Development of tidal energy generation modelling using Morecambe Bay as a case study under different environmental, storage and demand scenarios.

By

Simon Baker

Supervisors:

Prof George Aggidis

Dr David Howard

In collaboration with

Northern Tidal Power Gateways

Signed Declaration

Lancaster University

Faculty of Science and Technology

Engineering Department

Signed Declaration on the submission of a dissertation

"I declare that this dissertation is my own work and has not been submitted in substantially the same form towards the award of a degree or other qualification. Acknowledgement is made in the text of assistance received and all major sources of information have been appropriately referenced. I confirm that I have read and understood the publication Guidance on Writing Technical Reports published by the Department."

S. P. Baker

Signed …………………………………………………………

9th October 2020
Date ……………………………………………………………

Abstract

Electricity generation is a major source of greenhouse gas emissions. Renewable energy mitigates those emissions but poses different problems for the use of the power. This project examines the potential of using a barrage across Morecambe Bay to capture tidal range energy. The tidal system for a specific location is complex and requires multiple levels of robust detailed modelling.

The optimum barrage operational parameter values (e.g. generating starting head and turbine speed) vary with the height of the tide. They are also influenced by the design and should be adjusted whenever the design changes. Using 0-D modelling, the energy from real tides can be modelled effectively by applying a set of linear functions to relate the operational parameters to the tidal range. The function coefficients were determined via an optimisation process, specific to each new design, to maximise the generated energy.

The timing of electricity generation is dictated by the tide with little scope for adjustment. Electricity cannot be stored and is effectively lost if not needed. To store all the surplus energy from the barrage would require a dedicated facility on a similar scale to the barrage itself and is deemed impractical. With electricity usage patterns expected to change as the UK transitions to a low carbon economy, the solution is perhaps one of energy balancing in combination with storage and conversion to different energy carriers.

At a global level the barrage addresses sea-level rise by generating large amounts of green energy. At a local level the barrage can mitigate the threat of increased flooding and can preserve the tidal range to help protect the inter-tidal zone.

It is difficult to value all the aspects of a barrage with confidence and to reliably balance the costs of building a barrage against not building it. The environment plays a crucial part in this assessment and must be incorporated into the performance estimates of the barrage at all stages of the design.

Table of Contents

List of figures

List of tables

1 Introduction

The transition from fossil fuels to green energy is driven by the need to reduce the county's CO_2 emissions and mitigate the effects of global warming. Tidal range generation is capable of supplying large amounts of clean and reliable electrical energy making it a candidate technology to help achieve sustainability. In addition to generating power, barrages also offer improved or novel transport links, the creation of new jobs and can provide recreational facilities. Increasingly important, with the prospect of significant sea-level rise, a barrage can protect property from flooding and provide a level of stability to assist with ecological and environmental conservation. A tidal range barrage also has a major advantage over alternative renewable energy solutions; it uses tried and tested technology.

The UK is at a critical point where long-term decisions are being made about our future energy supply, and there is a danger of excluding viable solutions if they are based on incomplete data and out-dated perceptions. This study is one piece in the overall design process and its intention is to provide reliable information to help decide whether a barrage here is a potential option to address the country's energy needs and if the design should proceed to the next phase.

This project was funded through the Centre for Global Eco-Innovation (CGE) supported by an industrial partner Northern Tidal Power Gateways (NTPG). The company's aim was to design and build barrages across both Morecambe Bay and the Duddon Estuary. NTPG was actively in pursuit of financial support from the Government to fund the next phase.

A Preliminary Investigation Memorandum (PIM) [1] and Strategic Outline Business Case (SOBC) were commissioned and completed, in accordance with the UK Government's appraisal process. The PIM and SOBC specify details of a candidate barrage design along with the expected energy returns, related construction and financial information. Although necessary for Government process, the designs are preliminary. The costings are useful but not reliable since the location and operational requirements had not been established. The PIM and SOBC were undertaken by Mott MacDonald (MM) using their own tidal range energy generation models; details of which were not published within the reports nor released to NTPG.
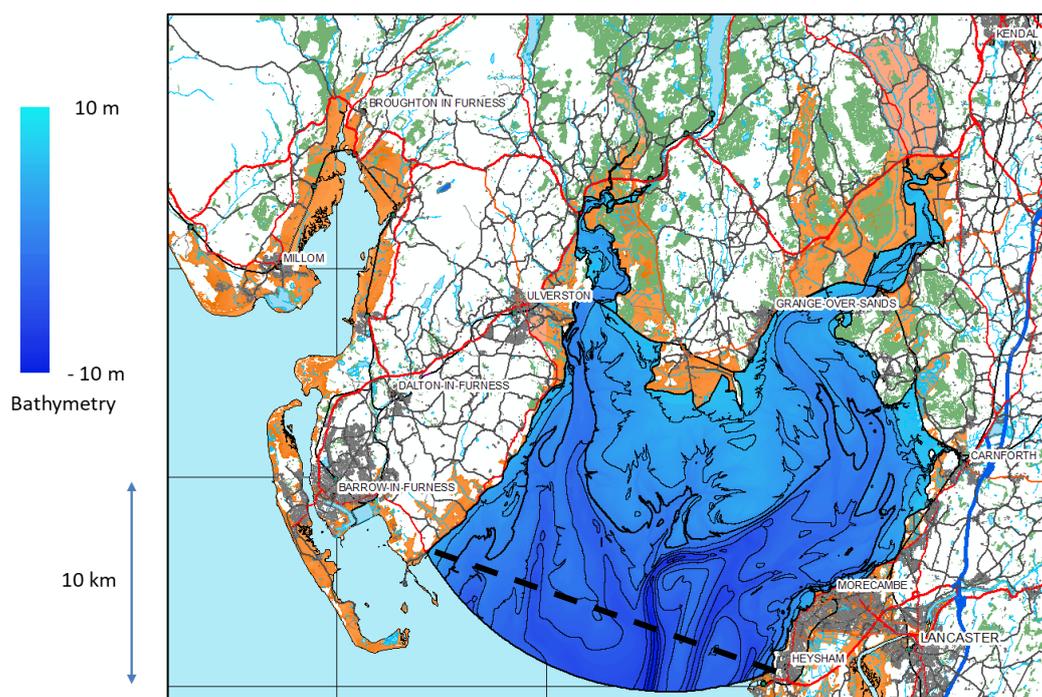
Figure 1.1 Morecambe Bay barrage area.

Morecambe Bay is situated in the NW of England with the South Lakelands of Cumbria to the north and Lancashire to the south. Barrow-in-Furness is at the north-western end of the barrage, and Heysham, Morecambe and Lancaster at the south-eastern end. The nuclear power plant (NPP) and ferry port at Heysham lie just outside the barrage wall, and the location with the NPP provides a good connection point to the National Grid. The majority of the land between Barrow and Carnforth to the east is characterised by hills punctuated by the Levens and Kent river estuaries with low-lying fluvial valleys extending inland. From Carnforth south to Heysham the land is lower with the tidal reaches of the river Lune running roughly parallel to the Bay creating a strip of land approximately 2 km wide.

Figure 1.1 shows a map of Morecambe Bay and the surrounding area. The graduated darker blue area and the contour lines (at 2m intervals) represent the bathymetry. It also denotes the extent of the reservoir, limited landward by the topography and flood defences, and seaward by the location of the barrage wall. Two different barrage configurations were analysed (both sharing the same end points). The straight version (dashed line) was taken approximately from a map in the PIM to enable a comparison to be made between the PIM estimated AEP values and the values modelled here. The curved option was added to explore the possible uplift in AEP available by extending the barrage reservoir volume with relatively little increase in the length of the barrage wall and maintaining the same end points and connections to the road network.

Designing a barrage is an iterative process, and the sophistication of the modelling and the quality of the input data evolve and develop over time. For example, the cost of conducting a detailed engineering geology survey is high and is not undertaken until a preferred location of the barrage is identified from more basic modelling. Progression to the next stage only takes place after the necessary agreement to proceed is reached. Previous proposals have established estimates of the nominal Annual Energy Production (AEP) value and how the timing of the power generation corresponds to other potential tidal range schemes around the UK [2]. MM selected a particular design and calculated a new set of detailed predictions of the AEP presented in the PIM and SOBC. Their reports did not include details of the model, the exact location or character of the barrage, description of the inter-tidal bathymetry or projections on the operation of the barrage with changing sea-levels. Thus the modelling undertaken in this study repeats and extends the design process but without testing agreement.

The aims of the study were to:

1. advance the modelling capability of tidal range energy generation;
2. determine the design and operational implications imposed by sea-level rise and environmental protection requirements;
3. investigate the options available for electrical demand power balancing.

1.1   Modelling methods
There are several methods typically employed for tidal barrage modelling, classified by the number of spatial dimensions considered, ranging from 0-D through to 3-D. As the number of dimensions increases the complexity of the modelling and the run-time increases substantially. Neil *et al.* [3] and Angeloudis *et al.* [4], provide a good overview of the different methods.

In summary, 0-D models are computationally efficient but ignore any hydrodynamic effects, i.e. assume there is no impact of the barrage on the tide, and changes to the reservoir volume are distributed across the wetted area instantly. Commonly, 0-D overestimates the energy available, and in the case of large schemes, this can be quite significant [5]. 1-D modelling divides the estuary into a

series of connected 1-D cross-sectional slices and allows for some degree of regional hydrodynamics, but is not considered particularly accurate for larger schemes [4].

2-D modelling is an extension of 1-D with the estuary represented by a mesh of depth-averaged cells. This approach is preferred for more reliable simulations of the impact on the tide and estimates of available energy [5, 6]. 3-D modelling introduces a further extension, and is more suited to understanding complex flow around structures and the effects of high velocity regimes up and downstream of the turbines and sluice gates [4]. In the case of Morecambe Bay and the Duddon, 3-D modelling will be essential at a later stage as it will allow the complex movement of sediment and dynamics of the bathymetry to be included.

Predominantly, 0-D modelling is the method of choice for establishing the main design and operational parameters [7-9]. The small computational effort allows many thousands of runs to be executed, enabling the importance or sensitivity of different parameter variables in the model to be assessed. 2-D modelling, in contrast, is computationally expensive and is limited to a few runs with few variables. It typically follows on from 0-D modelling and is used to generate more realistic energy values for the 0-D derived designs, and to model the regional impact on the tide of these designs [5].

The focus of this study is on the initial design phase; the sensitivity of the AEP on the design and operational parameters, and the investigation of how environmental and ecological considerations might dictate the design. For this reason, all the modelling has been carried out in 0-D. It does not attempt to specify a preferred design by incorporating component costs into the energy modelling. As yet the design requirements are still open to discussion, and the optimum means of operating the barrage remain to be properly understood. The focus here remains predominantly on process. Although this is targeted at Morecambe Bay, the process applied here is applicable elsewhere.

The following objectives were set to address the main aims:

1. develop a fit-for-purpose 0-D model (written using MATLAB) capable of design and operational optimisation;
2. use the model to explore the effects of sea-level rise on the design, operation and energy generating capacity of the barrage;
3. specify a proxy environmental protection requirement in the form of a barrage operating condition to explore the effects on the design, operation and energy generating capacity of the barrage;
4. model the power and energy characteristics of the barrage and investigate local energy storage options and barrage operational scenarios to enable delivery of the generated energy to meet demand.

The term "fit-for-purpose" indicates that the modelling had to be of sufficient sensitivity to be able to make reliable decisions about proceeding to the next phase of development. It is only when the model is fit-for-purpose that meaningful carbon reduction estimates can be made; in detail this meant being robust, capable of giving estimates of power output and costs with known levels of confidence and enabling design parameter sensitivity analysis. In practice this has necessitated the optimisation of the operational parameters every time the design parameters change. Without this, the difference in the predicted net energy between designs can be smaller than the differences between different levels of optimisation. Using generalised operational parameters for all designs results in poor barrage characterisation, and as the parameters become increasingly suboptimal, this quickly leads to a breakdown in the operation altogether. For example, if the generating starting head is not reached the generating phase is skipped altogether, which has a knock-on effect and severely reduces the calculated net energy value.

1.2   Literature review

All of the issues related to the aims and objectives have been addressed to some degree in the literature to date, however, there is no single paper that addresses all of the issues together. Earlier attempts at optimisation addressed a limited number of variables independently. Burrows *et al*. [2] modelled how the generated energy varied with the number of installed turbines. Aggidis and Feather [8] derived the optimum generating starting head using a scanning approach for a barrage across the Solway Firth in ebb only mode with fixed installed capacity. Aggidis and Benzon [7] extended this by deriving a function of optimum starting head against tidal range for a Mersey barrage, similarly for ebb only and fixed installed capacity.

More recent papers have employed sophisticated optimisation schemes involving more parameters. Kontoleontos and Weissenberger [10] maximised energy generation for a fixed lagoon design (number of turbines and sluice gates) using an evolutionary algorithm: this was for a year-long tidal series, and determined the starting and stopping heads of each operating mode and the turbine speed and number of units for each time step. Angeloudis *et al.* [9] used an iterative optimiser to determine the start and stop times for both ebb and flood generation (including pumping) for a lagoon with a fixed number of turbines and multiple tide cycles. This was developed further [11] by maximising value instead of energy and employing a global basin-hopping algorithm. Xue *et al*. [12] employed a grid search scanning approach for a range of transition head values (including pumping) on a rolling pair of half tide cycles. This was later adapted [13] to use a genetic algorithm, which achieved the same result using half the processing cost.

In all cases the model is simplified in some way: using a constant area lagoon (to render the ebb and flood cycles to behave the same); a fixed tidal series or discrete set of fixed tidal amplitudes; a limited number of design parameters; and/or limited number of operational parameters. These simplifications severely limit the use of the models as a design/sensitivity analysis tool when applied to a specific barrage scheme.

1.3   Barrage operation and 0-D modelling theory

Figure 1.2 shows a diagram of a tidal range barrage illustrating the main elements. The barrage consists of a semi-permeable wall across an inlet or estuary that creates a reservoir. It encloses an area of water and controls the flow between the reservoir and the open sea. The usable reservoir is the volume of water created by the difference in water level inside and outside the barrage. This is regardless of which level is higher, since electricity can be generated by flow into or out of the barrage.
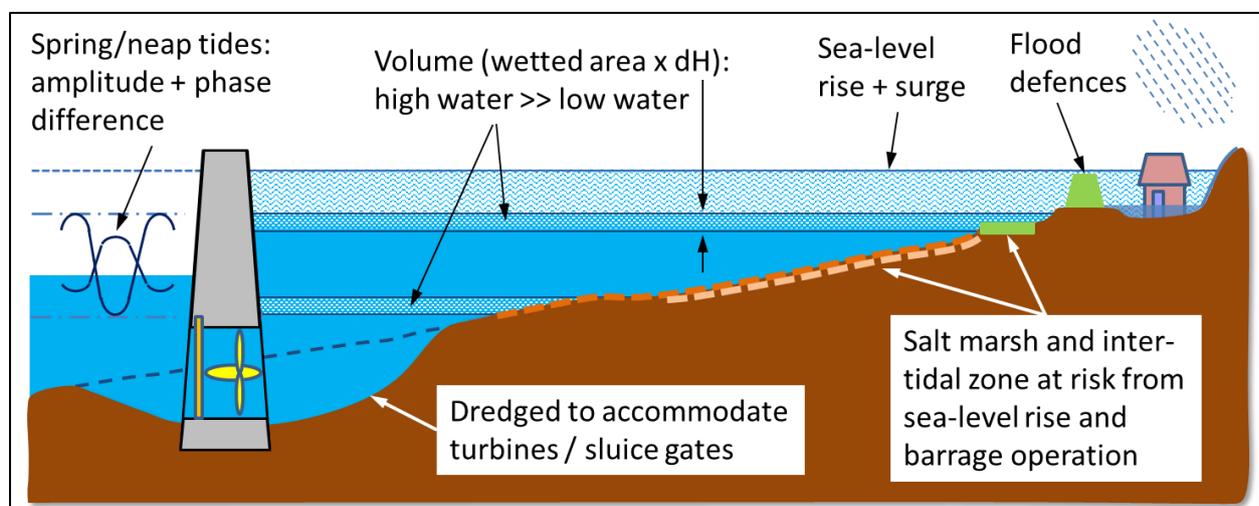
Figure 1.2 Simplified cross-section of a tidal range barrage.

Tidal range energy generation is calculated from the potential energy held by the water:

$$Energy = mgh \qquad (1)$$

where *(m)* is the mass of a volume of water, *(g)* is gravity and *(h)* is the height of its centre of mass above or below the level of the sea. In hydraulics

$$Power = \rho g Q H \qquad (2)$$

where *(H)* is the head (in this case the difference in sea and barrage levels) and *(Q)* is the volumetric flow across the barrage. For constantly varying *(Q)* and *(H)*, the energy can be calculated by summing over short time intervals, during which the values are assumed to be constant

$$Energy = \rho g \sum_{i=1}^{n} Q_i H_i \, \Delta t \qquad (3)$$

Maximum theoretical output energy is achieved if all the energy stored in the reservoir is converted instantly when the head is at a maximum [14].

As discussed above the 0-D modelling approach was adopted for this study. A common implementation of 0-D uses the backward difference method:

$$Z_{i+1} = Z_i + \frac{Q_i}{A_i} \Delta t \qquad (4)$$

The backward difference method determines the barrage reservoir level for the next time step from the previous one, i.e. the new reservoir height is calculated from the old height plus the volume change of the reservoir divided by the wetted (exposed) surface area. The volume change also includes any river inflows or precipitation falling directly into the reservoir.

The potential energy increases as the head develops, and this is achieved by holding the barrage reservoir level constant as the tide height changes. For an incoming tide, water is allowed to flow through the barrage until high tide is reached, at which point the flow is stopped and a head develops as the tide falls. When the head has built up sufficiently, the gates to the turbines are opened to allow the water to flow and generate electricity. This sequence describes three phases of the operation; sluice, hold and generation.

Whenever there is flow across the barrage the level of the reservoir will change, and whether that change is slower or faster than the tide, will determine if the head increases or decreases. It may increase initially when the generation phase first starts, but ultimately the head will fall as the low tide mark approaches and the tide slows.

Figure 1.3 shows a series of time series graphs that describe the operation and behaviour of a barrage system. The first shows the head, the tide and barrage levels; the second the flow rate across the barrage together with the natural tide flow rate assuming no impediment; and the third shows the power generated. The background colours indicate the operational phases with a distinction made between the ebb (E) and flood (F) cycles. The ebb cycle is defined here to occur when water leaves the reservoir (outgoing tide) and the flood cycle when water enters the reservoir (incoming tide). The turbines can also be used as pumps to increase the head prior to the hold phase. Pumping is used to perform two functions, as shown in the modelled case below. First, it increases the net energy generated by increasing the usable reservoir volume when the head is low and releasing that volume when the head is high. Second, it can maintain the tidal range within the barrage to minimise the environmental and ecological impact on the Bay. The sequences of operation are hold, generation, sluice and pump, repeated for both ebb and flood cycles. Strictly, the sluice and pump phases would

not take place if the following generation phase was not performed, and it could be argued the sequence should be sluice, pump, hold and generate.
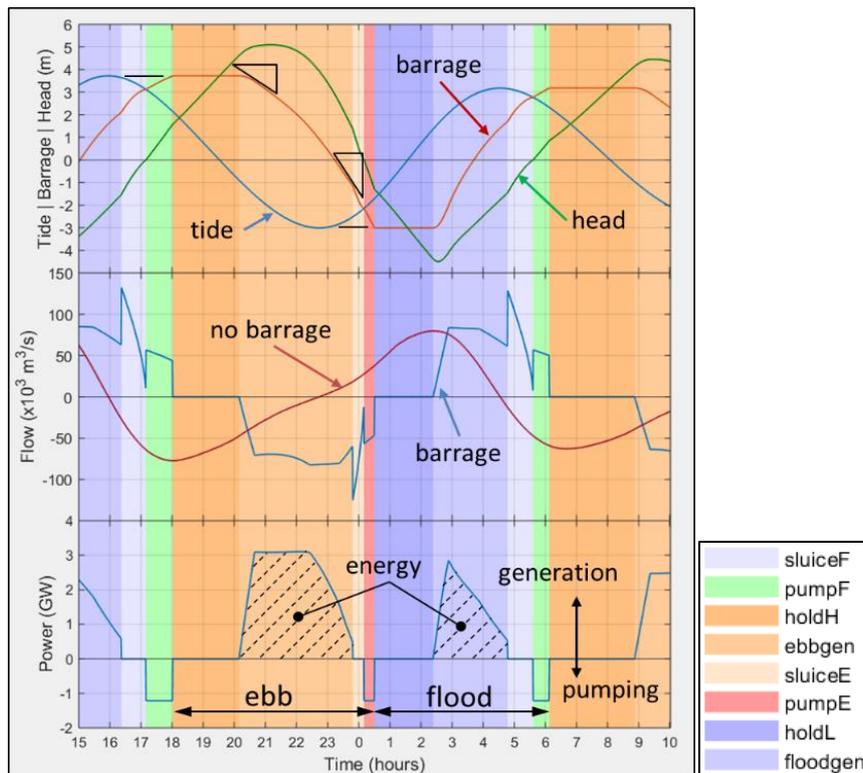
Figure 1.3 Annotated time-series graphs of 19 hours of barrage operation: barrage level, tide level and head (top); flow across the barrage with and without the barrage (middle); power (bottom). The background colours indicate the operational phase of the barrage.

A characteristic of a barrage, as opposed to a lagoon, is the variation in wetted area with the level of the water. A given volume of water distributed evenly over a large wetted area will have a smaller vertical thickness than an equal volume distributed over a small area. The consequence is that the barrage level falls more quickly for the same flow rate when at a lower level than at a higher level (this can be observed in the slope of the barrage level curve). At the start of an ebb generation phase the head increases slightly before dropping off; at the start of a flood generating phase the head drops rapidly. This has a significant impact on the amount of energy generated. A lower head generates less power. It also leads to lower flow, and lower flow also generates less power. On top of this the generation phase is shorter; less power over less time leads to even less energy.

The transitions between phases are triggered when certain conditions are met. The sluice phase stops when the tide and barrage levels are equal, i.e. when the head is zero. The pumping phase starts immediately following this and continues until the required reservoir level is reached. This is followed immediately by the hold phase. The generation phase start and stop times are determined in order to maximise the energy generated.

As stated above, the maximum theoretical energy assumes generation occurs instantly (infinite flow) at the maximum head, i.e. when the reservoir and tide levels are at opposite ends of the tidal range. In practice the flow is limited by the number and size of the turbines, and the generation phase takes a finite time. During this time the head is changing, and the turbines operate at something less than the maximum. More turbines result in a shorter generation phase and allow the timing to better coincide with the higher heads. Fewer turbines result in a longer generation phase and force operation to extend into periods of lower head. Too few turbines mean the reservoir cannot be exploited fully

before the head is lost as the tide reverses. This can also delay the transition to the sluice phase and in turn delay when the equilibrium point is reached. Any delay past the high/low tide mark means there is a shortfall in the barrage level with respect to the tidal limit, requiring more pumping to achieve the required reservoir level. To illustrate this Figure 1.4 shows a comparison of results for two simulations, with 120 and 160 turbines respectively. Using fewer turbines reduces the Annual Energy Production (AEP) from 9.84 TWh to 8.03 TWh; the maximum head is lower and the tidal range levels cannot quite be achieved with pumping.

(a)   *(b)*
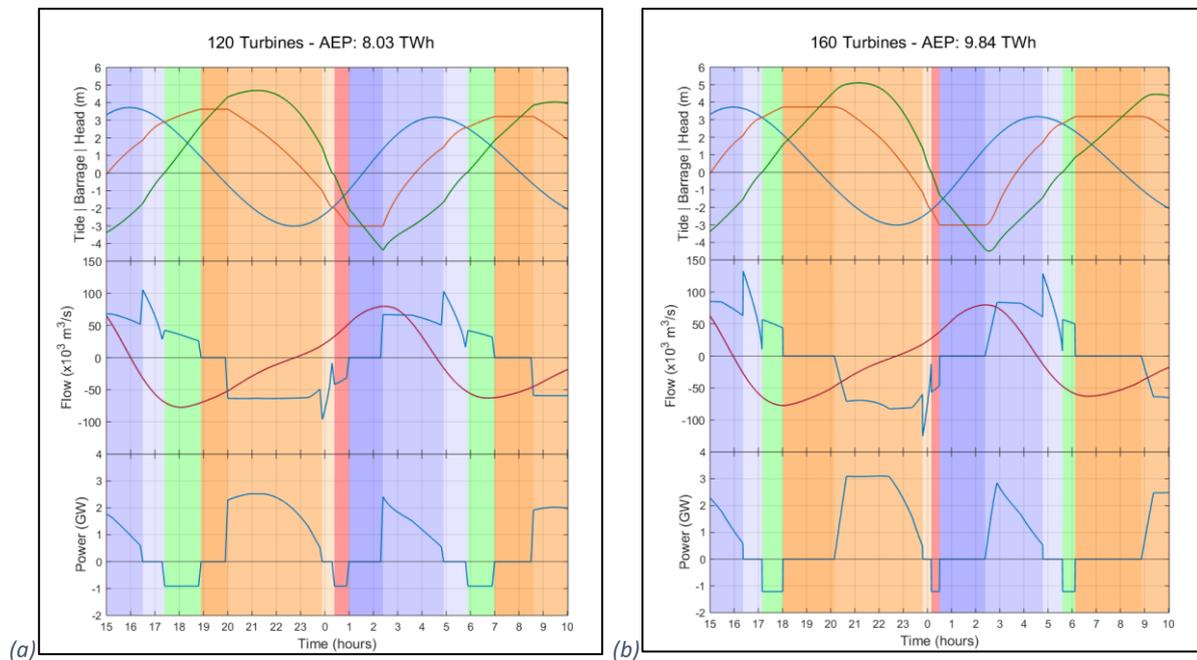
Figure 1.4 Comparison of the barrage operated for different numbers of turbines; (a) 120 turbines and (b) 200 turbines. More turbines allow more of the available energy to be extracted from the reservoir. A higher average head was achieved by delaying the start of generation.

1.4 Cyclic intermittency and power balancing

Tidal range schemes typically generate electricity in discrete blocks, each half or complete tide cycle, interspersed with periods of a similar duration with no generation or as a consumer of electricity for pumping (see Figure 1.3 - power curve). This cyclic intermittency of supply to the electricity grid has to be balanced against demand. One way to facilitate this is to use some means of energy storage to capture excess energy during generation, and to release it during the intervening interval, in order to smooth out the supply. Alternatively, it can be left to the National Grid to use its existing mechanisms to match supply with demand. The viability of both options was considered.

As a preliminary step, the power and energy characteristics of the barrage were determined, together with today's electricity demand profile from the electricity supply data. Potential energy storage scheme solutions were explored, ranging from large standalone schemes capable of matching the power and energy of the barrage to much smaller scale schemes for smoothing or power balancing. In addition to looking at storage options, the study considered the energy balancing demands placed on the grid by the large periodic step changes in power supplied and demanded by the barrage.

As part of the process of changing to a carbon-free economy, the demand for electricity is projected to increase significantly and the daily profile is expected to change with a shift to more overnight use [15]. Exactly how it will change, and how supply will change to accommodate it, is uncertain. Planning a large-scale dedicated storage scheme in the light of such uncertainty is a risky proposition

The idea of tailoring the barrage operation to favour the value of the electricity is not new [11], although there is only limited scope to advance or delay the timing of the generation cycle. Alternatively, at times of low demand, the operating cycle can be adapted; the pumping phase can be omitted, or even the cycle as a whole. However, this approach would be at odds with a requirement intended to protect the inter-tidal zone by matching the tidal range. This highlights the need to fully understand what is required from an environmental standpoint; is it absolutely necessary to match every tide cycle?

1.5 The environment and sea-level rise

The barrage is designed to have a nominal operating life of 120 years and consideration must be given to what might happen during that time. The most obvious consideration is the effect of climate change on sea-level. Sea-level rise is already a reality and the predictions over the next 120 years vary widely due to uncertainties in the mechanisms involved and on future net greenhouse gas emissions. IMechE (Rising Seas: the engineering challenge) [16] states "… prepare for a minimum of 1 m rise in sea level this century but plan for 3 meters of rise".

Sea-level rise brings with it two major concerns: an increased risk of terrestrial flooding, and a threat to the ecology of transient inter-tidal habitats. There are upwards of 75 km^2 of low-lying land surrounding Morecambe Bay that are protected by approximately 50 km length of embankments plus one-way river flow gates and pumping stations. Sea-level rise will threaten this infrastructure; embankments will need to be raised to prevent future breaches. Climate change is also predicted to increase the likelihood and severity of storms, whilst they become less predictable. The main catastrophic flood risk at present is from rivers following heavy rain [17]. The ability to drain the land is impeded by high tides and will become increasingly difficult with sea-level rise.
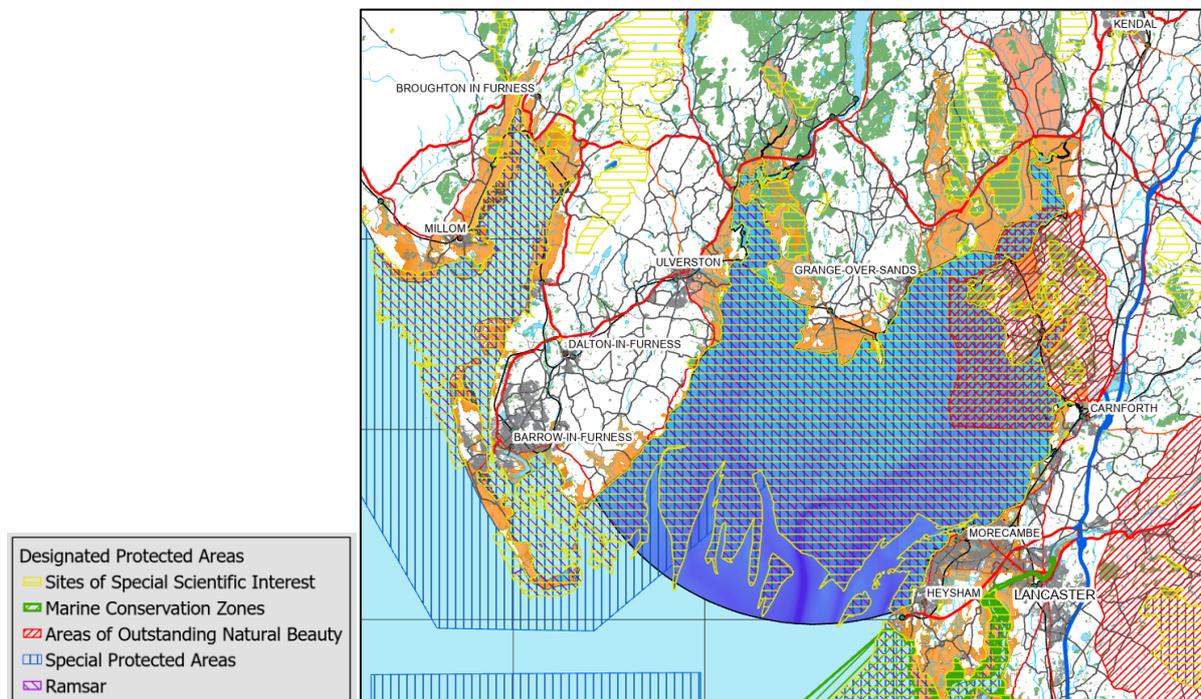
Figure 1.5 Major designated protection areas around Morecambe Bay.

There are multiple overlapping environmental and ecological designations covering Morecambe Bay, with the aim of protecting both the whole ecosystem and specific characteristics (see Figure 1.5 for a map of the major designations, [18]). In addition, there are several specific areas allocated to shell fishing, a long-standing traditional industry. Sea-level rise will force change on the inter-tidal zone,

which will be pushed further inshore, where it will meet man-made resistance, preventing its natural progression. The mud flats will shrink as they remain submerged longer, and the salt marsh areas and other inshore environments will diminish or be lost altogether. Saltmarsh acts as an important carbon sink and should be protected where possible. The presence of a barrage can change the nature of the inter-tidal zone within it, but it can also operate as an environmental management scheme. Importantly, it can limit the height of the high tides to alleviate tidal flooding, mitigate riverine flooding and maintain the current tidal range, thus preserving existing habitats. The criterion of maintaining the current tidal range has been applied as part of the study as a proxy for an environmental requirement.

The uncertainty in the extent of sea-level rise, poses a problem for the design of the barrage. Even if the barrage is not built initially to accommodate the more extreme predictions, it might be prudent to design for future augmentation. In addition to making sure the height of the barrage is adequate, a raised sea-level increases the reservoir volume (the wetted area increases with water level) potentially requiring additional turbines to fully exploit the extra resource. The Thames barrier serves as an example where sea-level rise was not given due consideration in the design and it is now considered likely to be breached.

A barrage, if designed and operated appropriately, can mitigate the impacts of sea-level rise on the inter-tidal zone and can help satisfy the Government commitment to the protection of designated ecosystems. If the ecological and environmental considerations dictate that the tidal range within the barrage must be maintained at pre sea-level rise levels, then how might this affect the design and/or AEP? Ultimately, how does this impact the cost of construction and the operational performance? Similarly, if the barrage is used for flood prevention, what is the cost of deviating from maximum energy generation?

Whether intentionally or otherwise a barrage will have an impact on the environment and the ecology, and it is important our decisions at all stages in the design and development process are suitably informed. Incorporating these considerations in the modelling from the outset will help to address valid concerns and facilitate the acceptance process. The environmental considerations should encompass the complete cycle of the barrage and continue all the way through to decommissioning. Because of the potential mitigating effect of the barrage on sea-level rise, it is no longer sufficient to dismiss such schemes out of hand on environmental grounds; the consequence of doing nothing must be considered and presented.

For the purposes of this study, the environment and ecological considerations were based on the assumption that changes to the tidal characteristics within the barrage should be kept to a minimum. No attempt was made to rank the various characteristics in terms of their impact.

1.6   Report structure

The body of the report is set out in three main sections followed by a discussion and conclusion:

Section 2 describes the barrage model: the bathymetry model and the reservoir definition; the tidal model; the turbine and pump models; turbine control; the model parameterisation; and the program structure and Graphical User Interface (GUI). It also covers the importance of design-by-design optimisation and optimisation schemes.

Section 3 addresses the scenarios considered/modelled: how the increased risk of flooding from sea-level rise can be mitigated by the barrage and at what cost; how the number of turbines determines the amount of energy that can be extracted from the reservoir; and how the position of the barrage and sea-level rise affect both the available and recoverable energy.

Section 4 looks at energy storage: the generated energy characteristics of the barrage; energy storage technologies; current UK demand; electricity energy balancing; and asks if conventional energy storage is necessary, particularly considering expected changes to demand and usage patterns.

The Appendices provide information as follows:

Appendix A Provides details of the design-by-design optimisation functionality of the modelling program.

Appendix B Provides details the Graphical User Interface (GUI) and execution of the modelling program.

Appendix C A tabulated list of selected design scenarios, with the optimised operational function coefficients and the generated energy values.

Note: All maps, plots and photographs are the authors own work unless otherwise stated. Base map data were downloaded from the Ordnance Survey [19, 20].

2 Data and model

The section provides details of all aspects of the modelling, including the bathymetry model and reservoir definition; the tidal model; the turbine and pump models; control of the turbine; model parameterisation; operational parameterisation; optimisation of the operational parameters; program structure and the program's Graphical User Interface (GUI). It also covers the importance of design-by-design optimisation and optimisation schemes.

2.1 Bathymetry and reservoir modelling

An important element of the model was the reservoir. For a tidal range scheme this is defined in terms of a function of the wetted area with elevation. Integrating over this function with respect to depth yields the volume. This relationship provided a key element of the backward difference modelling method where the new surface elevation was calculated from the change in volume.

Such a function was constructed from the cumulative sum of the bathymetry within the confines of the barrage. The elevation was not restricted by the level of the highest astronomical tide. Increasing the water level above this would generally result in a larger volume, unless the area was restricted. Figure 2.1 shows that above about 6m elevation (OS datum) the area continues to increase only slightly and reflects where the extensive use of flood prevention embankments (see section 3.1) and steeper topography act to confine the bay extents.
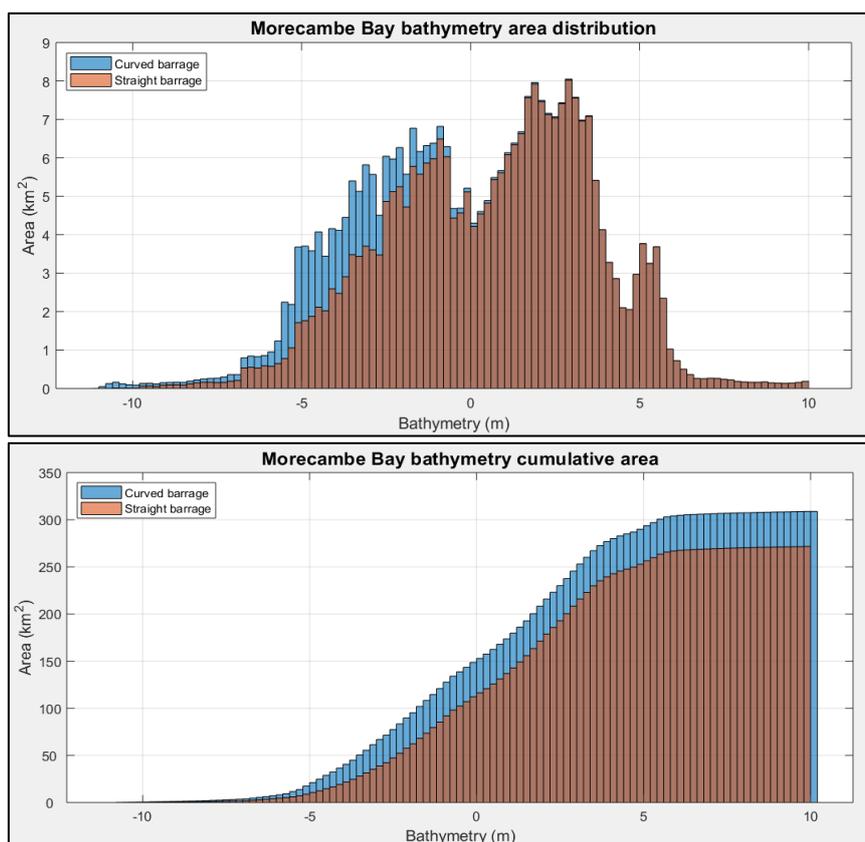
Figure 2.1 Bathymetry distribution for the confined extents of Morecambe Bay (curved and straight barrages): area distribution (top) and the cumulative distribution (bottom).

