt{|h|}, \tag{3.3}$$

where ρ is the seawater density ($1024 kg/m^3$), Q_{11} the specific unit discharge and η_{eff} the total efficiency, accounting for the losses suffered by the turbine during its operation. The derivations for S_p, Q_{11} and η_{eff} will be shown in the next section. Following Eq. 3.2, once power output reaches P_C, and the water head continues to increase, Q_t needs to be reduced in order for P remain constant. The shape of the Q_t flow chart in this phase (Fig. 3.6b) is explained by the fact that once $P = P_C$, flow through the turbine becomes inversely proportional to the head difference [71].

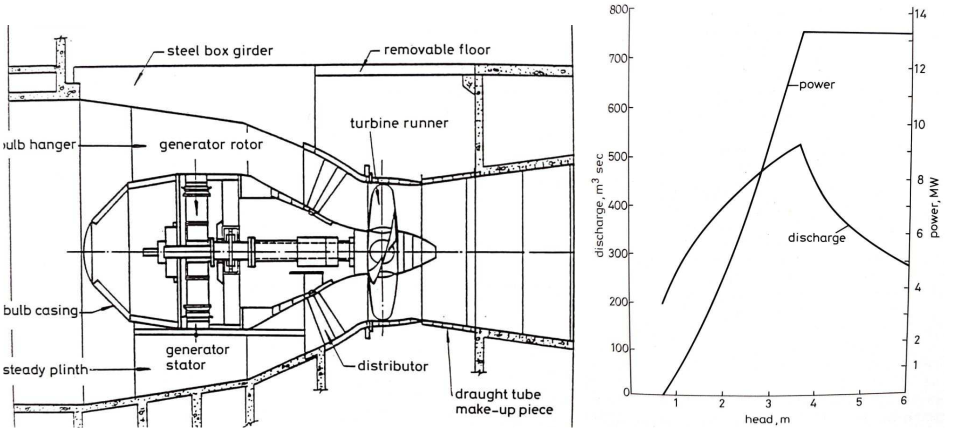

(a) Detailing of units that compose a double-regulated bulb turbine.

(b) Power and flow rate chart for a double-bulb unit with fixed diameter and capacity.

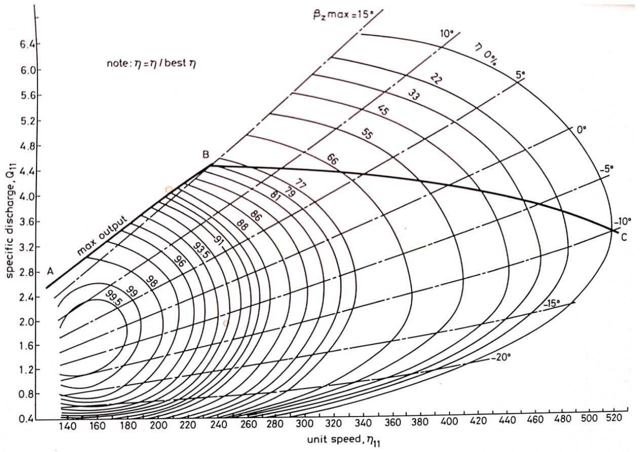

(c) Performance hill chart provided by Sulzer Escher Wyss of Zurich.

Figure 3.6. Components and chart's parametrisation detailing for a double-regulated bulb turbine (from [6]).

Regarding pump operational mode of turbines, literature research has been very limited. In fact, most recent studies have simulated pump operation either by fixing the efficiency (for any head difference) to 70% or 40% [71; 76], or by directly following experimental pump efficiency curves from [76] (Fig. 3.7).

For negative head differences, the literature has also adopted an idealised pump efficiency interpretation [23] (disregarding efficiency variations with pump rotation), cast as:

$$\eta_p = \frac{P_{out}}{P_{in}}, \tag{3.4}$$

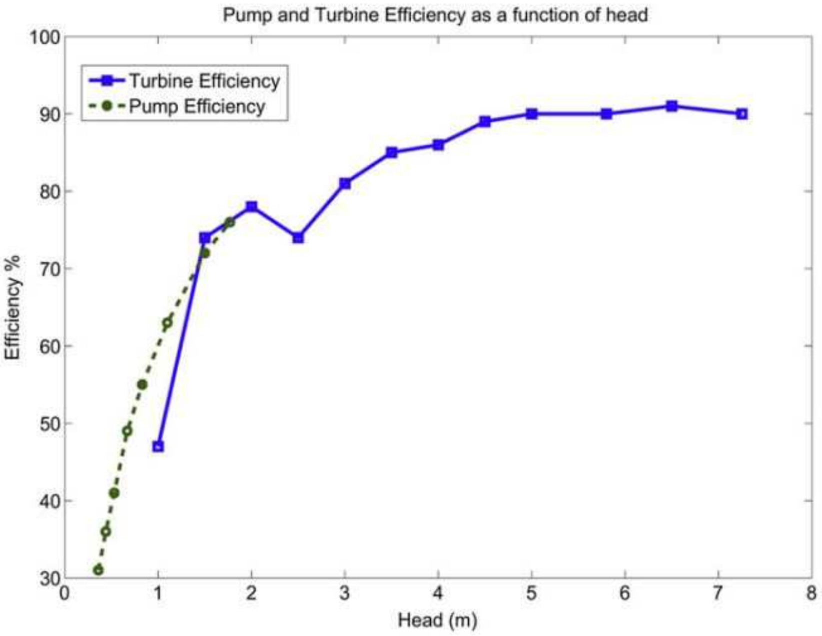

Figure 3.7. Experimental turbine and pump efficiency curves from a hydraulic turbine with unknown specifications (from [76]).

where P_{in} is the electrical power available to the pump and P_{out} is the rate of work exerted from the pump to the fluid:

$$P_{out} = \rho g Q_p |h_p|. \tag{3.5}$$

Also, Q_p is the flow-rate through the pump and h_p the negative head surpassed during pumping.

From Equations 3.4 and 3.5 we obtain:

$$Q_p = \frac{\eta_p P_{in}}{\rho g |h_p|}. \tag{3.6}$$

Although Eq. 3.6 can help to estimate pump flow rate, given a negative head and input power, it cannot explain maximum pump shutoff heads (where $Q_p = 0$), or maximum pump flow rates that occur when $h_p = 0$. These issues, combined with the fact that turbine specifications (e.g. diameter, capacity, applied power input) were omitted in the [76] study, and that pump efficiencies (Fig. 3.7) are only available for negative heads in the small range $[\approx .3\ m, \approx 1.9\ m]$, indicate the urgent need of more accurate and generalizable TRS pump models. Supporting this claim are the measurements from the La Rance TRS [67], showcasing pump operation in both positive and negative water head scenarios, with pump shutoff negative heads up to 6 m. In order to fill this gap, a more generalizable TRS pump model theory is presented in this work in Section 4.2.5.

Finally, when operating in “idling” mode, turbines act as sluices, aiding the sluicing stage of operation. For estimating flow rates in this mode, the orifice equation Eq. 3.1 is employed. Although generally a $C_d = 1$ has been utilised for several studies [17; 68; 19] a $C_d = 1.36$ has been set as a standard for modern SBL studies, considering fixed turbine specifications [23; 71]. The modified value for C_d scales the calculated flow rate from the orifice equation, smoothing flow rate transition when switching from turbine generation to sluicing modes.

Although parametrisation allows for simplifying turbine representation, all three dimensional flow characteristics are lost in the process. Therefore, the use of such methods is more adequate when dealing with preliminary analysis (such as TRS) or when blade/turbine scale hydrodynamics are not of immediate interest.

3.2.2.1 Andritz Performance Hill Chart

The most recent (and utilised) equations for describing flow rate and power for low head bulb turbines are based on experimental results from Andritz Hydro [27]. The edited Andritz chart shown in Fig. 3.8 demonstrates how turbine unit speed n_{11} and specific unit discharge Q_{11} (obtained experimentally) are related. The graph also shows wicket gate and running blade openings α and β, in degrees), and iso-efficiency curves E_f.

By specifying the parameters of the turbine: diameter D, number of generating poles G_p and grid frequency f, the turbine rotation S_p (rpm) is obtained from $S_p = 120f/G_p$. Furthermore, unit speed n_{11}, turbine flow rate Q_t and power output P are calculated as:

$$n_{11} = S_p D/\sqrt{|h|}, \tag{3.7}$$

$$Q_t = Q_{11} D^2 \sqrt{|h|} \tag{3.3}$$

$$P = |\rho g Q_t h \eta_{eff}|, \tag{3.8}$$

where C_E is the product of other efficiencies shown in Table 3.3, and $\eta_{eff} = E_f C_E$.

When h is available, n_{11} is estimated directly from Eq. 3.7. For calculating Q_t and P, Q_{11} and E_f are obtained experimentally by adjusting the opening of the wicket gates (α), the pitch angle of the runner blades (β) and crossing the values with the obtained n_{11} (see Fig. 3.8). In order to choose appropriate values for α and β, a parameterised curve of maximum power output was drawn over Fig. 3.8 (blue line) by

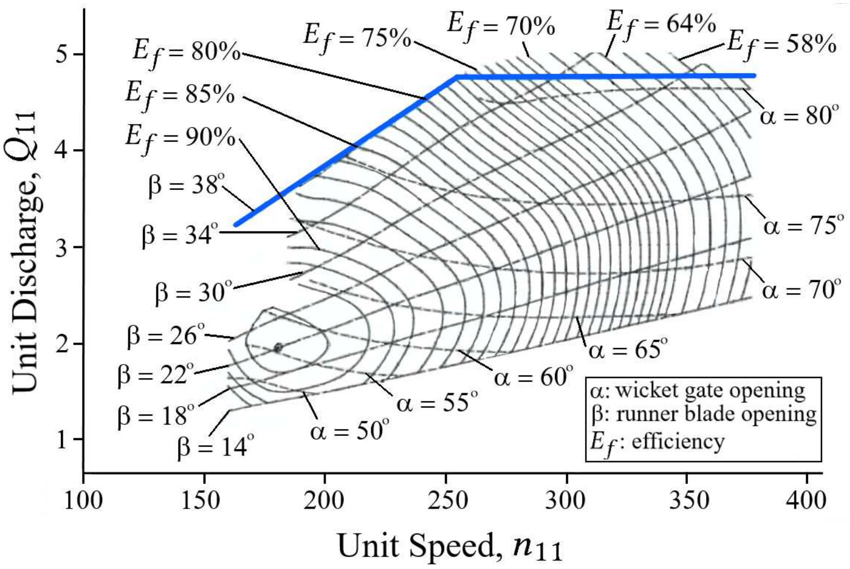

Figure 3.8. Edited Andritz Performance Hill Chart for a double regulated turbine (varying α and β angles), adapted from [27]. The blue line represents the parametrised maximum power output curve.

Table 3.3. Other efficiency considerations for TRS turbines [68].

TRS Efficiencies	(%)
Generator	97
Transformer	99.5
Water friction	95
Gear box/drive train	97.2
Turbine availability	95
Turbine orientation (Flood Gen. only) [23]	90

following the path where the product between E_f, Q_t and h is maximised. If we assume α and β are automatically adjusted to always be in the maximum power output curve, then Q_{11} and E_f become functions of n_{11}, as shown in Eq. 3.9 and (3.10).

$$\begin{aligned} Q_{11} &= (0.0166)n_{11} + 0.4861;\ \text{(when } n_{11} \leqslant 255) \\ Q_{11} &= 4.75;\ \text{(when } n_{11} > 255), \end{aligned} \tag{3.9}$$

and

$$E_f = (-0.0019)n_{11} + 1.2461. \tag{3.10}$$

Therefore, during power generation, Eq. 3.3 is used for estimating flow rate through the barrage.

3.2.3 0D & 2D Simulation

In order to estimate lagoon water motion inside the impounded area and power production of TRS, numerical models (0D to 3D) can be considered. When the goal is the optimisation of TRS operation for maximising energy generation, 0D models are usually chosen, given their computational efficiency, and the fact that for "small-scale" projects, such as the Swansea Bay Tidal Lagoon, 0D models present good agreement with more complex finite-element 2D models [71; 23; 25; 2; 37]. 0D models are derived from conservation of mass:

$$\frac{dL}{dt} = \frac{Q_T}{Al(L)}, \tag{3.11}$$

where L is the water level (in meters) inside the lagoon, Q_T is the total directional water flow rate (m^3/s) from both sluices and turbines and $Al(L)$ is the variable lagoon area (m^2). From Eq. 3.11, the lagoon water level at the following time-step (L_{t+1}) can be calculated by a backward finite difference method:

$$L_{t+1} = L_t + \frac{Q_T}{Al(L)}\Delta t, \tag{3.12}$$

where L_t is the water level at time t and Δt the discretized time step.

2D simulations of TRS are special boundary conditions implemented in shallow water numerical models. Such models (e.g. DG-SWEM, Thetis) are capable of simulating wave physics, such as resonance and reflection, while also being sensitive to coastal funnelling effects and variable bathymetry (Section 2.1.2). Nevertheless, at the location of the boundary condition, the flow rate through hydraulic structures is simplified as a pure function of parametrised equations, as showcased in Sections 3.2.2 and 3.2.1. Examples of simulated 2D TRS in the literature are shown in Fig. 3.9. A more in depth explanation for 2D TRS models will be presented in our implementation in Section 4.3.2.

When starting or stopping either turbines or sluices (for 0D and 2D TRS models), the literature has used heuristic sinusoidal ramp functions $r(t)$ as a function of time t:

$$r(t) = sin[(\pi/2)(t - t_m)/t_r], \tag{3.13}$$

where $t \in [t_m, t_m + t_r]$, t_r = transition time (around 15 min to 20 min to fully start or stop hydraulic structures [23; 17]), and t_m is the time when the current operation (e.g. $E.G$, Sluicing) was triggered. $r(t)$ is then multiplied by turbines and sluice equations (Eq. 3.1 and 3.3), smoothing flow rate outputs.

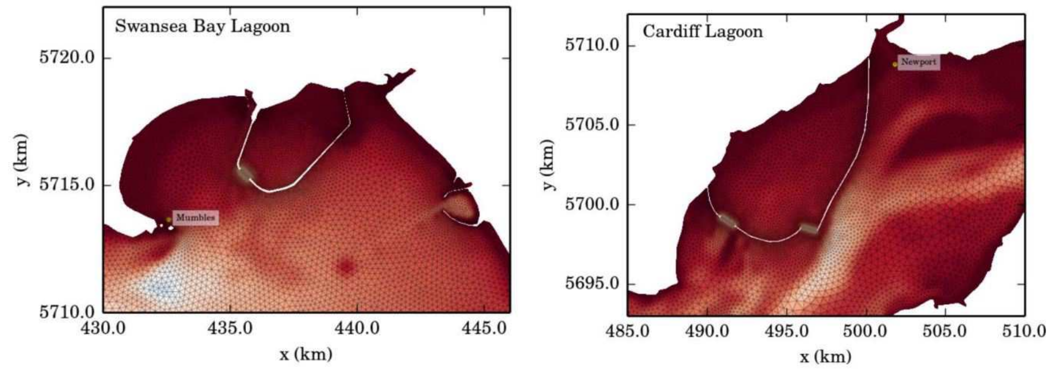

(a) Swansea Bay Tidal Lagoon. **(b)** Cardiff Lagoon.

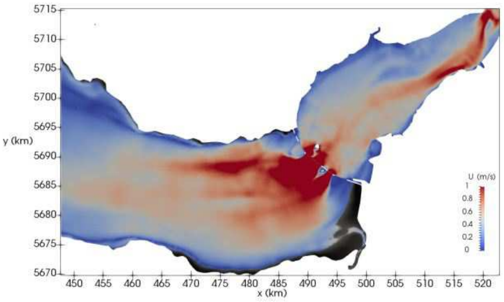

(c) Severn Barrage.

Figure 3.9. 2D Thetis finite element examples of TRS representation (from [23; 71]).

3.2.4 Lagoon Wetted Surface Area

As shown in Eq. 3.12, in the previous section, 0D models require an approximation for the lagoon wetted (or surface) area in order to predict lagoon water levels at every time-step. For small scale lagoons, research has shown that the plane area obtained from bathymetric data (Fig. 3.2) can be utilised in 0D models, yielding good agreement of results (maximum predicted energy deviation $\approx 6\%$ [25]) with more complex 2D models, when utilising the same operational strategy. A good example of this are the several comparison studies between 0D and 2D models for the SBL project [23; 25; 71], where the impounded area oscillates around $8\ km^2 \lesssim Al(L) \lesssim 12.5\ km^2$.

Conversely, as the scale of TRS increases (e.g. Cardiff Lagoon, Severn Barrage), the plane surface area obtained from bathymetry becomes an oversimplification due to wave reflections inside the impounded lake. Indeed, a study from [71] compared energy estimations between 0D, 1D and 2D models utilising the SBL and the Severn Barrage as case studies, where 1D and 2D models are obtained from the shallow water equations. For the 0D model, the wetted area estimate was considered a constant, using the enclosed bathymetry as reference. Results for the yearly energy estimate for all models (assuming fixed H_{start} and H_{min} operational heads) are showcased in Table 3.4. We note that the agreement between 1D and 2D models are closer than 0D to 2D, for both SBL and Severn Barrage case studies. This is explained by the fact that both 1D and 2D models can simulate wave resonance and reflections along the length of the estuary. Then, thanks to the elongated design of the Severn Barrage (Fig. 3.10), an equivalent 1D behaviour for both 1D and 2D simulations is observed.

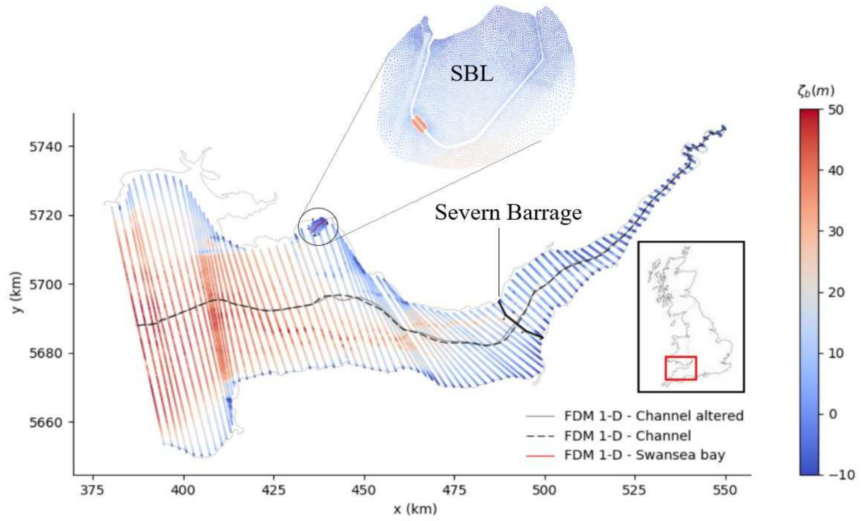

Figure 3.10. 1D finite difference and 2D finite element models for SBL and Severn Barrage case studies (edited from [71]).

Even though the 0D model showed reasonably close results to the 2D model for the SBL project, the deviation between 0D and 2D models for the Severn barrage was ≈ 100%. Since 0D models are required for optimisation routines, the significant errors associated with 0D representations of large scale TRS means that, to date, large scale TRS have not had their operation optimised with the best available optimisation routines.

Table 3.4. Annual Energy Estimates for Swansea Bay and Severn TRS (edited from [71]).

Simulation	Swansea Lagoon (11.6 km^2) (TWh/yr)	Severn Barrage (573 km^2) (TWh/yr)
0D − TW	0.445	20.71
1D − TW	0.436	11.50
2D − TW	0.411	11.73
0D − TWP	0.530	22.41
1D − TWP	0.512	11.57
2D − TWP	.475	11.18

3.3 TRS Optimisation – Baselines

In this section, the theory behind analytical and numerical techniques for TRS operational optimisation from the literature will be presented.

3.3.1 Analytical Methods

By choosing the TW scheme of operation, [7] estimated an analytical upper bound for the available energy E_{TW} in a general TRS. For calculating E_{TW}, a series of simplified assumptions are made:

- The incoming tide is sinusoidal with amplitude A_m.
- Operational heads are equal $H_{start} = H_{min}$.
- Flow rate through operating turbines is constant throughout the operation.
- Impounded lagoon area Al is constant with depth.

So that,

$$E_{TW} = 0.37 E_{Pmax}, \text{ where } E_{Pmax} = 4\rho g {A_m}^2 Al \tag{3.14}$$

In Eq. 3.14, E_{Pmax} is the maximum energy that can theoretically be extracted by instantaneously emptying the enclosed body of water (from high tide to low tide) and A_m the ocean tidal range amplitude.

Another recent take on analytical upper bound for TRS was proposed by [16], where the following assumptions were made:

- The incoming tide is sinusoidal with amplitude A_m.
- Turbine flow rate and total efficiency losses η_{eff} are constant throughout the operation.
- $E.G$ and $F.G$ time periods are centered at periods of minimum and maximum ocean water levels, respectively.
- Impounded lake water volume ΔV between low and high tides is known.

The upper bound for TW can be estimated, with a higher degree of accuracy (for considering turbine efficiency and impounded water volume), as

$$E_{TW} = \rho g \eta_{eff} 2 {A_m}^2 \Delta V. \tag{3.15}$$

Although both expressions give a reasonably good estimate for TRS energy capabilities, they are unable to consider more complex scenarios, such as the variant TW scheme operation (Section 3.1), pump operational modes and variable turbine efficiency with water head.

For simulating the potential energy of TRS for these complex scenarios, numerical optimisation routines are coupled with 0D models of TRS. The state-of-art methods of TRS optimisation are shown and explained in the next sections.

3.3.2 Constant Head Methods

Widely used as a baseline for several TRS optimisation studies [23; 77; 25; 78; 28], constant head (CH) optimisation methods maximise the energy extraction of a TRS over a set period of time (day, month, year) by finding best, fixed TRS operational heads, i.e. setting [H_{start} , H_{min}] for classical TW operation (Section 3). This naive approach is commonly optimised with brute-force algorithms due to the reduced number of parameters to be optimised.

The resulting operation suffers from a null power production when the head difference between ocean and lagoon fail to meet the required H_{start} (Fig. 3.11). Therefore, in order to best exploit the various tidal ranges provided from the ocean, flexible operation methods, explained in the next section, have been developed.

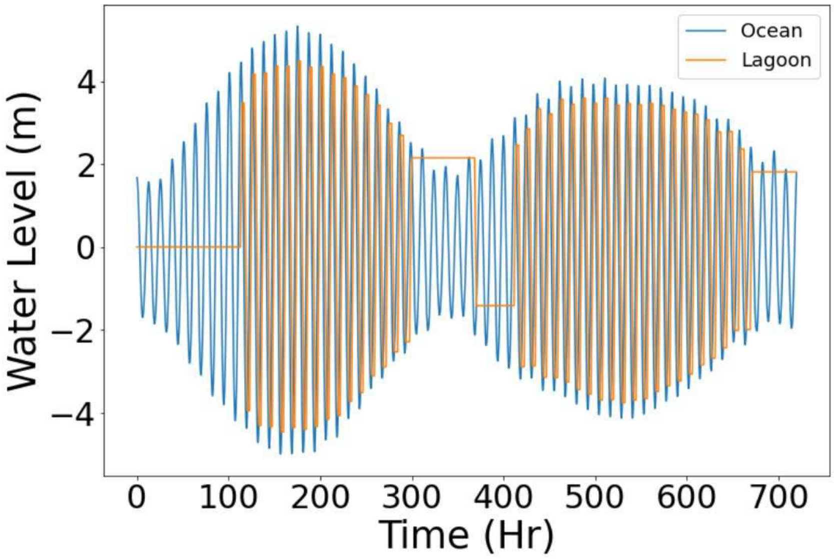

Figure 3.11. Example of a month CH optimisation for a TW classical operation.

3.3.3 Flexible Operation Methods

Flexible operation methods are the literature's state-of-art strategy for estimating the available energy in TRS. The idea behind this type of approach is to time (through HC or TC methods) the sequential operation of hydraulic structures according to the current (and following, when applicable) tidal cycle's amplitude [25; 23]. With the assumption of well predictable tides, flexible operation of turbines and sluices can be inferred by "looking–ahead" through harmonic or numerical tidal prediction methods [79] and applying the acquired operation to the real, measured ocean – a procedure that needs to be repeated for every new tide. In fact, and to the best of our knowledge, the requirement of accurate future tidal predictions (concurrent with measured data and up to a multiple of half-tidal cycles into the future) has been the basis for all optimisation routines developed for enabling flexible operation [77; 25; 23; 78; 24; 2; 28]. This constraint can be a problem when future tidal predictions are unavailable, unreliable, have some associated validation cost [80] or are regulated by private companies or government agencies.

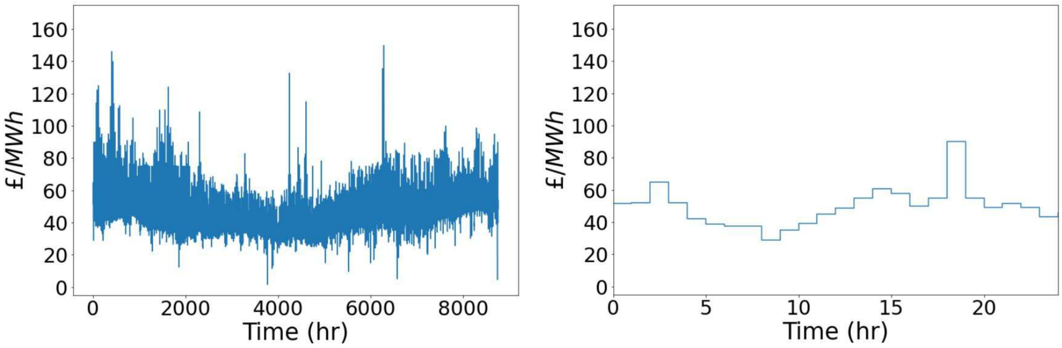

(a) 2017 day-ahead-market values.

(b) Day-ahead-market values for the first day of 2017.

Figure 3.12. UK day-ahead-market measurements for 2017, acquired from [81].

Beyond maximising energy extraction of TRS, the objective of flexible operation methods can be modified as to maximise revenue. For instance, in a SBL case study by [24], revenue-oriented optimisation was possible by assuming a perfect forecast of tides and the 2017 day-ahead-market (Fig. 3.12). With these assumptions, income based optimisation led up to 10 % increase in revenue, while generating $4-5$ % less energy. The day-ahead-market (used in UK, US and several European countries) is defined as the value of electricity ($\pounds/MWh$), fixed for every hour. The value is stipulated for the next 24 hours (starting at midnight) through an auction process set at midday.

State-of-art optimisation routines in the literature utilise either grid search (a brute-force approach), gradient-based, global optimisation [25; 23; 24] and, more recently, genetic algorithm methods [28] to optimise the operation of TRS.

It is important to emphasise that, thus far, all flexible operation studies from the literature have considered an idealised perfect forecast scenario, where tidal predictions have a perfect match with tidal measurements – characterising an "upper bound" estimate. For improving upon this, we develop in this work a "real-time" estimate, where the acquired TRS operation through tidal predictions is applied to tidal measurements. The uncertainty deviations between "upper bound" and "real-time" methods are discussed in Section 5.3. Furthermore, the capabilities of revenue-oriented control of TRS are also explored in this work (for both HC and DRL methods), in Section 5.3.

3.3.3.1 Head-Controlled Flexible Operation Methods

Developed by [25], HC flexible operation methods break tidal measurements into sequential tidal (or half tidal) cycles. Since for every half tidal cycle, $E.G$ or $F.G$ will occur (given head differences above H_{mt}), optimisation routines can be employed to find best operational heads. For a fixed tidal (or half tidal) signal, this is done by running 0D TRS models multiple times and exploring, for each run, different operational head combinations with the goal of maximising power generation. For all case studies in [25], TW and TWP classic operation of TRS (Sections 3.1.1, 3.1.2) were considered. In this setting, brute-force optimisation methods were employed, guaranteeing a global optimal solution, given a precision of $\approx 1\ cm$ for the operational heads.

For the TW operational scheme case study, the strategy for finding the optimum pair of H_{start} and H_{min} for every tidal cycle was named ET (Every Tidal cycle method). Similarly, when optimising every half tidal cycle, the process was named EHT (Every Half-Tidal cycle method).

With the goal of further exploiting the predictability of tides, EHT methods can be augmented by aiming to find not only the best H_{start} and H_{min} pair for the half tidal cycle, but also the H_{start} and H_{min} pair for the following half tidal cycle. When the best quadruplet is found, only the first pair of H_{start} and H_{min} is used to run the 0D model, and the process repeats a half-tide forward in the future. The idea of this method is to "prepare" a better starting point (initial lagoon water level) for the following half-tide operation. The developed strategy is named EHN (every half-tidal cycle and next). By following the same logic, augmented ET methods are named ETN (every tidal cycle and next).

The described ET, EHT, ETN and EHN methodologies can also be applied when considering pump operation. For this augmentation a starting and finishing pump operational head is added to each half tidal cycle. Finally, when using a pump model (Eq. 3.6), either pump flow rate or power needs to be fixed during the pumping process. Fig. 3.13 shows the considered optimisation period for all flexible operation methods discussed.

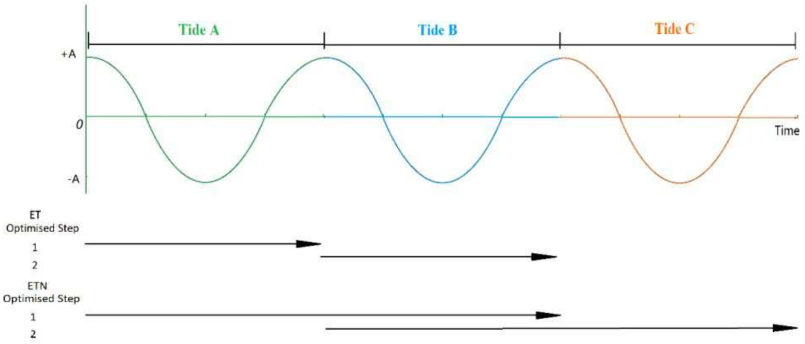

(a) ET and ETN methods.

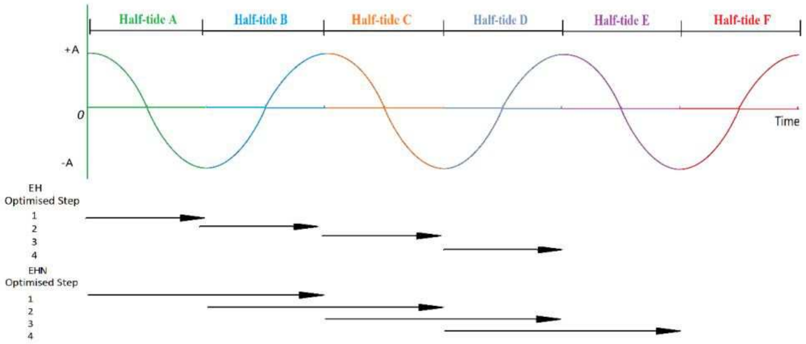

(b) EHT and EHN methods.

Figure 3.13. Schematic illustrations of various HC flexible operation methods. For each method, arrows indicate the time interval where optimal pairs (or quadruplets) of H_{start} and H_{min} are chosen (from [25]).

3.3.3.2 Time-Controlled Flexible Operation Methods

As shown in the work of [23], another way of performing flexible operation involves optimising the time length of each sequential mode of operation in a TW or TWP scheme. As a TWP example for this method, the sequential time length of the each operational mode in a tidal cycle is added to a vector $\vec{v} = t_i, i = 1, ..., N$, where each t_i corresponds to a different TRS operation, as shown in Fig. 3.14.

The total energy per tidal cycle can then be optimised by finding a best vector $\vec{v}$. Due to the computational cost of TC methods, literature solutions for this problem

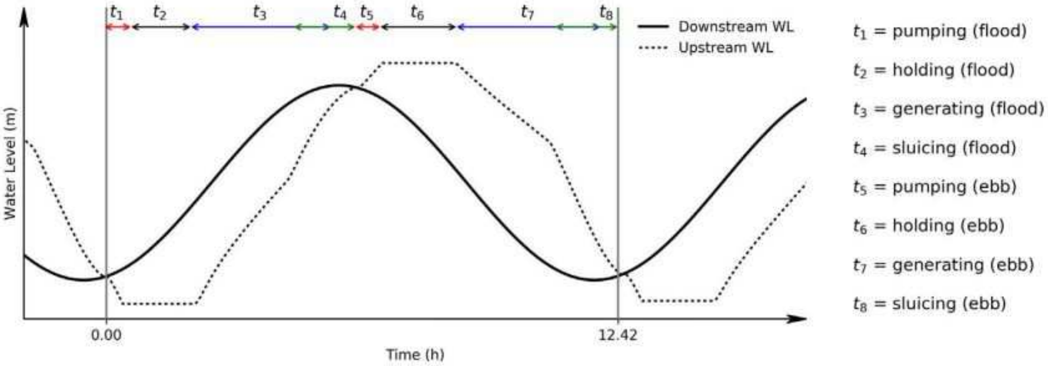

Figure 3.14. TC flexible operation detailing (from [23]).

were obtained with either Broyden-Fletcher-Goldfarb-Shanno with bounds (L-BFGS-B) [23], or basin-hopping [24] optimisation algorithms, that balance performance without guaranteeing a global optimal solution. As can be seen in Fig. 3.14, the variant operation of TRS, with independent operation of sluices, was contemplated by [23], when introducing the method (displayed in Fig. 3.14 as the overlapping of sluicing and power generation time vectors).

While TC and HC methods share the same (i) sequence of operation of hydraulic structures per tidal period and (ii) pump operation constraints (thoroughly detailed in Section 3.1.2), TC methods have pre-determined (fixed) operational heads for automatically starting holding periods (during ebb and flood tides) and turbine idling stage [23]. Due to this reduced flexibility of operation, HC optimisation methods are chosen in this work as the state-of-art control method for TRS.

3.4 AI Applications for Renewable Energy Systems

The concept of AI is not unique, but it could be well described by the definition given by the European Commission in its Communication presented to the European Parliament in 2018: “AI refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals. AI-based systems can be purely software-based, acting in the virtual world or AI can be embedded in hardware devices” [82]. Due to technical advances, more than 60,000 articles are published annually with AI applications in computing technologies, as well as a large number of applications for AI, such as problem solving and optimisation, data analysis, image recognition, text processing etc [83; 84]. As far as the applications of AI to energy production are concerned, a significant number of publications only started to emerge around 2003, with an exponential increase since then ($15,368$ works were published on AI & Energy, from 1970 to 2019, according to Scopus) [84]. This increase is partially explained by advances in solar and wind technologies and applications [84; 85]. Besides a comprehensive review on AI & Energy, Reference [84] investigates and compares several AI techniques employed in wind turbines monitoring systems (Artificial Neural Networks; Fuzzy Logic; Genetic Algorithms; Particle Swarm Optimisation; Decision Making Techniques; and Statistical Methods), concluding that not all AI techniques are appropriate for renewable energy applications – as most of them are designed with one purpose –, and thus, that it is necessary to know the best suited algorithm for each application

In recent years, DRL techniques have represented a breakthrough for AI methods, performing optimally in several real-time problems that previously could only be solved by experienced human operators. The potential of DRL techniques for solving such tasks has received a significant attention in applications with games, such as Atari, Chess, Shogi, Game of Go and StarCraft II [47; 86; 87], where smart DRL agents were able to showcase a performance superior to human players. Amid the advancements of DRL algorithms, complex real-life optimisation problems also started to be tackled, such as trading and finance, self-driving cars, healthcare and energy efficiency (e.g. heating, ventilation, air-conditioning and datacentre cooling systems) [88; 89; 90; 91; 92]. In the context of renewable energy, reinforcement learning techniques have also been used for optimising the operation of smart grids [93; 94; 95; 96], wind turbines [97; 84; 98], solar panels [99; 100; 85; 101; 102], stream turbines [103], high frequency wave energy converters [104] and more recently, TRS [105]. In particular, the energy demand forecasting problem has been successfully investigated using Long Short-Term Memory (LSTM) neural network [106] for different renewable energies, e.g., smart grids [96], wind turbine power [98] and solar energy [101; 102]. The characteristic features of this approach could be interesting for some particular operational modes of TRS (maximising the revenue, for instance), being investigated in Section 5.3.

Chapter 4

Developed Methodologies

In this chapter, we firstly propose improvements to the state-of-art control optimisation of TRS, enabling higher energy/revenue outputs and real-time control. Secondly, we develop general techniques for parametrising TRS components through experimental data. For this, experimental measurements from the La Rance Tidal Barrage are utilised as a case study. Thirdly, pseudo-code implementations for enabling numerical model simulations (0D and 2D) for TRS, developed for this work, are presented. Finally, by utilising the SBL project as a case study, we develop our MDP within Unity ML-Agents using our 0D TRS model, so that our DRL agent can be trained.

4.1 Improvements to State-of-Art Control Optimisation of TRS

In this section we propose improvements to HC control optimisation of TRS. We have chosen HC methods instead of TC approach due to their (i) low computational cost, which allows for using brute-force algorithms and (ii) higher flexibility of operation, enabling the optimisation of pump and idling turbine stages (without pre-determined operational heads), therefore guaranteeing a global solution, (considering a given precision for operational heads), as discussed in Section 3.3.3.2.

4.1.1 Head-Controlled, Variant Operation of TRS

In this section we propose a modification of the discussed HC classical operation (Section 3.1), which allows for opening TRS sluice gates at the end of "flood" and "ebb" generation stages, independently of H_{min}. We show that this variant scheme of operation increases power generation due to higher lagoon tidal ranges when starting the

Table 4.1. TW, HC variant hydraulic structures operation.

Operational Mode	Turbines	Sluices	Power Gen.
$E.G$ (if $T_{Tr} == True$):	On	Off	Yes (if $h > H_{min}$)
$F.G$ (if $T_{Tr} == True$):	On	Off	Yes (if $h > H_{min}$)
Sl (if $Sl_{Tr} == True$):	On	On	Yes (if $h > H_{min}$)
Ho (if $Ho_{Tr} == True$):	Off	Off	No

h: Head difference between ocean and lagoon
T_{Tr}: Turbine power generation trigger ($h > H_{start}$)
Sl_{Tr}: Sluice trigger ($h > HS_{start}$)
Ho_{Tr}: Holding trigger ($h == 0$)

next "ebb" or "flood" stages. Disregarding pumping capabilities, the variant TW operation scheme requires 3 control variables every half-tide: H_{start}, H_{min} and HS_{start} (sluice gate starting head). A comparison of the water level variations within the lagoon, following classic and variant operations of hydraulic structures, can be seen in Fig. 4.1. Table 4.1 shows all possible combined operations of turbines and sluices, with resulting power generation, for each control stage in TW HC variant operation.

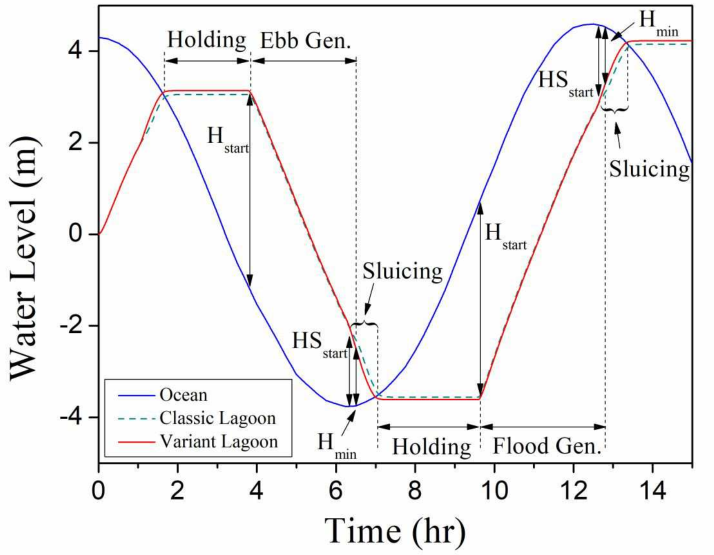

Figure 4.1. Classic and variant "two-way scheme" operation. Ocean level is represented by the blue line, while the lagoon level is shown in either green dashed lines or red, for classic or variant lagoon operations, respectively.

We also apply the independent operation of sluices to TWP schemes (Section 3.1.2). For simulating pumping capabilities in TW schemes, two extra control variables are added for every half-tide, namely (i) a maximum pumping head H_p and (ii) a fixed power input P_{in}. The pumping equations utilised in our model are drawn from our developed pumping model (Section 4.2.5).

Table 4.2. TWP, HC variant hydraulic structures operation.

