Applied Modelling and Computation Group
Department of Earth Science and Engineering
Imperial College London

Imperial College London

The modelling and design optimisation of tidal stream turbine arrays

Submitted in part fulfilment of the requirements for the degree of

Doctor of Philosophy in Computational Continuum Mechanics

of Imperial College London

Author:
David M. Culley

Supervisors:
Dr Simon Funke
Prof. Matthew Piggott

Declaration

I herewith certify that all material in this dissertation which is not my own work has been properly acknowledged

David M. Culley

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

Abstract

The modelling and design optimisation
of tidal stream turbine arrays
A thesis submitted for the degree of
Doctor of Philosophy
David Culley

Tidal stream turbines are devices that convert the momentum of tidally induced marine currents to electricity. Individual turbines typically have a rated power of up to a few megawatts, and thus power may be extracted on an industrial scale by installing arrays of dozens to hundreds of such devices. Accurately modelling such large arrays presents a formidable computational challenge, as the length scale of the problem varies from tenths of a meter (the scale of an individual turbine blade) to several kilometres (the scale of an array) to hundreds of kilometres (the scale of the ocean currents that govern the flow). In addition to this, the sites under consideration are chosen for their highly energetic, advection-dominated flow regimes. Since the power extracted by a turbine is proportional to the cube of the flow speed into which it is placed, the interaction between the turbines and the flow is both highly coupled and highly non-linear.

'Array design' concerns determining the *array size* – the number of turbines comprising the array – and the *micro-siting* – the locations of the turbines within the turbine site that has been leased by the array developer. For large arrays in real-world domains the micro-siting designer is faced with hundreds or thousands of control parameters that impact the farm performance through hydrodynamic interactions. Non-hydrodynamic factors such as technical, financial and practical engineering constraints can also have a significant impact upon the design decision.

Thus, going beyond modelling such arrays to *optimising* their design adds an additional dimension of complexity to the problem. Manual optimisation over such large parameter spaces with complex nonlinear constraints is difficult, and there is little practical industrial experience to guide designs. As such the array design problem requires the deployment of specially designed computational tools which can balance the need for accurate evaluation of the hydrodynamics of the interactions between the flow and the array with a suitable optimisation strategy to achieve an efficient and scalable array design paradigm. This thesis documents work to develop and expand the utility of one such software tool; *OpenTidalFarm*,

with the goal of developing a practical approach to array design optimisation and constructing a framework for the inclusion of wide ranging and disparate design criteria.

OpenTidalFarm solves the Shallow-Water and associated adjoint equations to evaluate the power extracted by the array and the sensitivity of that power to the posisitions of the turbines. This work is aimed demonstrating that industrial sized turbine arrays may be designed within a broad range of constraints and optimised for complex and occasionally conflicting objectives. This is achieved through exploration of a collection of optimisation approaches including mixed fidelity hybrid global/local schemes for turbine layout, variable fidelity Bayesian optimisation for array size and integrating integer, routing optimisation for inclusion of location-based costs.

The first chapter of this thesis introduces the engineering and practical aspects of the problem, while in the second chapter the numerical methods and optimisation techniques through which the problem can be tackled are presented. A family of alternate modelling and optimisation approaches are developed, tested and compared in Chapter 4, to understand how alternate approaches to the problem can be formulated by limiting the complexity of the problem, analysis or optimisation, and how this may successfully be done to suit the designer's available computational budget. In Chapter 5, a framework is presented in which costs over the lifespan of the array may be modelled and considered as part of the design optimisation process. Idealised scenarios are used to demonstrate the effects of a financial-return optimising design approach as contrasted with a power maximisation approach. The impact of non-hydrodynamic considerations on array design is shown to be potentially high. In Chapter 6 the broader question of array sizing is presented and a surrogate model is used to efficiently determine the optimal number of turbines to install on a given site based, once again, on a simple financial model. It is demonstrated that lower turbine densities on sites (between 20 % and 40 % of the maximum density the site could support) maximises the developer's profit on the installation and prevents diminishing returns on investment of additional turbines.

Acknowledgements

I must first and foremost acknowledge my supervisors Matt Piggott and Simon Funke. The former for the freedom, opportunities and guidance that facilitated this work and the latter for fielding even my most elementary questions patiently and without scorn. In addition, I thank my mother who has always been an inspiration and unwavering supporter.

I would also like, in advance, to thank my examiners for the time, effort and thought they have offered in both reading this work and giving me feedback.

Hofstadter's Law:
It always takes longer than you expect,
even when you take into account Hofstadter's Law.

To have reached this stage of my PhD thesis I must thank both my family and friends – those who have inspired this thesis and who, in thanks for which, I will never make read it.

Dave Culley

London, 2016

'Technology', as the computer scientist Bran Ferren memorably defined it, is 'stuff that doesn't work yet'. We no longer think of chairs as technology, we just think of them as chairs. But there was a time when we hadn't worked out how many legs chairs should have, how tall they should be, and they would often 'crash' when we tried to use them. Before long, computers will be as trivial and plentiful as chairs (and a couple of decades or so after that, as sheets of paper or grains of sand) and we will cease to be aware of the things. In fact I'm sure we will look back on this last decade and wonder how we could ever have mistaken what we were doing with them for 'productivity'.

From the News Review section of The Sunday Times on August 29th 1999.

Douglas Adams

Contents

1 Introduction **1**

1.1 Introduction . . . 1

1.2 A primer on tidal stream energy . . . 2

1.2.1 Tides . . . 3

1.2.2 Tidal stream turbines . . . 4

1.2.3 Tidal stream turbine arrays . . . 5

1.2.4 The tidal stream resource . . . 6

1.3 The hydrodynamics of tidal stream turbines . . . 7

1.3.1 Wakes . . . 8

1.3.2 Blockage . . . 13

1.3.3 Power extraction . . . 16

1.4 Motivation and scope . . . 17

1.5 Research Objectives . . . 21

1.5.1 Summary of research objectives . . . 22

1.6 Thesis summary and authorship of work presented . . . 23

1.6.1 List of the Author's publications . . . 24

2 Models and methods for array design and analysis **26**

2.1 A family of models . . . 26

2.1.1 Classifying complexity . . . 28

2.2 Hydrodynamic modelling for tidal stream turbine applications 33
2.2.1 Device scale models: From individual blades to individual turbines 34
2.2.2 Array scale models: From idealised- to more realistic domains 35
2.2.3 Regional scale models: Capturing turbines at large scale 38
2.3 Micro-siting design 40
2.3.1 Wind turbine array optimisation 42
2.3.2 From wind array modelling to tidal array modelling 45
2.3.3 Approaches to tidal array micro-siting design 45

3 OpenTidalFarm 52

3.1 The finite element method 52
3.1.1 Numerical methods for partial differential equations 52
3.1.2 The finite element method 54
3.2 OpenTidalFarm 57
3.2.1 Overview of OpenTidalFarm 57
3.2.2 Modelling turbine arrays 58
3.2.3 Turbine array optimisation 62
3.2.4 The adjoint approach 62
3.2.5 Code architecture 65
3.2.6 Problem work-flow 66
3.2.7 OpenTidalFarm optimisation algorithm 67
3.3 Summary 68

4 Comparing appoaches to tidal turbine array optimisation 70

4.1 Introduction 71
4.2 Problem Definition 74
4.3 Modelling approaches 76

4.3.1 Stochastic global optimisation 77
4.3.2 Blockage correction 81
4.3.3 Gradient-based local optimisation 82
4.3.4 Hybrid Approach 85
4.4 Single turbine comparison 87
4.4.1 Calibration 87
4.4.2 Comparison 88
4.5 Array comparison 91
4.5.1 Wake combination approaches 91
4.6 Optimisation 95
4.6.1 Test scenario 1 95
4.6.2 Test scenario 2 98
4.7 Optimisation comparison 100
4.7.1 Test 1 101
4.7.2 Test 2 104
4.7.3 Test 3 107
4.7.4 Discussion 107
4.8 Conclusions 108

5 Integration of cost modelling within the micro-siting design optimisation 111

5.1 Introduction 112
5.2 Problem Formulation 113
5.3 OpenTidalFarm 115
5.4 Cable length minimisation 116
5.4.1 Problem Formulation 116
5.4.2 Cable routing optimisation in literature 117
5.4.3 Approaches to solving the cable-routing problem 118

5.4.4 Notation . . . 120
5.4.5 Use of a gradient-based optimisation approach with the cable-routing . . . 121
5.5 Tidal Array Routing . . . 123
5.5.1 Chromosome selection . . . 123
5.6 Integrating the micro-siting and cable-routing optimisations . . . 127
5.6.1 Cost model . . . 128
5.6.2 Income model . . . 129
5.7 Test cases . . . 130
5.7.1 Idealised channel . . . 130
5.7.2 Idealised channel bend . . . 133
5.7.3 Pentland Firth . . . 134
5.8 Conclusions and summary . . . 136
5.8.1 Future work . . . 137

6 Sizing a tidal turbine array **139**

6.1 Introduction . . . 140
6.2 Problem Formulation . . . 142
6.3 Surrogate modelling . . . 143
6.3.1 Gaussian process regression . . . 144
6.3.2 Bayesian optimisation . . . 147
6.4 Combined Model . . . 148
6.4.1 Financial model . . . 149
6.5 Idealised example . . . 152
6.6 Pentland Firth example . . . 156
6.7 Conclusions . . . 160

7 Uncertainty estimation **162**

7.1 Introduction 163
7.2 Motivation 164
7.3 Studying output sensitivities to parameters using the adjoint 166
7.3.1 Model boundary conditions 168
7.3.2 Data assimilation 170
7.4 Bayesian methods 171
7.4.1 Other approximation techniques 173
7.5 Complexity and uncertainty quantification 174
7.6 Optimisation under uncertainty 177

8 Summary and future work **180**

8.1 Summary 180
8.2 Future work 182
8.2.1 More realistic test cases 182
8.2.2 Turbulence modelling 182
8.2.3 Operational optimisation 184

List of Figures

1.1 Schematic of a typical horizontal axis tidal turbine 5

1.2 Two-dimensional backward facing step 8

1.3 Flow through a domain with inflow separated by a large body 9

1.4 Flow past a fixed, solid body 10

1.5 Flow through a domain with a fast and slow inflow seperated with 'splitter plate' 10

1.6 Fast, slow, fast inflow conditions recreate wake structure behind actuator disc 11

1.7 Actuator disc model for a turbine in flow 11

1.8 Section of a turbine blade showing shed and tip vortices 12

1.9 Tidal stream turbine blade 13

1.10 Turbine wake with vortices 14

1.11 Illustration of blockage 15

1.12 Photograph of the *SeaGen S* tidal turbine installed in Strangford Lough, UK 19

2.1 Hierarchy of analysis, model and optimisation complexity 31

2.2 Graph of problem complexity 32

2.3 Schematic showing communication between different model scales . 33

2.4 Schematic showing communication between different model and optimisation scales 40

2.5 Comparison of prevailing resource direction for tidal currents and wind 46

2.6 Graph of problem complexity with example models 47

3.1 An example of an unstructured mesh 53

3.2 Definition of H, h and η 55

3.3 Default *OpenTidalFarm* thrust curve 61

3.4 Default *OpenTidalFarm* power curve 61

3.5 Flow chart showing main *OpenTidalFarm* classes 67

3.6 Schematic of the *OpenTidalFarm* optimisation procedure 68

4.1 Sketch of velocity deficit behind a 20 m wide turbine computed using the Jensen wake model. 79

4.2 Sketch of velocity deficit behind a 20 m wide turbine computed using the Larsen model. 80

4.3 Sketch of velocity deficit behind a 20 m wide turbine computed using the Ainslie model. With a higher thrust coefficient the wake would grow in width downstream of the turbine. 81

4.4 Wake behind a 20 m wide turbine in 2 ms^{-1} flow using the shallow water model with the turbine modelled as an area of increased bottom friction. 83

4.5 Plot of optimisation methodology demonstration function 86

4.6 Demonstration of global optimisation 86

4.7 Demonstration of local optimisation 86

4.8 Demonstration of basin-hopping optimisation 86

4.9 Wake model comparison to CFD 89

4.10 Wake model array comparison to CFD 92

4.11 Wake model array comparison to CFD 93

4.12 Flow through test domain of scenario 1 96

4.13 Flow through test domain of scenario 2 96

4.14 Shallow water spacing plot scenario 1 97
4.15 Larsen spacing plot, scenario 1 97
4.16 Shallow water spacing plot scenario 2 98
4.17 Larsen spacing plot, scenario 2 98
4.18 Optimised turbine layout for simple-island test scenario 99
4.19 Plot showing impact of global blockage on power extraction of array . 100
4.20 Bathymetry for test scenario 1 102
4.21 Ambient flow through test scenario 1 102
4.22 Best optimal array layout for scenario 1 (peak flood) 103
4.23 Best optimal array layout for scenario 1 (peak ebb) 103
4.24 Comparison of optimised layouts for test scenario 1 104
4.25 Bathymetry for test scenario 2 105
4.26 Ambient flow through test scenario 2 105
4.27 Best optimal array layout for scenario 2 (peak flood) 105
4.28 Best optimal array layout for scenario 2 (peak ebb) 105
4.29 Comparison of optimised layouts for test scenario 2 106
4.30 Shallow water, gradient-based optimised array layout for scenario 3 (peak flood) 107
4.31 Wake model, genetic optimised array layout for scenario 3 (peak flood) 107

5.1 Diagram showing change in routing with movement of a turbine . . . 122
5.2 Plot of routing length with a traversing turbine 122
5.3 Schematic of an idealised routing, $R = r_1 \cup r_2 \cup r_3$ with capacity constraint $C_{ap} = 4$ 124
5.4 Schematic of *TidalArrayRouting* developed for *OpenTidalFarm* 125
5.5 Map showing idealised channel set up 131
5.6 Optimisation of eight turbines with and without cable cost minimisation 132

5.7 Domain for right angle bend simple channel example 133

5.8 Comparison of optimisation for power maximisation only and for financial return 134

5.9 Domain, mesh and inflow boundary of the idealised Orkney example. 135

5.10 Results for inclusion of cable costs in Pentland Firth test case 138

6.1 Bivariate distribution as Gaussian process prior 145

6.2 Example posterior distribution for a bivariate Gaussian 145

6.3 Prior for multivariate Gaussian distribution over 10 variables 146

6.4 Posterior distribution for multivariate Gaussian of figure 6.3 146

6.5 Flow chart showing the ensemble *TidalArraySizer* model. 150

6.6 Unstructured computational mesh of the simple channel with circular island test example 153

6.7 Plot of the Profit(N) surrogate initialised with training set data points at Profit($N = 0$), Profit($N_{\max}$), Profit($N_{\max}/2$). 153

6.8 Plot of the Profit(N) surrogate before switching to high fidelity sampling. 153

6.9 GPR surrogate at end of optimisation. 154

6.10 Optimised micro-siting of optimally sized array for flood tide 155

6.11 Optimised micro-siting of optimally sized array for ebb tide 155

6.12 Map of Pentland Firth with domain mesh overlaid 157

6.13 Surrogate model initialised with starting training data. 158

6.14 Surrogate model at conclusion of low fidelity runs. 158

6.15 Final surrogate model of N to P mapping for the Inner Sound example.158

6.16 Optimised micro-siting for optimally sized array in flood tide 159

6.17 Optimised micro-siting for optimally sized array in ebb tide 159

7.1 Plot showing mean peak flow speeds in the Inner Sound of the Pentland Firth 163

7.2 Sensitivity of array power with respect to bathymetry 167

7.3 Sensitivity of array power with respect to bottom friction 168

7.4 Plot showing $\mathrm{d}P/\mathrm{d}H$ for simplified Inner Sound domain 169

7.5 Plot showing $\mathrm{d}P/\mathrm{d}c_b$ for simplified Inner Sound domain 169

7.6 Example regions of uncertainty 173

7.7 3D plot of functional response to uncertainty in depth 174

7.8 2D Plot of functional response to uncertainty in depth 175

7.9 Graph of problem complexity 176

7.10 Schematic comparing robust and traditional approaches to design . 178

Chapter 1

Introduction

Contents

1.1 Introduction 1
1.2 A primer on tidal stream energy 2
1.3 The hydrodynamics of tidal stream turbines 7
1.4 Motivation and scope 17
1.5 Research Objectives 21
1.6 Thesis summary and authorship of work presented 23

1.1 Introduction

In recent years, marine energy has become an emergent sector in the renewable power industry, with interest driven largely by the significant contribution tidal stream energy could potentially make as a part of the UK's renewable energy portfolio [Johnstone et al., 2013]. While methods for evaluating the exploitable tidal resource vary, it is generally considered that the national tidal power asset could amount to around a fifth of the UK's electricity consumption [Carbon Trust, 2011] and 10 to 15% of the harvestable global tidal power resource [Johnstone et al., 2013]. That notwithstanding, the global market is also developing, with notable potential tidal power assets in Canada, Argentina, France, Ireland, Russia, Australia and China [O'Rourke et al., 2010]. The rate of research increased early in the new millennium, led by academia and fuelled by private industry benefiting from offshore expertise from a North Sea oil and gas industry in decline [Charlier, 2003]. Since then several full-scale prototypical devices have been deployed for

field testing of the technologies and this year (2016) the first multi-device array is to be installed in the Pentland Firth in the UK [MeyGen, 2014]. As such, the industry stands on a tipping point: On one hand, failure could undermine investor confidence in similar projects and set the industry back significantly. On the other hand, success could bolster the sector, which could have a beneficial economic impact for the UK – both through facilitating energy autonomy by increasing domestic generation and through development of an industry with exportable expertise.

Tidal stream turbines (TSTs) use the momentum of tidal currents to generate electricity. While the high power density of the resource, its regularity and reliability with which it can be forecast makes it an attractive technology for development, currently the cost of energy delivered to the grid through TST installations is unfeasibly high as compared to alternate sources [LCICG, 2012]. This is due in no small part to the infancy of the industry given that, to date, only prototypical, demonstration TST devices have been installed in isolation. However, it must also be noted that design for durability is expensive, and issues such as high flow forces which significantly stress hardware and turbulence which causes structural fatigue, require costly engineering solutions. In addition to this, sites with high tidal flow velocities often present very remote and challenging environments in which to install and operate infrastructure.

Until this point, much effort has been focused on proving the turbine technology itself, but as the tidal stream energy industry enters a new phase in its genesis it has reached a 'make or break' milestone of having to prove itself feasible on an industrial scale. The word 'feasible' masks a host of criteria for success from a variety of stakeholders, but in general it means reliable power generation at a competitive price. The focus of this thesis is in developing an open-source software array design tool which will help facilitate the development of feasible, financially profitable marine energy installations.

1.2 A primer on tidal stream energy

Before delving into the motivation for the work let us first 'set the scene' for tidal stream turbines, the arrays which comprise them and the tidal stream resource they exploit. This primer is only intended as the most basic of introductions; to introduce the 'broad strokes' of technologies and concepts involved in tidal stream energy generation before they are abstracted later in this thesis for modelling pur-

poses.

1.2.1 Tides

Consider an oceanic planet with no land, in the presence of a celestial body which has a gravity field. This gravity field induces the formation of 'bulges' of ocean on opposing sides of the planet – one facing the source of the gravitational field, the other facing away. This sets up a long wave in the ocean, circumnavigating the whole surface. If the ocean-planet is spinning, then an observer on the surface of the planet would notice sea-level change over the course of the planet's rotation. This vertical change in surface height is called a tide. Any relative motion of the celestial body and the planet will affect the gravitational field which is driving the long wave and impact upon its characteristics.

Earth differs from this idealised model in two key ways. Firstly, there are multiple celestial bodies in our solar system each of which has a gravity field that is felt on Earth to a greater or lesser extent (the Sun and the Moon being by far the most important ones). This means that the driving force behind the tidal long wave is continually (though predictably) changing as the external gravitational forces interfere with each other constructively or destructively. Tidal prediction is achieved through combining a set of *tidal constituents*, each of which represents a different cyclical movement of the Sun or Moon relative to Earth.

The second major difference is the presence of land masses on the earth which interact with the tides to induce horizontal, tidal currents. Such currents can develop in different ways depending on bathymetric and geographic features. For example, when two large bodies of water are connected via a (comparatively narrow) strait, it is common for the tides at each end of the strait to be out of phase and therefore be unequal in surface elevation. This results in a head difference between one end of the strait and the other which induces a tidal current. On Earth, the fastest tidal current is found in Saltstraumen, a strait in Norway where flows reach over 11 ms^{-1}.

Finally, tides and tidal currents are impacted by meteorological mechanisms such as barometric pressure, wind effects and storm surges. In general this impact is minor, but in some cases the effects of a number of mechanisms can combine to produce more significant results. In the case of tidal turbines, such variable perturbations around the mean currents, while small, could be important for loading or fatigue of the structure or components.

1.2.2 Tidal stream turbines

Tidal power may be generated through exploitation of the head difference of tides (gravitational potential energy, P.E.), or through converting the momentum of tidal currents (kinetic energy, K.E.).

In the first case, a barrage is constructed which creates a basin or lagoon (or in the case of 'dynamic tidal power' the barrage is simply very long – thus causing a tidal phase difference from one side to the other). The head difference from one side of the barrage and the other is exploited by passing the water through low-head turbines – in much the same way as a hydro-electric dam (although these use turbines designed for high head differences). Such constructions are significant and costly civil structures and generally have a large impact upon the local environment.

It is the second case – converting the momentum of tidal currents into electricity – that is the focus of this thesis. Tidal stream turbines are devices which extract the kinetic energy of the tidal current's mass of moving water, converting it into mechanical work to drive a generator and output electricity. A number of different device types exist including horizontal axis tidal turbines (HATT), vertical axis tidal turbines (VATT), ducted turbines etc., however in general, much like with wind power, the tidal stream industry seems broadly to be coalescing around minor variations of the HATT design. A schematic for a generic HATT is shown in figure 1.1, indicative measurements are shown to the right of the figure and the turbine itself is based on Atlantis Resources Limited's (*www.atlantisresourcesltd.com*) 1 MW prototype turbine (rated at 2.65 ms^{-1}) which is similar to that which will be used in MeyGen's (*www.meygen.com*) array development in the Inner Sound of the Pentland Firth, UK in 2016.

In general, a tidal stream tubine device comprises:

- A blade / rotor assembly – converts the kinetic energy in the flow into rotational energy in/off the device
- A drivetrain - conditions the rotational energy to appropriate torque / speed for the generator
- A generator - converts the rotational energy into electrical energy
- A support structure - holds the device in place.