In practice, the backward difference calculation was modified to use the volume vs depth relationship (Figure 2.2), derived by integrating the cumulative area curve. The function was monotonic and was represented by a series of elevation-volume pairs. The new barrage volume was calculated by adding the discharge volume for this iteration to the volume at the start of the iteration. The new depth was

found by interpolating the curve in reverse. The reason for this was to address the "wind-up" observed with the backward difference method. This describes how the net volume discharge across the barrage accumulates over time (Figure 2.3). This does not occur with lagoons, and was more severe the longer the time step. The error (difference) between using a 0.1 hour time step and a 0.01 hour time step was reduced from ~10% to ~3%.
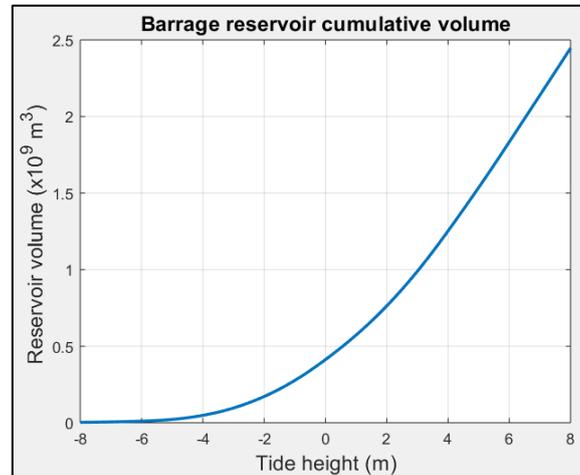
Figure 2.2 Cumulative volume curve for the curved barrage.
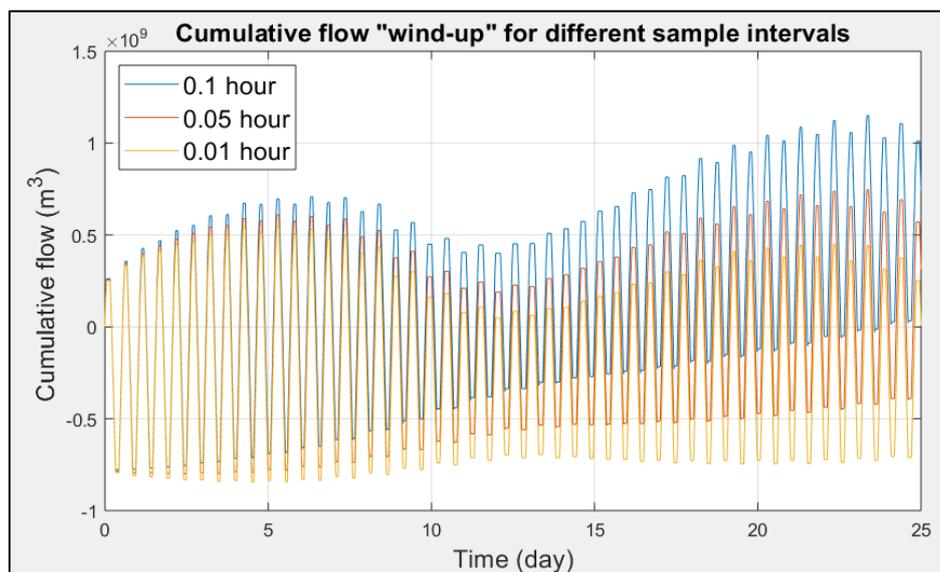
Figure 2.3 Cumulative "wind-up" of net flow across the barrage with the backward-difference method. This does not occur with a lagoon where the wetted area is constant.

The bathymetry, together with the position and the shape of the barrage determine the reservoir volume to elevation relationship. Each new barrage position requires its own function. For a lagoon, assuming the wetted area was constant over the complete elevation range, then the relationship would be a constant.

Separate functions were generated for the two different barrage locations considered in the study. The curved barrage substantially increased both the area and volume of the barrage reservoir compared to the straight barrage. Since the extended area was predominantly in deeper water there was a preferential benefit to the flood generating cycle. This was seen in the percentage increase in area at the different depths. At -3 m the area has increased by about 60% and at 5m by about 10%. If a whole range of different barrage positions were required to be modelled, then a more sophisticated way of building up the function from a series of segmented areas could be devised.

If pumping were used to increase the reservoir volume beyond the tide limit, then the substantial increase in volume below low tide created with the larger area would allow more efficient use of the pumping capacity.

The task of building an accurate bathymetry surface was not trivial. The primary source of information used was the 2017 1m DTM (Digital Terrain Model) data derived from LIDAR (light detection and ranging) data [21]. LIDAR cannot image below the sea surface and, although the data were acquired at low tide, the DTM consequently only extends down to -3m elevation OS. Bathymetry derived from the Admiralty chart (Morecambe Bay, 2010, based on data surveyed prior to 1970 [22]) was used to complete the model over the full extent of the barrage. The drainage channels at all scales migrate over time and the larger channels in the deeper water are no exception. The bathymetry below -3m was built by outlining contours in keeping with the characteristics from the Admiralty chart data and tying in with the channels from the DTM model.
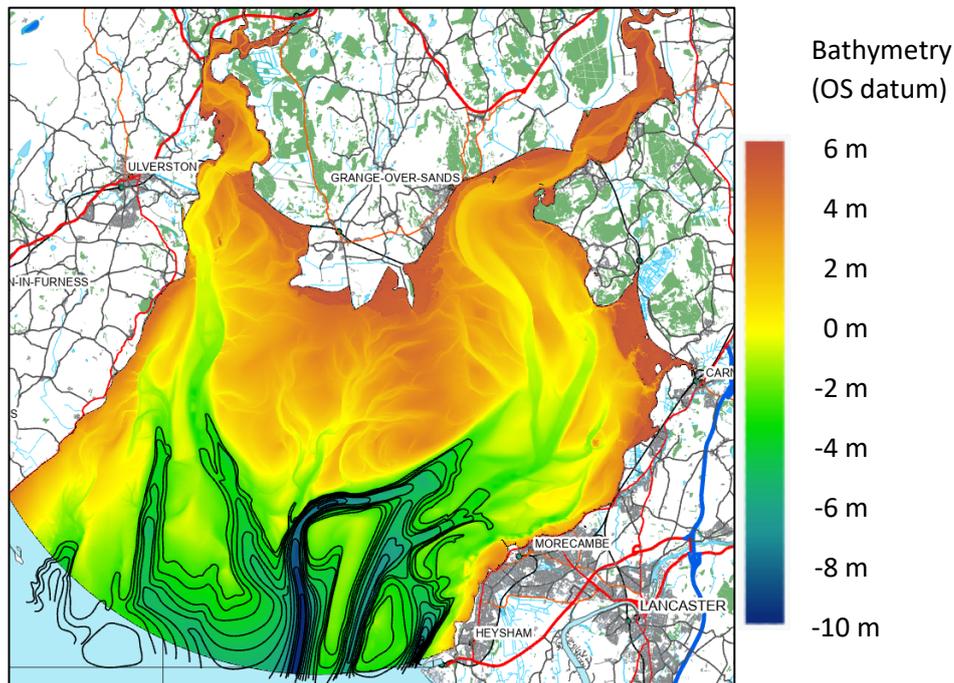
Figure 2.4 The bathymetry map and subsequent reservoir definitions, were constructed from LIDAR DTM data and manually digitised contours (annotated) based on the Admiralty chart.

ArcMap functions were used to resample, mask, surface-fit, smooth and merge surfaces, until a complete bathymetry map was constructed. Although the positions of the deeper channels were not particularly accurate, there was sufficient information from the DTM data and Admiralty chart to roughly approximate their positions. The contours in Figure 2.4 show the position of the manually digitised data. Since the purpose of the model was to derive an elevation-area function, the actual channel positions were not important. The part of the function relating to this part of the Bay would only be used if the barrage level drops below -3m. When operating within the tidal range this would only happen at spring low tides by ~0.5m. If pumping was used and the level was allowed to drop below this then the accuracy of this part of the model would be more critical.

The LIDAR data has an absolute height error of less than ±15 cm RMSE (root mean square error) and horizontal accuracy of ±40 cm RMSE. If the maximum vertical error was systematic across the whole area, this would be equivalent to a change in sea-level of the same amount, and the error in the energy predictions could be easily calculated.

Building the reservoir model required an additional step to exclude low-lying land protected from flooding by the embankments. Relying on the elevation alone would include these areas as part of the reservoir. The embankments, and in places the railway where it acts as a flood barrier, were digitised and applied as a boundary mask.

2.2 Tide modelling and datums

There are two mapping datums of relevance to the study, these are the Ordnance Survey (OS) Datum at Newlyn and the Admiralty Chart datum. Chart datums are set to the local "approximately lowest astronomical tide". The same chart datum applies to Morecambe Bay and the tide data at Heysham, where chart datum was 4.9m below OS datum [23]. Figure 2.5 show the datums and predicted tidal levels between 2008 and 2026.
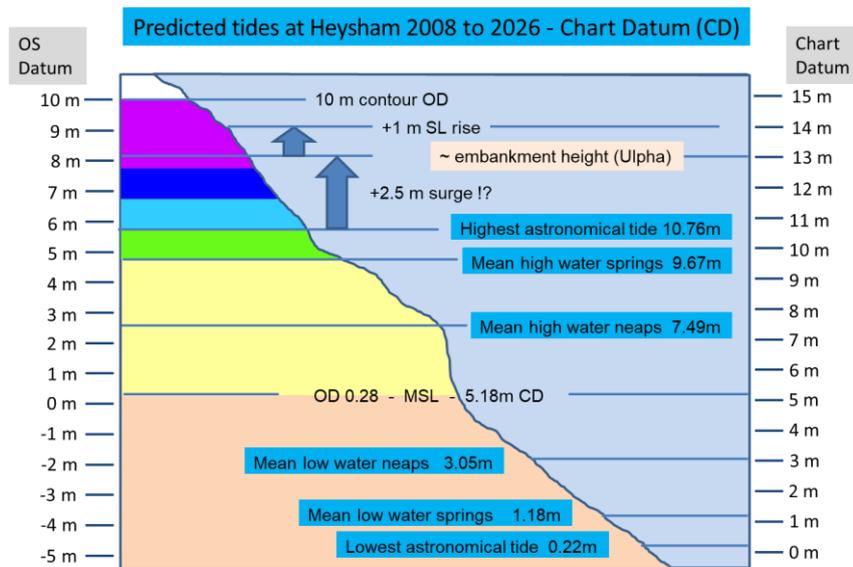
Figure 2.5 Ordnance Survey and Chart datums at Morecambe Bay with predicted tidal levels between 2008 and 2026.

The tide model used in the study was derived from the tide data at Heysham (National Oceanography Centre (NOC). National Tidal and Sea Level Facility [23]). The site's published measured data were at a 15 minute sample rate, together with a residual value. The residual was calculated as the difference between the measured value and the model data, where the model was the NOC model for Heysham. Removing the residual values from the measured values gives the model values. Short of the actual constituents (not freely available) this represented an adequate compromise. The NOC model is updated on a regular basis.

These data were run through a tidal analysis program in MATLAB called UTide [24]. The complete dataset for 2018 was analysed and the 10 tidal constituents with the highest amplitude (greater than 0.1m) were used in the tide series code. The UTide long term trend parameter value was not included in the model and although tides can be generated with any start date, the further it is removed from the start of 2018 the less accurate it will be.

Tidal predictions were made based on the superposition of these 10 main harmonic constituents using the following formula:

$$h = \sum M_2 \cos\left(\omega_{M_2} t + \phi_{M_2}\right) + S_2 \cos\left(\omega_{S_2} t + \phi_{S_2}\right) + \cdots \qquad (5)$$

where ω is the angular velocity (rad/hour), *t* is time (hours) from the start of the series, and ϕ is the phase at *t*=0. The 10 constituents are as follows (where ϕ_{t0} is the phase at 1 Jan 2018 00:00):

Constituent:	M2	S2	N2	K2	M4	NU2	L2	K1	MS4	O1	
Amplitude:	3.1638	1.0271	0.6042	0.2988	0.1999	0.1361	0.1279	0.1227	0.1164	0.1114	m
ω:	0.5059	0.5236	0.4964	0.5250	1.0117	0.4976	0.5154	0.2625	1.0295	0.2434	rads/hr
ϕ_{t0}:	1.0612	6.1366	1.6528	3.1241	2.9217	1.7977	3.7682	2.9775	1.5813	6.0057	radians

The tide model was assumed to act uniformly over the entire bay. In reality, the tidal range varies slightly; the Admiralty chart quotes the Mean High Water Spring (MHWS) and Mean Low Water Spring (MLWS) at Haws Point (just outside the barrage near Barrow) as 9.2m and 1.0m respectively, and at Heysham as 9.5m and 1.1m respectively. Although these values were not entirely insignificant, the error inherent in the 0-D approach, together with the assumption that the barrage has no impact on the behaviour of the tide, render any attempt to accommodate for this variation futile. A more accurate representation of the tide warrants a more sophisticated modelling method.

2.3 Turbine and pump models

2.3.1 Turbine model

The turbine model used in the simulation is based on the Andritz Hill chart (Figure 2.6) published by Aggidis and Feather [8]. A Hill chart is the typical way of representing turbine performance data; it represents the measured performance of a 1m diameter model at 1m head.
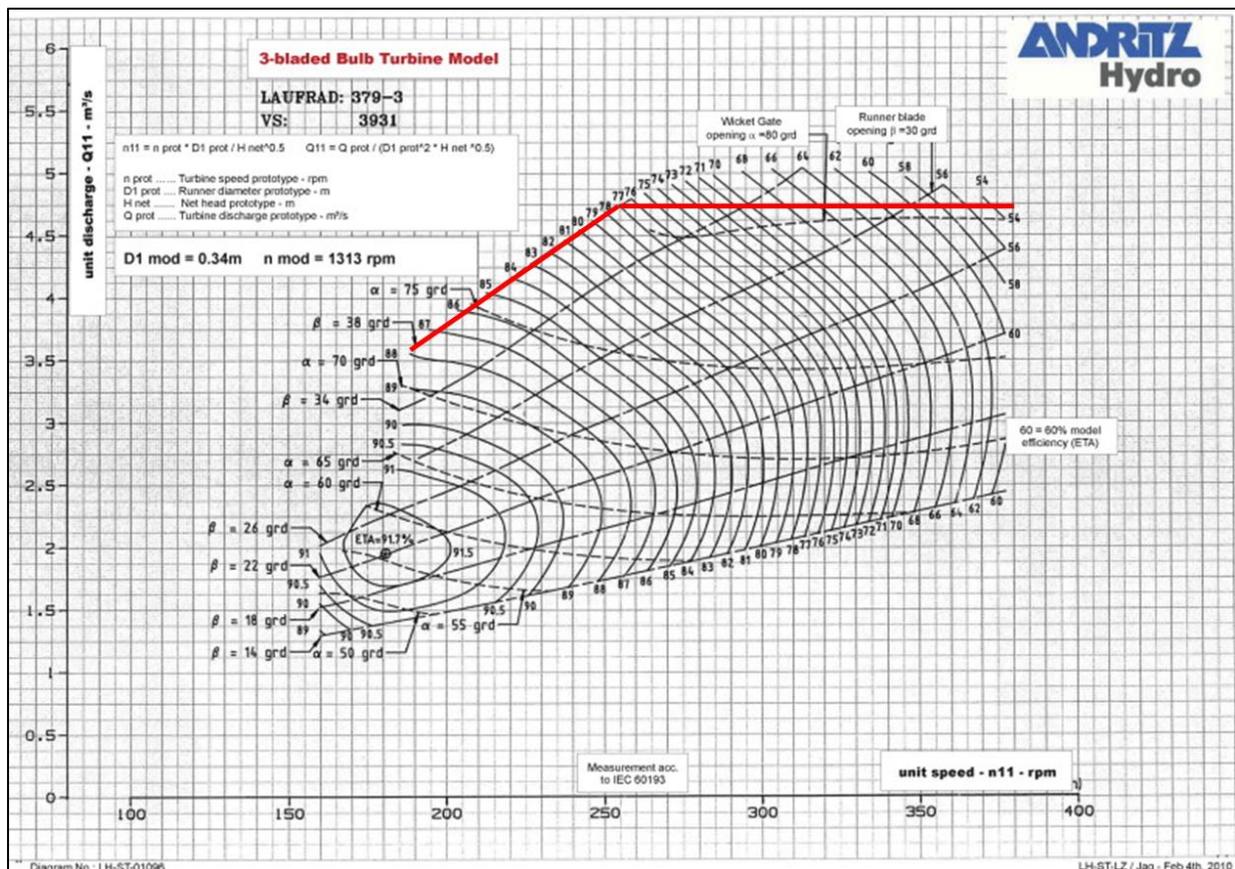
Figure 2.6: Andritz Hydro 3-bladed bulb turbine model hill chart. The red line indicates the operational envelope used in the modelling and was chosen to provide maximum discharge [8].

It enables the performance of any size of turbine to be determined using the following equations [25]:

$$n_{11} = \frac{S_P \times D}{\sqrt{H}} \qquad (6) \qquad \text{and} \qquad Q = Q_{11} \times D^2 \times \sqrt{H} \qquad (7)$$

15

where n_{11} is the model shaft speed *(rpm)*, D is the diameter *(m)*, H is the head *(m)*, Q is the turbine flow rate *(m³/s)*, Q_{11} is the model flow rate *(m³/s)*, and Sp is the turbine shaft speed *(rpm)*.

For double regulation where a fixed gear ratio, *Gp,* couples the turbine directly to the fixed electricity power grid frequency, *f (Hz)*, the shaft speed is given by

$$S_P = \frac{2 \times 60 \times f}{G_P} \qquad (8)$$

The gear ratio had to be chosen such that the operational range of heads expected at the barrage fell within the operating limits of the turbine.

The Hill chart also includes the efficiency characteristics and the operational envelope. The operational envelope is dictated by the minimum and maximum head, the maximum power and the cavitation limits. Inspection of equation (6) shows the lower limit of n_{11} corresponds to the maximum head, and the upper limit to the minimum head. Equation (7) shows the flow is proportional to Q_{11}.

From the Hill chart it can be seen that the efficiency is highest at lower n_{11} and lower Q_{11} values, with a larger variation over the range n_{11} compared to Q_{11}. Figure 2.7 shows the slope of the efficiency surface to be low around the maximum efficiency point creating a broad operating region of high efficiency.

The free flow through the turbines and sluice gates during the sluice phase is given by

$$Q = C_d A \sqrt{2gH} \qquad (9)$$

where C_d is the coefficient of discharge and A is the cross-sectional area. A C_d greater than one reflects there is a venturi effect because the feeder tubes are of a larger diameter than the turbine. The choice of 1.1 was decided in consultation with Prof. George Aggidis.
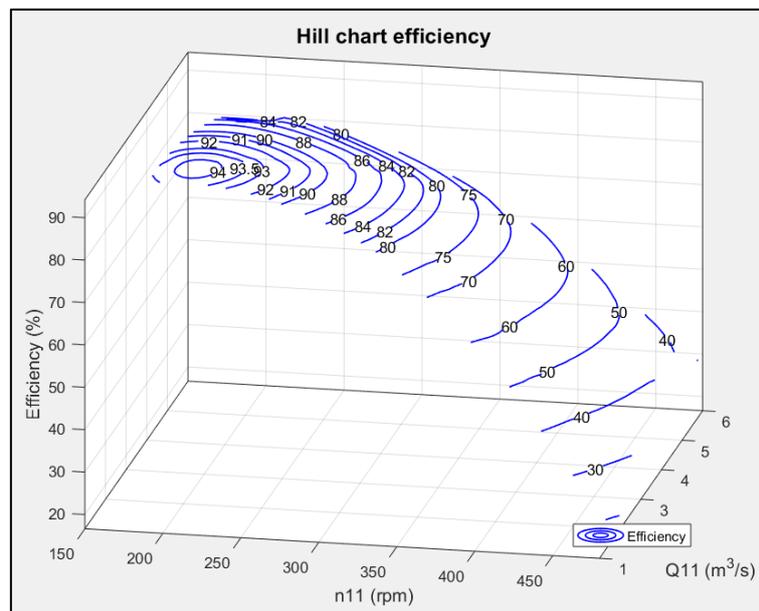
Figure 2.7 Andritz Hydro 3-bladed bulb turbine Hill chart efficiency plotted as a surface in 3-D showing a broad flat region around peak efficiency.

Section 1.3 discusses how the theoretical maximum energy is achieved if the reservoir is depleted instantly (with infinite flow) at the maximum head. In practice, the number of turbines is limited by cost and space, and the flow rate is similarly limited. In addition, how the turbine is operated would

also affect the flow rate. The Hill chart shows there is a trade-off between higher flow (Q_{11}) and higher efficiency (η). Halving the model flow (Q_{11}) from 4 m³/s to 2 m³/s resulted in an increased efficiency (η) from 85.5% to 91.7%. Higher flow allowed a shorter generating phase at a higher average head, but at lower efficiency. Conversely, operating at maximum efficiency and lower flow, lengthened the generating phase and lowered the average head. Doubling the flow doubled the power (equation (2)) and allowed generation to take place at a higher average head. The slightly higher efficiency for the slower flow rate case can only result in more energy if the generating time is short enough to effectively operate at the same average head. The only time where there may be an abundance of discharge capacity is during neap tides, and dropping the turbine operation point below the maximum flow to gain more efficiency may prove beneficial. If the discharge rate is too slow, there is not the ability to exhaust the reservoir before the tide reversers and the head is lost completely.

The size of the reservoir is dependent on the amplitude of the tide. In Morecambe Bay the difference in amplitude between neap and spring tide is roughly a factor of two. This corresponds to a factor of four for the available energy (roughly double the volume and double the head). There has to be a sufficient number of turbines to utilise the spring tide energy effectively which leaves some scope to reduce the flow during neap tides and operate at a higher efficiency. There is a competing factor however, which is that the barrage level can more closely approach the peak tide if the generating phase transition is not delayed. Prolonging the generating phase by operating at a lower flow and higher efficiency can limit the head for the next cycle.

The turbine speed and flow are controlled for a given head by the inlet guide vane and runner blade angles. As the head changes, so the vane and blade angles are adjusted to maintain the desired operating conditions. For double regulation, one of the conditions was constant speed, which together with the head dictates the model (n_{11}) speed. The flow was otherwise free to be adjusted along the given n_{11} line, to achieve anything between maximum power (and flow), and maximum efficiency. With triple regulation, where the turbine speed is decoupled from the National Grid phase velocity, there is more freedom to position the operating point anywhere within the performance envelope, and in practice this allows more scope to operate the turbine at higher efficiency.
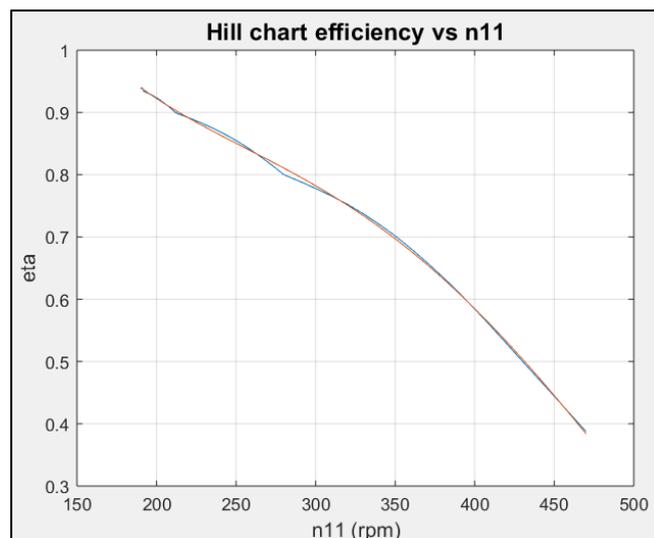
Figure 2.8 Function fitted to the digitised efficiency vs model speed (n_{11}) data from the hill chart.

For double regulation, the turbine operation was limited to the peak flow edge of the operational envelope on the Hill chart (red line Figure 2.6), i.e. for any given value of n_{11}, Q_{11} was then chosen to be the maximum available. Linear functions relate Q_{11} to n_{11} in a piecewise manner. A polynomial

function fitted to the digitised values along the line of operation related efficiency to n_{11} (Figure 2.8). Only when the maximum power limit was exceeded, was the Q_{11} value reduced below the line by an amount needed to maintain maximum power. This only occurred at very high heads and where the n_{11} value was at or close to its minimum. The efficiency value was not adjusted in the calculation since the change was quite small.

For triple regulation, the operation of the turbine was limited to just two options; constant speed and constant Q_{11}. These can be considered as end member options.

For constant speed, the value of n_{11} on the Hill chart increases proportionally to the square root of the decrease in head (equation 6). The speed was determined such that the range of n_{11} values, corresponding to the head range of the individual cycle, fell within the operating range of the turbine. The actual speed was determined via the optimisation process and was a function of the tidal range; otherwise the operation was the same as for double regulation. As the head falls, the operation point moves along the upper edge of the envelope on the hill chart from left to right and the efficiency is reduced as it does so.
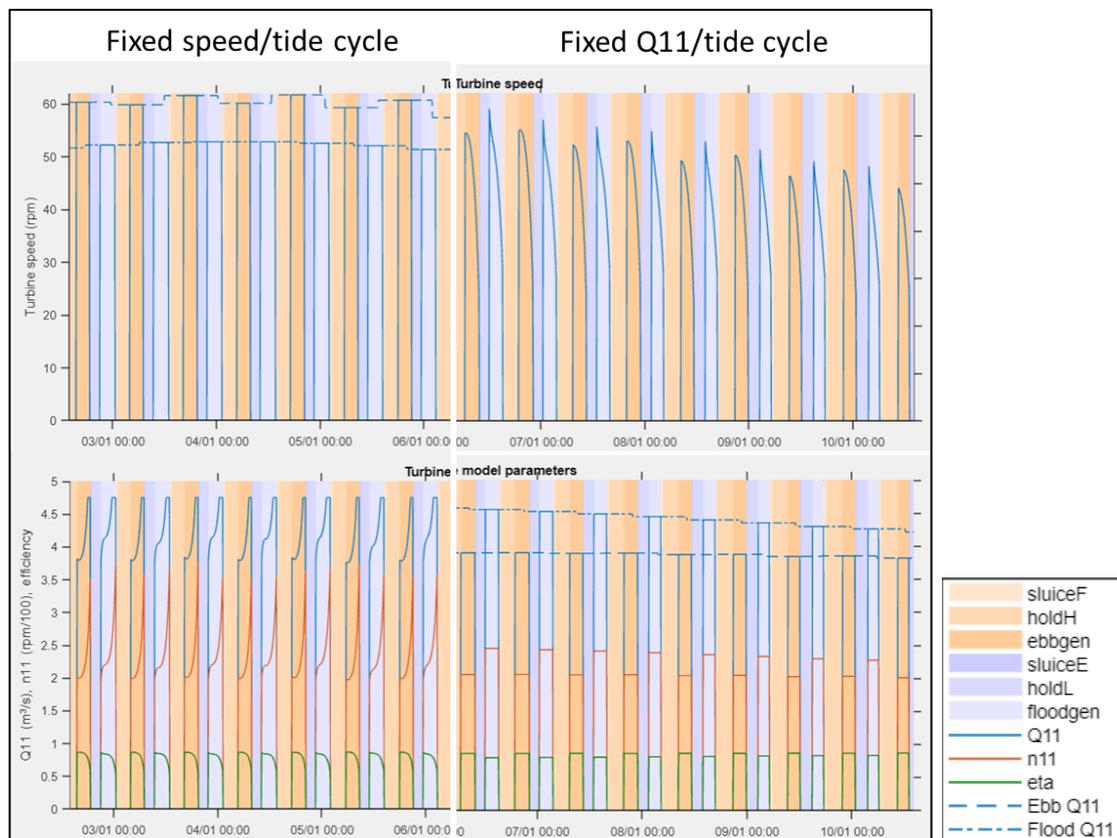
Figure 2.9 Turbine operation using the two options: constant turbine speed for each generating cycle (left hand side); constant model flow for each generating end cycle (right hand side).

For constant Q_{11}, the operation was constrained to a horizontal line on the Hill chart and the value of n_{11} was chosen to be the lowest value possible such that the turbine was operated at the maximum available efficiency.

The rationale for choosing these options was for constant speed to enable a more equivalent comparison with double regulation, and for constant Q_{11} to provide a means to maintain operation throughout a cycle at a higher efficiency, if lower flow. Figure 2.9 shows the turbine speed and Hill chart parameter values for turbine operation using both methods.

2.3.2 Pump model

Turbine design is industrially very sensitive and it is difficult to obtain state of the art performance data. It is normally only supplied in consultation with a developer when a scheme is at an advanced stage of design and looking to go into deployment. The pump model used here was developed in consultation with Prof George Aggidis with insights from the performance data at La Rance [26]. Consequently, the pump model was fairly simplistic with a predominantly linear flow vs head relationship when operated at constant power (Figure 2.10(a) and (b)) giving an efficiency represented by the purple curve shown in Figure 2.12(a).

The similarity laws (Figure 2.11), for turbines with geometric and hydraulic similarity, were used to scale the model to different diameters and speeds (see Nechleba [27] for a derivation of the equations).
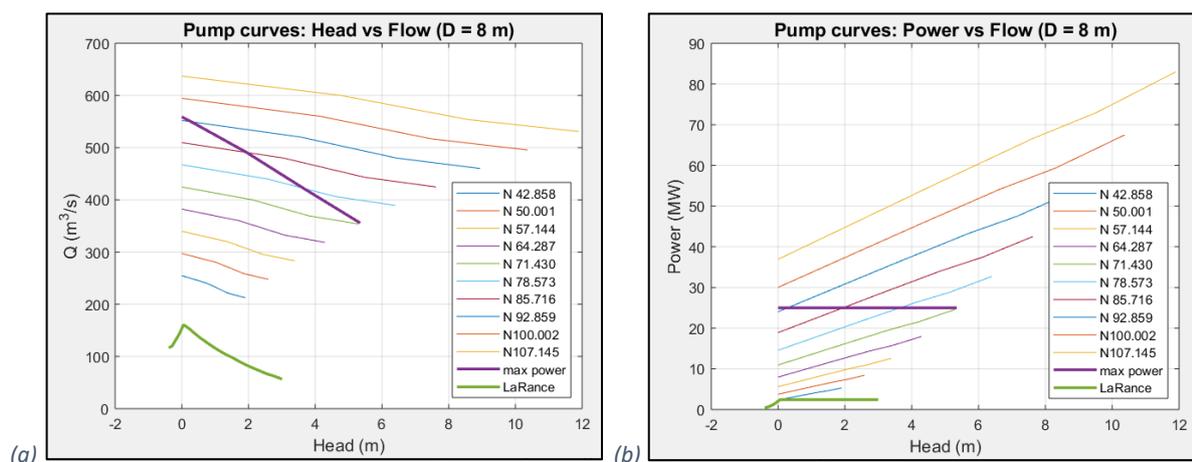
Figure 2.10 Pump performance curves at 25 MW power: (a) flow rate vs head; (b) power vs head – held constant.

When operated at 25 MW, approaching the maximum rated power of the turbine (30 MW), the expected increase in net energy when pumping was not achieved. From energy data at La Rance, Hillairet and Weisrock [28] reported an uplift on the order of 10%. Using 0-D modelling for a tidal range scheme at the Duddon, Yates *et al* [29] could only achieve energy gains of 11% for dual and 13% for ebb only modes.

$$\frac{Q_1}{Q_2} = \left(\frac{N_1}{N_2}\right)^1 \qquad \frac{Q_1}{Q_2} = \left(\frac{D_1}{D_2}\right)^3$$

$$\frac{H_1}{H_2} = \left(\frac{N_1}{N_2}\right)^2 \qquad \frac{H_1}{H_2} = \left(\frac{D_1}{D_2}\right)^2$$

$$\frac{P_1}{P_2} = \left(\frac{N_1}{N_2}\right)^3 \qquad \frac{P_1}{P_2} = \left(\frac{D_1}{D_2}\right)^5$$

Constant diameter — Constant shaft speed

where:
Q is the volumetric flow rate
D is the rotor diameter
N is the shaft rotational speed
H is the pressure head
P is the shaft power
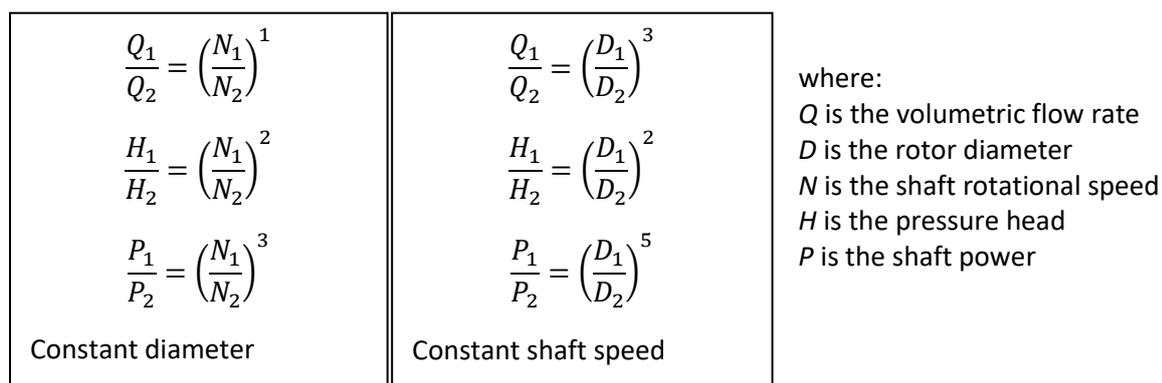
Figure 2.11 Similarity Laws for fixed proportions (see Nechleba [27] for a derivation).

A comparison of the pump efficiency curves between the modelled data for Morecambe Bay and the La Rance data, shows La Rance achieves a significantly higher efficiency at the same head (Figure 2.12(a)). It was observed that the pumps are operated at a quarter of the rated generating power at La Rance. When the power was similarly reduced at Morecambe Bay, the efficiency closely matched that at La Rance (Figure 2.12(b)) and net generated energy gains of 10% were achieved. The flow rate

achieved at zero head for 25MW and 7.5MW pump power was 560 m³/s vs 380 m³/s respectively - 30% of the power gave 68% of the flow (Figure 2.13 (a) and (b)). This was explored further by modelling the AEP for a range of pump power and maximum AEP occurred at around the 7.5MW rating (Figure 2.14). Limiting the power also limited the maximum pumping head.

(a)

(b)
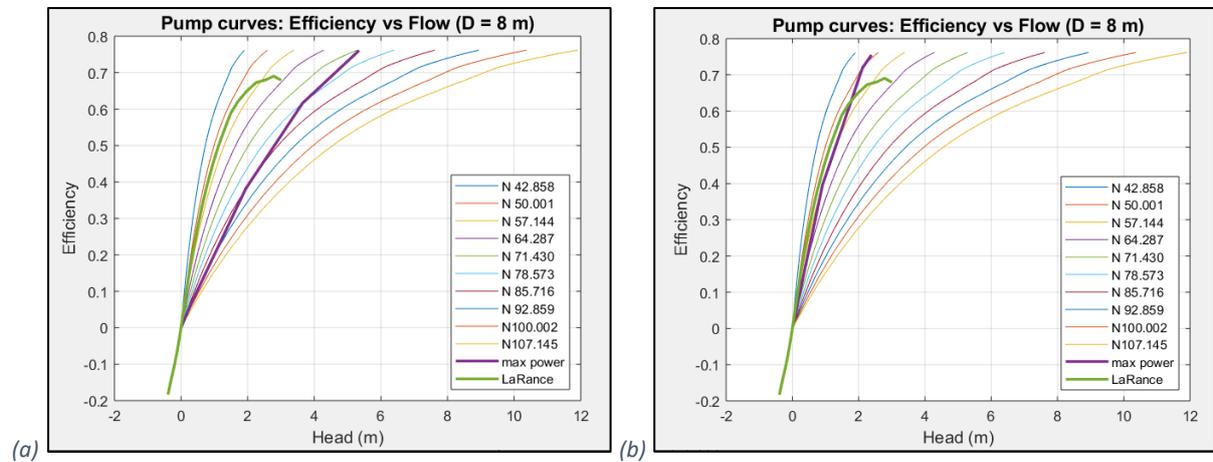
Figure 2.12 Pump efficiency curves: (a) operated at 25 MW power; (b) operated at 7.5 MW power – closely matches the efficiency achieved at La Rance.

(a)

(b)
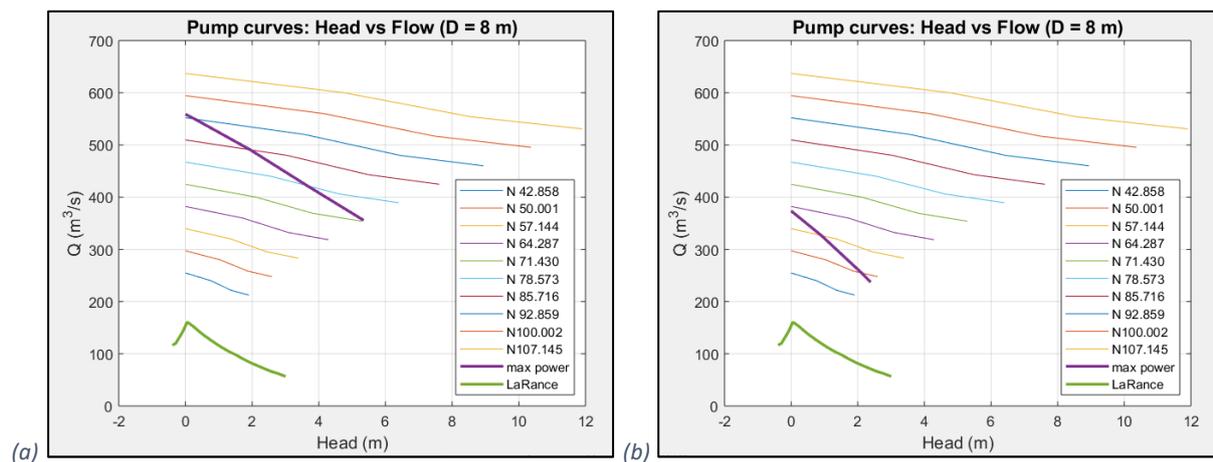
Figure 2.13 Pump performance flow rate curves: (a) 25 MW; (b) 7.5 MW.

The timing of pumping is critical to the performance of the barrage. Pumping was set to start when equilibrium between the tide and the barrage levels was reached, which was always at some time after peak tide. The rate of change of the tide increases up to mean tide, and the head constantly increases regardless of the performance of the pumping. On top of this, the pump flow rate decreases with increasing head, and it is possible that the desired barrage level cannot be achieved before the pump head limit is reached.

The model was in broad agreement with La Rance, i.e. it operated at constant power, and had similar efficiency and net energy gains. Where the pump models differed was at positive head. The pump model used here only operated from zero head and below. At La Rance, the pumps are started at approximately quarter power while there is still a positive head of 0.4m, and the power is ramped up to maximum power at zero head. Operating in this way would increase the flow during the end of the sluice phase and shorten the time to reach equilibrium. This time could be utilised to achieve a higher equilibrium level requiring less pumping thereafter to achieve the desired barrage level. Alternatively, it could achieve a higher barrage level to increase the energy available for the next generating phase.
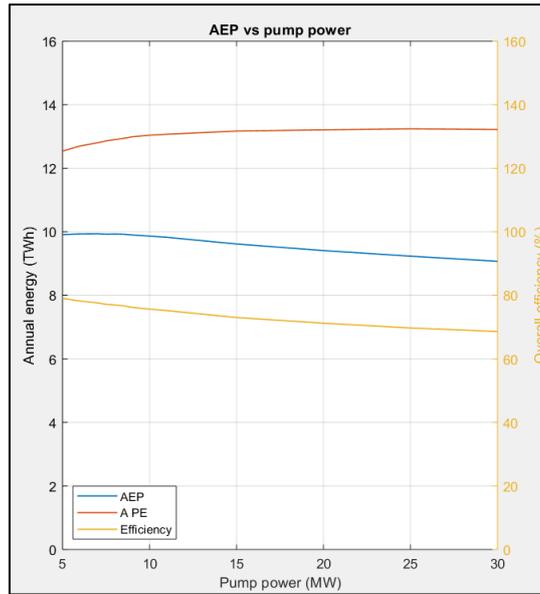
Figure 2.14 AEP and overall barrage operational efficiency for a range of pump power settings. Maximum AEP occurs around 7.5 MW power. APE is the annual potential energy and is calculated assuming no losses.