Operational Mode	Turbines	Sluices	Power Output/Input.
$E.G$ (if $T_{Tr} == True$):	On	Off	Yes (if $h > H_{min}$)
$F.G$ (if $T_{Tr} == True$):	On	Off	Yes (if $h > H_{min}$)
$E.P$ (if $P_{Tr} == True$):	On	Off	Yes (if $h < H_p$)
$F.P$ (if $P_{Tr} == True$):	On	Off	Yes (if $h < H_p$)
Sl (if $Sl_{Tr} == True$):	On	On	Yes (if $h > H_{min}$)
Ho (if $Ho_{Tr} == True$):	Off	Off	No

h: Head difference between ocean and lagoon
T_{Tr}: Turbine power generation trigger ($h > H_{start}$)
Sl_{Tr}: Sluice trigger ($h > HS_{start}$)
P_{Tr}: Pumping trigger ($h == 0$)
Ho_{Tr}: Holding trigger ($h == H_p$)

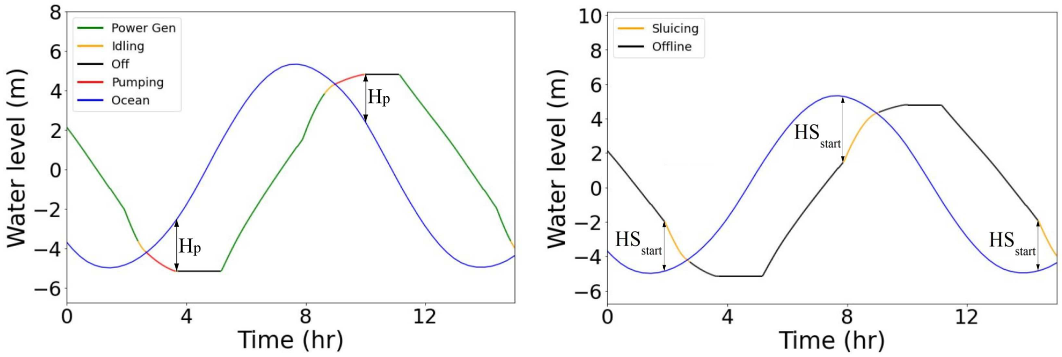

(a) Coloured lagoon water levels, following turbine operation.

(b) Coloured lagoon water levels, following variant sluice operation.

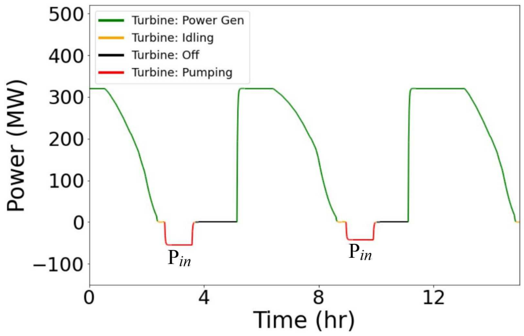

(c) Power generation.

Figure 4.2. Variant TWP HC scheme operation of TRS. Ocean is coloured in blue, while power generation and lagoon water levels are shown and coloured following turbines and sluices operational modes, with green representing power generation mode, orange – idling/sluicing mode, black – offline mode and red – pumping mode.

All possible combined operations of turbines and sluices for variant TWP HC operation are shown in Table 4.2. Coloured lagoon water levels, following turbine and sluice operation are shown in Fig. 4.2. Figs. 4.2a and 4.2c highlight the control variables H_p and P_{in} added to the TW turbine operational scheme, while Fig 4.2b highlights the independent operation of sluices through HS_{start}.

A comparison between power output estimates of classic and variant TW HC operation is shown in Section 5.1. Similarly, Section 5.3 presents a comparison between classic and variant power and revenue outputs, following TWP HC operation.

4.1.2 Real-time control of HC methods

So far, the literature on control optimisation of TRS has considered a perfect forecast of tidal predictions when estimating power generation and revenue – characterising an upper bound estimate. Therefore, we propose in this work the real-time control of TRS, through HC state-of-art methods.

In our method, both ocean prediction and ocean measurements signals are cropped at half-tides (Fig. 4.3). Then, using the EHT method for instance, best operational H_{start} and H_{min} are found for ocean predictions and then applied to ocean measurements. The cropping procedures occurs when the ocean signal's derivative changes signal. For avoiding the influence of high frequency waves (whose peak-to-valley period is $\approx 150s$) and for consistency with the opening and closing of hydraulic structures, we used $\Delta t = 15\ min$ during the cropping process.

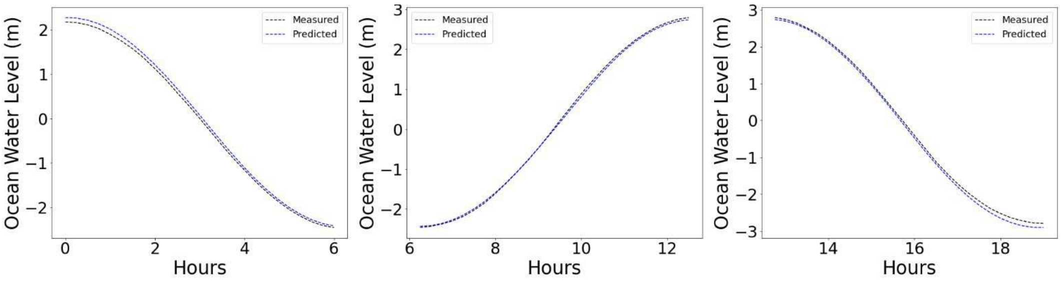

Figure 4.3. Ocean Measurements and Predictions, cropped for every half-tide. Data provided by BODC database.

Since the real-time method require forecast signal, ocean predictions are used when the optimisation goal is maximising energy. Naturally, in order for maximising revenue, the real-time method would require a form of forecasting the day-ahead-market (Section 3.3.3, for the UK case study). Although significant research has been performed on the field, due to time constraints we have not explored day-ahead-market forecast algorithms [107]. Therefore, our real-time approach for TRS operation, for both revenue or energy estimates, only uses forecast tidal signals (provided by BODC database). Nevertheless, upper bound estimates for energy and revenue are computed for comparison analysis.

Full code implementations for real-time control and upper bound estimates can be accessed in the GitHub page: `https://github.com/TmoreiraBR/PhD-Thesis-codes`. A comparison of the upper bound estimate against our DRL agent and real-time control through HC methods is presented in Section 5.3.

4.2 TRS Components Experimental Parametrisation – La Rance Case Study

In this section, we propose techniques to create parametric functions that emulate the behaviour of the hydraulic components and structures (turbines, sluices and enclosed estuary) that compose the La Rance tidal barrage, i.e., we reverse engineer a real TRS into a 0D parametric representation. The available data utilised in this section is mostly obtained from a 1975 study [67], executed in collaboration with Electricité de France (EDF), the company responsible for La Rance operation still to this day. Data from this study are utilised in order to create parametric models for turbines in ebb and flood operation (in both power generation and pump modes), sluices and equivalent impounded wetted area. Furthermore, a novel ramp function, named momentum ramp function, is proposed in order to help estimate flow rate and power variations when opening/closing TRS's hydraulic structures.

From [67], the observed variations of lagoon water level and power output/input for two days of observations are digitised, yielding Figs. 4.4a and 4.4b. In Fig. 4.4a, the observed operational modes correspond to a conventional ebb-only generation (E_oG) scheme, while Fig. 4.4b showcases a TWP scheme. The resolution of the digitised data is represented by a "×" label. The initials for the operational modes that were set for each scheme of operation are also presented at the top of each image. The detailings of each labelled initial, numbered in order of occurrence in Figs. 4.4a and 4.4b, are shown in Table 4.3.

It is worth emphasising that during $E.G$ at the E_oG scheme, turbines are expected to operate alone. This contrasts with the TWP scheme, where at the end of $E.G$ and $F.G$, the turbines are expected to operate together with sluices (i.e., the variant operation described in [105], and Section 3.1.1) [6; 67].

Design specifications for La Rance, taken from [67], are shown in Table 4.4. Also, [67; 108] present data showcasing turbine efficiency as a function of water head (for $E.G$ and $F.G$ modes of operation). From these data, $2nd$ order approximations are derived

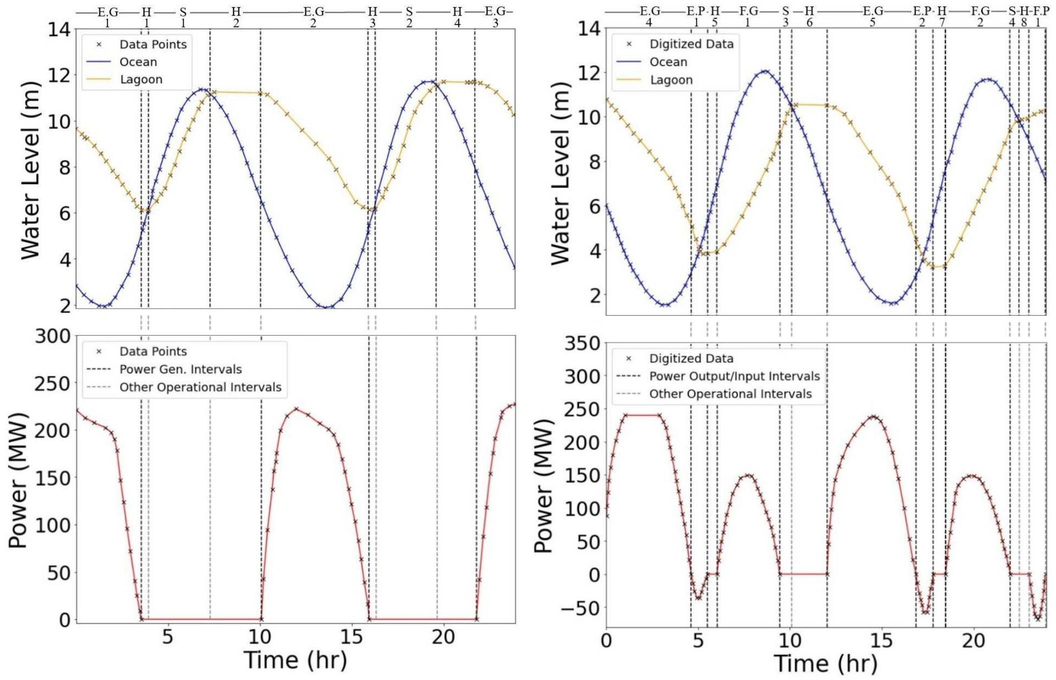

(a) Ocean and Lagoon Water Levels and corresponding Power generation for E_oG scheme.

(b) Ocean and Lagoon Water Levels and corresponding Power generation for TWP scheme.

Figure 4.4. Measurements of E_oG and TWP schemes of operation at La Rance. Digitised from [67].

Table 4.3. Operational Modes for E_oG and TWP La Rance's Schemes of Operation.

Label	Operational Mode	Operating Hydraulic Units[a]	
		E_oG	TWP
$E.G$	Ebb Generation	T: ON & Sl: OFF	T: ON & Sl: OFF $\rightarrow$ ON
$F.G$	Flood Generation	$N.A.$	T: ON & Sl: OFF $\rightarrow$ ON
Ho	Holding	T: OFF & Sl: OFF	T: OFF & Sl: OFF
Sl	Sluicing	T: ON & Sl: ON	T: ON & Sl: ON
$E.P$	Ebb-Oriented Pumping	$N.A.$	T: ON & Sl: ON $\rightarrow$ OFF
$F.P$	Flood-Oriented Pumping	$N.A.$	T: ON & Sl: ON $\rightarrow$ OFF

[a] For each scheme of operation [6; 67].
T: Turbines. Sl: Sluices. $N.A.$: Not Applicable.

and shown as dashed lines in Figs. 4.5a and 4.5b. Furthermore, data presenting the expected pump flow rates Q_p as a function of h_p for La Rance's bulb turbine in $E.P$ and $F.P$ operational modes, with a fixed (maximum) power input of $P_{in} = 6\ MW$, are available in Table 4.5. This table also showcases the pump shutoff head $h_s = -6\ m$ for both $E.P$ and $F.P$ modes [70]. From these data, $2nd$ order approximations are derived and shown as coloured dashed lines in Fig. 4.6. By reading Fig. 4.6, we

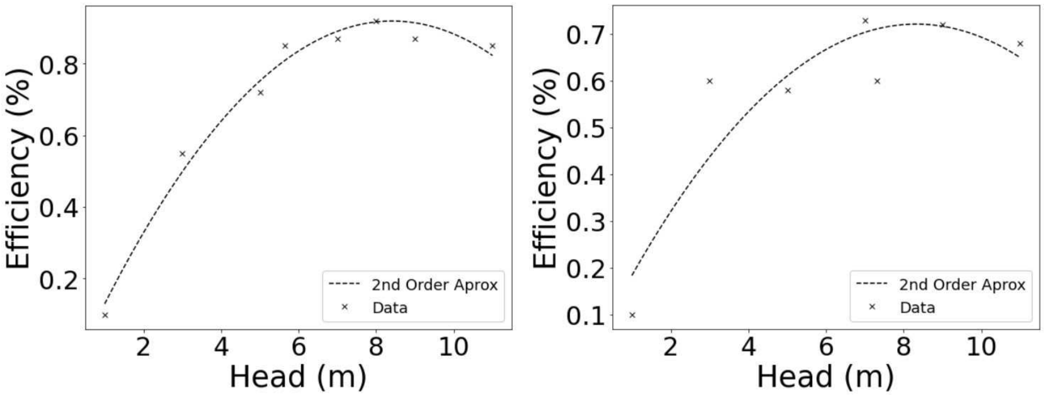

(a) $E.G$ Turbine Efficiency.
$(Eff_{E.G} = -.0144h^2 + .2417h + .0981)$

(b) $F.G$ Turbine Efficiency.
$(Eff_{F.G} = -.01h^2 + .167h + .0259)$

Figure 4.5. $E.G$ and $F.G$ turbine efficiency for La Rance (from [67; 108]).

Table 4.4. La Rance Barrage Design.

N^o of Turbines	24
Turbine speed (rpm)	94
Turbine Diameter ($metres$)	5.35
Turbine Capacity (MW)	10
Max. Pump Head ($metres$)	6
Sluice Area ($metres^2$)	900

Table 4.5. Measured $E.P$ and $F.P$ flow rates at La Rance, for a fixed power input $P_{in} = 6MW$ (from [67; 70]).

$Q_p(h_p, P_{in} = 6MW)$	Against $h_p(m) < 0$			
	-6	-3	-2	-1
Q_pEbb (m^3/s)	0	108	160	200
Q_pFlood (m^3/s)	0	100	168	175

note that maximum flow rates, when $h_p = 0\ m$, measure $Q_{MEbb} = 252.2\ m^3/s$ and $Q_{MFlood} = 215.8\ m^3/s$, respectively.

Reference [67] also provides various estimates for the volume (ΔV) of water stored in the La Rance estuary as ocean tide fluctuates to a height z, defined from the lowest tidal level in the equinoctial low tide ($z = 0$), reproduced here in Table 4.6. Finally, an upper bound for the maximum turbine flow rate (for any mode of operation) is set to 280 m^3/s, utilising site measurements from La Rance as reference [109].

In the following sections, the data presented here will be utilised for parametrising all elements that compose the La Rance TRS. As a first step, since sluices and turbines (in both power generation and pump modes) require a ramp function to be simulated,

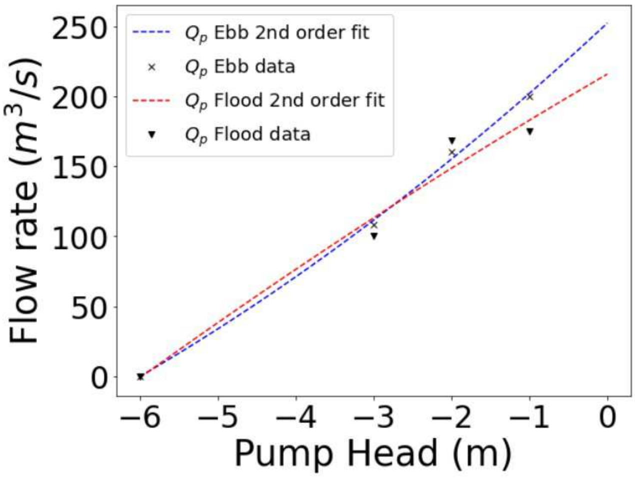

Figure 4.6. *E.P* and *F.P* flow rate estimates as a function of water head. ($Q_{pEbb} = 1.6{h_p}^2 + 51.8h_p + 252.2$) and ($Q_{pFlood} = -.6{h_p}^2 + 32.4h_p + 215.8$).

Table 4.6. Stored volume of water for the La Rance estuary (from [67]).

$\Delta V_i(10^6 m^3)$	0	65	110	150	184
$z_i(m)$	0	5	8.5	10.9	13.5

a novel momentum ramp function for the hydraulic structures is derived in the next section.

4.2.1 Momentum Ramp Function

In order to model flow rate variations from opening/closing the hydraulic structures that compose a TRS, a ramp function is required.

Given that solutions presented from the literature try to solve this problem with a heuristic approach (Section 3.2.2), we propose a new derivation for a ramp function more grounded on physical principles. In order to do so, we utilise an electro-hydraulic analogy. These types of analogies have been used extensively in the literature, with applications for both pipe fluid flow and open channel flow [110; 111].

Considering a pipe network – electric circuit analogy, we simulate the system: hydraulic unit (turbine or sluice) under a pressure differential (due to gravity or pumping system) as an *LR* circuit under a *DC* current (*LR-DC*), represented in Fig. 4.7.

In an *LR* circuit, *R* represents the electrical resistance inherent from the wire, while *L* represents the inductance. As current starts flowing due to a voltage difference applied from the battery, electrons reach the inductor (wire coil), which reacts by storing energy in an increasing magnetic field. This magnetic field increases until a

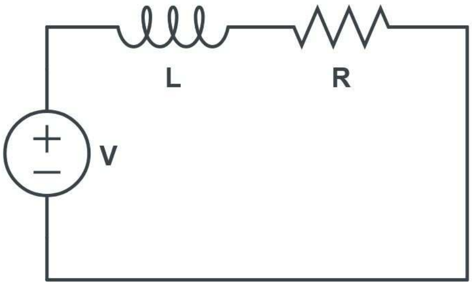

Figure 4.7. *LR* circuit representation under a *DC* current.

constant maximum value. Initially (after closing the circuit), when the electrons reach the inductor, the magnetic field is null. As the magnetic field increases, the inductor reacts with a back EMF (electromotive force), which opposes the circuit's voltage (Lenz's law). This opposite voltage ceases when the magnetic field stops increasing, which means that the current $i_{(t)}$ in the circuit slowly increases from a null value until a steady state current I_{ss} (Fig. 4.8)

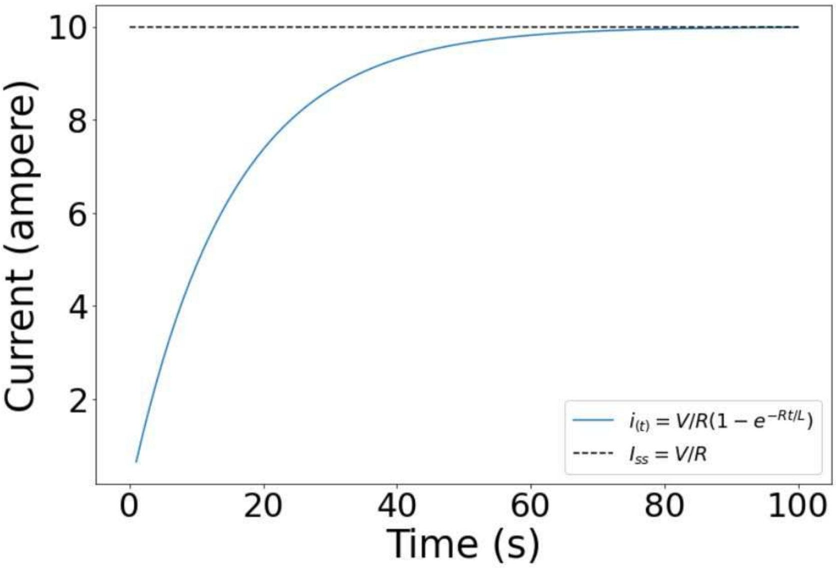

Figure 4.8. Electric current $i_{(t)}$ evolution in a *LR* circuit under a fixed voltage V.

In order to retrieve the $i_{(t)}$ current equation shown in Fig. 4.8 we utilise Kirchhoff's loop Rule, which states that: "the algebraic sum of the voltage differences is equal to zero". Therefore:

$$V - iR - L\frac{di}{dt} = 0, \tag{4.1}$$

where iR and $L\frac{di}{dt}$ are the voltage drop due to the resistor and inductor, respectively.

Assuming that $i(t=0)=0$:

$$V = i(t)R + L\frac{di(t)}{dt} \rightarrow V - i(t)R = L\frac{di}{dt} \rightarrow \frac{dt}{L} = \frac{1}{V - i(t)R}di, \tag{4.2}$$

Integrating on both sides:

$$\int_0^t \frac{dt}{L} = \int_0^i \frac{1}{V - i(t)R}di \rightarrow \frac{t}{L} = \int_0^i \frac{1}{V - i(t)R}di. \tag{4.3}$$

Utilising integration by substitution, $u = V - i(t)R \rightarrow \frac{du}{di} = -R$, therefore

$$\frac{t}{L} = \int_0^i \frac{-1}{Ru}du \rightarrow \frac{-Rt}{L} = ln(V - iR)|_0^i \rightarrow \frac{-Rt}{L} = ln\left(\frac{V - iR}{V}\right). \tag{4.4}$$

Taking the exponential of both sides and isolating the current term, the equation for the current in an *LR-DC* circuit is obtained:

$$i(t) = \frac{V}{R}(1 - e^{-Rt/L}). \tag{4.5}$$

Eq. 4.5 shows that when $e^{-Rt/L} \rightarrow 0$, as t increases, the steady-state current ($I_{ss} = V/R$) is reached. From Eq. 4.5, we now utilise electro-hydraulic analogy to substitute the circuit parameters by hydraulic ones in a pipe network.

The electrical resistance is substituted as a hydraulic resistance $R_h = \Delta p/Q(t)$, where Δp corresponds to the pressure drop observed in a conduit with some flow rate $Q(t)$ (substituting $i(t)$). Similarly, the battery is substituted either by a pumping system or a water head due to gravity – in both cases, a pressure difference Δp is applied to the closed hydraulic system, analogous to a voltage drop. The Inductor is then usually approximated as a water-wheel in a hydraulic system, where the inductance is renamed as inertance (L_{in}). When the system is closed, the pressure difference pushes water into the water-wheel. Initially the water-wheel is not moving, so $Q(t=0)=0$. As the pressure difference keeps providing force, the water wheel starts moving until it reaches a maximum constant rotational speed. When this speed is reached, flow resistance is minimum and $Q_{(t\rightarrow\infty)} = Q_{ss}$, where Q_{ss} is the steady-state flow rate.

With the showcased analogy, Eq. 4.5 can be recast as the momemtum ramp function in analytical form:

$$Q(t) = Q_{ss}(1 - e^{-t/\zeta}), \tag{4.6}$$

where $\zeta = L_{in}/R_h$.

In order to apply Eq. 4.6 in the context of 0D or 2D TRS models, a numerical form of the momemtum ramp needs to be developed. This can be acquired by discretising time $t(n) = n\Delta t$, where n is the number of iterations and Δt the time step:

$$Q(n) = Q_{ss}(1 - e^{-n\Delta t/\zeta}). \tag{4.7}$$

For a time $t_{(n+1)} = (n+1)\Delta t$:

$$Q(n+1) = Q_{ss}(1 - e^{-(n+1)\Delta t/\zeta}), \tag{4.8}$$

where the exponential term can be expanded into

$$Q(n+1) = Q_{ss}(1 - e^{-n\Delta t/\zeta} \times e^{-\Delta t/\zeta}). \tag{4.9}$$

From Eq (4.7), we isolate the exponential term:

$$- e^{-n\Delta t/\zeta} = \frac{Q(n)}{Q_{ss}} - 1. \tag{4.10}$$

Finally, by substituting Eq (4.10) into Eq (4.9), a numerical form of the momentum ramp function is obtained:

$$Q(n+1) = Q_{ss} + (Q(n) - Q_{ss})\, e^{-\Delta t/\zeta}, \tag{4.11}$$

where $Q(n+1)$ represents the flow rate at the next time-step and $Q(n)$ the flow rate at the present time. The term Q_{ss} represents the steady-state flow rate estimate for the hydraulic structure as a function of head difference. As an example: for sluices, Q_{ss} can be attained from the orifice equation (Eq. 3.1), while for turbines in power generation mode, Q_{ss} can be attained from parametrised turbine power charts (e.g. Andritz chart, Section 3.2.2). Furthermore, since in steady-state regime power output relates linearly to turbine flow rate (Eq. 3.8), the numerical form of the momentum ramp function can also be applied to predict the power output evolution:

$$P(n+1) = P_{ss} + (P(n) - P_{ss})\, e^{-\Delta t/\zeta}, \tag{4.12}$$

where $P(n+1)$ represents the power output at the next time-step and $P(n)$ the power output at the present time. The term P_{ss} represents the steady-state power output estimate for the turbine as a function of head, i.e. a power chart estimate.

The numerical form of the momentum ramp function in Eq. 4.11 simplifies the very complex phenomena of opening/closing the turbines and sluices that compose a TRS. Indeed, beyond the hydraulic resistance and inertance of these systems, the opening/closing of turbines and sluices also involve adjusting the pitching of runner blades/guide vanes (for turbines) and aperture of gates (for sluices). Nevertheless, we show in following sections that the developed ramp function can accurately help estimate power output and flow rate evolution for both starting/closing stages of turbines and sluices, given appropriate ζ values.

4.2.2 Turbines

Given power production and water level variations for the ocean and impounded lagoon (for the numbered $E.G$ and $F.G$ turbine modes of operation in Fig 4.4), we can draw interpolated (parametric) curves for the turbine power output as a function of head, with time "t" as hidden parameter:

$$P_{int}(EG_i) = P_{int}(h,t) \quad \text{or} \quad P_{int}(FG_i) = P_{int}(h,t), \tag{4.13}$$

where "i" corresponds to the ith occurrence of $E.G$ and $F.G$ turbine modes. The interpolated $P_{int}(EG_i)$ and $P_{int}(FG_i)$ curves, for every 'i" are shown as solid curves in Figs. 4.10a, 4.10b, 4.10c, 4.10d, 4.10e and Figs. 4.11a, 4.11b, respectively. These Figures showcase power production as a function of head difference, where time evolution is presented by labels over the interpolation.

Considering scenarios where the starting phase for the turbines is available ($P_{int}(EG_2)$, $P_{int}(EG_3)$, $P_{int}(EG_5)$ and $P_{int}(FG_1)$, $P_{int}(FG_2)$), the interpolated results show a two-step process where initially (i) power outputs rapidly increases from a starting non-zero head difference until reaching a plateau (maximum head difference and power output). (ii) Then, from this plateau, water head variations, therefore power generation, slowly decrease with time, until power production is ceased.

We interpret the first process as non-steady acceleration stage where the turbine, initially at rest, is submitted to a starting operational water head H_{start} , accelerating until maximum power generation is achieved. Conversely, the second process is interpreted as a quasi steady-state deceleration stage, where the turbine is submitted to a slowly decreasing water head, taking $\gtrsim 3$ hours to reach H_{min} .

With this interpretation, we take the average of all available quasi steady-state phases for $E.G$ and $F.G$ modes, considering power output to be a pure function of head difference "h", thus obtaining power charts ($P_{EG}(h)$, $P_{FG}(h)$) for both modes of operation. A comparison of our parametrised power charts against interpolated power charts from [67] is shown in Figs. 4.10f and 4.11c, for $E.G$ and $F.G$ modes of operation, respectively. The differences observed between these power charts could be due to a series of factors not available to us, such as the pitching of runner blades/guide vanes. Nevertheless, proceeding with our parametrised power charts enabled a satisfactory 0D simulation of turbines, as will be shown in the following sections.

In order to simulate the smooth power output evolution for both accelerating and decelerating stages of turbine operation, the momentum ramp function, developed in Section 4.2.1, is utilised. A schematic showcasing the result of augmenting the parametrised power chart with the momentum ramp function can be seen in Fig. 4.9.

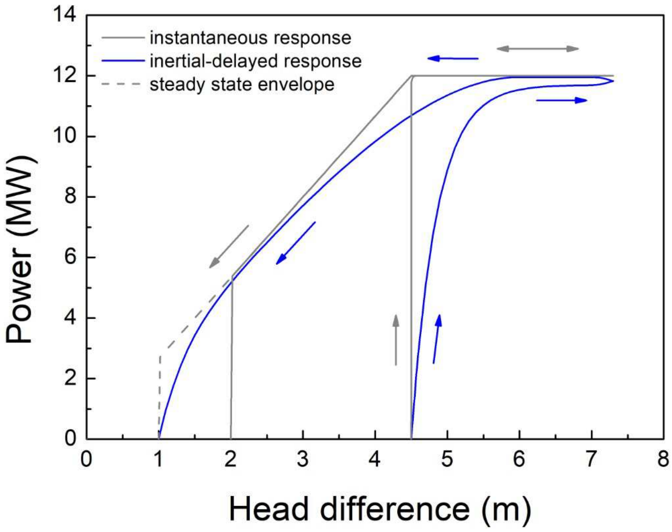

Figure 4.9. The effect of augmenting the parametrised turbine power chart (grey, instantaneous response) with the momentum ramp function is shown in blue, with time evolution displayed with arrows.

For best adjusting the predicted power output (from Eq. 4.12) as a function of water head, best ζ values for accelerating and decelerating stages of $E.G$ and $F.G$ modes of operation (ζ_{AE}, ζ_{AF} and ζ_{DE}, ζ_{DF} respectively) are found by minimising the sum of squared residuals SSR (Eq. 4.16) between each interpolated power chart ($P_{int}(EG_i)$, $P_{int}(FG_i)$) and the parametrised power chart ($P_{EG}(h)$, $P_{FG}(h)$) augmented with the momentum ramp function (Eq. 4.14, 4.15). It is worth noting that, when minimising the residuals, the water head input for both charts is also a function of time.

$$\hat{P}(h,t+1,\zeta_{jE}) = P_{EG}(h) + \left(\hat{P}(h,t,\zeta_{jE}) - P_{EG}(h)\right) e^{-\Delta t/\zeta_{jE}}, \tag{4.14}$$

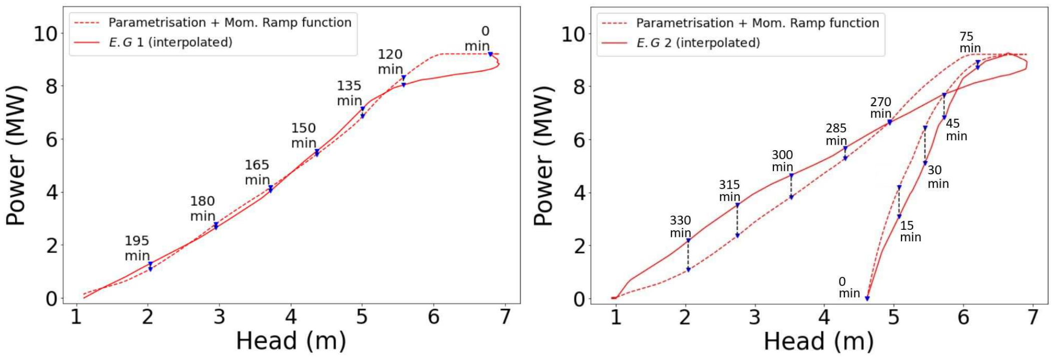

(a) $E.G$ 1 interpolation and parametrisation. **(b)** $E.G$ 2 interpolation and parametrisation.

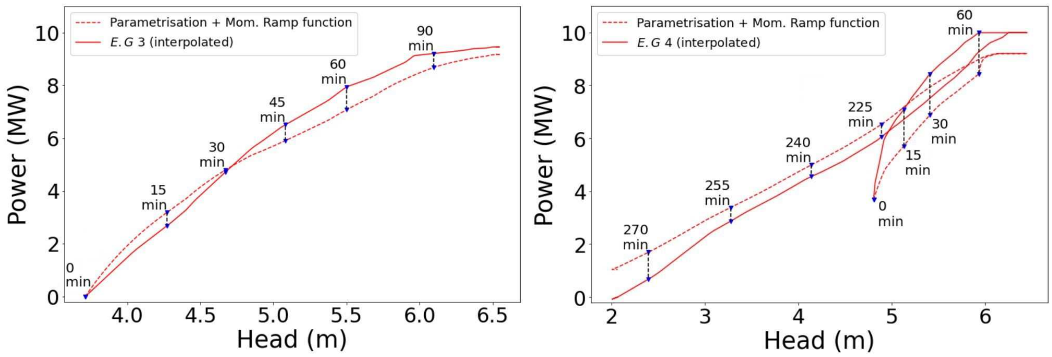

(c) $E.G$ 3 interpolation and parametrisation. **(d)** $E.G$ 4 interpolation and parametrisation.

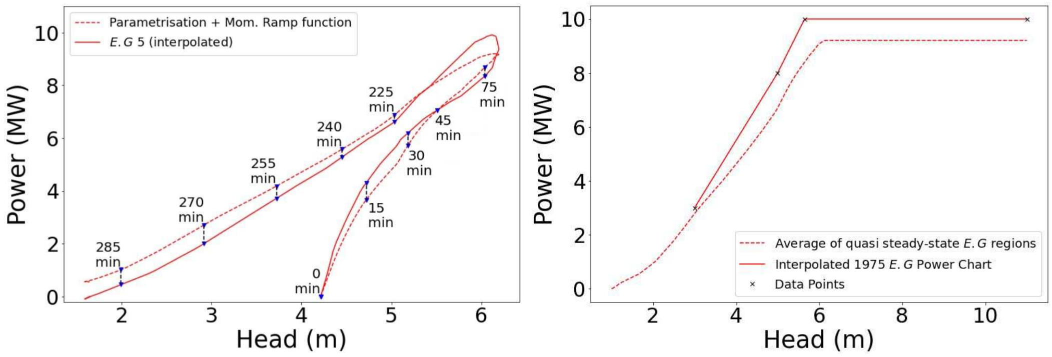

(e) $E.G$ 5 interpolation and parametrisation. **(f)** Parametrised (average of quasi steady-state $E.G$ regions) vs data interpolated power charts.

Figure 4.10. $E.G$ parametrisation verification for La Rance.

for $E.G$, and

$$\hat{P}(h,t+1,\zeta_{jF}) = P_{FG}(h) + \left(\hat{P}(h,t,\zeta_{jF}) - P_{FG}(h)\right) e^{-\Delta t/\zeta_{jF}}, \tag{4.15}$$

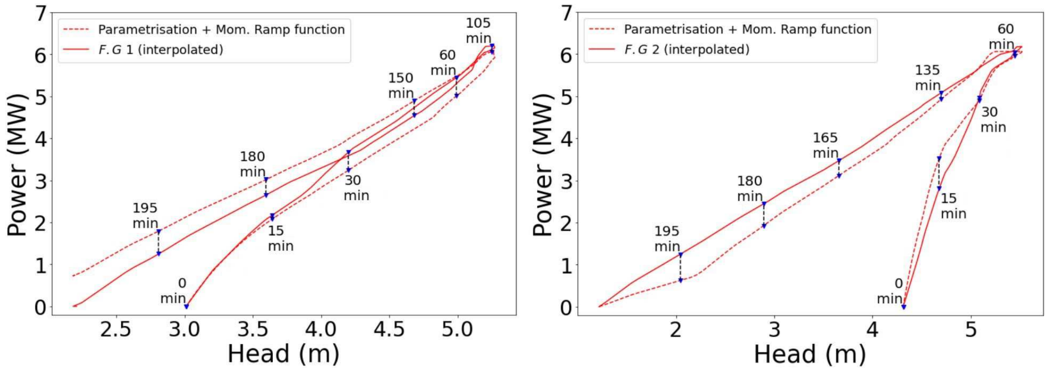

(a) *F.G* 1 interpolation and parametrisation. **(b)** *F.G* 2 interpolation and parametrisation.

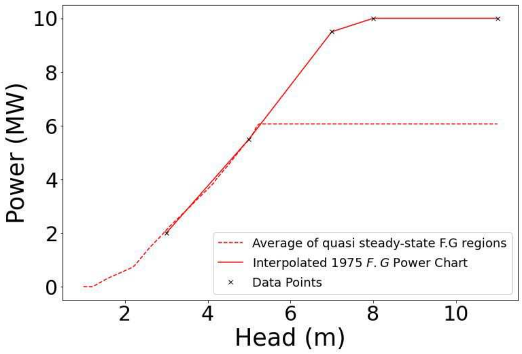

(c) Parametrised (average of quasi steady-state *F.G* regions) vs data interpolated power charts.

Figure 4.11. *F.G* parametrisation verification for La Rance.

for *F.G*, where "$j = A$" if $\hat{P}(t+1) > \hat{P}(t)$, else $j = D$.

$$SSR = \sum_{t=0}^{t=n\Delta t} (P_{int}(h(t),t) - \hat{P}(h(t),t,\zeta))^2. \tag{4.16}$$

Optimum $\zeta_{AE} = 14.2\ min$, $\zeta_{AF} = 11.257\ min$, $\zeta_{DE} = 1.355\ min$ and $\zeta_{DF} = 1.091\ min$ are then obtained, by only considering scenarios where the starting phase for the turbines were available, i.e. EG_2, EG_3, EG_5, FG_1 and FG_2.

Results of the parametrised power chart augmented with the momentum ramp function with optimum ζ values are shown as dashed curves in Figs. 4.10a, 4.10b, 4.10c, 4.10d, 4.10e and Figs. 4.11a, 4.11b. It is worth stressing that these results have been obtained with fixed sets [ζ_{AE}, ζ_{DE}] and [ζ_{AF}, ζ_{DF}] for *E.G* and *F.G* modes of operation, respectively.

4.2.3 Equivalent Lagoon Wetted Area

So far in the literature, and to the best of our knowledge, 0D numerical simulations of TRS have utilised bathymetric data for estimating a “flat” lagoon wetted area for TRS. Although this technique can enable accurate simulation of TRS for “small-scale” systems, such as the SBL, as the impounded area and length increases (like in Severn barrage), this technique starts presenting significant deviations when compared to more realistic 2D models (Section 3.2.4).

In this section, we showcase a simple methodology that can be applied to any coastal region to extract an equivalent lagoon wetted area representation. This methodology was applied in our La Rance case study, since bathymetric data for the barrage were not available. In fact, knowing the real impounded area of La Rance is even more complicated, because, like any estuary, La Rance presents a severe accumulation of silt and sand due to siltation (a natural phenomenon due to sedimentation), amounting to approximately $50,000m^3$ of sediment deposited each year in the estuary [112; 113] – altering La Rance´s bathymetry. In Section 4.2.6 we show that the method proposed here enabled an accurate prediction of lagoon water level evolution when operating turbines (in power generation and pump modes) and sluices, even though the geographical maximum lagoon wetted area spans $22km^2$ (twice that of SBL) and, given its narrow $720m$ barrage, has an approximate length $> 30km$ [114].

As shown in [16], analytical estimates for the available energy in a TRS can utilise the impounded volume of water instead of bathymetric representations. In a similar fashion, we utilise the stored water volume data shown in Table 4.6, to derive an equivalent lagoon wetted area for the La Rance 0D model. In the literature, this equivalent area is attained from the inter-tidal prism volume estimate, which have been used to estimate the impounded water volume of rivers that end in tidal regions [115]. In its simplest form, a constant wetted area Al is used, so the stored volume of water ΔV is proportional to variations in tidal level (tidal range) [115]. We note that this simplified model yielded fairly good results when applied to our La Rance 0D model. However, considering that there are several model approximations that can be used to estimate Al as a function of tidal level variations [116], we proceed our study with a linear approximation for Al. For the sake of simplicity, we use the tidal datum (height z from the lowest equinoctial low tide, defined as $z = 0$), instead of the tidal prism. So that,

$$Al(z) = 2s \times z + Al_0, \tag{4.17}$$

where Al_0 is the equivalent wetted area at $z = 0$ and $2s$ the slope. From this approximation, the stored volume ($\Delta V = V - V_0$) at a tidal datum z, obtained by integrating

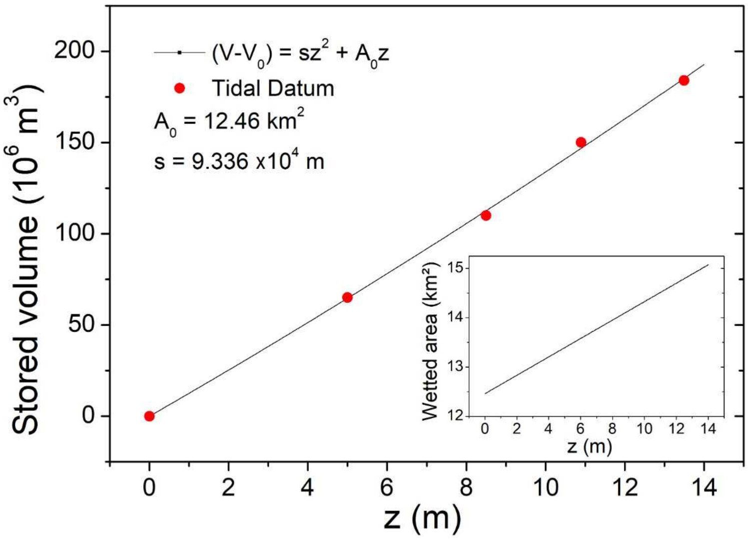

Figure 4.12. Quadratic adjust of the volume of stored water (ΔV), and linear equivalent wetted area estimate (insert) as a function of tidal datum z for the La Rance estuary.