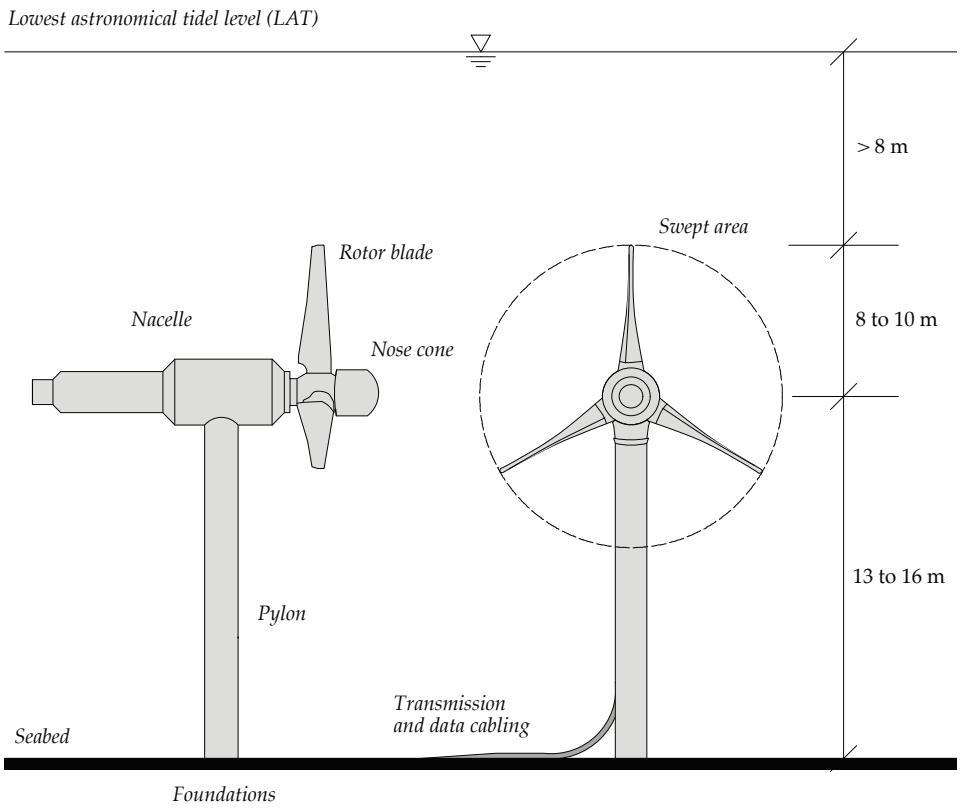


Figure 1.1: Schematic of a typical horizontal axis tidal turbine (HATT). Diagram based upon Atlantis Resources Limited's (*www.atlantisresourcesltd.com*) 1 MW prototype turbine – similar to that which will be used in MeyGen's (*www.meygen.com*) array development in the Inner Sound of the Pentland Firth, UK in 2016

1.2.3 Tidal stream turbine arrays

Much as with wind power, tidal stream power may be extracted on an industrial scale by installing many turbines (10's to 100's) together in an array on a site with highly energetic flow regimes. Apart from the obvious benefit of extracting more power, this facilitates economies of scale, reducing the per turbine costs of purchase, installation, maintenance and infrastructure – as compared to a turbine installed in isolation or in community scale (<10 turbines) schemes. The main properties of turbine arrays are the number and type of turbines and their individual locations within the site leased by the array developer.

The presence of turbines, however, impacts upon the flow both in the near-field as well as, potentially, further away. This interaction results in a coupled hydrodynamic system in which flow velocities vary on a wide range of spatial scales. Furthermore, for a significant portion of the turbine's flow velocity operating win-

dow, its power production is approximately proportional to the cube of the flow velocity at the turbine location and hence small changes in the velocity can have a large impact upon the power production of the turbine.

Understanding and predicting the behaviour of arrays of turbines is further complicated as, to date, no tidal arrays have been installed anywhere globally, only individual turbines have been tested in-situ and at full scale. Furthermore, laboratory experimentation with arrays of turbines (i.e. tank-scale) has only been conducted a handful of times and only with a relatively small number (3 to 5) of scaled turbines – most notably by Stallard et al. [2013]. Given the expense of real-world construction, installation of a full-scale test array is prohibitive, and few laboratory tanks of a necessary size exist to conduct scaled array experiments. Consequently computational modelling is fundamental to the design and analysis of turbine arrays – with no real opportunity for complete / extensive validation until the first array becomes operational.

1.2.4 The tidal stream resource

Some of the most attractive aspects of tidal power as compared to other renewable energy sources are its predictability, high power density and a reduced aesthetic and acoustic impact for local residents of installations. While tidal power is variable, it is predictable – a high quality survey of a site for 29 days can facilitate prediction of the resource for the lifetime of the turbine array with considerable accuracy [EPRI, 2006]. Coker et al. [2013] note that while tidal power exhibits strong daily variation, when viewed on monthly or annual timescales the resource is, in fact highly stable (as compared to wind and solar power which see seasonal variation) this is of great benefit for grid scheduling.

It does not, however, come without challenges – for example, assessing the size of the tidal resource at a local or global scale. This is difficult because of the highly coupled nature of power extraction and the resource. Each site that may be promising for tidal power development has a 'potential' – the maximum amount of power that may be extracted [Vennell, 2011a]. Adding turbines slows the flow velocity through the site and at a certain point, as more turbines are added, the amount of power being extracted diminishes. Furthermore, while the total power extracted by the array will decrease once more turbines are added after this threshold value, in actual fact long before this point the extracted power *per additional turbine* will be decreasing. That is, there is a diminishing return of power on investment in ad-

ditional turbines. Consequently, while there is a physical limit to the extractable power from a site there is also an economic one. Exploring these limits is one of the focuses of this thesis.

Having said that characterising the size of the resource is difficult, once an array is designed (that is, the number, type and placement of the turbines on a given site is fixed) then the key benefit of tidal power over other forms of renewable energy generation is the predictability of the resource. So-called variable renewable energy sources (VREs) are problematic for the management of electrical distribution networks (in the UK this is called the national-grid) due to their unpredictability. This variability means that VREs cannot be relied upon as part of the grid's so-called baseload (the always-on power stations whose supply fulfills the grid's permanent minimum load) [Cornelis van Kooten, 2009]. In fact, a sudden drop in the VRE supply will require the grid to respond by seeking supply from more traditional gas or coal sources. This means that such sources must be kept online and ready – though maybe idle for much of the time. This can cause some unexpected market effects; Hirth [2013] found that a wind penetration of 30 % of grid generation capacity would only be worth half the equivalent provision from a constant source of electricity (measured as the relative price of a given technology compared to the average wholesale electricity price). As such, as market penetration of VRE's rises, there is a strong possibility that their value will decrease. This problem is often complicated by energy policy – for example requiring renewable energy to be used in precedence of non-renewable energy – which can have strange effects such as a negative price of energy on sunny, windy days (for example in Germany on May 11th 2014) [EnergyTransition.de, 2014].

1.3 The hydrodynamics of tidal stream turbines

In the next chapter a variety of approaches for modelling applied to tidal stream power applications will be presented. Such approaches are wide and varied and a designer may choose one over another dependent upon the scale at which they are interested and physical phenomena they are hoping to study. Before introducing these models which are, by their nature, simplifications of the real-world physical processes at work, it is instructive to gain a high-level, intuitive understanding of the hydrodynamic behaviour of a tidal turbine placed in energetic flow. Developing this understanding is the focus of this section.

1.3.1 Wakes

It is helpful to build up a picture of the wake structure of a single turbine piece by piece, by starting with simple conceptual models which may be combined to gain a feeling for the physical processes which generate it.

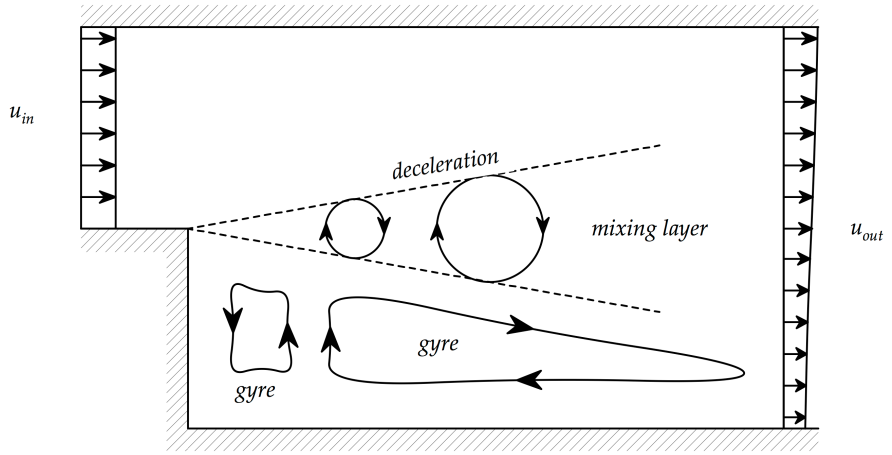


Figure 1.2: Two-dimensional flow passed a backward facing step (plan-view). Stationary fluid at the face of the step interacts with fluid flowing past it, a mixing layer develops which transfers momentum across the domain.

The first conceptual model is the backward facing step problem, in which a two-dimensional flow passes along a simple channel the width of which suddenly increases on one side in a step configuration (see figure 1.2). The fluid at the face of the step is stationary while the fluid flowing over and past the step has momentum from its velocity and mass. A layer of shear forms with the faster flowing water decelerating as momentum is transferred to the stationary water, which accelerates. A mixing layer develops which grows like a cone from the corner of the step until the momentum is well mixed across the entire cross-section of the channel. The differential velocity induces a turning moment in the flow which helps to pull momentum across from fast to slow flow. Immediately behind the step and out of the cone of the mixing layer slow moving recirculation zones or gyres are formed.

Reflecting this domain across the x axis (figure 1.3), so that now two inflows are separated by a large body, mixing layers form in the same place and these meet downstream of the body. Once again with slow moving gyres immediately behind the body. While this advances the understanding of downstream effects, there are

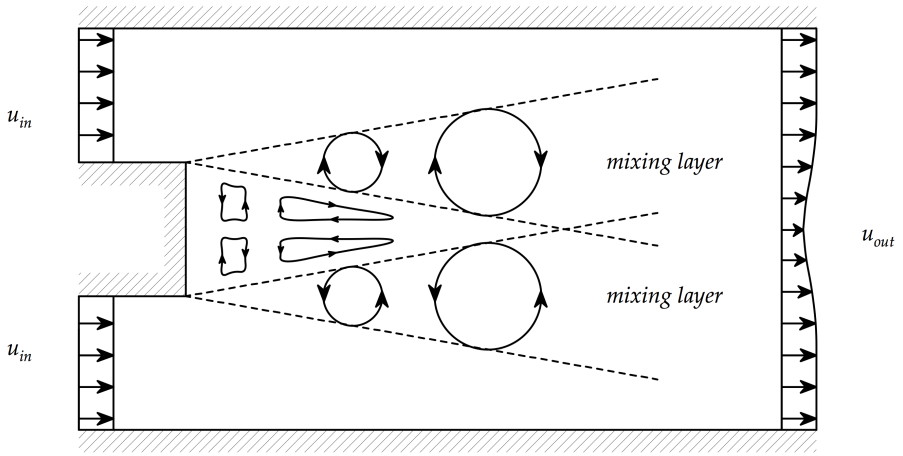


Figure 1.3: Stationary flow behind the obstruction is energised by momentum transfer from two sides, and consequently two mixing layers form and expand outwards. These meet and overlap downstream.

also physical processes of significance upstream and outside the wake of a turbine.

These can be seen by extending the domain upstream and reconfiguring the obstruction into a triangle (which may represent, for example, an island in a river), the obstruction forces the oncoming flow to part on either side, figure 1.4. Conservation of mass (and the incompressibility of water) requires that the same volume of water that once occupied the full cross-section of the channel must now occupy the reduced cross-section of the obstructed channel, hence the flow speed of this bypassing flow increases compared to the incoming flow.

While this is the effect on the flow of placing fixed, solid bodies such as the turbine nacelle and support structure, the effect of placing the spinning blades of a turbine rotor in the flow has a somewhat more complex effect. To reduce this complexity, often-times turbines are modelled, not as rotating blades, but instead as an *actuator disc*. This is a disc whose diameter matches that of the swept diameter of the turbine blades but is instead a stationary, porous plate. The concept can be simplified by adjusting the backward-facing step in figure 1.2 to a domain in which a fast and slow inflow are separated by a 'splitter plate', figure 1.5.

Much the same physical processes take place in this example as did in the backward facing step however, stationary gyres do not form as before, as this rotation is continually advected downstream by the slower-moving flow. Once again a more

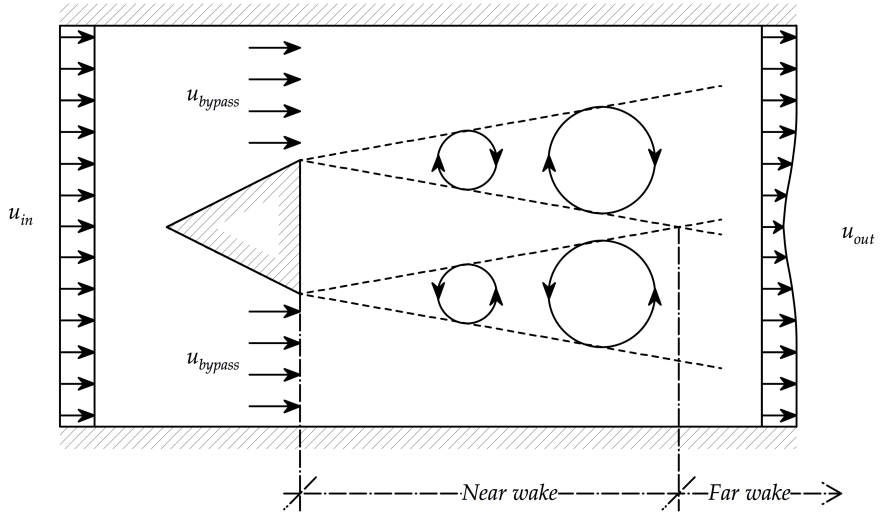


Figure 1.4: Flow past a fixed, solid body. Uniform inflow across the domain is forced to part either side of the obstruction, this leads to a zone of high pressure and deceleration upstream of the obstruction, and acceleration of the flow as it bypasses either side of it.

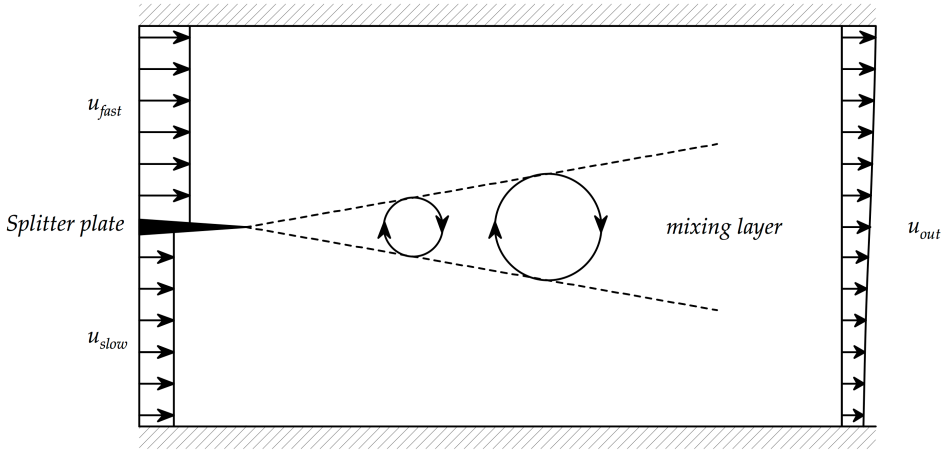


Figure 1.5: Flow through a domain a fast and slow inflow separated by a 'splitter plate'.

complete picture of the wake is obtained by reflecting the domain along the x axis so that now two inflows on the outside at equal high velocity are divided by a low velocity inflow in between, all three separated by 'splitter plates', figure 1.6.

As before, the mixing zone eventually meets in the middle of the channel, downstream of the central lower-velocity inflow. In the context of turbines, the area upstream of where these mixing layers meet is called the 'near-wake region' and the

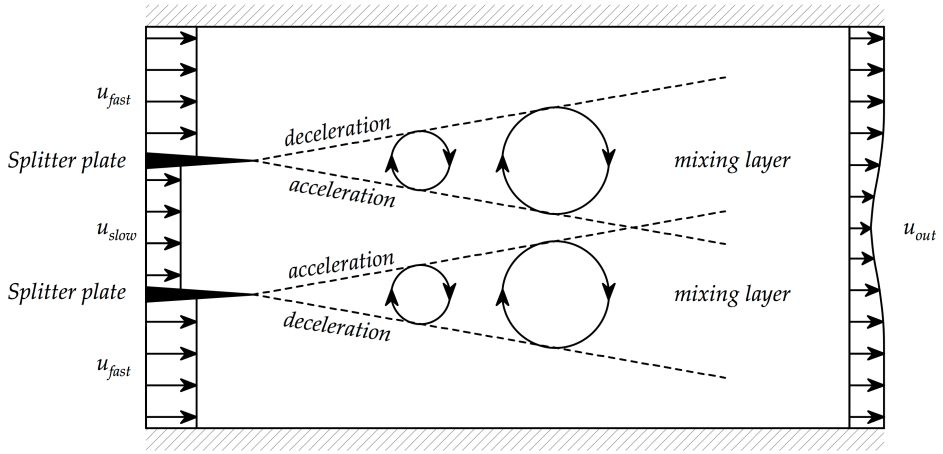


Figure 1.6: Fast, slow, fast inflow conditions recreate wake structure behind actuator disc.

area downstream the 'far-wake' region. These labels are shown in figure 1.4.

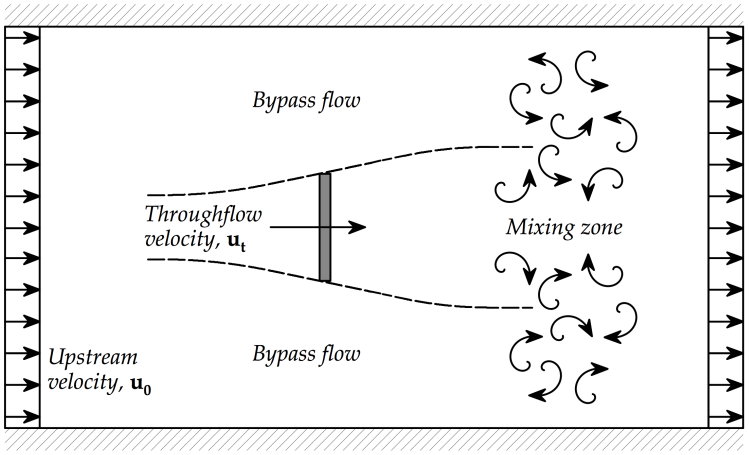


Figure 1.7: Actuator disc model for a single turbine showing idealised stream tube expansion of flow passing through disc.

Finally, rather than a split inflow or an island-obstruction as in figures 1.6 and 1.4, a momentum sink (a strip in two-dimensions or an actual disc in three-dimensions) is placed in the flow, figure 1.7 shows the 2D representation of the actuator disc. Flow upstream of the 'turbine' decelerates, while the flow bypassing the 'turbine' accelerates and re-energises the slower flow downstream via turbulence in the

mixing layer. Once again, this mixing of momentum ensures that, if observed sufficiently far downstream, the cross-section will be at constant velocity.

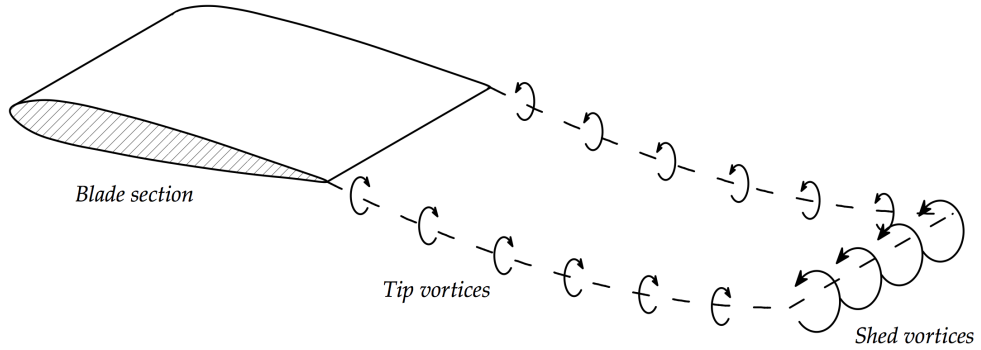


Figure 1.8: Section of a turbine blade showing shed and tip vortices

This picture of a turbine wake is complicated enormously by consideration of the rotating blades. Turbine blades are essentially wing shapes which twist axially along their length. The turbine blade is twisted because the tips of the blades are travelling much faster than the roots of the blade and the angle of attack of the blade must alter to compensate for this. The blade is composed of aerofoil cross-sections such as in figure 1.8. Aerofoils are structures designed to generate lift perpendicular to incoming flow while minimising drag on the structure in line with that flow ('incoming' and 'in line' being relative to the support structure). This is because aerofoil geometry induces an area of high-pressure on one side and an area of low pressure on the other, the pressure differential perpendicular to the flow results in a lift force on the wing.

While gyres were observed forming downstream of the backward-facing step, so too do vortices develop along the trailing edge of the aerofoil. The high-pressure flow on one side of the aerofoil, once past the trailing edge, is no longer separated by the wing from the low-pressure flow that has developed on the other side. This induces the formation of a 'shed vortex' which is advected by the flow downstream and away from the aerofoil. Vortices also form at the tip of the blade as high-pressure fluid spills over the tip of the blade, reducing its efficiency. As with shed vortices, tip vortices are advected downstream, figure 1.8.

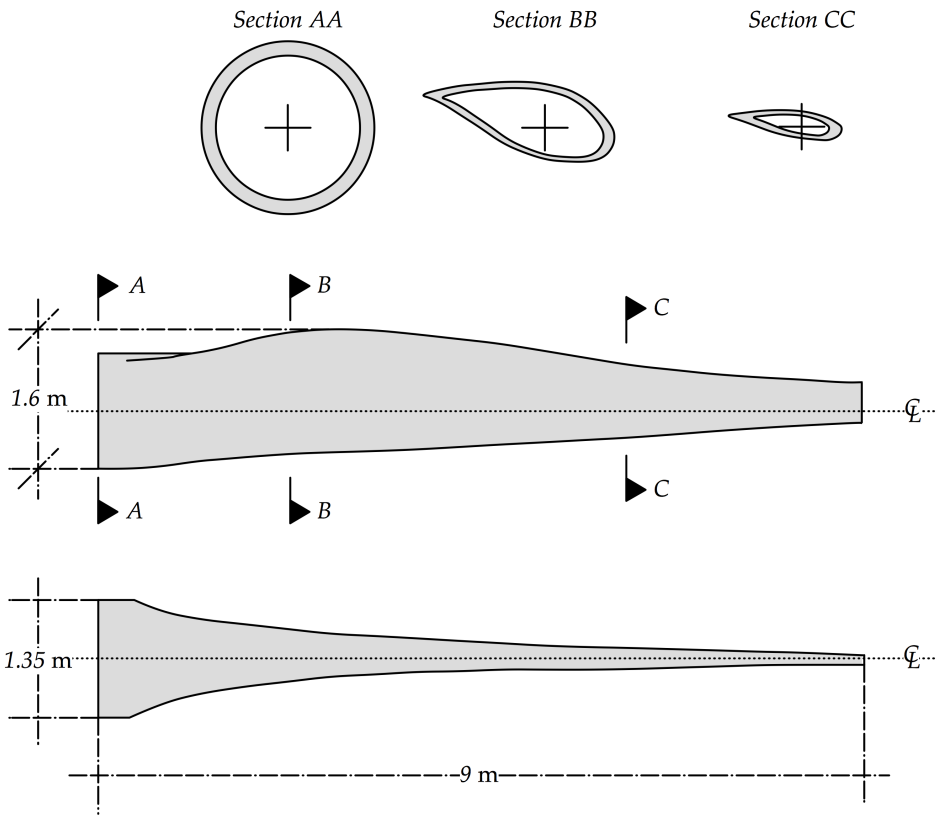


Figure 1.9: Andritz Hydro Hammerfest HS1000 tidal turbine blade.

Tip vortices help momentum transfer into the mixing zone, while shed vortices help mixing in the near-wake region [Mungar, 2014]. Consequently, the structure of the near-wake is a device specific (dependent upon the shape of the blades – figure 1.9 – and the vortices they induce), highly complex, three-dimensional structure (figure 1.10). By the far-wake, the swept diameter of the blades and thrust force they exert, become dominant and the turbine geometry-specific vortices play a reduced role.