As a QC of the simulation pumping code, the annual potential energy (APE) value was also calculated to compare with the AEP value. The APE ignored any generating/pumping efficiency losses and simply summed up the potential energy (*mgh*) values for each time step and scaled to a full year. This proved useful when testing the pump model to verify that net energy gains were being achieved. In Figure 2.15 the blue curve was the turbine and pump power and the red curve was the power ignoring the turbine/pump efficiency; calculated from the potential energy change per time step.
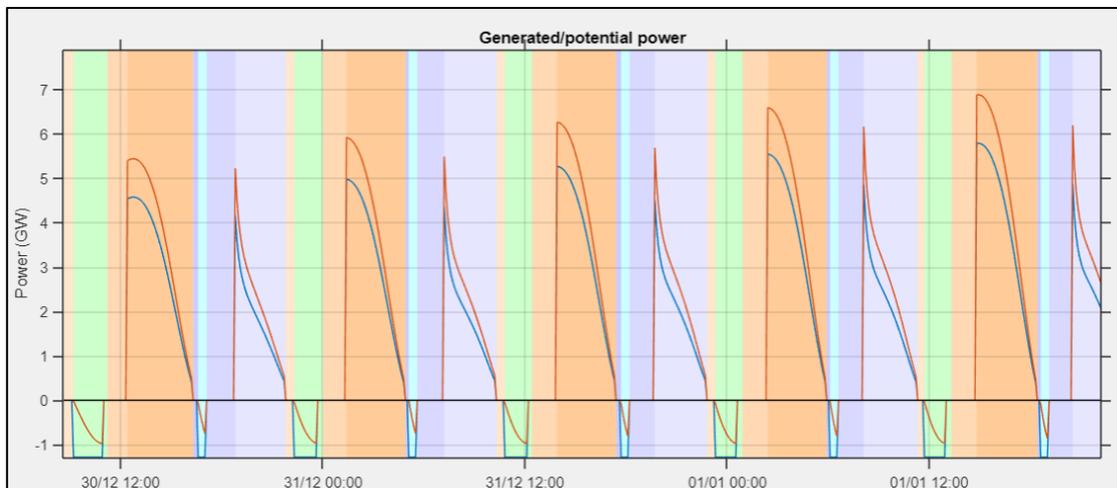
Figure 2.15 Comparison of turbine power using the turbine performance curves and assuming 100% efficiency. This proved to be a good QC of the modelling code and can be used to determine the overall efficiency of the turbine performance.

2.4   Sequenced-based modelling vs broad range scanning

There are two main approaches using the backward difference method for controlling when the transitions between the different phases of the barrage operation occur. These are: sequence-based and time-based. Although the modelling for this study used the sequence-based approach, a comparison of the two serves to justify that decision and illustrate its weaknesses.

The sequence-based approach uses head values to trigger the transitions, i.e. generation starts when a certain head value is reached, transition to sluice occurs when the minimum generating head is

reached, pumping starts at zero head (tide and barrage levels are equal), pumping goes to hold at a predetermined barrage level or when the pump head limit is reached, and the hold state continues until the next generation phase starts. The time-based approach uses set transition times within the tide cycle as the trigger.

In both cases, the triggers are set in advance. The sequence-based approach was preferred because it allowed the program to react to the prevailing conditions and was tolerant to any tidal series input. The actual timing of the transitions within the tidal cycle are not fixed and it is possible for any sub-optimal timing in one cycle to be corrected later on. A pure time-based approach is blind and was set up to expect a particular tidal series. It has been used in the past in optimisation studies and has the advantage of complete flexibility in the sequencing of the operation; for example, overlapping the generating and sluicing phase can be easily configured [9].

When optimisation was employed with the sequence-based approach, how sure are we that the optimum result has been sampled? The time-based approach makes no assumptions about when the transitions should take place and it was a fairly simple task to write a program to scan over all possible cases. Since operational optimisation using sequence-based modelling was a major part of this study, an exercise using a simplified model was constructed and the results of both methods were compared. Separate standalone MATLAB programs were written to find the operation which produced the maximum energy in each case.
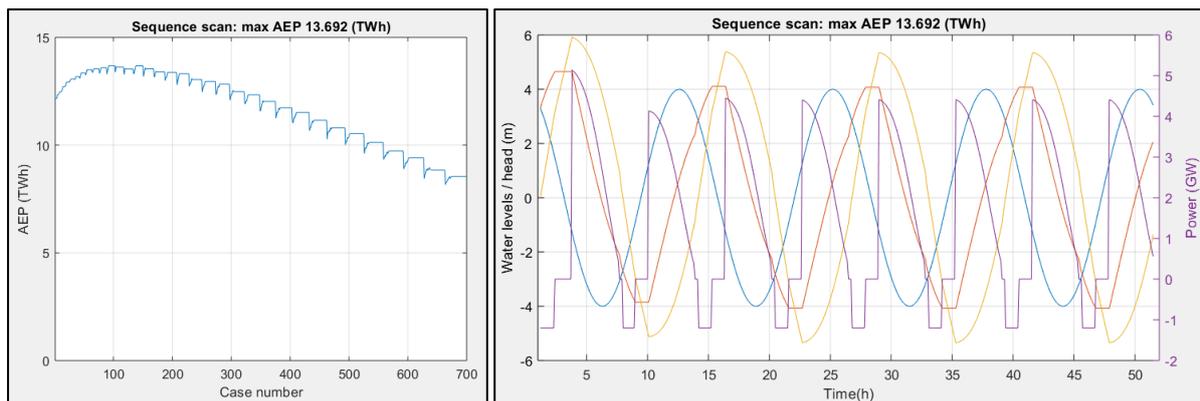
Figure 2.16 Sequenced-based optimum operating cycle at 4m tidal amplitude: AEP per scanned scenario (left), time series of the optimum cycle (right).
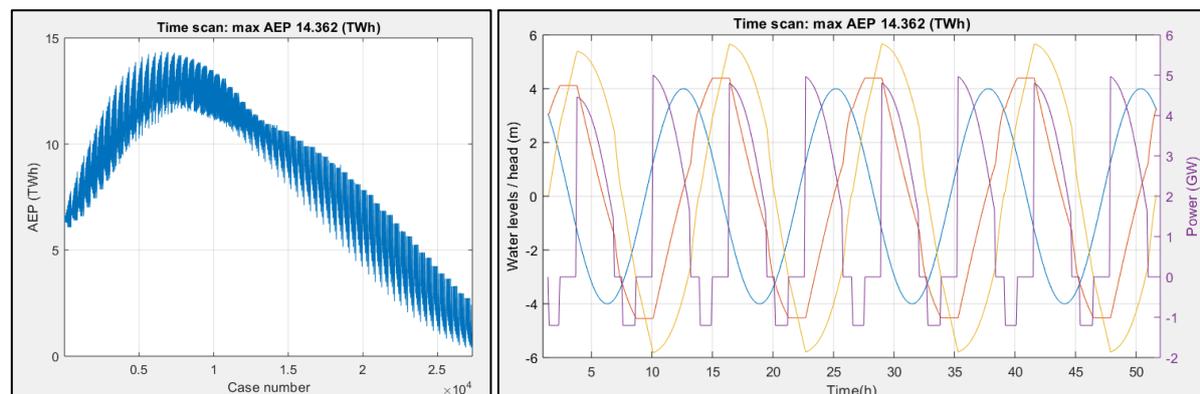
Figure 2.17 Time-based optimum operating cycle at 4m tidal amplitude: AEP per scanned scenario (left), time series of the optimum cycle (right).

The demonstration model was for a 150 km^2 lagoon (symmetric ebb and flood cycles), with 160 turbines and 80 sluice gates, similar in scale to Morecambe Bay barrage. A fixed, 4m tidal amplitude series was input. The turbine generating and pump models were very simplistic and probably over-

efficient. Four (4) complete tide cycles were modelled (the sequence-based approach quickly settles down to a steady repeated pattern) and the energy was calculated over the last full cycle and scaled to AEP.

The sequence-based approach still requires the generating start and pump stop times to be determined and this was handled in the same way as the time-based approach by scanning over a range of times. Time increments of 0.1 hour were used and this led to about 700 cases, which executed in less than a second. The time-based approach used minimum and maximum time window lengths for each operation phase, starting after an initial lag time. This resulted in 27,342 valid cases running at about 5,000 cases per second.

Both approaches gave similar optimum cycles, with the time-based modelling giving slightly more energy (14.4 TWh compared to 13.7 TWh). Closer inspection of the time series plots (Figure 2.16 and Figure 2.17) revealed the time-based approach transitions to the sluice phase well before the minimum generating head limit used in the sequenced approach (NB. this would not necessarily be the case for different tidal ranges). This allows the lagoon levels to more closely approach the tidal limits providing higher head for generating. Overall, this demonstrated that the assumptions made with a sequenced-based approach do not limit its applicability for optimising the transition points to maximise energy output. It also demonstrated that the generating stopping head should be treated as an operational parameter and included as part of the optimisation process.

2.5 Operational parameterisation

In principal the use of a barrage to generate electricity is a relatively simple concept, with the amount of electricity generated dependent on the tidal range, the size of the reservoir and the number and power of the turbines. The reservoir volume is determined by the position of the barrage, and the number and size of turbines control how much of the available energy can be converted to electricity. In practice it is more complicated than this. A major factor is the variation in the tidal range; a barrage at Morecambe Bay must contend with a range that varies from less than 4 m to over 10 m.

As described in section 1.3 above, the generating cycle consists of hold, generating, sluice and pump phases per half tide cycle, requiring strict control over the transition times and operation of the turbines. In the modelling undertaken, the parameters that control the transitions and the turbines were classified as operational parameters. Since successive tide cycles are always different, the operational parameter values need to change each cycle to maintain efficient operation. To achieve this a series of functions were generated relating the operational parameters to the tidal range or tidal amplitude (see section 2.6.2).

The following is a list of the operational parameters included in the modelling, each with separate functions for ebb and flood cycles:
- Generating starting head
- Turbine speed
- Turbine model flow
- Pumping limits

The two turbine parameters are mutually exclusive and the pumping parameter was only used if pumping had been selected.

The turbine speed and flow are controlled for a given head by the inlet guide vane and runner blade angles, and these can be adjusted continually throughout a generating cycle. For this study, control of the turbine has been modelled using two alternative parameters; the speed or the model flow (Q_{11}). In either case, they are held constant during a generating cycle (regardless of how the vane and blade

angles change to achieve it). This was a simplification for the ease of modelling. See section 2.3 for more details.

The pumping limits define the target barrage level that should be achieved during the pumping phase, and the different options require a more detailed explanation (see Figure 2.18).

- Lowest-Highest: all high/low tide limits set at the lowest/highest tide in the tide series
- Cycle-by-Cycle: high/low tide limits set to the individual high/low tides
- From equilibrium: defined from the equilibrium position – determined during run time
- Head: defined relative to the tide height – determined during run time
- Force C-b-C: ensure cycle-by-cycle was achieved before pumping head limit reached

The pumping limits are operational parameters and the values derived from the tidal range functions are added to the limits described above. To keep just these limits, the functions should be set to zero.

Figure 2.18 Design parameter settings panel with list of pumping tide limit options.

The cycle-by-cycle operational parameter values were calculated from the input tide series as part of the initialisation phase of the modelling program (see section 2.7). These were represented by vectors, and since all the parameters were constant during each generating phase, they appeared as a sequence of steps when plotted as a time series (see Figure 2.19). In the example shown here, cycle-by-cycle pumping was activated and the pump functions were set to zero; the limits duly follow the individual tidal amplitudes.
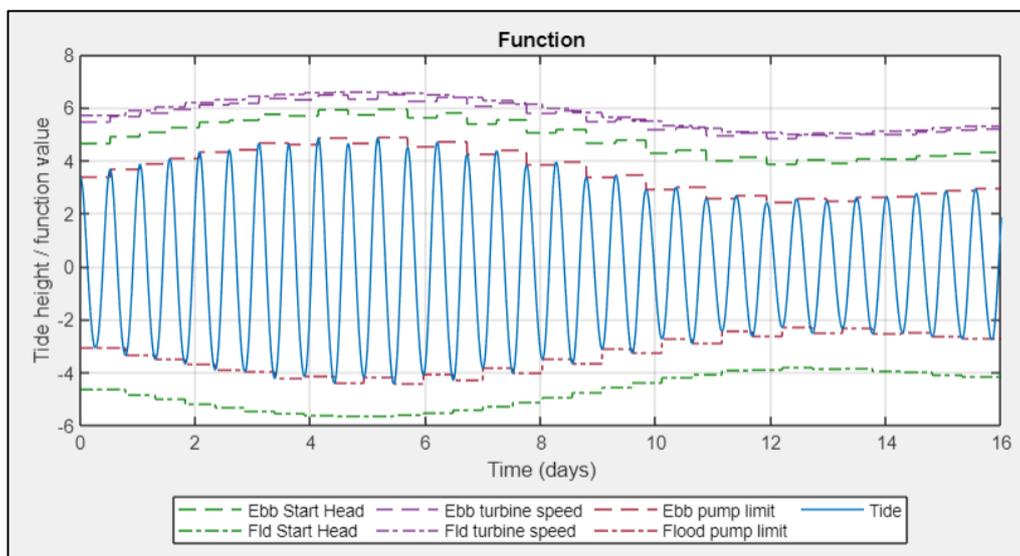
Figure 2.19 Run-time operational parameter values calculated from the function definitions for the input tidal series.

It was useful to make the distinction between modelling to optimise the operation once the barrage was built, and modelling to perform sensitivity analysis on multiple designs to establish the actual design. The number of turbines is governed by the trade-off between cost and generated energy. The efficiency at which the turbines and barrage are operated will have an effect on this trade-off and will ultimately influence the number of turbines chosen for the design. It was important therefore to

achieve a realistic estimate of the turbine performance and barrage operation at the design stage. This was accomplished through a process of optimisation.

For the purposes of this project, the ecological and environmental considerations were encapsulated as a requirement to preserve the tidal range. This could be either at pre sea-level rise levels or at the prevailing levels. The requirement as stated may not address all the ecological and environmental issues. More targeted requirements would require a much broader consultation and possibly separate studies. As an engineering project in an environmentally sensitive area it was deemed prudent to incorporate these considerations at an early stage. One argument for building a barrage was the ability to have some control over the inter-tidal zone and to mitigate the impacts of sea-level rise. This will impose restrictions on the way the barrage can be operated, and one aspect of the analysis was to model the cost in energy this might impose.

During the life of the barrage the operational controls would be continually updated. A rise in sea-level, for example, would necessitate changes to the starting heads (the reservoir increases because of an increase in the wetted area with elevation), and the starting head function would have to be updated to compensate.

2.6   Operational parameter optimisation

As set out in the introduction, the literature describes various methods for optimising the operation of tidal range schemes. In all cases the analysis has been carried out using 0-D modelling, and either a sequenced based or time based approach has been adopted (section 2.4 for a description). The level of sophistication varies in both the complexity of the model (the number of operational parameters) and the method (a simple grid search or some form of optimisation). The method may also be adaptive and allow individual generating cycles to be skipped altogether. In all cases, including here, the solutions are limited to a specific design or subset of designs. If the design or the mode of operation changes, the process should be repeated to determine a new set of operational parameters. Ideally, in order to enable the method to be effectively used as a design and parameter sensitivity analysis tool, the time to complete an optimisation run should be fairly short.

In this study, following the sequenced based approach, two methods for optimising the operational parameters were tested, both of which involved deriving functions to relate the parameters to the tidal amplitude.

From the input tidal series, the sequence of tidal ranges for each half cycle, either peak-to-trough or trough-to-peak, were calculated as part of the program initialisation. From this the half-cycle operational parameter values were then determined using the defined functions (also as part of the initialisation process).

The first approach derived the functions through a grid search process by scanning over the full range of tidal amplitudes. The second, developed in response to the difficulties and limitations of the first, derived intercept and gradient values for a set of linear functions through an optimisation process run on a design-by-design basis. Whenever a new design configuration was defined, a corresponding set of functions were determined.

2.6.1   Scan-based function optimisation

With this method the functions were derived in advance and integrated into the program to enable various design scenarios to be analysed. For this to work, the design parameters being tested had to be included in the scanning and function building process.

The two main design variables that allow a range of values to be entered, are the number of turbines and the number of sluice gates. When these were combined with the tidal amplitude, it resulted in a function of three variables for each operational parameter. The other design parameters remained fixed; dual/ebb only operation, barrage location, sea-level rise, river inflow and pump limits etc. Every design parameter combination would have to be subjected to the same process to complete the functionality of the program.

The functions were derived by scanning over the full range of both the design parameters, the tidal range and the operational parameters. Initially this was tested for ebb only operation without pumping, and required only the generating starting head and turbine triple regulation operating parameters to be included. Figure 2.20a illustrates the principle by reducing the problem to a single design parameter. For a fixed number of turbines and a single tidal amplitude, the range of starting heads and turbine parameter values were scanned by running the simulation once for every combination. The operational parameter values corresponding to the maximum energy point represent the optimum operating condition. This was repeated for a range of tidal amplitudes and a function fitted to the points (Figure 2.20b).
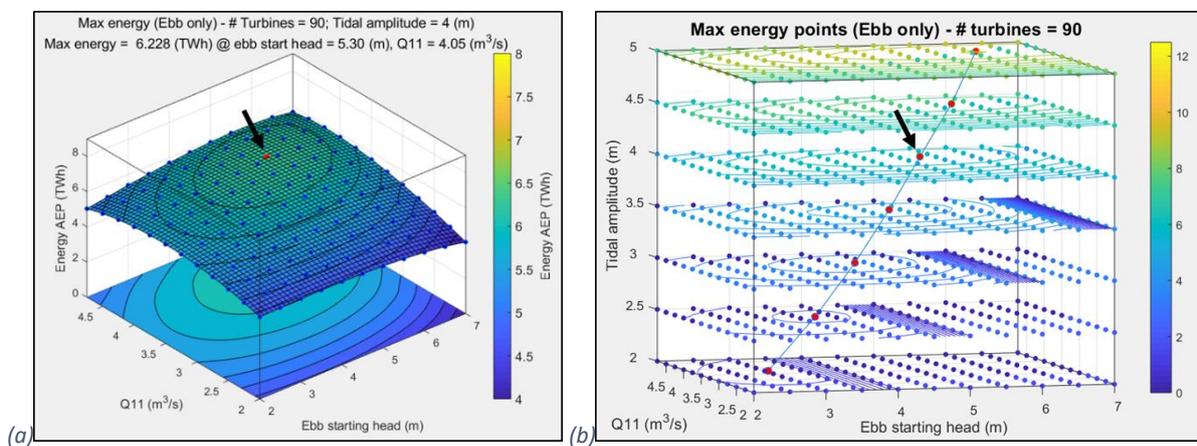
Figure 2.20 Scanning to determine design-based function: (a) maximum AEP over scanned operational parameters for a single design point (fixed number of turbines and sluice gates) and fixed tidal amplitude; (b) fitting a function through the maximum AEP points for a range of tidal amplitudes. The arrows indicate the same data point.

Figure 2.21 shows the extracted operational parameter values corresponding to a maximum AEP data point plotted in colour along the scanned parameter axes in 3-D. A separate volume was generated for each operational parameter. There are 350 data points, each a single point selected from the scanned range of the operational parameters. For the ebb only case above there were 132 scanned operational points per design/tide point. In total there were 46,200 simulation runs with an elapsed time of about 6 hours.

The next step was to fit 3-D functions to the data. The plotted contour slices show that the data are not entirely smooth in places, although the overall trend was reasonably well behaved. The task of fitting the volume functions was time consuming and difficult to QC. A lot of effort was expended to develop a reliable process with little success.

On inspection of the plot of max AEP (Figure 2.22), it was observed the number of sluice gates has a relative minor effect on the AEP compared to the number of turbines. Setting the number of sluice gates to half the number of turbines allowed it to be removed as a design parameter from the analysis. This reduced the problem to 2-D and meant the function could be derived by fitting a surface.
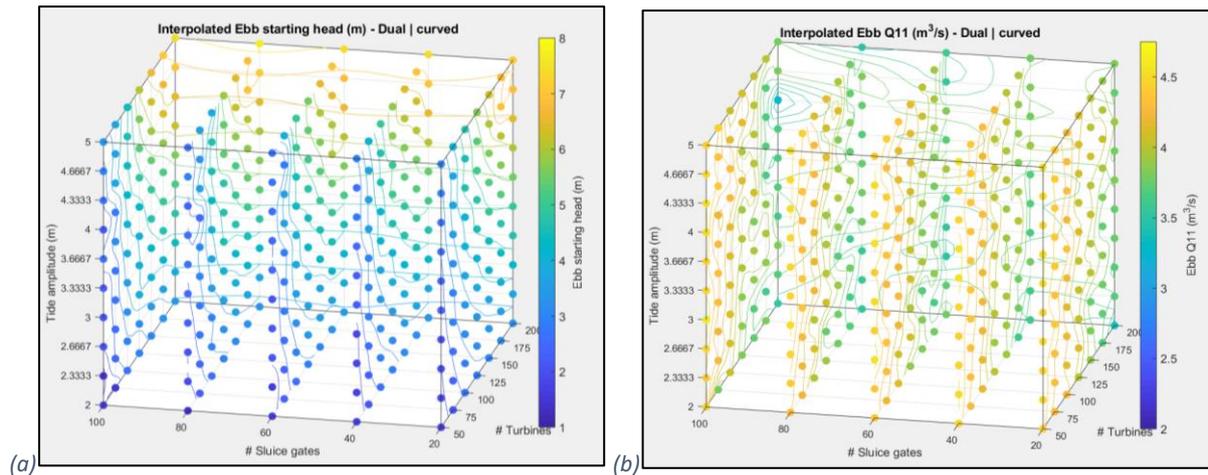
Figure 2.21 Optimum operational (max AEP) points plotted for each scanned design and tidal amplitude point: (a) Ebb strating head; (b) Ebb Q_{11} model turbine flow.
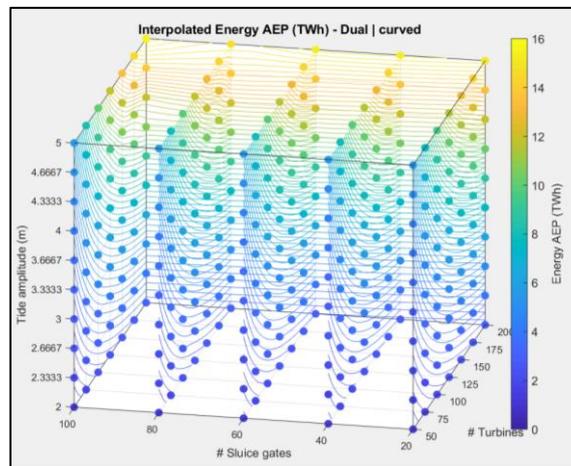
Figure 2.22 Optimum operational (max AEP) points plotted for each scanned design and tidal amplitude point: number of sluice gates has a secondary effect on the AEP.

For dual generation, the number of operational parameters doubles (there are separate starting head and Q_{11} values for ebb and flood cycles). Figure 2.23 shows the fitted surfaces and data points for the four operational parameters and the energy. The Q_{11} data were generally noisier. The energy estimate was the least noisy, and the energy misfit error was used as a means to weight the data for surface fitting of the other parameters, with little benefit.

Even after combining the number of turbines and sluice gates to a single variable, the number of cases to be scanned was significant. In the case above, where six variables were used, if each has just six discrete values this totals 46,656 scans. At half a second per run this takes about 7 hours to execute just to generate the data.

An alternative approach considered was to use a minimisation search engine with the four operational parameters as variables and repeat for each tidal amplitude and number of turbines. This removes the need to scan over the operational parameters. Using the previous result as the input point in each case meant the seed point was reasonably close to the final solution and the local search method was fairly efficient. The cost function was set up to minimise negative AEP. A conditional search engine *fmincon* was used as a first pass to constrain the parameter values within the lower and upper bounds, followed by the more robust *fminsearch* to improve the result further.

27
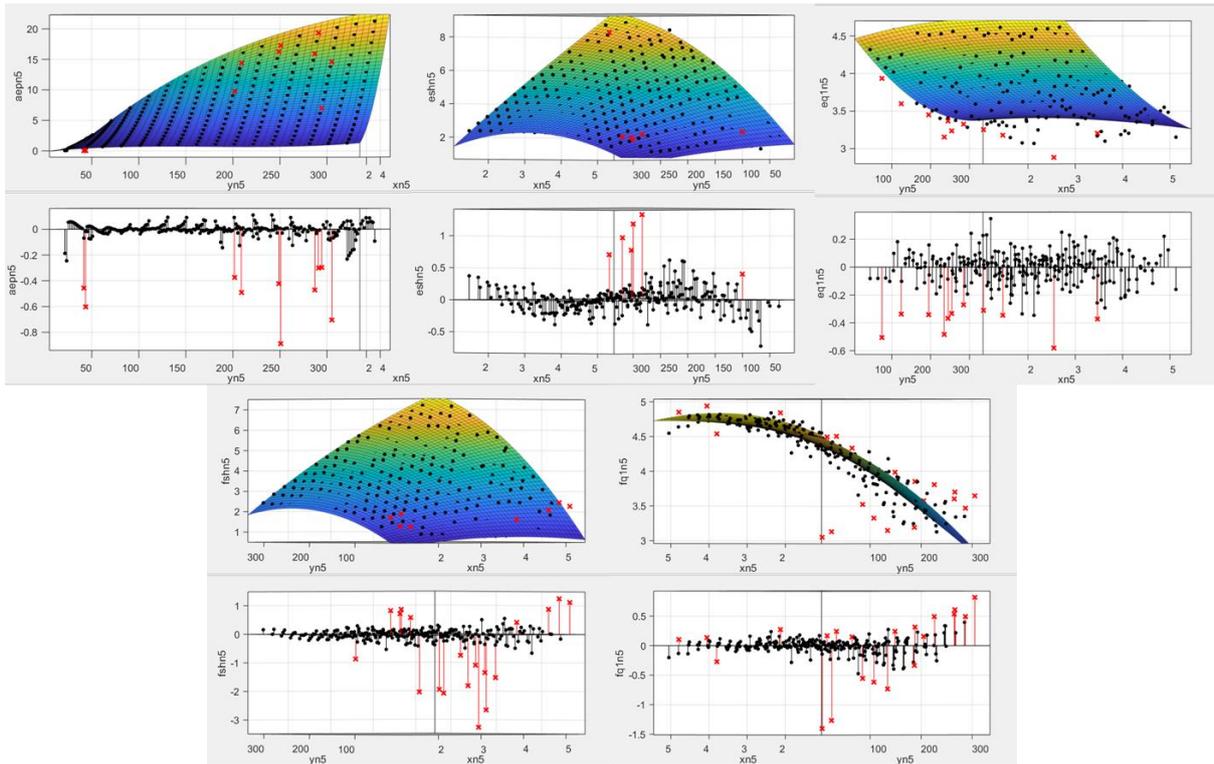
Figure 2.23: Fitted surface functions for the AEP and 4 dual-mode operational parameters after setting the number of sluice gates to half the number of turbines (from top left to bottom right: AEP, ebb turbine speed, ebb model flow, flood turbine speed and flood model flow)

The minimisation struggled where the slope in parameter values were close to zero (Figure 2.24). The design case with 440 turbines illustrates the point. By manually adjusting the operational parameter to the values represented by the blue dots, the calculated AEP rose from 7.9 TWh to 8.1 TWh.

This whole parameter scanning approach was deemed to be unreliable and too time consuming to be used to develop a fit for purpose simulation tool, and an alternative approach was sought.
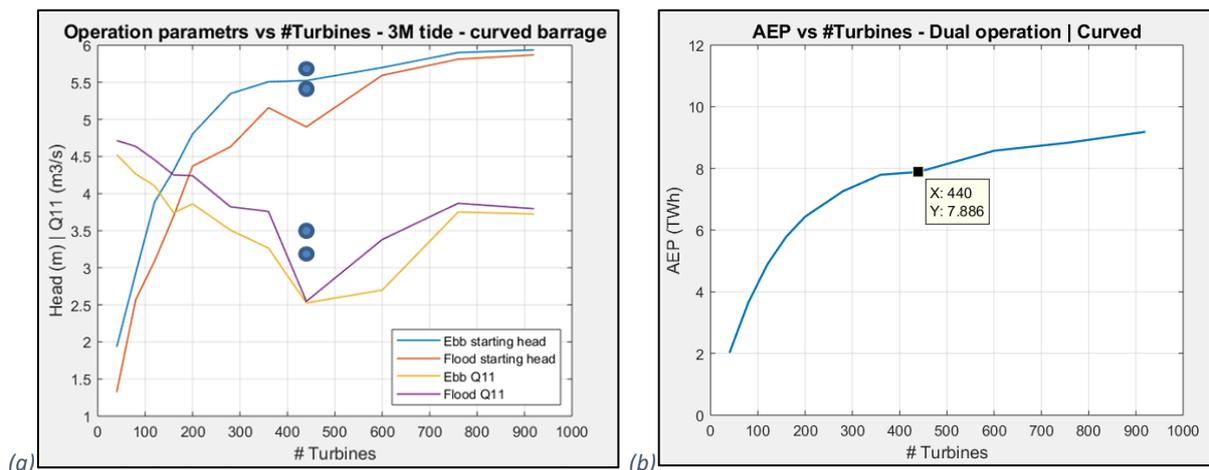
Figure 2.24 The minimisation algorithm struggled to find a reliable solution when the slope of one or more of the parameter valuse were small : (a) the curves are the optimised parameter values and the points represent manually "corrected" values; (b) the optimised AEP values without the correction . Applying the correction increases the AEP from 7.9 to 8.1 TWh.

2.6.2   Design-by-design optimisation

The second approach had to be more accommodating to the different design parameters and environmental variables, i.e. a method which would handle any and all changes in the design scenario.

The number of possible scenarios, even with the limited number of design parameters considered in the model, was still very large.

The approach adopted was to perform some level of optimisation for every design scenario considered on a design-by-design basis. The method also had to accommodate any arbitrary tidal series input. The idea was to relate each operational parameter to the tidal range with a simple linear function and to optimise for the intercept and gradient values. These would have to be solved simultaneously for all the operational parameters considered.

To test the applicability of using a linear function, the sequenced-based (lagoon) modelling code (see section 2.4 above) was used to find the optimum generating starting head for a range of tidal amplitudes and number of turbines. The results are plotted in Figure 2.25, together with least squares linear and quadratic functions fitted to the 160-turbine case.
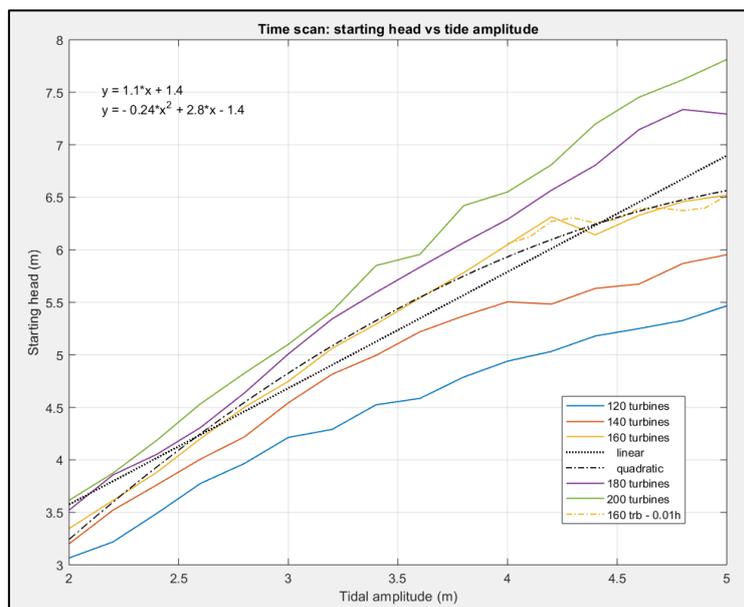
Figure 2.25 Validity of using linear functions to relate the operational parameters to the tidal amplitude.
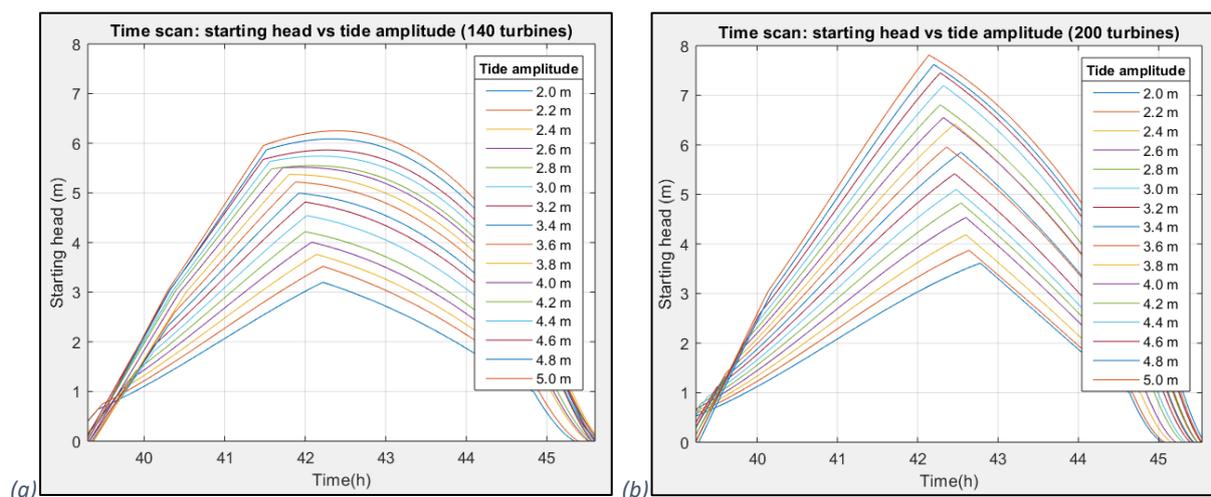
Figure 2.26 Time series plots of the starting heads used in Figure 2.25: (a) 140 turbines; (b) 200 turbines.

Figure 2.26(a) and (b) show the individual optimised head functions for the 140 and 200 turbines cases. The inflection point represents the point of the start of generation, and it was the value of the head at that point that was extracted and plotted in Figure 2.25. For 200 turbines, where the flow capacity

was sufficient to handle the peak tide flow rate, the trend was linear across the whole tidal range. For 140 turbines, the tidal amplitude reaches a point where the starting point has to be brought forward to discharge the reservoir in the time available, and maximum energy was not achieved solely by an increase in the starting head.

Inspection of the linear and quadratic trend lines shows that although the quadratic line was a better fit, the linear trend was reasonable. Increasing the function to a quadratic would increase the number of variables to solve by 50% and would tax the current hardware configuration. Note that all of the modelling was performed using a personal laptop with a 2.5 GHz i7 Intel CPU and 16 GB of RAM.

The optimisation was carried out using a realistic tidal series with the full expected tidal range. As discussed in section 1.3 the energy available roughly quadruples with a doubling of the tide. For a typical input tidal series, the higher tides with the higher energy would disproportionately influence the optimisation result. Note, the curves fitted in Figure 2.25 assume each point had equal weighting. Figure 2.27 shows the peak tide probability distribution (around the mean tide height) and Figure 2.28 the modelled energy per generating cycle for the tidal series over the whole of 2019.
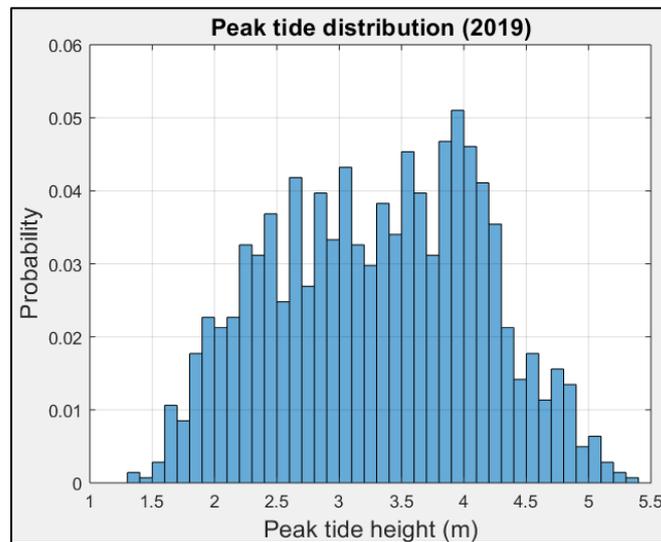
Figure 2.27 Peak tide probability distribution (around mean tide).
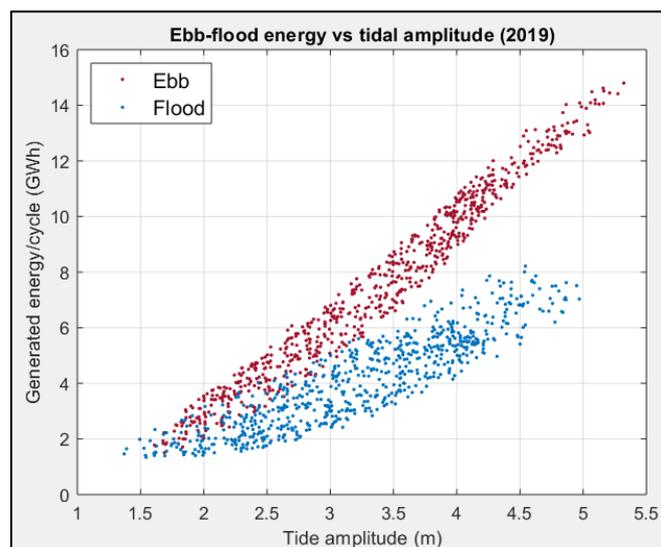
Figure 2.28 Energy per generating cycle. With this approach, the function optimisation could be performed on a fairly short but representative tidal series. The full years modelling displayed in figure 5 was based on the optimisation of a 16-day series.

To establish the magnitude of the effect of modelling with and without optimisation, a comparison in the AEP was made over a range of turbine and sluice gate numbers. The 160 turbines and 80 sluice gates design point was common to both, i.e. the optimised operational parameter values for this design were applied to all the non-optimised designs. Figure 2.29(a) shows the results plotted as surfaces for both cases and Figure 2.29(b) shows the difference. For 100 turbines the difference was approaching 1 TWh in AEP for a value of around 7 TWh; an error of approximately 15%, which increases to 20% for 80 turbines. Differences of this order were approaching those of comparing 0-D and 2-D modelling. This puts into question the validity of those differences when the 2-D design is a copy of the 0-D design without any further optimisation.
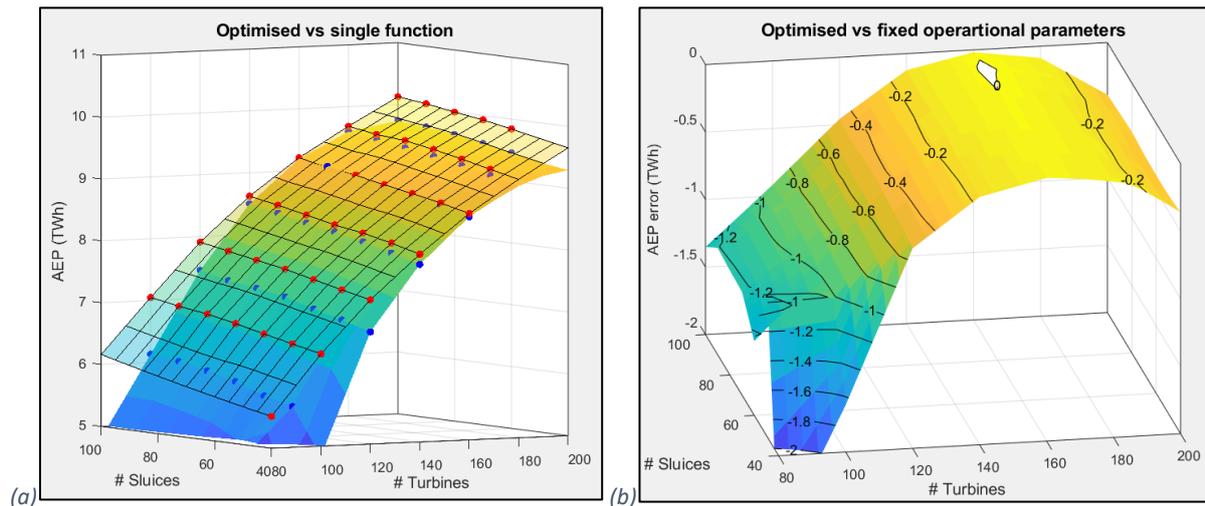
Figure 2.29 AEP values for a range of turbine and sluice gates numbers with and without design-by-design optimisation: (a) results plotted as separate surfaces; (b) difference.