Eq. 4.17 from 0 to z, becomes a second order polynomial:

$$\Delta V = s \times z^2 + Al_0 \times z. \tag{4.18}$$

Now the coefficients s and Al_0 can be estimated with a least square method, utilising the point measurements $(z_i, \Delta V_i)$ from Table 4.6, as a target. The obtained quadratic fit for ΔV is shown in Fig. 4.12, together found values for s and Al_0. With these parameters, the wetted area given by Eq. 4.17 is also plotted and shown as an insert in Fig. 4.12. The obtained linear $Al(z)$ gave improved results when applied to the 0D model of La Rance barrage (Eq. 3.12).

4.2.4 Sluicing

For parametrising the sluicing stages shown in Fig. 4.4a and 4.4b, we need to define appropriate (i) ζ values for our momentum ramp function (Eq. 4.11) and (ii) discharge coefficient (C_d) for the orifice equation (Eq. 3.1), responsible for estimating flow rates from the hydraulic structures.

In contrast with turbines (in power generation mode) that have accelerating and decelerating stages (Section 4.2.2), sluices allow for the free passage of water with minimum resistance, independent if head differences are increasing or decreasing. Due to this symmetric behaviour, we assume a constant ζ_s value for our momentum ramp func-

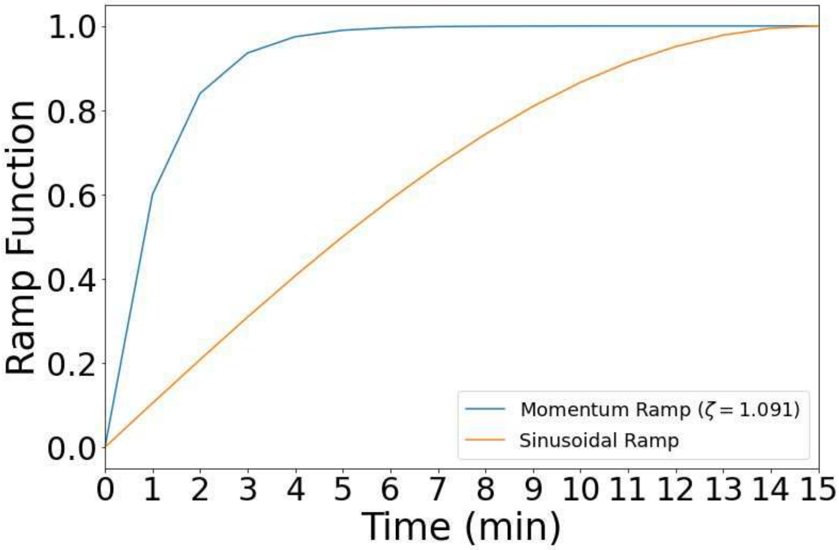

Figure 4.13. Comparison between momentum and sinusoidal ramp functions.

tion estimate. Since we do not have previous data to fit a best ζ_s value for sluices, we adopt $\zeta_s = 1.091$, which guarantees a precision of 10^6 for the complete opening/closing of hydraulic structures in a 15 min time interval [105]. A comparison of this ramp function, with $\Delta t = 1\ min$, against conventional sinusoidal ramp functions adopted in the literature is shown in Fig. 4.13. Although the momentum ramp converges more quickly than the sinusoidal ramp, we show in this section that the agreements between our model and measured data are satisfactory.

Beyond sluice gates, the sluicing stages shown in Fig. 4.4a and 4.4b, need also to account for turbine flow rates operating in “idling” mode. For simplification purposes, we assume that turbine operation in idling mode have a similar behaviour to sluices, since there is no electrical energy conversion. Therefore, the same $\zeta_s = 1.091$ will be used for turbines in “idling” mode. This similar behaviour is also supported by the literature, since flow rate estimates for sluices and idling turbines utilise the same “orifice” equation (Section 3.2.1).

For best fitting Eq. 3.1 to La Rance, appropriate C_d and C_{dt}, for sluices and idling turbines, respectively, need to be found. For sluice gates, we bound the C_d estimate between two experimental values from the literature. As a lower bound, we have $C_d = 1$, from experimental results of a large sluice gate experiment by [75]. As an upper bound, we utilise measurements from [67], which indicate that, at its maximum, sluice gates from La Rance provide a flow rate of $9600m^3/s$. Since the total sluice gate area for La Rance is $900m^2$, we attain an upper bound of $C_d = 1.077$.

Similarly, idling turbines flow rate is bound by experimental measurements, at La Rance, from [109]. In this work, flow rate measurements from a turbine unit under a fixed $4m$ head was recorded in the range $182.2m^3/s$ to $280m^3/s$. From [67; 109; 75] data, upper and lower bound discharge coefficients for turbines and sluices are defined as:

$$[1 \leq C_d \leq 1.077] \text{ and } [.91 \leq C_{dt} \leq 1.4]. \tag{4.19}$$

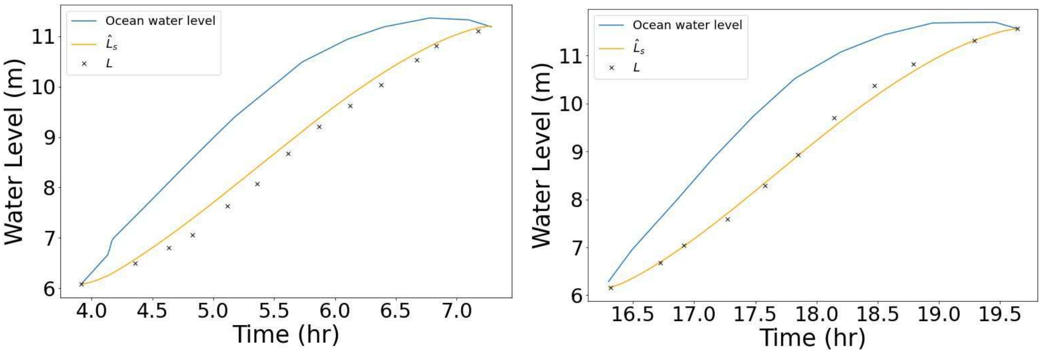

(a) Sluicing Region 1. **(b)** Sluicing Region 2.

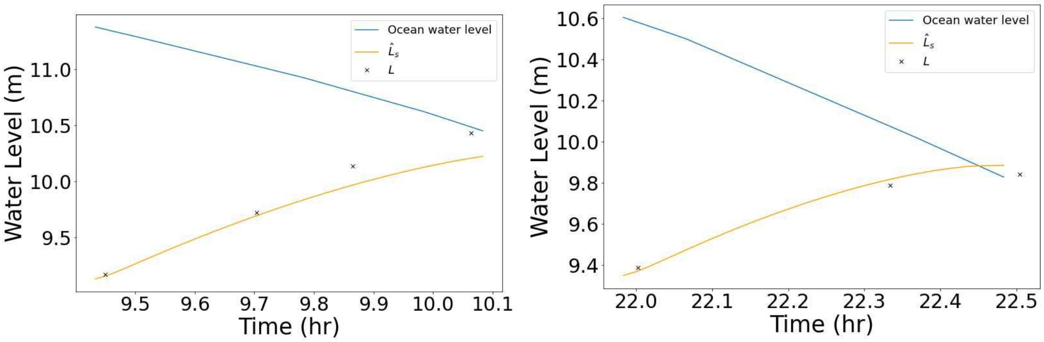

(c) Sluicing Region 3. **(d)** Sluicing Region 4.

Figure 4.14. Comparison between predicted ($\hat{L}_S$) and measured (L) lagoon water level variations during sluicing stage for La Rance.

For choosing best C_d and C_{dt} values within bounds (Eq. 4.19), we utilise the 0D model for La Rance (Eq. 3.12), with the equivalent lagoon wetted area derived in the previous section, to predict lagoon water level variations during each sluicing stage Sl. The predicted levels $\hat{L}_S$ are compared with measured data L and a time-normalised sum of squared residuals ($NSSR$) is calculated:

$$NSSR = \sum_{S=1}^{S=4} \frac{(\hat{L}_S - L)^2}{t_S} \tag{4.20}$$

The best C_d and C_{dt} that minimise $NSSR$ for all sluicing stages are $C_d = 1.017$ and $C_{dt} = .967$. However, since adopting $C_d = C_{dt} = 1$ does not change results significantly, for the sake of simplification, we resume our work with a discharge coefficient of "1", for both sluices and idling turbines. Comparison between $\hat{L}_S$ and L results, with $C_d = C_{dt} = 1$, are shown in Fig. 4.14.

4.2.5 Pumps

As shown in Section 3.2.2, current idealised pump models for TRS (Eq. 3.6) fail to explain expected behaviours such as maximum pump shutoff head h_s (when $Q_p = 0$), maximum pump flow rates Q_M (for $h_p = 0$) and pump operation under positive head scenarios – all observed in La Rance measurements [67].

For contemplating these behaviours, we utilise pump affinity laws in order to derive a general formulation for the pump flow rate as a function of head difference and applied power input: $Q_p = Q_p(h_p, P_{in})$. The affinity laws express the mathematical relationships between several variables involved in the performance of kinetic pumps (centrifugal or axial). These laws show that, under dynamically similar conditions, dimensionless parameters remain constant. They are useful for predicting pump performance changes when varying either (i) pump operational speed (ii) pump impeller diameter [117]. Since bulb turbines operating in reverse have a similar behaviour to axial-flow pumps [118; 119; 120; 121; 122], we continue this section with the assumption that affinity laws can be applied to bulb turbines in pump mode.

Each affinity law postulate can be expressed as a set of three equations, where pump flow rate Q_p, negative head h_p and input power P_{in} are estimated for a point "$p2$" using a known point "$p1$" as reference. Then, following the affinity laws, values for "$p2$" are obtained as a function of either pump rotation (N) or pump impeller diameter (D), according to the first or second postulate from the affinity laws, respectively. From the first postulate:

$$\frac{Q_{p1}}{Q_{p2}} = \frac{N_1}{N_2}; \quad \frac{h_{p1}}{h_{p2}} = \left(\frac{N_1}{N_2}\right)^2; \quad \frac{P_{in1}}{P_{in2}} = \left(\frac{N_1}{N_2}\right)^3. \tag{4.21}$$

For the relationship shown in Eq. 4.21 to be true, pump efficiency must remain relatively constant as we move from "$p1$" to "$p2$", for a fixed diameter pump [117]. Utilising the first postulate of pump affinity laws (and assuming a near constant efficiency for any change in rotational speed) is sufficient for estimating a general pump equation for La Rance. The first postulate can also be manipulated, so that Q_p and

h_p are obtained as a function of power input variation:

$$Q_{p1} = Q_{p2}\left(\frac{P_{in1}}{P_{in2}}\right)^{1/3}, \tag{4.22}$$

$$h_{p1} = h_{p2}\left(\frac{P_{in1}}{P_{in2}}\right)^{2/3}. \tag{4.23}$$

For both $E.P$ and $F.P$ operational modes, the second order approximation for flow rate $Q_p(h_p, P_{in})$ as a function of head for a fixed (maximum) power input of $P_{in} = 6MW$, shown in Fig. 4.6, is utilised. For simplification purposes, $\{P_{in} = 6MW\} = P_6$. Also, since pump operational head h_p depends on power input (Eq. 4.23), we have for the quadratic approximation $Q_p(h_p(P_6), P_{in} = P_6) = Q_p(h_p(P_6))$, so that:

$$Q_p(h_p(P_6)) = a{h_p}^2(P_6) + bh_p(P_6) + Q_M(P_6),\ \text{where } h_p(P_6) <= 0. \tag{4.24}$$

Eq. 4.24 expresses a characteristic pump curve, where $Q_M(P_6)$ is the maximum pump flow rate expected, for $h_p = 0$ and $P_{in} = 6MW$. Knowing that Q_p is zero at the shutoff head (h_s), the quadratic approximation can be expressed in terms of its roots, yielding a convenient form:

$$Q_p(h_p(P_6)) = a[h_p(P_6) - h_s(P_6)][h_p(P_6) - Q_M(P_6)/(ah_s(P_6))]. \tag{4.25}$$

For deriving a characteristic pump curve for any power input (P_{in}) and pump head ($h_p(P_{in})$), we expand all terms in Eq. 4.25, by utilising Eq. 4.23 and Eq. 4.22, thus obtaining Eq. 4.26. For simplification purposes, the maximum allowed pump power (P_6) over pump power input (P_{in}) is reduced to:$(P_6/P_{in}) = R$.

$$\begin{aligned} Q_p(h_p(P_6)) &= Q_p(h_p(P_{in}))R^{1/3}, \\ Q_M(P_6) &= Q_M(P_{in})R^{1/3}, \\ h_p(P_6) &= h_p(P_{in})R^{2/3}, \\ h_s(P_6) &= h_s(P_{in})R^{2/3}. \end{aligned} \tag{4.26}$$

Performing the substitution of terms from Eq. 4.26 into Eq. 4.25, returns

$$Q_p(h_p, P_{in}) = aR[h_p(P_{in}) - h_s(P_{in})][h_p(P_{in}) - Q_M(P_{in})/(aRh_s(P_{in}))]. \tag{4.27}$$

Eq. 4.27 allows for estimating pump flow rates for any given h_p and P_{in}, including regions of pump shutoff head h_s and maximum pump flow rate Q_M. Finally, in order to estimate pump flow rates for positive head scenarios, we assume that the maximum Q_M attained for $h_p = 0$ is summed with the gravitational flow rate estimate Q_g derived from the orifice equation for turbines (Eq. 3.1). The maximum flow rate allowed during pumping (upper bound "ub") is $ub = 280m^3/s$, as defined in Section 4.2. Examples of predicted pump flow rates, for $E.P$ and $F.P$ modes of operation are shown in Fig. 4.15.

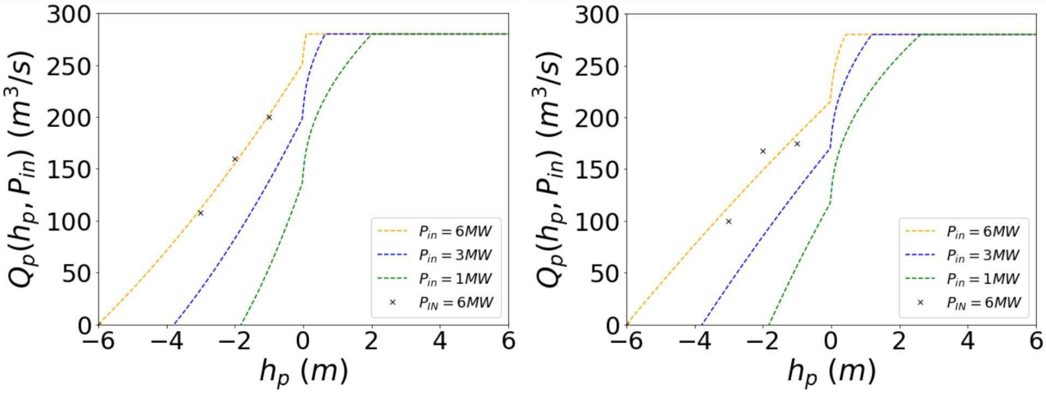

(a) $Q_p(h_p, P_{in})$ estimate for $E.P$. **(b)** $Q_p(h_p, P_{in})$ estimate for $F.P$.

Figure 4.15. La Rance $Q_p(h_p, P_{in})$ estimate for $E.P$ and $F.P$, for $P_{in} = 6, 3$ and 1MW.

For verifying the quality of our pump $Q_p(h_p, P_{in})$ estimate (Eq. 4.27), we utilise the 0D model for La Rance (Eq. 3.12), with the equivalent lagoon wetted area derived in Section 4.2.3, to predict lagoon water level variations $\hat{L}_P$ during each $E.P$ and $F.P$ pumping stage. The predicted $\hat{L}_P$ values, for a given (measured) power input P_{in} and varying ocean levels, are compared with measured lagoon water levels L and shown in Fig. 4.16, along with the input parameters.

As shown in Fig. 4.16, the derived pump model for La Rance presented good agreement of results when predicting lagoon water level variations against measured data. In order to generalise the developed pump equation for bulb turbines of various diameters and power capacity, we now make use of the second postulate of the Affinity Laws, that states:

$$\frac{Q_{p1}}{Q_{p2}} = \frac{D_1}{D_2}; \quad \frac{h_{p1}}{h_{p2}} = \left(\frac{D_1}{D_2}\right)^2; \quad \frac{P_{in1}}{P_{in2}} = \left(\frac{D_1}{D_2}\right)^3. \tag{4.28}$$

Such as in Eq. 4.21, for the relationship above to be true, pump efficiency must remain relatively constant as we move from "$p1$" to "$p2$" [117]. The required D_1 and D_2 pump diameters, for La Rance and SBL, are shown in Tables 4.4 and 4.8, respectively.

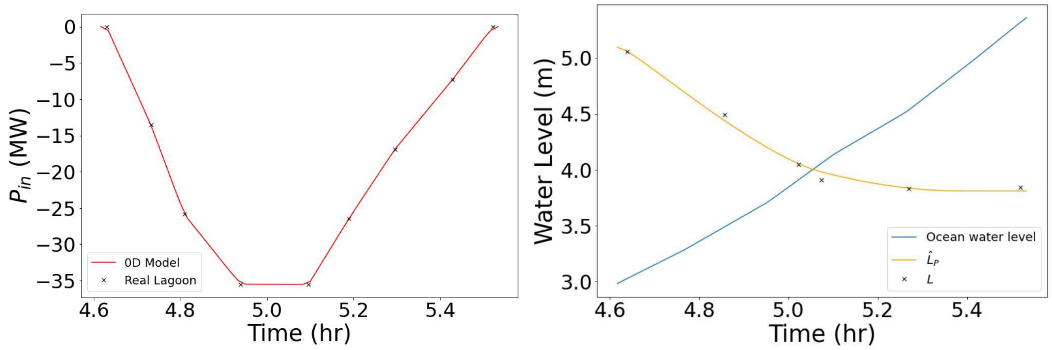

(a) P_{in} for $E.P$ 1.

(b) $\hat{L}_P$ and L for $E.P$ 1.

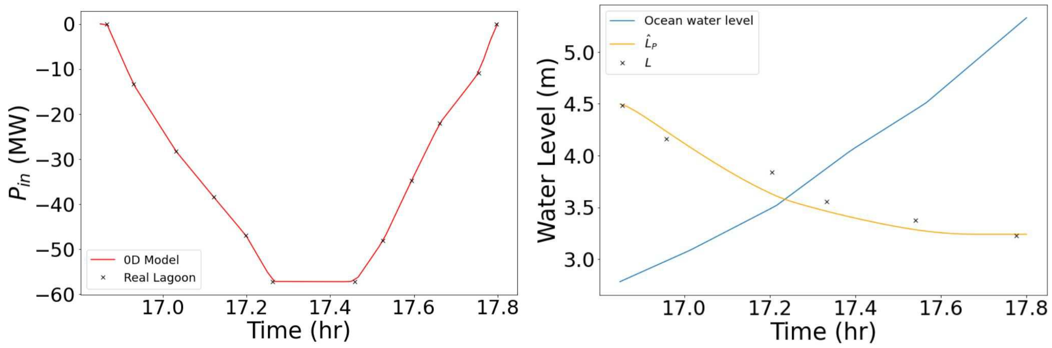

(c) P_{in} for $E.P$ 2.

(d) $\hat{L}_P$ and L for $E.P$ 2.

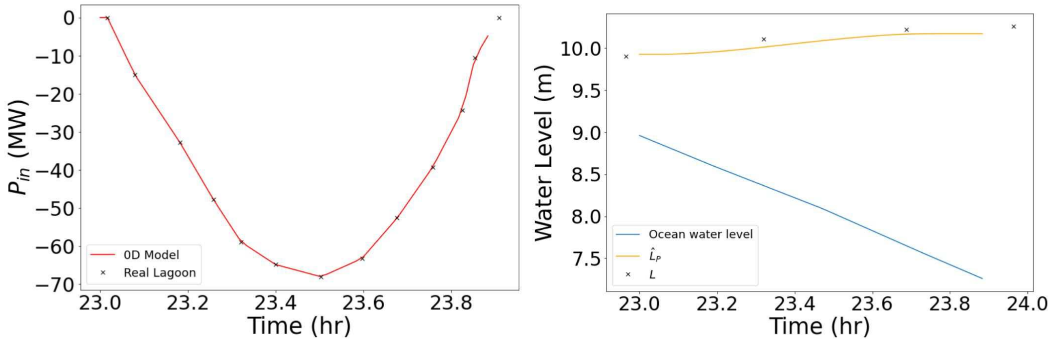

(e) P_{in} for $F.P.$

(f) $\hat{L}_P$ and L for $F.P.$

Figure 4.16. Comparison between predicted $\hat{L}_P$ (yellow solid curves) and measured L (×) lagoon water level variations during pumping stages, along with measured ocean water levels and P_{in}, used as inputs for the 0D La Rance model.

The second postulate allows for estimating pump characteristic curves for the SBL project, by considering the following assumptions:

- The bulb unit in SBL have similar efficiency as La Rance, when under the same head difference.

Table 4.7. Estimated $E.P$ and $F.P$ flow rates for SBL, utilising the second postulate from the affinity laws and a fixed power input $P_{in} = 12MW$.

$Q_p(h_p, P_{in} = 12MW)$	Against $h_p(m) < 0$			
(m^3/s)	-9.52	-4.76	-3.17	-1.59
Q_pEbb	0	136.07	201.59	251.98
Q_pFlood	0	125.99	211.67	220.49

- The maximum power input allowed during pumping for the SBL project equals 60% of the turbine capacity, such as La Rance [67] (i.e. $12MW$).

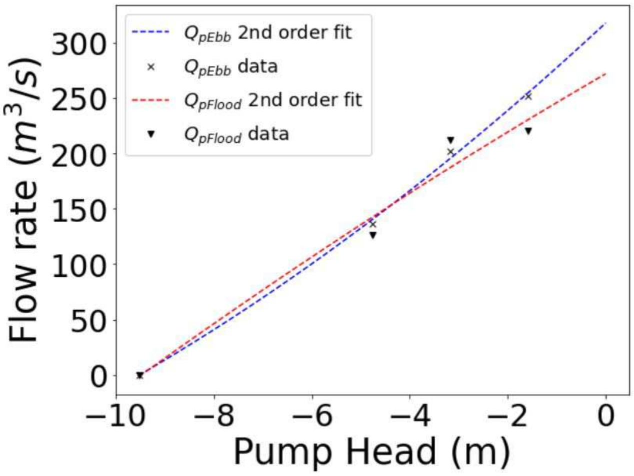

Figure 4.17. SBL $E.P$ and $F.P$ flow rate estimate (from the second postulate of the affinity laws) as a function of water head.
($Q_{pEbb} = 0.8{h_p}^2 + 41.1h_p + 317.8$) and ($Q_{pFlood} = -.3{h_p}^2 + 25.7h_p + 271.9$).

Following these assumptions, we utilise experimental measurements of pump flow rate $(Q_pEbb(h_p), Q_pFlood(h_p))$ for La Rance (Fig. 4.6) in conjunction with the second postulate, to estimate $(Q_pEbb(h_p), Q_pFlood(h_p))$ point values for SBL (Table 4.7). With these values in hand, quadratic expressions for pump flow rate as a function of head (Fig. 4.17), and the full pump equation as a function of P_{in} and h_p (Eq. 4.27) are obtained following the same steps as for La Rance. Furthermore, the maximum flow rate (upper bound ub) allowed during pumping is set to $ub = 487m^3/s$. This value was obtained by following the Andritz´s chart prediction (Section 3.2.2.1), for a $20MW$ turbine with SBL specifications (Table 4.8). Examples of characteristic curves estimated for the SBL are shown in Fig. 4.18.

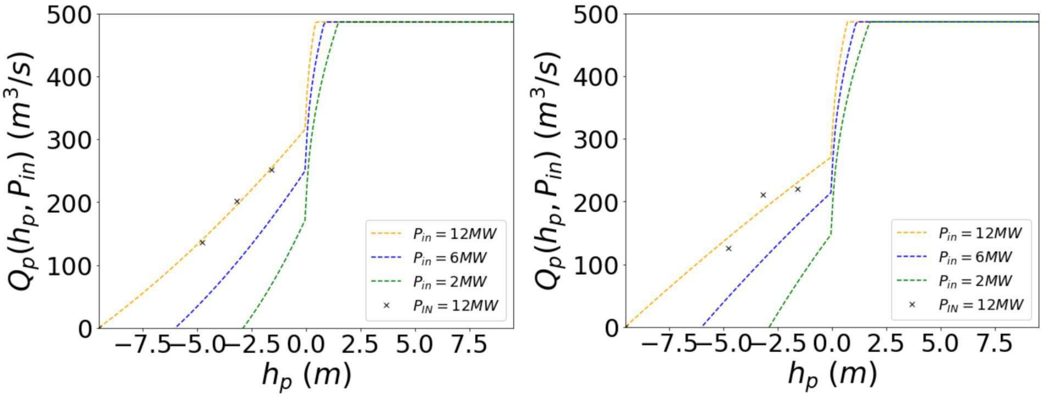

(a) $Q_p(h_p, P_{in})$ estimate for $E.P.$ **(b)** $Q_p(h_p, P_{in})$ estimate for $F.P.$

Figure 4.18. SBL $Q_p(h_p, P_{in})$ estimate for $E.P$ and $F.P$, for $P_{in} = 12, 6$ and 2MW.

4.2.6 Validation of the Parametrised 0D La Rance Model

With the developed parametric models for turbines (in power generation and pump modes), sluices, equivalent lagoon area and momentum ramp function, we can now verify the accuracy of our 0D La Rance model in predicting lagoon water level variations and power output.

Utilising Fig. 4.4 as reference, we set initial predictions for lagoon water levels $\hat{L}$ and turbine output power $\hat{P}(h, t, \zeta)$ to the same initial measured values (L and P, respectively) at La Rance. Henceforth, we vary the operational mode for the hydraulic structures following the same timing as in Fig. 4.4a and 4.4b. When operating turbines in pump mode, the measured P_{in} is applied to our pump model, so that pump flow rates can be predicted. A comparison of the 0D La Rance model predictions with measured data is shown in Fig. 4.19.

As noted in Section 4.2, sluice operation is expected at the end of power generation stages for the TWP scheme (Figs. 4.19b, 4.19d and 4.19f). Since the timing for starting sluice operation was not provided by [67], the showcased results for the TWP scheme simulation assume a best fit between predicted and measured lagoon water levels. Nevertheless, the agreement between predicted and measured power outputs for both E_oG and TWP schemes is shown to be satisfactory. Indeed, by integrating predicted and measured power, energy deviation is only 4.7% and 2.1%, for E_oG and TWP schemes comparisons, respectively.

While results in (Figs. 4.19b, 4.19d and 4.19f) aim to validate the parametrisation techniques applied in reverse engineering La Rance into a 0D model, an optimal and comparable strategy for the control sequence of turbines and sluices is still required. In Section 5.2, as a second validation step, we show that our trained DRL-Agent is

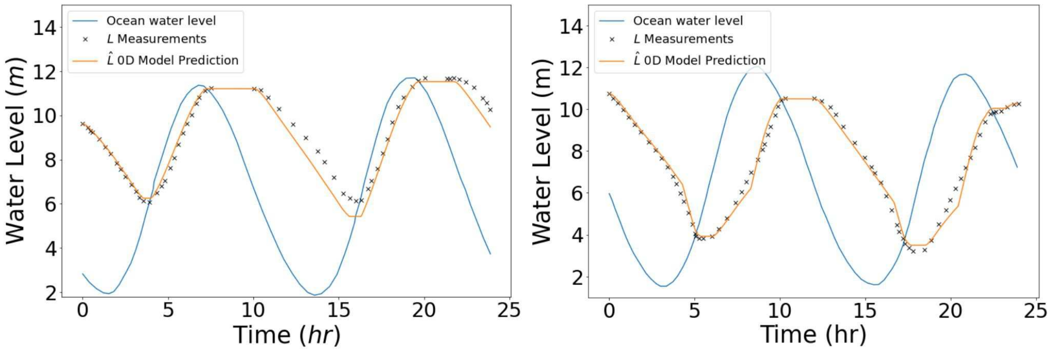

(a) $\hat{L}$ vs L comparison for E_oG. **(b)** $\hat{L}$ vs L comparison for TWP.

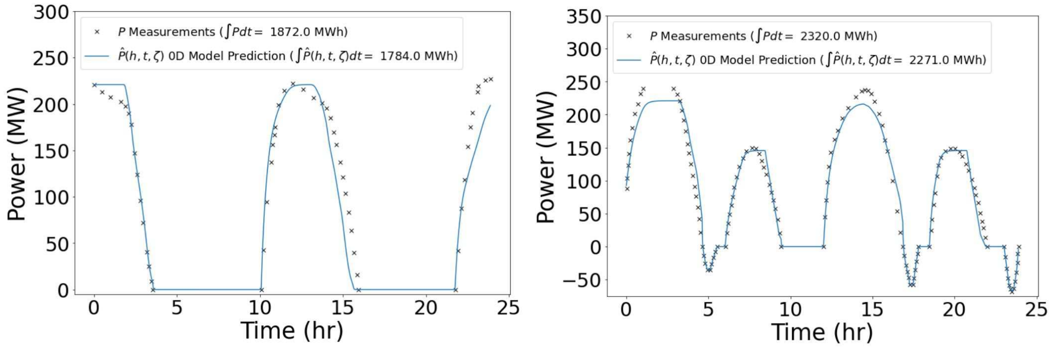

(c) $\hat{P}(h,t,\zeta)$ vs P comparison for E_oG. **(d)** $\hat{P}(h,t,\zeta)$ vs P comparison for TWP.

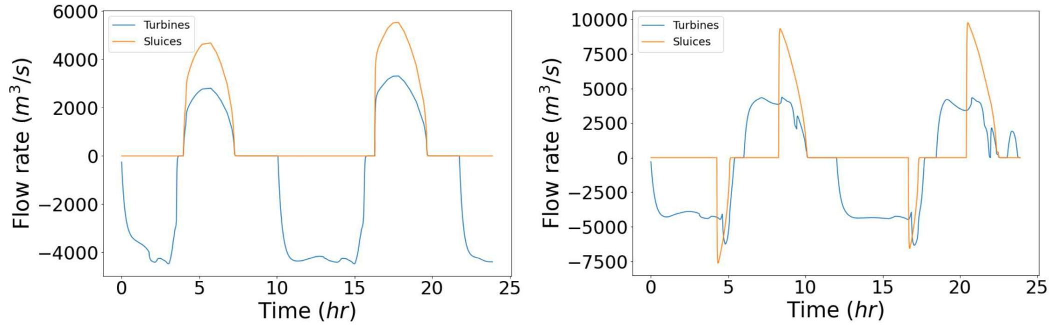

(e) Turbine and sluice flow rates estimates for E_oG. **(f)** Turbine and sluice flow rate estimates for TWP.

Figure 4.19. Validation of 0D La Rance model predictions against site measurements from [67].

capable of operating La Rance with such strategy.

4.3 TRS Numerical Implementation

In this Section, the pseudocode implementations developed for simulating La Rance and SBL (in 0D and 2D representations, when applicable), will be presented and discussed.

4.3.1 0D Modelling

For simulating TRS hydraulic structures in our 0D model representations for La Rance and SBL, algorithm depictions for turbines (in power generation, idling and pump modes) and sluices were developed. For the SBL, the pseudocodes for sluices and turbines (in power generation and idling modes) were inspired from the literature (reviewed in Section 3.2), being showcased in Algorithms 1 and 2. In order to simulate turbine pump modes, the developed method based on the affinity laws (Section 4.2.5) is showcased in pseudocode format (Algorithm 3).

Algorithm 1 Turbine Power and Flow Rate – Andritz representation.

```
Require: $g, \rho, n_t, S_p, D, T_o, C_E, P_C$, as fixed parameters.
function PowerGen($h, T_{op}$)
    if $|h| >= 1$ and $T_{op} ==$ "$powerGen$" then // Assuming $H_{mt} = 1$
        $n_{11} \leftarrow Sp \cdot D/\sqrt{|h|}$
        if $n_{11} <= 255$ then
            $Q_{11} \leftarrow 0.017 \cdot n_{11} + 0.49$
        else if $n_{11} > 255$ then
            $Q_{11} \leftarrow 4.75$
        $Eff = (-0.0019 \cdot n_{11} + 1.2461) \cdot CE$
        if $T_o == 1$ then // Flood Oriented Turbine
            if $h < 0$ then // $E.G$
                $Eff \leftarrow .90 \cdot Eff$
        else if $T_o == 2$ then // Ebb Orientation
            if $h > 0$ then // $F.G$
                $Eff \leftarrow .90 \cdot Eff$
        $Q_t \leftarrow sign(h) \cdot Q_{11} \cdot D^2\sqrt{|h|}$ // Andritz chart
        $P \leftarrow |\rho \cdot g \cdot h \cdot Q_t| \cdot Eff$
        if $P > P_C$ then
            $P \leftarrow P_C$ // $P_C = 20MW$ for the Swansea Lagoon
            $Q_t \leftarrow sign(h) \cdot P/(\rho \cdot g \cdot |h| \cdot Eff)$
    else
        $P \leftarrow 0$
        $Q_t \leftarrow 0$
return $n_t \cdot P, n_t \cdot Q_t$
end function
```

For simulating La Rance's turbines in power generation mode, the pseudocode shown in Algorithm 4, developed with parametric models (Section 4.2.2), is utilised. Similarly to the SBL, La Rance also makes use of Algorithms 2 and 3 to simulate sluices and turbines in idling and pump modes.

The pseudocode enabling the 0D model simulation for SBL and La Rance is shown in Algorithm 5 – being called every time iteration. The developed 0D model has also been modified from what is usually observed in the literature, accommodating the novel momentum ramp function (Section 4.2.1).

Algorithm 2 Idling Turbines and Sluice Gates Flow Rate Estimate – Orifice Equation.

```
Require: $g, \rho, A_T, A_S, n_t, C_d, C_{dt}$, as fixed parameters.
function IdlingTurbines($h, T_{op}$)
    if $T_{op} ==$ “$sluicing''$” then
        $Q_t \leftarrow C_{dt} \cdot Sign(h) \cdot A_T \cdot \sqrt{2 \cdot g \cdot |h|}$
    else
        $Q_t \leftarrow 0$
return $n_t \cdot Q_t$
end function
function Sluices($h, S_{op}$)
    if $S_{op} ==$ “$sluicing''$” then
        $Q_s \leftarrow C_d \cdot Sign(h) \cdot A_S \cdot \sqrt{2 \cdot g \cdot |h|}$
    else
        $Q_s \leftarrow 0$
return $Q_s$
end function
```

Algorithm 3 Pump Flow Rate.

```
Require: $g, \rho, A_T, n_t, P_{inMax}, h_s, ub, a_E, Q_{MEbb}, a_F, Q_{MFlood}$, as fixed parameters
// $h_s = 6m$ for La Rance and $h_s = 9.52m$ for SBL (Ebb and Flood)
// Maximum pump power input $P_{inMax} = 6MW$ for La Rance and $= 12MW$ for SBL
// $ub = 280$ for La Rance and $ub = 487$ for SBL
// For $a_E, Q_{MEbb}, a_F$ and $Q_{MFlood}$ values for La Rance, see Fig.4.6
// For $a_E, Q_{MEbb}, a_F$ and $Q_{MFlood}$ values for SBL, see Fig.4.17
function PumpFlowRate($h, T_{op}, P_{in}, P_o$) // $P_o$ is the pumping orientation
    if $T_{op} ==$ “$pumping''$” then
        $h_s = h_s \cdot (P_{in}/P_{inMax})^{2/3}$ // Re-estimate $h_s$ for a given $P_{in}$
        $R = (P_{inMax}/P_{in})$
        if $P_o ==$ “$Ebb''$” then
            $h \leftarrow -h$ // Since $h = ocean - lagoon$ water levels
            $a = a_E; Q_M = Q_{MEbb}$
        else if $P_o ==$ “$Flood''$” then
            $a = a_F; Q_M = Q_{MFlood}$
        if $h \geq 0$ then // Pumping at positive head
            $Q_g \leftarrow C_{dt} \cdot A_T \sqrt{2 \cdot g \cdot h}$
        else
            $Q_g \leftarrow 0$
        if $h \leq h_s$ then // Pump does not work
            $Q_p \leftarrow 0$
        else if $h < 0$ and $h > h_s$ then // Pumping at negative head
            $Q_p \leftarrow a \cdot R \cdot (h - h_s) \cdot (h - Q_M \cdot (P_{in}/P_{inMax})^{1/3}/(a \cdot R \cdot h_s))$
        else if $h \geq 0$ then // Pumping at positive head
            $Q_p \leftarrow Q_M \cdot (P_{in}/P_{inMax})^{1/3} + Q_g$
        if $|Q_p > ub|$ then
            $Q_p \leftarrow Sign(Q_p) \cdot ub$
    else
        $Q_p = 0$
return $n_t \cdot Q_p$
end function
```

Algorithm 4 Turbine Power and Flow Rate – La Rance Parametrisation.

```
Require: $g, \rho, n_t, ub, P_{EG}, P_{FG}, Eff_{E.G}, Eff_{F.G}$, as fixed parameters
function LaRancePowerGen($h, T_{op}$)
    if $|h| >= 1$ and $T_{op} ==$ "$powerGen$" then // Assuming $H_{mt} = 1$
        if $h > 0$ then // $F.G$
            $P \leftarrow P_{FG}(h)$
            $Q_t \leftarrow P/|\rho \cdot g \cdot h \cdot Eff_{F.G}(h)|$
            if $Q_t > ub$ then //
                $Q_t \leftarrow ub$
                $P \leftarrow \rho \cdot g \cdot h \cdot Q_t \cdot Eff_{F.G}(h)$
        else if $h \leq 0$ then // $E.G$
            $P \leftarrow P_{EG}(|h|)$
            $Q_t \leftarrow P/|\rho \cdot g \cdot h \cdot Eff_{E.G}(|h|)$
            if $Q_t < -ub$ then //
                $Q_t \leftarrow -ub$
                $P \leftarrow |\rho \cdot g \cdot h \cdot Q_t \cdot Eff_{E.G}(|h|)|$
    else
        $P \leftarrow 0$
        $Q_t \leftarrow 0$
return $n_t \cdot P, n_t \cdot Q_t$
end function
```

The required fixed parameters for running the pseudocodes, for SBL and La Rance, are presented in Sections 3.2 and 4.2. Beyond these, we add the parameters A_T, T_o, T_{op}, S_{op}, and P_o, which stand for turbine area, orientation, operation, sluice operation and pumping orientation, respectively. The addition of T_{op}, S_{op}, P_{in} and P_o into Algorithm 5 allows for easier control when coupling the 0D model with any of the developed optimisation routines.

4.3.2 2D Modelling

In order to include the effects of sluices and turbines for 2D TRS simulation, a barrage boundary is implemented into DG-ADCIRC. This barrage model was created based on a culvert model already present in DG-ADCIRC [34; 35], which adds an internal barrier into the finite element mesh. The internal barrier acts as a reflective boundary throughout its length (zero perpendicular velocity), except at the node pairs where hydraulic structures are simulated (Fig. 4.20a). The flow through these node pairs is calculated from the water head developed between "front" and "back" sides, instead of the utilising the numerical model's shallow water equations (Eq. 2.1, 2.2 and 2.3).