1.3.2 Blockage

A very important concept in the study of tidal turbines in flow is that of blockage. Vennell et al. [2015a] notes that because of blockage effects, harvesting 100 MW of power from a site is not as straightforward as simply installing 100 one MW rated turbines. When, as in a simple channel or tidal strait the boundaries

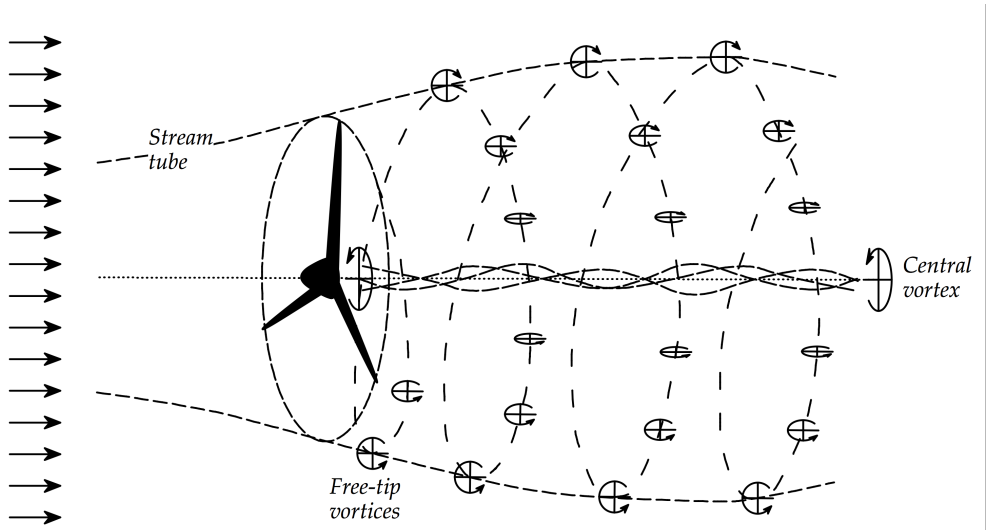


Figure 1.10: Turbine wake with vortices.

are such that lateral displacement of the flow is prevented, the effect of placing an object in the flow is different as compared to what would happen in an unlimited stream. Clearly these effects will have an impact upon the formation and structure of the wake, however in the tidal stream power application the effects are also felt strongly outside of the wake. *Local* blockage effects result in some of the phenomena already encountered – upstream high pressure and the acceleration of flow bypassing the turbines. *Global* blockage, on the other hand, deals with the effects of the turbines on the wider domain. For example, were an east-west orientated tidal strait to have an island situated in the middle and a turbine was placed to be south of the island, the increased blockage due to the presences of that turbine would redirect flow around the north of the island. This will become an important consideration when thinking about arrays of large numbers of turbines – in such a case the designer may (seemingly reasonably) expect more power from more turbines, however if flow is easily redirected and the velocity of the flow through the turbine site decreases as a result, the designer is likely to observe a reduction in the power extracted per turbine as more are added to the site. This trade-off forms the focus of chapter 6.

Blockage is one of the key differences between the behaviour of wind turbines in flow versus tidal turbines in flow. While, on the surface, the two technologies look very similar, wind turbines are very rarely located in laterally constrained flow and, perhaps more fundamentally, never in vertically constrained flow (due to the height of the atmosphere). Conversely, tidal turbines may be placed in laterally

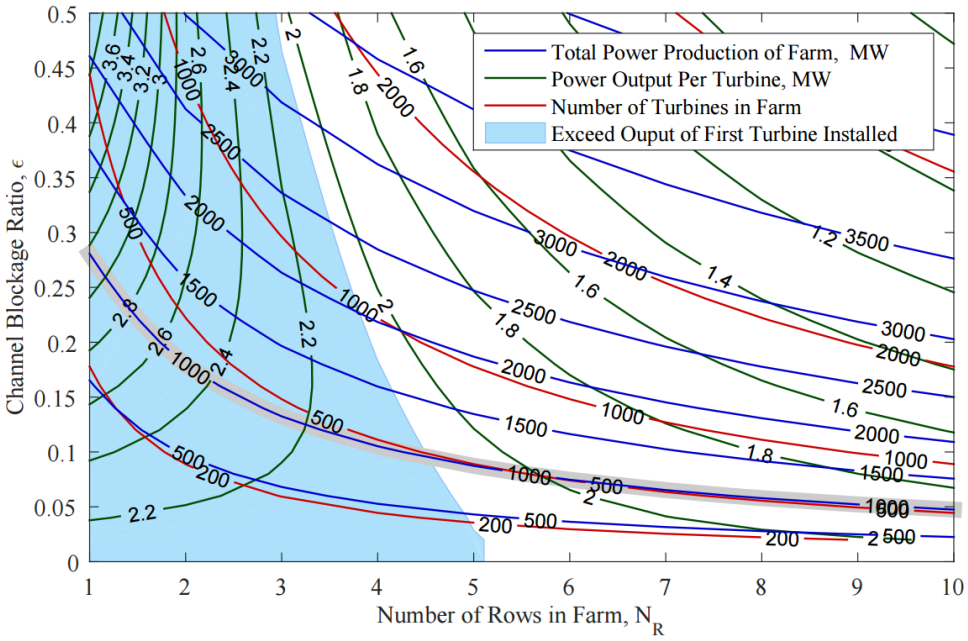


Figure 1.11: Plot from Vennell et al. [2015b] showing the interplay between power extracted from the flow versus the array design (here the number of rows of turbines and the percentage of the channel each of those rows obstruct).

constrained channels or near to headlands which will constrain flow on one side. Additionally, they will almost always be severely constrained vertically, placed as they are between the seabed and the water surface, potentially occupying a significant portion of the water column, see figure 1.1.

While that is not to say that wind turbines are not subject to blockage effects, they are less certainly dominant a mechanism. In the examples so far in this section, only two dimensional plan-view cases have been shown. However in the three-dimensional case the mixing layer becomes a three-dimensional structure. As such, in the xy plane the mixing layer is pulling momentum into the wake laterally, while in the yz plane, momentum is being pulled in from flow bypassing above and below the turbine. The rotors of a wind turbine are located generally over 60 m from the ground and the flow is essentially unconstrained to the upper atmosphere. This means there is an abundance of energy to be drawn in to re-energise the wake. The bounding provided by the free-surface and bottom bathymetry means that this re-energisation mechanism is much less effective for tidal turbines. As a result the wakes are longer for tidal turbines than for their wind analogues. As a result, *tidal* turbine array design optimisation is arguably more challenging to carry out, and more important in terms of the array's performance.

A parameterised representation of the relationship between power extraction and array design (array design is parameterised through the number of rows of turbines and the percentage of the channel they obstruct) is shown in figure 1.11.

1.3.3 Power extraction

One of the effects of blocking the flow with a solid body or indeed a turbine is that a pressure differential will form between high-pressure on the upstream side of the blockage and low pressure on the downstream side. This effect causes drag – a force felt by the body in the downstream direction. Momentum theory links this force to the freestream velocity, u_{fs}, the velocity far upstream of the turbine where the flow is undisturbed

$$F_D = \frac{1}{2}\rho||u_{\text{fs}}||u_{\text{fs}}C_d A, \tag{1.1}$$

where ρ is the density of the fluid medium, A is the area of the orthographic projection of the object (or in the case of turbines the swept area of the blades) and C_d is the drag coefficient – a dimensionless coefficient which depends on the object geometry. Multiplying by velocity gives the power extracted from the flow

$$P_D = \frac{1}{2}\rho||u_{\text{fs}}||^3 C_d A. \tag{1.2}$$

For turbines the drag coefficient is substituted in (1.2) by the power coefficient, C_p which seperates the power lost to turbulent mixing from that which does useful work on the turbine. It additionally represents a measure of the turbine's efficiency; that is the ratio of energy extracted from the flow to the total energy in the flow passing through the turbine. Consequently, the power extracted from the flow by the turbine is

$$P = \frac{1}{2}\rho||u_{\text{fs}}||^3 C_p A. \tag{1.3}$$

Note that for engineering reasons C_p is variable and dependent upon u_{fs}, this shall be discussed in the next chapter when turbines are parametrised for inclusion in a numerical model.

1.4 Motivation and scope

In the previous section, some of the mechanisms underlying the hydrodynamic interactions between turbines and the flow into which they are placed were explored. Two key points are the highly coupled interactions between turbine and flow, and the high sensitivity of the power with respect to the velocity of that flow. As has been noted, this means that obtaining a power estimate for multiple turbines via modelling is particularly challenging; as inaccuracies in capturing the turbine / flow interactions will lead to inaccuracies in the flow, which will be cubed to give even larger inaccuracies in the power estimate. The likelihood that turbines are to be deployed in highly energetic flow environments exacerbates the modelling problem, as the flow will be highly dynamic and advection-dominated or turbulent.

Bryden et al. [2004] demonstrated the importance of designing arrays based on the *resultant* flow through the turbine site – that is to say that the coupling between turbines and flow cannot be overlooked if an accurate power estimation is to be achieved. As such, straightforward *analysis* of an arrangement of several turbines in a flow is challenging, while improving that arrangement in the design process adds additional complexity. Four key factors mean that without a formal optimisation approach the designer is unlikely to produce optimally performing micro-siting designs for realistic sites with multiple turbines in close proximity.

1. The size of the design space, is very large: Assuming that the plan-view location of N turbines (their micro-siting design) are the only variables the designer can adjust, the space of feasible designs has $2N$ continuous control parameters (two Cartesian location coordinates for each of the N turbines). In practice the designer can have many additional controls (type of turbine, size of turbine etc.)

2. The turbine site is likely chosen because of its energetic flow regime and the turbine and flow are very closely coupled; the resulting hydrodynamics are complex and turbulent.

3. The power estimate (which is a natural indicator of an array design's performance) is highly sensitive to those hydrodynamics.

4. Practical constraints exist which limit the possible solutions, for example there is a limit to how close turbines can be placed to each other, and tur-

bines must be sited with the area leased by the array developer at an appropriate water depth with the necessary seabed conditions.

In the following chapter, some of the available methodologies will be explored in detail, however there are two main approaches to micro-siting optimisation: Either use a global optimisation algorithm which tries a large number of different micro-siting designs, computes the flow and then estimates the power for each. Or use a local optimisation algorithm which uses gradient information (in this case the gradient of the power with respect to the turbine locations) to optimise the design in far fewer flow solves (in the order of 100 as compared to upwards of 10^6). The former methodology is based upon approaches developed for wind turbine array micro-siting optimisation, while the latter is an approach developed by Funke et al. [2014]. One of the aims of this thesis is to integrate these two approaches and determine whether in combination the strengths of each method will be complementary and their weaknesses minimised.

The starting point for this survey of methods is to investigate the approaches to array design developed for the wind industry. While there are clearly parallels between the concept of tidal stream power and wind power, the differences in both technology and sector maturity have a noteworthy impact upon design decisions [LCICG, 2012]. That is not to say, however, that lessons from the wind industry cannot be applied to the tidal sector. While the underlying physical dynamics may not be common between the problems, the form of the optimisation problem is, and hence wind industry array optimisation approaches will provide a useful starting point for approaching the tidal analogue.

One notable factor in wind farm design – and an observation that motivates much of the work to come – is that there has, in the past, been a focus across the renewables industry on design optimisation based upon maximisation of energy yield [Thomson et al., 2011]: For example micro-siting to optimise power extraction of the array has generally been separated from in-depth consideration of other design factors such as minimisation of environmental impact and optimisation of project economics. However, any large infrastructure project has a wide variety of stakeholders, often with divergent objectives and motivations. Consequently, a framework in which these varying concerns could be modelled and integrated is critical in determining whether a proposal is viable and if so, in developing a design approach in which these differing interests could be balanced and their effects understood. Most notably, those developing or holding tidal current assets will almost certainly be businesses motivated largely by profit. While energy yield

is an indicator of the potential income of a project, this must be offset against the potential costs in order to fully understand the proposal.



Figure 1.12: SeaGen Device by MCT is a prototypical large scale TCT installation in Strangford Lough, Northern Ireland. The device consists of an arm with two 16 m diameter rotors – here shown raised for maintenance – that can be lowered on a monopile into the tidal flow. The device generates its rated power of 1.2 MW for a large part of the day.

In short, while several prototypical projects, such as the SeaGen device in Strangford Lough, Northern Ireland (figure 1.12), have demonstrated the technical viability of tidal stream power technology, what remains thus far unproven is the viability (in the broadest sense) of the concept on an industrial, array scale. As with any engineering design problem, the quality of any solution is measured in a variety of ways against often disparate criteria; the most environmentally friendly solution may not be the most economically advantageous.

While thus far the discussion has broadly centred around micro-siting optimisation, the design of an array has potentially many more parameters – for example (and perhaps more fundamentally) the question of array *size* – or, how many turbines should comprise the array. This is a question that is often overlooked in literature, especially in resource assessments which frequently assume a size and micro-siting design for an array without justifying the choice made. A number of papers have highlighted the significant increases in power that may be achieved

by micro-siting optimisation, so it is highly likely that array design optimisation will become an increasingly important part of the resource assessment methodology. One interesting example of where this may be important is in the leasing of sites for development of tidal turbine arrays by the UK's Crown Estate. The Crown Estate calculate the value of a tidal site in the same way they would for off-shore wind array sites – based upon the area of the site and the value of the power that a *wind* array could produce on it. Clearly, however, with tidal stream power there are two dominant flow directions which means that the shape and *orientation* of the site play an enormous role in the potential resource that could be extracted by tidal turbines on it. Understanding the role of array design – especially array size – on resource assessment by developing efficient tools is one of the key objectives of this work.

As with any commercial venture there is risk attached to the development of tidal power arrays. The role of the designer is to ensure that the technological risk associated with the venture is minimised, quantified and communicated. Models are frequently designed to produce a straightforward output – for example the exploitable resource of a given channel in mega Watts – but if there is uncertainty surrounding that figure which has not been calculated and communicated alongside it, then it is of very limited value. Sensitivity studies are an important tool for understanding some of the aspects of these uncertainties, however (depending on the underlying model) they are often computationally expensive and time-consuming to produce and the control parameters with respect to which the sensitivity is being computed often have to be greatly simplified for the endeavour to be computationally tractable at all. Development of a tool which facilitates sensitivity studies to be conducted efficiently and computed at a cost independent from the number of control parameters is an important first step to introducing uncertainty quantification methodologies into the tidal industry, with the ultimate goal of working towards optimisation while accounting for uncertainty.

The scope of this work is limited in some key ways: Most notably only two-dimensional models will be considered, though obviously there are some inherently three-dimensional mechanisms that this assumption will fail to capture. In addition, the spatial scales of the problems under consideration require the use of simplified, parametrised, turbine representations. These simplifications are a product of the need to balance the computational complexity of the problem under consideration and assign computational resources where they will have the most benefit answering the specific questions that this thesis aims to answer. This

balancing of complexity will be considered in some detail in the following chapter. Another limitation on the scope of this work is the, while financial modelling is a recurring theme in the results chapters of this thesis, only simple financial models are employed in this thesis. This is because real-world business information is highly proprietorial and therefore inaccessible for academic presentation. The assumptions made in the simplifications of these models will be addressed in later chapters as they arise.

1.5 Research Objectives

At its most broad, the aim of this project is to develop numerical tools and computational processes to better facilitate a more holistic and flexible approach to tidal stream turbine array design – both in terms of array size and micro-siting design.

The first objective (presented in the following chapter) is to analyse and categorise; firstly the different modelling approaches for tidal power applications and secondly, how these approaches are (or may be) integrated into an optimisation framework to automate array design. The infancy of the tidal stream power sector means that, in general, such frameworks are under-developed so this survey of methods will also pull in techniques and approaches from the wind industry.

The aim is then to combine different design optimisation approaches into a single coherent software package which is composed of a family of models that may be deployed dependent upon the objective of interest, desired accuracy and computational budget of the designer. *OpenTidalFarm*, the framework upon which these models are to be built, is introduced in chapter 3. The family of models and optimisation approaches are then discussed, implemented and compared as presented in chapter 4.

Since such array developments will be led by business for profit and judged on the merits of their viability, in chapter 5 consideration is given to the incorporation of economic factors within the model. Through this work a framework should develop in which the beneficial and detrimental effects of design decisions can be judged. The aim of this portion of the work is to demonstrate how location-based costs may be integrated into the micro-siting design optimisation process, with the goal of comparing a power-maximisation objective against a profit-maximisation objective.

While the micro-siting problem has received some limited attention [Thomson et al., 2011, Sullivan and McCombie, 2013, Roc et al., 2013, Funke et al., 2014], the more fundamental question of how many turbines there should be in the first place, has been given seemingly less consideration and is tackled in chapter 6. Resource assessment methodologies often make unrealistic assumptions as to how the resource may be harnessed – for example assuming that turbine barrages stretch across tidal straights from shore to shore, not leaving space for shipping (for example [Adcock et al., 2013]) or acknowledging that nearer shore there is insufficient depth to install turbines. While these models serve a useful purpose in identifying the upper bound of power available in a given region, for a developer – who has a clearly defined turbine deployment area leased from the land-owner – they are less helpful, and a more practical methodology is called for. One of the aims of this work is to develop this methodology.

Finally, while the ability to optimise an array is a vital ingredient for the wider deployment of tidal power generation, investment will only materialise if there is confidence in the outputs of design tools. Various uncertainties need to be quantified, minimised, and incorporated within the design process. An important first step towards this end, and goal of this work is to develop the ability to facilitate the production of sensitivity analyses which can be computed in an efficient and control-parameter-size invariant manner. This is discussed in chapter 7.

1.5.1 Summary of research objectives

The goals of this project are to:

1. Develop and combine array micro-siting design techniques into a single software tool.
2. Develop a framework in which objectives other than power maximisation may be included in the micro-siting design optimisation.
3. Develop an efficient approach to resource assessments that more realistically captures the array design considerations.
4. Develop an efficient and practical sensitivity analysis tool as a first step toward uncertainty quatification and ultimately array design optimisation under uncertainty.

1.6 Thesis summary and authorship of work presented

Since much of the work in this thesis is constructed around the development of an existing code; *OpenTidalFarm* [Funke et al., 2016] in the context of this thesis some clear delineation must be made between existing work, collaborative work and my own individual work. In laying out an overview of the remainder of this thesis I will identify the publications on which each chapter is based and which contributions are my own.

In the following chapter is presented a survey of existing methods for turbine array optimisation – both for wind and tidal applications. This is structured around a framework which has developed from Culley et al. [2014]. The underlying mathematical principles of the finite element method and adjoint model are then introduced culminating in an explanation of *OpenTidalFarm* in chapter 3. The core idea of *OpenTidalFarm* (which is described in section 3.2) is largely the work of Simon Funke as outlined in his doctoral thesis [Funke, 2012], however since then the code has been almost entirely refactored – a project on which, as a core developer, I was heavily involved. The product of this refactoring is presented in Funke et al. [2016] to which I was a major contributor and am second author, this section is largely based on the work I contributed to that paper. The functionality of realistic turbine parametrisation (section 3.2.2) (implementing thrust and power curves) was my own work. Beyond chapter 3 all work is my own.

Following this, chapter 4 which is based upon Culley et al. [2016c] documents implementation of analytical wake models in *OpenTidalFarm* and compares them and a shallow-water model with a high-fidelity CFD code solving the 3D Navier-Stokes equations with a RANS turbulence model, in which the turbine is represented as an actuator disk. Having compared the hydrodynamics computed by the various models for individual turbines, the exercise is repeated for a staggered array formation. Once the differences in models is understood, different optimisation methodologies (global, local and hybrid) are tested and compared on two scenarios, with the quality of solutions returned by the various approaches evaluated. Conclusions are then drawn based upon how the models may be used in concert on a real-world array design problem.

In chapter 5 which is based upon Culley et al. [2016a] the second goal from section 1.5 is tackled – to develop a framework in which broader objectives than just power maximisation may be included in the micro-siting optimisation. Here, a model for the cost of submarine cabling was developed and integrated into *Open-*

TidalFarm. Cost of submarine cabling is a particularly interesting objective to integrate as the optimal cable layout for a given micro-siting arrangement must, itself, be found as the result of an integer optimisation, similar to the famous travelling salesman problem. Furthermore, to be incorporated within the *OpenTidalFarm* framework, the cabling function must be differentiable with respect to turbine location. The setup is then demonstrated on several test examples.

In chapter 6 which is based upon Culley et al. [2016b] and Culley et al. [2015] a surrogate model is used to efficiently estimate the optimal number of turbines on a given site. This model is built on top of *OpenTidalFarm* which is largely treated as a black-box non-linear solver computing the maximum power that may be generated by a given number of turbines. Bayesian optimisation is used to optimise that number.

In the final results chapter, the importance of uncertainty quantification is discussed and the difficulties in combining optimisation and uncertainty quantification are explored. A methodology within *OpenTidalFarm* to efficiently produce senstivity analyses of even very large control parameters with respect to the power produced by the array is presented. This enables the designer to, for example, establish that the power produced by the array is sensitive to the bathymetric conditions in certain areas of the domain but not others. This leads into discussions of potential future work concerned with (a) optimisation under uncertainty and (b) developing data assimilation techniques in *OpenTidalFarm*.

The final chapter sums up and draws conclusions on the work contained with this thesis and with especial reference to the objectives laid out in section 1.5.

1.6.1 List of the Author's publications

Lead authorship

Culley et al. [2014], D.M. Culley, S.W. Funke, S.C. Kramer and M.D. Piggott (2014). A hierarchy of approaches for the design of tidal turbine arrays. *5th International Conference on Ocean Energy, ICOE, Halifax*. (9 pages)

Culley et al. [2015], D.M. Culley, S.W. Funke, S.C. Kramer and M.D. Piggott (2015). Tidal stream resource assessment through the optimisation of array design with quantification of uncertainty. *Proceedings of the 11th European Wave and Tidal Energy Conference, EWTEC, Nantes*. (9 pages)

Culley et al. [2016a], D.M. Culley, S.W. Funke, S.C. Kramer and M.D. Piggott (2016). Integration of cost modelling within the micro-siting design optimisation of tidal turbine arrays. *Renewable Energy*, 85, 215-227.

Culley et al. [2016b] D.M. Culley, S.W. Funke, S.C. Kramer and M.D. Piggott (2016). A surrogate-assisted approach for optimising the size of tidal turbine arrays. *Submitted to Internation Journal of Marine Energy*.

Culley et al. [2016c] D.M. Culley, S.W. Funke, S.C. Kramer and M.D. Piggott (2016). Comparison of approaches to tidal stream turbine array optimisation. *Submitted to Renewable and Sustainable Energy Reviews*.

Secondary authorship

Funke et al. [2016] S.W. Funke, D.M. Culley, S.C. Kramer and M.D. Piggott (2016). *OpenTidalFarm*: A software package for the optimisation of tidal turbine arrays. *Submitted to Geophysical Model Development*.

In the above paper I was the main author of the introduction, as well as the sections on code design and tidal turbine specific features, and the sensitivity study example.

du Feu et al. [2016] R. du-Feu, D.M. Culley, S.W. Funke, B. Halpern, S.C. Kramer and M.D. Piggott (2016). The trade off between tidal-turbine array yield and environmental impact: a multi-objective optimisation problem. *Submitted to Renewable Energy* Jan. 2016.

Chapter 2

Models and methods for array design and analysis

Contents

2.1 A family of models 26

2.2 Hydrodynamic modelling for tidal stream turbine applications 33

2.3 Micro-siting design 40

This chapter is structured around a framework which has developed from Culley et al. [2014].

2.1 A family of models

All models are wrong, but some are useful

George E. P. Box, 1976

This section sets the scene for computational modelling for engineering applications in general, and for tidal turbine array design in particular. The goal is to motivate a computational approach to array analysis and design, to establish a framework that enables classification of models and optimisation approaches and then to populate that framework with some examples of computational modelling and design approaches from the existing literature.