2.7   Program implementation and parameterisation

The 0-D backward difference method was implemented programmatically using a loop structure. The simulation time and time interval were defined by the tidal time series data input to the model. The loop was executed once per sample and the program execution time was proportional to the number of samples. The tide series can be of any length and was typically a representative series forecasts for the area. The tidal series was adjusted to always start at high tide. At the start of the simulation, the operation was set to sluice mode with a small head differential. Although the head was unlikely to be the optimum value, the method was fairly insensitive to the initial condition.

Table 2-1 lists all the parameters used in the model, together with the setup options available. They have also been classified as either design, operational, design and operational, or environmental parameters. The design parameters are deemed fixed when the barrage was built, e.g., the position of the barrage and the number of turbines. The operational parameters control the operation of the barrage. The environmental parameters refer to external factors, and include river inflow, the tidal sequence, surge and sea-level rise. A few parameters are classified as both design and operational. From the design perspective this reflects the way the barrage was expected to be operated. From the operational perspective, the values can be changed at will once the barrage is built. In practice, it is not essential, or might not be desirable, to always operate in a particular way; for example, generating every flood cycle.

Group	Parameter	Options	Design	Operation	Environ	Comment
Design Tab	Reservoir (barrage)	Straight	✓			
		Curved	✓			
		Lagoon: area (km^2)				Constant wetted area
	Number of turbines	Integer value	✓			
	Number of sluice gates	Integer value	✓			
	Operation mode	Ebb only	✓			
		Dual				Demand dependent
	Turbine regulation	Double	✓			Fixed ratio gearbox
		Triple(speed)	✓	✓		Power electronics
		Triple(Q11)				
	Pumping (no/yes)	Yes | No	✓	✓		
	Pumping limit mode	Lowest - Highest	✓	✓		Highest/lowest tide in tide series
		Cycle-by-Cycle				Each high/low tide
		From equilibrium				Height relative to equilibrium point
		Head				Head relative to tide height
		C-b-C force limits				Ensure cycle-by-cycle is achieved
	Sea-level rise	Value (m)			✓	
	Pre Sea-level rise	Check box				Apply pumping limits pre SL rise
	River inflow	Value (m^3/s)			✓	
Operational	Ebb generation starting head	Function		✓		
	Flood generation starting head	Function		✓		
	Ebb turbine speed	Function		✓		Coupled if double regulation -
	Flood turbine speed	Function		✓		Ebb≡Flood; gradient=0
	Ebb turbine Q11	Function		✓		
	Flood turbine Q11	Function		✓		
	Ebb pump stopping head	Function		✓		within operating limit
	Flood pump stopping head	Function		✓		within operating limit
Tide series	Tide model	M2 only			✓	M2 only Heysham model
		Heysham model				10 tidal constituent Heysham model
		M2 user amplitude				M2 only user amplitude
	M2 amplitude	Value (m)			✓	
	Amplitude scale factor	Value			✓	Possible 2-D calibration mechanism
	Sample interval	Value (hours)			✓	
	Start date	Date selector			✓	Adjusted to nearest peak tide
	Duration	Value			✓	
	Duration units	Days | Cycles			✓	
Auxiliary	Turbine diameter	Value (m)	✓			
	Turbine coefficient of discharge	Value	✓			
	Turbine availability	Value (0-1)	✓			
	Turbine phased start	Value (hours)	✓	✓		Generation ramp start duration
	Turbine power rating	Value (MW)	✓			
	Pumping maximum power	Value (MW)	✓	✓		
	Sluice width	Value (m)	✓			
	Sluice gate height	Value (m)	✓			
	Sluice coefficient of discharge	Value	✓			
	Mechanical/transformer efficiency	value (0-1)	✓			Includes all transmission losses
	Maximum n11	Value (rpm)	✓	✓		
	Generation stopping head	Value (m)		✓		Earlier to achieve pumping limit
	Analysis starting head	Value (m)				Initial tide-barrage head (t=0)
	Sea water density	Value (kg/m3)				
	Ordnance-Chart datum shift	Value (m)				Area specific: bathy@OS; tide@CD
	Electricity generation emissions	Value (kgCO$_2$/kWh)				CO2 emissions current supply mix
	Pump starting head	Hard coded - 0m		✓		Starts at equilibrium - 0m head
Tide maps	Map type	Tide cycles				Number annual cycles
		Days exposed/sub				Days exposed/submerged per year
	Data to map	Natural tide				Input tide series
		Barrage				Simulated barrage tide
		Difference				Difference

Table 2-1 List of model parameters.

32

As discussed in section 1.3 the operational phases follow a set sequence and transition to the next phase was triggered when the head or barrage level matched (or exceeded) the relevant operational parameter value. A series of *if* statements determined the active mode for the current execution of the loop. This was followed by a *switch* statement to select the relevant section of code to calculate the flow rate, power and turbine speed etc. The output from the modelling was a set of time series of the various parameters. The energy was calculated by summing over the power time series and was scaled to represent the equivalent Annual Energy Production (AEP) expressed in units of TWh.

Prior to the main loop section of the modelling, during the program initialisation phase, the sequence of tidal ranges for each half cycle, either peak-to-trough or trough-to-peak, were calculated from the input tidal series. From this, the half-cycle operational parameter values were then determined using the defined functions (see sections 2.5 and 2.6).

A key feature of the simulation modelling is the requirement to run through an optimisation step for each new design configuration. This process creates design-specific function coefficients for the operational parameter values. To track this, a feature to store and load these design scenarios was added to the program.

The modelling program and Graphical User Interface (GUI) (see Appendix B), developed as part of this study, involved a significant amount of work and were based on an existing model that was written as part of a previous study at Lancaster University [7, 8, 30]. The modelling code was benchmarked against the PIM [1] and found to give similar results. Although the program is set up for Morecambe Bay, it is a fairly simple operation to modify it for any other tidal range scheme with the appropriate reservoir definition (bathymetry) and tide model

2.8 Tide cycle plots

Two map types were created to provide a means to assess the impact of the current run simulation on the inter-tidal zone. One is a count of the number of tide cycles in a year, the other is the cumulative time exposed/submerged (in days) over a year.

(a) (b)
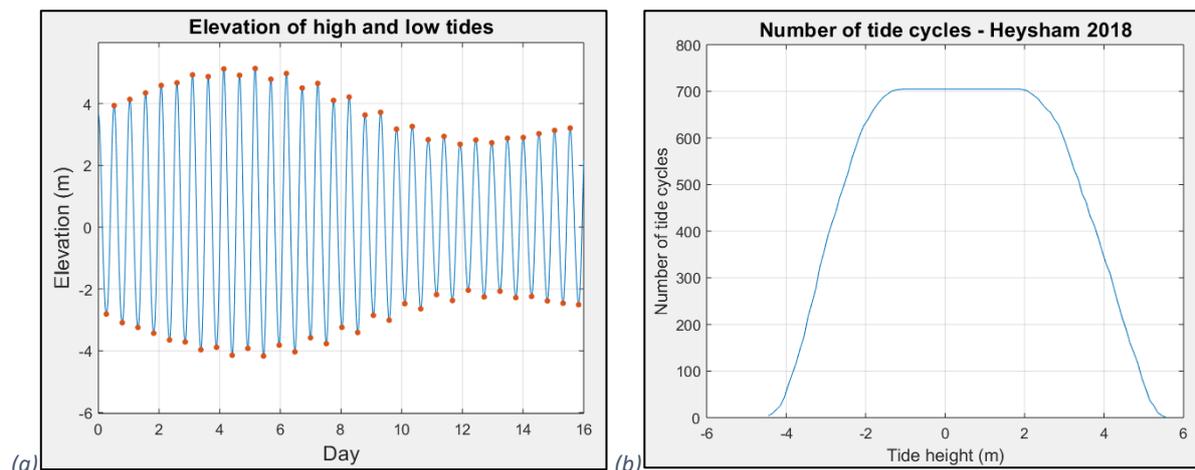
Figure 2.30 The tide cycle plot was a count of the number of tide cycles: (a) typical spring/neap tidal sequence; (b) count of the number of times exposed at low tide mark and the number of times submerged at the high tide mark.

There are nominally 706 tide cycles a year, a little less than two cycles per day. The areas of the Bay where the bathymetry is at mean tide level will experience every tide cycle. Areas at the highest astronomical tide mark will experience very few, if any, cycles in a year; similarly, at the lowest astronomical tide. At high water it is a count of the number of times submerged, and at low water, the number of times exposed. There is a band about the mean tide level which is exposed/submerged

every cycle and this will be represented by the full 706 cycles. Since there is no ambiguity between the high and low water regions on the map, the colour scale was chosen to mirror about the mean tide mark. Areas with a low number of cycles indicate the extent of the spring tides. Figure 2.30(a) shows the typical range of high and low tides and Figure 2.30(b) the cumulative count.

Matching the cycle-by-cycle tides has been applied in fact as part of the study as a proxy for an environmental requirement. Achieving this does not mean the character of the tide will be the same, however. The second map type was designed to provide more insight into the change. Figure 2.31 shows that although the tidal ranges match, significant time was spent at the tidal limits during the hold phase after pumping stops and before generation starts. The second map shows the amount of time in days per year spent submerged at high water and exposed at low water, i.e. it follows the same rationale as the tide cycle plot in mirroring about mean tide.
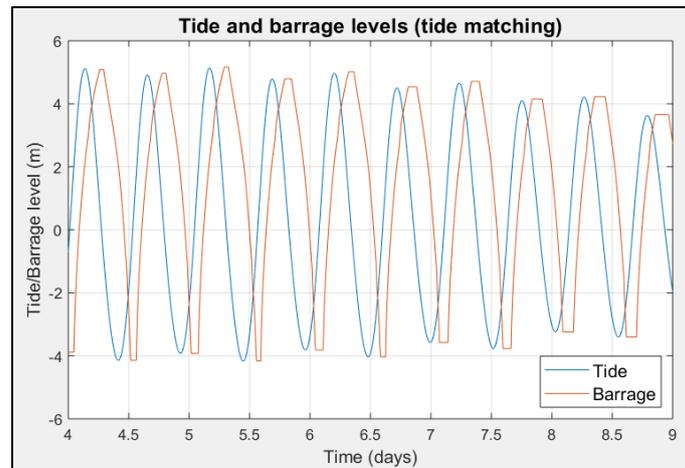
Figure 2.31 Inter-tidal characteristics will be changed by barrage operation. The hold phases subject the high / low tide areas to longer times submerged / exposed.

Three options were created for both map types, these are for the natural tide (no barrage), with the barrage, and the difference between the two.
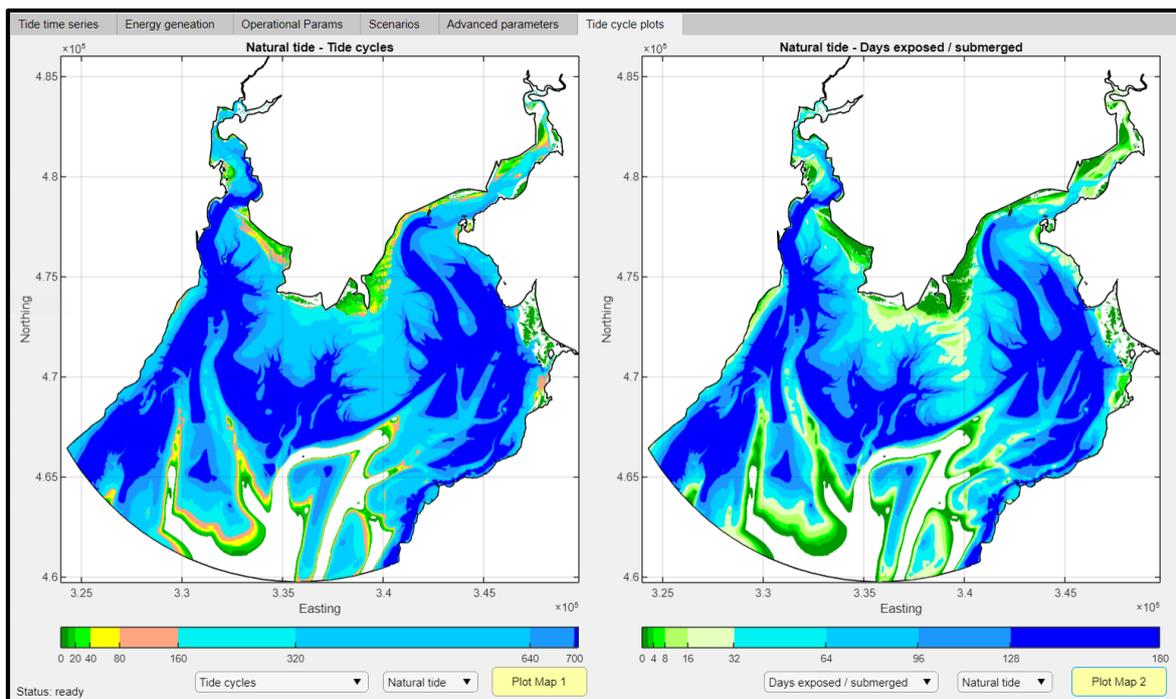
Figure 2.32 Inter-tidal QC plots: (left) natural tide annual cycle count; (right) natural tide days exposed/submerged.

3 Scenario modelling

This chapter details the results from the barrage scenario modelling. The modelling program described in the previous chapter was used for the analysis. There are two sections, the first covers the interaction between the barrage and the environment; the second how sea-level rise, barrage position and the number of turbines determines the amount of energy that can be generated.

The objective of addressing how the barrage might impact the environment, and conversely, how environmental considerations might influence the design and operation of the barrage, was to switch the focus away from simply maximising energy or value generation, and to assess the pros and cons of other factors related to the scheme, for example the role of the barrage in flood prevention, and preserving the inter-tidal zone.

The second section serves to illustrate there are still fundamental design decisions to be made; the position of the barrage greatly affects the size of the available reservoir and the turbine capacity determines the extent to which the reservoir can be exploited. This analysis also acts as a precursor for the energy balancing chapter (section 4).

3.1 Ecology and the environment

The primary driver for the deployment of a tidal barrage is to generate electricity by exploiting an inexhaustible and freely available energy source. Tidal range energy is considered capable of offering both security of electricity supply and to contribute to meeting the UK's decarbonisation targets [31]. For Morecambe Bay, there are two other reasons, as yet uncosted; namely flood prevention and conservation of the inter-tidal zone, both of which are driven by sea-level rise.

Height (m)	Area (km²)	Cum. Area (km²)
<= 4.8m (MHWS)	14	14
4.8 - 5.8m (HAT)	19.6	33.6
5.8 - 6.8m (+1m)	10.5	44.1
6.8 - 7.8m (+2m)	4.9	49
7.8 - 10m	9.8	58.8
No Lidar data	17.7	76.5

MHWS
Mean High Water Spring

HAT
Highest Astronomical Tide
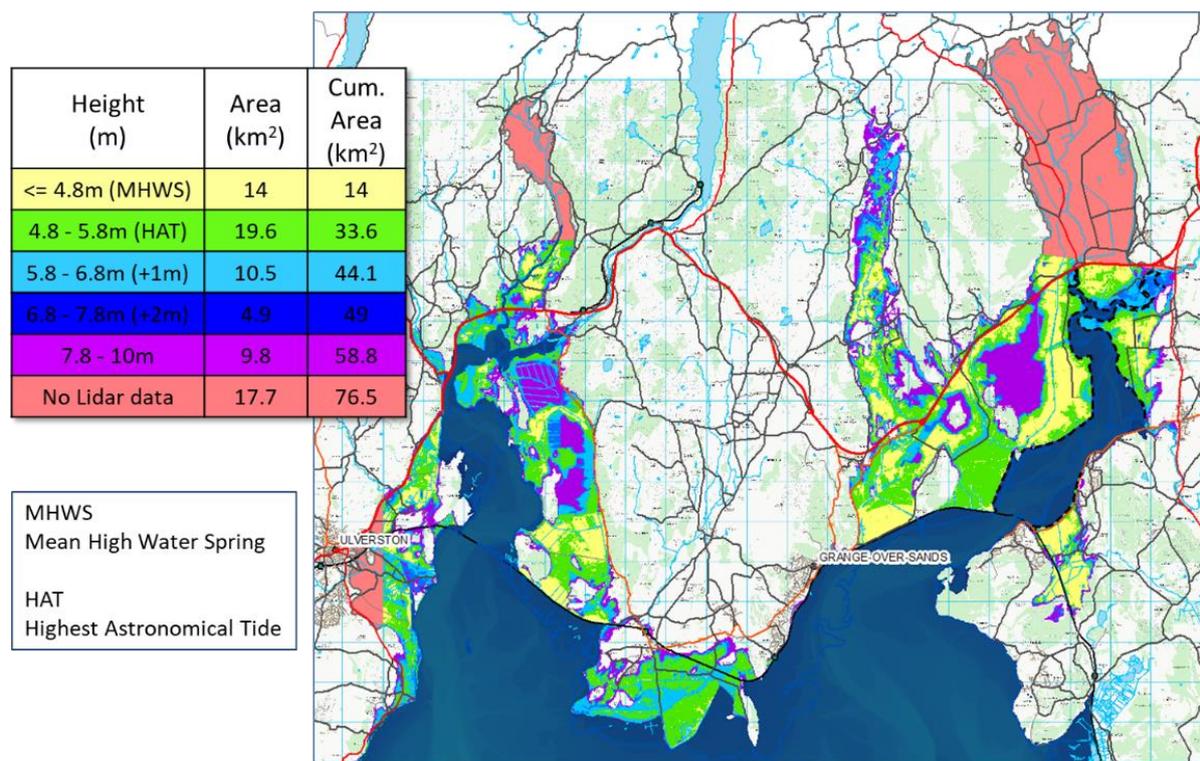
Figure 3.1 Map of the low-lying areas to the north of Morecambe Bay currently protected by sea flood defences or which are at risk of sea and river flooding.

The Bay is currently protected by multiple conservation designations ranging in status from global to local. It may seem contrary that building a barrage, with inevitable disruption to the ecology, may be beneficial overall. The balance of ecological costs and benefits needs close examination. The threats

are uncertain in magnitude since climate change is modelled and forecasts into future centuries. However, sea-level rise is widely accepted and already being observed. As a consequence, the inter-tidal zone will be pushed further inshore, where it will meet man-made resistance, in the form of flood prevention embankments, preventing its natural migration. The mud flats will remain submerged longer as will the salt marsh and other inshore environments forcing a reduction in extent or becoming lost altogether. Saltmarsh acts as an important carbon sink and should be protected where possible.

The presence of a barrage can change the nature of the inter-tidal zone within it, but it can also operate as an environmental management scheme. Importantly, it can limit the height of the high tides to alleviate tidal flooding and mitigate riverine flooding, and it can maintain the current tidal range, thus preserving existing habitats. This criterion, of maintaining the current tidal range, has been applied as part of the study as a proxy for an environmental requirement.

3.1.1   Flood prevention

There is extensive tidal flood prevention infrastructure around sections of Morecambe Bay. Figure 3.1 shows a map with land areas highlighted that are currently protected from tidal inundation, or areas that are increasingly under threat. The coloured areas represent all the land below 10m elevation (OS datum). The yellow and green areas are below the current highest astronomical tide and would be flooded periodically if not for the flood prevention schemes. In adverse conditions of low pressure and high surge then an even greater area is at risk. Figure 3.2, Figure 3.3 and Figure 3.4 illustrate the flat fluvial depositional geomorphology of the valleys at several places around the Bay.

Figure 3.2 Photograph of the Lyth Valley looking south towards the Kent estuary.

Figure 3.3 Photograph of the Levens estuary.

Figure 3.4 Photograph of the upper reaches of the Kent estuary.
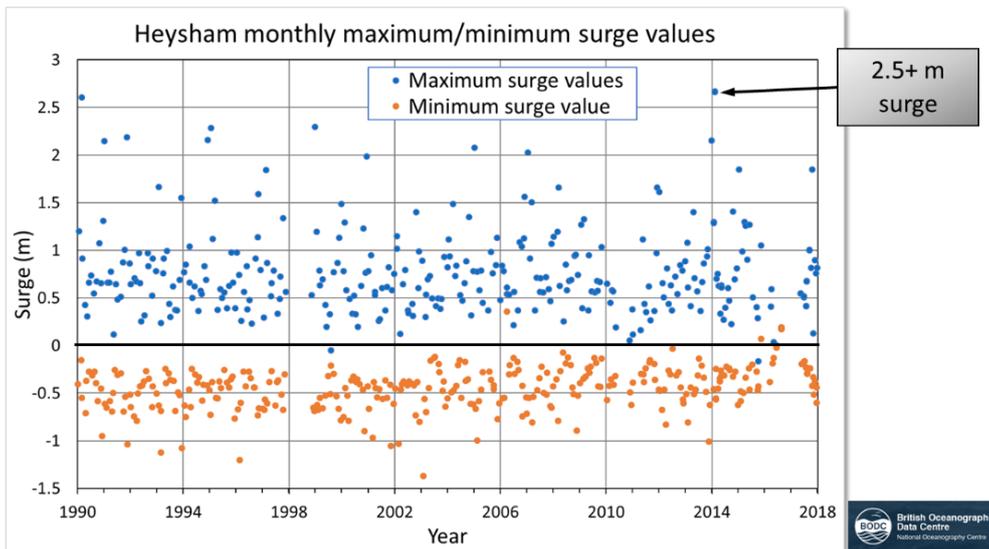
Figure 3.5 Plot of the maximum and minimum surges per month between 1990 and 2018.

Figure 3.5 illustrates how much the tide can be raised by sea surges driven by the prevailing winds. The data was extracted from the difference between the measured and forecast tide levels at Heysham between 1990 and 2018 [23]. A high surge coinciding with a very high spring tide maximises the threat. Add to this the predicted sea-level rise, then in the long term the flood defences will be stretched and probably inadequate, requiring upgrading or areas abandoning.

Figure 3.6 shows the projected sea-level rise using data from the UK Climate Protections 18 [32]. The shaded area represents the range of projections assuming the Representative Concentration Pathway (RCP) 4.5; the dashed green line is the 50th percentile or median prediction. The 50th percentile lines are also shown for a lower emissions pathway (RCP 2.6) in blue and a higher one (RCP 8.5) in red. The IMechE report Rising Seas: [16], using models such as these, recommends ".... prepare for a minimum of a 1 metre rise in sea level this century but plan for three metres of rise."

The design life of the barrage is expected to be 120 years or more, and in that time, sea-level is expected to rise significantly. Just how much is not known, and will depend in part on future greenhouse gas emissions.
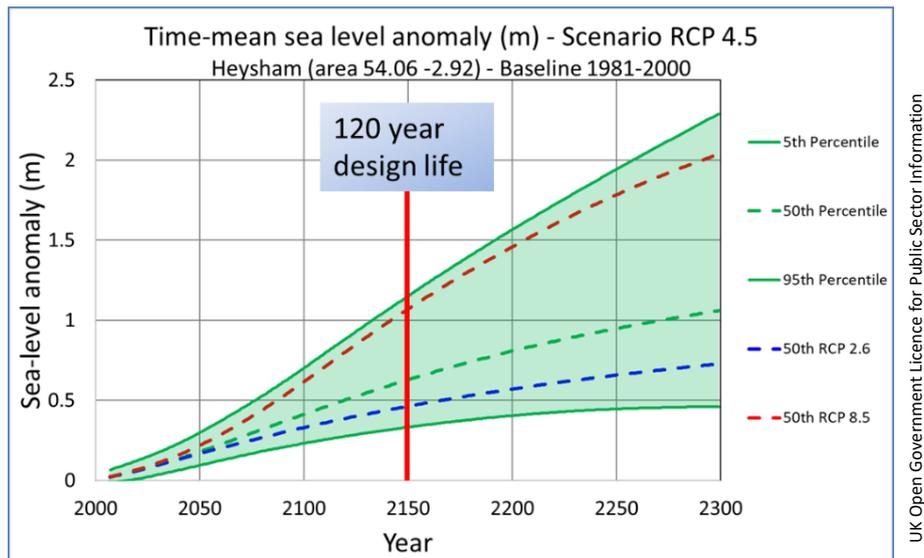
Figure 3.6 Projected sea-level rise for various IPCC RCP scenarios.
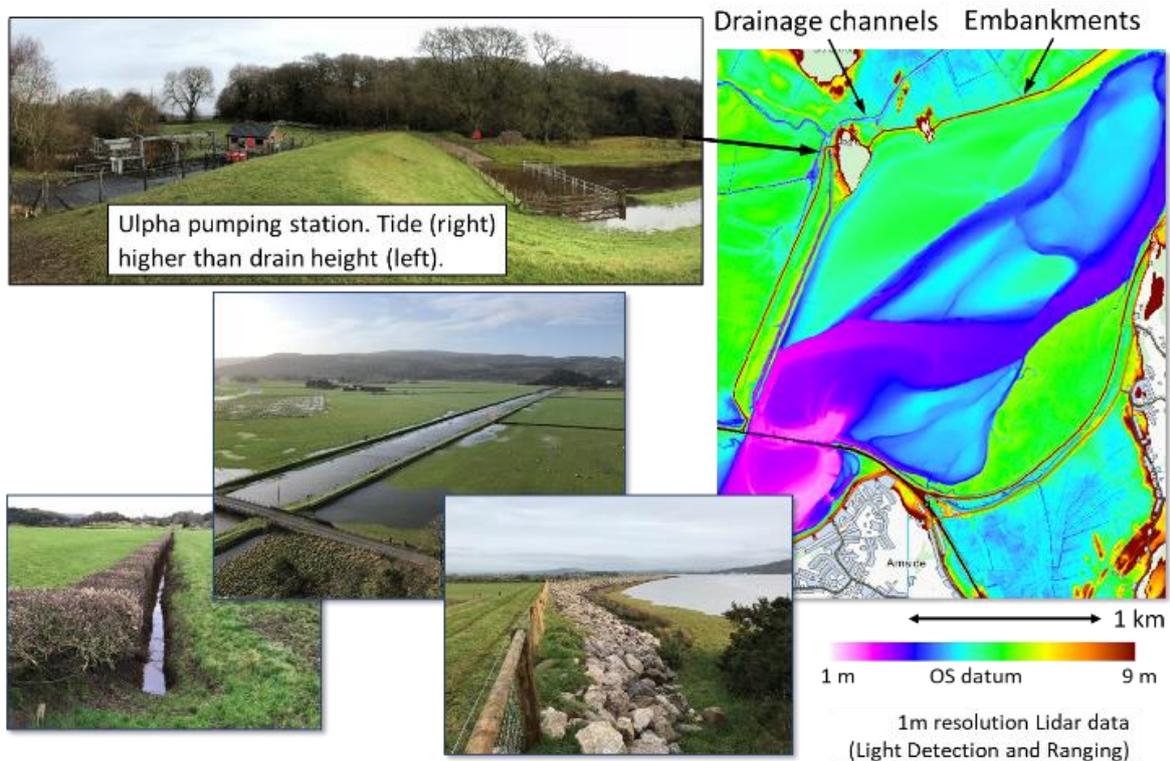
Figure 3.7 Flood defences around Morecambe Bay.

Figure 3.7 show examples of flood prevention features, comprising embankments, drainage ditches and pumping stations. At high tide the sea level is above the surface water level in the drainage channels and the pumping station at Ulpha, Cumbria is used to maintain the flow and help prevent the land from flooding. There are approximately 22km of embankments around the Kent estuary, a similar amount around the Levens estuary, and yet more on the southern side of the Bay south of Silverdale.
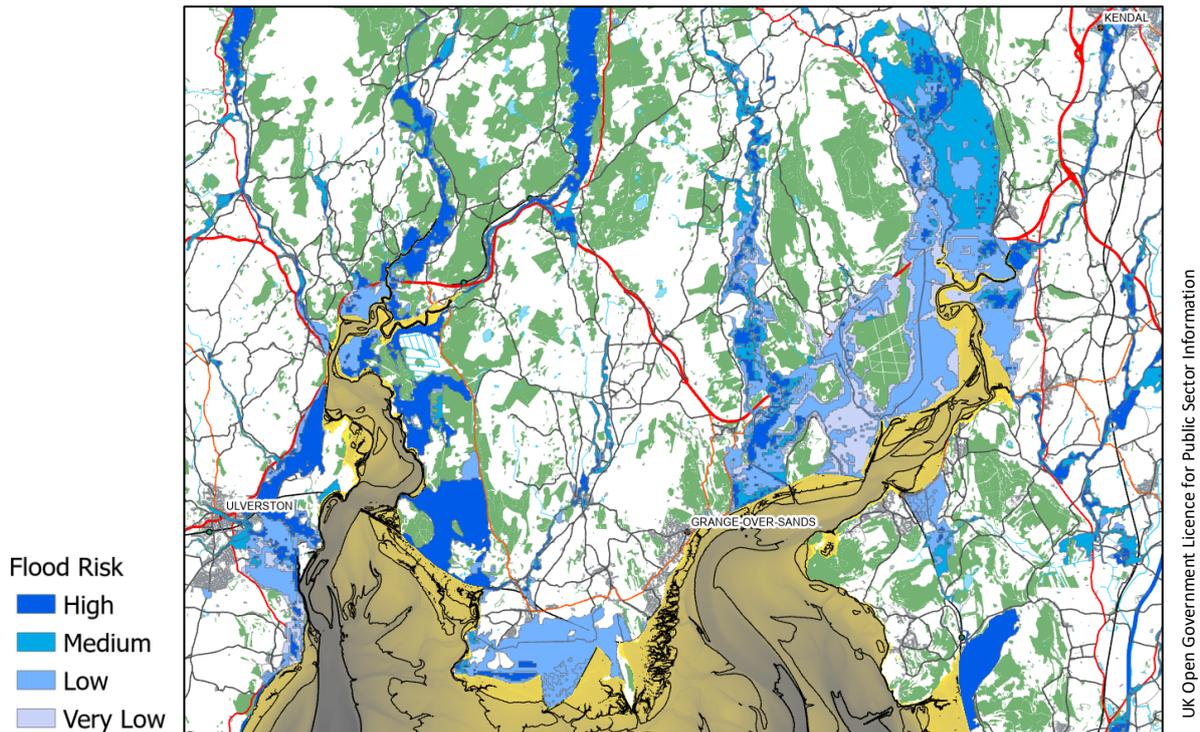
Figure 3.8 Environment Agency flood level risk assessment from rivers and the sea, north of Morecambe Bay.

At present, the main risk of flooding around the Bay is not from the sea, but from rivers. Figure 3.9 shows the flood level risk from rivers and the sea as determined by the Environment Agency [33], taking into account the flood defences. Figure 3.9 shows an extract from the Lyth Valley Flood Investigation Report [17] from the Environment Agency following storm Desmond in 2015. The report states "The volume of floodwater combined with the restriction caused by the high tide, resulted in overtopping of the right bank of the River Kent in the area around Levens Hall and Levens Moss." Although it occurred at high tide, it was also a neap tide (Figure 3.9(b)). Had it occurred a few days earlier at the peak of the spring tide, the sea level would have been nearly 3m higher and the consequences would have been worse.
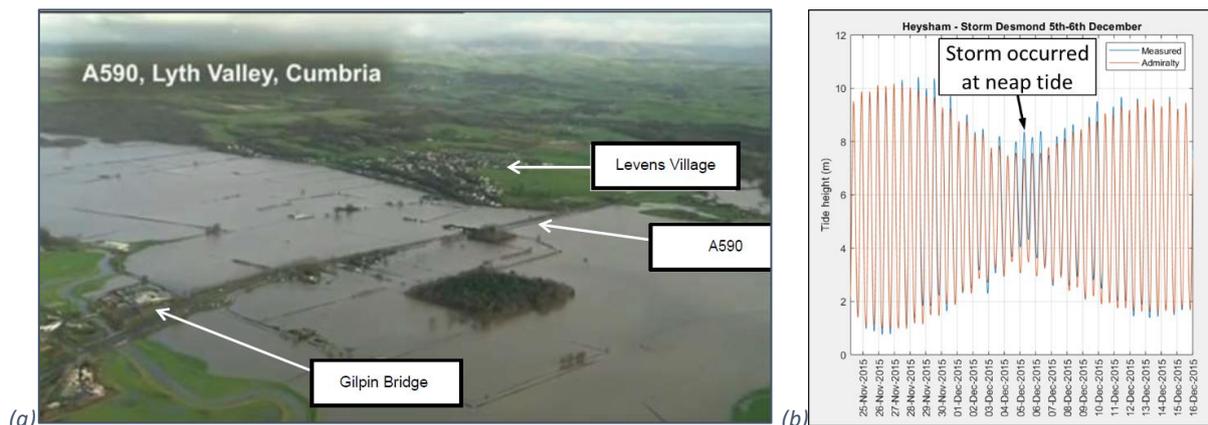
Figure 3.9 Storm Desmond Lyth valley flooding in 2015: (a) photograph from the Environment Agency Investigation Report where high tide was cited as a contributing factor; (b) the storm occurred during a neap to and the situation would have been worse if it had occurred during spring tides.

To get an idea of how significant the river inflows into the Bay are from a flood prevention perspective, the catchment area and measured flow rates were obtained [34], where available, and scaled to represent the whole Bay. Figure 3.10(a) shows the catchment areas around the Bay, with a total area of 1,057 km². Figure 3.10(b) shows the peak river flows measured at the main flow gauge stations,

which covers a catchment area of 617 km^2. Scaling this to represent the whole catchment area, and adding in a precipitation rate of 25 mm/hour over the Bay area of 316 km^2, gives a rough estimate of the inflow to the Bay of 3,000 m^3/s. This is insignificant when compared to the natural tidal flow in and out of the Bay. Figure 3.11 shows the simulated flows across the barrage with and without 3000 m^3/s of river inflow. There is no need to use pumping to ease flooding, it is sufficient to control the barrage levels by adjusting the operational phases.

(a)   (b)
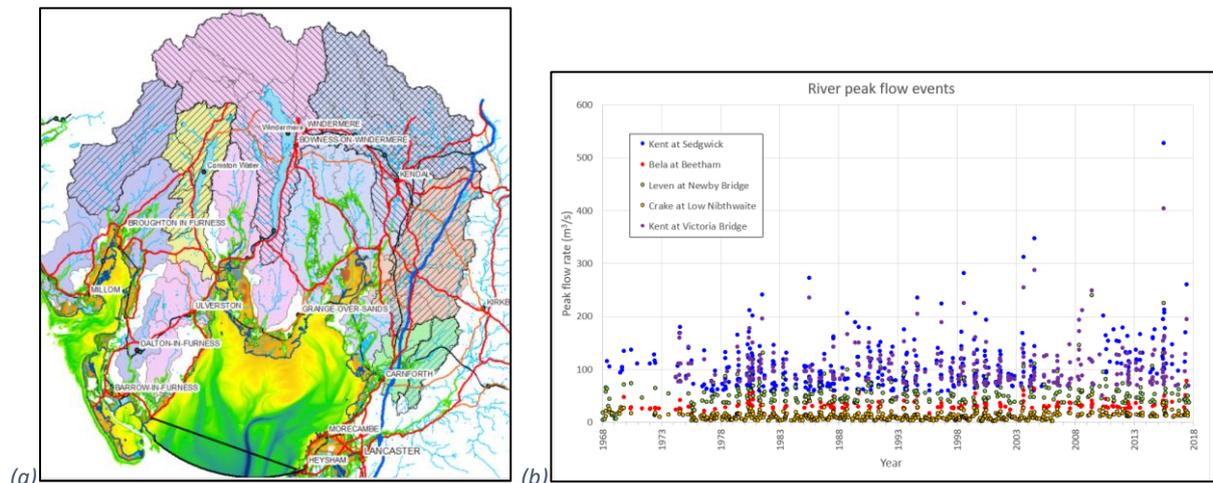
Figure 3.10 River catchment areas and flow rates around Morecambe Bay: (a) total catchment area for the bay is 1,057 km2; (b) measured peak flow = 900 m^3/s for a catchment area = 617 km^2.

3000 m^3/s inflow: AEP = 9.223 TWh   No river inflow: AEP = 9.126 TWh
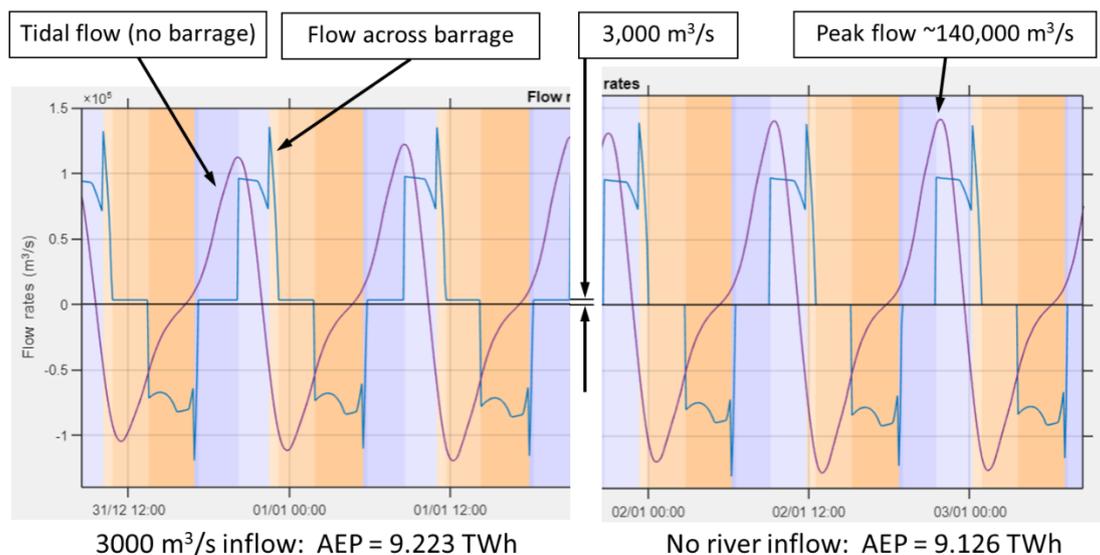
Figure 3.11 River inflows are negligible compared to the bay flows. No need to use pumping to ease flooding, it is sufficient to control the barrage levels by adjusting the operational phases.

3.1.2   Preserving the intertidal zone

The inter-tidal zone is the area exposed and inundated by tides usually defined as between low water and high tide. It is characterised by a progression of unvegetated mud-flats, and rock outcrops through to areas where vascular plants can colonise ending in full salt-marsh intersected by drainage channels at a range of scales. It is a dynamic environment with a tidal bore and meandering channels. Figure 3.12 shows a series of four satellite images from Google Earth taken between 2004 and 2018, and shows the main Kent River drainage channel as it migrates towards the shore at Grange-over-Sands, eroding the salt marsh margin in the process. Figure 3.13 shows the tidal bore; high energy flow and erosion at Grange-over-Sands.