For verification purposes, the new 2D barrage boundary is implemented in the rectangular domain shown in Fig. 4.20b, where an internal barrier connects the node pairs between Planes 1 and 2. In this case study, 15 node pairs are selected to simulate sluice gates, following Algorithm 2, where flow is allowed to pass when the head dif-

Algorithm 5 0D TWP scheme TRS Model.

```
Require: $g, \rho, A_T, A_S, n_t, C_d, C_{dt}$, as fixed parameters. // $j$ is the time iteration index
function 0DTRSControl($Oj, Lj, P_j, Q_tj, Q_sj, TRS$) // Reveive variable values for the $j-th$ iteration
    $h \leftarrow Oj - Lj$ // where, $O$ is the ocean water level
    if $turbineIsReady$ or $sluiceIsReady$ then
        $T_{op}, S_{op}, P_{in}, P_o \leftarrow RequestDecision()$ // Control Policy
    if $TRS ==$ “$LaRance$” then
        $P, Q_t \leftarrow LaRancePowerGen(h, T_{op})$
        $Q_{ti} \leftarrow IdlingTurbines(h, T_{op})$
        $P_{in}, Q_p \leftarrow PumpFlowRate(h, T_{op}, P_{in}, P_o)$
// Steady-state (head-only dependent) estimates for flow rates and power
        $Q_s(h) \leftarrow Sluices(h, T_{op})$
        $P(h) \leftarrow P + P_{in}$
        $Q_t(h) \leftarrow Q_t + Q_{ti} + Q_p$
        $PowDev \leftarrow Sign(Pj_N - Pj)$ // Utilise $\zeta_A$ and $\zeta_D$ appropriately
        if $T_{op} ==$ “$powerGen$” and $h < 1$ then // $E.G$
            if $PowDev > 0$ then // Turbine is accelerating
                $\zeta \leftarrow \zeta_{AE}$
            else if $PowDev < 0$ then // Turbine is decelerating
                $\zeta \leftarrow \zeta_{DE}$
        else if $T_{op} ==$ “$powerGen$” and $h \geq 1$ then // $F.G$
            if $PowDev > 0$ then // Turbine is accelerating
                $\zeta \leftarrow \zeta_{AF}$
            else if $PowDev < 0$ then // Turbine is decelerating
                $\zeta \leftarrow \zeta_{DF}$
        else// Idling, Sluicing or Pump modes
            $\zeta \leftarrow 1.091$
    else if $TRS ==$ “$SBL$” then
        $P, Q_t \leftarrow PowerGen(h, T_{op})$
        $Q_{ti} \leftarrow IdlingTurbines(h, T_{op})$
        $P_{in}, Q_p \leftarrow PumpFlowRate(h, T_{op}, P_{in}, P_o)$ // Remove for $TW$ scheme simulation
// Steady-state (head-only dependent) estimates for flow rates and power
        $Q_s(h) \leftarrow Sluices(h, T_{op})$
        $P(h) \leftarrow P + P_{in}$
        $Q_t(h) \leftarrow Q_t + Q_{ti} + Q_p$
        $\zeta \leftarrow 1.091$
// Momentum ramp estimates for flow rate and power for the next iteration $j_N$
    $Q_sj_N \leftarrow Q_s(h) + (Q_sj - Q_s(h)) \cdot e^{-\Delta t/\zeta}$
    $Pj_N \leftarrow P(h) + (Pj - P(h)) \cdot e^{-\Delta t/\zeta}$
    $Q_tj_N \leftarrow Q_t(h) + (Q_tj - Q_t(h)) \cdot e^{-\Delta t/\zeta}$
// For time-step lagoon water level estimate
    $Lj_N \leftarrow Lj + (Q_sj_N + Q_tj_N) \cdot \Delta t/Al(Lj)$ // $Al$ is a function of bathymetry for SBL and a linear approximation for La Rance
return $Q_sj_N, Pj_N, Q_tj_N, Lj_N$
end function
```

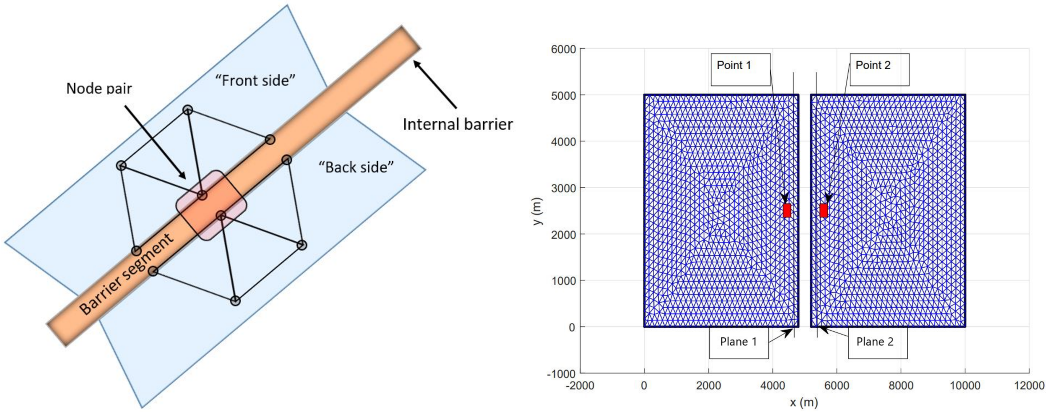

(a) Top view of Internal barrier and triangular element representation (based on: [35; 37]).

(b) Idealised rectangular model used for mesh convergence analysis and verifying mass conservation.

Figure 4.20. Barrage boundary representation and idealised rectangular mesh model.

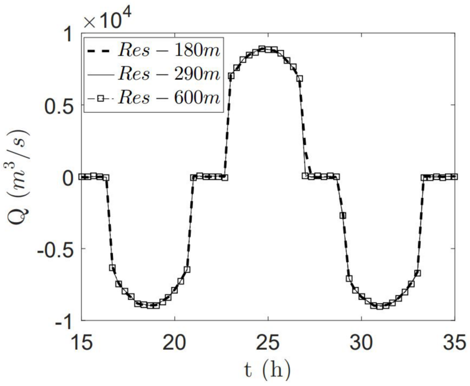

Figure 4.21. Flow rate comparisons of different mesh-resolution models of Fig. 4.20b.

ference h is greater than 1.5 m. For driving the flow through the sluices, a sinusoidal component, with 3 m amplitude and 12.42 h period (same as M_2 tide) is applied to the rectangular domain on the left, while the right side is set to have a constant water level (Points 1 and 2, Fig. 4.22a). Toward setting up the total sluice area, each node is approximated as an orifice with 9 m in diameter. For stability purposes, a quadratic bed friction of $\tau = 0.0025$ is applied to the model. With this setup in place, we can evaluate whether flow-rate outputs are dependent on mesh resolution (element edge length) by comparing flow-rates through Plane 1 from three different mesh-resolution models. The results are illustrated in Fig. 4.21, showing that different model flow rates agree well within the studied resolution range.

In order to verify mass conservation between Planes 1 and 2 (Fig. 4.20b) for vari-

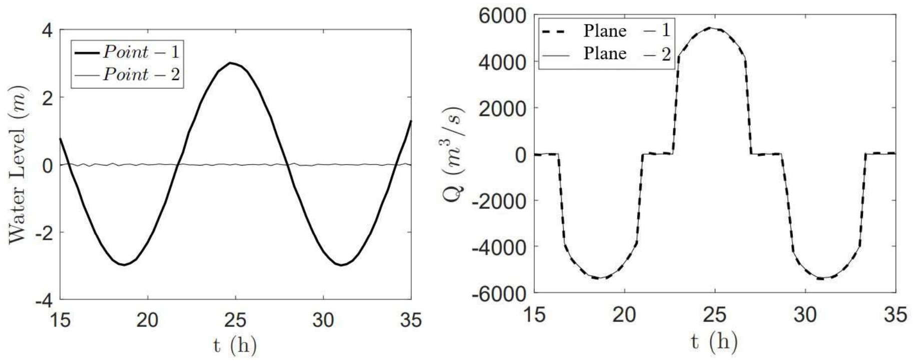

(a) Water level comparison between Points 1 and 2.

(b) Flow rate comparison between Planes 1 and 2, for 7m diameter "sluices´´.

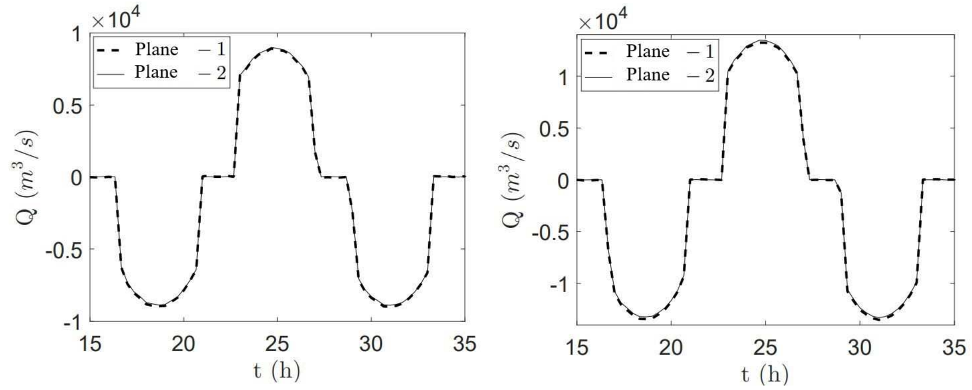

(c) Flow rate comparison between Planes 1 and 2, for 9m diameter "sluices´´.

(d) Flow rate comparison between Planes 1 and 2, for 11m diameter "sluices´´.

Figure 4.22. Flow-Rate comparison between Planes 1 and 2 [(b),(c) and (d)], for the water head difference shown in (a).

able flow-rate magnitudes, diameters of 7, 9 and 11 m are used for another comparison. The results are plotted against each other in Figs. 4.22b, 4.22c and 4.22d, showing that, for the different magnitudes of diameter presented, mass is accurately conserved.

The developed barrage boundary has also been verified for mass conservation in [37]. Furthermore, the model has been utilised for simulating SBL and Severn Barrage in [20], [21] and [36], utilising the Andritz turbine representation showcased in Algorithm 1. In these studies, the TRS were evaluated in their capability of protecting coastal regions during storm surges and their effect on varying the natural tidal resonance of the Bristol Channel. Finite element mesh representations developed for both the Severn Barrage and SBL can be seen in Fig. 4.23.

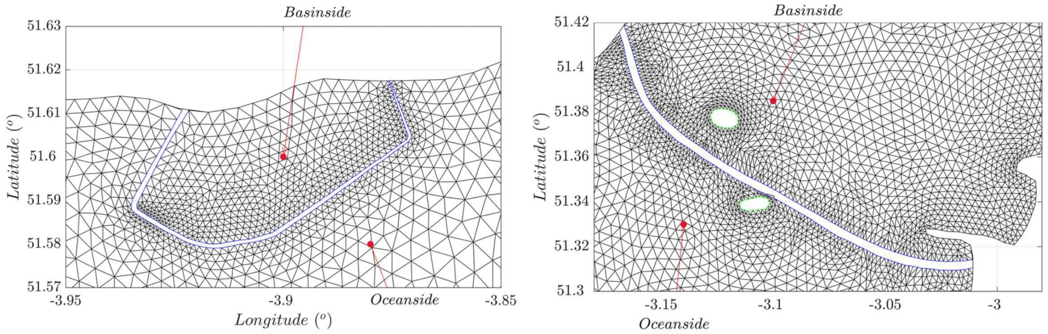

(a) Mesh representation and barrage boundary for the SBL.

(b) Mesh representation and barrage boundary for the Severn Barrage.

Figure 4.23. Developed mesh representations for the Severn Barrage and SBL (from [20], [21] and [36]).

4.4 Markov Decision Process Setup in Unity ML-Agents – SBL Case Study

In this section, we will go through the steps for developing our Markov Decision Process (MDP) in a Unity ML-Agents project. As a case study, we will utilise the SBL, operated with a TW scheme (without pumping), where the agent's goal is to maximise power generation. Furthermore, both actor and critic neural networks are of the type: MLP. Comparison results between our trained agent and optimisation baselines from the literature can be found in Section 5.1. Sections 5.2 and 5.3 showcase trained DRL results, for La Rance and SBL, respectively, when operated in a TWP scheme.

To formalise TRS operation as a RL problem (and subsequently solve the RL problem through PPO) we need to design an MDP in Unity ML-Agents with environment, agent, actions, states and reward components.

By creating simple representative 3-D models for turbines, sluices, ocean and lagoon, a training environment simulating TRS for our MDP is created in Unity3D and then imported to a Unity ML-Agents project. In this environment, the equations simulating power and flow output through the lagoon, when operating turbines and sluices, are extracted from the 0D model representation, detailed in Algorithms 1, 2 and 5. In order to choose appropriate parameters for operating our environment, we follow literature representations suggested by [23; 71; 25], for the SBL project. The chosen parameters are shown in Table 4.8. A variable lagoon surface area, digitised from [25] (shown in Fig. 3.2), is also utilised.

For ease of visualisation and debugging, the 3-D representations of sluice and turbine change colours depending on the operational mode picked by the agent. For

Table 4.8. Swansea Lagoon design.

N^o of Turbines	16
N^o of Gp	95
Grid frequency (Hz)	50
Turbine Diameter ($metres$)	7.35
Sluice Area ($metres^2$)	800

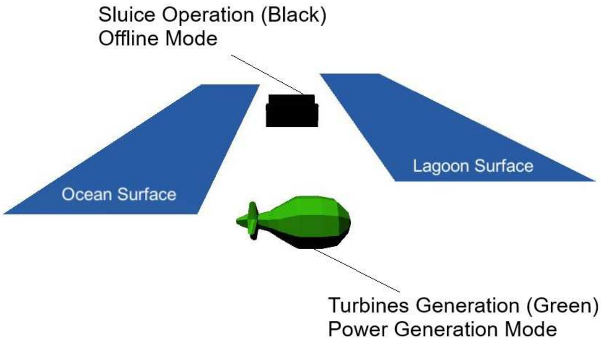

Figure 4.24. Unity ML-Agents MDP environment for a 0D model of the Swansea Bay Tidal Lagoon during ebb generation.

the turbine, green represents power generation mode, orange – idling mode and black – offline mode (zero flow rate). Similarly, sluices change colour between orange and black for sluicing and offline modes, respectively. Fig. 4.24 shows a capture of the Unity ML-Agents MDP environment representation for the SBL during ebb generation, with the representative models for sluice and turbines in offline and power generation modes, respectively. Ocean and Lagoon surface level motion are also represented.

Additionally, in this proposed MDP, the actor-critic agent is defined as an operator responsible for controlling turbine and sluice operational modes through actions (policy network node outputs n_o), according to a vector of input states s_t. n_o outputs can be discrete or continuous. In this case study, continuous outputs are chosen, reducing the number of nodes in the last layer, and consequently, the complexity of the neural network. There are 3 node outputs that determine turbine and sluice operation every 15 min of the environment simulation. The 15 min window (MDP time-step) was selected for this work since the time usually associated with the opening/closing of hydraulic structures lies in the [15 min – 20 min] range [23; 17]. Each node in the last layer outputs a value between 0 and 1, and the resulting actions are computed in a hierarchical fashion. The first node determines the number of turbines set to power generation mode (0 or 16), depending if the node output is below or above a threshold (0.5), i.e. if the node outputs a value below the threshold, no turbines will be generating energy, otherwise, all 16 turbines are set to power generating mode. Therefore, if no turbines are set to power generation, 16 are available for other operational modes (idling or offline).

The second node selects the number of idling turbines just as the first node, if the number of turbines available is 16. Otherwise the number of idling turbines is 0, independent of this node output. If no turbine is selected for power generation or idling modes, all turbines are set offline. Therefore, the first two nodes control turbines through discrete actions.

The third and final node outputs the % opening area of the sluice gates. Since any value between $[0, 1]$ can be chosen by the neural network, the momentum ramp function (Section 4.2.1) is applied to the outputted flow rate every time-step, ensuring smooth flow rate transitions, independently of the opening sluice area set by the agent.

Beyond reducing the number of node outputs, this configuration also allows for having the sluice operation independent of turbine operation. All possible operational modes for turbines and sluices as a function of node output (n_o) are shown in Tables 4.9 and 4.10, respectively.

Table 4.9. Possible turbine operational modes.

Node 1	Node 2	Discrete Turbine Control
$n_{o1} < 0.5$	$n_{o2} < 0.5$	Offline Mode
$n_{o1} < 0.5$	$n_{o2} \geq 0.5$	Idling Mode
$n_{o1} \geq 0.5$	$n_{o2} < 0.5$	Power Generation Mode
$n_{o1} \geq 0.5$	$n_{o2} \geq 0.5$	Power Generation Mode

Table 4.10. Possible sluice operational modes.

Node 3	Continuous Sluice Control
$0 < n_{o3} \leq 1$	Sluicing Mode (Available sluice area = $n_{o3}A_S$)
$n_{o3} = 0$	Offline Mode (Available sluice area = 0)

The actions selected by the agent are a function of the input states s_t. In this case study, these states are the water levels of ocean and lagoon, plus current operational mode of turbines and sluices, for current and previous MDP time-steps (Table 4.11). Finally, the reward received by the agent equals the accumulated energy generated by the turbines, for every 15 *min*.

For stabilising and speeding up training, parallel training is performed with 64 copies of the environment (Fig. 4.25), all being operated by a single agent (actor neural network), with episodes set to 1 month of duration. During training, each environment instance requires a representative ocean input at the location where the

Table 4.11. Input states for PPO neural network.

States (at times t and $t-1$)	Units
Ocean water level	"Normalised" $[0, 1]$ (float)
Lagoon water level	"Normalised" $[0, 1]$ (float)
Number of online turbines	0, or 1 (integer)
Number of idling turbines	0, or 1 (integer)
Sluice gate opening area	0 to 1 (float)

Figure 4.25. 64 instances of the environment during parallel training. For the turbines, green represents power generation mode. For turbines and sluices orange represents idling/sluicing mode and black – offline mode.

Swansea Lagoon is planned to be constructed ($51°35'58.9"N$ $3°53'42.4"W$). Ideally, ocean measurements could be used as training data. However, due to the lack of sufficient measured data (Section 5.1.1), it is not possible to train the agent until reasonable performance is reached. Instead, an artificial tide signal to simulate the ocean is created by summing the major sinusoidal tide constituents (due to gravitational pull of the Moon and Sun). Although we are not accounting for other less predictable local wave motions (e.g. wind waves), the artificial ocean input representation is sufficient for enabling the agent to converge to an optimal policy. Two major advantages of this approach is the fact that (i) we can generate any amount of input data required for training the agent and (ii) we reduce the risk of over-fitting by training with synthetic data [123].

The tide constituent's amplitudes of the simulated ocean utilised in this work (Table 4.12), were obtained from a numerical simulation, at the location of Swansea Bay Tidal Lagoon, by [23]. The periods for each constituent were obtained by [124]. The final equation for simulating the ocean can be seen in Eq. 4.29.

$$\begin{aligned} o(t) = A_{M2} sin(\omega_{M2} t + \phi_{M2}) + A_{S2} sin(\omega_{S2} t + \phi_{S2}) + \\ A_{N2} sin(\omega_{N2} t + \phi_{N2}) + A_{K1} sin(\omega_{K1} t + \phi_{K1}), \end{aligned} \tag{4.29}$$

where ω_{M2}, ω_{S2}, ω_{N2} and ω_{K1} are angular frequencies (rad/s) of each tidal component, and ϕ_{M2}, ϕ_{S2}, ϕ_{N2} and ϕ_{K1} are random phase lags in the range $[0, 2\pi]$, generated for each environment instance during parallel training when starting an episode, which

Table 4.12. Simulated tide constituents at Swansea Bay.

Ocean tide constituent	Amplitude (m)	Period (hr)
$M2$	$A_{M2} = 3.20$	$T_{M2} = 12.42$
$S2$	$A_{S2} = 1.14$	$T_{S2} = 12$
$N2$	$A_{N2} = 0.61$	$T_{N2} = 12.66$
$K1$	$A_{K1} = 0.08$	$T_{K1} = 23.93$

allow for learning more generalised scenarios.

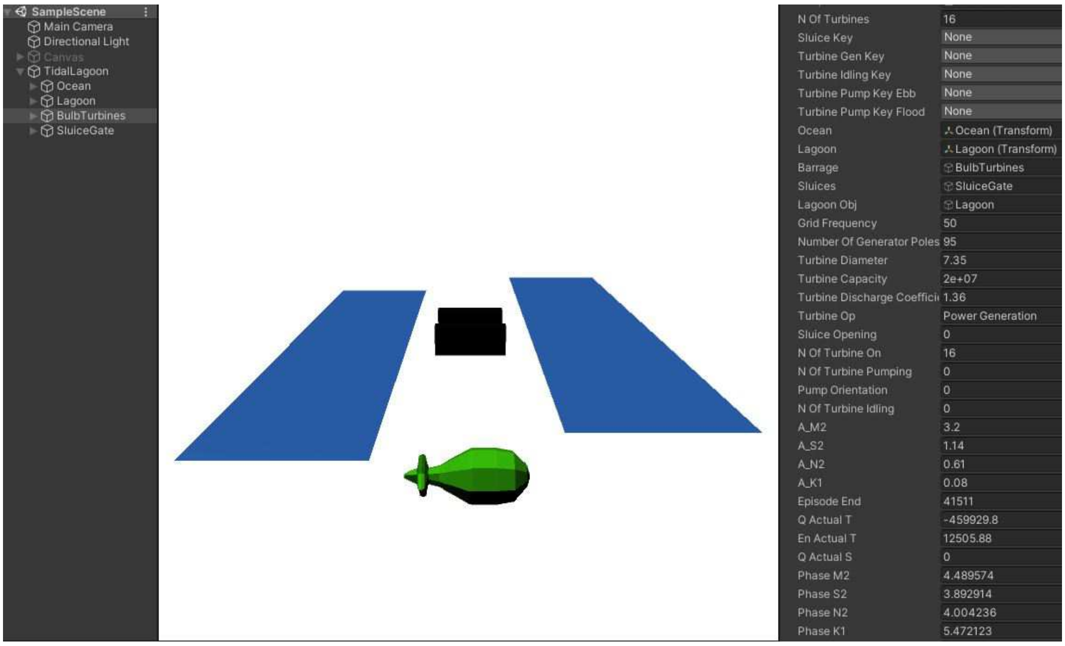

N Of Turbines	16
Sluice Key	None
Turbine Gen Key	None
Turbine Idling Key	None
Turbine Pump Key Ebb	None
Turbine Pump Key Flood	None
Ocean	Ocean (Transform)
Lagoon	Lagoon (Transform)
Barrage	BulbTurbines
Sluices	SluiceGate
Lagoon Obj	Lagoon
Grid Frequency	50
Number Of Generator Poles	95
Turbine Diameter	7.35
Turbine Capacity	2e+07
Turbine Discharge Coeffici	1.36
Turbine Op	Power Generation
Sluice Opening	0
N Of Turbine On	16
N Of Turbine Pumping	0
Pump Orientation	0
N Of Turbine Idling	0
A_M2	3.2
A_S2	1.14
A_N2	0.61
A_K1	0.08
Episode End	41511
Q Actual T	-459929.8
En Actual T	12505.88
Q Actual S	0
Phase M2	4.489574
Phase S2	3.892914
Phase N2	4.004236
Phase K1	5.472123

Figure 4.26. Unity visual framework for (i) setting up TRS and PPO parameters and (ii) visualising training and testing steps of the DRL agent.

The designed MDP is solved through the PPO algorithm (Section 2.4.1). A flow-chart, illustrating the agent's training stage through the PPO algorithm in Unity ML-Agents, is presented in Fig. 4.28. Training occurs until a "max-steps" number of observations are sampled.

After training, the policy (actor) neural network receives input states s_t and outputs optimum n_o values, following a policy that maximises energy generation. During testing, this means that the agent receives real ocean measurements as inputs and predicts the optimum action to take (for any ocean amplitude). The predicted operation does not require an external model for tidal predictions and enables the agent to perform real-time flexible control of turbines and sluices.

The TRS specifications and the number of policy network node outputs n_o and in-

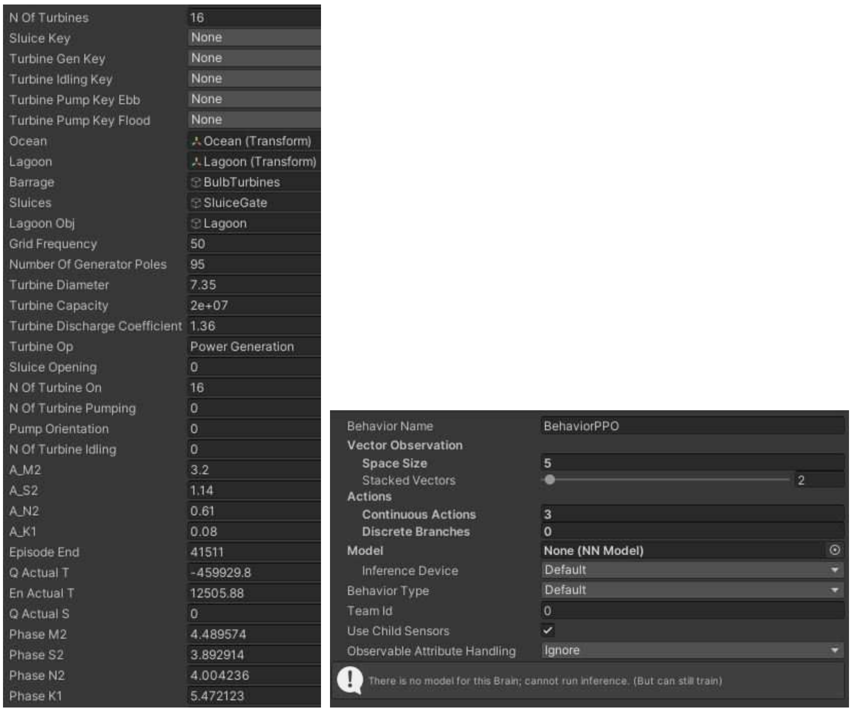

N Of Turbines	16
Sluice Key	None
Turbine Gen Key	None
Turbine Idling Key	None
Turbine Pump Key Ebb	None
Turbine Pump Key Flood	None
Ocean	Ocean (Transform)
Lagoon	Lagoon (Transform)
Barrage	BulbTurbines
Sluices	SluiceGate
Lagoon Obj	Lagoon
Grid Frequency	50
Number Of Generator Poles	95
Turbine Diameter	7.35
Turbine Capacity	2e+07
Turbine Discharge Coefficient	1.36
Turbine Op	Power Generation
Sluice Opening	0
N Of Turbine On	16
N Of Turbine Pumping	0
Pump Orientation	0
N Of Turbine Idling	0
A_M2	3.2
A_S2	1.14
A_N2	0.61
A_K1	0.08
Episode End	41511
Q Actual T	-459929.8
En Actual T	12505.88
Q Actual S	0
Phase M2	4.489574
Phase S2	3.892914
Phase N2	4.004236
Phase K1	5.472123

Behavior Name	BehaviorPPO
Vector Observation	
Space Size	5
Stacked Vectors	2
Actions	
Continuous Actions	3
Discrete Branches	0
Model	**None (NN Model)**
Inference Device	Default
Behavior Type	Default
Team Id	0
Use Child Sensors	✓
Observable Attribute Handling	Ignore

There is no model for this Brain; cannot run inference. (But can still train)

(a) TRS parameter setup inside Unity.

(b) PPO parameter setup inside Unity.

Figure 4.27. TRS and PPO parameter setup in Unity. TRS editable parameters can be added into the setup, by declaring variables as public.

put states s_t, are all defined in a convenient visual framework representation developed inside Unity (Fig. 4.26). Detailing for TRS and the PPO neural network configuration are shown in Fig. 4.27a and Fig. 4.27b, respectively. Hyperparameters setup during training are described in the next section.

4.4.1 Hyperparameter Setting

The final hyper-parameters used for training our agent for the SBL, no pumping, case study (Section 5.1) are shown in Table 4.13. They are ordered according to a $trainerconfig.yaml$ file, available from the Unity ML-Agents package (version specification in Table 4.14). The hyperparameters batch size (sampled number of trajectories m), α, β, ϵ, λ, time horizon H and γ are all detailed in Section 2.4.1. Training occurs until a "max-steps" number of observations (Eq. 2.7) are sampled. During training, learning rate schedule is set to "linear", which means that the learning rate decays in a linear fashion from an initial value $\alpha = 1.5\text{e-}4$ to $\alpha = 0$.

Time horizon H represents each trajectory size. In the case of this work, since the

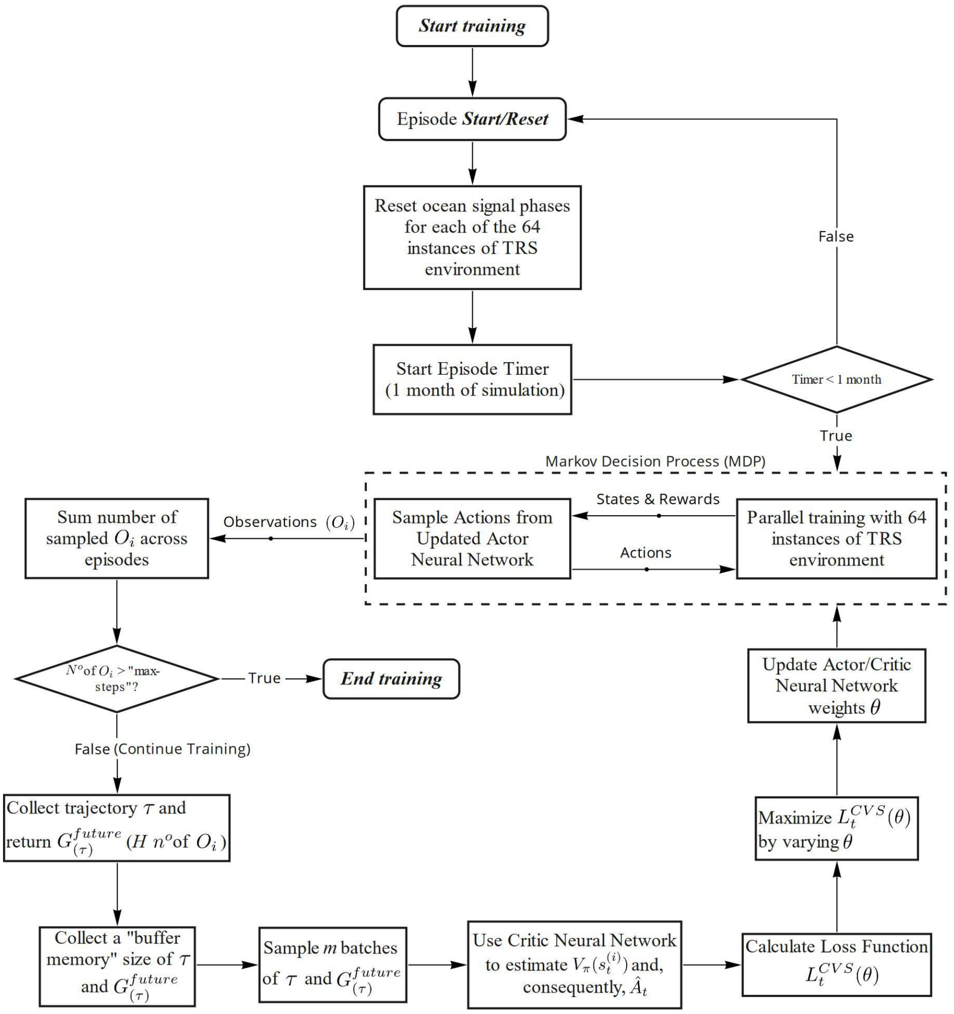

Figure 4.28. Flow-chart with detailing of the PPO algorithm in Unity ML-Agents, during training. For a given trajectory $\tau^{(i)}$, $\hat{A}_t$ is the advantage estimate and $V_\pi(s_t^{(i)})$ the value function.

MDP time-step $= 15\ min$, a time horizon of 54 equals to 13.5 h – slightly more than the tidal period ($M_2 =\ 12.42\ h$), and enough so that all possible operational states (during ebb and flood tides) can be observed every trajectory. When a “buffer size” number of these trajectories is stored during training, gradient ascent is performed a “num epoch” number of times by sampling random “batch size” (number of trajectories m) and differentiating the loss function in Eq. 2.28. After “num epoch” of training occur, the buffer is reset and the process repeats.

“normalize” set to “true” means that input states are normalised during training, which for continuous control helped the agent to converge to an optimal policy. During

Table 4.13. Hyperparameters for training in Unity ML-Agents, for the SBL (no pumping) case study

Hyperparameters:	
batch size (m)	256
buffer size	4096
learning rate (α)	$1.5e-4$
learning rate schedule	linear
β	$4.6e-2$
ϵ	0.2
λ	0.98
num epoch	6
Neural network configuration:	
normalize	true
hidden units	48
num layers	4
Trainer Configuration:	
max steps	$8.0e7$
time horizon (H)	54
summary freq	40000
Reward Signal:	
strength	1
γ	.97

Table 4.14. Unity ML-Agents (version 0.22.0), for the SBL (no pumping) case study.

Version Information	
ml-agents	0.22.0
ml-agents-envs	0.22.0
Communicator API	1.2.0
Pytorch	1.6.0

testing (real-time control), real ocean measurements are normalised by the sum of the tidal constituent's amplitudes of our artificial ocean signal (Table 4.12). "summary frequency" only determines the granularity of training results.

Both policy and critic networks follow an MLP architecture, with 10 input states s_t, 3 node outputs and 4 hidden layers (each with with 48 units).

In the Chapter 5, the developed methodologies of this Chapter (enhanced state of art and DRL approach for operation optimisation of TRS) will be applied to the chosen case studies in this work: a tidal lagoon and tidal barrage TRS project types, namely the SBL pathfinder project and La Rance tidal barrage.

Chapter 5

Performance Evaluation of TRS Operation through DRL and Enhanced State-of-Art Methods

5.1 SBL (no pumping) Case Study

In this section, we compare our DRL trained agent performance against baselines: state-of-art and enhanced state-of-art optimisation routines. For this case study, we utilise the SBL, without pumping capabilities, with the goal of maximising power generation. While the operational strategy for the DRL agent is learnt during training, the baselines utilise the TW operational scheme. The MDP and hyperparameters utilised for training our DRL agent were presented in Section 4.4.

5.1.1 Test Data and Baselines Optimisation

For comparing our DRL agent performance against conventional (and state-of-art) optimisation routines, we model six HC baselines – three from the literature [25] and three "variant" methods with improvements inspired from [23] –, and compare the energy generated in a month for each method. All baselines in this work consider the operation of the SBL either through classic or variant "two-way scheme" methods, as detailed in Section 3.1. We highlight that the showcased HC variant operation is an augmentation proposed in the present work.

Regarding test data for baselines and trained agent, we utilise all tide gauge ocean measurements available from the British Oceanographic Data Centre (BODC) at Mumbles Station [125], located at the edge of Swansea Bay. The obtained measurements of ocean elevation are recorded every 15 min in a table, for the years of 1993 and the range [1997 – 2019]. Before utilising the data, a preprocessing step is performed so that data flagged as "improbable", "null value" and "interpolated" by BODC are not considered. After this step we retain 26 months of usable, non-overlapping, test data. The preprocessing step ensures a conservative comparison between baselines and our trained agent, since it considers scenarios where tidal predictions had a good match with measured data.

Tidal predictions for the same 26 months are also provided by BODC in the same data-set. For each month, baseline optimisation routines utilise tidal predictions for capturing operational head values H_{start}, H_{min} and HS_{start} (when considered) that optimise power generation. These operational head values are then applied to the measured ocean test data, so that comparisons between baselines and trained agent can be made. Differences between this method and direct upper bound estimates (Section 3.3.3) were shown to be negligible ($< 0.1\%$), being discussed in Section 5.3, for TWP scenarios.

Baselines, in increasing order of optimisation complexity, are described next:

- CH (Constant Heads): Best, constant H_{start} and H_{min} are picked for extracting energy during a whole month [77].
- CHV (Constant Heads, with variant operation): Best, constant heads H_{start} , H_{min} and HS_{start} are picked for extracting energy during a whole month.
- EHT (Every Half-Tide): optimised pairs of H_{start} and H_{min} are picked for every consecutive half-tide. Proposed by [25].
- $EHTV$ (Every Half-Tide, with variant operation): optimised H_{start} , H_{min} and HS_{start} are picked for every consecutive half-tide.
- EHN (Every Half-Tide and Next): optimised H_{start} and H_{min} are picked for every half-tide, considering the best H_{start} and H_{min} for the next half-tide as well. Proposed by [25].
- $EHNV$ (Every Half-Tide and Next, with variant operation): optimised H_{start} , H_{min} and HS_{start} are picked for every half-tide, considering the best H_{start} , H_{min} and HS_{start} for the next half-tide as well.

All variant optimisation methods are augmented through the addition of independent sluice head operation HS_{start} . This modification we are introducing in this thesis is inspired by the work of [6; 23], from TC methods. CH and CHV perform non-flexible operation, while EHT, $EHTV$, EHN and $EHNV$ perform flexible operation. A summary detailing each baseline operational heads and method is shown in Table 5.1. A full description of CH, EHT and EHN methods can be seen in Section 3.3.

Table 5.1. Simplified reference table for baselines.

	Constant Head		Every Half-Tide		Every Half-Tide and Next	
	CH	CHV	EHT	$EHTV$	EHN	$EHNV$
H_{start}	✓	✓	✓	✓	✓	✓
H_{min}	✓	✓	✓	✓	✓	✓
HS_{start}		✓		✓		✓
non-flexible operation	✓	✓				
flexible operation			✓	✓	✓	✓

All baselines, except $EHNV$, are optimised with a grid search optimisation algorithm, which iteratively increases its search resolution until convergence. Initial search resolution starts with 1 meter, with optimisation heads H_{start} , H_{min} and HS_{start} (when considered) within ranges $[1m - 6m]$, $[1m - 3m]$ and $[1m - 5m]$, respectively. After the first run, search resolution is halved and the algorithm performs a brute-force search around the best previous configuration attained. The latter procedure is repeated until final search resolution is lower than $1cm$.

$EHNV$ requires a different optimisation approach due to its high computational time when utilising the previous grid search method. For this case, we utilise the stochastic global optimisation algorithm basin-hopping [126] from Scipy package [127], with COBYLA as a local minimizer [128]. Basin-hopping was chosen for its efficiency when solving smooth function problems with several local minima separated by large barriers [129]. The local minimizer COBYLA is a nonlinear derivative–free constrained optimisation that uses a linear approximation approach. Even though basin-hopping is not guaranteed to converge to a global optimum, $EHNV$ is shown to be, on average, the best baseline method for energy generation.

5.1.2 Agent Performance Evaluation

Following hyperparameter tuning, we trained the agent for $8E7$ steps, until convergence. The cumulative reward (energy) per month (episode) during parallel training (with the artificial ocean signal in Eq. 4.29), averaged for the 64 instances of the lagoon environment, is shown in Fig. 5.1. The log-representation insert highlights the two-step plateau that is observed when converging to an optimal strategy. After starting in a total random strategy, the cumulative reward received by the agent increases until reaching an intermediate plateau at around $2E6$ steps, where the agent learns

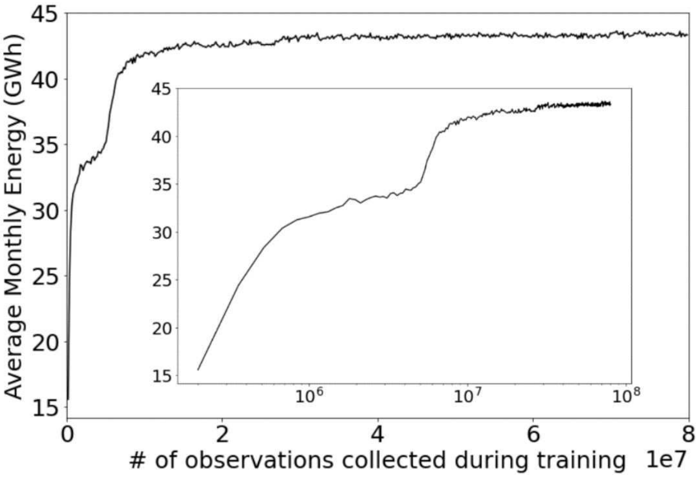

Figure 5.1. Monthly cumulative energy (reward) in GWh, averaged for all 64 environments during parallel training. The log-representation insert highlights the two-step plateau.

the strategy of operating mostly the turbines (during ebb and flood tides), while keeping sluices practically offline during ebb generation. Then, after about $5E6$ steps, the cumulative reward starts increasing again. The second plateau stabilises around $4E7$ steps, with a cumulative reward approximately 25% higher than the first plateau – a gain allowed by (i) a TW flexible operational strategy learnt by the agent, that adjust TRS operation according to tidal range (ii) the smart usage of the sluicing mode, as discussed below in test results.

For test data, we utilise 26 months of real ocean measurements from BODC. These months are presented and numbered in Table A.1, while Table A.2 compares the amount of energy obtained in the numbered months between our trained agent (performing real-time flexible control) and the baselines estimate. A block diagram representing the trained agent test stage for this work, while using ocean water level (WL) measurements, is shown in Fig. 5.2. The averaged monthly energy attained for all methods is shown in Fig. 5.3.

For the baselines, CH and CHV present the worst performance, since constant operational heads cannot account for the varying ocean amplitudes in a month (about $\approx 2m$ to $\approx 4.5m$ in our test set). Furthermore, baselines with variant operation outputted more energy in average than their classical counterparts. This fractional gain in respect to each baseline can be seen in Fig. 5.4 ($\approx 2\%$ gain for $EHTV$ and $EHNV$ over their classic counterparts, and $\approx 1\%$ gain for CHV over CH). Finally, "half-tide and next" approaches showed very small improvements ($< 0.2\%$) when compared to "half-tide" methods, while requiring much greater ($\gtrapprox 20\times$) computational time (Tables

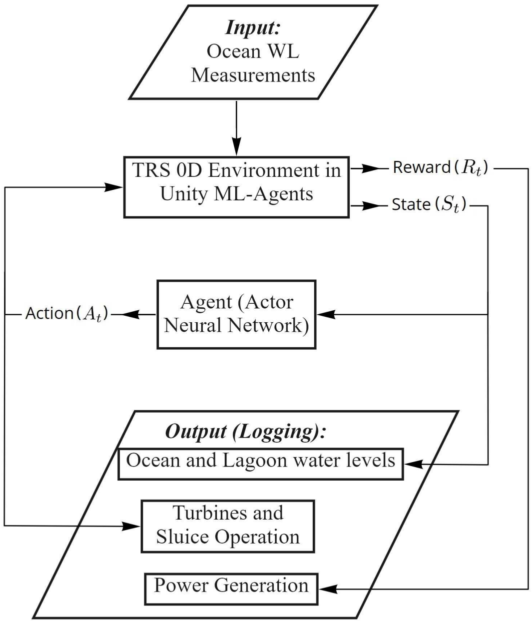

Figure 5.2. Block diagram with detailing of the trained agent test stage.