Computational models are vital for design and analysis for tidal power applications. The needs of different designers working on different aspects of the problem vary widely, for example a component designer, turbine designer and array designer will have vastly different requirements of their design and analysis tools. This variety manifests not just in the spatial and temporal scales at which they are interested and the physical phenomena which are driving the analysis, but also in the constraints and objectives which parametrise the design / optimisation. For example, wave-current interactions may cause differential loading on the surface-most blade as the turbine rotates, this may cause bending out of the rotor plane which loads the drive bearings eccentrically. Consequently such interactions are of high importance to component designers, while they are of less interest to the array micro-siting designer.

For design-work, laboratory and field testing simply cannot give a sufficiently complete picture of the hydrodynamic processes at work. One driving reason for this is the difficulty of taking measurements of the flow – for example its speed, direction and depth variation. In a field setting, the best that can be achieved is by use of an *Acoustic Doppler Current Profiler* (ADCP) which uses sound waves scattered back from particles in the water to measure the flow velocity in a water column. Best practice suggests that profilers should be deployed and gather data for over a month [IEC, 2015] – getting 'good' data from such devices requires expertise and patience and consequently such data tends to be both temporally and spatially patchy as well as financially expensive. Even at its best, ADCP data provides little more than a short snapshot of the velocity of the column of water directly in line with the ADCP at the time it was recording. Torrens-Spence et al. [2015] illustrates some of the difficulties in generating reliable data – showing that even collocated devices may record very different data.

Laboratory experimentation is also limited for practical reasons – large fast-flowing tanks are not ubiquitous and so recreating large domains with many turbines is not straightforward. Furthermore, recreating realistic boundary conditions in such a set up is nearly impossible. Consequently laboratory testing is generally limited to the study of an individual or small groups of turbines in rectangular (i.e. tank-shaped) domains.

As such, laboratory and field testing must be used in concert with computational modelling. Where real-world data can give isolated snapshots of phenomena and processes, it is the job of the computational model to firstly recreate this snapshot (validation) then 'fill in the blanks' to give a more complete picture as to what

activity is occurring.

While the ultimate goal of computational modelling may be to recreate reality in as accurate a way as possible, limits in computational resource render this an impossible goal. Rather, engineers must identify their objective, recognise the limitations of their computational resource and select models and strategies which (within budget) will help illuminate the behaviour of interest. And so, a range of design questions, coupled with a range of computational budgets begets a family of analysis models which differ in scope and complexity.

Venkataraman and Haftka [2004] proposed a framework for classifying model complexity which helps illuminate the relationships between the plethora of disparate models and modelling strategies in use in hydrodynamic design for tidal power applications. By arranging this set of models into a coherent framework, it becomes clear that certain parts of this family enjoy much more research focus and scrutiny than others. This helps to illustrate the niche into which the work presented in this thesis fits, along with its importance and its novelty.

2.1.1 Classifying complexity

The framework presented by Venkataraman and Haftka [2004] considers 'problem complexity' as a function of the analysis, model and optimisation methodology and illustrates how high complexity in one of these areas must be balanced by simplification in another in order to keep the overall problem complexity computationally tractable.

It is very difficult to be absolutely precise in measuring and comparing 'complexity' between models and approaches. Only rarely can models be directly compared on the same problem; even models which purport to tackle the same type of problem (for example coastal modelling codes) may use different discretisations, require differently constructed meshes or parametrise variables differently, all of which makes inter-model comparison a somewhat subjective endeavour. For this section, 'complexity' will be based upon the author's subjective assessment of computational cost.

Analysis complexity

Analysis complexity is concerned with what it is that one is hoping to discover, which directs the required fidelity of the model output and the extent of its spatial

or temporal domain. To compute the loading from turbulence on turbine blades, for example, requires the flow to be solved at high temporal and spatial resolution and may even involve consideration of the fluid-structure interactions –potentially limiting domain sizes to the order of metres. Conversely resource assessments involve consideration of the interaction between arrays and regional hydrodynamics, requiring domain sizes in the order of 100's to 1000's of kilometers. The demands for high fidelity, large domains or additional post flow-solve modelling (e.g. structural response to hydrodynamic loading) are dictated by the analysis we are hoping to carry out.

Model complexity

Perhaps the most obvious form of complexity is the model choice itself. Developed from first principles, analytical one-dimensional models as by Bryden et al. [2004] may be solved by hand, but provide important insight into the behaviour of the more complex models. At the other end of the spectrum, three-dimensional direct numerical simulation (DNS) of the Navier-Stokes equations captures all the physical processes but at a cost that is infeasible for all but the smallest of problem sizes. Between these two extremes lie a host of models, figure 2.2(b). There is, of course, some inherent linking between model and problem complexity – as certain models will require, for example, high resolution meshes.

Optimisation complexity

It is the third axis of complexity which has received far less attention in tidal power applications; optimisation complexity. Optimisation is concerned with finding the values of control parameters which will either maximise some quantity of interest or provide the most desireable balance between competing design objectives. Funke et al. [2014] demonstrated that the micro-siting design of turbines within an array can dramatically affect the power extracted by the array, however micro-siting design optimisation remains an under-scrutinised area.

Many papers concerned with resource assessment or on the effects of arrays on regional hydrodynamics overlook turbine placement as a factor all together. For example Fairley et al. [2015] who use MIKE3 [DHI, -2016b], look at the effects on sediment transport of arrays in the Pentland Firth using assumed array layouts – the array design may have a notable impact upon this transport [du Feu et al.,

2016]. Conversely Karsten et al. [2013] makes more effort in locating turbines in hydrodynamically logical locations by placing turbines into the flow 'greedily' (turbines are placed in the domain one at a time at the location of highest flow speed) to model an array of 1000 turbines in the Midas Passage. This provides a plausible array layout – but the impact of these chosen turbine locations on the conclusions that are drawn regarding the performance and impact of the array is not explored because doing so in such a model would be "extremely difficult, if not impossible" [Karsten et al., 2013].

Throughout this thesis several optimisation approaches will be employed including global, local and hybrid methods. These will be considered in detail later in this chapter, however figure 2.2(c) shows where they may sit on an axis of complexity.

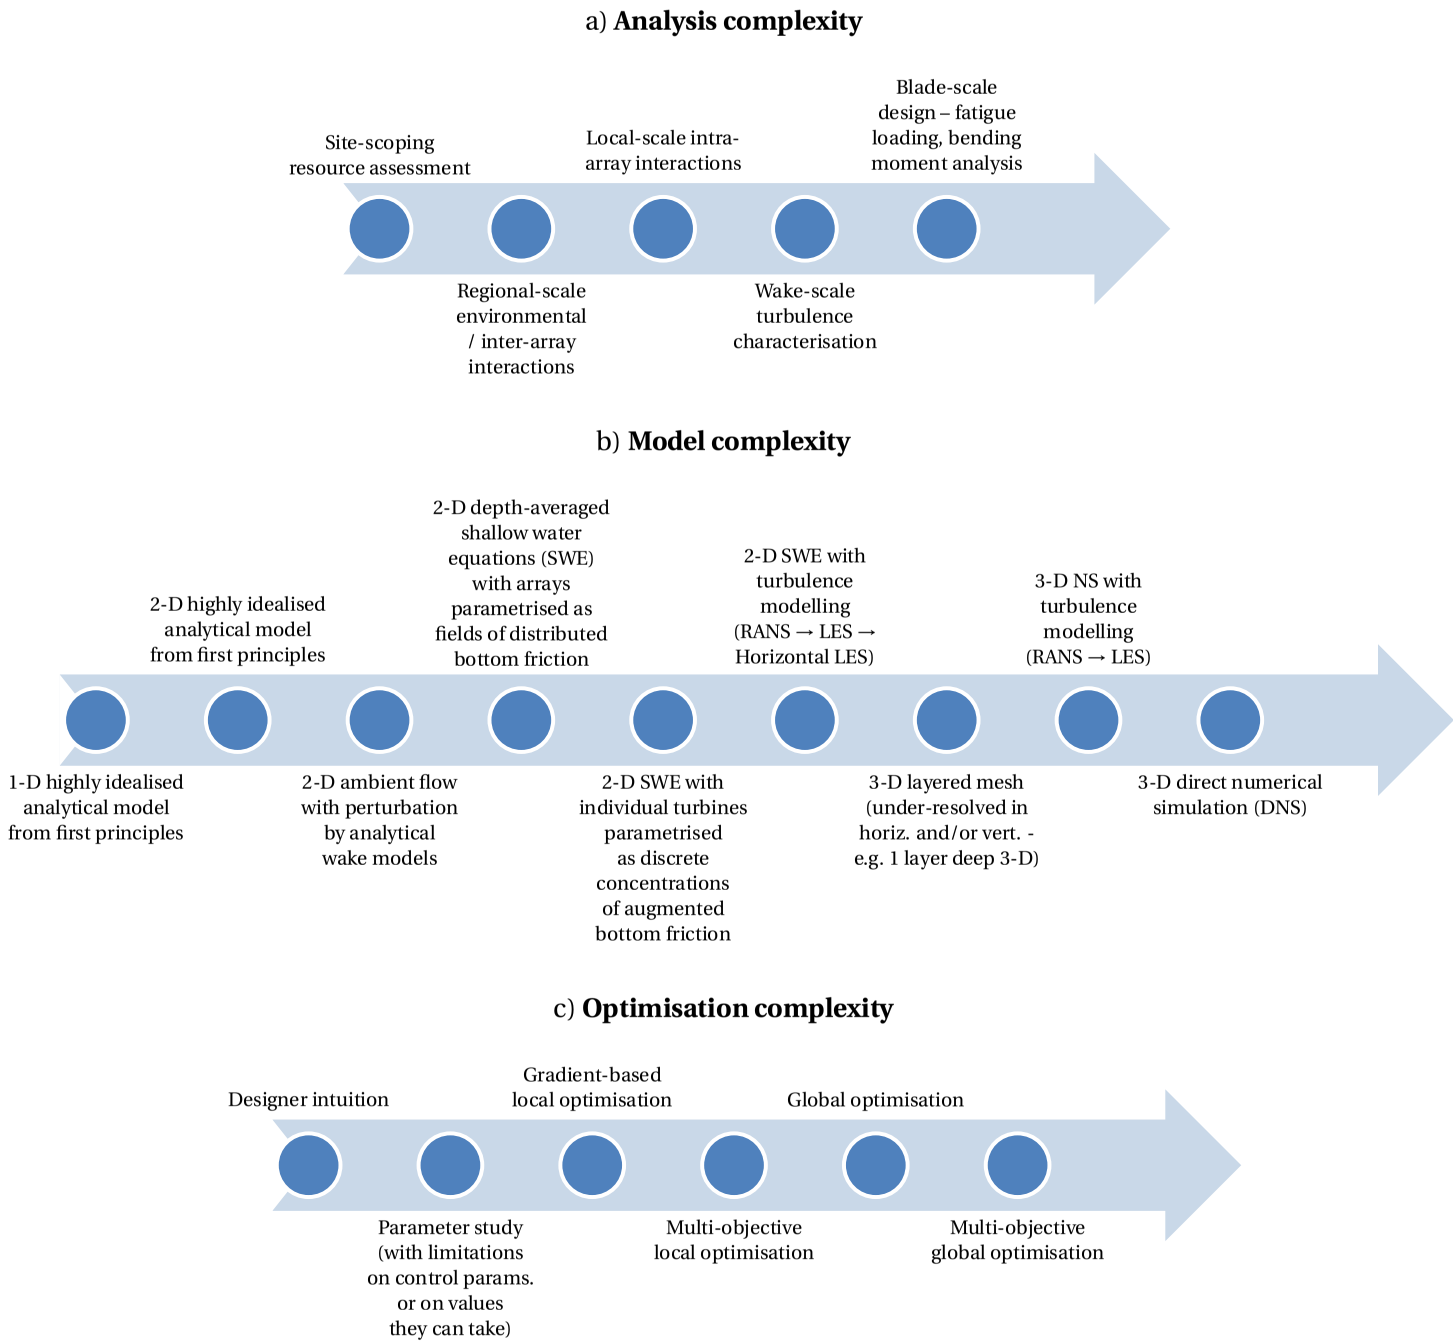


Figure 2.1: A framework for characterising problem complexity as a function of analysis, model and optimisation complexity (presented in figures (a), (b) and (c)). Examples of each are arranged (with a degree of subjectivity) in increasing order of complexity.

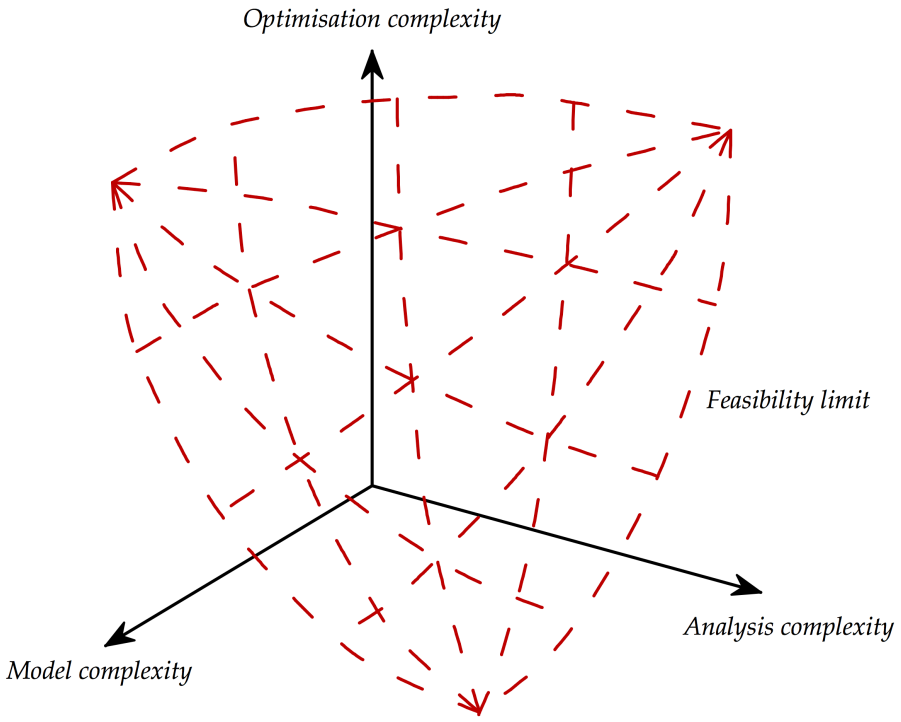


Figure 2.2: Illustrative graph of problem complexity as a function of model, analysis and optimisation complexity with limits of computational budget shown by red surface.

Problem complexity can, therefore be thought of as a three-dimensional space (figure 2.2) with a front (shown in red on the figure) representing the limit of available computational resources. This illustrates the trade-off which must be considered when choosing a modelling approach for a given design problem. This approach to classification mimics similar discussions from within the climate-system modelling community, where models of different physical phenomena must be coupled and run for large time-scales over large (global) domains. In Claussen et al. [2002] the three axes of complexity are processes (the number of processes being modelled), detail of description (the level of detail the model for each process provides) and integration (the extent to which the individual models are coupled). A fourth axis is also common which is related to the number of ensembles to be run for the scenario. These classification systems provide a useful formalism for choosing a solution strategy for a given problem.

2.2 Hydrodynamic modelling for tidal stream turbine applications

Before looking at design optimisation in the next section, this section is concerned with the methodologies for *evaluating* a given design. Of course this can be done physically in the field or at scale in the lab, however as has been established this is generally impractical, so this review will be limited to computational methodologies.

Culley et al. [2014] and Adcock et al. [2015] organised models into a hierarchy based upon the spatial scales at which models are employed, and explored the transfer of knowledge between these scales that is needed to achieve a holistic understanding of how an array comprising particular types of turbines will interact with the hydrodynamics at a site in question. This requires blade-scale models to feed understanding into larger scales (so that simplified turbine representations in larger-scale models are well parametrised) and for regional scale models to feed understanding into smaller-scale models (so that the boundary conditions of array-scale models are accurate). This communication is represented in figure 2.3.

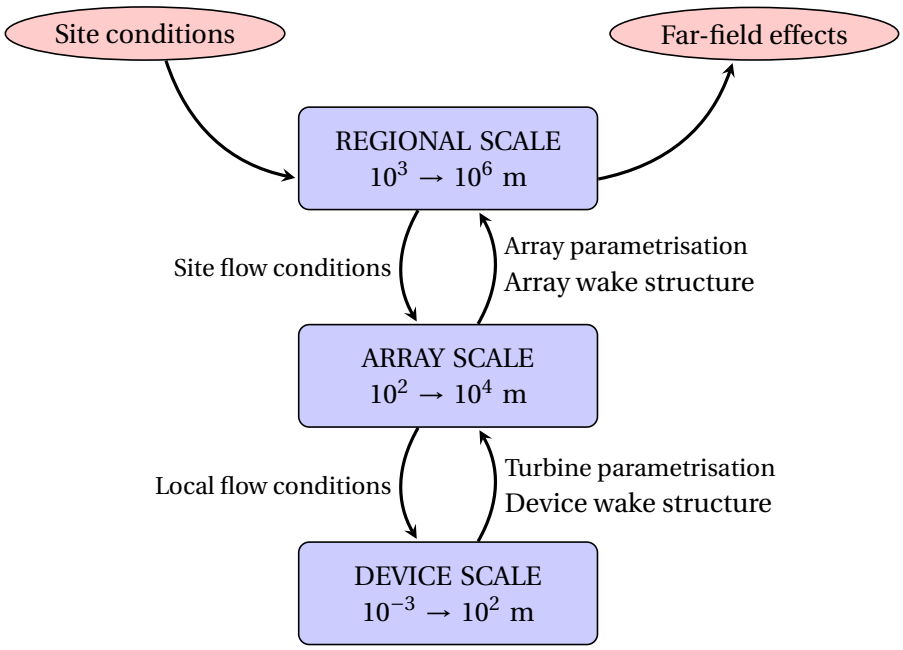


Figure 2.3: Schematic showing communication between different model scales.

2.2.1 Device scale models: From individual blades to individual turbines

Analysis of the rotor blades of tidal turbines is focused around (though not limited to) two main methodologies. The lowest fidelity methodology – using blade element momentum (BEM) theory – is used for two main reasons; to understand the structural loading on the blade (and transferred through it to the structure) or to optimise the shape of the blade. The model is a combination of momentum theory with blade element theory. In the former, the swept area of the spinning blades is subdivided into annuli, and the mass flow rate through each annulus can be computed as a function of the annulus geometry and flow speed. Blade element theory characterises the blades in each annuli as a single blade section and thus can resolve the forces acting on the section based upon its lift and drag properties. Optimisation of the turbine blade using this model or just blade element theory proceeds by defining the blade sections as Bezier curves linking a number of control points [Sale et al., 2009]. Optimisation methodologies are concerned with how adjustment of the control points in two dimensional space alters the section shape and affects the objective.

When the analysis question is more complex, and rather than determining the loading on the structure the designer is interested in the hydrodynamic structures the blades induce, a much higher fidelity model is required. This generally involves solving the Navier-Stokes (NS) equations using numerical approaches (discussed in section 3.1.1). The Navier-Stokes equations cannot be solved analytically on such complicated problems, so the domain is discretised and an approximate solution is computed over the discrete approximation of the domain. The chord Reynolds number of turbine blades in a tidal environment is high $Re_c = 2.1 \times 10^5$ [Bachant et al., 2014] and so advection-dominated, highly turbulent flow is expected. This means that energy in the flow cascades to smaller and smaller length-scales which will result in instability in the numerical methods if not captured by coarse discretisations of the numerical model or through some numerical stabilisation scheme. Direct numerical simulation (DNS) tackles this directly by ensuring that the domain discretisation is sufficiently fine so as to capture and model these length-scales. For realistic applications this results in very large numbers of degrees of freedom in the discretised domain and an intractably large problem for even modern high-performance computer (HPC) architectures.

A solution is to introduce an additional, dedicated, turbulence model to deal with sub-grid-scale phenomena. The two main classes of turbulence model are large eddy simulation (LES) and Reynolds averaged Navier-Stokes (RANS). In RANS

modelling, turbulence is treated as random velocity variations from a stable mean, the idea is to define the viscosity parameter in the NS equations as the sum of the background viscosity and the turbulent viscosity. In regions of the domain where turbulence would form, the turbulent viscosity increases to prevent this but still accounts for the impact of the turbulent motion on the mean flow. The model remains stable, and the distribution of turbulent viscosity over the domain can be used as a proxy for intensity of turbulence that would be forming. The idea behind large eddy simulation is similar, however larger eddies are allowed to form (where they can be properly resolved) and only the smaller-scale turbulence is damped out.

Within these RANS and LES classes are various turbulence models which are implemented in different ways and add different numbers of additional transport equations or closures to the standard NS equations. Hence, while the turbulence modelling methodology requires less high resolution meshes than DNS, it is still complex and computationally expensive. Examples of this sort of study are exemplified by Hee Jo et al. [2012], Shi et al. [2013] who use the RANS equations with a shear stress transport (SST) turbulence model to study hydrodynamic performance and flow separation on the rotors of HATTs. Turbulence modelling approaches continue to be combined with other advanced numerical techniques, for example McNaughton et al. [2014] who uses a combination of both LES and RANS modelling with a sliding-mesh interface at the turbine and see improved correlation with experimental results, but at substantial computational cost.

2.2.2 Array scale models: From idealised- to more realistic domains

As the spatial scale and number of turbines increases, so the designer must reduce their ambitions on the complexity of analysis that may be performed. At the array scale, it is no longer tractable to model individually spinning blades, and so the turbine representation must be abstracted. The quality of these abstractions is validated against lab data and evaluated through comparison to the fine-scale models discussed in the last section and thus knowledge is transferred from different model scales.

One methodology for modelling turbines at this scale is the use of an actuator disk placed in a 3D RANS model [Batten et al., 2013]. An actuator disk is a circular momentum sink which corresponds to the area swept by the turbine blades. This momentum sink may capture the blockage effects of the turbine and model

its effects on the ambient velocity field accurately, however the vortical structures observed in figure 1.10 will not be replicated. This may, to some extent, be addressed by adding sources of turbulent kinetic energy (TKE) at the location of the disk Nishino and Willden [2014]. Determining how much momentum should be removed and how much TKE should be added is done through validation against either laboratory tests on appropriately scaled porous disks or the kind of high fidelity computational modelling discussed already.

As an indication of the expense of this methodology, Abolghasemi et al. [2016] solved the flow for a site-scale domain over circa 7 hours of a tidal cycle taking approximately 48 hours of computational time on 240 cores. Lab-scale validation simulations against data from Stallard et al. [2013] (three rotors in line) took approximately 31 hours on 16 cores and used mesh adaptivity to ensure the most efficient spatial discretisation at each time step.

Clearly this level of resource-usage for a solve is impractical if iterative optimisation is required. One way to dramatically decrease the complexity of the analysis is to simplify from fully three-dimensional modelling to simplified analogues. For example the depth can be discretised into a small number of layers by extruding the plan-view elements into prisms, this means that the 3D NS equations can still be used and vertical mixing is still, to some extent captured.

Probably more common is to fully depth average the NS equations to arrive at the Shallow Water (SW) equations. In such models, three-dimensional turbines in the vertical plane are typically represented by two-dimensional patches of locally increased bottom friction [Divett, 2014]. Compared to the horizontal length-scale, the water depth is very small and thus vertical mixing processes are limited – the size of the flow structures responsible for mixing is restricted. Having said that, in certain situations vertical effects can influence larger-scale mixing, so in situations where one might expect significant shear in the vertical plane or secondary flows to be induced, the validity of the depth-averaging assumption must be inspected and questioned. Unfortunately these conditions arise in the wakes of obstructions in the flow – for example from islands, headlands, steep bathymetric features or indeed turbines – which are the sort of features one would expect when modelling tidal stream turbine arrays located near coastlines. These limitations are the product of a finite computational budget and while they must be acknowledged openly, results from such models still provide great insight and are of great utility – especially in conjunction with some of the other modelling approaches discussed here.