Figure 3.12 Sequence of satellite images near Grange-over-Sands, Cumbria. The area has experienced rapid migration of the drainage channel and erosion of the salt marsh flats in the last few years.

Figure 3.13 Photographs near Grange-over-Sands: the approach of the bore (top left); high energy flow over a rocky section of bed (bottom left); rapid salt marsh erosion (right).

One question is whether a barrage will affect the tidal bore and impact the natural dynamics of the Bay. How important it is for the ecology has not been considered, other than to assume minimal change is preferred. To answer this with any certainty would require more sophisticated modelling than the 0-D used in this study. It is assumed that if the flow rate with a barrage is comparable to the flow rate without a barrage, then the bore would potentially still occur. Figure 3.14 shows the modelled flow rates with and without the barrage represented by the blue and purple curves respectively (positive values represent flow into the Bay). The spike occurs during the sluice phase. The model is for a curved barrage with 160 turbines and 80 sluice gates. It is not at all clear from this if the bore wave would occur or not. Adding more turbines would allow a shorter generation time at a higher flow. This is another example where an environmental requirement might dictate the design.
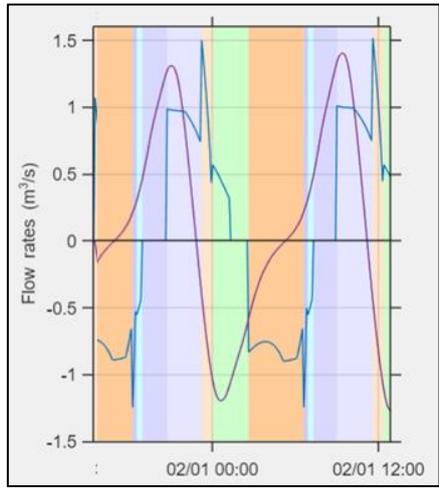
Figure 3.14 Modelled flow rates across the line of the barrage, with (blue) and without (purple) the barrage.

Tide cycles
per year

- 0
- 10
- 20
- 40
- 80
- 160
- 320
- 640
- 700
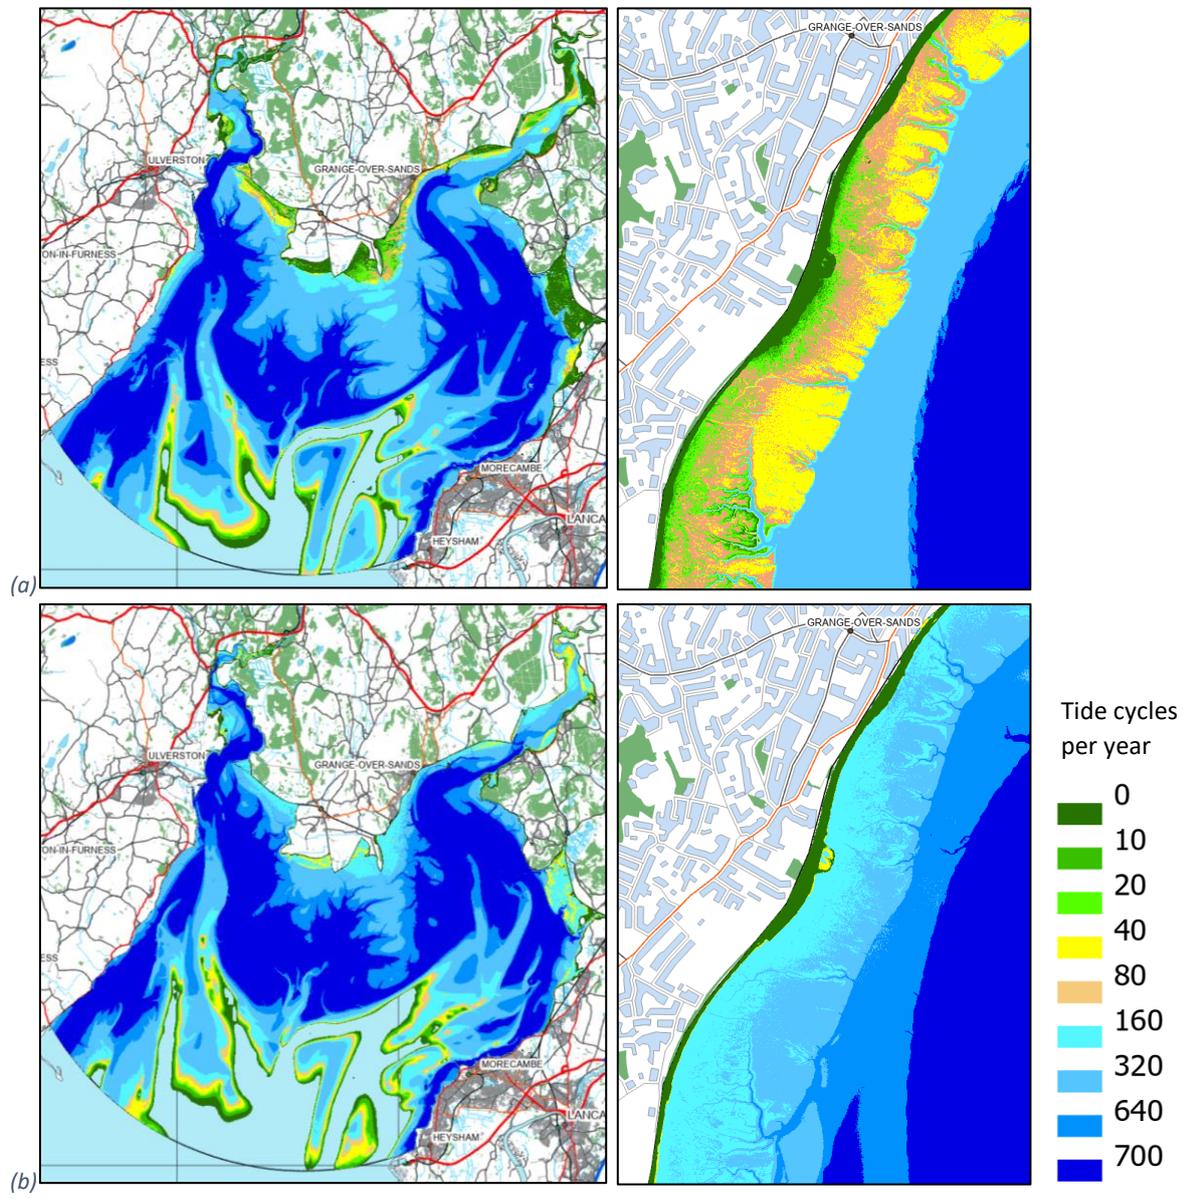
Figure 3.15 Maps of the tide cycle count (no barrage) with detail around Grange-over-Sands: (a) at current sea-levels; (b) with 1.0 m sea-level rise The salt flats would experience an increase of 4 times the number of inundations a year and may be lost altogether: (a) at current sea-levels; (b) with 1.0m sea-level rise.

42

The concept of using minimal change as the environmental and ecological design requirement, was also applied to the tidal characteristics of the inter-tidal zone. The characteristics considered are the tidal range and the time distribution. In reality, these are simple proxies for what is an extremely complex system. As illustrated in Figure 3.14 above, the flow rate is significantly changed. The flow will also be concentrated around the turbine and sluice gate positions along the barrage leading to highly localised flows. This will change the sediment transport behaviour and impose some influence on the channel positions. To gain any understanding of the effects on the inter-tidal zone and sediment transport will require more detailed study.

Two map types were developed to assess the impact of the barrage design on the inter-tidal character (see section 2.8 for a detailed description). The tide cycle count map has been used to assess the impact of a 1m sea-level rise (without a barrage). Figure 3.15 show the maps before and after with details of the salt marsh near Grange-over-Sands. At present the salt marsh is inundated around 80 times per year; with a 1m rise this increases to around 320 times. It is unclear what effect this will have, but it is possible the salt marsh will disappear altogether.

3.1.3 Sea-level rise and tidal matching

The idea of using a barrage to limit the adverse effects of sea-level rise on the environment and ecology needs further investigation. Here the assumption was made that preserving the tidal range at today's levels will achieve this.

Two questions were posed: Can the current tidal range be matched using the barrage to mitigate future sea level rise; and what will this cost in terms of lost electricity generation?

With a rising sea-level, it becomes easier to achieve the high tides as less pumping is required - the point of equilibrium (zero head) rises as the sea-level rises. Conversely, the low tide levels become more difficult to achieve, although this is less of a problem; the relatively small wetted area at low tide means the pump discharge capacity is capable of achieving a greater height change for the same volumetric flow than at high tide.

Figure 3.16 shows that it is possible to match the tide over a complete tide cycle up to a sea-level rise of 1.6m with 160 turbines. This was only achieved after a modification to the program to force an early transition to the sluice phase, in order to bring the equilibrium point earlier and enable the required pump level to be reached before the pumping head limit is exceeded. This was accomplished by working backwards from the point where the tide level falls below the required pump level by an amount equal to the pumping head limit. By running the backward difference in reverse from this point, it was possible to define a boundary that must not be crossed if the pump level is to be reached. Figure 3.17 shows the boundary curves for the ebb and flood cycles, brown and yellow respectively, together with the tide. To ensure the boundary is not crossed, the barrage height is checked every time step during the generating phase, and if it reaches the boundary, the operation transitions to sluicing.

Inspection of the curves in Figure 3.16 shows the hold phase reduces with increasing sea-level rise at high tide, and increases at low tide. This reflects the change in flow rate due to a decrease in the average head during the ebb cycle and an increase for the flood cycle.
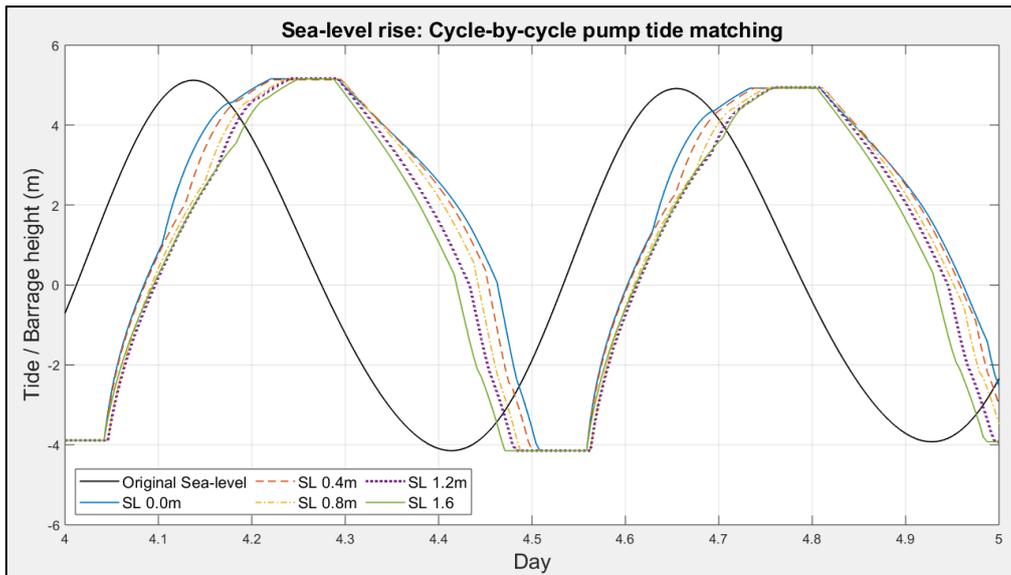
Figure 3.16 Barrage levels at spring tide demonstrating it is possible to achieve pre sea-level rise tide matching.
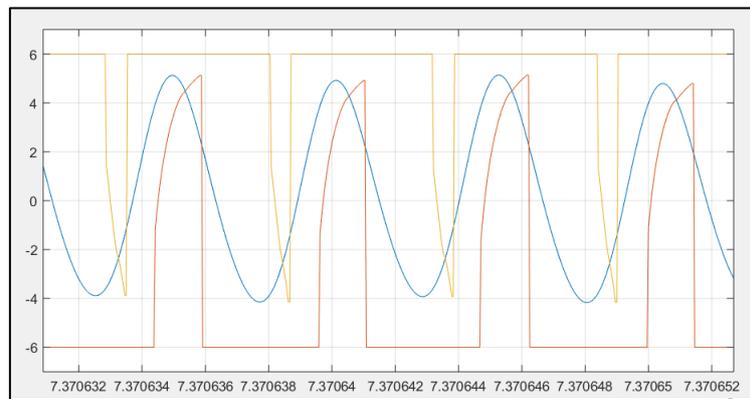
Figure 3.17 Ebb and flood time series defining the regions where it becomes impossible to achieve the required barrage reservoir levels with pumping.

The other question was what is the cost in lost energy if pre-sea-level (pre-SL) rise tidal limits are to be preserved? An answer was obtained by comparing the modelled AEP values when matching the pre-SL rise tidal limits and when matching the prevailing tidal limits. Figure 3.18 shows the results. There are four contributing factors at work:

1. The reservoir volume increases with sea-level rise and there is more energy available (the theoretical maximum energy rises by 25%, see Figure 3.18(b).
2. It is easier to match the lower pre-SL rise high tide levels requiring less pumping.
3. It is harder to match the post-SL rise levels with the same number of turbines requiring the ebb generating phase to terminate early.
4. The head available for the ebb-generating phase is reduced when limited to pre-SL rise levels. Note, although some of the above works in reverse at low tide and for the flood generating cycle, the difference in wetted area between high and low tides means the net effect is dominated by the effects at the high tide levels.

Exactly how the factors interrelate is not entirely clear. There is a marginal if not insignificant cost in lost energy of nearly 5%. This is relatively small when compared against the magnitude of the increase in the available energy and the drop in overall efficiency from 60% to less the 50%. This energy can be

captured, however, with the addition of more turbines. It raises the prospect of designing for the future, either during the initial build or providing an easy means to expand at a later time.

(a)

(b)
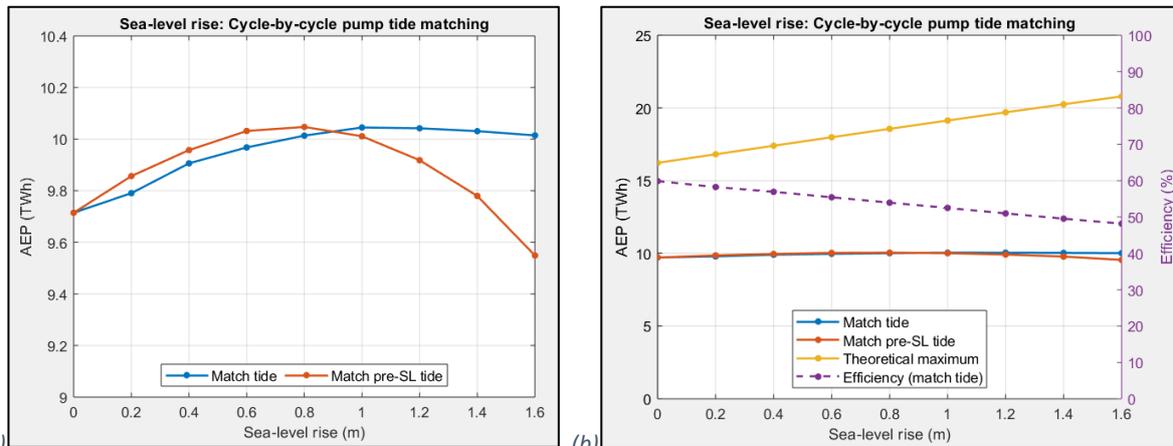
Figure 3.18 Generated AEP with sea-level rise: (a) match to the prevailing tide-limits (blue), current tide limits (red); (b) with the addition of the maximum theoretical energy and over efficiency (matching prevailing levels).

3.2   Barrage design

3.2.1   Number of turbines and sluice gates

The generating discharge capacity is a fundamental parameter affecting the operation and performance of the barrage, and is determined principally by the number and size of the turbines. For this study the turbine size was fixed at 8m diameter, in line with the PIM, and the number of sluice gates set equal to half the number of turbines. This was illustrated by modelling the AEP for a range of turbine numbers and fixed tidal amplitudes (Figure 3.19). Also calculated and plotted is the maximum theoretical AEP for each tide, which remains constant regardless of the number of turbines since discharge is assumed to take place instantly.
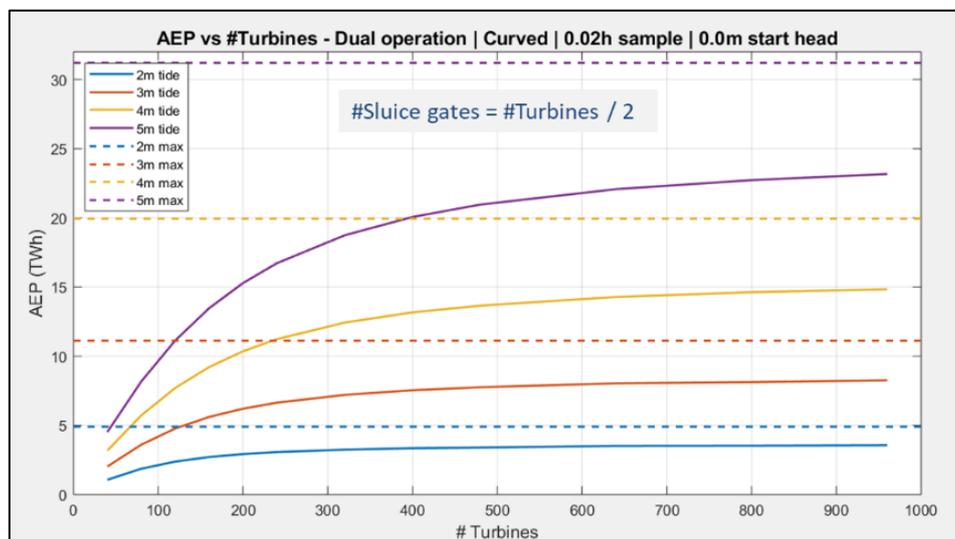
Figure 3.19 AEP vs number of turbines – for tidal amplitudes of 2m, 3m, 4m and 5m. The corresponding maximum theoretical AEP values are also shown. At high turbine numbers the curves level off and become asymptotic to a maximum efficiency line.

Evident from the plot is that with too few turbines the AEP is limited, and at the other extreme, adding too many turbines has no appreciable effect. The trend becomes asymptotic at a level below the maximum theoretical value dictated by the operational efficiency. For higher tides where there is increased flow, more turbines are required before the trend levels off. Although, higher tides demand more turbines their full capacity would only be utilised during spring tide cycles.
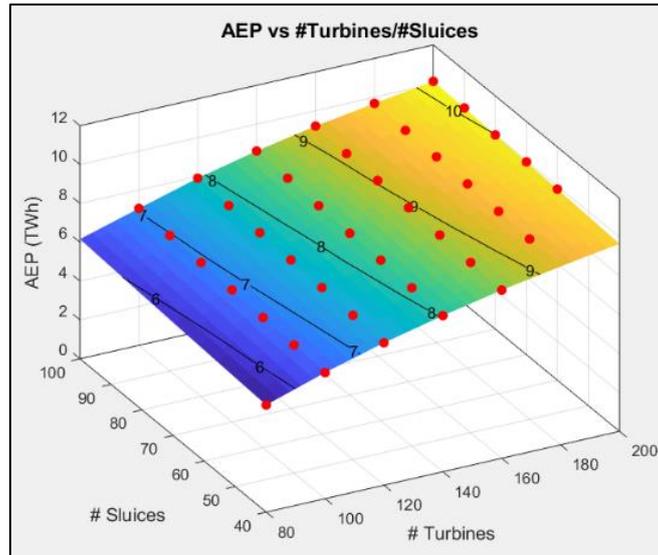
Figure 3.20 AEP for varying numbers of turbines and sluice gates with a realistic tide series.

Figure 3.20 shows the analysis for a realistic tide series at Morecambe Bay and hence reflects the whole tidal amplitude range. It is modelled over a more realistic number of turbines, and also treats the number of sluice gates independently from the number of turbines. It roughly follows the 4m tidal amplitude curve from Figure 3.19, and demonstrates that the number of sluice gates has a secondary effect compared to the number of turbines. This is in part due to the high open flow discharge capacity of the turbines when they are used during the sluice phase.

Ultimately, the number of both turbines and sluice gates that appear in the final design will be dictated heavily by the cost and by constraints on their position along the barrage. The highly localised flows will be constrained to the deeper water areas along the main channels and away from the shore.

3.2.2   Barrage route / reservoir volume

The impact of the reservoir size is illustrated following a similar approach to the number of turbines analysis above. The reservoir size is determined by the bathymetry, position of the barrage and the tidal range. Short of whole scale dredging of the Bay it was assumed the bathymetry is fixed (although this may change over time with altered current patterns). The geographic location of the barrage is a key design parameter and is one of the variables in the modelling. For this scenario the tidal range is fixed, although it is expected to change as a result of sea-level rise. Figure 3.21 shows the magnitude of the possible uplift in AEP for two viable candidates for the position, together with a possible increase in sea level of 1.2 m. Extending the area by building a curved barrage rather than a straight barrage (from the same endpoints) gives an increase in theoretical AEP of ~29%; and a sea-level rise of 1.2m provides a further increase in theoretical AEP of ~24%.

These numbers are significant and feed into the number of turbines required. To what extent should predicted sea-level rise be incorporated into the initial design, i.e. should more turbine capacity be built in from the outset? At the sort of numbers envisaged (less than 200), it is far short of the asymptotic part of the curve, and more turbines will generate more energy. Alternatively, the design can incorporate future expansion by building empty turbine caissons at some upfront cost. It is quite possible these would not be required before upgrades to the existing turbines are needed. Designing for the future runs the risk of it being completely obsolete or forcing unnecessary restrictions on future solutions. Would the potential money saved in the future justify the upfront costs, given the risk it might prove futile?
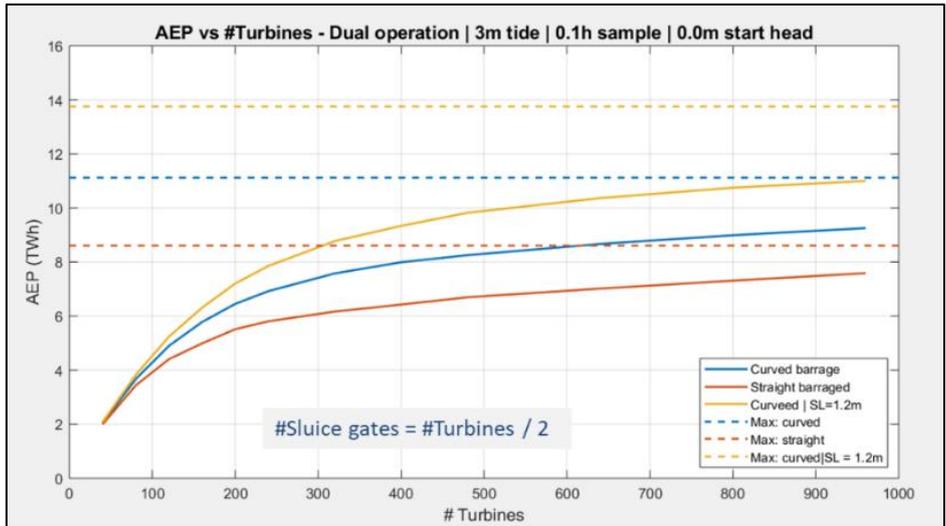
Figure 3.21 AEP vs number of turbines for the straight barrage, curved barrage and curved barrage with 1.2m sea-level rise (all at 3m tide), with corresponding maximum theoretical AEP values.

One approach in deciding on the number of turbines is to look at the cost benefit, i.e. how much energy is gained by adding more turbines. The slope of the curve on the above plot can determine the increase in energy per additional turbine. Where the slope is high the cost benefit is high. With increasing numbers of turbines, the slope and cost benefit reduce to a point where it becomes uneconomic to add more turbines. The actual value of the slope will be determined by the value of the energy and the cost of building and running the turbines. Figure 3.22 shows the detail for the two barrage configurations and illustrates the trade-off between the number of turbines and the generated energy. A larger barrage will generate more energy even without increasing the number of turbines. To exploit the larger barrage on the cost benefit basis requires a proportional increase in the number of turbines.
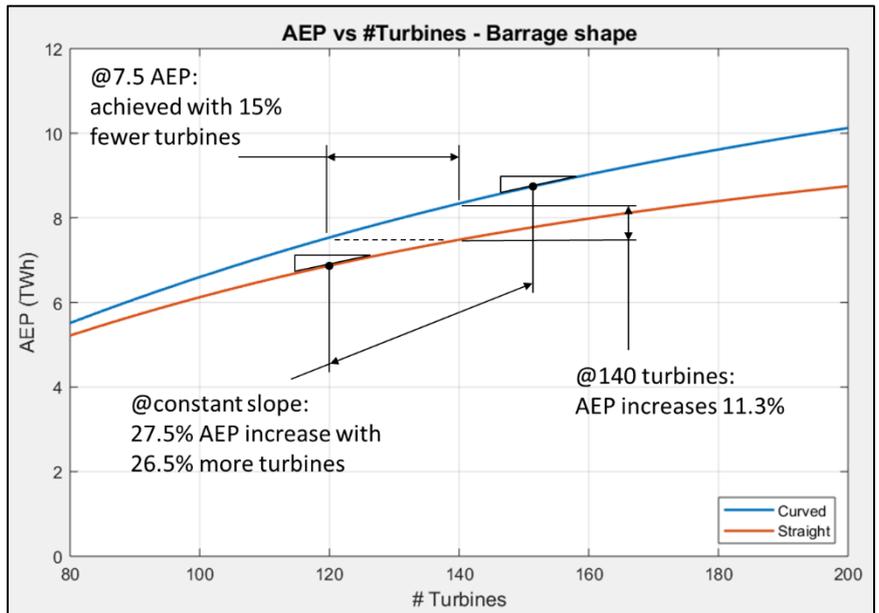
Figure 3.22. Comparison of the AEP between the straight and curved barrages for different choices in the number of turbines.

47

4   Energy balancing

Although tidal range energy is both predictable and reliable, it is also cyclically intermittent, and the delivery of supply does not always match the timing of demand. The original aim of the second part of the study was to investigate and model energy storage options with the objective of increasing the value and usefulness of the energy.

This approach posed several questions:
- How much electricity can the barrage generate and what are its temporal characteristics?
- What is the current electricity demand and how might it change in the future?
- How can the generated electricity be valued to inform generation and storage options?
- How will the generated electricity be handled by the electricity grid and what is needed to facilitate it?
- What would the functional requirements be for a scheme capable of storing all the excess electricity until it is required?

Power consumption varies dramatically on daily, weekly and seasonal cycles making design, operational strategies and value uncertain. This is compounded by electricity demand and usage profiles being expected to change over future decades [15] – but with even less certainty. This uncertainty renders making design decisions tailored to predicted future requirements fraught with risk.

Our study focuses on analysing the supply of electricity to the grid under current conditions, and only briefly looks at future demand to highlight the problem of making long-term decisions in a rapidly changing environment.

For all the analysis in this section, the barrage design selected was the curved barrage, with 160 turbines, 80 sluice gates, dual generating mode, cycle-by-cycle pump tide matching, triple turbine regulation at constant speed, and calculated over a 16 day representative tide series. This is a model design used extensively throughout the study and was derived from the PIM design, upgraded to reflect the increased reservoir volume of the curved barrage.

4.1   UK electricity supply and demand

The demand for electricity can be reliably determined from the supply since the national grid is tasked with maintaining a net balance between the two over a 24-hour period. Figure 4.1 shows the supply by generation type for 2019 (the sequence of generating source, Nuclear through to Combined Cycle Gas Turbine (CCGT), is bottom to top on the plot) [35]. The data is sampled every 5 minutes and there is a 24-hour moving average applied. The regular high frequency peaks reflect the weekly variation between weekdays and weekends. At this scale, the standout features are the annual cycle with higher demand in winter, the fairly consistent supply of nuclear and biomass, the large and irregular variation in wind power, the seasonal variation in solar and the almost complete elimination of coal during the summer.

At a more detailed level (Figure 4.2 shows 3 days from the same dataset without smoothing), power held in pumped storage (pale blue) can be seen to be used at times of peak power demand and most significantly, CCGT (top) provides by far the majority of power balancing, i.e. it is rises and falls to balance the difference in supply and demand and accommodates the large fluctuations in wind (purple) and solar (green).

Figure 4.3 shows the monthly average daily demand profiles based on supply data over the six years from 2011 to 2016. The daily trend is fairly consistent from month to month with an overall increase

from summer to winter of between 25-35%. There is a wider variation during the day compared to the night, and the early evening peak is greatly reduced during the summer.
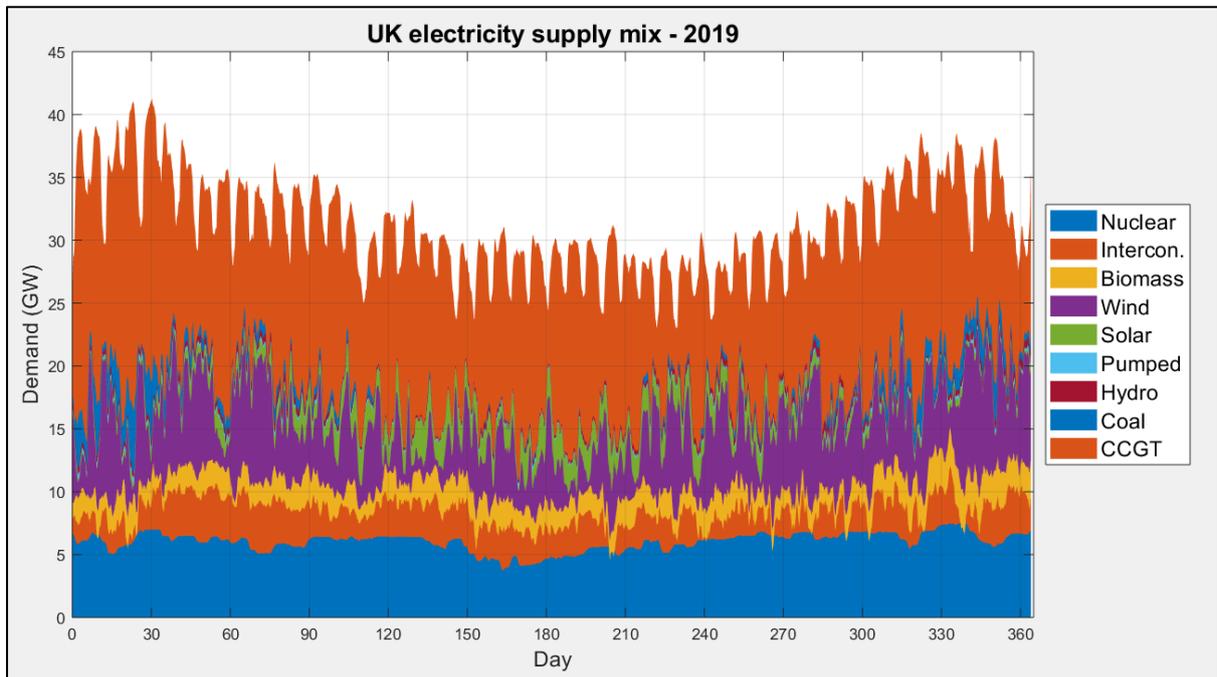
Figure 4.1. UK electricity supply by generation type 2019 (nuclear at the bottom to CCGT at the top). The data is sampled every 5 minutes and there is a 24-hour moving average applied.
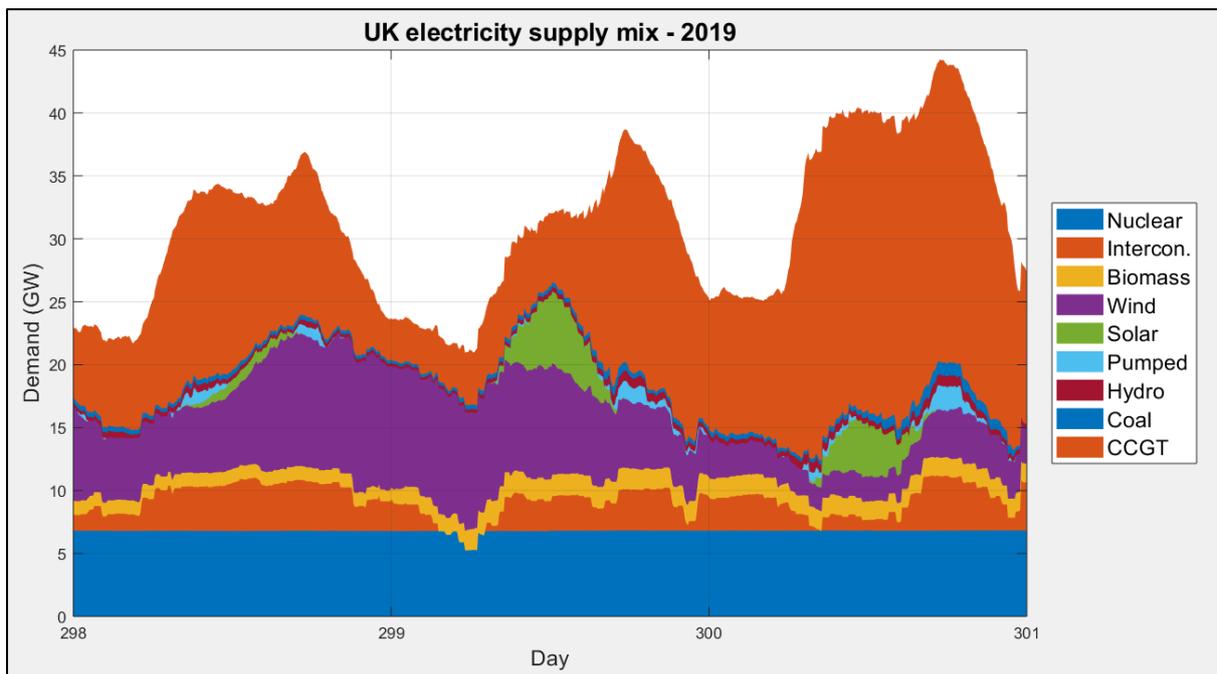
Figure 4.2. As in Figure 4.1 at 3 days where there is large variation in wind and solar generation and showing how CCGT is used as the primary balancing energy source.

The demand for electricity is expected to rise significantly over the next few decades while the UK transitions to a carbon neutral economy by 2050, as set in statute following international agreements. This will require a shift from fossil fuel use for transport, in the home for heating and for power generation. Not all of the energy will necessarily be provided directly through electricity, and biofuels or hydrogen may be used as a direct replacement for hydrocarbon-based fuels. At present most future

pathways indicate electricity will be involved in the delivery of energy. The carbon emissions may be reduced by using clean technologies or by carbon capture and storage (CCS). In whichever case, the electricity demand is expected to rise. Figure 4.4 shows a graph of the predicted daily demand in 2040 [15], broken down into electric vehicle (EV), heating and base power demand.
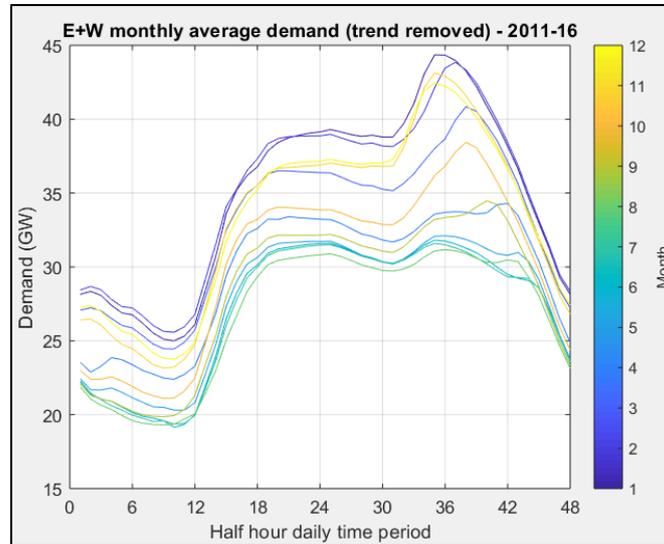
Figure 4.3 Average daily demand profiles per month. There is an overall increase in demand from summer to winter with a wider variation during the day compared to the night, and the early evening peak is greatly reduced during the summer.

Heating follows the daily demand profile we see today, with peak demand early to mid-morning and early evening. The EV demand is somewhat reversed with highest grid electricity demand during the night when batteries are charged; car batteries are a form of electricity storage. The ratio of maximum to minimum demand is approximately 1.4:1 compared to 1.7:1 today. If demand were more evenly distributed throughout the day then electricity from some sources could be used directly as it is generated without the need for storage. This would benefit tidal schemes where energy is generated in blocks at a little over 6-hour intervals and where the cycle times are retarded progressively each day.
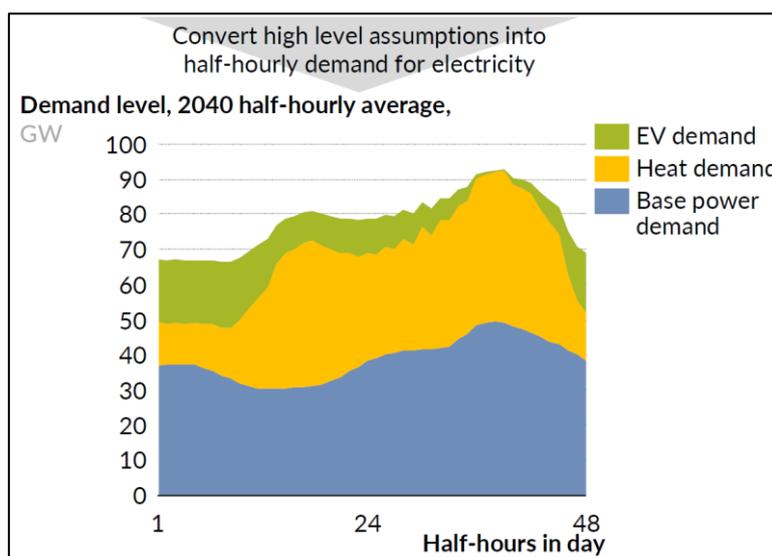
Figure 4.4 Predicted 2040 daily electricity demand (ref Aurora-Energy-Research). The electric vehicle demand is reversed high demand during the night.

4.2   Maximising supply value

The cyclical intermittency of barrage power means that at certain times the generation cycle does not coincide with high demand. The idea of adjusting the operation of the barrage to shift the timing to maximise electrical value has been addressed before [11]. They achieved an annual uplift of 10% based on day-ahead spot pricing at the expense of the total amount of energy. Note, this was for a Swansea Bay Lagoon where the phase of the tide is different from Morecambe Bay. The timing of the spring tide relative to the daily peak in demand has a significant impact on the ability to increase the electrical value via manipulation of the generating cycle (see section 4.3).

It is impossible to store energy as electricity; hence if it cannot be used it should be converted into a form that can be stored, or not generated in the first instance. Consequently, there is a requirement to store the energy until such time that it is needed. In the absence of a storage facility, is there a way to inform the model to help optimise the schedule of the barrage operation to favour power production when it matches demand? In other words, is there a way to use the existing optimisation process to determine the operational parameters to maximise the value-based energy.

This was addressed by assigning a pseudo value to the electricity based on the level of demand. This was achieved simply by using the demand profile, derived from the supply mix data to apply a value weighting to the energy. Figure 4.5(a) shows the average demand profile (using the same data as Figure 4.3) normalised by the overall mean value, and shifted to set the hours of negative demand between 11:30 pm and 6 am. This provides a time of day scale factor that is applied to each time step of the generated electricity. A value-based pseudo AEP was calculated by summing over the time series and dividing by the mean weighting value of 1.1673.