5.2 and A.3).

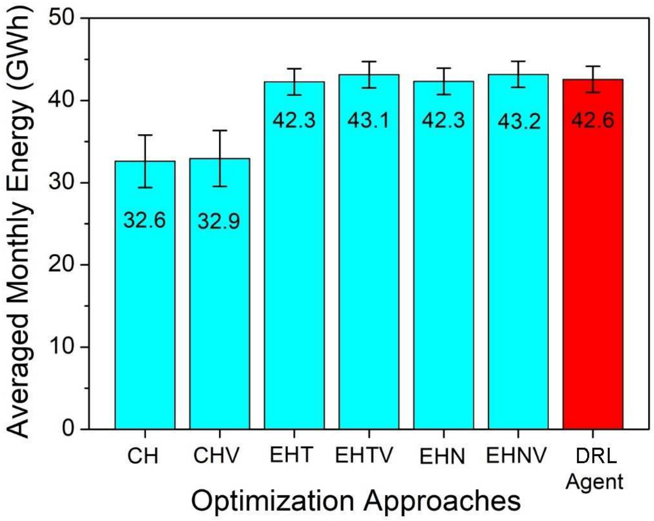

Figure 5.3. Averaged monthly energy comparison between baselines and trained agent utilising test data. Sample standard deviations for the various months are also shown as error bars.

For the trained agent, Fig. 5.5 show operational test results of power generation and lagoon water levels for one month of measured ocean data (starting with initial lagoon water level at mean sea level).

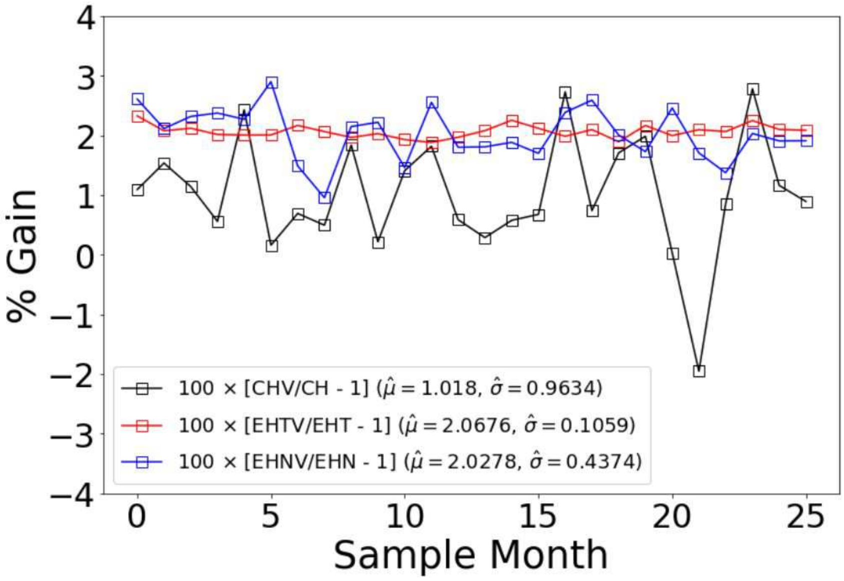

Figure 5.4. Fractional % gain of monthly energy generation for each baseline, comparing classic and variant operation.

Table 5.2. Computational time for all methods when optimising over 26 months of tidal measurements.

Optimisation Approach	Computational time (hh:mm:ss)
DRL Agent	**Not Applicable**
CH	00:29:27
CHV	01:04:02
EHT	00:31:23
$EHTV$	01:45:26
EHN	10:05:30
$EHNV$	84:47:07

We note that the agent quickly converges to an optimal energy generation strategy for sequential tidal cycles, independent of tidal range input – a characteristic of state-of-art flexible operation [25]. Furthermore, Figs. 5.6 and 5.7 showcase detailed results of real-time control on test data. Apart from ocean water levels, results are coloured according to actions taken by the agent for turbines and sluices, respectively, as defined in Section 4.4. More specifically, Figs. 5.6a and 5.7a show lagoon water level variations, while Fig. 5.6c, 5.6b and 5.7b show power generation, turbine and sluice flow rates.

From the sequence of actions taken, we see that the agent arrives at a policy with independent operation of sluices, i.e. the variant operation of TRS, which was shown to be a better strategy than the classical operation in our fractional gain comparison (Fig. 5.4). A summary of our method accomplishments in comparison with state-of-art baselines is shown in Table 5.3.

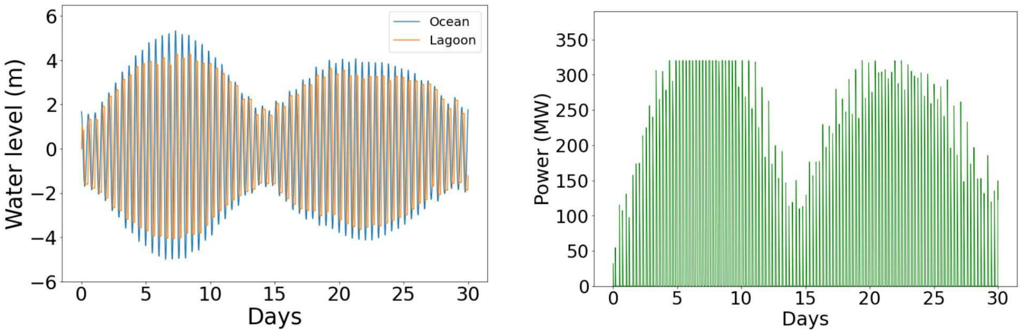

(a) Ocean and lagoon water levels. **(b)** Power generation.

Figure 5.5. Lagoon water levels and power generation results for the trained agent performing flexible control in a month, with measured ocean data only.

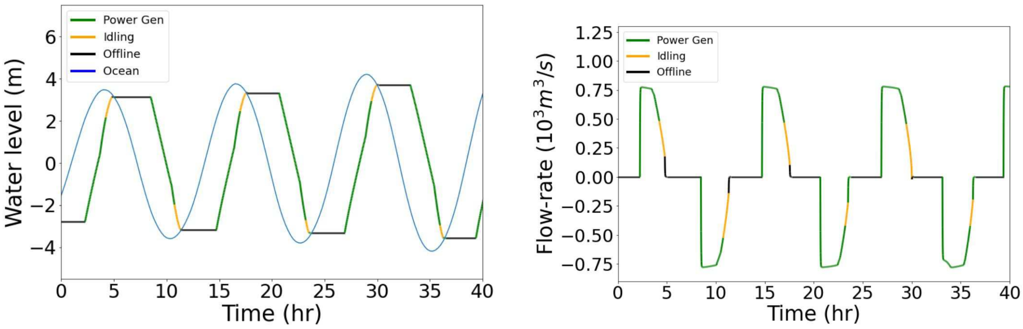

(a) Ocean (in blue) and lagoon water levels. **(b)** Flow rate from the 16 turbine units.

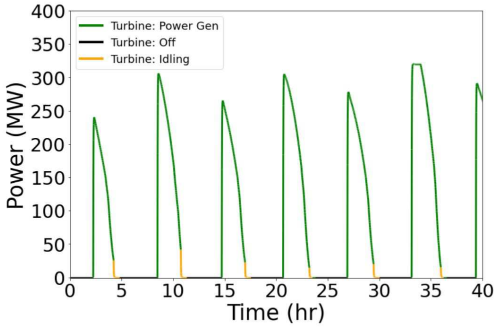

(c) Combined power output from turbines.

Figure 5.6. Lagoon water levels, turbine flow rates and power output are shown and coloured following turbine operational mode chosen by the trained agent. Green represents power generation mode, orange – idling mode and black – offline mode.

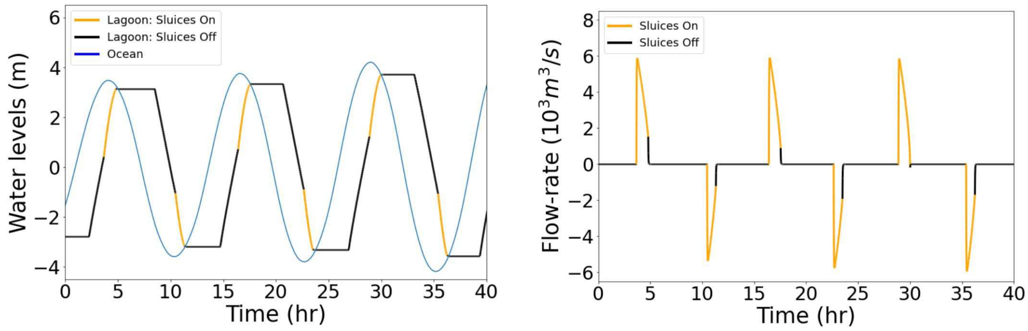

(a) Ocean (in blue) and lagoon water levels. **(b)** Flow rate from sluice gates.

Figure 5.7. Lagoon water levels and sluice flow rates are shown and coloured following sluice operational mode chosen by the trained agent. Orange represents idling (i.e. sluicing) mode and black – offline mode.

Table 5.3. Comparison of state-of-art baselines with our proposed DRL Agent.

	CH [77]	CHV^a	EHT [25]	$EHTV^a$	EHN [25]	$EHNV^a$	**DRL Agent (our work)**
real-time flexible control							✓
prediction-free approach							✓
variant lagoon operation		✓		✓		✓	✓
state-of-art performance[b]			✓	✓	✓	✓	✓

[a] Optimisation routines with novel HC variant operation of TRS.
[b] Equivalent outputs, in average, within the error bars (Fig. 5.3).

Our agent managed very competitive energy outputs, staying on average within 1.4% of the best baseline (our augmented $EHNV$ method) and slightly higher to that outputted by the state-of-art EHN baseline ($\approx +0.7\%$) . Indeed, for all months tested, our agent performed optimally, outputting better results than the current literature state-of-art EHN method for 22 out of 26 months, within a 1.1% margin in worst scenarios (Table A.2). We note that this novel result was obtained by training the agent once with a simple artificial ocean input, in contrast with the baselines that require being re-run for every new tide.

5.1.3 Remarks

For our SBL case study, we have shown that Proximal Policy Optimisation (a DRL method) can be used for real-time flexible control of Tidal Range Structures after training with artificially generated tide signals. Our DRL agent implementation was then compared against state-of-art optimisation approaches devised (or augmented) from the literature, yielding competitive results for all test data utilised. Our results were obtained in a conservative setting, i.e., when available tidal measurements had good agreement with tidal predictions. Furthermore, we showed how our proposed improvements to HC methods, with independent sluice operation, can consistently increase power generation capabilities of state-of-art methods.

We have chosen the Swansea Bay Tidal Lagoon for our analysis, given its status as a pathfinder project for larger tidal lagoon projects. We show that our DRL approach obtains optimal energy generation from measured tidal data only, through an optimised control policy of turbines and sluices, which are operated independently by the trained DRL agent. Our method shows promising advancements over state-of-art optimisation approaches since it (i) performs TW real-time flexible control with equivalent energy generation, (ii) does not require future tidal predictions, and (iii) needs to be trained a single time only. Results of this work are published in [105].

5.2 Operational Optimisation of the AI-Driven La Rance Model

In this section, we insert the complete parametrised 0D La Rance model (Section 4.3.1) as an MDP environment into Unity ML-Agents and train the TRS representation. We show that the obtained AI-Driven La Rance model learns a TWP operational strategy that is comparable in (i) energy extraction capabilities and (ii) detailed scheme of operation to the actual strategy utilised in La Rance. For our comparison analysis, we utilise measurements of the yearly net energy measured in [130] and the observed sequence of operation of hydraulic structures from [67]. It is important to highlight that the actual operation in La Rance, in contrast with the AI-Driven strategy, has the objective of maximising revenue instead of energy [131]. However, as noted in [24], deviations in energy extraction when comparing revenue and energy based TRS optimisation are expected to be around $4-5$ % only. Therefore, the comparisons we are showcasing are technically sound. In order to train and test our DRL agent, representative tidal data at the location of La Rance, are required. In the next section, we present a free software utilised for tidal predictions that can provide such data.

5.2.1 JTides – Tidal Prediction Software

JTides is a free, worldwide, tidal and current prediction software that utilises harmonic decomposition techniques for predicting ocean tides in several locations of the planet [132]. Although JTides have been utilised for research [133; 134], its application in the field of tidal power has not yet been explored.

From JTides interface, the user can insert coordinates for any location on Earth. From this location, JTides looks for a nearest point of reference, for which it can extract tidal predictions from its database. By providing the coordinates of the La Rance tidal barrage, JTides returns tidal predictions for the location of St. Helier, Jersey, the largest of the Channel Islands in the English Channel. This island is located around $80km$ of La Rance and, as the latter, experiences one of the highest tidal ranges on the planet (up to $10m$ [135]).

In order to assess if tidal predictions from St. Helier are appropriate for La Rance, tidal predictions provided by EDF, the company responsible for operating La Rance, are utilised as reference. The tidal predictions from EDF were available through the web-page (https://www.edf.fr/usine-maremotrice-rance/marees-en-rance), being updated every week, providing 3 weeks of forecast. A whole year of tidal predictions were collected manually, for comparison with JTides predictions.

By comparing JTides and EDF predictions, we note that results have the same pattern, although with a small deviation at tidal range's extremes, with tidal predictions by EDF consistently predicting higher tidal amplitudes. By assuming that EDF and JTide's waves are similar, a deviation coefficient between tidal predictions can be obtained through the simple method of root mean square differences (RMS) [136]. With this assumption, we can estimate a correction factor C_f to be applied to JTides's prediction ($JTides \times C_f$), in order to reduce the deviation between JTides and EDF estimates:

$$C_f = \frac{\sum_{i=1}^{i=N} O_{Ji} O_{Ei}}{\sum_{i=1}^{i=N} O_{Ji}^2}, \tag{5.1}$$

where O_{Ji} and O_{Ei} are the ocean predictions for JTides and EDF (oscillating around mean water level), respectively, and N is the number of data points for a whole year. With this method, a $C_f \approx 1.10$ was obtained. A comparison of tidal prediction elevations from EDF and $JTides \times C_f$ is showcased in Fig. 5.8, for one year round. From the observed residuals, we see that the agreement between $JTides \times C_f$ and EDF is consistent throughout the year, apart from very few isolated spikes. Therefore, we resume our work with the assumption that $JTides \times C_f$ tidal predictions are representative of expected ocean water levels at La Rance. With this assumption, we utilise the $JTides \times C_f$ software capabilities to generate training data (from 2013 until 2038) for our DRL agent.

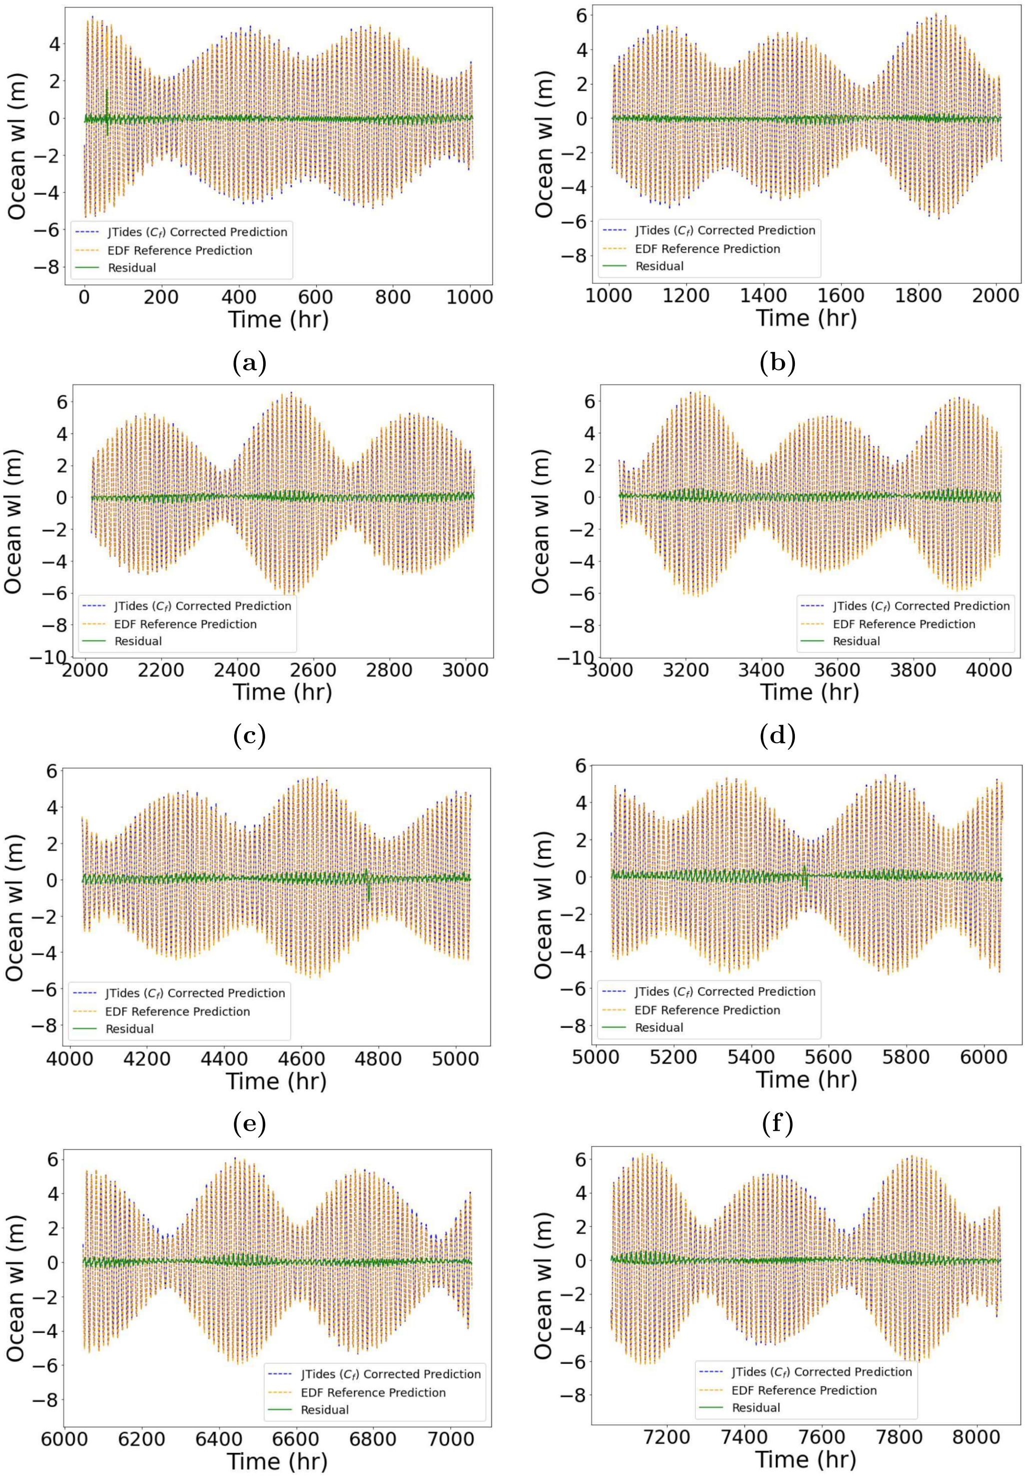

(a) (b)

(c) (d)

(e) (f)

(g) (h)

Figure 5.8. Comparisons between C_f corrected JTides and EDF ocean predictions, for the La Rance tidal barrage, for one year-round ($June/06/2020$ up to $June/05/2021$).

5.2.2 Unity ML-Agents implementation

With training data in hands, we implement the La Rance parametrised hydraulic structures into the 0D TRS representation in Unity, following the TWP algorithmic representation from Section 4.3.1. Although La Rance's DRL implementation is nearly identical to the SBL case study in Section 4.4, there are important differences that need to be discussed. Most notably, the previously adopted continuous control strategy was not capable of operating turbines in pump mode. For solving this, La Rance's DRL approach utilises a discrete control solution (covered in Section 2.4.1), when utilising the PPO algorithm. The discrete control solution outputs 3 branches in the last layer of the policy network. The first branch n_{oS} is responsible for controlling sluices, while the second branch n_{oT} is responsible for setting the turbine operational mode. Finally, the third branch n_{oP} controls the input power P_{in} available to turbines in pump mode. Each branch ramifies into possible discrete actions that can be taken (i.e. a probability mass function), allowing for the full range of combinations between turbines and sluices to be explored during training.

Table 5.4. Discrete control options for La Rance.

Branch	Discrete Action
n_{oS} (sluices)	Offline Online
n_{oT} (turbines)	Offline Power Generation (Ebb, or Flood) Idling Ebb Pumping Flood Pumping
n_{oP} (pump P_{in})	$\{P_{in} : P_{in} = 0.25MW \times n_{oP}, n_{oP} \in \{0, 1, ...16\}\}$

The possible actions for each branch are showcased in Table 5.4, while the required input states for the policy neural network are presented in Table 5.5. The PPO setup implementation in Unity is shown in Fig. 5.9a.

Table 5.5. Input states for PPO neural network, for La Rance.

States (at times t and $t-1$)	Units
Ocean water level	"Normalised" $[0, 1]$ (float)
Lagoon water level	"Normalised" $[0, 1]$ (float)
Sluice Mode	$n_{oS} \in \{0, 1\}$ (integer)
Turbine Mode	$n_{oT} \in \{0, 1, 2, 3, 4\}$ (integer)
Pump Power Input	$n_{oP} \in \{0, 1, ...16\}$ (integer)

Regarding the modelled MDP environment in Unity3D, the 0D model representation was augmented in order to consider pumping capabilities (Section 4.2.5). Furthermore, the 3-D visual representation of hydraulic structures operation in Unity3D for La Rance's implementation follows the same colour scheme as for SBL, with the addition of turbines changing colour to red when in pumping mode (Fig. 5.9b).

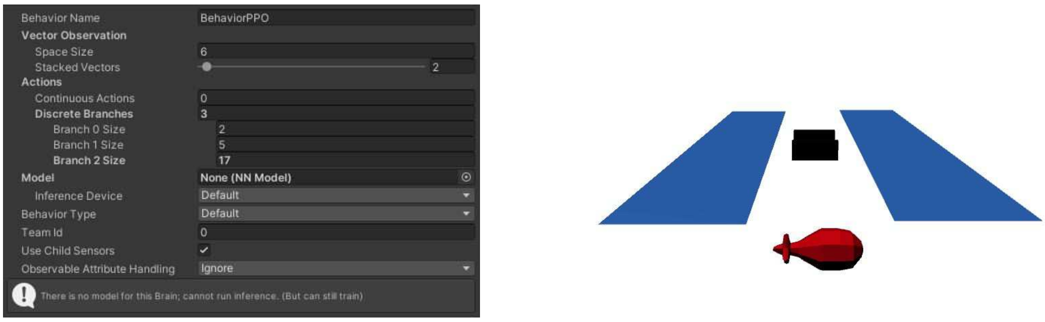

Behavior Name	BehaviorPPO
Vector Observation	
Space Size	6
Stacked Vectors	2
Actions	
Continuous Actions	0
Discrete Branches	3
Branch 0 Size	2
Branch 1 Size	5
Branch 2 Size	17
Model	None (NN Model)
Inference Device	Default
Behavior Type	Default
Team Id	0
Use Child Sensors	✓
Observable Attribute Handling	Ignore

There is no model for this Brain; cannot run inference. (But can still train)

(a) Unity's PPO parameter setup for La Rance.

(b) La Rance pump operation representation in Unity.

Figure 5.9. PPO parameter setup in Unity for a TWP scheme operation and the visual pump representation for La Rance.

As shown in Table 5.4, P_{in} values have been discretised with a $0.25MW$ resolution and upper bound of $4MW$, for each bulb unit. This upper bound have been selected, given that initial training sessions for La Rance never surpassed $4MW$ power input, for each of the 24 bulb units available (Table 4.4). Furthermore, as discussed in [67], even though each unit can receive $6MW$ of power input, optimal results were obtained with smaller values, with pumping sometimes being restricted to $50MW$ (for all 24 units), depending on tidal range [67].

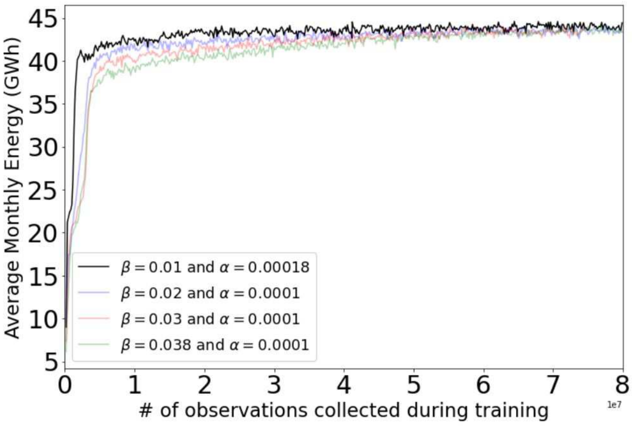

Figure 5.10. Monthly cumulative energy (reward) in GWh, averaged for all 64 environments during parallel training, for the La Rance tidal barrage.

Parallel training is performed with 64 copies of the environment in Unity3D and observing the average monthly energy reward obtained. By starting with the initial hyperparameter settings from the optimised model in [105] as reference, we note a significant improvement of results when increasing the number of units and hidden layers of the policy neural network. After fixing the complexity of the neural network, converging to a stable plateau was possible by tuning hyperparameters α and β for multiple runs (Fig. 5.10), with α being the learning rate and β a linear decaying hyperparameter responsible for controlling how much the agent explores the environment during training (i.e., increasing β leads to more random actions at the beginning of training) [105]. From initial $\beta = 0.038$ and $\alpha = 0.0001$, an optimal policy was obtained by decreasing β and increasing α. The final hyperparameter setting for acquiring the optimal policy are shown in Table 5.6. The used Unity version is indicated in Table 5.7.

Table 5.6. Hyperparameters for training La Rance environment in Unity ML-Agents.

Hyperparameters:	
batch size (m)	256
buffer size	4096
learning rate (α)	$1.8e-4$
learning rate schedule	linear
β	$1e-2$
ϵ	0.14
λ	0.98
num epoch	6
Neural network configuration:	
normalise	false
hidden units	64
num layers	4
Trainer Configuration:	
max steps	$8.0e7$
time horizon (H)	54
summary freq	40000
Reward Signal:	
strength	1
γ	.97

Table 5.7. Unity ML-Agents (version 0.26.0).

La Rance's Version Information	
ml-agents	0.26.0
ml-agents-envs	0.26.0
Communicator API	1.5.0
Pytorch	1.7.1+cu110

After training, the DRL agent performs real-time optimal control of the hydraulic structures, without the need of future tidal inputs (in contrast with state-of-art methods) [105].

5.2.3 Agent Performance Evaluation

With our trained DRL agent, we proceed to compare the yearly energy extraction capabilities of our real-time AI-Driven model against measured data by EDF [130]. In order to do so, tidal predictions for the reference years in [130] are produced with $JTides \times C_f$ and used as inputs for the AI-Driven La Rance model. A comparison between the predicted and measured yearly energy generated by La Rance is shown in Fig. 5.11, where we can see a satisfactory agreement of results, with La Rance's DRL prediction and measurements averaging $521.5MW$ and $507.4MW$, respectively. The 2.6% average gain from the DRL operation is expected, given that energy-oriented optimisation schemes have been shown to attain up to 5% more energy than revenue-oriented optimisation strategies [24].

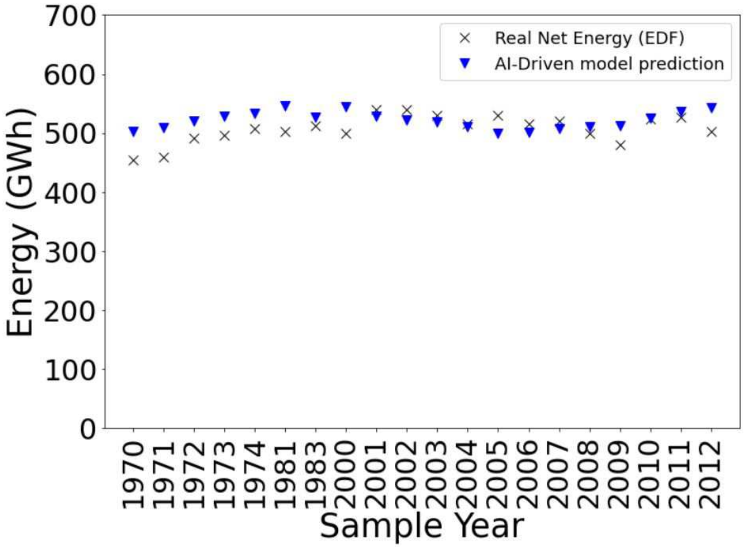

Figure 5.11. Comparison between measured and predicted yearly energy outputs for the La Rance Barrage.

From the collected data provided by the AI-Driven model, we observe that the resulting optimal policy chooses to operate La Rance with a TWP scheme strategy, independently of tidal range (Fig. 5.12). The obtained optimal policy consistently starts turbine pump mode when a positive pump head is still available (Figs. 5.12a and 5.12b), which is (i) not possible for state-of-art methods and (ii) the same strategy adopted in the actual operation of La Rance (Fig. 4.4b). Also exclusive to the AI-Driven model is its capability of fine-tuning power input for turbines in pump mode

(Figs. 5.12c and 5.12d), in a similar manner to what is observed in La Rance (Fig. 4.4b). Furthermore, we note that for tidal amplitudes above $\approx 3\ m$, turbine idling mode is not utilised by the agent (i.e. turbine in power generation mode directly switch to pumping mode), indicating the DRL agent capability of adjusting its strategy according to the observed tidal range.

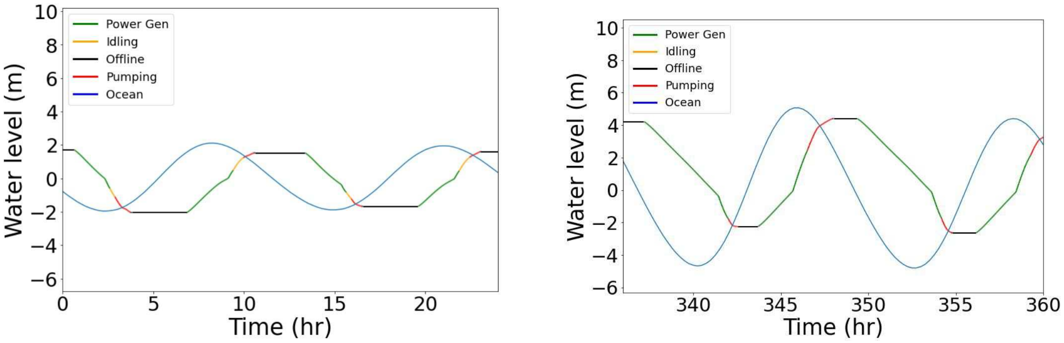

(a) Ocean and lagoon water level variations. **(b)** Ocean and lagoon water level variations.

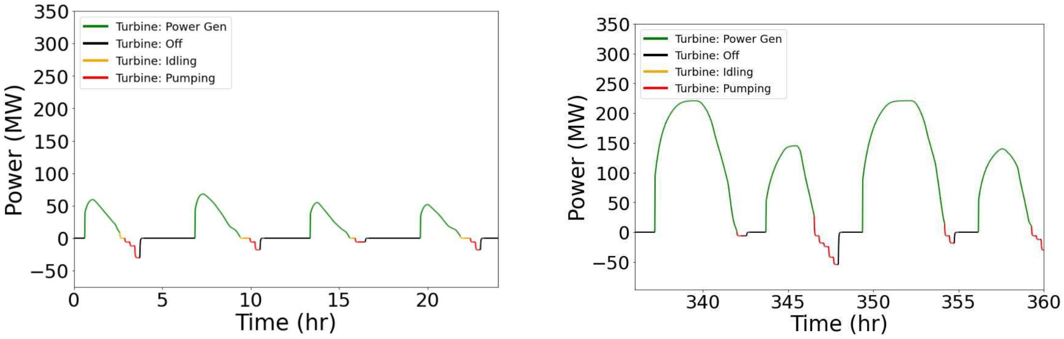

(c) Power output. **(d)** Power output.

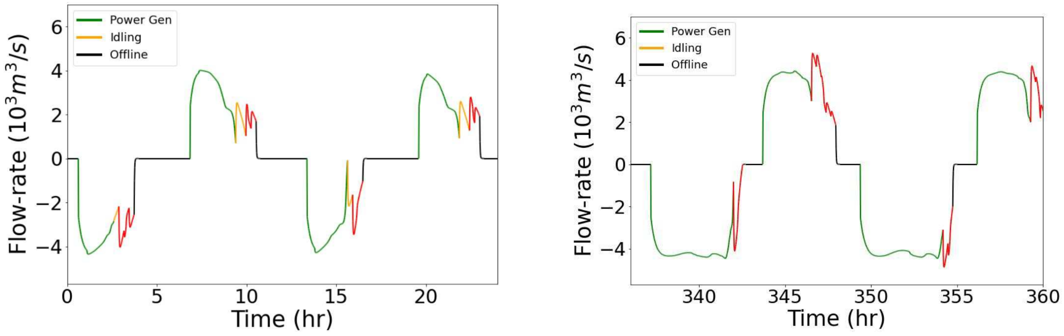

(e) Turbine flow rate. **(f)** Turbine flow rate.

Figure 5.12. Lagoon water level variations, power output and predicted flow rate coloured according to turbine operation from the AI-Driven La Rance model, for small (a, c, e) and large (b, d, f) tidal ranges. In (a, b), ocean is coloured in blue.

Finally, by observing sluice operation in Fig. 5.13, we note that sluices are oper-

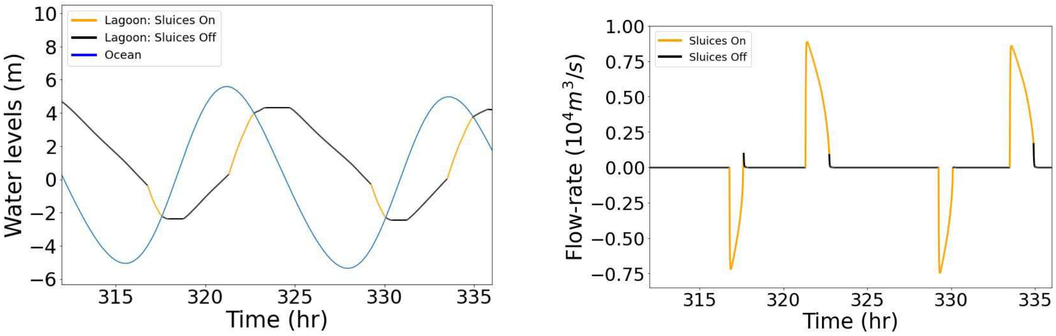

(a) Coloured lagoon water levels following sluice operation by the DRL agent.

(b) Sluice flow rate by the DRL agent.

Figure 5.13. Lagoon water level and sluice flow rate determined by the DRL agent for the La Rance tidal barrage.

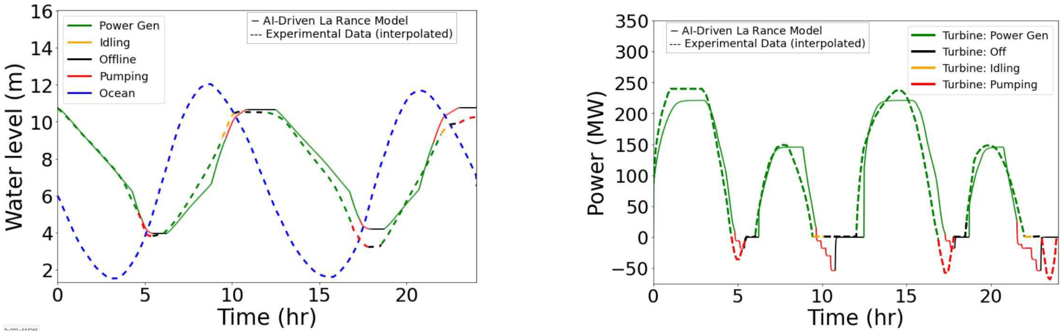

(a) Measured and predicted lagoon water levels for the same real ocean signal.

(b) Measured and predicted power output/input for the same real ocean signal.

Figure 5.14. Comparisons between measured lagoon water level variations and power output/input from La Rance's $T.W.P$ scheme operation against learned operational strategy by the AI-Driven La Rance model.

ated independently from turbines, characterising the variant operation of TRS (shown to be superior to the classical operation of TRS when the goal is maximising power generation [105]).

A comparison of the predicted and measured operation of La Rance, for the same measured tide and starting lagoon water levels from [67], is showcased in Fig. 5.14. From the results, we see how the strategy and predicted lagoon water levels of our real-time AI-Driven model closely resembles the results from the TWP scheme observed in La Rance, with minor differences in operation probably due to the different goal of optimisation (revenue vs power generation). Together with the verified parametric models for La Rance's hydraulic structures in Section 4.2.6, Fig. 5.14 completes our model validation against a real TRS.

5.2.4 Remarks

In this Section, an artificial intelligence (AI) driven representation of a constructed TRS was developed using a 0D TRS parametrised model operated through DRL techniques. Both the 0D TRS model and the AI-Driven representation were then validated, with model predictions showcasing good agreement of results with measured data. For our case study, we utilised the La Rance tidal barrage – the oldest and most successful constructed TRS.

The developed methodologies for reverse engineering the hydraulic structures at La Rance (assembled into the 0D TRS model) are generalisable and can be applied to other constructed TRS. Furthermore, novel representations for the lagoon wetted area, turbines in pump mode and momentum ramp functions can be applied for simulating future TRS projects (e.g., Swansea Bay Tidal lagoon).

Once trained with the goal of maximising the net energy output, the DRL agent extracted as much energy from the 0D TRS model as those reported by La Rance's measurements for a series of analysed years (with an average gain of 2.6%). Such results were possible thanks to the quality of our parametrised 0D TRS model and the operator (DRL agent) capability of delivering: (i) fine-tuned power input for turbines in pump mode; (ii) pumping with positive head differences (aided by gravity); and (iii) independent operation of sluices, which are operational characteristics observed in real operation of TRS. Our AI-Driven model is the first operational optimisation method for TRS capable of delivering this combined control strategy, if compared to state-of-art methods. Also unique to the AI-Driven approach is the real-time flexible control of the hydraulic structures without the need of future tidal predictions, after training a single time only. Results of this work are published in [72].

5.3 SBL (with pumping) Case Study

In this section, we resume the SBL case study, now considering the addition of pumping capabilities – using our developed pump model and the second pump affinity law (Section 4.2.5). Furthermore, we investigate the capabilities of our developed models (variant TWP HC scheme and DRL with pumping method) for operating TRS with the goal of maximising energy or revenue.

For comparison analysis, the results of our models are compared with those obtained using state-of-art methods from the literature. Since the real-time approach for HC methods (developed in Section 4.1.2) rely on forecast data, we need to select a time frame where tidal measurements, high quality tidal predictions and day-ahead market data are available. Considering BODC and the day-ahead Nord Pool available databases, we acquire 11 consecutive months (approximated to 30 days per month) from January 1^{st} 2013 (Fig. 5.15). As in Section 5.1, these data were selected considering scenarios where tidal predictions had good agreement with tidal measurements, according to analysis by BODC.

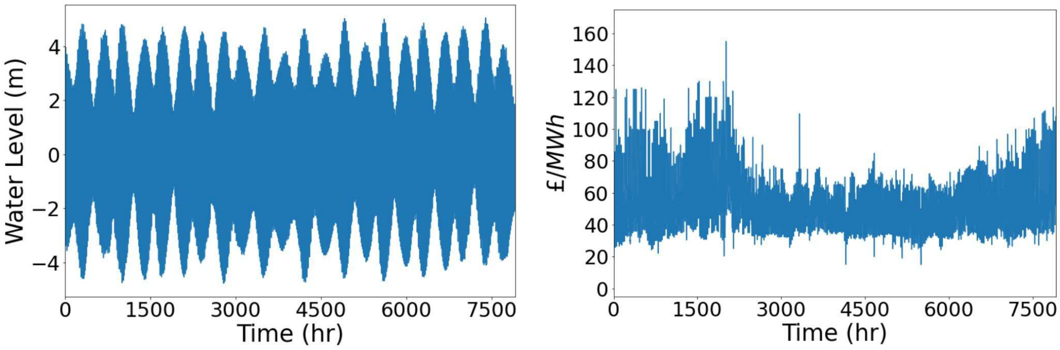

(a) 2013 Ocean Measurements. **(b)** 2013 day-ahead-market values.

Figure 5.15. Swansea Bay ocean measurements and UK day-ahead market measurements for 2013, acquired from [125; 81].