The reduction in complexity of such depth-averaged models facilitates the in-

vestigation of much larger domains modelled for much longer periods, however the turbines themselves still need to be adequately resolved – or at least the manner in which turbines are represented in the *discretised* domain must be carefully considered. In *MIKE 21* [DHI, -2016a], a two-dimensional flow solver for coastal and fluvial environmental modelling applications, the augmented bottom friction which simulates the turbine thrust is spread over the single element in which the centre of the turbine is located. Thus, while the mesh at the turbine location does not need to be sufficiently fine to resolve the turbine, the representation of the turbine is mesh dependent. This becomes problematic, as described by Kramer and Piggott [2015] when the size of the element containing the turbine tends to the size of a turbine and below; as the friction that is applied over it increases (so as to apply the same thrust over a shrinking area). This effects the hydrodynamic interactions between turbine and array, as well as the computation of the power extracted by the turbine. In *OpenTidalFarm*, conversely, the patches of friction representing each turbine can be summed to give a turbine friction field for the array which is zero over the majority of the domain and non-zero at the turbines, this field is then discretised on the same mesh as is the domain. This approach avoids the problems suffered by the MIKE 2D approach (turbines may be spread over multiple elements – but must be resolved), however does require a sufficiently fine mesh at the location of the turbines (a mesh element size of a quarter to a half the turbine radius would be typical).

Others have used the actuator disc modelling approach in large-scale two-dimensional models, whether it be using a patch of drag or simply a line; in Lalander [2009] this took the form of the addition of a shear stress while in Draper et al. [2009] it took the form of a drag force.

An alternate methodology which was used extensively for resource assessment and some early array modelling work [Triton Consultants, 2002, Tarbotton and Larson, 2006, Polagye and Bedard, 2006, Polagye et al., 2006, Hagerman and Bedard, 2006] and is still used in some modern software [Thomson et al., 2011] was the *perturbation approach* (also variously called the *undisturbed flow approach* or the *flux method*). This methodology was used by consultants Black and Veatch to produce for the UK Crown Estate a technical report identifying areas of high tidal resource with potential for development. The approach involves use of a standard coastal model to compute the ambient (i.e. without any turbines) flow regime through a region of interest. Turbines are modelled via analytically computed wake deficit models which have been developed in the wind industry and

tweaked for tidal application. These wakes are combined with the ambient flow field via a process of combination modelling (where wakes overlap) and superposition. While wake models often do a sound job of matching the experimental data for the velocity deficit in a turbine's wake, they are unable to capture local or global blockage effects as the turbines are modelled entirely uncoupled from the flow. This methodology has its detractors [Garrett and Cummins, 2007, Whelan et al., 2009, Vennell, 2010, Nishino and Willden, 2012], with Bryden and Couch [2006] asserting that "for high extraction rates the general validity of this assumption [of uncoupled modelling] is questionable". However, for applications where large areas need to be surveyed or a global optimisation approach is to be used for array design optimisation, there may be applications of such a simplified modelling approach – so long as it is employed with understanding of its significant limitations. We will meet this methodology in the context of array design optimisation later in this chapter and see it compared to the *OpenTidalFarm* approach discussed above in chapter 4.

One final note on the reduction of complexity for modelling larger spatial scales is that as noted by the IEC, "it is not currently known what scale, frequency and magnitude of current variability (resulting from eddies/turbulence) are important" (to feed into array design) [IEC, 2015]. Detailed high-fidelity modelling, along with analysis of field-deployments of turbines will begin to fill these knowledge gaps, and as this happens, lower-fidelity 'engineering' models should be developed tailored to the specific requirements of array-scale modelling for tidal turbines. This is another example of knowledge transfer between modelling scales – just as parametrisation of turbines in lower fidelity models is developed from higher fidelity models working at smaller spatial scales. Of course, increased data from field deployments will also improve validation of models and develop understanding as to the limits of modelling for tidal-stream power capture design; another key unknown [IEC, 2015].

2.2.3 Regional scale models: Capturing turbines at large scale

To fully capture the regional hydrodynamic impact of a turbine array, ideally the domain boundaries should extend to the edges of, or beyond, the continental shelf [Adcock et al., 2011]. Unstructured meshes (such as in figure 3.1) allow small element sizes where high resolution is needed and large elements away from regions or features of interest.

Once again the turbine representation can be abstracted, this time rather than individual turbines being modelled as discrete areas of increased friction, the whole array may be modelled as one combined area of friction. This means that the mesh resolution required at the site of the turbines is greatly reduced and the turbine site can be captured with a handful of elements. This is a popular choice for resource assessments using numerical models, for example Sutherland et al. [2007] and Karsten et al. [2008] who use it to evaluate the resource in the Johnstone Strait, Canada with the *TIDE2D* finite element model and the Minas Passage, Canada, with the *FVCOM* finite volume code respectively. The results from both of these models showed good agreement with an application of the idealised one-dimensional model of Garrett and Cummins [2005] and its extension Garrett and Cummins [2004], although Sutherland et al. [2007] noted that in situations where the flow was able to divert away from the tidal site, the analytical theory was no longer applicable and thus agreement was less good.

Though comparatively simple, the one-dimensional models of Garrett and Cummins [2005] and Garrett and Cummins [2004] are used extensively in resource assessments [Polagye et al., 2009, Blanchfield et al., 2008] and are useful at preliminary quantification of the impact of in-stream energy extraction on barotropic flow.

Rather than using a constant friction across the whole channel as in the papers discussed above, the natural extension is to apply variable friction (so as to model the fences of turbines) to portions of the width of the strait (for example to simulate clear lanes for navigation). This approach was implemented by Funke et al. [2015b], in which optimisation of the continuum of friction which represented the farm enabled the most profitable array size and general layout to be established over the actual areas available for leasing for tidal farms in the Pentland Firth, UK. Four farms were modelled in the same domain *and optimised* to study the benefits of farm designers on nearby sites coordinating their designs. The domain comprised over 320,000 elements and covered the Scapa Flow and Pentland Firth in Scotland, the minimum mesh element size in the turbine areas was 25 m – if this value were decreased (so as to resolve and model individual turbines) the overall number of elements in the domain would rapidly grow to the limit of computational tractability.

2.3 Micro-siting design

Figure 2.3 demonstrated the communication and linkage that must exist between different modelling scales of a tidal array. Optimisation can also occur at each of these levels to improve the design of the object driving the hydrodynamics being studied at that scale. For example, in high fidelity modelling of an individual device, the blade and nacelle are the objects driving the hydrodynamic structures being studied – and may be optimised. Conversely at the regional scale it is the location and size of the array affecting the hydrodynamics. Once again, regional modelling must feed downwards to smaller scales and these smaller scales must feed back up to facilitate parametrisation of turbines and arrays in the larger models. With an optimisation process at each scale, the modelling process of figure 2.3 becomes the *design* process of figure 2.4. Henceforth, this section is concerned with design optimisation at the array scale.

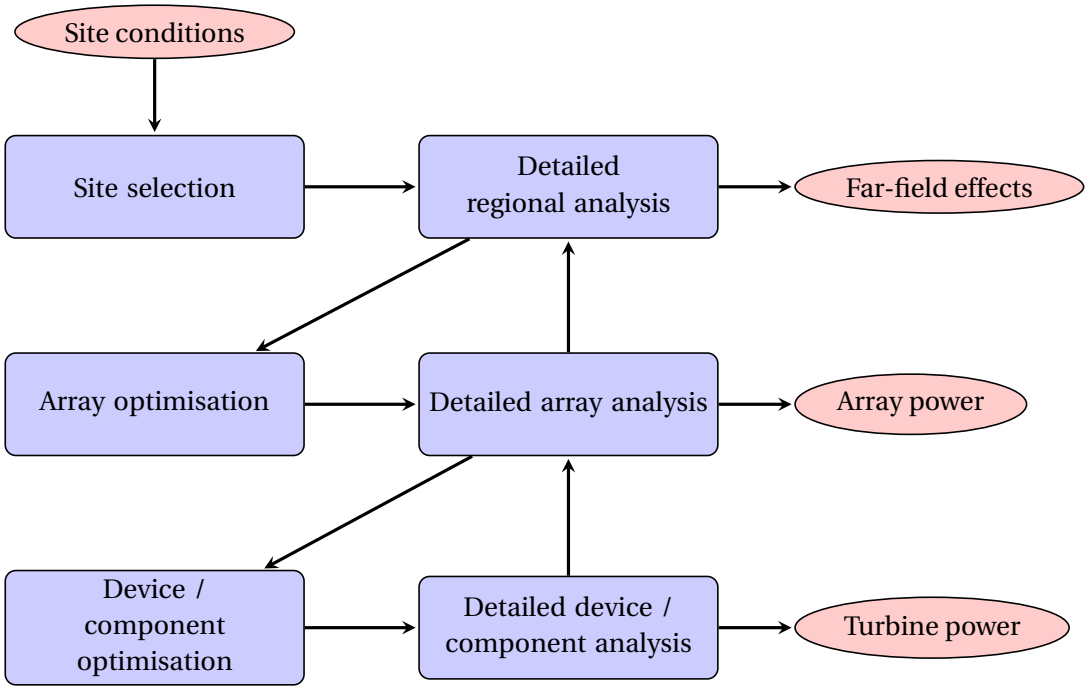


Figure 2.4: Schematic showing communication between different model and optimisation scales.

For the reasons discussed in the introductory chapter, the optimisation of the micro-siting design of tidal turbine arrays is a formidable computational challenge. These reasons can be summarised as:

- Problem length scale varies from individual turbine blades (potentially 10^{-3}m) to whole coastal regions (circa 100 km).
- Turbines are highly sensitive and strongly coupled to the flow resulting in non-linear behaviour, meaning that the model needs to be reevaluated whenever the optimisation changes the design variables.
- Computational budget of designers is limited – increase in the complexity of optimisation requires a decrease in the fidelity of the underlying model.
- Large numbers of control variables – two-dimensional turbine coordinates give $2N$ control parameters, however the designer may be interested in additional parameters (for example individual turbine thrust).
- Large number of performance metrics – different stakeholders in the project will have different concerns so environmental, technical and financial considerations need to be defined quantitatively.
- These performance metrics are not always directly comparable / trade-able requiring multi-objective optimisation or conversion to a common intermediate such as currency.
- Lack of industrial experience and validation data, not just for the hydrodynamic interactions between turbines and flow but also for costings, maintenance scheduling and a whole host of other concerns that may feed into the performance metric. This is because no turbine arrays have yet been commissioned.

Thus micro-siting optimisation is a complex, multi-faceted problem, and as no large scale arrays have yet been installed (and proposed demonstration arrays compose a handful of turbines – which do not require too much optimisation) it is one that has received limited attention to date.

Conversely, an analogue of the problem has been considered for a number of years in the comparatively mature wind energy sector. There are some strong parallels between the problems of wind farm and tidal current farm micro-siting optimisation, which make study of the former problem a sensible place to begin consideration of the latter. However, while many of the tools and approaches may be applicable, the key physical and technological differences discussed in the introductory chapter of this thesis make the problems sufficiently divergent as to merit independent consideration.

The goal of this section is to present the current state of the art in tidal array micro-siting design. It will begin by looking at the wind industry – tracing through the development of array layout design decision making to the present day, before turning to the tidal analogue and determining which methods may be cross-adopted.

2.3.1 Wind turbine array optimisation

Despite differences in the physics driving wind power as compared to tidal stream power, wind farm design has the benefit of many more years of development effort; both in the models and tools for design and in actual array deployment. Indeed, if the rate of academic publications is used as a metric, wind energy seems to lead tidal energy by circa 20 years. While always being mindful of where the nature of the resources and technologies to capture them are divergent, there are valuable insights and methodologies that can be migrated from wind to tidal.

An early attempt to automate the process of wind turbine placement within a farm was presented by Mosetti et al. [1994]. Prior to this, turbine spacing had been determined empirically, with spacings between turbines given as multiples of turbine diameters parallel and perpendicular to the prevailing wind direction. In their study, a 10 by 10 grid was used, with each cell allocated a binary value relating to whether a turbine was present or not. A simple wake model was applied to adjust the ambient flow field for each turbine and the power it generated found as a function of the turbine characteristics and the wind characteristics (direction, intensity and the time over a year that that wind type occurs).

The algorithm found the optimum number of turbines and their location by solving the optimisation problem

$$\min_{N,\mathbf{m}} \quad \frac{1}{P_{\text{tot}}(N,\mathbf{m})}\omega_1 + \frac{C_{\text{tot}}(N)}{P_{\text{tot}}(N,\mathbf{m})}\omega_2, \tag{2.1}$$

[Mosetti et al., 1994]

where P_{tot} is the total power output per year of the farm – a function of the number of turbines, N and their locations, $\mathbf{m}$; C_{tot} is the total cost per year of the farm (calculated solely in relation to the number of turbines); and ω_1 and ω_2 are weighting constants so as to properly balance cost and income from power.

A genetic algorithm was used to search the solution space to find the optimum

number of turbines and their locations. Genetic algorithms will be more fully explored in chapter 4 where they are contrasted in detail with gradient-based and hybrid methods. Simply described, in a genetic algorithm, a possible solution is represented by a *chromosome*; a vector representing an encoding of that solution – for example a list of the binary values for each cell (0 indicating that no turbine is present, otherwise 1). A *population* of random chromosomes is initialised and each one is evaluated against the functional of interest – in this case (2.1). The best performing few solutions are selected and combined (or *bred*) together and *mutated* in various different ways to generate a new population, this population is evaluated to begin the loop anew. The process continues until a given convergence criteria is met or a limit on the number of iterations is reached. The Mosetti et al. [1994] algorithm used a population of 200 to 300 chromosomes and converged in a few hundred iterations.

The computational cost of employing a genetic algorithm scales poorly with the size of the solution space (a 10 by 10 grid in this example). Thus the available computational resources enforced a certain level of crudity on the model developed by [Mosetti et al., 1994] and limited its practical usefulness. However, the work served as a proof-of-concept, showing that the computational optimisation of turbine placement was feasible. As computational resources developed, so too could the number of independent variables considered and the complexity with which they were modelled [Grady et al., 2005]. In González et al. [2010] a genetic algorithm was used which combined a detailed commissioning, operational and decommissioning cost model with a more intricate turbine generator and wake model. This built upon the improvements made by Grady et al. [2005] to the Mosetti et al. [1994] model. While the results compared favourably to the preceding work, the approach retained the attribute of scaling poorly as the number of potential turbine location cells is increased.

Of course genetic algorithms represent just one subset of a whole host of optimisation tools, many of which have also variously been applied to the problem of array optimisation. While the majority of these also formulated the optimisation in a discrete manner (i.e. restricting turbine placement to a grid of cells), Aytun Ozturk and Norman [2004] developed a greedy heuristic for a continuous formulation of the problem, and Eroğlu and Seçkiner [2012] applied an ant colony algorithm, building on work by Kusiak and Song [2010] both of which also used a continuous formulation. This meant that solution chromosomes were encoded as vectors in $\mathbb{R}^c$ rather than $\mathbb{Z}^c$ (where c is the number of control variables) making their breed-

ing and mutation operators much easier to formulate and thus those operations more efficient to perform. The continuous formulation facilitates greater realism in representation of the problem, improving the scope for finding better optima and simplifying chromosome manipulation, however the search space also increases and constraints must be reformulated; both of which result in increased computational cost.

In general, the approaches to wind farm layout optimisation utilise global search algorithms, which combine an element of randomness with some capability to focus toward the globally optimum solution (given a sufficient number of iterations to do so). The main drawback to such approaches is that they scale poorly as both the number of control variables (i.e. the number of turbines and their possible locations) and size of the problem being considered (i.e. the size of the domain and the resolution at which it is modelled) are increased. Additionally, the models generally assume there is little or no feedback on the flow from the presence of the turbines – beyond the super-positioning of a wake model onto it. This is because the number of iterations required precludes full simulation of the fluid dynamics – the optimisation complexity has increased necessitating a reduction in the modelling complexity. As has been outlined, wake recovery is much more rapid for a wind turbine than for a tidal current turbine, consequently this modelling assumption requires scrutiny before application to the tidal current case.

While the focus of array design optimisation across the renewables industry has generally been on maximisation of power yield [Thomson et al., 2011], some work has been undertaken tackling design optimisation with alternate objectives. One example of this is the work by Kwong et al. [2013] who used a multi-objective optimisation approach to consider the turbine layout with respect to noise propagation and energy generation or Tong et al. [2013] who traded land use against capacity factor. Conversely, Lumbreras and Ramos [2013a] looked at optimisation of the electrical layout for wind farms to maximise objectives relating to cost and reliablility and Hofmann [2011] reviews a large number of models which focus on the modelling and optimisation of operations and maintenance (O&M) strategies for offshore wind farms based on a cost / benefit approach. The conclusions of that latter review were that models and optimisation approaches existed for most aspects of offshore wind farm design but that none covered a broad range of aspects and many were commercial and proprietorial. For the most part, array layout decisions remain based upon power maximisation strategies.

2.3.2 From wind array modelling to tidal array modelling

The core differences between wind and tidal *turbine* modelling are listed below and the effects of these differences are also felt in *array* analysis and design:

- Wind turbines are less affected by blockage effects from lateral flow constraints (as these are geographically uncommon) and are unconstrained vertically by a free-surface and so benefit from atmospheric wake re-energisation.
- The direction of the tidal resource is essentially limited to two dominant modes – flood and ebb – and can to a large extent be precisely predicted. Meanwhile the wind direction varies and may come from any direction and is often difficult to predict.

The constraints on the flow and resulting blockage effects change the behaviour of the turbine/flow coupling and interaction, and thus the required models and analysis are difference. In addition, the simpler directionality of the tidal resource (figure 2.5) has a reduced complexity of the formulation of the micro-siting optimisation: Evaluating the array performance over an average (i.e. neither spring nor neap) tidal cycle may be achieved simply by choosing a representative tidal cycle for the site in question and modelling it. In contrast, the wind can blow from any direction, so sites are characterised by wind rose plots which chart the probability that the wind will blow from a given direction at a given speed on the site in question. This variability is generally modelled via a 'binning' approach; the wind rose is discretised into representative bins of directions and speeds, and the power results from the turbine farm are then weighted by the probability of occurrence of those conditions.

2.3.3 Approaches to tidal array micro-siting design

Micro-siting optimisation for tidal turbine arrays is far less developed. While the first industrial tidal array is to be installed later this year (2016) by MeyGen in the Orkney Islands, this will initially only comprise four turbines and is therefore not of sufficient size to require a formal optimisation approach. In addition, it is unlikely that power maximisation is the foremost driving design objective – rather ensuring the robustness, reliability and durability of the devices that are installed. Consequently, few alternate approaches have thus far been developed and these are yet to proven by industrial application. Assuming that the initial MeyGen

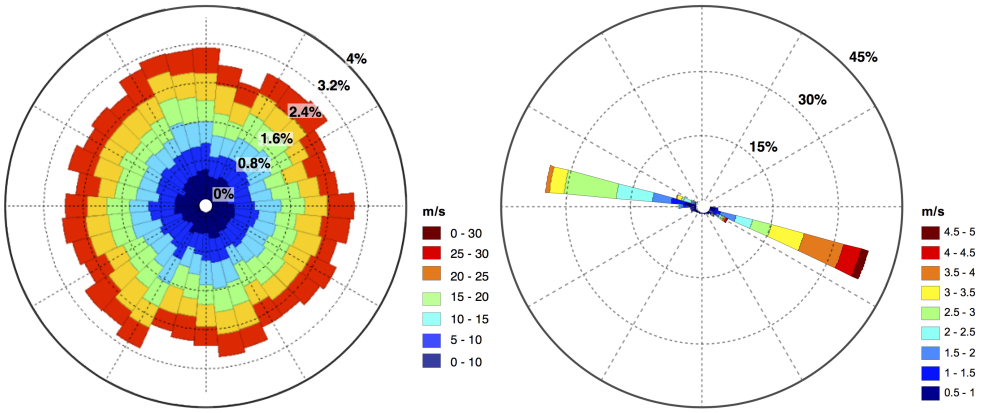


Figure 2.5: Comparison on prevailing resource direction for wind (left) and tidal currents (right). The wind data is taken from Boston Harbour, USA, while the tidal data is taken from ADCP measurements in the Midas Passage, Canada.

project is successful, the industry will begin to focus on array design with more urgency. Given this, the sparsity of these array optimisation methods illustrates the timeliness of this research.

Armed with the complexity framework for classifying models we can begin to consider existing turbine micro-siting optimisation methodologies in a coherent way. For analysis of fixed designs, modelling approaches sit in the xz plane of the graph in figure 2.2 and to date, the majority of research effort in computational approaches for the design of tidal turbine arrays has been limited to analytical modelling rather than rigorous design optimisation. Moving from analysis to design and optimisation thus requires a compromise in analytical and model complexity, necessitating certain simplifications and assumptions.

Existing approaches to the optimisation of turbine arrays can be split into four categories;

- Designer intuition with high fidelity model (e.g. Fairley et al. [2015]).
- Parameter study of row spacing with medium fidelity model (e.g. Divett [2014]).
- Local optimisation of turbine locations with medium fidelity model (e.g. Funke et al. [2014]).
- Global optimisation of turbine locations with low fidelity model (e.g. Thomson et al. [2011]).

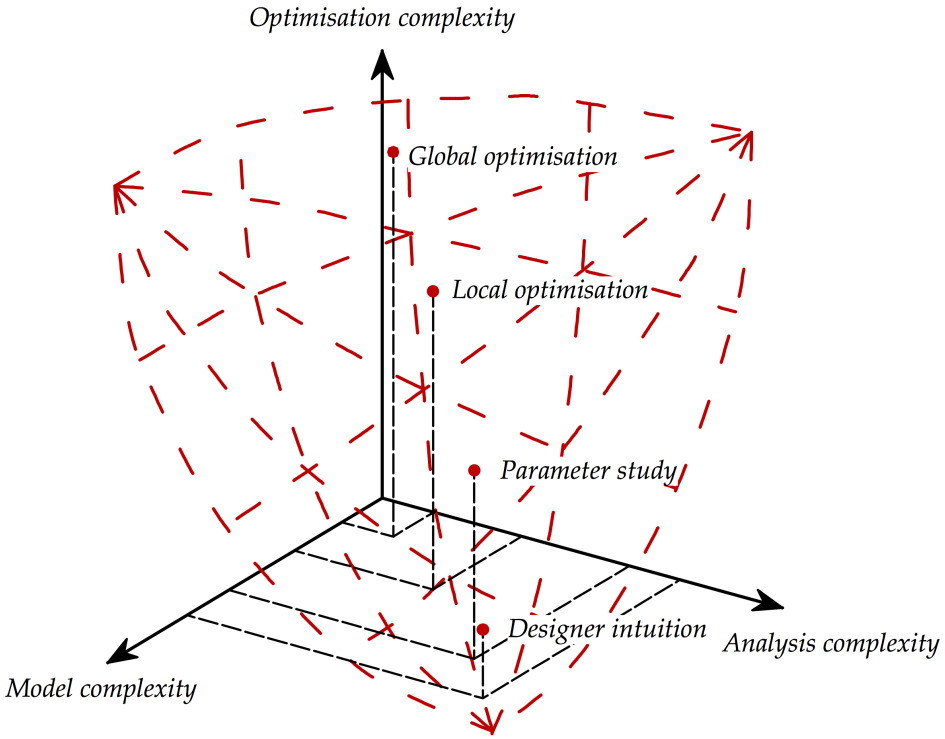


Figure 2.6: Graph showing how different modelling approaches use computational resources in different aspects of problem complexity. Note that model complexity and accuracy of model prediction are not necessarily related linearly – the minimum model complexity required to resolve the key dynamic features may be dependent upon the flow field – which may change in the course of an optimisation.