Figure 4.5(b) shows the scaling factor profile plotted over the power output from the barrage for a two-day period with the data coloured to highlight periods of demand and surplus. The negative power values are for pumping and represent a demand on the electricity grid; it is beneficial for these to occur during periods of surplus.
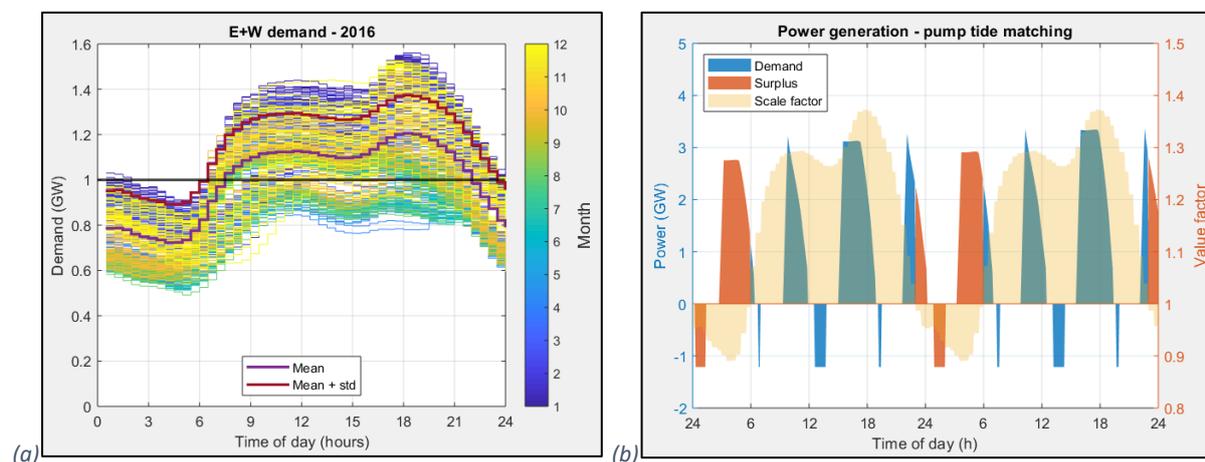
Figure 4.5 Using the daily demand curve to weight the generated energy and determine the surplus: (a) pseudo value scalar derived by normalising the demand curve and shifting ; (b) scalar applied to generated power curve.

Value based optimisation is achieved by ticking the check box on the Parameter Optimisation Tab in the GUI before execution. The modelling code calculates the pseudo AEP value automatically and the calculated value is displayed in the results panel of the Design Tab. A comparison between optimising for maximum energy and maximum value was made with very little difference in the resulting total energy generation (Figure 4.6). This reflects the limited scope to adjust the barrage operation within

the optimisation process, (it is impossible to revert to ebb only or no pumping for example), and may be in part due to the coincidence of the spring tide with the main demand peak (see section 4.3).

Generated energy		Generated energy	
Max theoretical TWh	16.224	Max theoretical TWh	16.224
This model TWh	0.396	This model TWh	0.396
Overall efficiency %	55.700	Overall efficiency %	55.687
AEP equiv. TWh	9.037	AEP equiv. TWh	9.035
APE equiv. TWh	11.142	APE equiv. TWh	11.132
Value scaled TWh	9.091	Value scaled TWh	9.096

Figure 4.6 Comparison of maximising the energy vs maximising the value. There is very little difference in this case reflecting the limited scope to adjust the barrage operation with the optimisation process and the coincidence of the spring tide with the main demand peak (see section 4.3).

4.3   Barrage energy and power generation characteristics

Before different energy storage strategies could be considered it was necessary to understand the nature and magnitude of the problem to be solved. Specifically, to understand the expected level of demand when the energy is transferred to storage, the level of demand that triggers retrieval, how much energy needs to be stored and for how long. The relationship between power, energy and time means any constraints on two variables imposes a constraint on the third. If there is a time limit on emptying a given size reservoir then this will impose a power limit.

The cyclic intermittency is dictated by the semidiurnal timing of the tides. A second, equally significant variation, is in the amplitude of the tide and is a result of the spring-neap cycle. Figure 4.7 shows the tide model for the whole of 2020, and a more detailed view of a section with extreme spring and neap tidal ranges. Spring tides occur when the sun and moon are aligned on either opposite or the same side of the earth; neap tides occur when the sun and moon form a right angle with the earth. There are two spring and neap tide cycles every lunar cycle.

(a) (b)
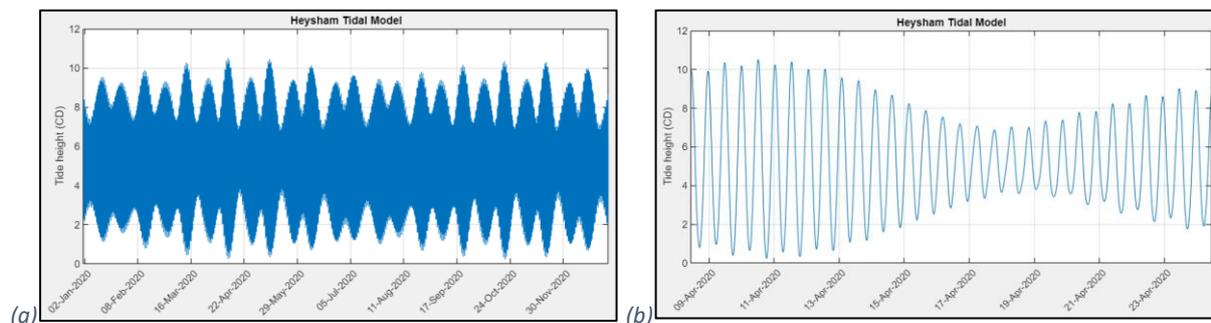
Figure 4.7. Tide series for 2020 (a) and a 16-day detail showing a section with unusually high/low spring/neap tides (b).

An inevitable consequence of this, if not obvious, is that spring tides occur around the same time of day each cycle. The actual timing depends on many factors and varies significantly around the UK. The Severn estuary tide is approximately 4.5 hours out of phase with Morecambe Bay [2]. At Morecambe Bay the timings happen to coincide with when the sun is either overhead or on the opposite side of the earth (conceptually as expected) and hence occurs around noon and midnight respectively. This is repeated every spring cycle. Figure 4.8 shows the daily timing and amplitude of all the high and low tide cycles for 2020 (blue dots) with 16 complete cycles traced out (orange dots).

This is significant because the amount of energy generated increases with the amplitude. Figure 4.9(a) shows the energy generated per cycle plotted against the amplitude (the marked difference between the ebb and flood cycles is due to the asymmetry of the reservoir-wetted area with elevation).

When viewed in terms of power, where the energy generation is spread over several hours, the timing pattern is still evident (Figure 4.9(b)). When this is averaged over the year it translates into an uneven power distribution. Figure 4.10 shows the average distribution together with the average daily demand profile. There is a reasonable alignment between the two factors, one of the barrage generation peaks coincides well with the early-evening demand peak.

Pumping is shown to increase the average power significantly and at times of low demand then it seems sensible to revert to a non-pumping operation mode. This might not be an option if strict tidal matching controls are in place.
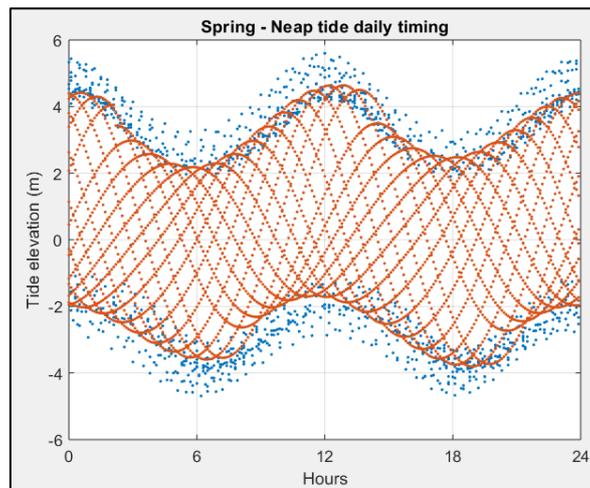
Figure 4.8. Spring tides occur at the same time of day each time.

(a) *(b)*
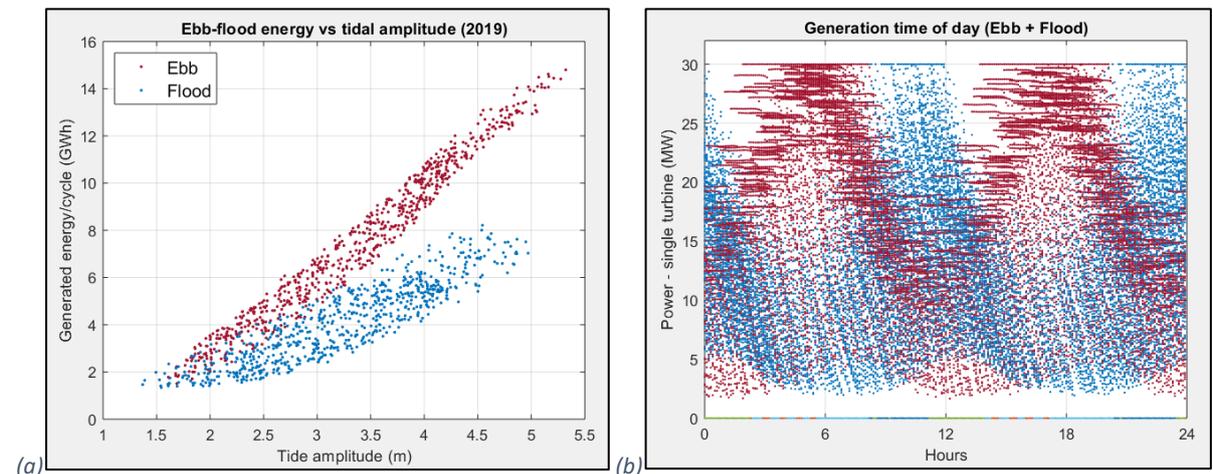
Figure 4.9. Generated energy and power: (a) energy per cycle vs tidal amplitude; (b) power per 0.1 hour time step vs time of day (2019 simulation).
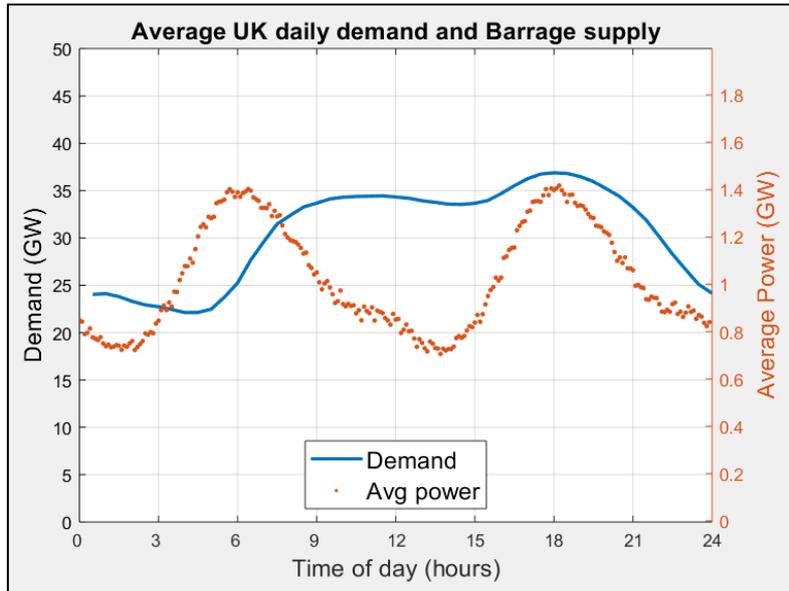
Figure 4.10. Average daily power generated for 2019 (per 0.1 h time step - compared against the daily demand profile. The fluctuations in the power reflect the spring/neap tide timing. The second daily ebb spring generating cycle coincides with the evening demand peak.

The maximum power output of the barrage would dwarf all the other individual generating facilities in the UK. This value is dependent on the number of turbines and is calculated from the maximum power per turbine multiplied by the number of turbines and the availability factor. Maximum power of 160 x 30 MW turbines x 0.95 availability = 4.56 GW. Figure 4.11 shows the peak power plotted against the generated energy for every cycle during a simulation for the whole of 2019 (the colour represents the cycle duration). As a comparison, Table 4-1 lists the top 26 highest power generating plants in the UK, including all the nuclear stations. Even the lowest peak power value from the barrage would make it into the top 10.

The ability of the National Grid to handle this amount of power may impose a limit on the size of the barrage scheme.
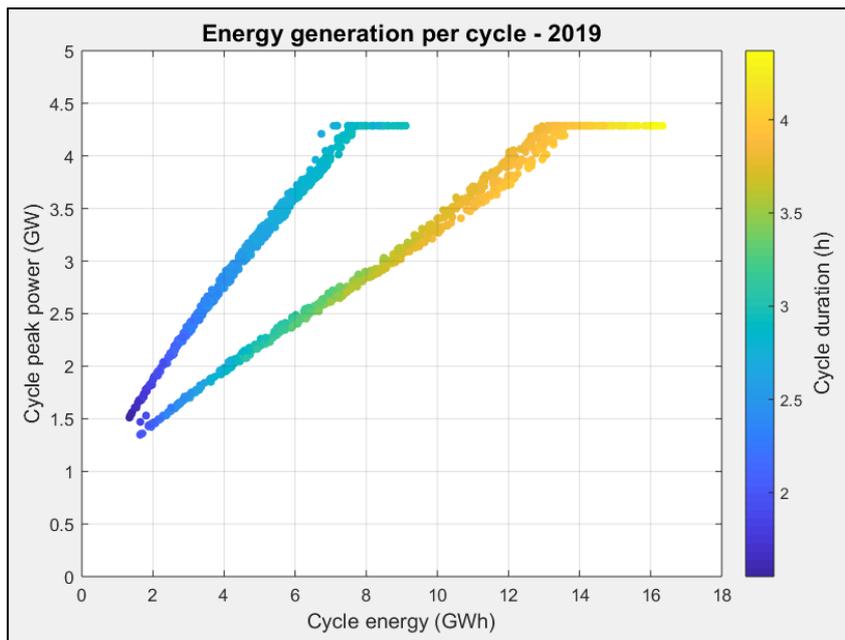
Figure 4.11. Cross plot of cycle peak power against energy coloured by cycle duration.

Power Stations in the United Kingdom (operational at the end of May 2020)				
Station Name	Fuel	Type	Installed Capacity (MW)	Location
Drax - biomass units	Biomass (wood)	Bioenergy	2,640	Yorkshire and Humber
Pembroke	Natural Gas	CCGT	2,199	Wales
Ratcliffe	Coal	Conventional steam	2,021	East Midlands
West Burton	Coal	Conventional steam	2,000	East Midlands
Dinorwig	Pumped Storage	Pumped Storage	1,800	Wales
Staythorpe C	Natural Gas	CCGT	1,772	East Midlands
Grain CHP*	Natural Gas	CCGT	1,517	South East
Didcot B	Natural Gas	CCGT	1,450	South East
Connahs Quay	Sour gas	CCGT	1,380	Wales
South Humber Bank	Natural Gas	CCGT	1,365	Yorkshire and Humber
West Burton CCGT	Natural Gas	CCGT	1,332	East Midlands
Drax - coal units	Coal	Conventional steam	1,320	Yorkshire and Humber
VPI Immingham*	Natural Gas	CCGT	1,252	Yorkshire and Humber
Heysham 2	Nuclear	AGR	1,240	North West
Seabank	Natural Gas	CCGT	1,234	South West
Hornsea 1	Wind (Offshore)	Wind (Offshore)	1,218	Yorkshire and Humber
Torness	Nuclear	AGR	1,200	Scotland
Saltend*	Natural Gas	CCGT	1,200	Yorkshire and Humber
Sizewell B	Nuclear	PWR	1,198	Eastern
Hartlepool	Nuclear	AGR	1,185	North East
Peterhead	Natural Gas	CCGT	1,180	Scotland
Dungeness B	Nuclear	AGR	1,090	South East
Heysham 1	Nuclear	AGR	1,060	North West
Hunterston B	Nuclear	AGR	985	Scotland
Hinkley Point B	Nuclear	AGR	965	South West
Spalding	Natural Gas	CCGT	950	East Midlands

Table 4-1 UK's largest power generating stations [36].

In addition to energy generation, there are also the pumping requirements to increase the head. Pumping takes place at constant power at a quarter of the maximum turbine power i.e., 1.14 GW. The duration of the pumping phase varies with the cycle and Figure 4.12 shows the distribution in the amount of energy required for pumping during 2019. In all but extreme cases this is less than 1.5 GWh. Because pumping is also used as the means used to satisfy the tide matching requirement, it may not be permissible to simply omit the pumping phase at times when the demand is low.
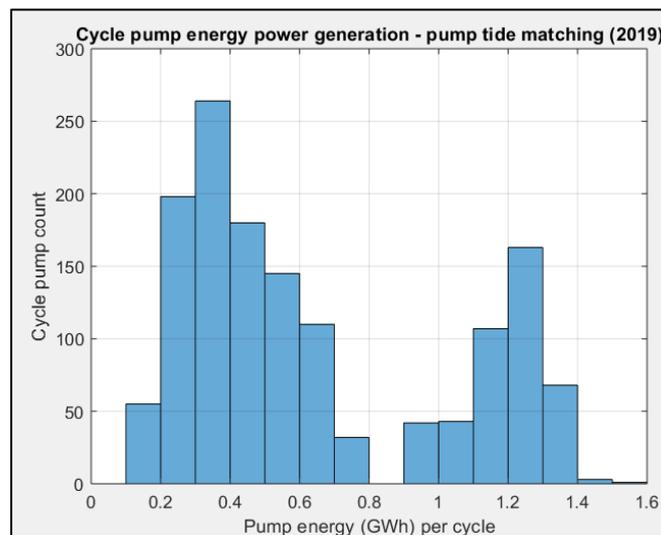
Figure 4.12. Distribution of pump energy per cycle. The pump power is fixed so the variation is due to the pump phase duration.

4.4 Storage options

This section explores the energy storage options available that are capable of handling the surplus energy from the barrage. The surplus energy, as defined by the low demand period overnight, may span more than one generating phase. Figure 4.13 shows the surplus energy per generating phase and the cumulative overnight energy with a zoomed section for clarity. In practice the low demand period is similar in length to a half tide cycle, and when two surplus generating phases are involved, they are partial phases, and the total is similar to a complete single cycle. Figure 4.14 shows the energy distribution for 2019.
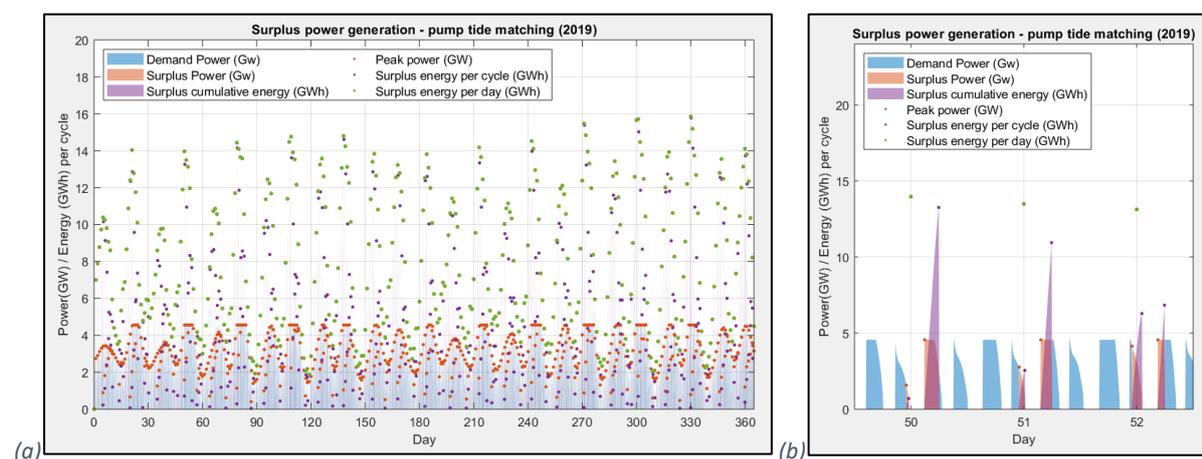
Figure 4.13. Cumulative surplus energy (green dots) from one or two generating phases per night, with detail for clarity.
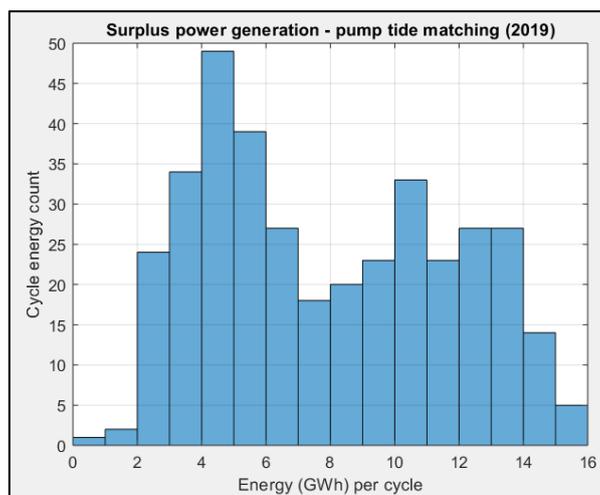
Figure 4.14. Distribution of surplus (overnight) energy.

To put the size of the energy and power storage requirements into perspective, the data plotted in Figure 4.11 above, are superimposed on a plot (Figure 4.15) of cycle energy storage capacity vs cycle power transfer rate for different energy storage technologies [37]. The Dinorwig Pumped Hydroelectric Storage (PHS) scheme in Wales has a power rating of 1.8 GW and energy capacity of 9.1 GWh, sufficient to supply electricity for a little over 5 hours at full power (this point is also plotted). Note the logarithmic scale. The ideal energy storage capacity capable of accommodating all the surplus energy from the barrage is on the order 16 GWh – a 75% increase on Dinorwig. Much more significant is the power requirement. For conventional PHS, the energy release is designed to occur over the peak demand periods with ample time to recharge overnight. For the barrage, the surplus energy is delivered over a relatively short time frame and it is this that will dictate the power requirements. In fact, the power requirements would need to match the power output of the barrage

if all the energy is to be utilised. Such a storage facility would have to be on a comparable scale to the barrage itself, and it is this that prompted alternative approaches be considered.
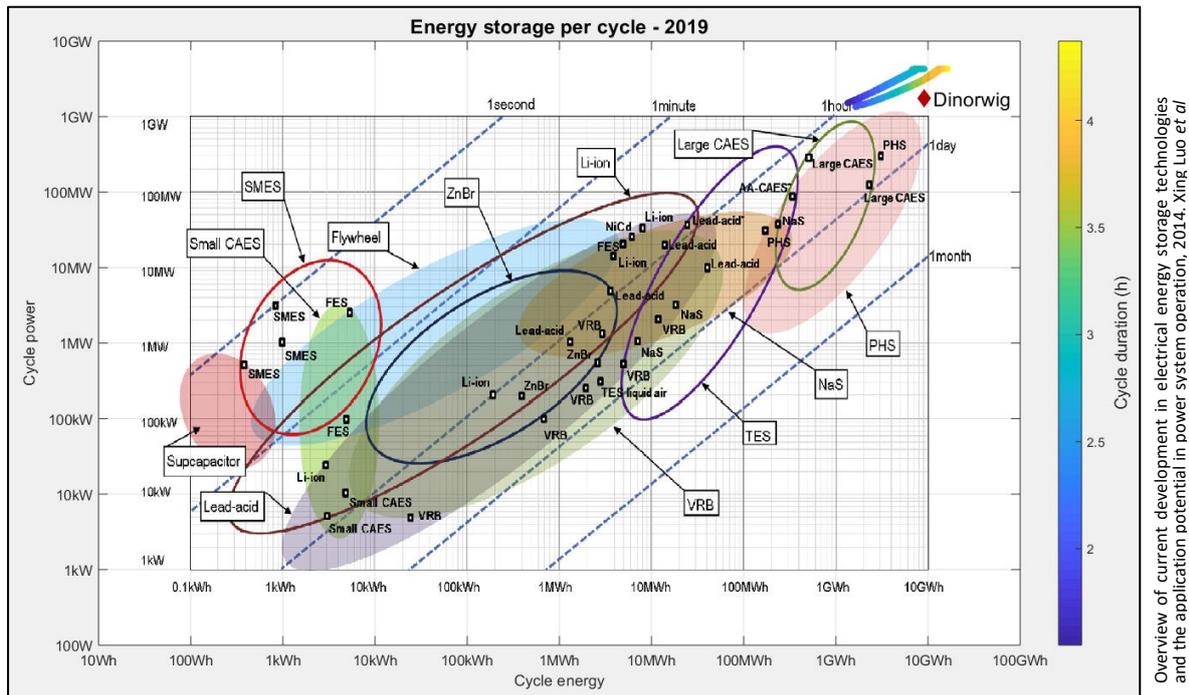
Figure 4.15. Morecambe Bay cycle peak power superimposed on a plot of cycle energy vs cycle power for different energy storage technologies [37].

An idea proposed by NTPG was to repurpose a quarry, used to provide material for the barrage construction, as a storage reservoir for a PHS scheme. As an exercise to get a feel for the energy storage capacity of such a scheme, the Kirkby Quarry (Figure 4.16) near Kirkby-in-Furness was chosen as a potential candidate. Figure 4.17 shows a location map and satellite image respectively. Figure 4.18 shows details of the site with suggested areas annotated and elevation data.

For an average depth of 50m, average area of 120,000m² and average elevation of 220m, the gross energy would be 2.9GWh. To achieve this would require significant further excavation, although if this was the source of building material that may be achieved as a matter of course. The major issue is there is no lower reservoir; here it is assumed water is drawn from and returned directly to the sea, which is unlikely to be permissible in the area.

Figure 4.16. Photograph of Kirkby Quarry, example candidate site for a pumped hydroelectric storage scheme.
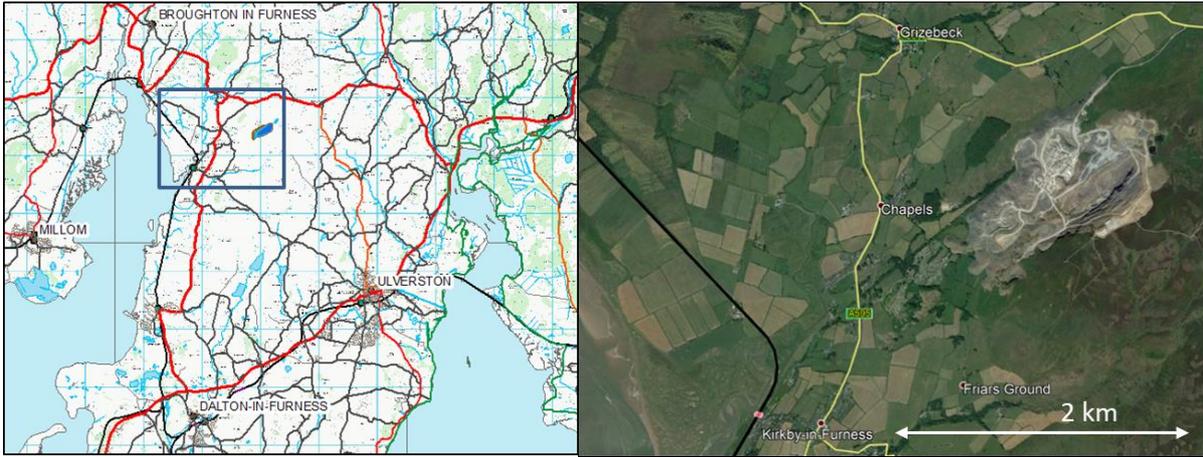
Figure 4.17. Location map and satellite image of Kirkby Quarry.
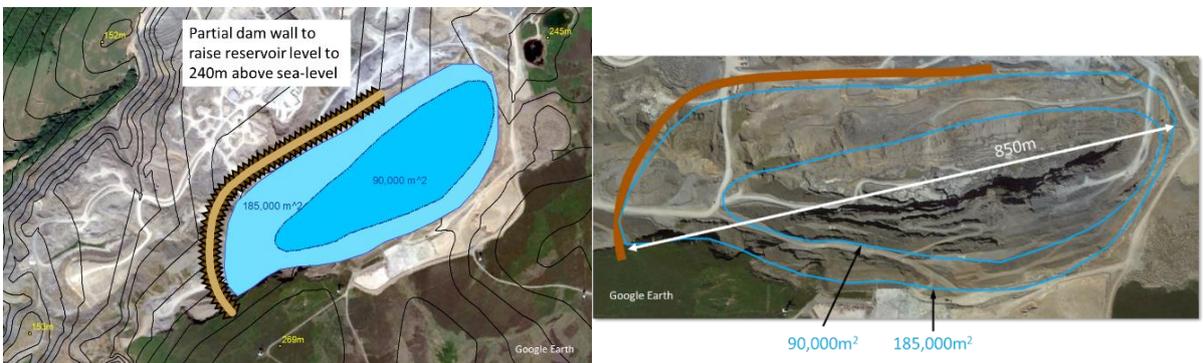
Figure 4.18. Details of the site with elevation, dam wall and reservoir volume.

A storage capacity of 3 GWh would impose a significant limit and result in a large percentage of discarded energy. A power limit could also result in further discarded energy if it was insufficient to provide a total 3 GWh of energy over the duration of a generating cycle. Figure 4.19 shows the distribution of the duration of the generating cycle for 2019. To store 3 GWh of energy in 1.5 hours requires 2 GW of power, more than the capacity at Dinorwig.
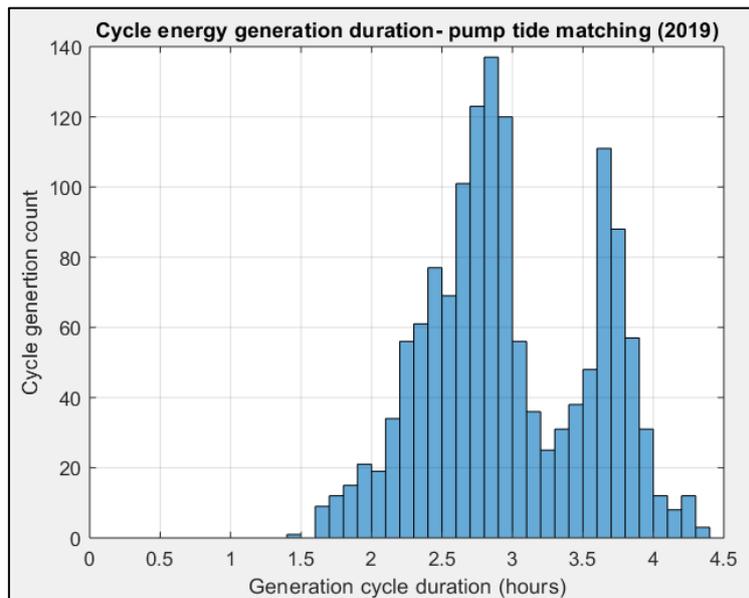
Figure 4.19. Distribution of generating cycle duration. The total energy stored is the product of the storage power and the duration over which it is stored.

58

4.5   Energy balancing

In the absence of a storage facility tailored to handle the cyclic intermittency of the barrage, then the burden is placed on the National Grid to balance the energy (supply must match or exceed demand). The terms energy and electrical power can almost be used interchangeably here since the power imbalance has to be kept within a very tight tolerance at all times, and there has to be no energy imbalance over a 24 hour period on average, i.e. any energy imbalance due to a power imbalance has to be compensated for over a 24 hour period.
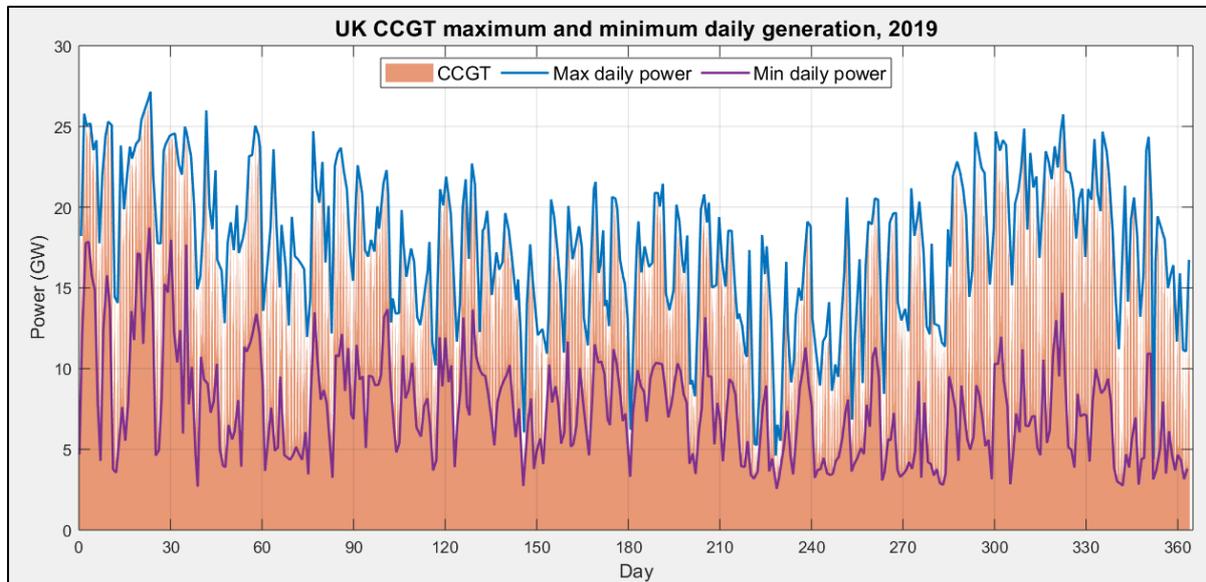
Figure 4.20. Maximum and minimum daily UK power generation using CCGT. This fluctuates to accommodate the difference between the other supply sources and demand.

As detailed in section 4.1, by far the bulk of energy balancing is achieved by ramping up and down the electrical generation of CCGT. Figure 4.20 shows the minimum and maximum daily values during 2019 and it illustrates the range and irregular pattern in response to a combination of the change in demand and the variability in supply of wind and solar power.

To illustrate the impact of relying on the current mechanism to handle the power demands of the barrage, Figure 4.2 is repeated here (Figure 4.21) with the barrage power curve subtracted from the CCGT curve (the original CCGT curve is plotted as a solid black line). Considering the fluctuations already taking place, the addition of the barrage power looks as if it can be accommodated within the existing mechanism. This follows the same approach as for wind and solar where the "*green*" energy option takes precedence over gas. The predictability of tidal range power means CCGT generation can be controlled proactively. There are times where the CCGT power output is less than the maximum output of the barrage, and the barrage output can be reduced accordingly.
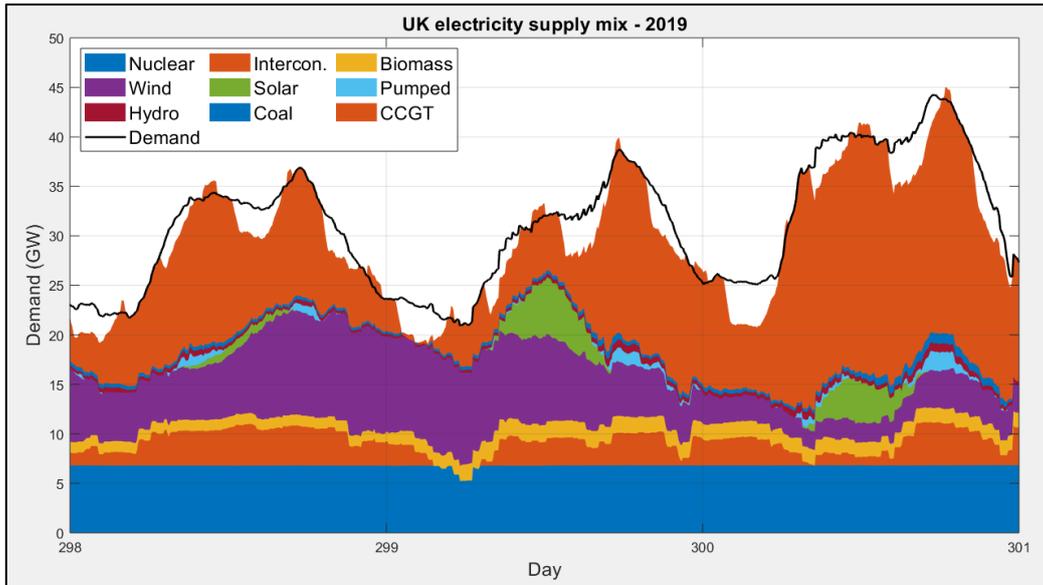
Figure 4.21. Supply mix data from Figure 4.1 with the CCGT values adjusted by the barrage power values. The black line represents the unadjusted values.

Potentially more problematic is how abruptly the power changes. Figure 4.22 shows the power curves for wind and solar, together with the barrage power. Although there are large changes in power of wind and solar it is by no means instantaneous. Power generation distributed over the whole country will not be subjected to even abrupt variation in solar or wind energy all at the same time. The modelling code written as part of this project initially assumed that transitions between operating phases occur instantaneously over a time step, effectively acting like a switch. In reality it will take more time to open and close the sluice gates to allow flow through the turbines.
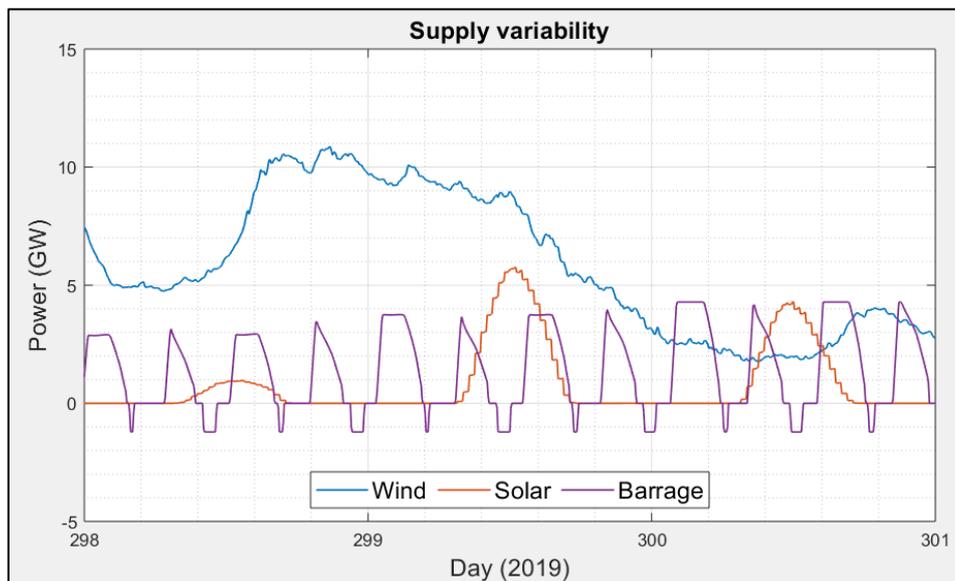
Figure 4.22. Wind, solar and the barrage power. Although there are large fluctuations in wind and solar the changes are not very abrupt compared to the barrage sudden onset.