5.3.1 Comparison of Prediction Dependent Methods

With our test data defined, we first compare our variant TWP HC method against the literature state-of-art TWP HC approach by [25]. For both approaches, we consider the EHT optimisation method evaluated in Section 5.1, which presented state-of-art energy extraction capabilities and the most consistent energy fractional gains when subjected to variant TRS operation (Table 5.3 and Fig. 5.4). In this context, our approach and the literature's state-of-art are named $EHTV - P$ and $EHT - P$, respectively.

With the goal of maximising energy generation, a comparison between our $EHTV - P$ method against the literature $EHT - P$ is shown in Fig. 5.16, where the real-time and upper bound estimate of $EHTV - P$ and $EHT - P$ methods are plotted together. In the real-time estimate for both methods, the operational heads acquired from tidal predictions are applied to ocean measurements, while for the upper bound estimate a perfect forecast of the ocean signal is assumed. We observe in Fig. 5.16a that the extracted energy by our method ($EHTV - P$) is systematically and consistently higher than that obtained by the literature ($EHT - P$), either for upper-bound or real-time estimates. The fractional gain for the real-time $EHTV - P$ over the real-time $EHT - P$ estimate is $\approx 2.71\%$, as summarised in Fig. 5.16b for all months sampled. Furthermore, differences between real-time and upper bound estimates (for the same method) are nearly insignificant ($\approx 0.12\%$ and $\approx 0.03\%$ for $EHTV-P$ and $EHT-P$ methods, respectively), which is an indicative of the robustness of the methods during actual operation (and when tidal predictions have a good match with tidal measurements). Finally our real-time $EHTV-P$ method also significantly surpassed the upper bound $EHT-P$ estimate by $\approx 2.68\%$.

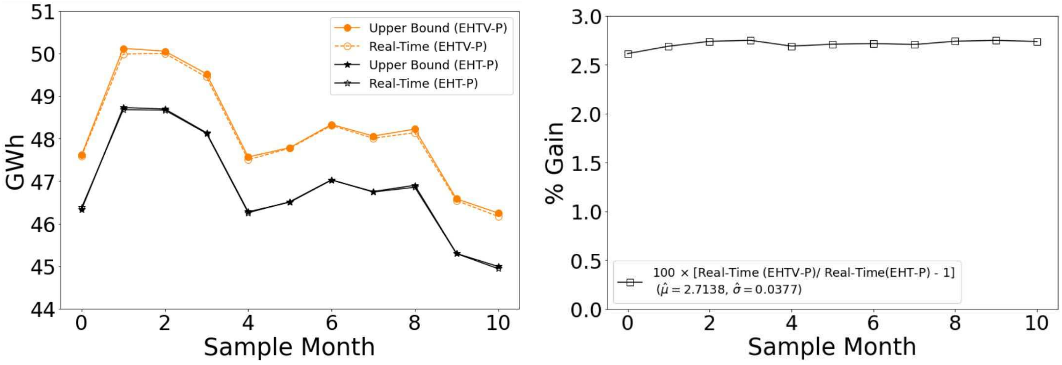

(a) Monthly energy estimates for $EHTV-P$ and $EHT-P$ methods.

(b) Fractional energy gain of $EHTV-P$ (real-time) over $EHT-P$ (real-time) methods.

Figure 5.16. Comparison of energy extraction capabilities between $EHTV-P$ (this work) and $EHT-P$ (literature) methods, for 11 months starting on January 1^{st} 2013.

The same comparison of results is repeated, but considering a revenue-oriented operation (Fig. 5.17). In this case, since we do not have day-ahead market predictions for the analysed period, the real-time estimates for $EHTV-P$ and $EHT-P$ consider an energy-oriented operation. For both upper bound estimates, a perfect forecast of day-ahead market and ocean signal is assumed. As in the preceding case (and shown in Fig. 5.17a), our $EHTV-P$ method gives higher revenue outputs (either for real-time or upper bounds paired comparisons) than the literature's state-of-art $EHT-P$ approach. As shown in Fig. 5.17a, the fractional revenue gain for the real-time $EHTV-P$ method over the real-time $EHT-P$ approach is $\approx 2.72\%$. Furthermore, differences between real-time and upper bound estimates (for the same method) are now significant ($\approx 2.33\%$ and $\approx 2.46\%$ for $EHTV-P$ and $EHT-P$ methods, respectively). This % difference represents the room for improvement for prediction dependent methods, when using accurate day-ahead market forecast. Finally, even though we lack day-ahead market predictions for real-time control, our real-time $EHTV-P$ method slightly surpassed the upper bound $EHT-P$ estimate by $\approx 0.29\%$.

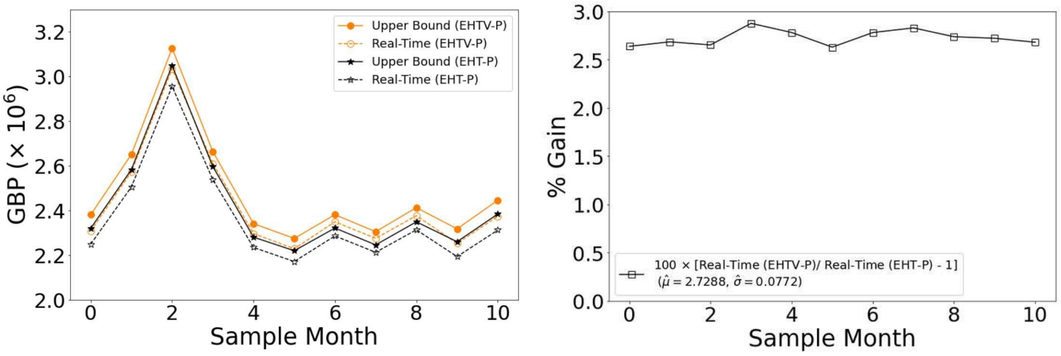

(a) Monthly revenue estimates for $EHTV-P$ and $EHT-P$ methods.

(b) Fractional revenue gain of $EHTV-P$ (real-time) over $EHT-P$ (real-time) method.

Figure 5.17. Comparison of revenue extraction capabilities between $EHTV-P$ (this work) and $EHT-P$ (literature) methods, for 11 months starting on January 1^{st} 2013.

From Figs. 5.16b and 5.17b, we see a significant % gain in both energy and revenue estimates for our $EHTV-P$ method, in comparison with the $EHT-P$ approach. For verifying the statistical significance of these results, we calculate a 95% confidence interval ($C.I$) with a paired t test:

$$C.I : \hat{\mu} \pm t_{[1-\alpha/2;n-1]} \frac{\hat{\sigma}}{\sqrt{n}}, \tag{5.2}$$

where $\hat{\mu}$ and $\hat{\sigma}$ are the sample-based estimates for the mean and standard deviation, respectively, $\alpha = .05$ is the significance level and n the sample size.

For utilising Eq. 5.2, the normality assumption needs to be verified [137]. Therefore, we perform a visual inspection with a Quantile-Quantile (QQ) plot for both case studies (Fig. A.1). Since the absolute and % gain deviations (for energy and revenue case studies, respectively) seem normally distributed, we proceed with the calculation of the $C.I$. Following [137], we have considered the absolute deviation for the energy case study since the normality assumption did not hold for the % deviation. Considering a 95% $C.I$, the gain interval for our variant, real-time $EHTV-P$ method relative to the real-time estimate of the literature $EHT-P$ approach is $[1.2458, 1.2977] GWh$ and $[2.6744, 2.7832]\%$, for energy and revenue extraction, respectively. Projected over a year (and adjusting for 2022 inflation [138]), the revenue increase is estimated to be $967,267£$, lying on the interval $[947,984 - 986,550]£$, with 95% $C.I$.

5.3.1.1 $EHTV - P$ Sensitivity Study

Having proven the superior results of our $EHTV - P$ method, we proceed to perform a sensitivity study, i.e., we compare the deviation between upper bound and real-time estimates for the $EHTV - P$ method. The % and absolute deviation for energy and revenue-oriented scenarios, respectively, can be seen in Fig. 5.18. As with our previous analysis, we have considered the absolute deviation for the revenue-oriented scenario since the normality assumption did not hold for the % deviation [137]. The corresponding QQ plots for both analysis are shown in Fig. A.2. With the normality assumption verified, the 95% $C.I$ for the % deviation between upper bound and real-time estimates, considering an energy-oriented operation with $EHTV - P$, is $[0.0758, 0.1679]\%$. Furthermore, the 95% $C.I$ for the absolute deviation between upper bound and real-time estimates, considering a revenue-oriented operation with $EHTV - P$, is $[42.543 - 70.435]\pounds$.

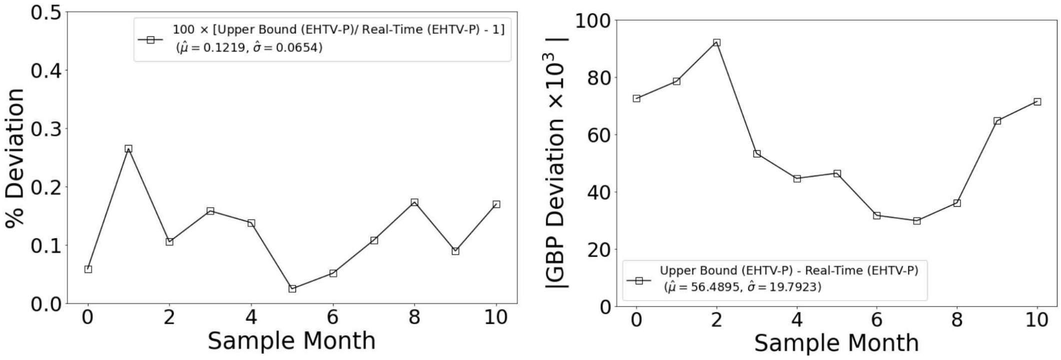

(a) % energy deviation between upper bound and real-time estimates.

(b) Absolute revenue deviation between upper-bound and real-time estimates.

Figure 5.18. Sensitivity study (upper bound vs real-time operation) for the $EHTV - P$ method.

5.3.2 Agent Training Performance Study

In this section, we investigate the capabilities of our DRL agent of operating the SBL with the addition of pumping capabilities. For our initial setup, the MDP implementation for the SBL with pumping follows the same discrete control strategy (and input states) adopted for the La Rance case study (Tables 5.4 and 5.5), the only difference being the power input range for turbines in pump mode, now established as

$$P_{in} : P_{in} = 0.5MW \times n_{oP}, n_{oP} \in \{0, 1, ...16\} \tag{5.3}$$

The P_{in} range in Eq. 5.3, with maximum P_{in} of 8 MW, was chosen since power input never surpassed 8 MW in our initial investigations during training. With the MDP established, the 0D model representation for the SBL with pumping follows the developed Algorithms 1, 2, 3 and 5.

With our initial setup defined, we proceed with a sequence of experiments, with the goal of optimising (i) hyperparameters, (ii) type of neural network representation and (iii) MDP input state/action space optimisation. Due to the time consuming nature of the training sessions ($\approx$ 32 hr, for every run), a thorough exploration of the hyperparameter space is not possible. Therefore, we proceed with a “greedy” approach, where if a variable change yields better results, we keep it for the next training session. Finally, we use our fine-tuned model for investigating the capabilities of a revenue-oriented DRL agent, comparing its performance against our best energy-oriented counterpart.

5.3.2.1 Hyperparameter Tuning

For starting our investigation, the neural network representation and hyperparameters are identical to those used for training La Rance (Section 5.2, published in [72]). Furthermore, similar to the SBL without pumping case study (Section 5.1), the training session utilises 64 instances of the environment for parallel training, with each instance receiving a random generated artificial ocean signal (Eq. 4.29) every episode (simulated month). Initial results show a stable learning experience for the agent, with a learning plateau appearing at the $6e7$ number of collected experiences (vertical dashed line in Fig. 5.19a, for $\beta = 0.01$). With this result, our first goal is to reach a more conclusive training plateau. Similar to our study with La Rance, we start our investigation by varying the β hyperparameter, who is responsible for controlling the magnitude of random explorations performed by the agent during training.

The tuning of β results are shown in Fig. 5.19a, where lower values of β ($\beta =$ 0.007) seem to help the agent converge to a stable plateau more quickly. Since β decreases linearly with the number of observations collected during training, we choose to increase training time – thus increasing the time which the agent is subjected to smaller values of β – to see if a stable plateau can be acquired. Furthermore, since learning rate (α) and policy update limiter (ϵ) hyperparameters (also detailed in Section 2.4.1) are also set to decrease linearly with the number of sampled observations, we have additional supporting evidence to investigate larger training periods. The results of this second attempt, compared to the best performing training scenario of Fig. 5.19a is shown in Fig.5.19b, with both training runs set with ($\beta = 0.007$).

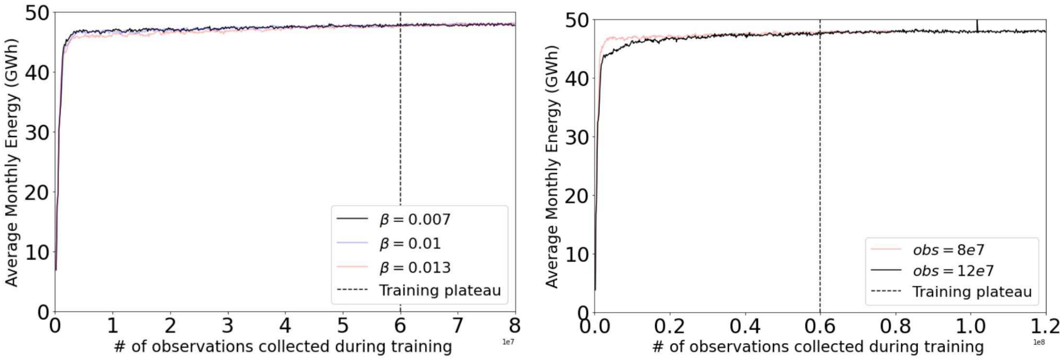

(a) β values for hyperparameter tuning.

(b) Plateau convergence by increasing training time (= 0.007, for both curves).

Figure 5.19. Monthly cumulative energy (reward) in GWh, averaged for all 64 environments during parallel training, for the SBL with pumping environment.

Observing Fig.5.19b, we note that when training starts (and up to $4e7$ collected observations) the training run with $8e7$ collected observations surpasses the longer training run (probably due to a faster decrease of β). However, since the longer training run (with $12e7$ number of collected observations) in Fig.5.19b managed to reach the same stable plateau (vertical dashed line) in half the training time, we consider stopping our hyperparameter tuning of β, α, ϵ and training time duration. For our next stage, we investigate different values for "batch size" and "buffer size" hyperparameters.

While the "buffer size" determines the number of training trajectories to be sampled during training before a "num epoch" number of gradient ascent updates are performed, "batch size" is the number of sampled trajectories randomly sampled from the buffer, for each gradient ascent update. Therefore, "buffer size" and "batch size" are interdependent hyperparameters whose increase in magnitude should lead to more stable (but slower) training. A comparison of our initial values of [batch size, buffer size] $= [256, 4096]$ against an exploratory analysis varying "buffer size" and "batch size" is shown in Figs. 5.20a and 5.20b, respectively.

From Fig. 5.20a, a slight increase of the training plateau is observed when increasing "buffer size" to 8192, therefore we keep the new configuration of [batch size, buffer size] $= [256, 8192]$. It is also noted that further increases of "buffer size" (to 16384) or "batch size" (to 512, in Fig. 5.20b) do not lead to increased performance during training.

Having specified optimum values for β, training time duration (which also influences α and ϵ), "batch size" and "buffer size", we now investigate if further improvements are possible by changing the type of neural network, from an MLP to an LSTM.

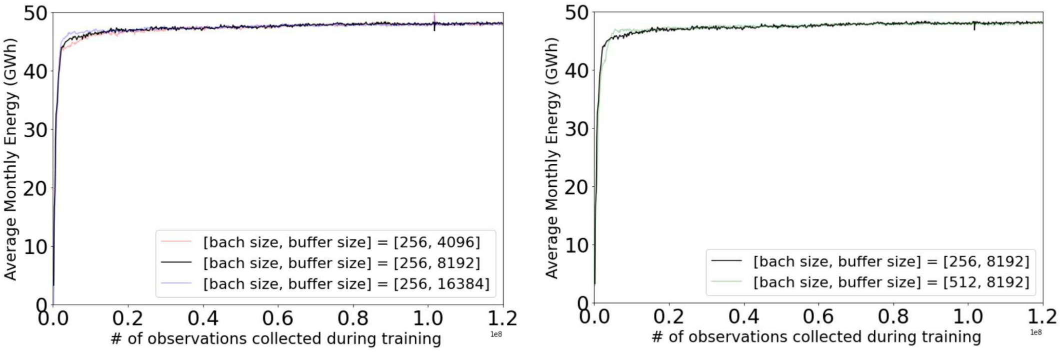

(a) Buffer size values for hyperparameter tuning.

(b) Batch size values for hyperparameter tuning.

Figure 5.20. Monthly cumulative energy (reward) in GWh, averaged for all 64 environments during parallel training, for the SBL with pumping environment.

5.3.2.2 Neural Network Representation

In order to investigate if our problem can take advantage of an LSTM neural network representation, we first inspect how the MLP training performance varies when receiving 1, 2 or 3 input state (s_t) stacked vectors. Although not as efficient as an LSTM, the stacking of s_t at different time-steps effectively increases the policy neural network's memory, so that environment states a "stacked" number of time-steps into the past are also taken into account by the policy neural network when choosing the next action. A comparison between 1, 2 and 3 s_t stacked vectors can be seen in Fig. 5.21.

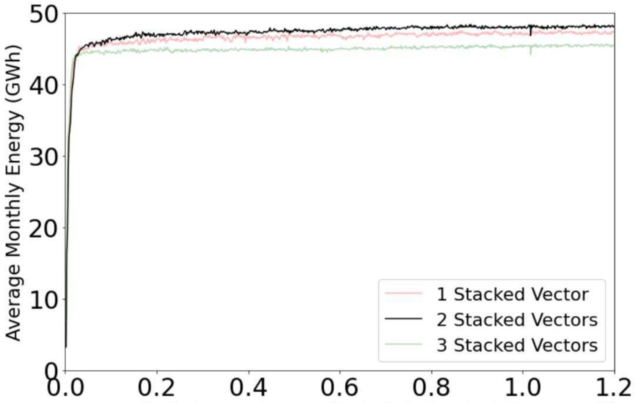

Figure 5.21. Comparison of training performance for 1, 2 and 3 input state stacked vectors. In the "y-axis" the monthly cumulative energy (reward) in GWh, averaged for all 64 environments during parallel training, for the SBL with pumping environment is shown.

From Fig. 5.21, we notice how training results improve as we move from 1 to 2 stacked s_t vectors. However, when 3 s_t vectors are used ($[s_t, s_{t-1}, s_{t-2}]$), a significant drop in performance is observed. This indicates that the optimal policy for maximising energy generation would not rely on input states older than s_{t-1}. For the sake of confirming this hypothesis, we substitute the MLP neural network representation for an LSTM (considering the same number of units and hidden layers). For exploring our hyperparameter space, we focus on adjusting "sequence length" and "memory size" hyperparameters. While "sequence length" determines the number of time-steps ($\Delta t = 15\ min$) into the past that will be considered by the neural network when determining the next action, "memory size" is the amount of memory (size of the array of floating point numbers) used by the LSTM for storing useful information from previous states. Results from our experiments can be seen in Fig. 5.22.

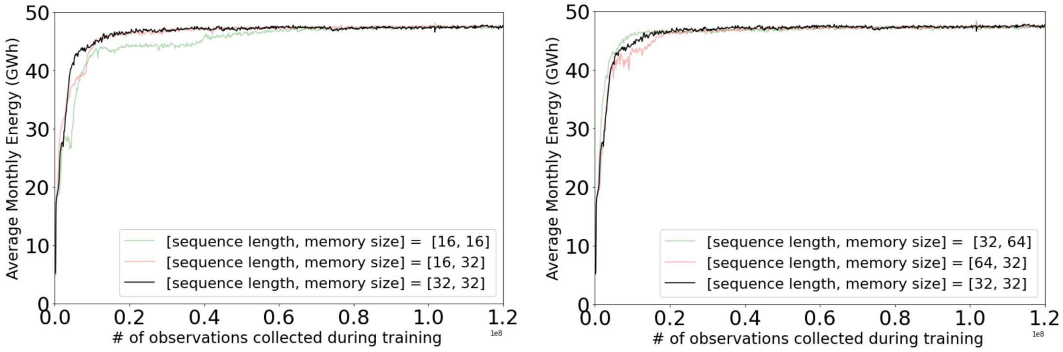

(a) **(b)**

Figure 5.22. Fine tuning of LSTM "sequence length" and "memory size" hyperparameters. In the "y-axis" the monthly cumulative energy (reward) in GWh, averaged for all 64 environments during parallel training, for the SBL with pumping environment is shown.

From Fig. 5.22, the best results when training the LSTM occur when [sequence length, memory size] $= [32, 32]$. However, as seen in Fig. 5.23, training performance is still inferior to what is observed when utilising the MLP neural network. The superior performance of the MLP neural network is an indicative that, although we are using temporal data (e.g. ocean water levels) as input states, these are highly predictable. Therefore, we keep the MLP neural network representation for our next analysis.

5.3.2.3 MDP input state/action space optimisation

Having specified the MLP neural network representation, we now investigate if there are more optimum MDP configurations beyond that defined in our approach. We are motivated to return to our MDP once our current model can only provide a specific range of discrete power input values to the turbines during pump mode. For tackling

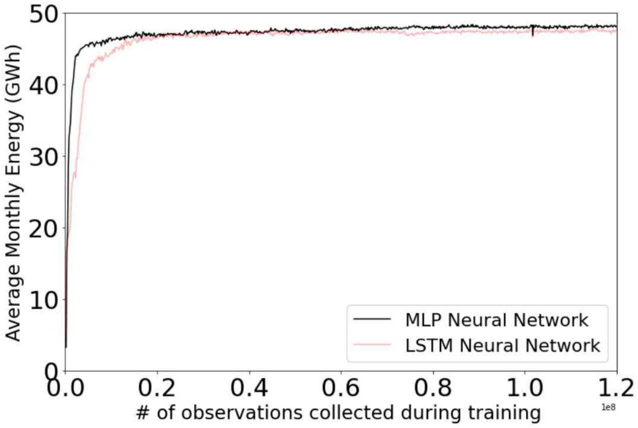

Figure 5.23. Training performance comparison between optimised MLP (2 stacked vectors) and LSTM ([sequence length, memory size] = $[32, 32]$) neural network representations.

this, we propose a hybrid control configuration, with both discrete and continuous actions at the last layer of the MLP neural network. Since initial trial results with this method showcased poor performance, we investigate modifications to the input state s_t (shown in Table 5.5), while keeping the hybrid control option. A comparison of results, when varying s_t can be seen in Fig. 5.24.

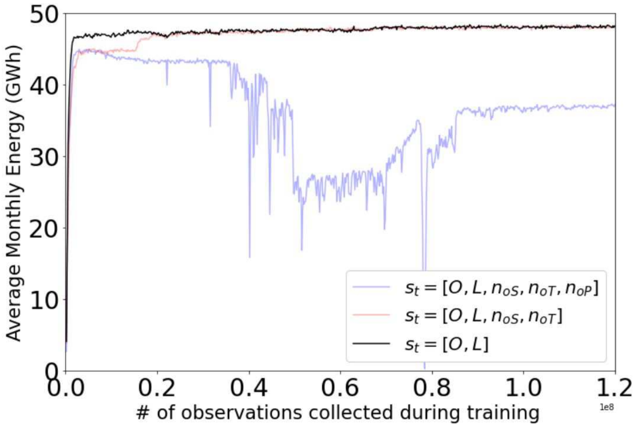

Figure 5.24. Hybrid control training performance comparison considering different input state s_t configurations (MLP neural network).

As seen in Fig. 5.24, a significant jump in performance is observed for the hybrid control strategy once the input state is reduced to only consider ocean and lagoon water levels (from the current and previous time-steps). This is an interesting result, since it means that the agent does not require knowing its previous selected actions for optimally operating the TRS (differently from some operational optimisation methods from the literature, such as the EHN method). A comparison of the optimised hybrid control option, against the best discrete control option is seen in Fig. 5.25. We notice that even though training convergence is faster with the hybrid control option, the final training plateau reached for both strategies is the same.

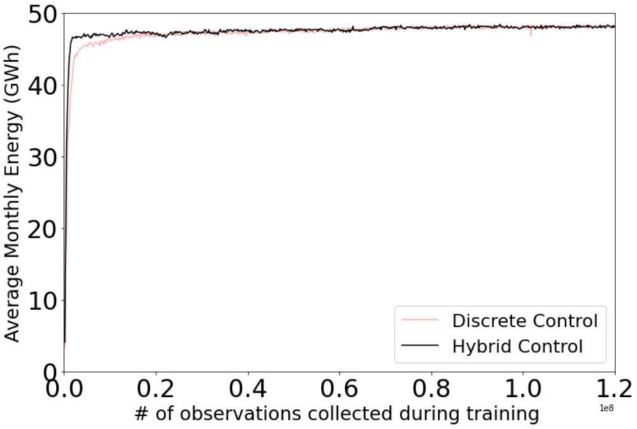

Figure 5.25. Training performance comparison between hybrid and discrete control strategies.

Further investigations, such as using $s_t = [O, L]$ for the case of discrete actions, or if the hybrid control strategy could benefit from an LSTM neural network representation are shown in Fig. 5.26. We note that, under the new input state ($s_t = [O, L]$) a extremely similar training behaviour is observed for either hybrid or discrete control options, when using an MLP neural network. Furthermore, the usage of an LSTM neural network (with [sequence length, memory size] = $[32, 32]$) reduces the hybrid model training performance.

Although no improvements to the training plateau are observed after the "buffer size" optimisation step in Fig. 5.20a, we consider the MLP hybrid control strategy with $s_t = [O, L]$ as the best model for energy-oriented optimisation, due to its faster plateau convergence and the flexibility of providing any real value of power input for turbines in pump mode. The final MDP input space and action space for the hybrid control TRS model are shown in Tables 5.8 and 5.9, respectively. In the next session, we investigate the possibility of revenue-oriented control by the agent, starting from the optimised hybrid control strategy.

Table 5.8. Stacked input states s_t for the PPO neural network, for the SBL with pumping and energy-oriented control.

States (at times t and $t-1$)	Units
Ocean water level	"Normalised" $[0, 1]$ (float)
Lagoon water level	"Normalised" $[0, 1]$ (float)

Table 5.9. Hybrid control options for the SBL, with pumping.

Branch	Action
n_{oS} (sluices) Discrete Action	Offline Online
n_{oT} (turbines) Discrete Action	Offline Power Generation (Ebb, or Flood) Idling Ebb Pumping Flood Pumping
n_{oP} (pump P_{in}) Continuous Action	$P_{in} = n_{oP} \times P_{inMax}$, where $0 \leq n_{oP} \leq 1$ and $P_{inMax} = 8\ MW$

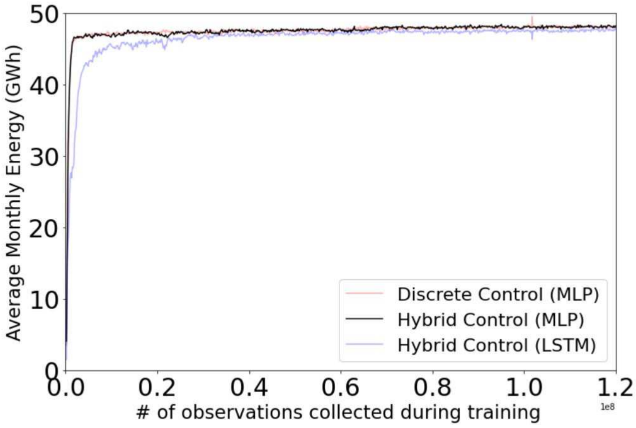

Figure 5.26. Training performance comparison between optimised hybrid (MLP), discrete (MLP) and hybrid (LSTM) configurations, considering the reduced input state space $s_t = [O, L]$.

5.3.2.4 Investigation of Revenue-Oriented Control

Considering the optimal hybrid control method acquired for energy-oriented operation, we investigate the DRL method's capability of operating TRS with the goal of maximising revenue.

For doing so, we start by utilising the day-ahead market measurements from 2019 (Fig. 5.27, acquired from [81]) as training data, together with our random artificially generated ocean signal for the SBL, used for energy-oriented optimisation. The 2019 day-ahead market data was initially chosen among a 2014 – 2021 database (Fig. A.5) for presenting relatively consistent price variations (apart from a single spike in energy price, which is substituted by the second highest day-ahead market peak for that year, during training). When starting a training episode each environment instance receives one randomly sampled month of the modified day-ahead data (Fig. 5.27b). Then, for every 15 *min* of TRS simulation, the current value of electricity, plus our random generated ocean signal and current lagoon water level are fed to the policy neural network (Table 5.10). For normalising the day-ahead market data, the maximum recorded value from "peak-removed" 2019 (Fig. 5.27b) is used as reference. Finally, since we want our agent to perform actions based on the TRS's revenue, the updated reward for the agent equals the 15 min sum of a minute-wise product between the energy generated times the value of that energy (in $£/MWh$), as dictated by the day-ahead market.

Table 5.10. Stacked input states s_t for the PPO neural network, for the SBL with pumping and revenue-oriented control.

States (at times t and $t-1$)	Units
Ocean water level	"Normalised" $[0, 1]$ (float)
Lagoon water level	"Normalised" $[0, 1]$ (float)
Day-Ahead Market ($£/MWh$)	"Normalised" $[0, 1]$ (float)

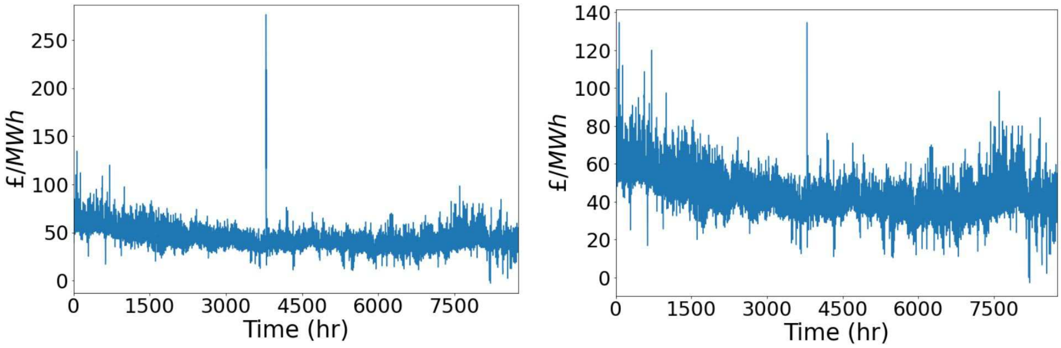

(a) Original 2019 day-ahead data. **(b)** "Peak-removed" 2019 day-ahead data.

Figure 5.27. 2019 day-ahead market measurements. Obtained from [81] database.

Similarly to our investigation with the energy-oriented approach, we see significant improvements during training by reducing β values from $\beta = 0.007$ to $\beta = 0.0007$ (Fig. 5.28a). Furthermore, the acquired $\beta = 0.0007$ appears to be an optimal value, since further increasing/decreasing β slightly decreases training performance (Fig. 5.28b). After finding an optimal β value, we investigate in Fig. 5.29 if our problem can benefit from an LSTM architecture with hyperparameters [sequence length, memory size] = $[32, 32]$. Since training performance decreases when using an LSTM, we continue our study with an MLP representation.

For our next evaluation, we compare the training performance of the model when utilising (i) our random ocean signal generator, versus (ii) tidal predictions from JTides from 2019 (concurrent with day-ahead market data), at the location of the SBL. We

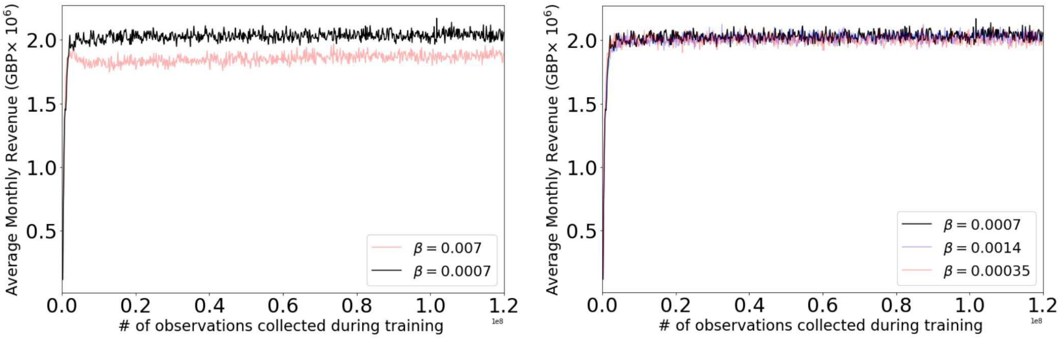

(a) **(b)**

Figure 5.28. β hyperparameter optimisation for our revenue-oriented DRL agent.

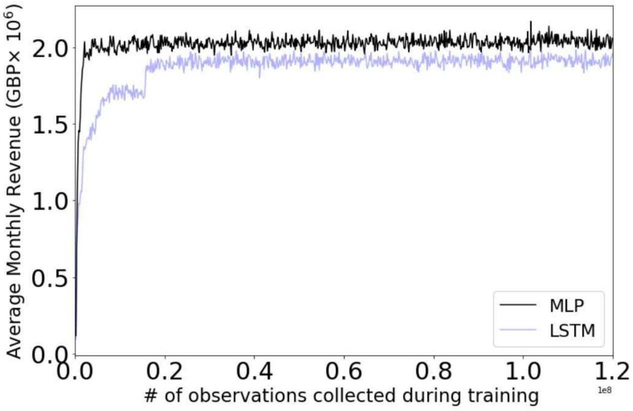

Figure 5.29. Training performance comparison between optimised MLP and LSTM neural network architectures (2019 day-ahead data).

have used JTides, since only "improbable" (i.e. not reliable) labelled measurements from BODC were available for the year of 2019 (or for any year from June 2014, onward). Results from this comparison are shown in Fig 5.30a. Even though training results show reduced performance when using JTides 2019 prediction, we also compare the outputted results for the 2013 test scenario (Fig. 5.30b), where the trained JTides model showcased consistent worse results, averaging 2.17% less revenue when compared with the random ocean model approach.

For our final evaluation, we see if we can observe improvement of results when using all day-ahead market data available (from 2014 – 2021, as shown in Fig. A.5), while using the random ocean representation. For this, we compare both MLP (with 2 stacked input vectors) and LSTM neural network representations (with [sequence length, memory size] = $[32, 32]$). Consistent with our previous analysis, the MLP representation showcases superior results (Fig. 5.31).

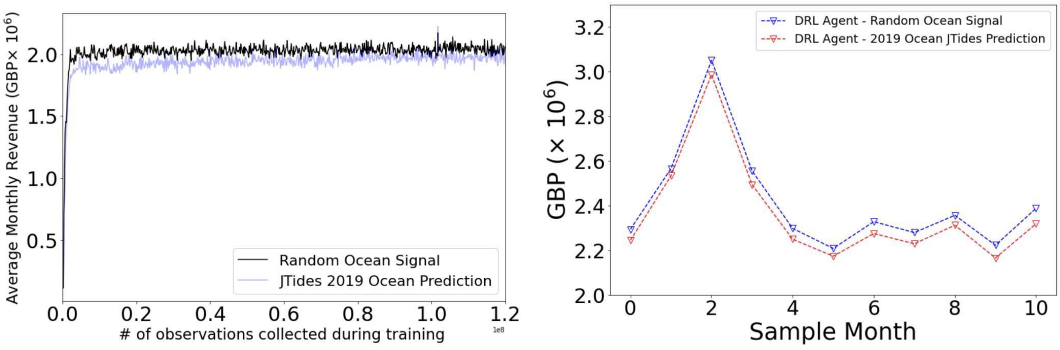

(a) Comparison during training.

(b) Comparison during testing (2013 ocean and day-ahead data)

Figure 5.30. Training (and testing) performance comparison between our random ocean signal representation and 2019 ocean prediction from JTides.

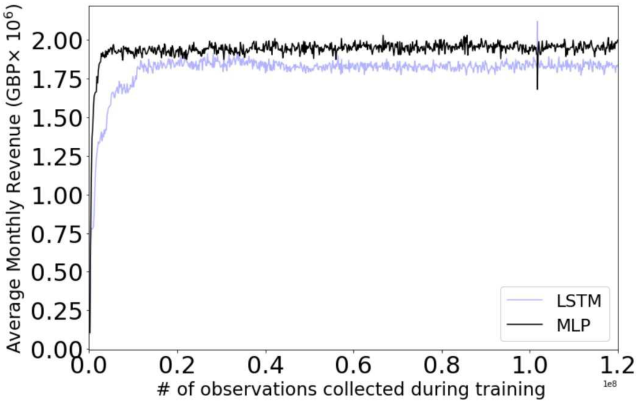

Figure 5.31. Training performance comparison between optimised MLP and LSTM neural network architectures (2014 to 20121 day-ahead data).

Even though our last evaluation significantly increases the available day-ahead training data, the revenue-oriented agent showcases reduced performance for the 2013 testing phase, in comparison with the previous "peak-removed" 2019 data (Fig. 5.32). A possible explanation for this is the very inconsistent price variations of some of the years considered as training data for the final analysis (Fig. A.5). With this result, we keep the DRL agent trained with the "peak-removed" 2019 data (and MLP neural network representation plus random ocean input) as the best revenue-oriented model.

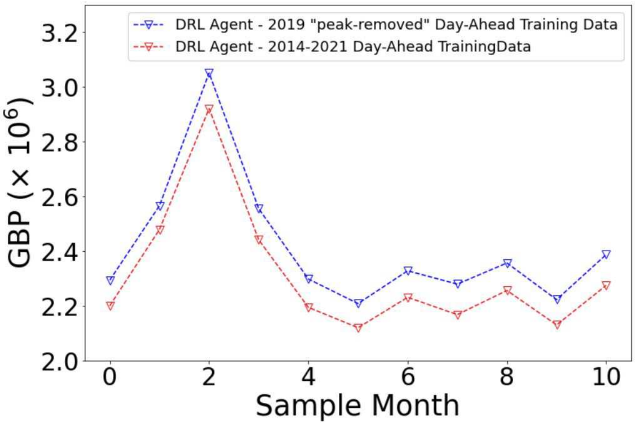

Figure 5.32. Test performance comparison when using two different day-ahead market training data, namely (i) "peak-removed" 2019 data versus (ii) $2014-2021$ data.

5.3.3 Agent Test Performance Study

5.3.3.1 Energy-Oriented Vs Revenue-Oriented DRL Agent.

With our trained (and optimised) policy neural network representations for both energy and revenue-oriented control for the SBL (with pumping) case study (see Table 5.11, for full detailing), we compare the energy and revenue extraction capabilities of our developed approaches (Fig. 5.33).

Table 5.11. Energy and Revenue-Oriented DRL Agent Characteristics.

Model Characteristics	Energy-Oriented	Revenue-Oriented
Training Data	Random Ocean Signal	Random Ocean Signal and 2019 "Peak-Removed" Day-Ahead Market
Neural Network Representation	MLP	MLP
Type of Operation (Table 5.9)	Hybrid Control	Hybrid Control
Input State s_t	$[O, L]$	$[O, L, V]$, where V is the Day-Ahead Market Energy Value in $£/MWh$
# of Stacked s_t Vectors	2	2

As seen in Fig. 5.33, our energy-oriented agent surpasses both energy and revenue extraction capabilities of the revenue-oriented approach. Indeed, while the revenue-oriented agent slightly surpasses the energy-oriented revenue in test months $0, 1, 7$ and 8 (Fig. 5.33b), on average it generates 0.27% less revenue than the energy-oriented counterpart. When looking at energy production, the energy-oriented agent generates 3.05% more energy on average than the revenue-oriented counterpart (Fig. 5.33a). From this comparison of results we note that the trained revenue-oriented agent's strategy is capable of yielding similar revenue outputs as the energy-oriented agent, while generating significantly less energy. However, since the energy-oriented agent is the winner on both scenarios, we keep it for our comparison against state-of-art and enhanced state-of-art methods.

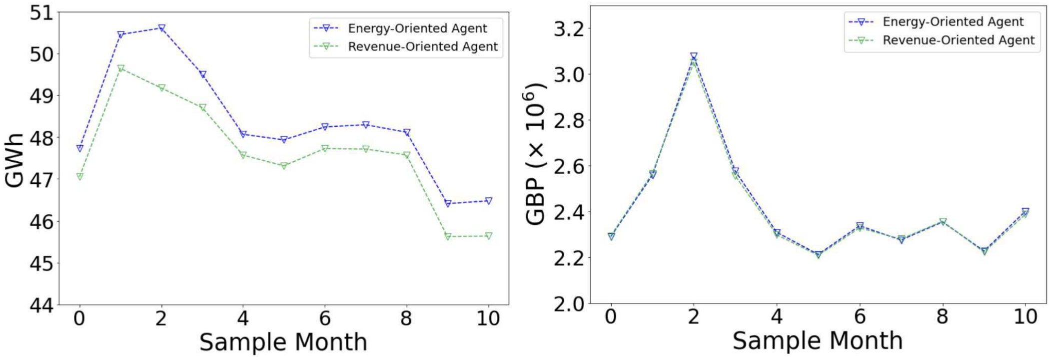

(a) Monthly energy generation during testing.)