In the first case, the designer might use a high fidelity coastal hydrodynamic model and run it for a month (model time) to capture the spring-neap variations on a 3-dimensional domain. In such a model, turbines are often modelled as actuator disks and sub-grid scale turbulence models are often used. Such a model working on a coastal-scale spatial domain and month-long temporal domain is a significant task even for a large high-performance computer (HPC). Thus the designer's capacity to optimise is severely limited and they can only realistically try one or two different array designs. Consequently 'designer intuition' is the optimisation mechanism here and so the quality of designs will be dependant upon the simplicity of the scenario and the experience of the designer.

In the second case, turbines are assumed to be arranged in (often staggered) rows and a parameter study may be conducted to determine the optimal spacing between turbines in a row and between the rows themselves. Exploration of the effect of these two control parameters requires multiple runs of the model

with the row spacing altered each time. This increase in optimisation complexity means that the simpler two-dimensional shallow water model is often used, with turbines parameterised as areas of augmented bottom friction. Alternatively if three-dimensional models are used, the domain is often significantly simplified to a rectangular channel over a time-scale that models a single sinusoidal 'tide'. Perhaps best known for such work is Divett [2014] who uses a 2-D model with large-eddy simulation and an adaptive mesh to explore array design under the constraint that turbines are arranged in rows. The high quality of this approach and the variety of layouts considered has provided a good basis for comparisons with *OpenTidalFarm* in Jacobs et al. [2016].

The concept of strict turbine rows is a popular one in the literature, however in realistic sites with irregular bathymetry and complex coastal geometry – such as headlands, islands and inlets, flow through the turbine site often curves and may have significant local variation. Such features may promote curved – rather than straight – rows, or placement of individual or small groups of turbines out of rows so as to exploit a high-flow 'hot-spot' in the turbine domain. This is an idea that is explored in some detail in chapter 4. Consequently, these parameter studies may not be directly applicable for commercial use, but they are important in furthering our understanding into the behaviour of interacting turbines. For example, in Draper and Nishino [2014] the authors extend linear momentum actuator disc theory to study how best to arrange two fences of turbines in laterally confined and unconfined flow. This helps formulate conclusions about spacing in multi-row turbine farms which can help inform a designer and enable them to better evaluate array designs. Certainly, experience has shown that the array design returned by local and global optimisation (the third and fourth cases in the list above) is not always readily interpretable – that is to say it is not always easy to understand why the optimal arrangement looks like it does. High order computational approaches with low order optimisation (parameter studies) like Divett [2014] and analytical work on row spacing and blockage, for example Vennell et al. [2015b] and Draper and Nishino [2014], are very useful in making sense of the products of more complex optimisation approaches – for example the channel-scale dynamics and ducting effects which arise.

If turbines are not forced into conforming rows, but instead allowed, during the optimisation, to move more freely, the number of control parameters grows, as does the optimisation space. There are two approaches to dealing with optimisation problems in such large parameter spaces – to use global or local optimisation.

In the third categorised approach to array micro-siting design, a shallow-water model is also used to solve the flow, but computationally economical access to sensitivity data enables gradient-based optimisation approaches to be employed. Due to the difficulty of deriving and implementing the adjoint model which provides the access to sensitivity data, this methodology is unique to *OpenTidalFarm* [Funke et al., 2014]; the further development of which is the focus of this thesis. A detailed discussion of the strengths and weaknesses of this approach is given in chapter 4.

The fourth case represents (to the author's knowledge, at the time of writing) the only other array design optimisation methodology that goes beyond two-parameter row spacing optimisation. Here, gradient data is not available, so the user is forced to use global optimisation which requires orders of magnitude more iterations than do gradient based optimisation approaches. This massive increase in optimisation complexity requires a similarly significant reduction in modelling complexity. The result is the adoption of a perturbation approach in which the precomputed ambient flow over the domain is adjusted by analytically computed wake-deficit models (sometimes with 'blockage correction'). One of the big assumptions that facilitates this simplification is that the turbines are not fully coupled to the flow and exploring to what extent this simplifying assumption may compromise the value of such a model is one of the focuses of chapter 4.

This approach is used in *TidalFarmer*, a commercial tool from DNV-GL [Thomson et al., 2011] which incorporates basin scale, array scale and device scale modelling. Unfortunately since the product is commercial and proprietary very little low-level detail about the workings of the methodology are available, although a high-level exposition is given in Thomson et al. [2011].

The broad methodology of *TidalFarmer* is to start by determining the tidal flow field over the domain. Long term temporal variation is determined using a standard tidal harmonic methodology, while spatial variation of the flow over the site is provided by a shallow-water or 3-D flow solver (though the details of these solvers are unclear). The key assumption is then that "the flow field governed by variations in seabed bathymetry and roughness can be decoupled from the inter-array governed flow field and hence that a perturbation approach is sufficient for the purpose of a design tool" [Thomson et al., 2011]. This means that the array hydrodynamics are modelled by superposition of simple blockage models and analytically computed wake models (adapted from the wind industry) onto the ambient flow field. A genetic algorithm is then used to optimise the array micro-siting de-

sign.

Very recently, Stansby and Stallard [2016] presented a similar approach in which a simple gradient-based optimisation approach was employed to optimise the locations of a small array of turbines modelled via the superposition of analytical velocity deficit profiles. A (local) blockage correction factor (developed from laboratory tank testing) was used to adjust the analytical wakes – though the extent to which this factor is dependent upon the domain-geometry is unclear – and with realistic bathymetry and coastal geometry it is difficult to see how it will be employed. Nonetheless, the modelling approach gives good agreement with laboratory results for small arrays of turbines and demonstrates the benefits of array optimisation beyond the row-spacing approach. One major issue with the method is the absence of consideration of global blockage effects – i.e. bulk flow redirection due to the presence of turbines. As noted in Sutherland et al. [2007], in larger scale projects there may be an appreciable change to the underlying flow because of the turbine array. That notwithstanding, the model is very fast and lightweight – the paper states that it can be run on a laptop computer in a matter of minutes.

It is the focus of chapter 4 to compare the approach taken by *OpenTidalFarm* and the perturbation approach adopted in *TidalFarmer* (unfortunately the Stansby and Stallard [2016] approach was published after this work was completed so was not compared), with the aim of determining how the approaches inter-relate and may be used in combination. Given the closed-source nature of *TidalFarmer* this required an open-source equivalent to be constructed, the implementation of which is one of the key achievements of this work.

As with wind, some thought has been dedicated to array optimisation with respect to non-power-maximisation objectives. In some of the cases in which array layout has been considered, the model choice has precluded any in-depth investigation into the impact of the micro-siting design alterations. One example of this approach was Ahmadian and Falconer [2012] wherein a large array of 1000 turbines was modelled in the Severn Estuary, UK. Three alternate formations were considered – each with turbines distributed evenly over rectangular areas. Two of those areas covered 7.2 km^2 in rectangles with aspect ratios of 1:5 (each area was modelled centered in the same place but orientated perpendicularly to the other) the other was modelled with ten times lower turbine density. The purpose of the study was to establish whether array formation had an impact on environmental solute transport around the estuary and established that high turbine densities resulted in greater but more localised environmental impact. However, the work

did not make further conclusions about how environmental impact should be included in a more realistic design scenario with an array of a size that is likely to be seen in such regions. Conversely, du Feu et al. [2016] has demonstrated that environmental considerations can be included in the array design process, by utilising some of the multi-objective optimisation functionality that was developed in *OpenTidalFarm* as part of the work presented in chapter 5 of this thesis.

An alternate approach to array design is espoused by the DTOcean project – a large, pan-European effort centered in Edinburgh University which focuses less on the hydrodynamic constraints and more on engineering considerations – such as operational constraints, time to failure, etc. However one work package is focused on modelling and optimising the turbine array hydrodynamics with a view to balancing effects on environmental impact, power performance and cost uncertainties. For this work package, the conclusion of a review of methods [DTOcean, 2014] was to employ a perturbation approach to array modelling for optimisation – akin to the *TidalFarmer* approach. This was largely because, of all the methods surveyed, this was the only approach which was meant for use in an optimisation application – the other models considered were deemed "too computationally expensive as well as unnecessarily accurate" (as measured against the criteria defined by the objectives of that project) [DTOcean, 2014]. Though it should be noted that gradient-based optimisation methods were not considered in the review, this conclusion supports the discussion on the balancing of modelling and optimisation complexity that has been a theme in this chapter.

Chapter 3

OpenTidalFarm

Contents

3.1 The finite element method 52

3.2 OpenTidalFarm 57

3.3 Summary 68

Much of this chapter is based upon my contributions to Funke et al. [2016].

3.1 The finite element method

3.1.1 Numerical methods for partial differential equations

While equations describing physical processes may be well known and understood, they frequently cannot be solved analytically on realistic problems and domains. For the Navier-Stokes and shallow water equations, there exist few analytical solutions, and in scenarios where the non-linear terms dominate or the domain or parameters are complex, they must be solved algorithmically. Thus, rather than finding the *exact* continuous solution over the domain, we instead look for the *approximate* discrete solution. Various approaches exist to achieve this, including Finite Differences, (FD); Finite Volumes, (FV); and Finite Elements, (FE).

While finite differencing will briefly be employed for a single application in a later chapter, the majority of the work in this thesis is developing a framework founded upon computing the solution of differential equations using the finite element method.

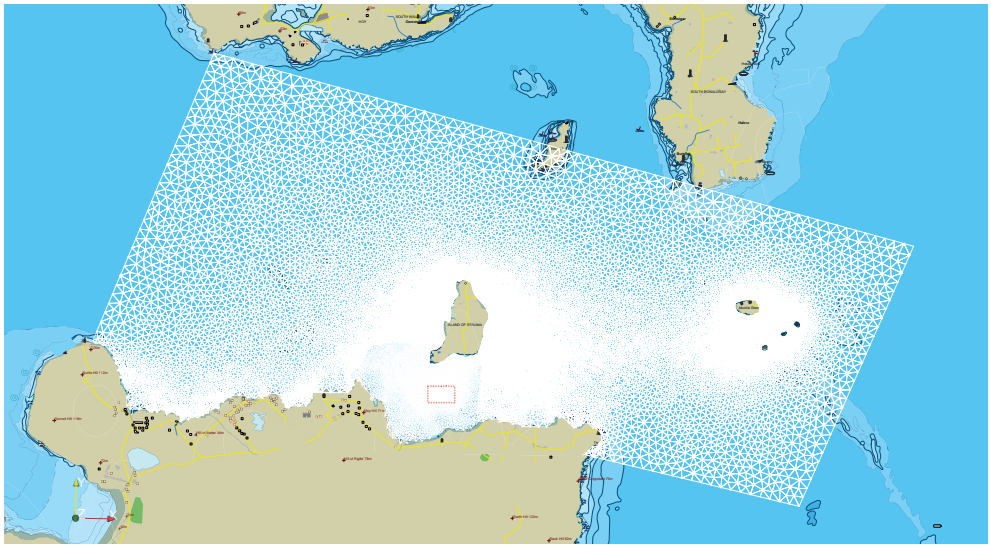


Figure 3.1: An example of an unstructured finite element mesh in the region of the Pentland Firth in the UK, with an emphasis on the Inner Sound. Note the varying element size in the unstructured mesh allows accuracy to be increased in areas of interest (for example the turbine area shown by the dotted box to the south of the island of Stroma).

At its most broad, the concept underpinning the finite element method is to break down the complex domain for which we can not directly solve the equations into an approximately equivalent tesselation of sufficiently simple and finitely numbered elements on which we can. This network is called the finite element mesh and represents a discretisation of the spatial domain. Each element is defined by its nodes – or degrees of freedom, the number of which depends on the chosen discretisation scheme. At the domain boundaries, conditions are specified, completing the description of the problem. Using these known boundary values and recognising that nodes between adjacent elements are shared, a system of simultaneous equations is constructed which can be solved using numerical linear algebra algorithms to approximate the unknown(s) over the domain. Broadly speaking the accuracy of the approximation is linked to the number of elements, which in turn is linked to the size of the linear system and therefore the computational expense of its construction/assembly as well as its solve.

The choice of finite element method over (arguably simpler to implement) finite difference schemes is that FE gives flexibility in the order of the method and makes few assumptions regarding the ordering, placement and connectivity of the nodes of the elements. This means that unstructured meshes may be used which are significantly advantageous in that the element size can vary (dramatically) over the domain to increase the modelling accuracy in areas of particular interest (for ex-

ample the location of a turbine array) or decrease the computational expense of the solve as appropriate, including in areas of less interest (figure 3.1). From an alternate point of view, this is also the down-side of the finite element method due to the corresponding increase in complexity required to support the flexibility in mesh structure. This involves creating and using maps and graphs which store the connectivity of the nodes and enable movement from global coordinate systems in the whole domain to local coordinate systems in the individual elements. Conversely, structured meshes have regular connectivity and mappings are implicit, but while varying resolutions can be achieved by nesting or 'telescoping' grids, interfaces of grids of different resolutions require special attention as cell connectivity is no longer implicit.

While for coastal modelling applications the finite volume method is frequently a popular choice, due to the simple upwinding schemes used to evaluate fluxes at the boundaries of cells, deriving high-order accuracy schemes is challenging and as such first-order accuracy is often the limit of these models [LeVeque, 2002, Cai et al., 2003].

In summary a great strength of the finite element method, and a major reason it is used in this application is the flexibility of the underlying mathematical formulation which facilitates varied choices of discretisation and finite element spaces providing broad applicability to a diverse range of problems. The other major reason is borne of practicality; the *FEniCS* [Logg et al., 2012] and *dolfin-adjoint* [Farrell et al., 2013] frameworks make implementation of FE very straight-forward, this will be examined in more detail in section 3.2.5. Of course, the solutions returned by FE should hold if another method were substituted for it.

3.1.2 The finite element method

For much of this work, the governing partial differential equations (PDEs) are the shallow water equations over the domain, Ω;

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} - \nu\nabla^2\mathbf{u} + g\nabla\eta + \frac{c_b}{H}\|\mathbf{u}\|\mathbf{u} = 0 \quad \text{in } \Omega\times(0,T], \tag{3.1a}$$

$$\frac{\partial \eta}{\partial t} + \nabla\cdot(H\mathbf{u}) = 0 \quad \text{in } \Omega\times(0,T], \tag{3.1b}$$

Here, the viscosity coefficient, acceleration due to gravity, and water density are

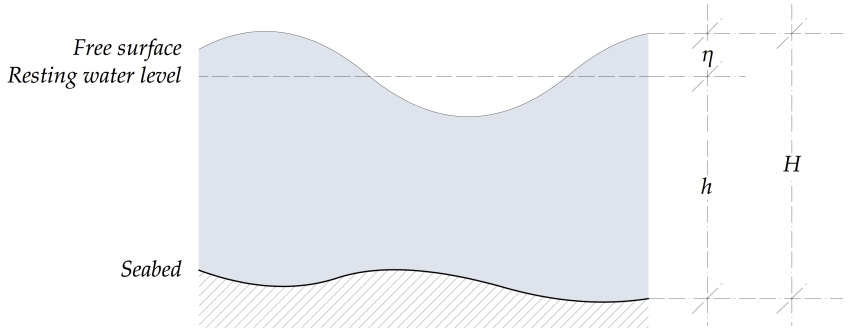


Figure 3.2: Section through the water column showing the relationship between the total water depth, $H : \Omega \times [0,T) \to \mathbb{R}$, the bottom depth, $h : \Omega \to \mathbb{R}$ and the free surface displacement $\eta : \Omega \times [0,T) \to \mathbb{R}$.

represented by ν, g, and ρ respectively. The total water depth, $H : \Omega \to \mathbb{R}$ is given by $H = h + \eta$ where $h : \Omega \to \mathbb{R}$ is the bottom depth and $\eta : \Omega \to \mathbb{R}$ is the free surface displacement (figure 3.2). The quadratic bottom friction parameter is represented by c_b and $\mathbf{u} : \Omega \times [0,T) \to \mathbb{R}^2$ is the depth-averaged flow velocity.

Our goal is to manipulate (3.1) – continuous equations for which a solution is sought over Ω – into discrete equations which can be approximated computationally at the nodes of the mesh in a finite-dimensional subspace. This is achieved by using the weak form of the problem, i.e. multiplying (3.1) with test functions Ψ and Φ, and integrating over the domain so as to find $\mathbf{u} : \Omega \times [0,T) \to \mathbb{R}^2$ and $\eta : \Omega \times [0,T) \to R$ such that for all $\Psi : \Omega \to \mathbb{R}^2$ and $\Phi : \Omega \to \mathbb{R}$,

$$\left\langle \frac{\partial \mathbf{u}}{\partial t}, \Psi \right\rangle_\Omega + \langle \mathbf{u} \cdot \nabla \mathbf{u}, \Psi \rangle_\Omega + \nu \langle \nabla \mathbf{u}, \nabla \Psi \rangle_\Omega - \nu \langle \nabla \mathbf{u} \cdot \mathbf{n}, \Psi \rangle_{\partial\Omega} + g \langle \nabla \eta, \Psi \rangle_\Omega + \left\langle \frac{c_b}{H} \|\mathbf{u}\| \mathbf{u}, \Psi \right\rangle_\Omega = 0, \tag{3.2a}$$

$$\left\langle \frac{\partial \eta}{\partial t}, \Phi \right\rangle_\Omega - \langle H\mathbf{u}, \nabla \Phi \rangle_\Omega + \langle H\mathbf{u} \cdot \mathbf{n}, \Phi \rangle_{\partial\Omega} = 0. \tag{3.2b}$$

where $\langle \cdot, \cdot \rangle_\Omega$ denotes the L^2 inner product over Ω. The viscosity term in (3.2a) and

divergence term in (3.2b) have been integrated by parts. Dirichlet boundary conditions on the velocity and free surface are strongly enforced by restricting the solution spaces to functions that satisfy them. No-normal flow boundary conditions are weakly enforced by neglecting the boundary integral in the weak continuity equation (3.2b).

To obtain the spatially discretised equations discrete function spaces are chosen constructed from a triangulation of Ω. For the entirety of this work Taylor-Hood ($\mathscr{P}_2\mathscr{P}_1$) element pairs [Taylor and Hood, 1973] are used; thus the velocity and free-surface displacement are represented by piece-wise quadratic and piece-wise linear functions respectively. Taylor-Hood pairs were chosen for practical reasons: While this choice of dicretisation is not point-wise divergence free and has no element-wise mass conservation, it is relatively compact and is robust and stable (more precisely it is a stable discretisation with viscosity – it is not stable for the Poisson equation). In contrast, an alternate popular choices for coastal modelling, for example $\mathscr{P}_{1\mathrm{DG}}\mathscr{P}_2$ (used in *Fluidity* [AMCG, 2010]) will result in the assembly of larger matrices and because of the requirement for surface integrals and often slope limiting for stability, the adjoint becomes more difficult to formulate (although a more practical reason for not using this element type was that it has only recently been supported by *FEniCS*).

For the temporal discretisation a θ-method is applied, that is; given the solution for the previous timestep $\mathbf{u}^n$ and η^n the solutions at the new timestep $\mathbf{u}^{n+1}$ and η^{n+1} are computed with:

$$\begin{aligned}\left\langle \frac{\mathbf{u}^{n+1}-\mathbf{u}^n}{\Delta t}, \Psi \right\rangle_\Omega + \left\langle \mathbf{u}^{n+\theta}\cdot\nabla\mathbf{u}^{n+\theta}, \Psi \right\rangle_\Omega + \nu \left\langle \nabla\mathbf{u}^{n+\theta}, \nabla\Psi \right\rangle_\Omega \\ -\nu \left\langle \nabla\mathbf{u}^{n+\theta}\cdot\mathbf{n}, \Psi \right\rangle_{\partial\Omega} + g\left\langle \nabla\eta^{n+\theta}, \Psi \right\rangle_\Omega + \left\langle \frac{c_b}{H}\|\mathbf{u}^{n+\theta}\|\mathbf{u}^{n+\theta}, \Psi \right\rangle_\Omega = 0, \qquad (3.3a)\end{aligned}$$

$$\left\langle \frac{\eta^{n+1}-\eta^n}{\Delta t}, \Phi \right\rangle_\Omega - \left\langle H\mathbf{u}^{n+\theta}, \nabla\Phi \right\rangle_\Omega + \left\langle H\mathbf{u}^{n+\theta}\cdot\mathbf{n}, \Phi \right\rangle_{\partial\Omega} = 0. \qquad (3.3b)$$

with $0 \le \theta \le 1$ and $\mathbf{u}^{n+\theta} = \theta\mathbf{u}^{n+1} + (1-\theta)\mathbf{u}^n$ and $\eta^{n+\theta} = \theta\eta^{n+1} + (1-\theta)\eta^n$. Setting $\theta = 1$ for fully implicit first order Euler, and $\theta = 1/2$ for second-order Crank-Nicolson. The former – implicit Euler – has the benefit of unconditional stability and simplicity.

Since all operators in (3.3) are continuously differentiable, once assembled, the

(coupled) non-linear systems are solved using Newton's method, here to a residual tolerance of 10^{-9}.

Having discretised the partial differential equation on paper, much of the development 'effort' involved in solving them is automated by the way in which the problem is implemented in *OpenTidalFarm*. This final step – obtaining the solution – will thus be discussed in terms of *OpenTidalFarm*'s code architecture in the following section.

3.2 OpenTidalFarm

3.2.1 Overview of OpenTidalFarm

Before delving into the details of model structure and code implementation, it is useful to have an overview of the *OpenTidalFarm* framework. *OpenTidalFarm* is an open-source software tool for scientists and industrialists to simulate and optimise the performance of tidal turbine farms. The performance metric is user-specific and can be the power production, the farm cost, the environmental impact of the farm or some combination of these. The farm's power production is estimated from the (time-varying) flow velocities at the turbine locations. These velocities are computed by numerically solving the shallow water equations, taking into account the nonlinear interactions between the turbines, coastline and bathymetry. These equations are an appropriate compromise between physical fidelity and affordable computational cost [Jacobs et al., 2016]. The farm itself is parameterised either by individually 'resolved' turbines (for micro-siting design optimisation) [Funke et al., 2014], or as a continuous turbine density function (for large-scale resource assessment and multi-farm optimisation) [Funke et al., 2015a].

OpenTidalFarm can currently be used in three modes. In the *simulation mode*, it solves the shallow water equations and predicts the farm's performance. The performance is judged by a metric such as the power production. In the *sensitivity mode*, *OpenTidalFarm* additionally solves the adjoint shallow water equations, enabling the computation of *sensitivities* (derivatives) of the farm performance with respect to all model inputs. This information is of enormous utility to the designer, as it shows how model inputs may be varied to improve the performance metric. The possible inputs include *design* variables the turbine positions, the turbine sizes, the pitch of the turbines or spatial *model* variables (which there may be some

uncertainty over) such as the viscosity, bottom friction or basin depth (examined in chapter 7). In the *optimisation mode*, which is the most advanced mode, *OpenTidalFarm* is used to design the array via the solution of an optimisation problem. In this mode, *OpenTidalFarm* applies derivative-based optimisation techniques to maximise or minimise the performance metric by systematically varying the user-specified control parameters. The necessary derivatives are computed using the adjoint model. The derivative-based optimisation methods employed are very efficient and enable optimisation of large farms with hundreds to thousands of turbines.

3.2.2 Modelling turbine arrays

As has been noted, modelling tidal stream turbine arrays spans a large range of length scales. *OpenTidalFarm* operates on the scale of a few turbines (domains of a few square kilometres) all the way up to multiple arrays (domains of whole coastal regions). To achieve this, the tidal turbines may be represented using one of two abstractions: Individual turbines may be modelled as discrete bumps of drag friction, or whole arrays may be represented by a continuous drag function bounded by the extent of the farm area that has been leased by the developer (henceforth the 'turbine area').

Discrete turbine representation

Where it is computationally tractable to refine the finite element mesh within the turbine area to sufficiently resolve individual turbines (a minimum mesh element size of a quarter of the turbine diameter is a good guide), then the turbines in the array may be represented individually.