A simple way to slow down the transition is to progressively ramp-up the number of turbines over a period of time. This was applied at the start of the generation phase. A short section of the power time series for different ramp time models is shown in Figure 4.23(a), and the impact on the net AEP shown in Figure 4.23(b). The power naturally ramps down as the head progressively falls and the abrupt change at the end of generation only occurs after the turbine minimum operating head is

60

reached at a much lower power level. A transition time of 1-1.5 hours would facilitate the energy balancing process with a relatively small loss of AEP. Given the predictability and control over barrage generation, the precise timing of the transitions can be communicated directly to the National Grid in advance.

(a)

(b)
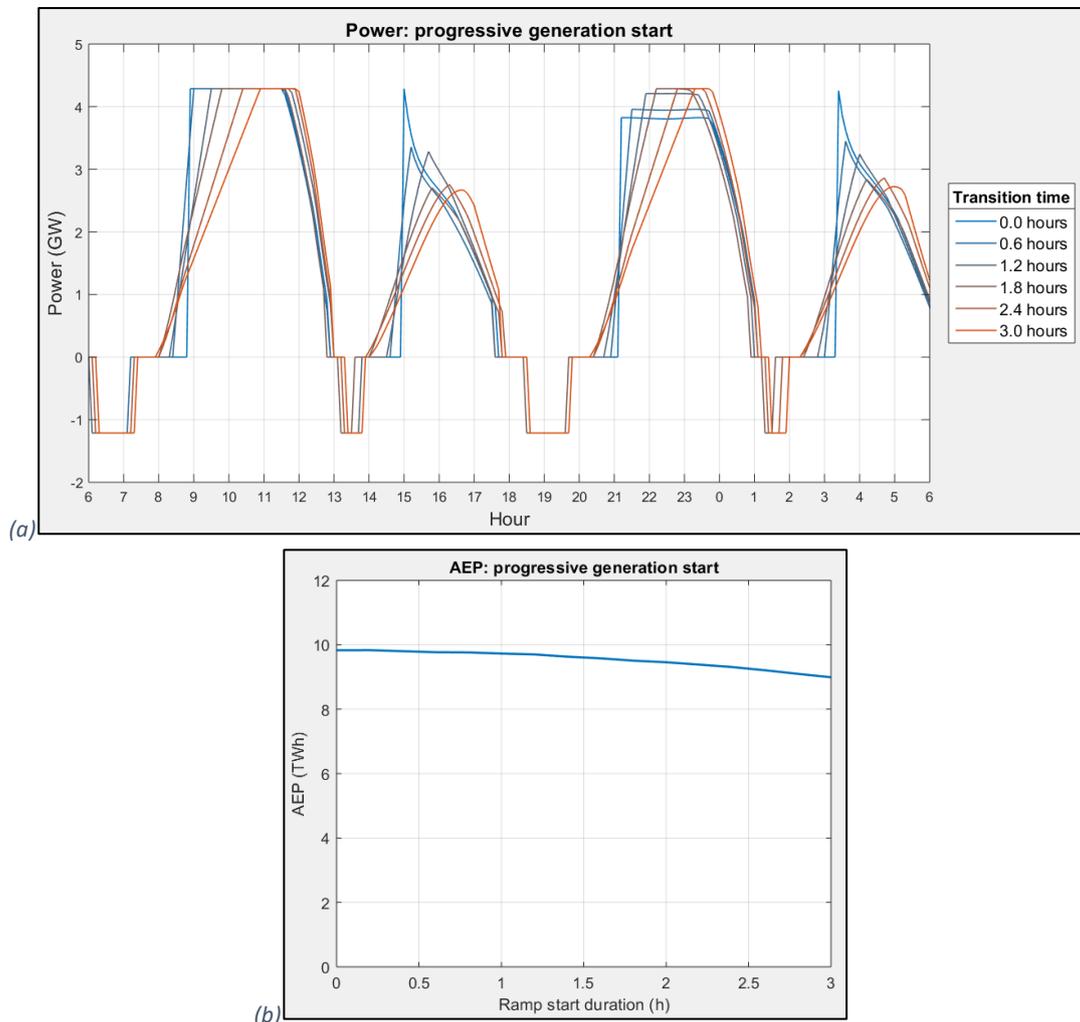
Figure 4.23. Power and energy after applying different ramp-up times – achieved by a progressive introduction of turbines: (a) power time series; (b) AEP against the ramp-up time.

As mentioned in section 4.1, the demand patterns and the way we meet those demands are expected to change. The charging of electric vehicles and the production of hydrogen using electricity are best performed when there is a surplus of electricity and when the price is low. The predictable nature of tidal energy means consumers can satisfy their energy requirements at a time when the barrage is generating power. There should be no need to curb supply to achieve energy balancing.

5 Discussion

Several key findings have been identified during the study. The data and model section, primarily demonstrated the importance of optimisation; the need for robust and consistent turbine and pump models; and the inclusion of the generating stopping head as an optimised operational parameter. The turbines can be controlled in a variety of ways that influence both the efficiency and timing of generation. Whilst our study was clearly at an early phase in the assessment, turbine operation deserves more attention, particularly determining the potential uplift between double regulation and triple regulation operation.

Traditionally, the main purpose of a barrage is to generate large amounts of energy. The renewable and sustainable character of the power mean that it can contribute to the UK's power supply as part of a zero-carbon economy. One of the main drivers for development is to limit the concentrations of greenhouse gases in the atmosphere and in turn mitigate climate change and sea-level rise. However, locally, the barrage can reduce the threat of increased flooding by controlling the water levels within the barrage reservoir, and similarly regulate the tide to moderate the impact of sea-level rise on the inter-tidal zone. Modifying the operation of the barrage to satisfy these requirements will shift away from operating to maximise energy or value output, and will have associated costs. Sea-level rise will increase these demands and the barrage design must not only be capable of maintaining this function but must also have clear management structures to direct the operation to meet different objectives.

Upon examination, the option to store all the surplus energy from the barrage during periods of low demand with a single dedicated facility is unfeasible. With electricity usage patterns expected to change as the country transitions to a low carbon economy, the solution is perhaps one of energy balancing, and using surplus energy to charge electric vehicles or produce hydrogen?

5.1 Modelling and operational parameter optimisation

A requirement for the modelling program was that it should be suitable for design parameter sensitivity analysis. Although it is accepted that 0-D modelling is subject to large clearly defined errors, it can be a perfectly robust and unbiased approach to take, provided the errors are constant or proportional over the value range of the parameters being analysed. The next question is whether the algorithm itself behaves in a consistent manner. A detailed review of how the 0-D method has been implemented can help answer this question.

At the core of the simulation, the program loops through each time step of the tide series, and the transition to the next operating mode is triggered when a certain head or barrage level is reached. These trigger values control the operation of the barrage, and the optimum levels vary with the tidal range. If a level is not reached, the trigger does not occur, and the sequence breaks down. The sequencing may or may not restart at some subsequent tide cycle; in either case there is a loss in energy generated and a reduction in the calculated AEP value. This may occur, for example, where a reduced discharge capacity prevents the barrage from reaching its expected hold level and the generating starting head is never achieved. It is important to quality control (QC) the result; this can be done easily by inspecting the head-levels plot with the background phases displayed - any breaks in the sequencing are clearly visible. The trigger values are termed operational parameters.

Well before a breakdown point is reached the trigger values will start to become suboptimal, and the calculated AEP value will suffer. If the operational values remain fixed as the design changes then any variation resulting from the change will be indistinguishable from the variation due to the suboptimal operation. Therefore, optimisation should be performed for every design point with the aim of maximising the structural variation in AEP to the design change alone.

Early optimisation studies [7] looked to derive functions to relate the operational parameter values to the tidal range. The same approach was tried here (section 2.6.1) on the basis that the results should be accurate since they would be derived from a grid search scan over the design, tidal range and operational parameter space. The problem came from the scale, complexity and non-linearity of the task, and that it was applicable to a limited number of design parameters (number of turbines and sluice gates in this case). Any other design change would require the whole process to be repeated.

The second approach was to develop a method that could perform the optimisation on a design-by-design basis. Although this removes the design variables from the functions, there is still the tidal range variation to contend with. The solution was to directly solve for the polynomial coefficients that relate the operational parameter value to the tidal amplitude. A simple linear relationship was adopted to minimise the number of variables to solve, i.e. to an intercept and gradient for each variable. Standard optimisation algorithms were employed and the cost function was set as the negative of the AEP. When a representative tidal series was used for the modelling, a new simulation could be run directly without repeating the optimisation step, for any tidal series of any length and sample interval.

Since there are a great many possible design variations, the design-by-design optimisation approach is very appealing, provided it is reliable and relatively quick. These two requirements are counteractive. Reliability is improved using a shorter time interval (results in a smoother AEP function) but takes proportionately longer to run.

As discussed in section 2.4, greater barrage operational efficiency can be achieved by varying the generating stopping head with the tidal range. The AEP increased from 13.7 TWh to 14.4 TWh when optimised for a 4m tidal amplitude. This is somewhat higher than the average amplitude and is probably an overestimation of the effect. None the less, it is not insignificant. Promoting the ebb and flood generating starting heads from fixed values to variable operational parameters would increase the number of operational parameters to optimise from four to six (excluding pumping limit levels). It is not clear how much difference this would make in practice, however. The program will end the generating phase early (above the set stopping head) if the turbine minimum head limit has been reached. Optimisation of the turbine control parameters will favour the higher head part of the generating cycle and cutting short the generating cycle may be beneficial overall.

The accuracy of the modelling may also benefit from using a second order polynomial (as discussed in section 2.6.2), increasing the number of optimisation variables by a further 50%. It is not clear if this would be significant when it comes to sensitivity analysis but is worthy of further investigation. Implementing both the stopping head and higher order polynomial changes would be prohibitive in terms of the optimisation runtime as currently implemented.

The absence of detailed and up-to-data turbine and pump models was another source of inaccuracy. There is an outstanding question related to the turbine model: the free flow discharge rate, calculated using a coefficient of discharge (Cd) of 1.1 and the runner diameter (equation 9), is less than the discharge calculated using the Hill chart when generating. A Cd greater than one reflects there is a venturi effect because the feeder tubes are of a larger diameter than the turbine. The choice of 1.1 was decided in consultation with Prof. George Aggidis. A lack of comprehensive and accurate information in the turbine model calls into question the confidence of the predictions.

The situation is no better for the pump model. The simple model implemented in the program, was controlled to behave in a similar way and with similar efficiency characteristics, as the pumps at La Rance. This was achieved without changing the model, just the way it was operated and the power

level used. The transition from sluice to pumping is handled differently at La Rance however; pumping begins while there is still a positive head in the direction of water flow. In the model, pumping does not start until after the equilibrium level is reached.

The magnitude of the errors due to the pumping model remain unknown. Provided the sensitivity analysis results are valid, then at this early stage in the design process the absolute errors can be ignored. The main concern is the lack of consistency between the turbine and pump models. The calculated AEP value is the net energy between the pump energy used and the energy generated. Any discrepancy in the efficiency between the pump and generating cycles will directly affect the results.

Another area of interest, related to turbine operation, is in the flexibility afforded by triple regulation, where power electronics allows the speed of the turbine to vary independently of the electricity grid phase velocity. This enables the turbine to be operated anywhere within its operational envelope (see section 2.3.1). The program provides two different ways to operate the turbine and these can be considered as end member options; constant speed at maximum flow (for maximum power), and constant Q_{11} (turbine Hill chart model flow) for higher turbine efficiency when a high flow rate is not critical. In both cases the parameter values are held constant throughout the individual generating phase. Neither option fully exploits the freedom offered by triple regulation.

Operating at constant Q_{11} was not fully explored during the analysis. Higher turbine efficiencies are achieved at lower flow rates, and this can only result in a higher overall efficiency if there is ample turbine discharge capacity. This is because the reservoir must be fully recharged in preparation for the next cycle in the time available. Analysis using more turbines would help to test this idea further. The benefits of using triple regulation also warrant a more complete investigation.

For larger tidal ranges the generating head range can test the operating limits of the turbine, and triple regulation becomes increasingly advantageous the greater the variation between spring and neap tides. Keeping the speed constant for any given generating cycle was applied as an artificial constraint to simplify the model; triple regulation imposes no such constraint.

The large difference in estimated generated energy values between 0-D and 2-D modelling raises the question, would a design optimised using 0-D modelling also be the optimum design for the equivalent 2-D model? Sub-optimal operational parameters significantly reduce the amount of generated energy demonstrating the sensitivity of the model. Optimising the design and operation in 2-D could significantly increase the generated energy and reduce the differences currently observed between 0-D and 2-D modelling.

5.2   Power balancing vs energy storage

The storage capacity required to accommodate the power and energy levels experienced during a spring tide generating cycle would exceed anything that is available today. Typically, storage facilities are designed to provide electricity over several hours during the day and are recharged during the night. The relatively short duration of a generating cycle, of around 3-5 hours, requires the storage facility to match the power level of the barrage if it is to capture all the energy generated. This would force the scheme to be of a comparable size to the barrage itself. On this basis the focus switched to energy balancing.

The cyclic power characteristics of generation and consumption (pumping) imposes a burden on the electricity supply grid to maintain the energy balance. Although the maximum power from the barrage is potentially very large (>4 GW), the national grid already deals with fluctuations in wind and solar power 2-3 times higher. The difference is, wind and solar generation is distributed all around the

country and the fluctuations on average are not very rapid. The advantage with tidal power is its predictability, and the problem with abrupt transitions can be mitigated to a degree at source by ramping up and down the generation over an hour or more.

In our bid to become carbon neutral by 2050, there is a need to reduce our dependence on burning fossil fuels for power, personal transport and heating. Currently, the options to replace fossil fuels are renewables and nuclear delivering predominantly electricity, with in future a growing proportion of hydrogen. Quite how hydrogen will fit into the mix is not clear, however it will still require production, which requires energy. It is still considered acceptable to use fossil fuels provided the carbon released is captured and sequestered, although this has not been applied at scale and is costly. The current predictions are that the UK will require around 60% more electricity by 2040 [15]. Also predicted is a change in the daily demand profile with proportionally higher demand overnight for the charging of electric vehicles. Whether hydrogen is used purely as a fossil fuel replacement or also as an energy storage medium in its own right, it can be produced using surplus electricity. The uncertainty in the future daily demand profile and level of demand, together with the changing supply mix expected over the next few decades makes decisions based on today's conditions very risky, and embarking on expensive conventional storage schemes could prove to be costly white elephants.

This raises the questions; is there an optimum size for the barrage and if so, what is it? The initial default approach is, whatever maximises the amount of energy, which quickly evolves into, whatever maximises the energy value. The location and design of the barrage ultimately determines the size of the reservoir and hence the amount of energy and power that can be generated. As a standalone tidal scheme, it poses energy balancing issues for the supply infrastructure as it is today. If operating in conjunction with other tidal schemes, then a more continuous supply is possible [2]. In short, anticipated future changes in both electricity demand and the supply mix mean there is every opportunity to fully utilise the output of the barrage without the need for a dedicated storage scheme.

5.3   Environmental considerations and sea-level rise

The environmental and ecological impact of a barrage must be valued at the outset. The constraint of matching the existing tidal range for each tide cycle was applied based on the assumption that this will be a good way of preserving the inter-tidal zone and therefore good for the ecology. The modelling found it is possible to maintain the existing tidal ranges under the dynamic conditions of sea-level rise driven by climate change. These restrictions impose limitations on the number of model variables by fixing the pumping barrage levels and serve to simplify the analysis. If the barrage levels were free to vary, then optimising the energy output would add another two operational parameters to the analysis. It may be possible to satisfy the environmental and ecological requirements and still allow some freedom to modify the tidal limits. Setting these limits requires a detailed understanding of the state and dynamics of inter-tidal ecosystems that can only be gained by wider consultation with informed parties.

Tied in with this is the potentially huge benefit of flood prevention and alleviation; flood prevention from sea inundation, and flood alleviation by facilitating river drainage. The scale of the existing flood prevention infrastructure is large with more than 50km of embankments, vastly more drainage ditches, one-way flow gates and pumping stations. Global warming not only brings sea-level rise, it also threatens more severe storms. Storm Desmond that flooded Kendal and the Lyth valley in December 2015 highlights the danger due to river flooding.

At this point, no attempt has been made to value this benefit. The concept of habitat loss compensation can work the other way if the barrage can protect the Bay from the damaging effects of sea-level rise. How is this valued? The cost of flood prevention is perhaps easier to estimate in

monetary terms. In addition, operating the barrage to achieve a level of tidal control (section 3.1.2) imposes a cost by limiting the amount of energy generated and lengthens the payback period.

The primary goal at this point in the development process is to pursue government to fund a detailed feasibility study of the barrage. Any tidal range scheme of this nature in the UK will be the first of its kind. A decision to proceed with the study is not a decision to build it. Given the necessity to reduce carbon emissions there is an urgency to develop our knowledge and understanding of all aspects of such schemes. Without this, it is impossible to make informed decisions. The clean energy credentials of tidal range energy have to be fairly appraised against the alternatives. The appraisal process follows strict guidelines and the strike price (cost per MWh of electricity) for tidal range schemes is considered to be too high. This is calculated over a fraction of the operating life and the true value is far lower. In the long-term tidal range generation is capable of providing a huge amount of carbon free electricity, the problem is how to get it accounted for properly in the overall assessment process.

5.4 Recommendations

As the design process progresses then the accuracy of the predictions is expected to improve. At the current stage, where 0-D modelling is still the primary tool, it would be of little benefit to spend significant effort on second order effects. Nevertheless, an appreciation of the contribution to the error in AEP of the different sources would be useful. Where practicable, sensitivity analysis should be performed. For example, the error in the LIDAR data is up to 0.15m – if this were systematic to the full extent then this is equivalent to a drop or rise in sea-level of 15cm, which is easily modelled.

While still following the 0-D approach implemented here, the following are recommendations for further study:
- Examine the generating stopping head as an optimised operational variable.
- Determine the maximum AEP possible with and without environmental constraints.
- Understand the gains available using triple regulation over double regulation.
- Explore the effect of tracing different turbine operational paths across the hill chart.
- Explore the potential gains of optimising for value; is this dependent on the timing of the spring high tide?
- Check if using a second order polynomial for the operational parameter functions will materially impact the sensitivity analysis results and how the calculated energy is affected.
- Obtain improved turbine generating and pump models – at least use self-consistent models.
- Check the sensitivity to errors in the efficiency of the turbine generating and pump models.
- Explore relaxing the tidal matching requirement – e.g., extending the range at neap tides.
- Calculate AEP sensitivity due to LIDAR data tolerances.

A general recommendation is not to delay starting the discussion about the impact sea-level rise will have on the environment and ecology of the Bay. Discussions should include how a barrage might mitigate this and what characteristics of the tide are important. How might the benefit of protecting the inter-tidal zone be valued? Examine the Environmental Agency Shoreline Management Plans [38]; their objective and ambition, their cost, and if possible estimate their likely effectiveness. Do they fully cost raising the height of all the embankments, for example?

The issue of how and to what extent the barrage is future proofed against sea-level rise should also be addressed early in the design process. The potential cost implications may be significant, particularly if the more extreme levels are used. There is a great deal of uncertainty in the predicted values over the design life of the scheme. The idea of designing for future augmentation, either in terms of height or the number of turbines, should be considered.

6 Conclusion

The following six main conclusions were reached from the analyses performed during this project.

6.1 Optimisation

The set of operational parameters (starting head, turbine speed etc.) that produce the maximum energy vary with each tidal cycle. They are also influenced by changes to the barrage design (position, number of turbines, turbine operation mode etc.) and with the way the barrage is operated. The optimised values for one design are unlikely to be optimal for another. It is important therefore when comparing different designs that the operational values used in each case are suitably optimised.

Relating the operational parameters to the tidal range with a set of simple functions is an effective way of handling real tidal series in the 0-D modelling. Since the backward difference method is purely numeric, it forces any optimisation approach to be derived through repeated runs of the program. One way to build up the functions is to scan over a range of design parameter values and tidal ranges and fit a function through the points that produce the maximum energy. There are a large number of possible design configurations however, and setting up all the functions in advance is an almost impossible task.

Applying a set of linear functions for the operational parameters (section 2.6.2) makes the process of optimising the function coefficients specific to each new design a much simpler task, and can be achieved by minimising a cost function - set equal to the negative of the energy.

Once the global minimum solution has been found for a design, then using a local minimum search can successfully track the global minimum when small incremental changes to the design are made. This is ideal for sensitivity analysis and can be performed fairly quickly, typically on the order of 2-3 minutes using a standard laptop computer. The initial global search can take significantly longer however and was most reliably found using a grid scanning approach.

6.2 Turbine regulation

Triple regulation is made possible by the use of power electronics. It allows the speed of the turbine to differ from the generator speed, and affords far greater flexibility in the control of the turbine. Two options were implemented in the model, one using turbine speed, the other model flow (Q_{11}). These were held constant for the individual cycles but were otherwise optimised for the tidal series. The constant speed option is the closest to double regulation and the model flow option was designed to allow the turbine to operate in regions of the Hill chart with higher efficiency.

The analysis was not taken to its full conclusion and it is recommended alternative ways of controlling the turbine within the operational envelope be investigated with the aim of finding potential improvements in the overall turbine performance. A reliable estimate of the uplift available in changing from double to triple regulation should also be ascertained. The absence of a complete set of turbine generating and pumping performance data limits the scope for this work however. The generating and pump models used in the study have different origins and any errors from using both for the same model are unquantified. The uplift observed using pumping is consistent with other studies and the results are considered to be reliable.

6.3 Operational model parameters

It was found that varying the generating stopping head with the tidal range could produce more energy than when using a fixed head. Consequently, the generating stopping heads should be added to the list of operational parameters for inclusion in the optimisation process if desired. This would increase the number of parameters by 50% and significantly increase the execution time.

6.4   Energy storage

The option to store all the surplus energy from the barrage during periods of low demand with a dedicated facility is unfeasible - the power and energy levels of the barrage are too great. For the standard barrage design (curved barrage, 160 turbines and pumping to cycle-by-cycle tide limits), the maximum energy generated for a cycle was 16.34 GWh, the average energy per cycle was 6.87 GWh, with a peak power of 4.29 GW and an average power of 2.28 GW. Dinorwig by contrast, which is the largest pumped hydroelectric storage facility in the UK, has a capacity of 9.1 GWh and maximum power of 1.8 GW.

The problem of cyclic intermittency can be addressed through energy balancing. Tidal power is predictable, and the problem of abrupt transitions for energy balancing can be mitigated, with a small energy penalty, by slowly ramping up and down the generation. As the country transitions to a low carbon economy and electricity and usage patterns change, it should be possible to fully utilise all the surplus energy e.g., for charging electric vehicles or producing hydrogen.

6.5   Protecting the environment from change

Sea-level rise sharpens the focus keenly on to the environment. The main purpose for building a tidal range barrage is to generate large amounts of renewable, low emission energy to address the environmental damage being done through greenhouse gas emissions. At a local level, sea-level rise threatens the Bay directly, with wholescale change to the inter-tidal zone, and the loss of high-water margin habitats, in particular the saltmarsh, which is an important carbon sink. Furthermore, it poses an increased risk to the low-lying land of the fluvial valleys feeding into the Bay. Sea-level rise also poses the problem of future-proofing the design - should the number of turbines be chosen to suit today's condition or for predicted future conditions, for example. The problem is exacerbated by the level of uncertainty in the predictions.

Protecting the environment is a major driver for the development of the scheme; indeed, with the proposed changes to the legislation defining conservation designations, it may be the only defendable reason for construction. Maintaining the current tides and extents should protect the ecologically valuable habitats. Modelling showed that this can be achieved with the use of pumping, but places demands on the design and limits the energy that can be generated.

6.6   Decision making

A decision to build a barrage at Morecambe Bay must be founded on economic principles, together with a wider understand of its positive and negative consequences. The environment is crucial in this assessment and should be incorporated into the performance estimates of the barrage at all stages of the design. Although it is difficult to cost and value all aspects of a barrage, it is essential in order to reliably balance the pros and cons of building a barrage against not building one.

References

1. Northern Tidal Power Gateways Preliminary Information Memorandum. Mott MacDonald; 2017. Report No.: 381470 | 1 | C.
2. Burrows R, Yates N, Chen D, Hedges T, Holt J, Walkington I, et al. Tidal energy potential in UK waters. Proceedings of The Institution of Civil Engineers-maritime Engineering - PROC INST CIVIL ENG-MARIT ENG. 2009;162:155-64.
3. Neill SP, Angeloudis A, Robins PE, Walkington I, Ward SL, Masters I, et al. Tidal range energy resource and optimization – Past perspectives and future challenges. Renewable Energy. 2018;127:763-78.
4. Angeloudis A, Piggott M, Kramer S, Avdis A, Coles D, Christou M. Comparison of 0-D, 1-D and 2-D model capabilities for tidal range energy resource assessments2017.
5. Angeloudis A, Falconer RA. Sensitivity of tidal lagoon and barrage hydrodynamic impacts and energy outputs to operational characteristics. Renewable Energy. 2017;114:337-51.
6. Angeloudis A, Falconer RA, Bray S, Ahmadian R. Representation and operation of tidal energy impoundments in a coastal hydrodynamic model. Renewable Energy. 2016;99:1103-15.
7. Aggidis GA, Benzon DS. Operational optimisation of a tidal barrage across the Mersey estuary using 0-D modelling. Ocean Engineering. 2013;66:69-81.
8. Aggidis GA, Feather O. Tidal range turbines and generation on the Solway Firth. Renewable Energy. 2012;43:9-17.
9. Angeloudis A, Kramer SC, Avdis A, Piggott MD. Optimising tidal range power plant operation. Applied Energy. 2018;212:680-90.
10. Kontoleontos E, Weissenberger S. Annual Energy Production (AEP) optimization for tidal power plants based on Evolutionary Algorithms - Swansea Bay Tidal Power Plant AEP optimization. IOP Conference Series: Earth and Environmental Science. 2016;49:102009.
11. Harcourt F, Angeloudis A, Piggott MD. Utilising the flexible generation potential of tidal range power plants to optimise economic value. Applied Energy. 2019;237:873-84.
12. Xue J, Ahmadian R, Falconer R. Optimising the Operation of Tidal Range Schemes. Energies. 2019;12(15).
13. Xue J, Ahmadian R, Jones O. Genetic Algorithm in Tidal Range Schemes' Optimisation. Energy. 2020;200.
14. Prandle D. Simple theory for designing tidal power schemes. Advances in Water Resources - ADV WATER RESOUR. 1984;7:21-7.
15. Power sector modelling: System cost impact of renewables. Report for the National Infrastructure Commission. Aurora Energy Research Limited; 2018 22 May 2018.
16. RISING SEAS: THE ENGINEERING CHALLENGE. Institution of Mechanical Engineers; 2019 12 Nov 2019.
17. Lyth Valley: Draft Flood Investigation Report Flood Event 5th – 6th December 2015. Environment Agency / Cumbria County Council; 2017.
18. Natural England. MAGIC mapping data 2020 [Available from: https://magic.defra.gov.uk/magicmap.aspx. Accessed:12 Aug 2020.
19. Ordnace Survey. OS VectorMap District data 2019 [Available from: https://osdatahub.os.uk/downloads/open/VectorMapDistrict. Accessed:21 Oct 2019.
20. Ordnace Survey. OS Terrain 50 vector data 2019 [Available from: https://osdatahub.os.uk/downloads/open/Terrain50. Accessed:21 Oct 2019.
21. Environment Agency. LIDAR 1m DTM 2017 [Available from: https://environment.data.gov.uk/DefraDataDownload/?Mode=survey. Accessed:8 Oct 2019.
22. UKHO, cartographer ADMIRALTY Chart - 2010 Morecambe Bay. Scale 1:37,500: The UK Hydrographic Office (UKHO); 1984.
23. The National Oceanography Centre. National Tidal and Sea Level Facility 2019 [Available from: https://www.ntslf.org/data/uk-network-real-time. Accessed:5 Nov 2019.

24. MathWorks. UTide Unified Tidal Analysis and Prediction Functions 2017 [Available from: https://www.mathworks.com/matlabcentral/fileexchange/46523-utide-unified-tidal-analysis-and-prediction-functions. Accessed:19 Nov 2019.

25. Baker AC. Tidal Power: IET; 1990.

26. Libaux A. La rance - Alstrom bulb turbine and pump performance data. In: Aggidis G, editor.: EDF; 2019.

27. Nechleba M, EVANS AG, MAYER C. Hydraulic Turbines. Their Design and Equipment. (Translated from the Czech Edition by Charles Mayer and A.G. Evans.): Prague; 1957.

28. Hillairet P, Weisrock G, editors. Optimizing production from the Rance tidal power station. Proceedings 3rd international symposium on Wave, Tidal, OTEC, and Small Scale Hydropower; 1986.

29. Yates N, Walkington I, Burrows R, Wolf J. Appraising the extractable tidal energy resource of the UK's western coastal waters. Philosophical transactions Series A, Mathematical, physical, and engineering sciences. 2013;371:20120181.

30. Baker S. Morecambe Bay tidal barrage: Modes of operation and their effect on energy and value generation.: Lancaster University; 2018.

31. Hendry C. THE ROLE OF TIDAL LAGOONS, Final report 2016 6 Dec 2016.

32. Met Office Hadley Centre Climate Programme. UK Climate Projections (UKCP) 2018 [Available from: https://www.metoffice.gov.uk/research/approach/collaboration/ukcp/download-data. Accessed:28 Nov 2019.

33. Environment Agency. Indicative flood risk areas shapefiles 2019 [Available from: https://data.gov.uk/dataset/7792054a-068d-471b-8969-f53a22b0c9b2/indicative-flood-risk-areas-shapefiles. Accessed:1 Nov 2019.

34. National River Flow Archive (CEH). Peak flow data 2019 [Available from: https://nrfa.ceh.ac.uk/peak-flow-dataset. Accessed:4 Nov 2019.

35. G.B. National Grid Status 2020 [Available from: https://www.gridwatch.templar.co.uk/. Accessed:15 June 2020.

36. Department for Business, Energy & Industrial Strategy. Digest of United Kingdom Energy Statistics (DUKES). 2020 [Available from: https://www.gov.uk/government/collections/digest-of-uk-energy-statistics-dukes. Accessed:5 Sept 2020.

37. Luo X, Wang J, Dooner M, Clarke J. Overview of current development in electrical energy storage technologies and the application potential in power system operation. Applied Energy. 2015;137:511-36.

38. Environment Agency. Shoreline management plans (SMPs) 2019 [Available from: https://www.gov.uk/government/publications/shoreline-management-plans-smps. Accessed:21 Nov 2019.

Appendix A   Optimisation functions

For a full description of the program interface refer to Appendix B. The primary way to run a simulation is by pressing the Model energy button in the main Design tab. This uses all the design settings (Figure A.1), including the intercept and gradient values in the Operational Params Tab (Figure A.2).

As stated in section 2.5 the cycle-by-cycle operational parameter values were calculated from the input tide series as part of the initialisation phase. These were represented by vectors, and since all the parameters were constant during each generating phase, they appeared as a sequence of steps when plotted as a time series (see Figure A.3).

In the example shown here, cycle-by-cycle pumping was activated and the pump limits were calculated using the intercept and gradient values in the table. All of these were set to zero and the limits duly follow the individual tidal amplitudes. For the starting head parameters, where the gradients were close to one, the lines closely follow the trend of the tidal amplitudes, offset by the intercept value.

Figure A.1 Main design parameter settings – (Graphical User Interface second tab).

	Intercept	Gradient	I:increment	#I	G:increment	#G	Skew	I Scale	G Scale	Scan
Ebb start head	1.7188	0.9059	0.5000	2	0.2000	2	-0.2000	1	1	✓
Fld start head	1.6948	0.8756	0.5000	2	0.2000	2	-0.2000	1	1	✓
Ebb turbine speed	31.1780	7.2475	5	2	0.5000	2	-0.2000	1	1	✓
Fld turbine speed	31.3613	7.6069	5	2	0.5000	2	-0.2000	1	1	✓
Ebb turbine flow	0	0	0.5000	2	0.2000	2	-0.2000	1	1	
Fld turbine flow	0	0	0.5000	2	0.2000	2	-0.2000	1	1	
Ebb pump head	0	0	0.5000	2	0.2000	2	0	1	1	
Fld pump head	0	0	0.5000	2	0.2000	2	0	1	1	

Figure A.2 Operational parameter function coefficient values and scanning grid definition – (GUI tab 3).
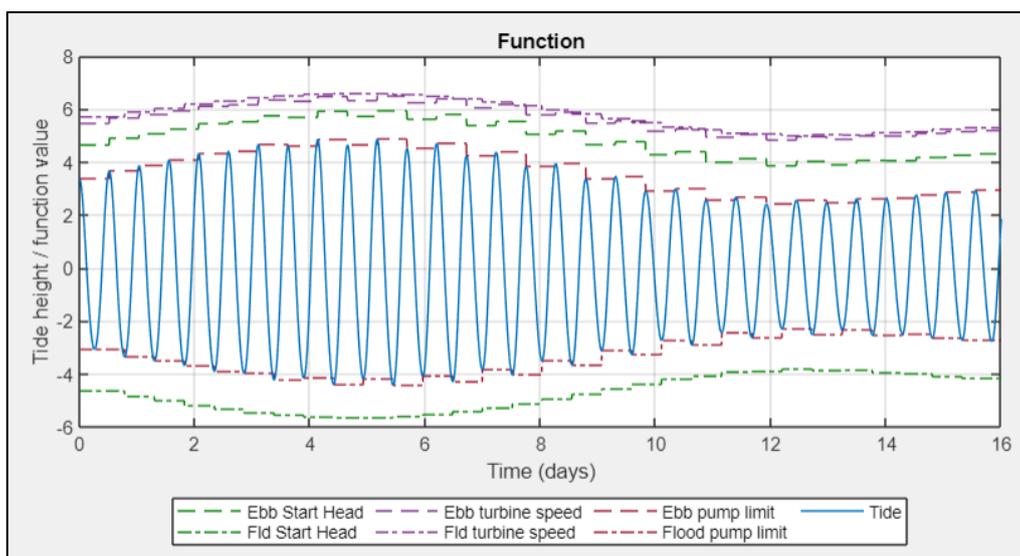
Figure A.3 Run-time operational parameter values calculated from the function definitions for the input tidal series.

The operational parameters in the table are activated or deactivated to reflect the design settings, i.e. if pumping is not selected the pump head parameters are greyed out. As described in the turbine performance section 2.3, the turbine could be operated in two different ways, constant speed or constant Q_{11}. These are listed separately in the operational parameter list and only the active set, as selected by the radio buttons on the main design panel of the GUI, were enabled. In the case of double regulation, where the turbine speed is coupled to the National Grid supply frequency, the speed was fixed for all cycles for both ebb and flood generation. In this case, the gradients were set to zero and the ebb and flood intercept values were forced to be the same.

Optimisation was performed from the Operational Params tab. Figure A.4 shows the options available and a description of their behaviour can be found in Table A-1.
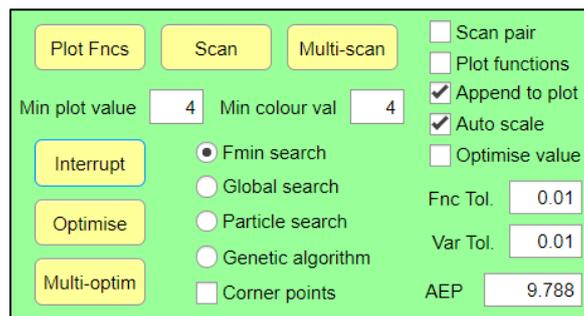
Figure A.4 Operational parameter optimisation options.

	Input	Function	Comment	
Buttons	Plot Fncs	Calculate and plot current functions		
	Scan	Scan a single parameter (redundant)	The parameter(s) selected in the table using the checkboxes	
	Multi-Scan	Scan multiple parameters		
	Optimise	Optimise a single function (redundant)		
	Multi-Optim	Optimise multiple parameters		
	Interupt	Interrupt optimisation		
Radio Butt.	Fmin search	Use *fminsearch* optimistion function	Local search from table intercept/gradient point	
	Global search	Use *global* optimistion function	Generates own bounded start points	Global search
	Partical search	Use *partical* optimistion function	Seeded with points defined using the scan range setuo	
	Genetic algorithm	Use *GA* optimistion function		
Checkboxes	Scan pair	Pair ebb/flood params. Force to be the same.	Applicable for lagoons.	
	Plot functions	Plot functions after each iteration		
	Append to plot	Do not clear plot at the start of the optim/scan		
	Auto scale	Set intercept and gradient scale factors	Scales parameter values to ≈ 1	
	Optimise value	Use value scalled AEP		
	Corner points	Use corner points only to seed optimisation.	Otherwise uses all scan defined points	
Text box	Min plot value	Sets the minimum Z-axis value	Can be changed during optimisation	
	Min colour val	Sets the minimum Z-axis colour scale value		
	Function tolerance	Convergence tolerance of the function value		
	Variable tolerance	Convergence tolerance of the parameters values	Uses the scaled values	
	AEP	The output AEP or value AEP in TWh		

Table A-1 Description of the operational parameter optimisation options.

There are two approaches available for finding the optimum intercept and gradient values: one was to scan over a range of values, the other was to use a cost function minimisation scheme. For the latter there were four options, one performed a local search, the others a global search.

The scan ranges were defined in the operational parameter table in the five columns after the Gradient column. These specify the number of increments and the increment values for both intercept and gradient about the central point, where the central point was the intercept and gradient specified in columns one and two. A value of 2 for the number of increments results in a range of 5 points. In

addition, a skew could be specified which modified the gradient values away from the central point. The value represented how much the gradient value was shifted as a proportion of the intercept offset (from the central point). This was introduced to allow a limited number of grid of points to better cover the general trend; invariably a higher gradient will have a correspondingly lower intercept and the skew value will be negative.

For multiple parameters the number of scanned scenarios is the product of the number of points of the individual parameters. Figure A.5 is a plot generated during a scanning run. Each time an AEP value was calculated the set of points for the variables were plotted with the z-value and colour corresponding to the AEP value, and linked by a grey line. After all the values had been computed, the scenario with the highest AEP value was uploaded to the intercept and gradient columns and the set of points re-plotted and connected by a black line.
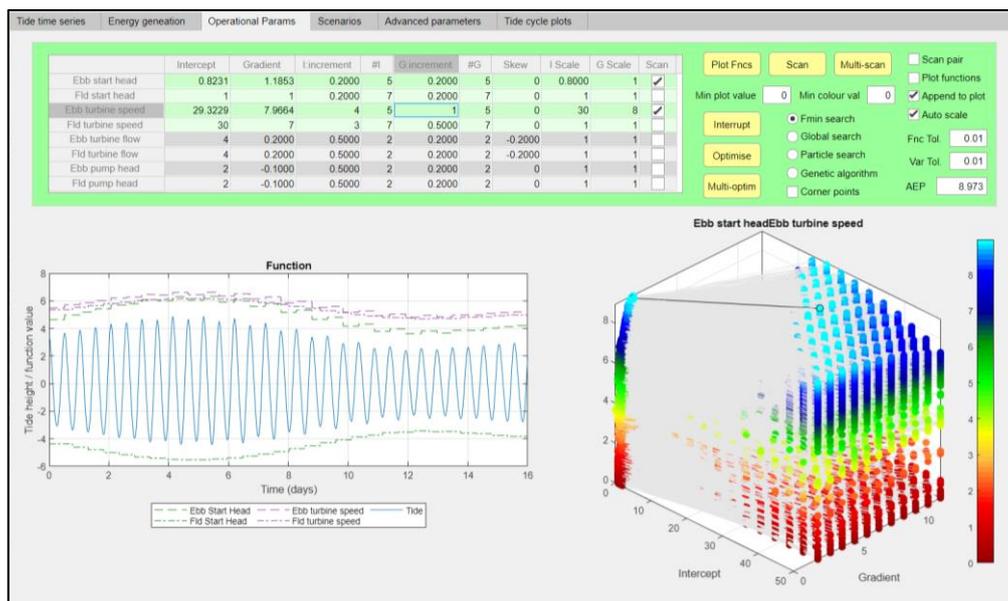
Figure A.5 Results from a scanning search.