(b) Monthly revenue during testing.

Figure 5.33. Test performance comparison between optimised revenue and energy-oriented trained DRL agents.

5.3.3.2 DRL Agent Vs Literature's State-Of-Art Estimate.

A comparison of the energy and revenue extraction capability of our energy-oriented DRL agent against the upper bound and real-time estimate (energy and revenue-oriented) of state-of-art methods with pumping ($EHT - P$) is shown in Fig. 5.34. Additionally, the respective fractional gain of our energy-oriented DRL agent, relative to the real-time estimate of $EHT - P$ is shown in Fig. 5.35. Since the % gain results for both case studies seem normally distributed (Fig. A.3), we proceed with the calculation of the $C.I.$ Considering a 95% $C.I$, the % gain interval for our energy-oriented agent relative to the real-time estimate of the literature $EHT - P$ approach is $[2.8190, 3.5238]\%$ and $[1.8286, 3.081]\%$, for energy and revenue extraction, respectively. Projected over a year (and adjusting for 2022 inflation [138]), the revenue increase is estimated to be $861,301£$, lying on the interval $[648,177 - 1,092,110]£$, with 95% $C.I.$ It is important to note that the superior results of our energy-oriented DRL agent were obtained without tidal predictions, after training a single time only and operating in real-time. Furthermore, even though our agent does not have access to tidal or day-ahead market forecasts, it manages to surpass the upper-bound of $EHT - P$ approach by 3.14% (on energy) and 0.03% (on revenue), considering energy-oriented and revenue-oriented operations, respectively.

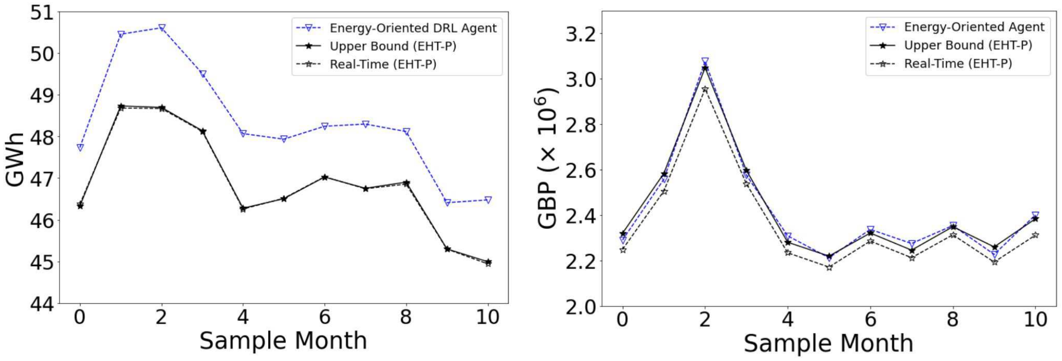

(a) Monthly energy during testing.

(b) Monthly revenue generation during testing.

Figure 5.34. Test performance comparison between our energy-oriented DRL agent and the upper bound and real-time estimates of the state-of-art method ($EHT-P$).

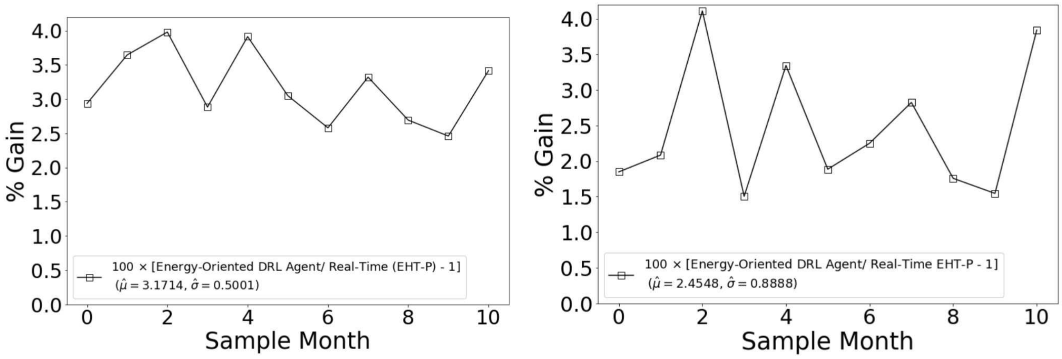

(a) DRL agent fractional energy gain over the real-time $EHT-P$ estimate.

(b) DRL agent fractional revenue gain over the real-time $EHT-P$ estimate.

Figure 5.35. Fractional gain of energy-oriented agent over real-time $EHT-P$ method.

5.3.3.3 DRL Agent vs Real-Time $EHTV-P$ Method

For a final evaluation, we compare energy and revenue extraction capabilities of the best control optimisation methods of this thesis, namely, both our developed (i) energy-oriented DRL agent and (ii) the variant, real-time $EHTV-P$ method.

A comparison of the energy and revenue extraction capability of the energy-oriented DRL agent against the real-time $EHTV-P$ approach is shown in Fig. 5.36, with the respective fractional gain (relative to $EHTV-P$) in Fig. 5.37. Since the % gain results for both case studies seem normally distributed (Fig. A.4), we proceed with the calculation of the $C.I$. Considering a 95% $C.I$, the % gain interval for our energy-oriented agent relative to the real-time $EHTV-P$ method is $[0.0978, 0.7932]\%$ and

$[-0.8877, 0.3546]\%$, for energy and revenue generation, respectively. From this result we see that our energy-oriented DRL agent surpasses energy extraction capabilities of the real-time $EHTV - P$ method (by 0.4455%, on average). Furthermore, we cannot reject the null hypothesis that both methods produce the same revenue, on average.

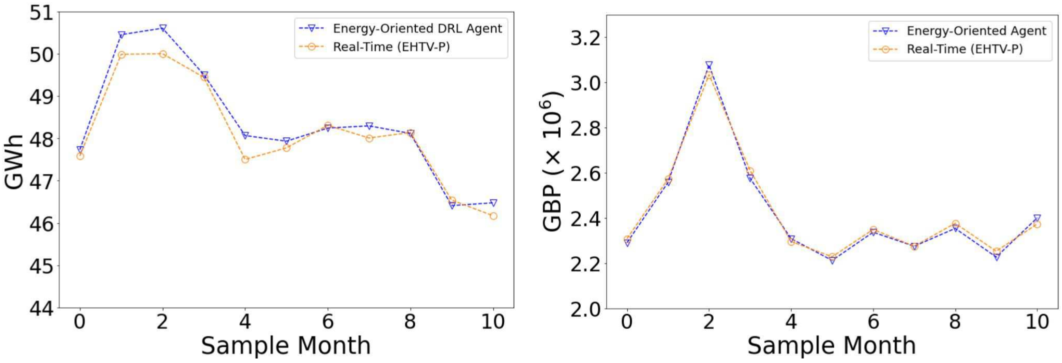

(a) Monthly energy during testing.

(b) Monthly revenue generation during testing.)

Figure 5.36. Test performance comparison between our developed (i) energy-oriented DRL agent and (ii) real-time $EHTV - P$ method.

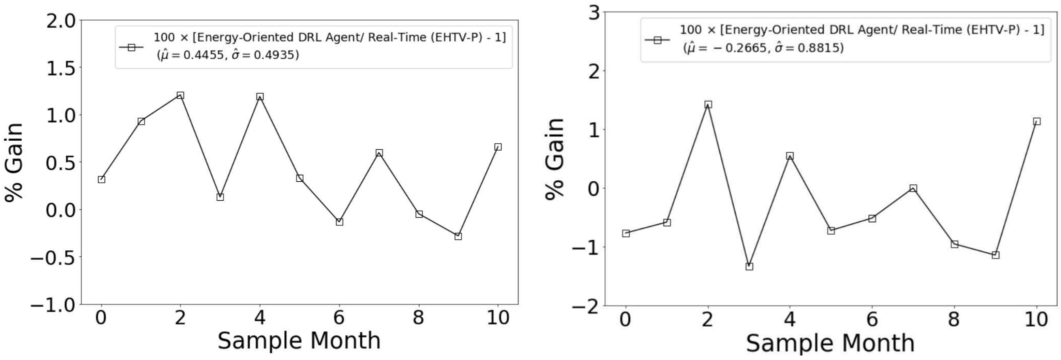

(a) Fractional energy gain.

(b) Fractional revenue gain.

Figure 5.37. Fractional gain of energy-oriented agent over $EHTV - P$ (real-time) method.

5.3.4 TRS Operation Comparison

Having evaluated the energy and revenue extraction performance of our $EHTV - P$ and DRL agent methods versus the state-of-art ($EHT - P$) approach, we analyse in this section the characteristics of the operational strategy of each method and its impact on lagoon water level evolution and power/revenue generation. It is important to highlight that, while for $EHTV - P$ and $EHT - P$ methods, the TRS scheme of operation (in this case TWP) needs to be specified beforehand, the DRL agent learns the operational scheme by itself during training.

In Fig. 5.38, lagoon water level variations following turbine operation for the DRL agent, real-time $EHTV - P$ method and the energy and revenue-oriented state-of-art upper bound ($EHT - P$) are presented for the same ocean signal.

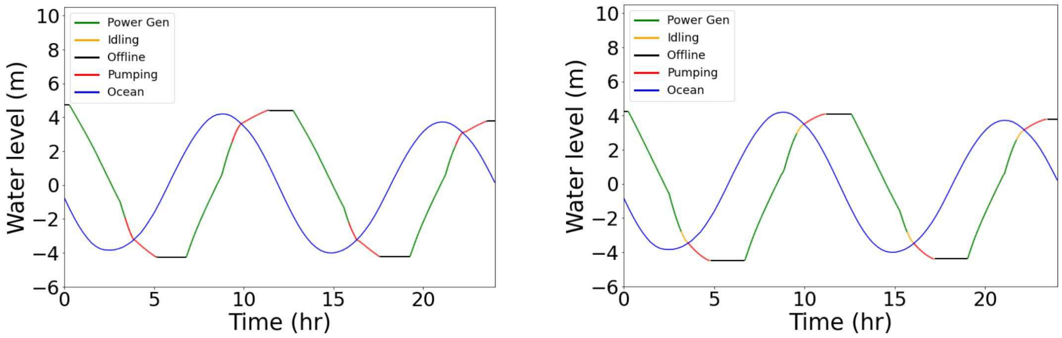

(a) DRL agent, hybrid control. **(b)** Energy-oriented $EHTV-P$ (real-time).

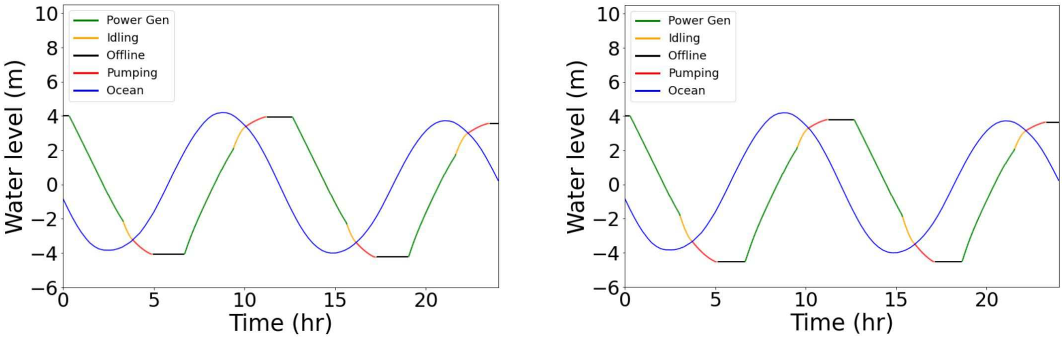

(c) Energy-oriented $EHT - P$ (upper bound). **(d)** Revenue-oriented $EHT - P$ (upper bound).

Figure 5.38. Lagoon water level variations coloured according to turbine operation. Ocean is coloured in blue.

While all HC methods showcase a similar pattern of operation (Figs. 5.38b, 5.38c and 5.38d), with pumping occurring after idling turbine stage and at negative head differences only, the DRL agent is able to completely skip the idling turbine stage and pump at positive head (gravity assisted) head differences. By starting the pumping stage earlier, the TRS experiences a greater lagoon water head, which leads to increased power generation, if compared to $EHTV - P$ and $EHT - P$ methods (Fig. 5.39). Indeed, by looking at Fig. 5.39a, we note that during ebb generation stages, power quickly plateaus at 320 MW, suggesting that the DRL agent could further extract more power of the TRS, if not limited by the current installed capacity of the SBL project. Furthermore, by comparing the DRL agent (Fig. 5.39a) turbine pumping stage against HC methods (Figs. 5.39b, 5.39c and 5.39d), we see the hybrid agent capability of dynamically providing fine-tuned power input to turbines in pump mode (Table 5.9), according to ocean and lagoon water levels input states (at times t and $t-1$, as presented in Table 5.8). Both fine-tuned operation of turbines in pump mode and pumping with positive (gravity-assisted) differences are characteristics observed in La Rance's measurements [67], but so far no TRS control method available to academics has been able to emulate this behaviour.

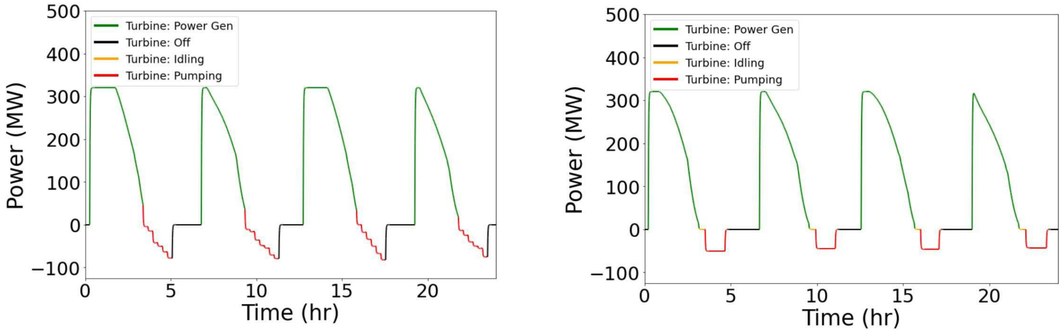

(a) DRL agent, hybrid control.

(b) Energy-oriented $EHTV-P$ (real-time).

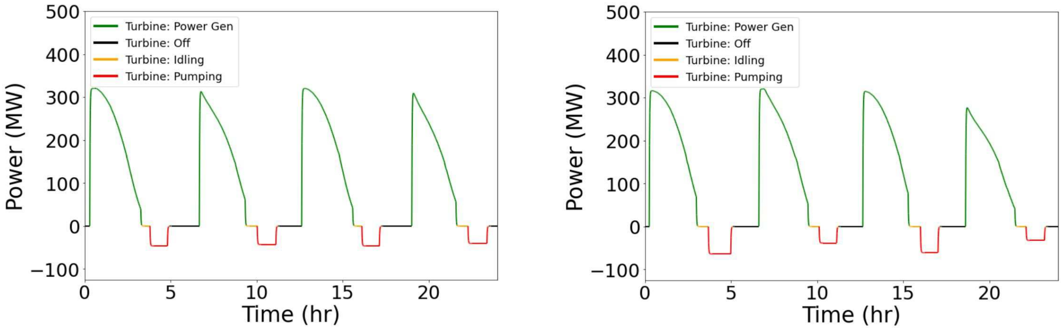

(c) Energy-oriented $EHT - P$ (upper bound).

(d) Revenue-oriented $EHT - P$ (upper bound).

Figure 5.39. Power output/input coloured according to turbine operation.

In Fig. 5.40, lagoon water level variations following sluice operation are presented. Comparing ours (i) DRL agent and $EHTV - P$ methods (Figs. 5.40a, 5.40b), against (ii) state-of-art $EHT - P$ approach (Figs. 5.40c, 5.40d), we can see how the variant (independent) operation of sluices in (i) can help increase lagoon water head differences for the following operation (therefore increasing power generation), by starting the sluicing stage while turbines are still in power generation mode.

Finally, Fig. 5.41, showcases the acquired revenue (in $£/min$) following turbine operation for each adopted strategy. Since our best real-time approaches (DRL agent

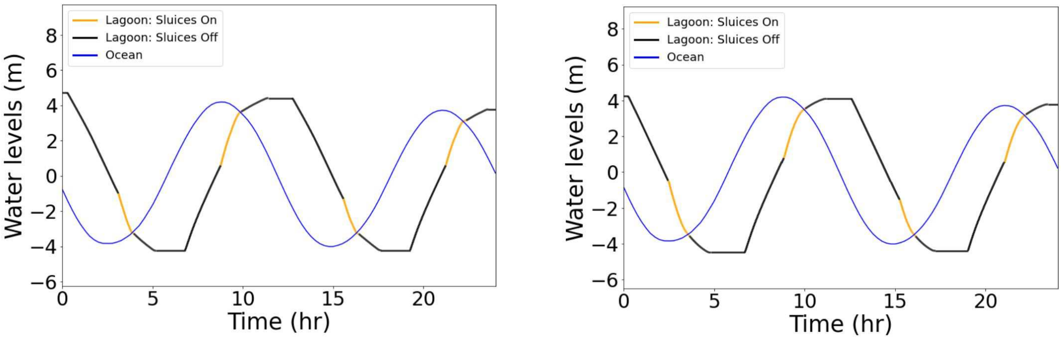

(a) DRL agent, hybrid control. **(b)** Energy-oriented $EHTV-P$ (real-time).

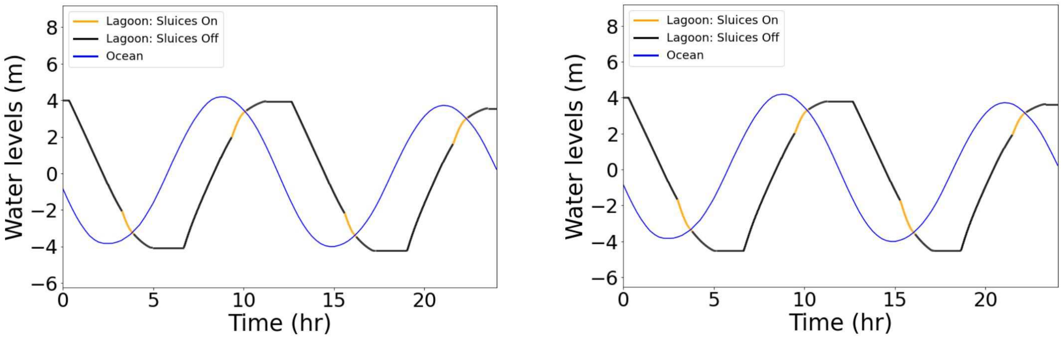

(c) Energy-oriented $EHT-P$ (upper bound). **(d)** Revenue-oriented $EHT-P$ (upper bound).

Figure 5.40. Lagoon water level variations coloured according to sluice operation. Ocean is coloured in blue. In a) and b), sluicing stage is independent of turbine idling stage, i.e. it can occur earlier than for c) and d), during turbine power generation or pumping stages.

and $EHTV-P$ method) are energy-oriented, day-ahead market fluctuations are not considered during operation. Conversely, the revenue-oriented upper bound $EHT-P$ (Fig. 5.41d) assumes a perfect tide and day-ahead market forecast, resulting in higher revenues than the energy-oriented upper bound counterpart (Fig. 5.41c). Although our current study did not find improvements in the DRL agent revenue-oriented operation, the fact that there is theoretical room for extracting more revenue when considering day-ahead forecast should be investigated. As a suggestion for future work, our DRL agent could be trained considering energy demand fluctuations (instead of day-ahead market), or a day-ahead forecast model. Furthermore, the real-time $EHTV-P$ capabilities should be tested, when utilising day-ahead market predictions.

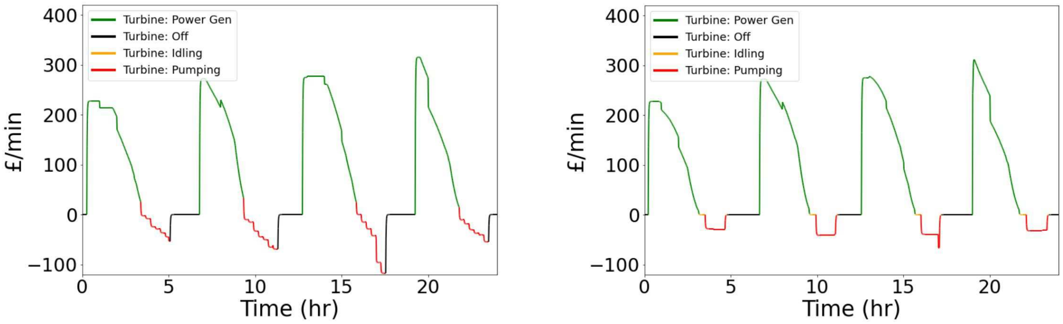

(a) DRL agent, hybrid control.

(b) Energy-oriented $EHTV-P$ (real-time).

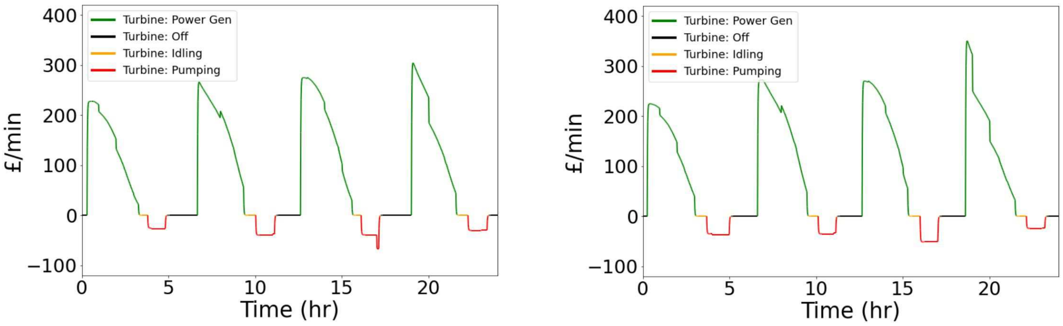

(c) Energy-oriented $EHT - P$ (upper bound).

(d) Revenue-oriented $EHT - P$ (upper bound).

Figure 5.41. Extracted revenue in $\pounds/min$ coloured according to turbine operation.

Chapter 6

Conclusion and Future Directions

In this thesis, we have proposed two new techniques for automatic (real-time) flexible control of Tidal Range Structures (TRS) that can surpass energy and revenue extraction capabilities of current (non-automatic) state-of-art methods. For our first approach (a), enhancements to the state-of-art operational optimisation approach from the literature were presented, enabling independent sluice operation and automatic control (utilising tidal prediction estimates). For our second approach (b) we used Deep Reinforcement Learning (DRL) techniques (more specifically Proximal Policy Optimisation) for training a smart agent using Unity ML-Agents. After training, our agent was shown to be capable of optimally operating TRS without tidal predictions, performing flexible control of TRS (dynamically adjusting its strategy, according to the experienced tidal range), after training a single time only. Furthermore, when applying the DRL technique to TRS with pumping capabilities, our agent managed to operate TRS in a more realistic manner than state-of-art (and enhanced state-of-art methods), due to its ability of (i) dynamically adjusting power input to turbines in pump mode and (ii) pumping with positive head differences.

As case studies for our work, we focused on the 0D modelling of two TRS projects: the constructed La Rance Tidal Barrage and the Swansea Bay Tidal Lagoon (SBL) pathfinder project. Through parametric and affinity laws techniques, a 0D model representation of La Rance was created, and both hydraulic representations and TRS operation (by the DRL agent) were validated against measured data, making this (to the best of our knowledge), the first 0D operated TRS model validated for a constructed TRS. Indeed, while the developed 0D La Rance model could be operated with state-of-art approaches, the operational strategy developed by our DRL agent enabled (i) fine-tuned pump power input and (ii) pumping with positive head differences. Characteristics (i) and (ii) are observed during the operation of real TRS (e.g. La Rance) but, so far, have not been possible to emulate using state-of-art methods. Furthermore, the development of this model enabled a generalisable description of pump operational modes and novel transition ramp functions (for switching between TRS operational modes), which were implemented into the 0D model of the SBL.

For the SBL case study (and considering an energy-oriented operation for all real-time methods), we compared state-of-art, (a) enhanced state-of-art and the (b) DRL agent methods, considering with and without pumping scenarios. For the "without pumping" scenario, our developed (a) and (b) methods showcased superior results in energy generation to state-of-art approaches, while operating the TRS in real-time. When including pumping capabilities to the TRS (therefore increasing power generation for all methods) approaches (a) and (b) surpassed the upper bound (perfect tidal forecast scenario) energy estimate of the state-of-art approach by 2.61% and 3.14%, respectively – yielding a revenue gain up to $967,267£$ per year. Comparing all approaches, the DRL agent was shown to be the best method, surpassing the enhanced state-of-art method by $\approx 0.45\%$ in energy generation, while being equivalent in revenue extraction capabilities and operating the TRS without requiring tidal predictions.

A comparison of the trained DRL agent performance against state-of-art and enhanced state-of-art methods, considering revenue-oriented operation was also undertaken by using day-ahead market data as both training and test data. In our investigations, the trained energy-oriented agent surpassed the revenue-oriented counterpart in both energy (by 3.05%) and revenue (by 0.27%) extraction capabilities, therefore becoming our best automatic model. Our results for revenue-oriented operation of TRS were obtained using a limited available day-ahead market database, therefore, further investigations utilising energy demand, or day-ahead market forecast models should be investigated in future work. Additionally, the joint optimum operation of TRS with other renewable energy sources (e.g. wind, solar), could be investigated in future work, given the capabilities of TRS for working as batteries, storing energy from more uncertain energy sources.

6.1 Future Research Directions

The presented thesis has focused its efforts in introducing real-time control optimisation of TRS, either as augmentations of state-of-art methods or with novel DRL approaches, with the goal of maximising energy or revenue generation.

A future research direction could involve the evaluation of our DRL agent in performing real-time optimal operation of TRS under extreme scenarios (e.g. sea level rise, storm surges). Such research would complement the literature, that has focused its efforts on verifying coastal protection capabilities of TRS (using the 2D model developed in this work) for storm surge events [20; 21; 36; 38]. As motivation for this analysis, we have performed an initial evaluation of the sea level rise in Brazil and UK, at locations with high tidal range, where TRS could be built. For estimating flood risk regions, we utilised Climate Central's Screening Tool [139]. In all showcased analysis, we have selected "sea level rise" projection only, not considering added sea levels due to annual or moderate floods. Also, we have fixed the "pollution pathway" (amount of heat-trapping pollution added to the atmosphere) to the current estimated trajectory and set the "luck" parameter to medium, which means that the worse mid-range result from sea-level projection range (50^{th} percentile) are showcased in the analysis. Furthermore, regions isolated by higher land are excluded from flood risk scenarios. With these inputs, the sea level rise projection is estimated by the 2021 worldwide local sea level projections from the Intergovernmental Panel on Climate Change (IPCC) [140], embedded into the screening tool.

Figure 6.1. Map of the UK, with red regions highlighting the land projected to be below tideline by 2030. Created by using [139].

For the UK case study, we can see the estimated flood risk region due to sea level rise (for 2030 prediction) coloured in red in Fig. 6.1. By far, the most affected region is the Wash estuary, circled in Fig. 6.1. A zoom of the Wash estuary region, together with the estimated flooded area of $\approx 4000\ km^2$ (using Google Earth) for 2030 are shown in Figs. 6.2a and 6.2b. For comparison purposes, the greater London covers $\approx 1579\ km^2$. A projection of the flooded area for 2150 is also shown in Fig. 6.2c.

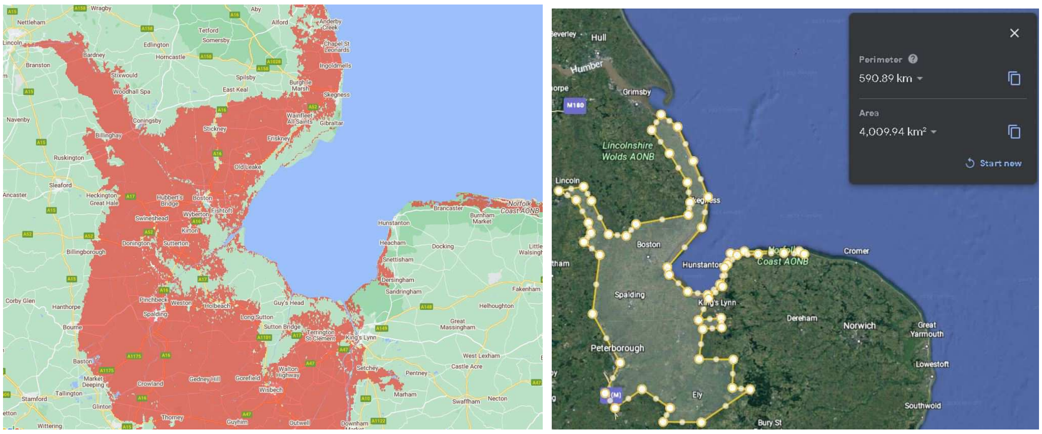

(a) 2030 flood risk area. **(b)** Area estimate for 2030 flood risk area.

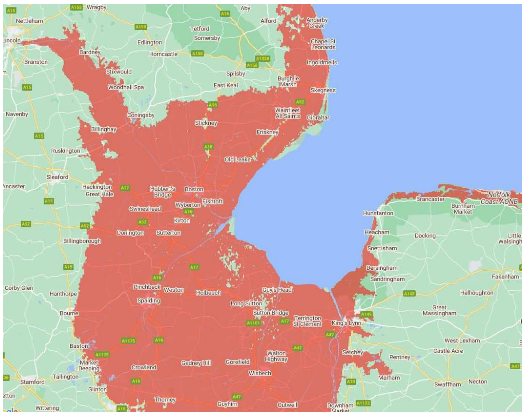

(c) 2150 flood risk area.

Figure 6.2. Maps of the Wash estuary. In a) and c) (created with [139]), red regions highlight the land projected to be below tideline by 2030 and 2150, respectively. b) showcases the approximate flooded area by 2030, through Google Earth estimate[141].

Nowadays, the primary flood defences along the Wash estuary's shoreline consist of relatively low earthen embankments [142]. However, these do not consider alternative scenarios with sea level rise, such as forecasted in Fig. 6.2. Therefore, and considering the human, environmental and economical scale of the predicted environmental catastrophe in the region, it is alarming that very few flood risk studies have been performed. Among these, a preliminary study by [142] was undertaken in 2012, analysing the necessity of developing a large scale coastal protection plan on the region.

Given the reasonably high mean tidal ranges in the Walsh estuary ($\approx 4.7\ m$) the study concluded that a TRS construction could be economically attractive, given the UK government participation (Fig. 6.3). The study assumed that the TRS would have 21 km in length, a 13050 m^2 total sluice area (225 sluices) and 97 turbines (where turbine diameter $D = 8\ m$, each) with a total installed capacity of 940 MW, generating 2945 GWh through a E_oG scheme. Adjusting for 2022 inflation and updated energy prices [138; 143], the yearly revenue would be worth $\approx$ £824.6 million per year (considering a non-flexible CH method). The total cost for constructing the TRS would be of $\approx$ £8.57 billion.

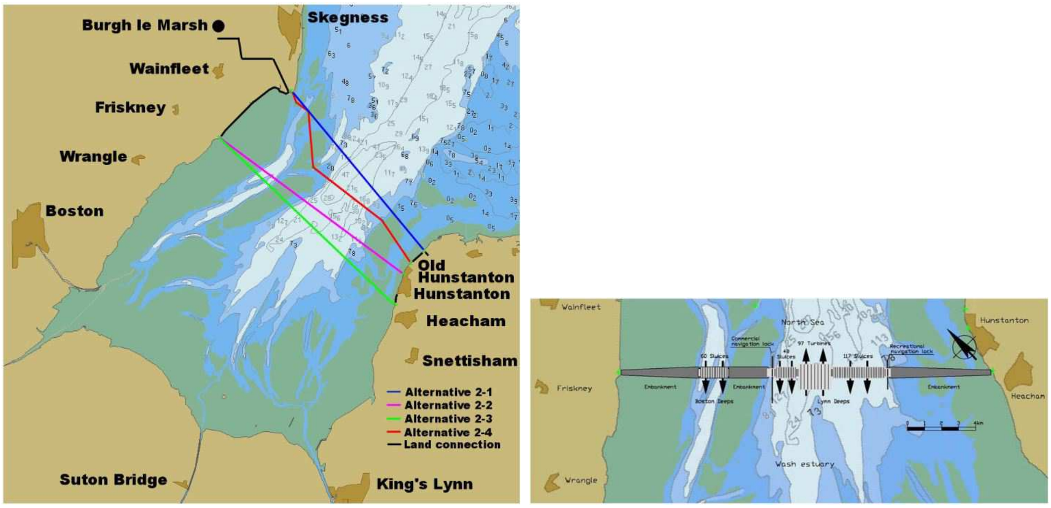

(a) Coastal barrier (or TRS) options. **(b)** Final TRS design.

Figure 6.3. Feasibility study of using a TRS for coastal protection in the the Wash estuary. From [142].

In Brazil a similar situation arises. As shown in Fig. 6.4a, starting in 2030 the State of Maranhão is under the risk of losing a significant portion of land due to sea level rise ($\approx 2.280\ km^2$ from Google Earth estimate in Fig. 6.4b). A projection of the flooded area for 2150 is also presented in Fig. 6.4c. For comparison purposes, the city of São Luís do Maranhão (also shown in Fig. 6.4) has $\approx 831.7\ km^2$.

Given the fact that the region showcases one of the highest tidal ranges in Brazilian coast (up to $\approx 6\ m$, as seen in Fig. 6.4d), one could argue that a TRS barrier could benefit the region, providing coastal protection and revenue. Strengthening this argument is the already constructed Bacanga dam, at the mouth of Bacanga estuary (Fig. 6.5). At this location, Figs. 6.5a and 6.5b showcase (in red) the scale of flood risk region due to sea level rise as we move from 2030 to 2150. Since the flood risk analysis does not consider the constructed Bacanga dam (circled in red in Fig. 6.5c), we can appreciate the environmental benefits that the region is already enjoying.

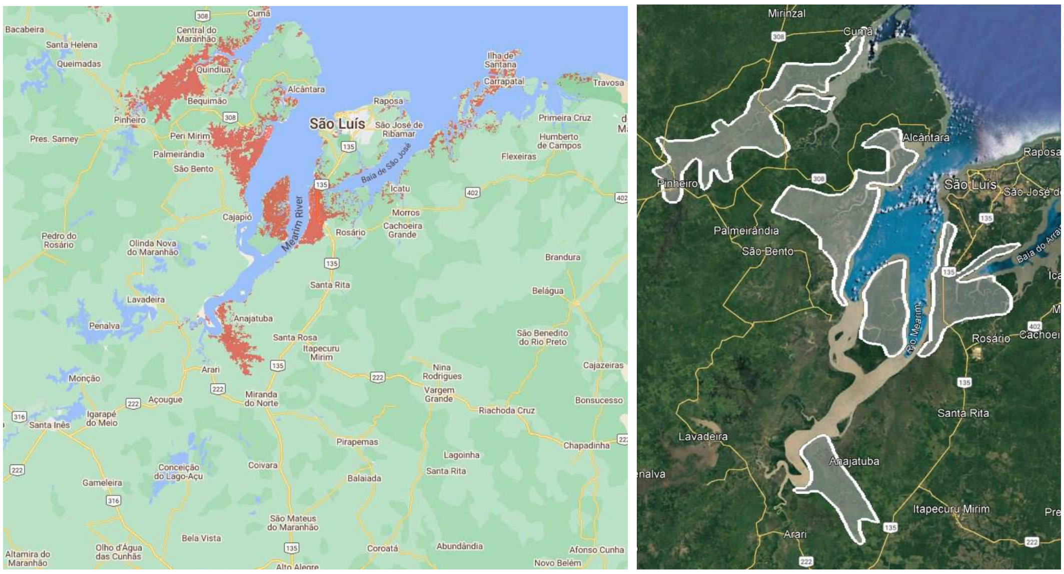

(a) 2030 flood risk area.

(b) Area estimate for 2030 flood risk area.

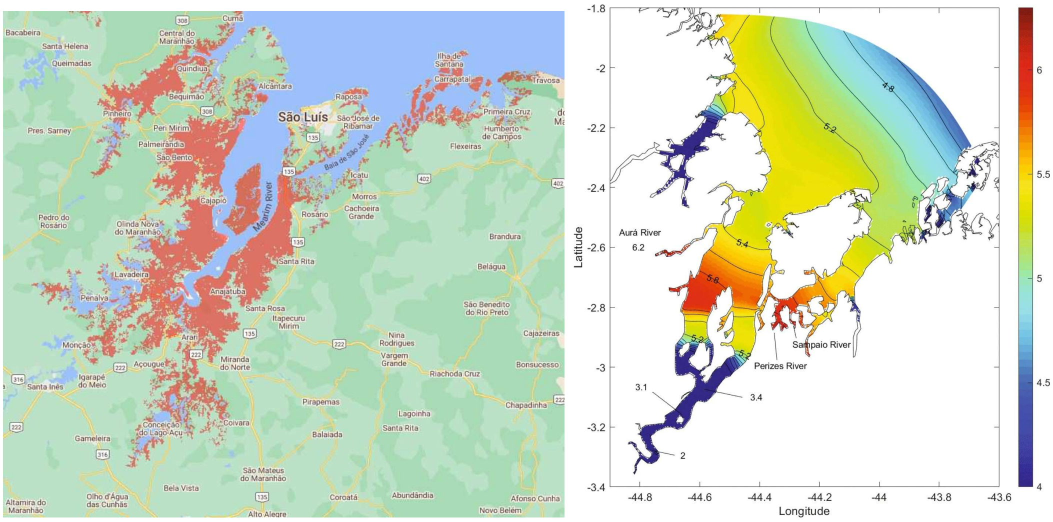

(c) Area estimate for 2150 flood risk area.

(d) Maximum tidal range colour map. Taken from [144].

Figure 6.4. Maps of the north region of State of Maranhão, Brazil. In a) and c) (created with [139]), red regions highlight the land projected to be below tideline by 2030 and 2150, respectively. b) showcases the approximate flooded area by 2030, Google Earth estimate [141]. d) showcases a colour map with the maximum tidal ranges observed in the region.

Considering the worldwide scale of the human, environmental and economical

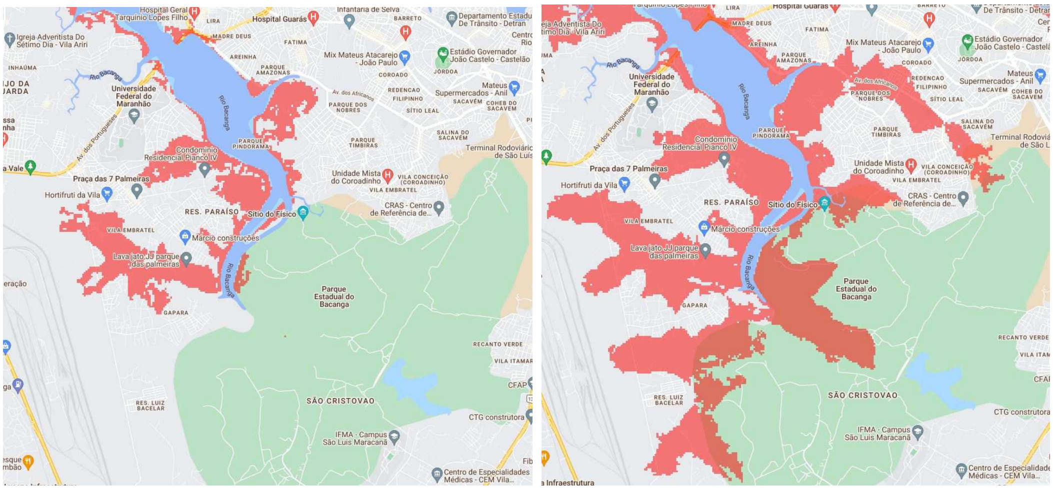

(a) 2030 flood risk area. **(b)** 2150 flood risk area.

(c) Baganca dam location (circled in red).

Figure 6.5. Maps showcasing the Bacanga estuary. a) and b) flood risk regions due to sea level rise (created with [139]), without considering the Bacanga dam, are coloured in red. In c) the Bacanga dam is highlighted in Google Earth [141].

damage that the sea level rise will have in the recent future, we believe TRS can become a viable option for coastal protection. Specially since that, given time, these structures can pay for their own construction costs, while generating clean, renewable energy and having a positive environmental impact by mitigating flood risks.

Bibliography

[1] P. IEA, Global Energy Review 2020 (April 2020, accessed January 10, 2021). URL `https://www.iea.org/reports/global-energy-review-2020`

[2] S. P. Neill, A. Angeloudis, P. E. Robins, I. Walkington, S. L. Ward, I. Masters, M. J. Lewis, M. Piano, A. Avdis, M. D. Piggott, et al., Tidal range energy resource and optimization–past perspectives and future challenges, Renewable Energy 127 (2018) 763–778.