The bottom friction term in (3.1), c_b/H, is augmented by a new turbine friction, c_t to give $(c_b + c_t)/H$. This turbine friction term depends on the vector, $\mathbf{m}$ which contains the parameters for the turbines in the farm. If the locations of the turbines are being optimised then $\mathbf{m} \in \mathbb{R}^{2n}$ and $c_t(\mathbf{m}) : \mathbb{R}^{2n} \to \Omega$ converts the two-dimensional Cartesian location coordinates of the turbines within the array to a friction field over the domain.

Recall that computing the force on the turbine and the power extracted by it required use of the freestream velocity, u_{fs}. Unfortunately, in a practical setting, determining this value is difficult. The usual procedure is to take the flow velocity far

upstream of the turbine – where the flow is undisturbed by presence of the turbine – however in a non-idealised setting this upstream value is likely to be affected by other turbines or the irregular geometry of the domain / bathymetry. Consequently, for the remainder of this thesis the velocity used to compute the thrust and power will be the velocity at the turbine. The extent to which this is a simplification – and potential procedures to correct it are explored in Kramer and Piggott [2015]. Hence the thrust applied to the flow by a three-dimensional turbine is

$$F_{3\mathrm{D}}(\mathbf{u}) = \frac{1}{2}\rho A_T C_t \|\mathbf{u}\|\mathbf{u}, \tag{3.4}$$

Where A_t is the swept area of the turbine blades C_t is the function describing how the turbine thrust coefficient varies with flow speed and ρ is the fluid density. In a depth-averaged sense, the thrust is defined to be

$$F_{2\mathrm{D}}(\mathbf{u}) = \rho \int_\Omega c_t \|\mathbf{u}\|\mathbf{u}\,\mathrm{d}x. \tag{3.5}$$

where c_t is the augmented bottom drag representing the turbine. Since (3.5) is depth-averaged and we wish to apply the same thrust in two dimensions as we have calculated in three, we can equate (3.4) and (3.5) to give

$$\frac{1}{2}\rho A_T C_t \|\mathbf{u}\|\mathbf{u} = \rho \int_\Omega c_t \|\mathbf{u}\|\mathbf{u}\,\mathrm{d}x, \tag{3.6}$$

or (assuming $\mathbf{u}$ is constant over the region c_t is non-zero)

$$\frac{1}{2}A_T C_t = \int_\Omega c_t\,\mathrm{d}x. \tag{3.7}$$

We need to design a (differentiable) function such that (3.7) is satisfied, thus we define

$$\Xi(\mathbf{p},\mathbf{x}) = \xi(\mathbf{p}_x, x)\xi(\mathbf{p}_y, y), \tag{3.8}$$

where $\mathbf{p}_x$ denotes the x-coordinate of the centre of the turbine,

$$\xi(p,x) = \begin{cases} \exp\left(1 - \dfrac{1}{1 - \left\| \dfrac{2(x-p)}{d} \right\|^2}\right) & \text{for } \left\| \dfrac{2(x-p)}{d} \right\| < 1 \\ 0 & \text{otherwise.} \end{cases} \tag{3.9}$$

and d is the plan-view diameter of the two-dimensional representation of the turbine. Rearranging (3.7) gives

$$c_t = \Lambda C_t(\mathbf{u}) \Xi(\mathbf{p}, \mathbf{x}), \tag{3.10}$$

where

$$\Lambda = \frac{A_t}{2 \int \Xi \, \mathrm{d}x}, \tag{3.11}$$

is a constant that can be fixed prior to the optimisation once the turbine dimensions are known.

For a multiple turbine array, a single turbine friction function is constructed from a summation of the individual turbines,

$$c_t = \sum_{i=1}^{n} c_t^i \tag{3.12}$$

The thrust coefficient, $C_t(\mathbf{u})$, is device dependent and is considered by device developers to be commercially sensitive information. *OpenTidalFarm* supports user defined thrust curves (via look-up tables), or defaults to

$$C_t(\mathbf{u}) = \begin{cases} c_t^{\text{rated}} \left(\dfrac{\tanh(10(\mathbf{u} - \mathbf{u}_{\text{in}})) + 1}{2} \right), & \text{if } \mathbf{u} \leq \mathbf{u}_{\text{in}} \\ c_t^{\text{rated}}, & \text{if } \mathbf{u}_{\text{in}} \leq \mathbf{u} \leq \mathbf{u}_{\text{rated}} \\ c_t^{\text{rated}} \dfrac{\mathbf{u}_{\text{rated}}^3}{\mathbf{u}^3}, & \text{if } \mathbf{u} \geq \mathbf{u}_{\text{rated}}, \end{cases} \tag{3.13}$$

which is adapted from Martin-Short et al. [2015] and produces the thrust and power characteristics as plotted in figures 3.3 and 3.4. Here, c_t^{rated} is the turbine's thrust coefficient at it's rated velocity. The corresponding power coefficient, C_p can be found via

$$C_p(\mathbf{u}) = \frac{1}{2}\left(1+\sqrt{1-C_t(\mathbf{u})}\right)C_t(\mathbf{u}). \tag{3.14}$$

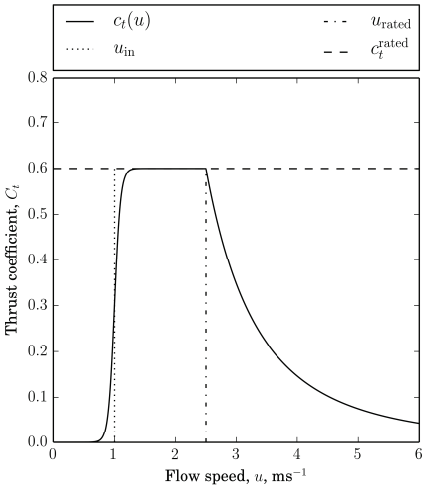


Figure 3.3: Thrust curve for tidal turbine with a cut-in speed of $u_{\text{in}} = 1.0$ m/s, and a rated speed of $u_{\text{rated}} = 2.5$ m/s.

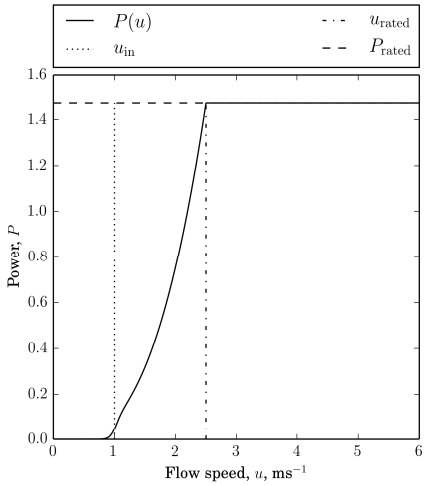


Figure 3.4: Power curve corresponding to thrust curve in figure 3.3.

Continuous turbine representation

The discrete turbine representation is not practical for tidal resource assessment, or large-scale simulations and optimisations that involve potentially multiple farms because the turbines must be explicitly resolved in the mesh.

To overcome these limitations, Funke et al. [2015a] proposed a continuous turbine approach in which turbines are represented as a continuous density function rather than individually resolved turbines. Of course, such a simplified turbine representation can not resolve the wakes of individual turbines, and hence the inter-array hydrodynamic interactions will not be captured. Nevertheless, experiments on simple domains have shown that the optimal power output, optimal number of turbines, and the optimal overall design of the farm is similar for both the discrete and the continuous approach. This array parameterisation is not used in the remainder of this thesis and is just mentioned for completeness.

3.2.3 Turbine array optimisation

OpenTidalFarm solves an optimisation problem constrained by the shallow water equations:

$$\begin{aligned} \max_{\mathbf{m}} \quad & P(\mathbf{m}) \\ \text{subject to} \quad & \mathbf{b}_l \leq \mathbf{m} \leq \mathbf{b}_u \\ & g(\mathbf{m}) \leq 0. \end{aligned} \tag{3.15}$$

The bounds $\mathbf{b}_l \leq \mathbf{m} \leq \mathbf{b}_u$ constrain the turbines to the turbine area (in this case rectangular in shape, for simplicity) and the inequality constraint, $g(\mathbf{m}) \leq 0$, implements a minimum distance spacing constraint between adjacent turbines. Turbines are modelled as distinct areas of increased friction, and at each optimisation iteration the performance of the turbine layout is evaluated as the power extracted by the turbines,

$$P(\mathbf{m}) = \int_\Omega \rho c_t(\mathbf{m}) ||\mathbf{u}||^3 \, \mathrm{d}x, \tag{3.16}$$

where ρ is the fluid density, $c_t(\mathbf{m})$ is the enhanced friction of the parametrised turbines, and $\mathbf{u}$ is the velocity of the flow which, along with the free-surface displacement, η, is the solution to the shallow water equations

The adjoint approach is used to efficiently evaluate the sensitivity of the power with respect to $\mathbf{m}$ and gradient-based optimisation approaches are employed to improve the power through manipulation of $\mathbf{m}$.

3.2.4 The adjoint approach

A key requirement of gradient-based optimisation is access to the gradient of the functional (e.g. (3.16)) with respect to input parameters. Given the crux of the OpenTidalFarm methodology is the ability to determine this gradient with computational economy, the technique employed to achieve this is worthy of consideration in some detail. In OpenTidalFarm, the 'adjoint approach' is used to determine the sensitivity of the output functional evaluation with respect to the input parameters.

There are three common ways in which the gradient of a function may be found; through finite differencing, the tangent linear mode and the adjoint mode:

The finite difference approach facilitates an approximation of the gradient by perturbing the input variables by a small amount and recording the effect of that perturbation on the output parameter. In the case of a problem with multiple inputs (i.e. an x & y coordinate for N turbines), each input must be perturbed individually. Consequently approximation of the gradient requires $2N+1$ functional evaluations – the original evaluation plus a re-evaluation for a perturbation applied to each of the $2N$ input variables.

With the tangent linear, and adjoint methods, the gradient can be determined directly and efficiently. These methods have a wide variety of applications across fields such as optimisation, sensitivity analysis and data assimilation. Consider a model made up of i component functions and j input parameters,

$$F_1(x_1,\ldots,x_j),\ldots,F_i(x_1,\ldots,x_j)$$

with the resulting $i \times j$ Jacobian matrix

$$J = \begin{bmatrix} \frac{\partial F_1}{\partial x_1} & \cdots & \frac{\partial F_1}{\partial x_j} \\ \vdots & \ddots & \vdots \\ \frac{\partial F_i}{\partial x_1} & \cdots & \frac{\partial F_i}{\partial x_j} \end{bmatrix}. \tag{3.17}$$

In tangent linear mode, the sensitivity of many output variables is found with respect to a single input variable – effect is found from cause. As such this matrix can be found using j forward sweeps through a linearised version of the model, each finding a column of the Jacobian. Conversely, the adjoint approach may be used to determine the sensitivity of a *single* output variable with respect to *many* input variables – cause is determined from effect, which would result in i backward sweeps through the linearised version of the model, each finding a row. The mechanics of this technique are best viewed through its application in OpenTidalFarm.

The adjoint equation as applied in OpenTidalFarm

Recall equation (3.15) in which the functional of interest, $P(\mathbf{m})$ is presented. Generalising this to a generic functional $J \equiv J(\mathbf{u}(\mathbf{m}),\mathbf{m})$, differentiation of the functional gives

$$\frac{\mathrm{d}J}{\mathrm{d}\mathbf{m}} = \frac{\partial J}{\partial \mathbf{u}} \frac{\mathrm{d}\mathbf{u}}{\mathrm{d}\mathbf{m}} + \frac{\partial J}{\partial \mathbf{m}}. \tag{3.18}$$

That is, the total sensitivity of the functional of interest is a combination of $\mathbf{m}$ affecting J directly, and $\mathbf{m}$ affecting J through its effect on the flow, $\mathbf{u}$. The effect of $\mathbf{m}$ on $\mathbf{u}$ is determined through the constraint (here the shallow-water equations) $F \equiv F(\mathbf{u}(\mathbf{m}), \mathbf{m}) = 0$, differentiation of which yields

$$\frac{\partial F}{\partial \mathbf{u}} \frac{\mathrm{d}\mathbf{u}}{\mathrm{d}\mathbf{m}} + \frac{\partial F}{\partial \mathbf{m}} = 0. \tag{3.19}$$

Equations (3.18) & (3.19) can be combined through substitution such that

$$\frac{\mathrm{d}J}{\mathrm{d}\mathbf{m}} = -\frac{\partial J}{\partial \mathbf{u}} \frac{\partial F}{\partial \mathbf{u}}^{-1} \frac{\partial F}{\partial \mathbf{m}} + \frac{\partial J}{\partial \mathbf{m}} \tag{3.20}$$

$$= -J_{\mathbf{u}} F_{\mathbf{u}}^{-1} F_{\mathbf{m}} + J_{\mathbf{m}}, \tag{3.21}$$

where subscripts have been used for clarity to denote partial derivatives. We define the adjoint variables via

$$\lambda^T = J_{\mathbf{u}} F_{\mathbf{u}}^{-1} \qquad or \qquad \lambda = F_{\mathbf{u}}^{-T} J_{\mathbf{u}}^T, \tag{3.22}$$

so now (3.20) can be expressed as

$$\frac{dJ}{d\mathbf{m}} = -\lambda^T F_{\mathbf{m}} + J_{\mathbf{m}}. \tag{3.23}$$

[Funke et al., 2016]

Since

$$-\lambda^T F_{\mathbf{m}} = \lambda^T F_{\mathbf{u}} \frac{d\mathbf{u}}{d\mathbf{m}} = (F_{\mathbf{u}}^T \lambda)^T \frac{d\mathbf{u}}{d\mathbf{m}} = J_{\mathbf{u}}^T \frac{d\mathbf{u}}{d\mathbf{m}}, \tag{3.24}$$

one can either evaluate $\lambda^T F_{\mathbf{m}}$ via the tangent linear approach: solving the linear system of (3.19), $F_{\mathbf{u}} \frac{d\mathbf{u}}{d\mathbf{m}} = -F_{\mathbf{m}}$; or through the adjoint approach: solving the linear system of (3.22), $F_{\mathbf{u}}^T \lambda = J_{\mathbf{u}}^T$. Since J is a scalar, $J_{\mathbf{u}}$ is a vector, while $F_{\mathbf{m}}$ is a matrix.

The major benefit of this methodology is that, in practice, the computational cost of determining the gradient via the adjoint approach is only about twice that of the functional evaluation (including the calculation of $\mathbf{u} \equiv \mathbf{u}(\mathbf{m})$). However, derivation of the adjoint system, $F_{\mathbf{u}}^T \lambda = J_{\mathbf{u}}^T$ is generally not straightforward. In *OpenTidalFarm*, this derivation is automated through the *FEniCS* framework using the *dolfin-adjoint* and *libadjoint* libraries [Farrell et al., 2013]. An additional noteworthy point is that, because the adjoint approach is determining cause from effect, for time-dependant problems, the adjoint appears as an integration backwards in time [Errico, 1997].

The effect of using this approach is that, when designing new modules for *OpenTidalFarm* framework the goal should be to devise a way in which a potentially multi-input design problem can be evaluated through inspection of a single scalar value. This enables techniques such as the adjoint approach to be exploited, and so the problem is formulated with inherent computational economy. Again, this has been demonstrated by Funke et al. [2014].

3.2.5 Code architecture

OpenTidalFarm is developed within a paradigm of high-level, open-source and object-orientated code design. Object oriented design has the benefit that individual component classes may be substituted or upgraded allowing adaptation of the base code with minimal intrusion. High-level code design facilitates a small dynamical core, easy extension to new equations and automated derivation of the adjoint model. Finally, being open-source facilitates greater adoption and promotes external development thanks to easy availability and customisation.

Once the discretised equations (3.2) have been derived, the *FEniCS* software is used as the framework within which the solution is computed. *FEniCS* is a collection of packages for the solution of differential equations using compact and flexible high-level code. The key idea is that the user constructs the program in *Python*, specifying the problem in *Unified Form Language* (UFL) which is very similar to its mathematical representation. At run-time, specialised low-level C++ code is generated and compiled, optimised for the problem being solved and potentially system architecture. The linear systems arising at each Newton step are solved using the *PETSc* library [Balay et al., 2015]. This approach has numerous benefits, not least of which is that the user is afforded a readable, high-level language in which to construct the computational model, which then runs with the efficiency

of a low-level compiled language. The second key benefit is the UFL objects representing the mathematical operations are retained at a high-level until the code is executed.

A key challenge of applying derivative-based optimisation is to implement the adjoint model to efficiently compute the derivatives of the farm performance with respect to the control variables. Traditional approaches to developing adjoint models require either the manual derivation and implementation of the adjoint equations or the application of low-level algorithmic differentiation tools. These approaches are technically challenging, with implementation requisite upon invasive low-level code alterations. The *dolfin-adjoint* library exploits the high-level representation of the equations in FEniCS to automatically derive the adjoint equations. It then instructs FEniCS to generate the code requisite for computing the adjoint solutions [Farrell et al., 2013]. In addition *dolfin-adjoint* provides the interface to various optimisation algorithms [Funke and Farrell, 2013] which are used by *OpenTidalFarm* to optimise the control parameters with respect to the designated objective.

The FEniCS and *dolfin-adjoint* libraries facilitate a compact and high-level implementation of *OpenTidalFarm* while achieving scalability and computational efficiency.

3.2.6 Problem work-flow

At its heart, *OpenTidalFarm* comprises a standardised work-flow from problem definition to solution. Each stage is implemented as a prototypical super-class in which the core and common functionality for that stage is defined. For example, each specific solver implementation (steady-state shallow-water, transient shallow-water, etc.) overloads the 'solver' super-class. The benefit of this approach is that more advanced solvers (e.g. development of a full three-dimensional solver is proposed) may be implemented with no intrusion into the remnant of the *OpenTidalFarm* code base. Figure 3.5 presents an overview of the *OpenTidalFarm* work-flow, showing the most important classes and their dependency upon each other.

The work-flow begins by instantiating a shallow water problem, *SWProblem*, with a *ProblemParameter* object. This parameter object stores all pertinent information about the shallow water problem, such as the domain and the boundary conditions to be applied. This problem is then used to create a solver for which solver-specific parameters may also be specified. For a sensitivity analysis or opti-

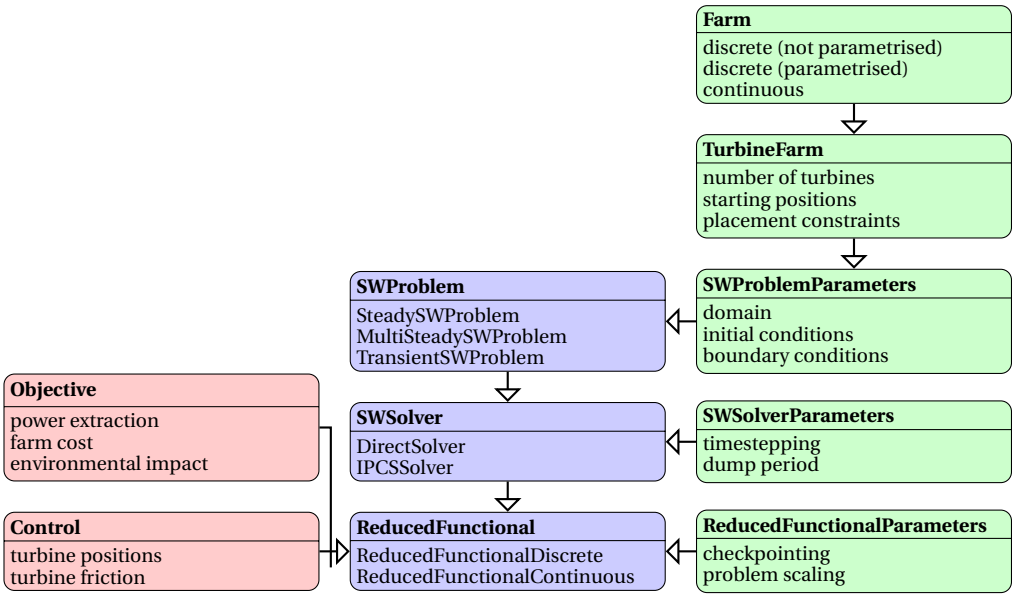


Figure 3.5: Overview of the main classes and their dependencies in OpenTidalFarm. The core work-flow classes are shown in blue with the prototypical class in bold and the overloaded variations beneath. Parameter classes are shown in green with example options listed beneath. Ancillary classes for optimisation and sensitivity are shown in red with example control parameters and objective functionals shown as appropriate.

misation, the user defines the objective functional via the *Functional* class and the control parameters through the *Control* class. These are passed to the *ReducedFunctional* which provides the interface with *dolfin-adjoint*.

3.2.7 OpenTidalFarm optimisation algorithm

At each iteration, the following steps are undertaken: The shallow water equations are solved using the finite element method on a triangular mesh and the Taylor-Hood element pair [Taylor and Hood, 1973]. The derivative of the power production with respect to the turbine locations, $\mathrm{d}P/\mathrm{d}\mathbf{m}$, is computed by solving the adjoint shallow water equations [Funke et al., 2014]. Based upon both P and $\mathrm{d}P/\mathrm{d}\mathbf{m}$, the locations of the turbines, $\mathbf{m}$ are improved for the next iteration using a gradient-based optimisation algorithm. The optimisation algorithm used is a variation on sequential quadratic programming, SLSQP, which is packaged as part of SciPy [Jones et al., 2001–] and outlined in Kraft [1988]. For the purposes of this work, SLSQP, is treated as a black box producing an improved $\mathbf{m}$ based upon previous functional and gradient evaluations. The *OpenTidalFarm* optimisation loop

is outlined in figure 3.6.

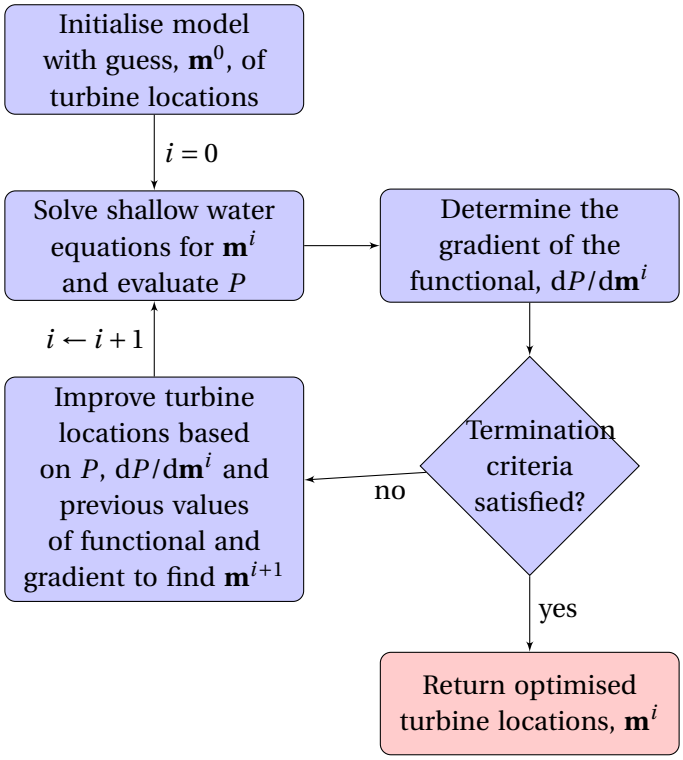


Figure 3.6: Schematic of the *OpenTidalFarm* optimisation procedure.