The column of checkboxes on the right of the table were used to select which parameters were to be optimised or scanned. Each time the model was updated, either by looping through the scanned ranges, or automatically via the optimisation algorithm, the new values were written to the intercept and gradient values in the table. These were read as part of the simulation run and the execution sequence was identical, whether run from the Design tab or from the Operational Params tab.

The reason for the scale factors was to present the optimisation engine with parameters of a similar value. The values are calculated by dividing the intercept and gradient values by the scale factor to give numbers around one. The optimisation algorithms perturb the parameters using a set limit that was common to all the variables. Similarly, the tolerance used to determine when a solution had been found was the same for all variables. The scaling significantly improves the robustness of the optimisation especially where the values became very small. The scale factors can be entered manually in the table or will be calculated automatically at the start of the scan/optimisation if the Auto scale checkbox was ticked. The scale factors calculated were the intercept and gradient values rounded to the nearest significant figure.

The local search optimisation was good for tracking the global minimum (maximum AEP) of small incremental changes in the design scenario; for example, when performing sensitivity analysis. This assumed the global minimum had been found in the first place. The global optimisation options were

provided to help find the global minimum when larger design changes were made. The scanning option enabled the parameter space to be visualised, providing a general QC tool and helping set the bounds for the global optimisation. The Particle search and Genetic algorithm methods require a set of seed points, which are generated from the set of range definitions for the scanned points in the same way. If the Corner points checkbox was selected it would limit the points to the central point and corner points only. The upper and lower bounds for the Global search method are also derived from the scanned points definitions. Note, any skew will effectively increase the gradient ranges.

Figure A.6 and Figure A.7 show the result from using a local search and global search respectively. The local search converged after relatively few iterations, whereas the global search could take several thousand.
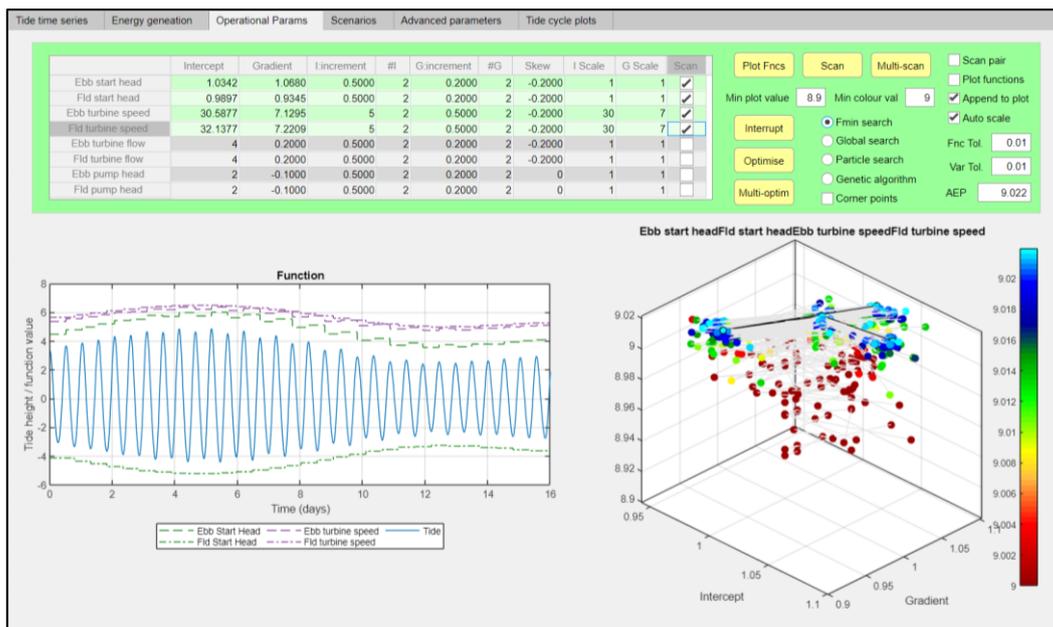
Figure A.6 Results from a local optimisation search (fminsearch).
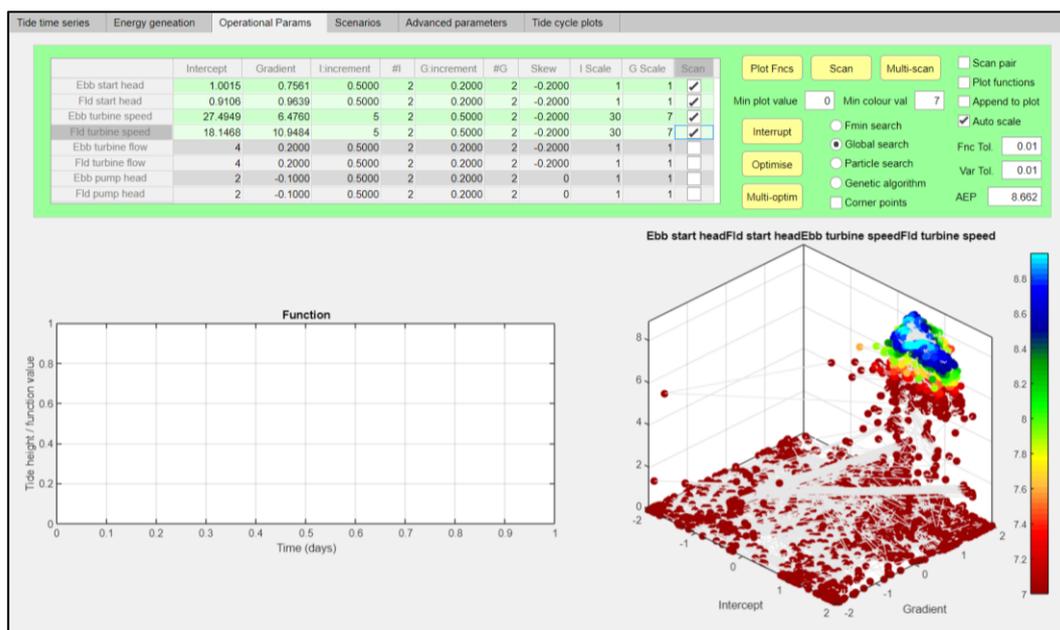
Figure A.7 Results from a global optimisation search (globalsearch).

All the optimisation algorithms struggle to find the precise minimum point and could be trapped in a very local minimum created by the roughness of the function. Figure A.8 shows the scanned values for the Ebb starting head parameter using different time sampling rates. The three surfaces from top to bottom are for 0.1, 0.05 and 0.01 hour time steps. Not only did the time step affect how smooth the function was, it also affected the magnitude of the AEP value. Where the slope was steep, away from the minimum point, the roughness was effectively irrelevant. It was when it approaches the minimum and the slope dropped that the roughness could capture the minimisation path.

Increasing the sample rate by a factor of 10 increased the runtime by a factor of 10. Even with a local incremental search, the number of iterations was typically on the order of 60-120. Increasing the run time of each iteration from 2 seconds to 20 seconds could increase the overall time from 3 minutes to 30 minutes.

In practice, a combination of scanning at different levels of refinement (increments) and local optimisation were enough to find the global minimum with a reasonable level of confidence. Small errors due to the roughness could be problematic and perpetuate to some extent through a sensitivity sequence of designs. This was evident if the sensitivity trend was inconsistent or there were anomalous points.
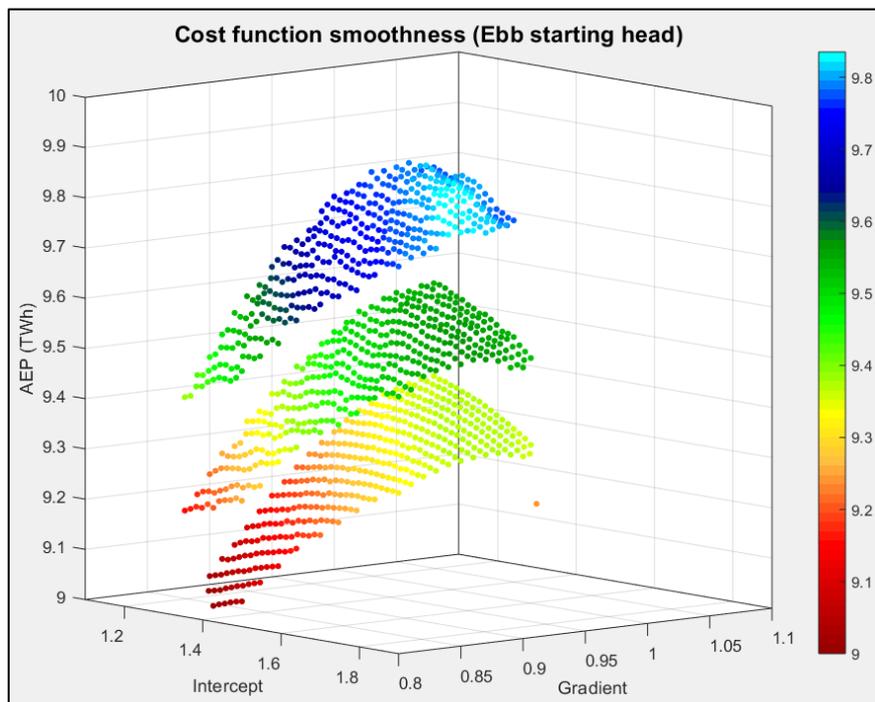
Figure A.8 The AEP values for a range of scanned ebb starting head coefficients with different time series sample intervals: 0.1 hour; 0.05 hour and 0.01 hour (top to bottom). The smoothness of the function improves with shorter time steps.

Appendix B   Graphical User Interface (GUI)

The simulation program was set up and run via a Graphical User Interface (GUI). This was written using App Designer in MATLAB. The interface consists of 6 pages in the form of tabs, which are described in the following sections. The final section details a typical execution sequence of the program.

A key feature of the simulation modelling is the requirement to run through an optimisation step for each new design configuration. This process creates design specific function coefficients for the operational parameter values. To audit this, a feature to store and load these design scenarios was added to the program. The complete definition for a simulation is spread over several tabs, namely the input tide series, the auxiliary parameters (e.g. maximum turbine power rating), and the main design parameters. The tide series definition and auxiliary parameters are defined and stored in their own configuration tables and are referenced from the main design scenario table.

B.1.   Tide series modelling

The tide series tab (Figure B.1) consists of an input parameter panel (blue), a graphical display and the tide scenario table with editing command buttons. The tidal parameters define the type of tide model, the start date/time, the sample interval and the duration. For the tidal model there is a choice between the 10 tidal constituent model at Heysham (see section 2.2), the M2 only (primary semidiurnal lunar tide) tidal constituent at Heysham, or user-defined M2 amplitude. The series is generated and displayed when the generate button is pressed. This also clears a flag that is interrogated before energy simulation is allowed.

Whenever any of the tidal parameters are changed, the whole setup is compared with those stored in the table, and if they match one of the entries, the number is written in the Tide model box in the entry panel. If there are no matches, zero is written.

There are four functions related to the scenario table: *Delete*; *Overwrite*; *Store model* and *Upload model*. The delete, overwrite and upload operations all require the row of an existing scenario to be selected using the checkbox in the rightmost column. *Store model* will save the current panel settings in a new row appended to the table and enter the new model number in the Tide model text box. *Upload model* will populate the panel settings with the selected table entry.

B.2.   Barrage design parameters and output display

The design parameter tab (Figure B.2) consists of the design setup panel (green), results summary panel (blue), plot options panel (purple) and the time series graph display and associated slider controls. Refer to Table 2-1 in section 2.7 for a description of the model parameters.

The results summary panel displays the following calculated values:
- **Max. theoretical TWh** – this assumes the whole reservoir volume is discharged instantly with the starting head at the limits of the tidal range, summed over all cycles and scaled to a full year
- **This model TWh** – the energy generated over the input tide series
- **Overall efficiency %** - the generated AEP / max theoretical AEP (Annual Energy Production)
- **AEP equiv. TWh** – the generated energy scaled to a full year
- **APE equiv. TWh** – Annual Potential Energy (APE) – ignores any efficiency losses, i.e. the sum of all the potential energy gains/losses per time step (useful to compare to the AEP to get the overall efficiency of the equipment)
- **Value scaled TWh** – uses the electricity demand profile (see section 4.2) to scale the generated energy on a sample-by-sample basis (at periods of high demand the scale factor was above one, during low demand it was less than one).

- **Annual CO2e (Mt)** – the nominal annual CO_2 emissions that would be produced to generate the simulated AEP using today's electricity supply mix in Mtonnes.

Several different data series display options are provided via the *Plot options* panel. The radio buttons select the data series to display, and the checkboxes can be selected to add additional information. *Limits* refers to the operational parameter values.
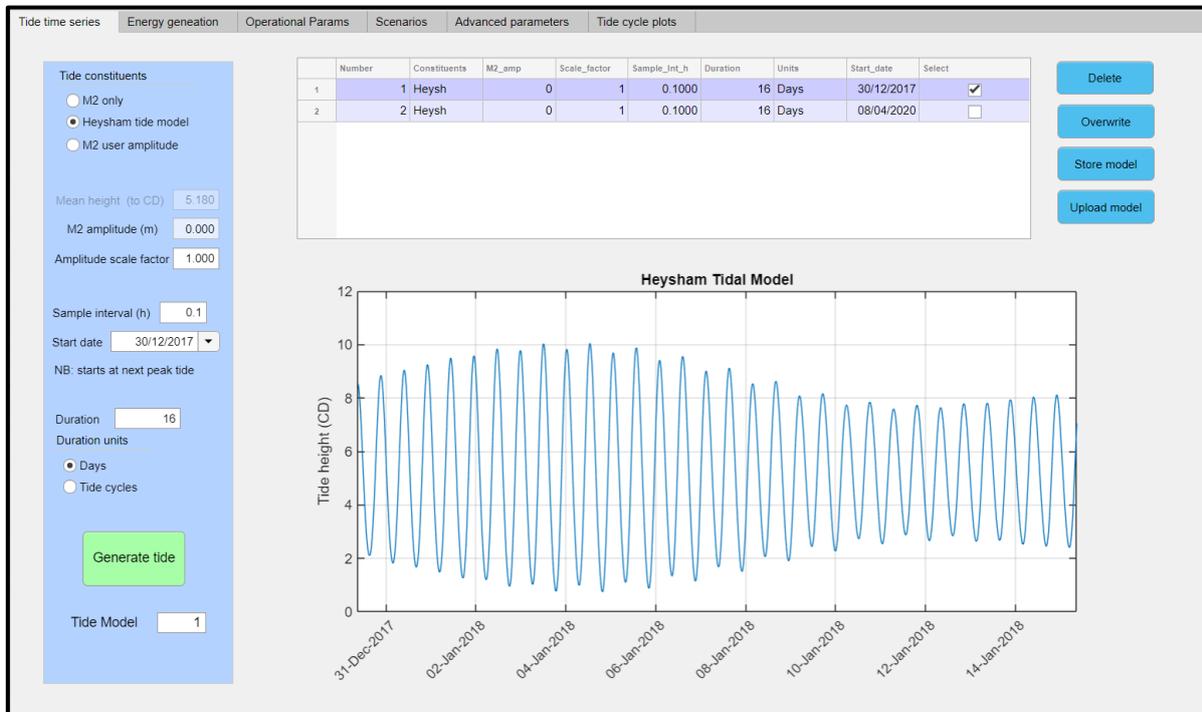
Figure B.1 Tidal series modelling tab.
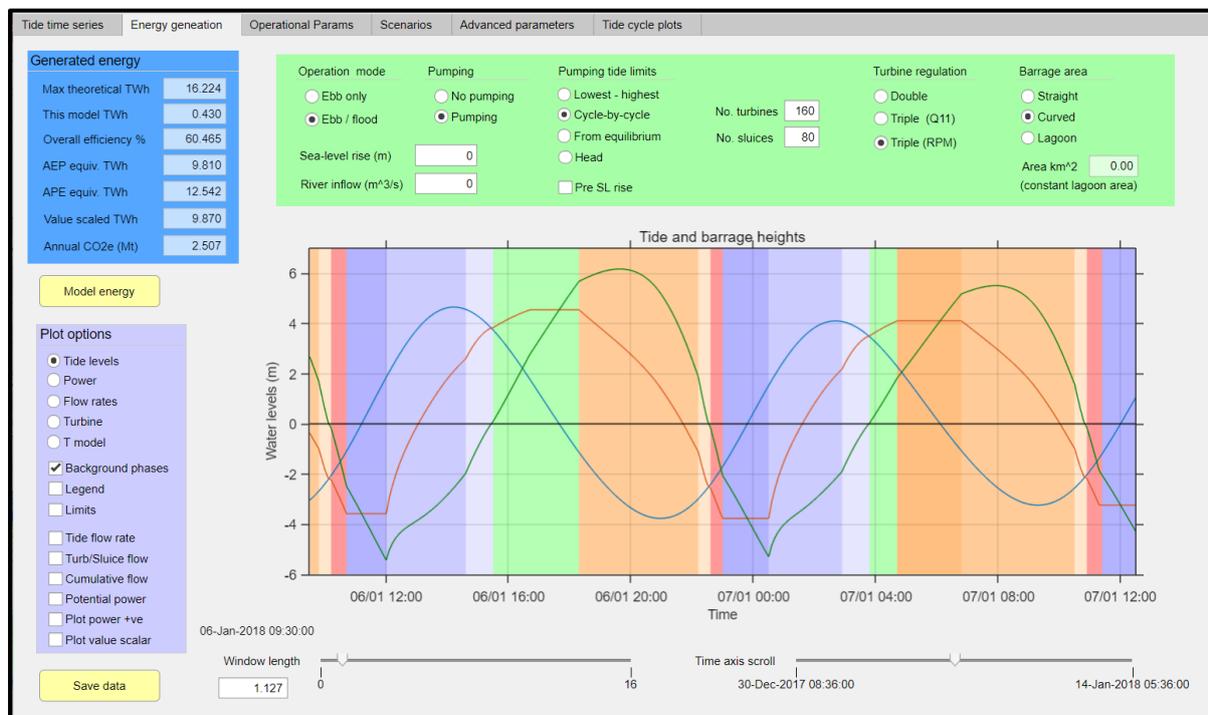
Figure B.2 Barrage design parameters and output display tab.

B.3. Auxiliary parameters

The auxiliary parameters tab is shown in Figure B.3. For parameter details refer to section 2.7. Similar to the tide series tab, the auxiliary setups can be stored in a table and referenced as part of a design scenario. The same functions are available and the same checks are performed. It is from this tab that the scenario tables can be saved or uploaded from file (purple panel). When saving to an Excel file, the file is checked to see if it already exists and a message is displayed to confirm if it is to be overwritten.
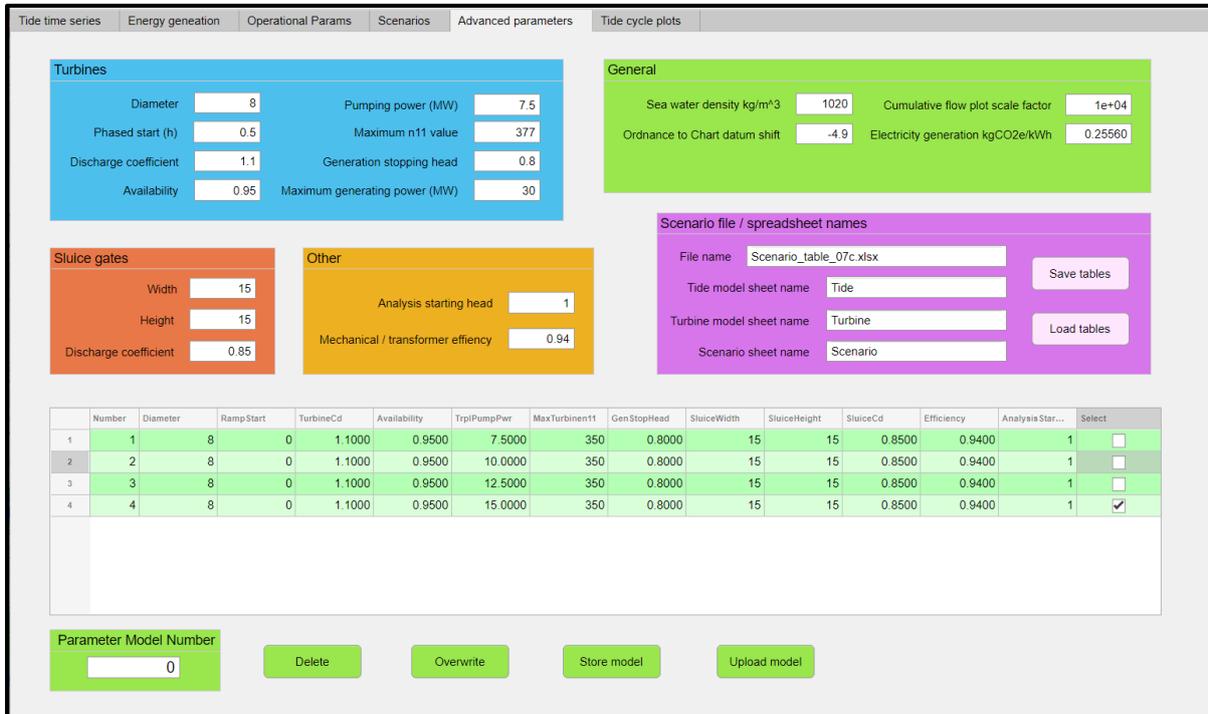
Figure B.3 Auxiliary parameters tab.
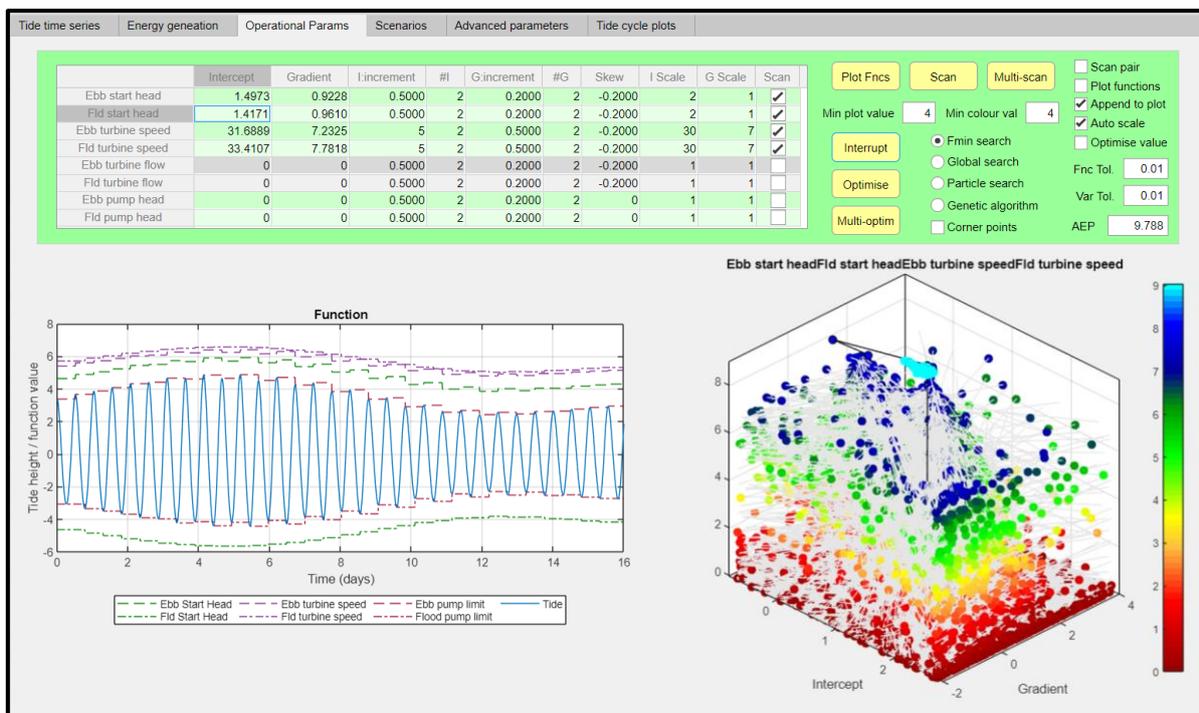
Figure B.4 Operational parameter selection and optimisation tab

B.4. Operational parameter selection and optimisation

The Operational parameters tab is shown in Figure B.4. Refer to Appendix A for a detailed description.

B.5. Scenario setup and recording

The design scenario tab (Figure B.5 comprises the table of scenarios and table edit functions. It shares the same four functions as the other setup tables and a fifth that checks if the current design matches one in the table. Unlike the tide and auxiliary tables, the scenario table is not checked when a parameter value is changed and there is no information box that displays the scenario number if the current settings correspond to a table entry.

Tide time series | Energy geneation | Operational Params | Scenarios | Advanced parameters | Tide cycle plots

	Number	AEP (TWh)	Tide series	Turb/sluice model	Mode	Pumping	Pump-limit mode	Pre Sea-level rise	Sea-Level rise (m)	River inflow (m³/s)	# Turbines	# Sluices	Barrage shape	Lagoon Area (km²)	Turbine regulation	Gen start head Ebb a_0	Gen start head Ebb a_1	Gen start head Flood a_0	Gen start head Flood a_1	Turbine speed / Q11 Ebb a_0	Turbine speed / Q11 Ebb a_1	Turbine speed / Q11 Flood a_0	Turbine speed / Q11 Flood a_1	Pump stop head Ebb a_0	Pump stop head Ebb a_1	Pump stop head Flood a_0	Pump stop head Flood a_1	Select
14	99	9.8307	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	1.4183	1.0356	1.4449	0.9201	31.8...	7.2137	31.7...	7.1708	0	0	0	0	☐
15	15	9.8290	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	1.6463	0.9300	1.6238	0.8935	31.5...	7.0147	32.7...	7.2953	0	0	0	0	☐
16	16	9.8309	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	1.8135	0.8720	1.7481	0.8665	31.2...	7.0600	32.4...	7.4413	0	0	0	0	☐
17	17	9.8129	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	1.9979	0.8003	2.0563	0.8175	31.2...	7.0580	32.5...	7.4392	0	0	0	0	☐
18	18	9.7664	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	0.8192	1.1704	0.7967	1.2039	31.8...	7.1101	32.7...	7.6388	0	0	0	0	☐
19	19	9.8301	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	1.7172	0.9026	1.7092	0.8717	31.2...	7.0355	32.0...	7.4469	0	0	0	0	☑
20	20	9.7890	1	1	Dual	Y	C-by-C	N	0	0	160	80	Curved	0	Trpl rpm	1.7035	0.9065	1.7255	0.8800	39.9...	5.0027	40.3...	5.0476	0	0	0	0	☐
21	21	9.9665	1	1	Dual	Y	C-by-C	N	0.2000	0	160	80	Curved	0	Trpl rpm	1.8041	0.8833	1.7927	0.8747	31.2...	7.1190	32.5...	7.5186	0	0	0	0	☐
22	22	10.0972	1	1	Dual	Y	C-by-C	N	0.4000	0	160	80	Curved	0	Trpl rpm	1.8387	0.8743	1.8144	0.8807	31.5...	7.2929	32.4...	8.1364	0	0	0	0	☐
23	23	10.2414	1	1	Dual	Y	C-by-C	N	0.6000	0	160	80	Curved	0	Trpl rpm	1.7944	0.8602	1.7603	0.8744	32.2...	7.2198	32.8...	8.0468	0	0	0	0	☐
24	24	10.3433	1	1	Dual	Y	C-by-C	N	0.8000	0	160	80	Curved	0	Trpl rpm	1.8089	0.8515	1.7510	0.8769	32.5...	7.2873	33.0...	8.1973	0	0	0	0	☐
25	25	10.4277	1	1	Dual	Y	C-by-C	N	1.0000	0	160	80	Curved	0	Trpl rpm	1.8055	0.8340	1.7385	0.8851	32.5...	7.1548	33.8...	8.4139	0	0	0	0	☐
26	26	10.5265	1	1	Dual	Y	C-by-C	N	1.2000	0	160	80	Curved	0	Trpl rpm	1.8592	0.7649	1.8893	0.8444	31.0...	7.5813	32.8...	8.6926	0	0	0	0	☐
27	27	10.6176	1	1	Dual	Y	C-by-C	N	1.4000	0	160	80	Curved	0	Trpl rpm	1.8240	0.7620	1.8334	0.8479	31.6...	7.6767	32.5...	8.7588	0	0	0	0	☐
28	28	10.6837	1	1	Dual	Y	C-by-C	N	1.6000	0	160	80	Curved	0	Trpl rpm	1.8568	0.7523	1.7993	0.8132	31.7...	7.7878	32.7...	8.9289	0	0	0	0	☐
29	29	9.8834	1	1	Dual	Y	C-by-C	Y	0.1000	0	160	80	Curved	0	Trpl rpm	1.7436	0.9013	1.7297	0.8909	31.7...	7.0530	31.3...	7.3862	0	0	0	0	☐
30	30	9.9228	1	1	Dual	Y	C-by-C	Y	0.2000	0	160	80	Curved	0	Trpl rpm	1.7255	0.8957	1.7639	0.8762	32.2...	7.0828	31.3...	7.5244	0	0	0	0	☐
31	31	10.0306	1	1	Dual	Y	C-by-C	Y	0.4000	0	160	80	Curved	0	Trpl rpm	1.7547	0.8430	1.8843	0.9346	32.4...	6.8250	30.6...	7.6562	0	0	0	0	☐
32	32	10.1006	1	1	Dual	Y	C-by-C	Y	0.6000	0	160	80	Curved	0	Trpl rpm	1.8004	0.8409	1.8697	0.9325	32.7...	6.8913	30.9...	7.7305	0	0	0	0	☐
33	33	10.1935	1	1	Dual	Y	C-by-C	Y	0.8000	0	160	80	Curved	0	Trpl rpm	1.7169	0.8152	1.9553	0.9469	32.5...	6.9462	31.1...	7.8828	0	0	0	0	☐
34	34	10.2765	1	1	Dual	Y	C-by-C	Y	1.0000	0	160	80	Curved	0	Trpl rpm	1.6862	0.7704	1.9769	0.9670	32.5...	7.0834	32.1...	7.9015	0	0	0	0	☐

Delete | Overwrite | Store model | Upload model | Check existing

Figure B.5 Scenario setup and recording tab.

When a scenario is saved, the main design settings from the design tab, the operational parameter function definitions, and the number of the active tide series and auxiliary setup configurations are written to the scenario table. If the tide series or auxiliary setup definitions are not already stored in the respective tables, an entry is automatically entered. The current AEP value as displayed in the Design tab is also saved as part of the scenario. It is necessary to run the simulation from the design tab after optimisation to update the AEP value before saving. Note: it is distinctly possible that different optimisation runs will produce different intercept and gradient values and different AEP values.

B.6. Tide cycle plots

Two map types were created to provide a means to assess the impact of the current run simulation on the inter-tidal zone (see section 2.8). These are available in the Tide Cycle plots tab (Figure B.6). One is a count of the number of tide cycles in a year, the other is the cumulative time exposed/submerged (in days) over a year.

For time series not exactly a year in length, the numbers are scaled accordingly. Note, generating maps from short time series can result in poor gradation and large errors. It was found to be preferable to simulate the barrage operation over a full year.
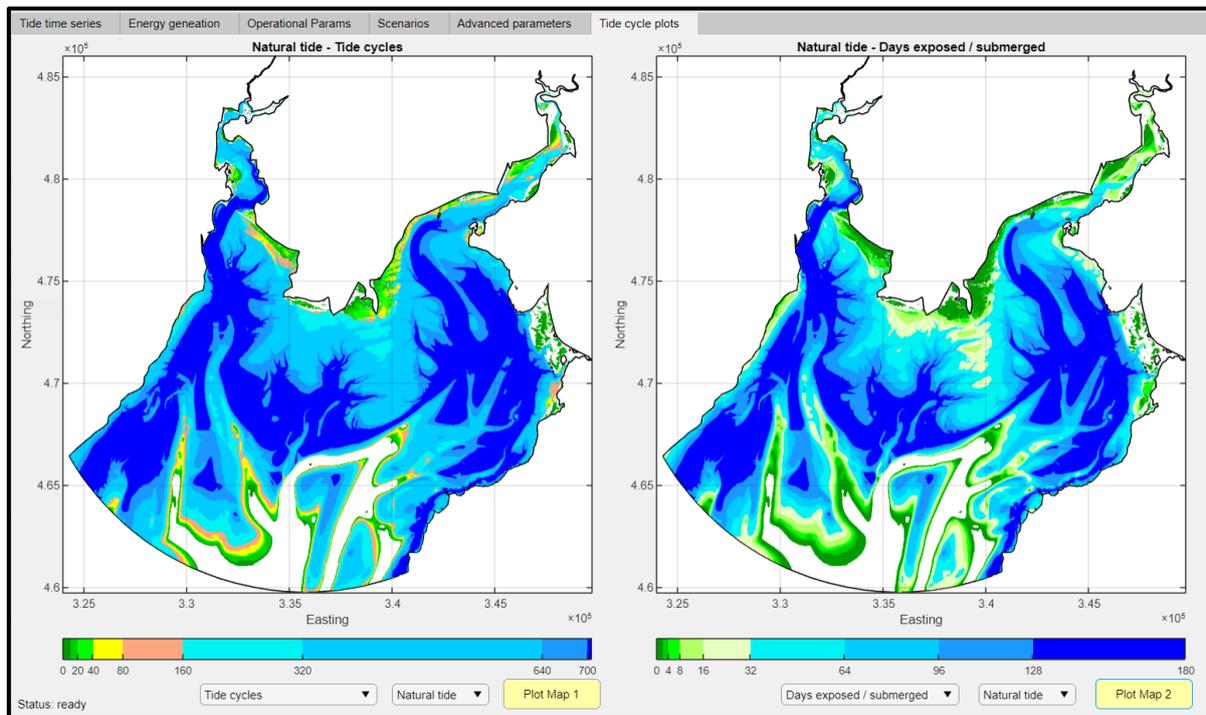
Figure B.6 Inter-tidal QC plots tab.

B.7. Program execution

The GUI consists of 6 pages in the form of tabs. A typical run will consist of the following program setup and execution (consult the relevant section above for more details):

- Set up and generate the tidal series (tab 1). Optionally store the setup in the tide table.
- Modify any of the auxiliary parameters if required – these relate to the turbine and sluice gate parameters, amongst other things, and are deemed to form fixed "building blocks" for the barrage design (tab 5). Optionally store the setup in the auxiliary table.
- Set the barrage design parameters (tab 2).
- Check/set the operational parameter linear function parameters (tab 3). Typically, these will be loaded in from the result of an optimisation of a previous scenario.
- Run a simulation and display the results (tab 2)
- Choose the method and perform one or more runs of optimisation (tab 3).
- Rerun and display the result (tab 2). Optionally export the time series to a MATLAB data file.
- Store the design and result in the main scenario table (tab 4).
- Optionally, save the scenario and setup tables to an excel file (tab 5).
- Optionally, generate inter-tidal maps (tab 6).

If the next design is very similar, changing from 160 to 170 turbines or from 0.4 to 0.5 m sea-level rise for example, and all the other parameters remain the same, there is no need to regenerate the tide series and the execution sequence is more straightforward:

- Make the design change and run to display and QC the result (tab 2).
- Choose a local optimisation search method and run (tab 3).
- Rerun and display the result (tab 2). Optionally export the time series to a MATLAB data file.
- Store the design and result in the main scenario table (tab 4).

Sensitivity analysis can be performed by repeating this sequence for a range of values of an individual design parameter. The final step is to save the complete scenario table to Excel, which can be loaded into MATLAB for analysis and plotted as required.

Appendix C   Sample model simulations: input and output

Table C-1 details the design settings, optimised operational function coefficients and calculated energy values for a sample set of design scenarios.

	Parameter		Units	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Design Tab	Reservoir			Curved	Curved	Curved	Curved
	Number of turbines			160	160	160	160
	Number of sluice gates			80	80	80	80
	Operation mode			Dual	Dual	Dual	Dual
	Turbine regulation			Triple(rpm)	Triple(rpm)	Triple(rpm)	Triple(rpm)
	Pumping			Yes	Yes	Yes	Yes
	Pumping limit mode			Cycle-by-cycle	Cycle-by-cycle	Cycle-by-cycle	Force C-by-C
	Sea-level rise		m	0	1	1	1
	Pre Sea-level rise			No	No	Yes	Yes
	River inflow		m³/s	0	0	0	0
Operational	Ebb generation starting head	Intercept	m	1.7172	1.80553	1.61268	1.64341
		Gradient	m/m	0.9026	0.83404	0.74873	0.71178
	Flood generation starting head	Intercept	m	1.7092	1.73846	1.98431	2.12503
		Gradient	m/m	0.8717	0.88508	1.05305	1.03096
	Ebb turbine speed	Intercept	rpm	31.2789	32.55362	29.45867	30.24863
		Gradient	rpm/m	7.0355	7.15476	7.60159	7.34445
	Flood turbine speed	Intercept	rpm	32.0036	33.85885	32.17215	33.78512
		Gradient	rpm/m	7.4469	8.41392	7.90005	7.36305
	Ebb turbine Q11	Intercept	m³/s				
		Gradient	m³/s/m				
	Flood turbine Q11	Intercept	m³/s				
		Gradient	m³/s/m				
	Ebb pump stopping head	Intercept	m	0.0000	0.00000	0.00000	0.00000
		Gradient	m/m	0.0000	0.00000	0.00000	0.00000
	Flood pump stopping head	Intercept	m	0.0000	0.00000	0.00000	0.00000
		Gradient	m/m	0.0000	0.00000	0.00000	0.00000
Tide series	Tide model			Heysham model	Heysham model	Heysham model	Heysham model
	Sample interval		hours	0.1	0.1	0.1	0.1
	Start date			30/12/2017 00:00	30/12/2017 00:00	30/12/2017 00:00	30/12/2017 00:00
	Duration		days	16 days	16 days	16 days	16 days
Auxiliary	Turbine diameter		m	8 m	8 m	8 m	8 m
	Turbine coefficient of discharge			1.1	1.1	1.1	1.1
	Turbine availability			0.95	0.95	0.95	0.95
	Turbine phased start		hours	0	0	0	0
	Turbine power rating		MW	30	30	30	30
	Pumping maximum power		MW	7.5	7.5	7.5	7.5
	Sluice width		m	15	15	15	15
	Sluice gate height		m	15	15	15	15
	Sluice coefficient of discharge			0.85	0.85	0.85	0.85
	Mechanical/transformer efficiency			0.94	0.94	0.94	0.94
	Maximum n11		rpm	377	377	377	377
	Generation stopping head		m	1	1	1	1
	Analysis starting head		m	1	1	1	1
	Sea water density		kg/m³	1020	1020	1020	1020
	Ordnance-Chart datum shift		m	-4.9	-4.9	-4.9	-4.9
	Electricity generation emissions		kgCO2/kWh	Value	Value	Value	Value
	Pump starting head		m	0.8	0.8	0.8	0.8
Output	Max. theoretical energy		TWh	16.224	19.135	19.135	19.135
	Energy		TWh	0.431	0.457	0.44	0.439
	Annual Energy Production		TWh	9.83	10.428	10.041	10.011
	Potential Energy		TWh	12.619	13.906	13.02	12.862

Table C-1 Sample model simulations detailing the input design settings, optimised operational function coefficients and the energy output.