[3] B. Alves, Net consumption of electricity worldwide in select years from 1980 to 2018, statista.
URL `https://www.statista.com/statistics/280704/world-power-consumption/#statisticContainer`

[4] S. Waters, G. Aggidis, Tidal range technologies and state of the art in review, Renewable and Sustainable Energy Reviews 59 (2016) 514–529.

[5] K. Hulsbergen, R. Steijn, R. Hassan, G. Klopman, D. Hurdle, Dynamic tidal power (dtp)-a new approach to exploit tides, in: 2nd International Conference on Ocean Energy, 2008, pp. 1–10.

[6] C. Baker, Tidal Power, Institution of Engineering and Technology (1711), (1991).

[7] D. Prandle., Simple theory for designing tidal power schemes, Advances in Water Resources 7 (1) (1984) 21–27.

[8] C. Hendry, The role of tidal lagoons, Final Report (2016) 67–75.

[9] S. Petley, G. Aggidis, Swansea bay tidal lagoon annual energy estimation, Ocean Engineering 111 (2016) 348–357.

[10] R. Burrows, I. Walkington, N. Yates, T. Hedges, J. Wolf, J. Holt, The tidal range energy potential of the west coast of the united kingdom, Applied Ocean Research 31 (4) (2009) 229–238.

[11] S. Evans, La Rance: learning from the world's oldest tidal project (October 4th, 2019, accessed January 10, 2021).
URL `https://www.power-technology.com/features/la-rance-learning-from-the-worlds-oldest-tidal-project/`

[12] B. Chang, K. Starcher, Evaluation of wind and solar energy investments in texas, Renewable Energy 132 (2019) 1348–1359.

[13] T. Guardian, Government rejects plan for £1.3bn tidal lagoon in Swansea (2018).
URL `https://www.theguardian.com/business/2018/jun/25/government-rejects-plan-for-tidal-lagoon-in-swansea`

[14] R. Horgan, Fresh tidal lagoon challenge launched in Wales (2021).
URL `https://www.newcivilengineer.com/latest/fresh-tidal-lagoon-challenge-launched-in-wales-26-03-2021/`

[15] R. Youle, Dozens of companies register interest in developing a tidal lagoon off Welsh coast (2021).
URL `https://www.walesonline.co.uk/news/wales-news/swansea-tidal-lagoon-wales-energy-20745132`

[16] A. Lisboa, T. Vieira, L. Guedes, D. Vieira, R. Saldanha, Optimal analytic dispatch for tidal energy generation, Renewable Energy 108 (2017) 371–379.

[17] A. Angeloudis, R. Ahmadian, R. A. Falconer, B. Bockelmann-Evans, Numerical model simulations for optimisation of tidal lagoon schemes, Applied Energy 165 (2016) 522–536.

[18] A. Cornett, J. Cousineau, I. Nistor, Assessment of hydrodynamic impacts from tidal power lagoons in the bay of fundy, International Journal of Marine Energy 1 (2013) 33–54.

[19] R. A. Falconer, J. Xia, B. Lin, R. Ahmadian, The severn barrage and other tidal energy options: Hydrodynamic and power output modeling, Science in China Series E: Technological Sciences 52 (11) (2009) 3413–3424.

[20] Q. Ma, T. M. Moreira, T. A. Adcock, Impact of the swansea bay lagoon on storm surges in the bristol channel, in: International Conference on Offshore Mechanics and Arctic Engineering, Vol. 58899, American Society of Mechanical Engineers, 2019, p. V010T09A018.

[21] Q. Ma, T. M. Moreira, T. A. Adcock, The impact of a tidal barrage on coastal flooding due to storm surge in the severn estuary, Journal of Ocean Engineering and Marine Energy 5 (3) (2019) 217–226.

[22] J. Zhou, S. Pan, R. A. Falconer, Optimization modelling of the impacts of a severn barrage for a two-way generation scheme using a continental shelf model, Renewable energy 72 (2014) 415–427.

[23] A. Angeloudis, S. C. Kramer, A. Avdis, M. D. Piggott, Optimising tidal range power plant operation, Applied energy 212 (2018) 680–690.

[24] F. Harcourt, A. Angeloudis, M. D. Piggott, Utilising the flexible generation potential of tidal range power plants to optimise economic value, Applied Energy 237 (2019) 873–884.

[25] J. Xue, R. Ahmadian, R. A. Falconer, Optimising the operation of tidal range schemes, Energies 12 (15) (2019) 2870.

[26] C. Gao, T. A. Adcock, On the tidal resonance of the bristol channel, International Journal of Offshore and Polar Engineering 27 (02) (2017) 177–183.

[27] G. A. Aggidis, O. Feather, Tidal range turbines and generation on the Solway Firth, Renewable Energy 43 (2012) 9–17.

[28] J. Xue, R. Ahmadian, O. Jones, R. A. Falconer, Design of tidal range energy generation schemes using a genetic algorithm model, Applied Energy 286 (2021) 116506.

[29] A. Ghaedi, H. Gorginpour, Generated power enhancement of the barrage type tidal power plants, Ocean Engineering 226 (2021) 108787.

[30] A. Juliani, V.-P. Berges, E. Vckay, Y. Gao, H. Henry, M. Mattar, D. Lange, Unity: A general platform for intelligent agents, arXiv preprint arXiv:1809.02627 (2018).

[31] R. A. Luettich, J. J. Westerink, N. W. Scheffner, et al., Adcirc: an advanced three-dimensional circulation model for shelves, coasts, and estuaries. report 1, theory and methodology of adcirc-2dd1 and adcirc-3dl.

[32] E. J. Kubatko, J. J. Westerink, C. Dawson, hp discontinuous galerkin methods for advection dominated problems in shallow water flow, Computer Methods in Applied Mechanics and Engineering 196 (1-3) (2006) 437–451.

[33] S. Bunya, E. J. Kubatko, J. J. Westerink, C. Dawson, A wetting and drying treatment for the runge–kutta discontinuous galerkin solution to the shallow water equations, Computer Methods in Applied Mechanics and Engineering 198 (17-20) (2009) 1548–1562.

[34] J. J. Westerink, R. A. Luettich, A. Militello, Leaky internal-barrier normal-flow boundaries in the adcirc coastal hydrodynamics code, Tech. rep., Engineer Research and Development Center Vicksburg ms Coastal and Hydraulics Lab (2001).

[35] C. Dawson, E. J. Kubatko, J. J. Westerink, C. Trahan, C. Mirabito, C. Michoski, N. Panda, Discontinuous galerkin methods for modeling hurricane storm surge, Advances in Water Resources 34 (9) (2011) 1165–1176.

[36] Q. Ma, T. A. Adcock, Modification of tidal resonance in the severn estuary by a barrage and lagoon, Journal of Ocean Engineering and Marine Energy 6 (2020) 171–181.

[37] A. M. Schnabl, T. M. Moreira, D. Wood, E. J. Kubatko, G. T. Houlsby, R. A. McAdam, T. A. Adcock, Implementation of tidal stream turbines and tidal barrage structures in dg-swem, in: International Conference on Offshore Mechanics and Arctic Engineering, Vol. 58899, American Society of Mechanical Engineers, 2019, p. V010T09A005.

[38] Q. Ma, Impact of barrages on extreme water levels in the bristol channel, Ph.D. thesis, University of Oxford (2020).

[39] S. Chandrasekhar, Newton's Principia for the common reader, Oxford University Press, 2003.

[40] A. T. Doodson, The harmonic development of the tide-generating potential, Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character 100 (704) (1921) 305–329.

[41] V. Deparis, H. Legros, J. Souchay, Investigations of tides from the antiquity to laplace, in: Tides in Astronomy and Astrophysics, Springer, 2013, pp. 31–82.

[42] R. G. Dean, R. A. Dalrymple, Water wave mechanics for engineers and scientists, Vol. 2, World Scientific Publishing Company, 1991.

[43] S. Fong, N. S. Heaps, Note on quarter wave tidal resonance in the bristol channel.

[44] C. Garrett, Tidal resonance in the bay of fundy and gulf of maine, Nature 238 (5365) (1972) 441–443.

[45] R. H. Karsten, J. McMillan, M. Lickley, R. Haynes, Assessment of tidal current energy in the minas passage, bay of fundy, Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy 222 (5) (2008) 493–507.

[46] R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction, MIT press, 2018.

[47] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing atari with deep reinforcement learning, arXiv preprint arXiv:1312.5602 (2013).

[48] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, arXiv preprint:1707.06347 (2017).

[49] P. Abbeel, Lecture 19 Off-Policy, Model-Free RL, UC Berkeley EECS (2019).
URL `https://www.youtube.com/watch?v=QASqaj_HUZw&t=2225s`

[50] J. Schulman, S. Levine, P. Abbeel, M. Jordan, P. Moritz, Trust region policy optimization, in: International Conference on Machine Learning, 2015, pp. 1889–1897.

[51] P. Abbeel, Lecture 4A: Policy Gradients, Deep RL Bootcamp Berkeley (2017).
URL `https://www.youtube.com/watch?v=S_gwYj1Q-44&t=2066s`

[52] P. Poupart, CS885 Lecture 7a: Policy Gradient, University of Waterloo (2018).
URL `https://www.youtube.com/watch?v=uoBii3rmyNk&t=1839s`

[53] E. Bøhn, E. M. Coates, S. Moe, T. A. Johansen, Deep reinforcement learning attitude control of fixed-wing uavs using proximal policy optimization, in: 2019 International Conference on Unmanned Aircraft Systems (ICUAS), IEEE, pp. 523–533.

[54] P. Abbeel, CS 287 Lecture 18 - RL I: Policy Gradients, UC Berkeley EECS (2019).
URL `https://www.youtube.com/watch?v=LsKXJl-_kzk`

[55] J. Schulman, Deep RL Bootcamp Lecture 5: Natural Policy Gradients, TRPO, PPO, Bootcamp Berkeley (August, 2017).
URL `https://www.youtube.com/watch?v=xvRrgxcpaHY`

[56] J. Schulman, P. Moritz, S. Levine, M. Jordan, P. Abbeel, High-dimensional continuous control using generalized advantage estimation, arXiv preprint arXiv:1506.02438 (2015).

[57] V. Pierre, Unity-technologies/ml-agents, `https://github.com/Unity-Technologies/ml-agents/blob/master/ml-agents/mlagents/trainers/ppo/optimizer_torch.py` (2020).

[58] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, K. Kavukcuoglu, Asynchronous methods for deep reinforcement learning, in: International conference on machine learning, 2016, pp. 1928–1937.

[59] M. G. Bellemare, Y. Naddaf, J. Veness, M. Bowling, The arcade learning environment: An evaluation platform for general agents, Journal of Artificial Intelligence Research 47 (2013) 253–279.

[60] M. Kempka, M. Wydmuch, G. Runc, J. Toczek, W. Jaśkowski, Vizdoom: A doom-based ai research platform for visual reinforcement learning, in: 2016 IEEE Conference on Computational Intelligence and Games (CIG), IEEE, 2016, pp. 1–8.

[61] E. Todorov, T. Erez, Y. Tassa, Mujoco: A physics engine for model-based control, in: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, 2012, pp. 5026–5033.

[62] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al., Soft actor-critic algorithms and applications, arXiv preprint arXiv:1812.05905.

[63] M. Rubak, Imitation learning with the unity machine learning agents toolkit, Ph.D. thesis, University of Applied Sciences (2021).

[64] B. Ospanov, Training intelligent tennis adversaries using self-play with ml agents.

[65] Unity-Technologies, Unity ml-agents learning environments, `https://github.com/Unity-Technologies/ml-agents/blob/main/docs/Learning-Environment-Examples.md` (2021).

[66] J. Xia, R. A. Falconer, B. Lin, Impact of different operating modes for a severn barrage on the tidal power and flood inundation in the severn estuary, uk, Applied Energy 87 (7) (2010) 2374–2391.

[67] C. H. Lebarbier, Power from tides—the rance tidal power station, Naval Engineers Journal 87 (2) (1975) 57–71.

[68] G. A. Aggidis, D. S. Benzon, Operational optimisation of a tidal barrage across the mersey estuary using 0-d modelling, Ocean Engineering 66 (2013) 69–81.

[69] R. Gibrat, Tidal energy.

[70] J. Bosc, Les groupes bulbes de la rance après trente ans d'exploitation-retour d'expérience, La Houille Blanche (3) (1997) 18–24.

[71] A. Angeloudis, M. Piggott, S. C. Kramer, A. Avdis, D. Coles, M. Christou, Comparison of 0-d, 1-d and 2-d model capabilities for tidal range energy resource assessments, EarthArXiv (2017).

[72] T. M. Moreira, J. G. d. Faria Jr, P. O. Vaz-de Melo, G. Medeiros-Ribeiro, Development and validation of an ai-driven model for the la rance tidal barrage: A generalisable case study, arXiv preprint arXiv:2202.05347.

[73] S.-H. Oh, K. S. Lee, W.-M. Jeong, Three-dimensional experiment and numerical simulation of the discharge performance of sluice passageway for tidal power plant, Renewable energy 92 (2016) 462–473.

[74] J. Xia, R. A. Falconer, B. Lin, G. Tan, Estimation of annual energy output from a tidal barrage using two different methods, Applied Energy 93 (2012) 327–336.

[75] R. H. Sellin, Severn Tidal Power Study, Preliminary Model Tests of Sluice Caissons, Department of Civil Engineering, University of Bristol, 1981.

[76] N. Yates, I. Walkington, R. Burrows, J. Wolf, The energy gains realisable through pumping for tidal range energy schemes, Renewable energy 58 (2013) 79–84.

[77] R. Ahmadian, J. Xue, R. A. Falconer, N. Hanousek, Optimisation of tidal range schemes, in: Proceedings of the 12th European Wave and Tidal Energy Conference, 2017, p. 1059.

[78] J. Xue, R. Ahmadian, O. Jones, Genetic algorithm in tidal range schemes' optimisation, Energy (2020) 117496.

[79] G. D. Egbert, R. D. Ray, Tidal prediction, Journal of Marine Research 75 (3) (2017) 189–237.

[80] E. Medina-Lopez, D. McMillan, J. Lazic, E. Hart, S. Zen, A. Angeloudis, E. Bannon, J. Browell, S. Dorling, R. Dorrell, et al., Satellite data for the offshore renewable energy sector: synergies and innovation opportunities, arXiv preprint arXiv:2103.00872 (2021).

[81] N. P. Spot, N2ex day ahead auction prices (2022).
URL `https://www.nordpoolgroup.com/en/search/?query=n2ex+day+ahead+auction+prices`

[82] C. Europea, Communication from the commission to the european parliament, the european council, the council, the european economic and social committee and the committee of the regions, Artificial Intelligence for Europe.

[83] S. Tobin, B. Jayabalasingham, S. Huggett, M. de Kleijn, A brief historical overview of artificial intelligence research, Information Services & Use 39 (4) (2019) 291–296.

[84] F. P. García Márquez, A. Peinado Gonzalo, A comprehensive review of artificial intelligence and wind energy, Archives of Computational Methods in Engineering (2021) 1–24.

[85] U. K. Das, K. S. Tey, M. Seyedmahmoudian, S. Mekhilef, M. Y. I. Idris, W. Van Deventer, B. Horan, A. Stojcevski, Forecasting of photovoltaic power generation and model optimization: A review, Renewable and Sustainable Energy Reviews 81 (2018) 912–928.

[86] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., Mastering the game of go with deep neural networks and tree search, nature 529 (7587) (2016) 484–489.

[87] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, et al., A general reinforcement learning algorithm that masters chess, shogi, and go through self-play, Science 362 (6419) (2018) 1140–1144.

[88] Y. Deng, F. Bao, Y. Kong, Z. Ren, Q. Dai, Deep direct reinforcement learning for financial signal representation and trading, IEEE transactions on neural networks and learning systems 28 (3) (2016) 653–664.

[89] A. E. Sallab, M. Abdou, E. Perot, S. Yogamani, Deep reinforcement learning framework for autonomous driving, Electronic Imaging 2017 (19) (2017) 70–76.

[90] A. Esteva, A. Robicquet, B. Ramsundar, V. Kuleshov, M. DePristo, K. Chou, C. Cui, G. Corrado, S. Thrun, J. Dean, A guide to deep learning in healthcare, Nature medicine 25 (1) (2019) 24–29.

[91] Z. Zhang, A. Chong, Y. Pan, C. Zhang, K. P. Lam, Whole building energy model for hvac optimal control: A practical framework based on deep reinforcement learning, Energy and Buildings 199 (2019) 472–490.

[92] Y. Li, Y. Wen, D. Tao, K. Guan, Transforming cooling optimization for green data center via deep reinforcement learning, IEEE transactions on cybernetics 50 (5) (2019) 2002–2013.

[93] R. Lu, S. H. Hong, X. Zhang, A dynamic pricing demand response algorithm for smart grid: reinforcement learning approach, Applied Energy 220 (2018) 220–230.

[94] A. Chaouachi, R. M. Kamel, R. Andoulsi, K. Nagasaka, Multiobjective intelligent energy management for a microgrid, IEEE transactions on Industrial Electronics 60 (4) (2012) 1688–1699.

[95] J. R. Vázquez-Canteli, Z. Nagy, Reinforcement learning for demand response: A review of algorithms and modeling techniques, Applied energy 235 (2019) 1072–1089.

[96] M. Kim, W. Choi, Y. Jeon, L. Liu, A hybrid neural network model for power demand forecasting, Energies 12 (5) (2019) 931.

[97] A. Saenz-Aguirre, E. Zulueta, U. Fernandez-Gamiz, J. Lozano, J. M. Lopez-Guede, Artificial neural network based reinforcement learning for wind turbine yaw control, Energies 12 (3) (2019) 436.

[98] J. Zhang, J. Yan, D. Infield, Y. Liu, F.-s. Lien, Short-term forecasting and uncertainty analysis of wind turbine power based on long short-term memory network and gaussian mixture model, Applied Energy 241 (2019) 229–244.

[99] S. Shresthamali, M. Kondo, H. Nakamura, Adaptive power management in solar energy harvesting sensor node using reinforcement learning, ACM Transactions on Embedded Computing Systems (TECS) 16 (5s) (2017) 1–21.

[100] B. C. Phan, Y.-C. Lai, C. E. Lin, A deep reinforcement learning-based mppt control for pv systems under partial shading condition, Sensors 20 (11) (2020) 3039.

[101] A. Gensler, J. Henze, B. Sick, N. Raabe, Deep learning for solar power forecasting—an approach using autoencoder and lstm neural networks, in: 2016 IEEE international conference on systems, man, and cybernetics (SMC), IEEE, 2016, pp. 002858–002865.

[102] F. Wang, Z. Xuan, Z. Zhen, K. Li, T. Wang, M. Shi, A day-ahead pv power forecasting method based on lstm-rnn model and time correlation modification under partial daily pattern prediction framework, Energy Conversion and Management 212 (2020) 112766.

[103] A. Nambiar, E. Anderlini, G. S. Payne, D. Forehand, A. Kiprakis, R. Wallace, Reinforcement learning based maximum power point tracking control of tidal turbines, in: Proceedings of the 12th European Wave and Tidal Energy Conference, Cork, Ireland, Vol. 27, 2017.

[104] E. Anderlini, S. Husain, G. G. Parker, M. Abusara, G. Thomas, Towards real-time reinforcement learning control of a wave energy converter, Journal of Marine Science and Engineering 8 (11) (2020) 845.

[105] T. M. Moreira, J. G. de Faria Jr, P. O. Vaz-de Melo, L. Chaimowicz, G. Medeiros-Ribeiro, Prediction-free, real-time flexible control of tidal lagoons through proximal policy optimisation: A case study for the swansea lagoon, Ocean Engineering 247 (2022) 110657.

[106] Y. Yu, X. Si, C. Hu, J. Zhang, A review of recurrent neural networks: Lstm cells and network architectures, Neural computation 31 (7) (2019) 1235–1270.

[107] J. Lago, G. Marcjasz, B. De Schutter, R. Weron, Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark, Applied Energy 293 (2021) 116983.

[108] H. Swane, Tidal power plant in saemangeum, MSc Thesis, Delft University of Technology, 2007.

[109] G. Rolandez, A. Abgottspon, T. Staubli, Discharge measurements at la rance tidal power plant using current meters method, IGHEM2014, Itajuba.

[110] J. Rodda, N. ALAS, Water for the future, IAHS publication, and references therein (164).

[111] F. Rodriguez, The analogy between fluid flow and electric circuitry, Chemical Engineering Education 13 (2) (1979) 96–98.

[112] RANCE-ENVIRONNEMENT, Estuaire de la rance – l'envers du décor, Le Bulletin de Rance-Environment, 2016/2, August 2016.
URL `https://www.canauxdebretagne.org/docs/files/rance-env.pdf`

[113] G. Caude, P. Clemént, D. Pillet, E. Vindimian, Gestion sédimentaire de l'estuaire de la rance, Rapport CGEDD n° 010860-01, CGE n° 2016-30, May 2017.
URL `https://cgedd.documentation.developpement-durable.gouv.fr/documents/Affaires-0009658/010860-01_rapport_publie.pdf`

[114] F. O. Rourke, F. Boyle, A. Reynolds, Tidal energy update 2009, Applied energy 87 (2) (2010) 398–409.

[115] V. C. Lakhan, Advances in coastal modeling, Elsevier, 2003.

[116] A. D'Alpaos, S. Lanzoni, M. Marani, A. Rinaldo, On the tidal prism–channel area relations, Journal of Geophysical Research: Earth Surface 115 (F1).

[117] M. Stewart, Surface Production Operations: Volume IV: Pumps and Compressors, Gulf Professional Publishing, 2018.

[118] N. Arshenevskii, S. Levina, E. Natarius, Characteristics of bulb turbines during their operation in reversible regimes, Hydrotechnical Construction 13 (10) (1979) 1004–1010.

[119] H. Ramos, A. Borga, Pumps as turbines: an unconventional solution to energy production, Urban Water 1 (3) (1999) 261–263.

[120] R. S. Stelzer, R. Walters, Estimating reversible pump-turbine characteristics, no. 39, Department of the Interior, Bureau of Reclamation, Engineering and Research . . . , 1977.

[121] M. Zu-yan, Mechanical Design and Manufacture of Hydraulic Machinery, Routledge, 2018.

[122] M. Simão, H. M. Ramos, Micro axial turbine hill charts: affinity laws, experiments and cfd simulations for different diameters, Energies 12 (15) (2019) 2908.

[123] Q. Wang, J. Gao, W. Lin, Y. Yuan, Learning from synthetic data for crowd counting in the wild, in: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, 2019, pp. 8198–8207.

[124] E. Wolanski, M. Elliott, Estuarine ecohydrology: an introduction. Section 2.1 The tides at sea, Elsevier, 2015.

[125] BODC, Download UK Tide Gauge Network Data from BODC (accessed January 10, 2021).
URL `https://www.wpc.ncep.noaa.gov/html/beaufort.shtml`

[126] D. J. Wales, J. P. Doye, Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms, The Journal of Physical Chemistry A 101 (28) (1997) 5111–5116.

[127] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al., Scipy 1.0: fundamental algorithms for scientific computing in python, Nature Methods 17 (3) (2020) 261–272.

[128] M. J. Powell, Advances in optimization and numerical analysis, in: Proceeding of the 6th Workshop on Optimization and Numerical Analysis, 1994, pp. 5–67.

[129] S. v0.14.0 Reference Guide, scipy.optimize.basinhopping (2015, accessed January 10, 2021).
URL `https://het.as.utexas.edu/HET/Software/Scipy/generated/scipy.optimize.basinhopping.html#r95`

[130] E. Sonnic, La rance, 50 ans de turbinage. et après? le statu quo est-il la seule option pertinente?, L'information géographique 81 (4) (2017) 103–128.

[131] M. Balls, The optimal selection of turbine-generators for tidal power projects and the optimization of their operation, University of Salford (United Kingdom), 1988.

[132] P. Lutus, Jtides tidal prediction software, Version: 5.3 build 368.
URL `https://www.arachnoid.com`

[133] R. G. Crockett, G. K. Gillmore, P. S. Phillips, D. D. Gilbertson, Tidal synchronicity of the 26 december 2004 sumatran earthquake and its aftershocks, Geophysical research letters 33 (19).

[134] N. P. Nezlin, K. Kamer, J. Hyde, E. D. Stein, Dissolved oxygen dynamics in a eutrophic estuary, upper newport bay, california, Estuarine, Coastal and Shelf Science 82 (1) (2009) 139–151.

[135] N. J. Cooper, J. S. Pethick, Sediment budget approach to addressing coastal erosion problems in st. ouen's bay, jersey, channel islands, Journal of Coastal Research 21 (1) (2005) 112–122.

[136] J. C. McNatt, A. Porter, K. Ruehl, Comparison of numerical methods for modeling the wave field effects generated by individual wave energy converters and multiple converter wave farms, Journal of Marine Science and Engineering 8 (3) (2020) 168.

[137] L. Zhang, K. Han, How to analyze change from baseline: Absolute or percentage change, D-level Essay in Statistics. Dalarna, Sweden: Dalarna University.

[138] I. Webster, Cpi inflation calculator, `https://www.in2013dollars.com/uk/inflation/2012?amount=279.78` (March/20/2022).

[139] Climate central's coastal risk screening tool - land projected to be below tideline. URL `https://coastal.climatecentral.org/map`

[140] S. Legg, Ipcc, 2021: Climate change 2021-the physical science basis, Interaction 49 (4) (2021) 44–45.

[141] Google-Earth, Wash estuary flood region estimate for 2030, `https://earth.google.com` (2022).

[142] B. Hofschreuder, Flood protection and marine power in the wash estuary, united kingdom: Technical and economical feasibility study.

[143] E. Yurday, Average cost of electricity per kwh in the uk 2022 - nimblefins, `https://www.nimblefins.co.uk/average-cost-electricity-kwh-uk` (March/20/2022).

[144] A. Czizeweski, F. M. Pimenta, O. R. Saavedra, Numerical modeling of maranhão gulf tidal circulation and power density distribution, Ocean Dynamics 70 (5) (2020) 667–682.

Appendix A

Appendix

This appendix presents selected BODC data used as test data for SBL case study (no pumping) in Section 5.1 (Tables A.1 and A.2). In detail, 26 "months" of 30 days with truthful data were selected in BODC measured data between the years 1993 to 2014 as indicated in Table A.1. Then, monthly energies were calculated and compared for these 26 months using state-of-art literature baselines (CH, EHT, EHN), improved baselines with variant (independent) sluice operation for HC methods introduced in this thesis (CHV, $EHTV$, $EHNV$) and the novel DRL automatic approach in Table A.2. For sake of completeness, the hardware specifications used in the major part of this thesis is presented in Table A.3. It is worth noticing that the barrage implementation with 2D model required supercomputer facilities, either from the Centro Nacional de Processamento de Alto Desempenho (CENAPAD – UFMG) or from the Advanced Research Computing (ArC – Oxford University).

Visual normality assumption (QQ) plots (verified with Shapiro Wilk test), required for performing confidence interval t-test, are also presented in Figs. A.1, A.2, A.3 and A.4. Furthermore, all day-ahead market data utilised for predicting the revenue extraction capabilities of our methods (and training the revenue-oriented DRL agent) are shown in Fig. A.5.

Table A.1. Numbered test data. For each data, year, month, day and hour are indicated.

Year	Start (M/D hr:min)	End (M/D hr:min)	Data #
1993	03/03 00:15	04/02 00:15	1
1993	04/02 14:15	05/02 14:15	2
1997	09/02 00:15	10/02 06:15	3
1998	06/03 13:00	07/03 19:00	4
1998	11/03 05:00	12/03 05:00	5
1999	07/05 04:00	08/04 10:30	6
1999	10/04 14:00	11/03 14:00	7
1999 → 2000	12/04 16:00	01/03 22:15	8
2000	06/04 13:30	07/04 13:30	9
2000	09/04 03:30	10/04 09:45	10
2003	05/08 03:30	06/07 03:30	11
2010	06/12 18:15	07/12 18:15	12
2012	04/13 17:00	05/13 17:00	13
2012	05/14 00:30	06/13 06:45	14
2012	07/14 02:00	08/13 02:00	15
2012	08/13 15:00	09/12 15:00	16
2012 → 2013	12/13 18:00	01/12 18:00	17
2013	01/13 01:00	02/12 07:15	18
2013	02/12 13:45	03/14 19:45	19
2013	03/15 02:00	04/14 08:15	20
2013	05/15 02:45	06/14 09:15	21
2013	06/14 15:15	07/14 21:45	22
2013	08/14 16:30	09/13 16:30	23
2013	09/14 06:30	10/14 06:30	24
2013	10/14 14:15	11/13 14:15	25
2013 → 2014	12/14 16:15	01/13 16:15	26

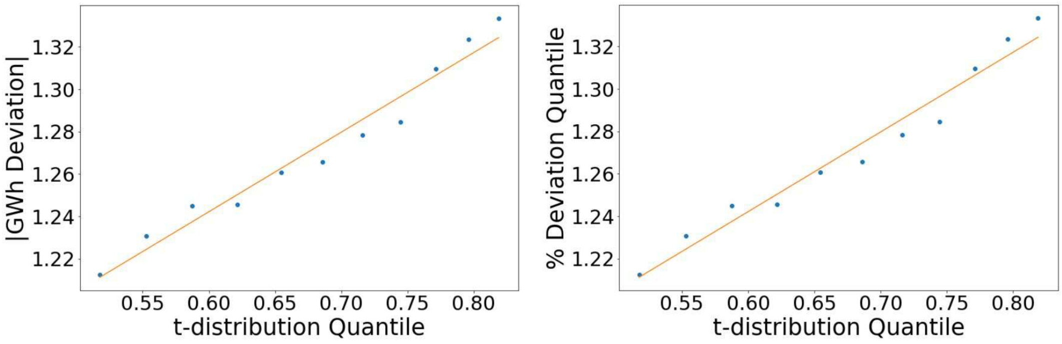

(a) QQ plot for Fractional energy gain. **(b)** QQ plot for Fractional revenue gain.

Figure A.1. Visual verification of normality assumption with QQ plots.

Table A.2. Monthly energy comparison (GWh) between baselines and our trained agent, based on numbered test data.

Data	CH	CHV	EHT	$EHTV$	EHN	$EHNV$	Agent
1	32.04	32.39	43.39	44.40	43.30	44.43	44.08
2	34.48	35.01	43.31	44.21	43.50	44.42	43.79
3	35.84	36.25	44.33	45.27	44.00	45.02	44.77
4	30.40	30.57	41.68	42.52	41.67	42.66	41.76
5	37.86	38.78	43.86	44.74	44.05	45.05	44.46
6	31.06	31.11	42.83	43.69	42.73	43.97	43.02
7	31.77	31.99	42.89	43.82	42.97	43.61	42.94
8	30.22	30.37	40.69	41.53	40.80	41.19	40.37
9	38.07	38.77	41.69	42.51	41.88	42.78	42.48
10	32.11	32.18	42.37	43.23	42.33	43.27	42.79
11	28.37	28.77	38.35	39.09	38.34	38.90	38.89
12	31.89	32.47	38.79	39.52	38.75	39.74	39.45
13	28.96	29.13	42.15	42.98	42.21	42.97	42.38
14	27.79	27.87	41.35	42.21	41.46	42.21	41.50
15	29.46	29.63	40.81	41.73	40.87	41.64	41.11
16	31.02	31.23	41.52	42.40	41.76	42.47	41.86
17	35.90	36.88	41.81	42.64	41.89	42.89	42.02
18	37.86	38.14	44.34	45.27	44.38	45.53	44.39
19	38.06	38.71	44.89	45.74	44.85	45.75	45.18
20	36.75	37.48	43.83	44.78	43.97	44.73	44.01
21	29.31	29.32	41.06	41.88	41.04	42.05	41.12
22	30.74	30.14	41.98	42.86	42.19	42.91	42.14
23	31.54	31.81	43.58	44.48	43.66	44.26	44.10
24	32.71	33.62	43.13	44.10	43.30	44.18	43.56
25	31.82	32.19	42.37	43.26	42.44	43.25	42.81
26	31.34	31.62	41.68	42.55	41.86	42.66	41.59

Table A.3. Hardware specifications for this work.

Hardware Components	
Operating System	Microsoft Windows 10 Pro
CPU	AMD Ryzen 7 3800X
GPU	GeForce RTX 2070 Super
RAM	16GB

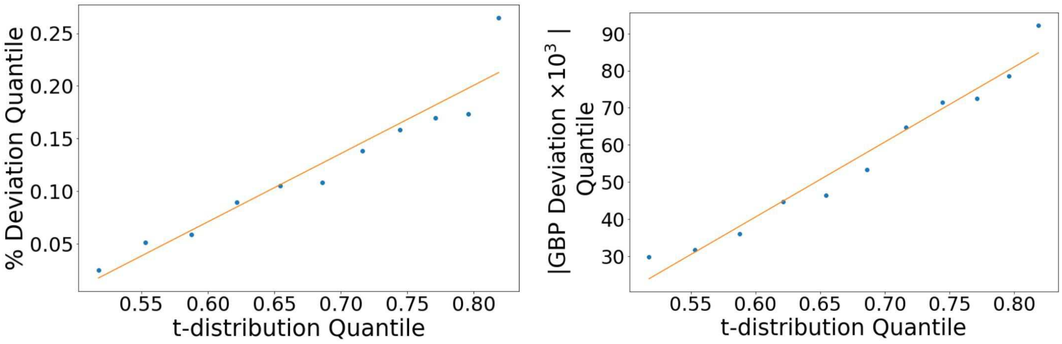

(a) QQ plot for upper bound and real-time estimate % energy deviation.

(b) QQ plot for upper bound and real-time estimate absolute revenue deviation.

Figure A.2. Visual verification of normality assumption for the deviation between upper and real-time estimates of the $EHTV - P$ method, considering a real-time energy-oriented operation.

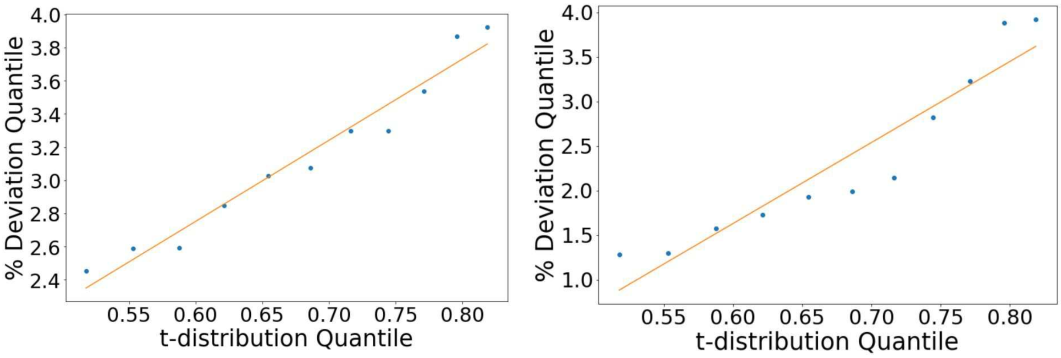

(a) QQ plot for fractional energy gain.

(b) QQ plot for fractional revenue gain.

Figure A.3. Visual verification of normality assumption for the deviation between our energy-oriented agent and the upper bound (energy and revenue-oriented) from the literature.

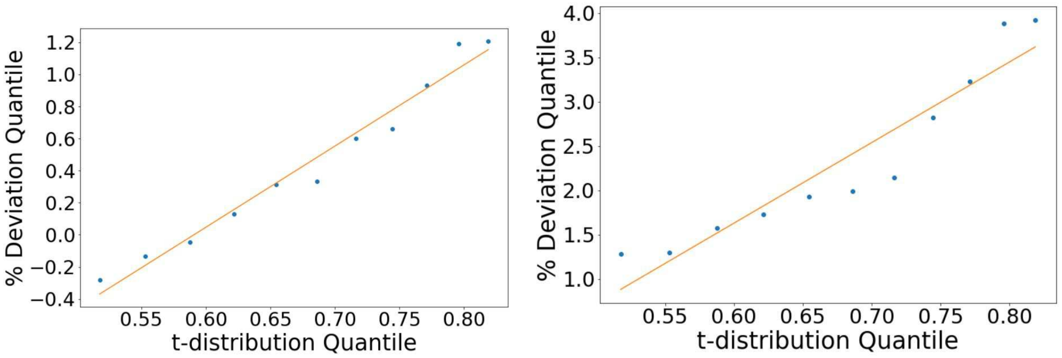

(a) QQ plot for energy-oriented DRL and real-time $EHTV - P$ % energy deviation.

(b) QQ plot for energy-oriented DRL and real-time $EHTV - P$ % revenue deviation.

Figure A.4. Visual verification of normality assumption for the deviation between our energy-oriented agent and the real-time $EHTV - P$ method.

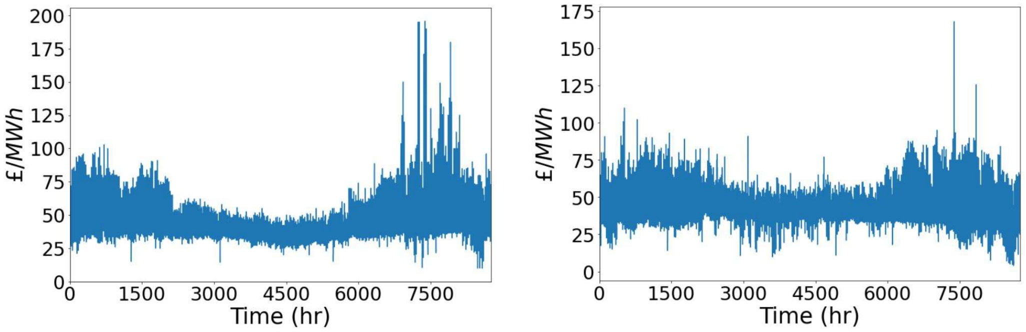

(a) 2014. **(b)** 2015.

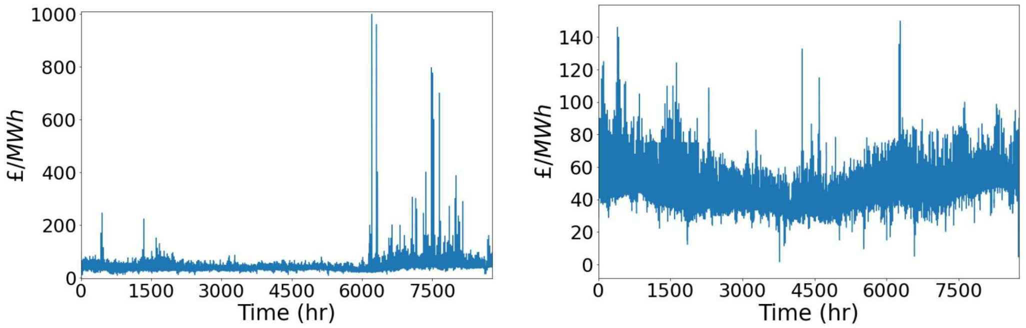

(c) 2016. **(d)** 2017.

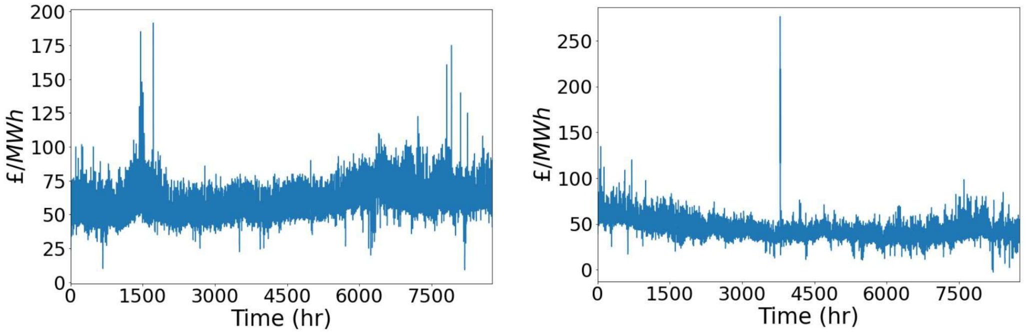

(e) 2018. **(f)** 2019.

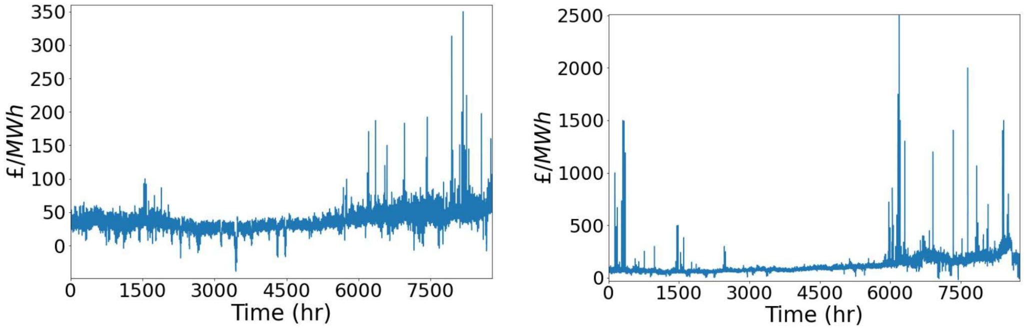

(g) 2020. **(h)** 2021.

Figure A.5. Day-ahead market data (in $\pounds/MWh$) from 2014 to 2021. Taken from [81].