3.3 Summary

OpenTidalFarm is a tidal current turbine array micro-siting optimisation tool that employs PDE-constrained optimisation using the adjoint approach [Funke et al., 2016]. The key idea is to use a gradient-based method to optimise the turbine layout. This is implemented to maximise the time-averaged power extracted due to increased friction in the tidal farm [Funke et al., 2014], this is the *functional of interest* (a functional is a function which acts upon a vector space – in this case the velocity field – but returns a scalar value). The efficiency with which the gradient of the solution can be found (via the adjoint approach) and then used to optimise the farm design means that a relatively complex – hence more realistic – flow model may be used. Unlike, for example, genetic algorithms, much fewer iterations are required and consequently the computational expense of calculating the functional evaluation at each iteration no longer needs to be so stringently economised. The functional of interest represents a measurement of a design fac-

tor against which the quality of a farm design may be judged. For the purposes of this section, this functional is representative of the amount of energy extracted from the flow, in later chapters this will be broadened to capture more of the the diverse design objectives hitherto discussed.

Chapter 4

Comparing appoaches to tidal turbine array optimisation

Contents

4.1 Introduction 71
4.2 Problem Definition 74
4.3 Modelling approaches 76
4.4 Single turbine comparison 87
4.5 Array comparison 91
4.6 Optimisation 95
4.7 Optimisation comparison 100
4.8 Conclusions 108

This chapter is based upon Culley et al. [2016c].

Abstract

This chapter considers and compares some of the approaches to tidal turbine array micro-siting optimisation; the process of determining where individual turbines should be placed within the site where the array is to be installed. The complexity and coupled nature of the advection-dominated and therefore turbulent hydrodynamic interactions which govern the flow means that a human designer's intuition is unlikely to find the best micro-siting layout design; instead it must be discovered through computational optimisation.

The Jensen, Larsen and Ainslie wake deficit models are presented, calibrated and compared with reference to a high resolution CFD model of both an individual turbine and an array. The Ainslie model performed best for an individual turbine, the Jensen model modelled an array of 32 turbines most accurately while the Larsen model had the best performance aggregated over the two tests. The performance of the shallow-water model is also benchmarked using this approach and while it is found to be less successful at capturing the wake deficit, it does capture the local and global blockage effects that are neglected in the perturbation-based approaches.

The models are then used in two test cases in which the longitudinal and transverse spacings of staggered rows of turbines are adjusted to investigate the effect on array performance. One of these cases is then used to demonstrate how allowing turbines to move individually during the optimisation (rather than constricting them to rows whose spacing is adjusted) facilitates exploitation of local flow features and can produce much-improved performing layouts.

Having established the benefits of optimisation approaches in which all turbines are free to move individually rather than being confined to rows global, local and hybrid optimisation approaches are compared on three idealised test scenarios designed to highlight the strengths and weaknesses of the approaches. Ultimately, it was concluded that for most realistic situations, gradient-based optimisation of the coupled shallow-water model produced the best performing designs. However, such an optimisation approach is liable to find local minima and is sensitive to the starting guess, so for highly non-convex domains, an approach which employs a genetic optimisation with an analytical wake deficit model to provide a starting guess for the gradient-based shallow-water optimisation may find the optimal layout.

4.1 Introduction

Tidal stream power production is a process in which arrays of tens to hundreds of turbines extract momentum from tidally induced currents to deliver electricity. From conception to construction, the processes by which the arrays are designed and specified are extensively modelled and optimised using computational techniques. This chapter considers and compares some of the approaches to array optimisation; the process of determining where individual turbines should be placed within the site that has been leased for the project. This ‘micro-siting’ design can

have a significant impact upon the local and regional hydrodynamics, affecting the power extracted by the array and therefore its profitability for developers.

To date there are several tidal stream energy extraction project proposals being developed globally. These involve installing arrays of turbines in sites where tidal characteristics and coastal geography result in hydrodynamically energetic flows. The power that the tidal turbines extract is proportional to the flow velocity cubed. Thus not only is the power extracted dependent upon the turbine's location with respect to bathymetric and coastal features that may affect the flow, but also with respect to the other turbines in the array – including its own presence and thrust characteristics.

'Tuning' a device for its intended use is a familiar engineering concept, the array micro-siting problem deals with tuning a whole array of devices to optimise their interactions and performance on a given site. The complexity and coupled nature of the advection-dominated and therefore turbulent hydrodynamic interactions which govern the flow means that a human designer's intuition is unlikely to find the best micro-siting layout design; instead it must be discovered through computational optimisation. One of the key difficulties in tackling this problem computationally is the orders of magnitude differences in the spatial scales involved: At the turbine scale, turbulence needs to be modelled at blade-width resolution. At the array scale, individual turbines need to be resolved to study the interactions between them. At the coastal scale, a sufficiently large domain is required so as to adequately capture the large-scale hydrodynamics which feed into the site and in order for imposed boundary constraints to remain valid with changes to the hydrodynamics induced by the array. Invariably, numerical models which perform well at one of these scales are generally deficient at the others, even if it were feasible to discretise a coastal-scale model in such a way to be able to resolve individual blades of turbines.

A second key difficulty is the highly coupled relationship between a turbine and the dynamics of the flow into which it is placed – and the high sensitivity that the turbine performance has to the velocity of that flow. As a result, finding the optimum design of an array for a given site requires a process of iterative optimisation. Such a process requires the model which analyses array performance (the analysis code) to be re-run at each iteration. Clearly with computationally expensive codes and inefficient optimisation approaches this can become infeasible.

The first of these key difficulties is currently tackled through employing a hierarchy of models – from low fidelity with a large spatial scale to high fidelity with a

small spatial scale – which feed into each other in as coupled a manner as possible [Culley et al., 2014, Adcock et al., 2015]. In the future, it may be that a single multi-scale model will become computationally tractable. However currently, modelling arrays of actuator disks in a regional model, as per Abolghasemi et al. [2016], represents the state of the art – and optimisation using such an analysis code is not currently viable. In tackling the second difficulty – the computational challenge of finding the optimal design for this coupled problem – as has been alluded to previously there are, in practice, a few approaches:

The first approach is for the designer to use their intuition to propose a design which is then tested using a high-fidelity code. This has the benefit that the physics of the problem are modelled to the highest possible degree of realism – with the obvious benefits of a reliable solution. The key assumption is that the designer's intuition is good. However, with complex domain geography and large numbers of turbines it is unrealistic to expect a human to be able to visualise the processes and return the optimal micro-siting design. Furthermore the computational expense of the high-fidelity code means that repeatedly exploring alternate configurations is not possible.

The second approach is to use global optimisation approaches (such as genetic algorithms) to explore a large number of alternate designs and converge towards one which performs the best of those tested. This has the benefit that a large number of potential designs are explored and evaluated, however there are several key limitations. Firstly, these techniques scale poorly with problem size – for example as the number of turbines increases the search space expands and with it the required number of design evaluations (this will be explored in detail in following sections). Since a (very) large number of different options must be evaluated, the computational expense of the endeavour becomes prohibitive for all but the cheapest analysis codes. Computationally cheap, low fidelity analysis codes invariably fail to accurately capture many of the physical processes that may be key to understanding the array's performance and hence provide a less reliable performance evaluation for a given array.

A third approach is to use local, gradient-based optimisation approaches to improve the array design in a minimal (around 100) number of design iterations – and therefore analysis code runs. This facilitates use of a higher fidelity and hence more realistic analysis code which should evaluate array performance, and thus improve array design, more reliably. Such approaches scale well with the size of the control parameter and provide a guarantee of (local) optimality. The key limi-

tations are that such approaches require initialisation with a starting 'guess' of the control parameters and in non-convex solution domains the optimised solution is sensitive to that guess – i.e. the solution that is returned by the algorithm will be locally rather than globally optimal. The difference between global and local optimality may or may not be large in terms of the performance of the array.

The contribution of this chapter is in illustrating that the design (i.e. optimisation) methodology for tidal stream turbine arrays is intrinsically linked to the choice of underlying modelling approach. Consequently, choice of one drives the choice of the other. Furthermore, the goal is to demonstrate in which scenarios the modelling and optimisation approaches are strong (or not) and to help the array designer determine the most valuable or appropriate tools available to them.

In this chapter a family of optimisation approaches is presented and a robust paradigm for resource assessment and array design is developed which exploits the strengths of each of the above methodologies. In the next section, the turbine array design problem is formulated such that it can be tackled within a standardised mathematical optimisation framework. In section 4.3 global and local optimisation techniques are considered in more detail, and suitable analysis codes for each presented and discussed. The performance of the analysis codes will then be compared in sections 4 and 5 for an individual turbine and an array respectively. In sections 6 and 7, the performance of the optimisation approaches will be compared for some idealised test scenarios designed to show up the strengths and weaknesses of the various methodologies. Finally, some conclusions may be drawn on which may be the best performing methodology for a given array design scenario.

4.2 Problem Definition

The performance of an array is strongly affected by the hydrodynamic interactions between turbines and the flow. Through these interactions turbines are affected by each other and by the bathymetry and other characteristics of the site and environment. Tuning the array micro-siting design can, to some extent, promote these interactions where desirable and mitigate them where unwelcome. Various control parameters may be adjusted to achieve this tuning, however this chapter will focus on the locations – within the bounds of the site leased by the developer – of the individual turbines comprising the array. This is referred to as the array's micro-siting design.

We have seen in the previous chapter how the micro-siting design is encoded through the vector $\mathbf{m} = (x_1, y_1, x_2, y_2, \ldots, x_N, y_N)$ where N is the number of turbines comprising the array and x_N and y_N represent the two dimensional Cartesian coordinates of the N^{th} turbine. As such, the optimisation problem may be expressed as

$$\begin{aligned} \max_{\mathbf{m}} \quad & P(\mathbf{m}) \\ \text{subject to} \quad & b_l \leq \mathbf{m} \leq b_u \end{aligned} \tag{4.1}$$

Where the constraint, ensures that turbines remain within the area leased for development of the array. In this case, the area is rectangular and b_l and b_u represent the lower and upper boundaries respectively, though in general other turbine area shapes may be applicable. An additional inequality constraint which may be expressed $g(\mathbf{m}) \leq 0$ could be used to enforce a minimum distance between the turbines if such a constraint is possible within the optimisation framework of choice.

For this work, the quantity of interest (QoI) is – for steady problems – the instantaneous power extracted from the flow by the array, defined by

$$P = \frac{1}{2}\rho A||u||^3 C_P, \tag{4.2}$$

where ρ is the water density, A is the swept area of the turbine blades and C_P is the non-dimensional coefficient of power of the turbine. For more realistic transient problems, the quantity of interest is the energy extracted by the array over the model's time window, T

$$E = \int_0^T P(t)\ \mathrm{d}t. \tag{4.3}$$

Various approaches exist for the modelling of tidal turbines arrays, the fidelity of which are constrained by the size of the domain (at least in the order of square kilometers to cover the leased array site) and the designer's computational budget. For optimisation problems – which require multiple iterations (and therefore multiple model solves) – the model fidelity is further constrained by the efficiency of the optimisation approach chosen.

The work presented in this chapter seeks to demonstrate the importance of full array micro-siting design optimisation beyond simple adjustment of the longitu-

dinal and transverse spacing – a commonly used methodology for array tuning, for example by Malki et al. [2014] or Divett [2014]. Having established the importance of having many degrees of freedom for array optimisation, the goal is then to examine the benefits and costs of three optimisation approaches (global, local and hybrid based) and compare the modelling paradigms that couple with those approaches. The aim is to answer the following questions:

- How do the modelling paradigms compare – in what circumstances is one superior to another?
- Can lower fidelity modelling paradigms be combined with higher fidelity approaches to produce more efficient solves?
- How do the different optimisation approaches compare – for what problems may one approach be chosen over another?
- Can these optimisation approaches be combined in a hierarchy to create a coherent and more efficient approach to tackling the array design problem?

4.3 Modelling approaches

Since the array design and hydrodynamics are a two-way coupled system, optimisation requires multiple iterations in which the quantity of interest (here the power extracted) is evaluated for an adjusted set of control parameters (here the turbine positions). Through analysing the response of the quantity of interest to the manipulation of those parameters, the optimisation algorithm converges to an optimum. The 'analysis code' evaluates the quantity of interest given the control parameters. The choice of optimisation approach is, to some extent governed by the characteristics of the analysis code, for example a gradient-based optimisation approach requires a differentiable analysis code. Additionally, as has been noted, the fidelity of the analysis code is clearly limited by the nature of the optimisation algorithm (if millions of iterations are required for convergence, the analysis code must be computationally cheap). Thus, optimisation method and modelling paradigm (analysis code) are linked and choice of one to some degree dictates the choice of the other.

In this section, two main categories of nonlinear numerical optimisation methodology will be considered and coupled with analysis codes suitable for them. These categories are global and local:

Global optimisation methods are a subset of optimisation methods concerned with finding the optimum over all possible input values. Many are non-deterministic, (for example evolutionary approaches which were pioneered in the late 1970's), using some random or stochastic element to 'jump' around the search space and evaluate the function for a broad range of input vectors, generally with a good global spread – hence the name 'global'. While they do not offer any mathematical guarantee of optimality, if the method has been well implemented they do guarantee a 'good' solution – the best of a large, broad and varied sample of inputs.

Local methods (often termed 'hill-climbers') are gradient-based methods which improve the input vector from an initial 'guess'. Gradients can be obtained from finite differencing, algorithmic differentiation or an appropriate adjoint model. The key benefit of such methods over global approaches are the rapidity of convergence (generally requiring orders-of-magnitude fewer iterations). The trade-off for this is that the optimised solution is sensitive to the initial input vector guess; consequently highly non-convex functionals can be problematic because often parameters are accepted based on greedy criterion – i.e. if and only if the functional is improved. Despite this, one of the benefits are that it can be mathematically proven that these methods will return a (local) optimum rather just a 'good' solution as with global methods [Keane and Nair, 2005]. Depending on the problem these may or may not be close in terms of design as well as QoI value.

4.3.1 Stochastic global optimisation

For this work, a genetic algorithm has been selected to perform global optimisation. Genetic algorithms are search meta-heuristics which emulate evolution through natural selection and hence much of the terminology is jargon borrowed from the field of genetics. Input vectors are encoded in 'chromosomes' and the optimisation process is begun by randomly seeding a (large) 'population'. This population is split into small communities of equal size. Each chromosome's 'fitness' is evaluated using the objective functional (in this case (6.1)) and the fittest chromosome in each community is retained, while the rest are destroyed. A new population (of equal size to the previous) is then generated by 'breeding' and 'mutating' the surviving, fittest, chromosomes. The breeding and mutation operations generally have a random element which ensures that the search space is well explored while basing each population on the best members of the preceding generation ensures that beneficial qualities are nurtured and improved as the algorithm progresses.

The benefits of this somewhat 'scatter-gun' approach is that the greedy criterion on which parameters are accepted is relaxed by a heavily parallelised sampling of the search space. This ensures that parameters are sampled broadly and in a varied manner. The disadvantage of this is that a very large number of samples must be considered and as the search space grows (e.g. through increasing the dimension of the input vector) the number of samples required to achieve a good solution rapidly becomes (potentially unmanageably) large. One of the earliest uses of a genetic algorithm for wind turbine array optimisation was by Mosetti et al. [1994] who used a ten by ten 'checkerboard' couple with a 100-bit chromosome with binary values representing the presence or otherwise of a turbine in each square (i.e. $2^{100} = 1.26\times10^{30}$ possibilities). This problem severely taxed the available computer resources of the time. This illustrates that if such an optimisation approach is to be employed, one must be capable of evaluating chromosomes very cheaply and thus out of necessity the problem must be modelled using a low fidelity analysis code, with this in mind.

Global optimisation methods have been used previously on turbine arrays in the wind industry Mosetti et al. [1994], González et al. [2010] and are used in the commercial code TidalFarmer [Thomson et al., 2011]. The usual approach is to model the problem using a fast engineering 'perturbation' approach: The ambient flow (with no turbines) is computed using a high fidelity fluids code and then stored. An analytical wake model is used to approximate the velocity deficit in the wake of a turbine due to energy extraction, and the deficit due to the array over the whole domain is obtained from the superposition of the wakes onto the ambient flow. The power extracted by the array can then be computed from the resultant velocity field.

Two suitably cheap wake models for this application are the so-called Jensen model [Jensen, 1983] – an empirically based expanding wake model, and the Larsen model [Larsen et al., 1996] which is based upon Prandtl's turbulent boundary layer equations. A third slightly more computationally expensive based upon a shear-layer approximation of the Navier-Stokes equations – the Ainslie model [Anderson and Stuart, 2009] will also be presented.

Jensen model

The Jensen model [Jensen, 1983] assumes a linearly expanding wake with a velocity deficit that depends only on the distance behind the turbine (figure 4.1).

Hence the shear layer is assumed infinitely thin. The model is defined by the diameter of its wake:

$$d_w(x) = D + 2kx, \tag{4.4}$$

and the velocity reduction factor:

$$r(x) = 1 - \frac{1 - \sqrt{1 - C_t}}{\left(1 + \frac{2kx}{D}\right)^2}, \tag{4.5}$$

where C_t is the dimensionless turbine thrust coefficient (the thrust force per unit of frontal area per unit of incompressible dynamic pressure), k is often referred to as the wake decay constant and is equal to the angle of the wake as measured from the line parallel to the flow direction at the point of the turbine. The diameter of the turbine is given by D whilst x is the distance downstream of the turbine parallel to the flow at the point of the turbine.

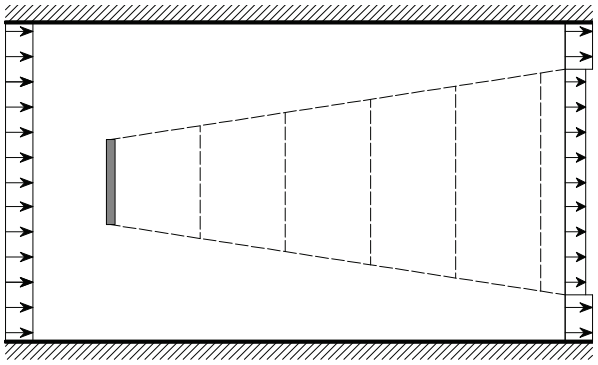


Figure 4.1: Sketch of velocity deficit behind a 20 m wide turbine computed using the Jensen wake model.

Larsen model

The Larsen model [Larsen et al., 1996] (figure 4.2) derives from Prandtl's turbulent boundary layer equations. Both first and second order approximate solutions to the boundary layer equations exist, but it has been demonstrated that in practice there is little difference between the solutions as discrepancies only lie in the near wake field [Sørensen et al., 2008] where it would anyway be inefficient to place a further turbine. The diameter of the wake in the Larsen model is defined as:

$$d_w(x) = 2\left(\frac{35}{2\pi}\right)^{\frac{1}{5}} (3c_1^2)^{\frac{1}{5}} (c_t A x)^{\frac{1}{3}}, \tag{4.6}$$

where c_1 is a constant based on the Prandtl mixing length (defined in Appendix) and A is the rotor disc area. The first order Larsen reduction factor is:

$$r(x,y) = 1 - \frac{1}{9}\left(c_t A x^{-2}\right)^{\frac{1}{3}} \left\{ y^{\frac{3}{2}} \left(3c_1^2 c_t A x\right)^{-\frac{1}{2}} - \left(\frac{35}{2\pi}\right)^{\frac{3}{10}} (3c_1^2)^{-\frac{1}{5}} \right\}^2, \tag{4.7}$$

where x and y are the distance downstream from the turbine parallel to the flow at the turbine and the orthogonal distance to the centerline of the wake, respectively.

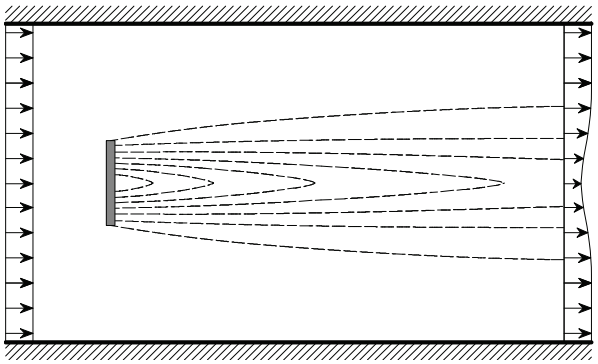


Figure 4.2: Sketch of velocity deficit behind a 20 m wide turbine computed using the Larsen model.

Ainslie model

The simplified thin shear layer approximation of the Navier-Stokes equations is

$$u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial r} + \frac{1}{r}\frac{\partial (r\overline{u'v'})}{\partial r} = 0, \tag{4.8a}$$

$$\frac{1}{r}\frac{\partial r v}{\partial r} + \frac{\partial u}{\partial x} = 0, \tag{4.8b}$$

where u is the axial velocity, v the radial velocity, x the axial distance from the turbine, r the radial distance from the turbine and $\overline{u'v'}$ is the Reynolds stress which is modeled with the eddy viscosity, $epsilon$

$$-\overline{u'v'} = \epsilon(x)\frac{\partial u}{\partial r}, \tag{4.9}$$

where ϵ is the sum of the ambient eddy viscosity ϵ_a and the wake eddy viscosity,

$$\epsilon_w(x) = kb(u_{\text{fs}} - u_c(x)), \tag{4.10}$$

where $u_c(x)$ is the wake centreline velocity, b is the wake width and k is an empirically-found constant.

The solution to these gives a wake profile that is Gaussian in shape – the scaling of which changes with distance from the turbine. The system is solved by finite-differencing using a centre-differencing scheme for a half-wake, then reflected along the turbine axis. Boundary conditions are imposed along this centreline (where the radial velocity is zero, and at the edge of the wake where the velocity is known from the ambient. This approach is outlined in Anderson and Stuart [2009].

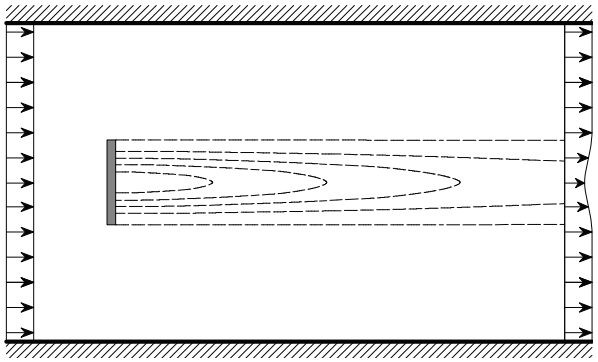


Figure 4.3: Sketch of velocity deficit behind a 20 m wide turbine computed using the Ainslie model. With a higher thrust coefficient the wake would grow in width downstream of the turbine.

4.3.2 Blockage correction

Very recently, Stansby and Stallard [2016] presented a similar approach in which a simple gradient-based optimisation approach was employed to optimise the locations of a small array of turbines modelled via the superposition of analytical velocity deficit profiles. A (local) blockage correction factor (developed from laboratory tank testing) was used to adjust the analytical wakes. Such a blockage-correction is not uncommon, however the extent to which this factor is dependent

upon the domain-geometry is unclear – and with realistic bathymetry and coastal geometry it is difficult to see how it will be employed. For this work, blockage correction has not been considered, though it is an area of interest for future work to improve the utility and accuracy of these uncoupled modelling approaches.

4.3.3 Gradient-based local optimisation

As with global, local optimisation is a set comprising numerous methods. For this work, the SLSQP algorithm [Kraft, 1988] has been selected as it supports both boundary and inequality constraints – which can be used to enforce a minimum distance between individual turbines. The implementation in the *SciPy* package [Jones et al., 2001–] was used.

In general, gradient-based optimisation works by evaluating the quantity of interest (QoI) for a starting guess of the control parameter(s) and then computes the derivative of the QoI with respect to those control parameters (the functional gradient). This can be used to determine a 'direction' in which to adjust the control parameters to improve the QoI. A step length for this adjustment is chosen algorithmically and the new control parameters are evaluated and accepted if they are an improvement on the previous functional evaluation. Gradient evaluations are only made once parameters are accepted as these are generally more computationally costly than QoI evaluations.

The benefits of this approach are the orders of magnitude fewer iterations required, meaning that, with the same computational budget the analysis code can be far higher fidelity (though must be differentiable) and thus represent the underlying physics better. The key drawback is the local nature of the optimisation – meaning the optimised parameters are sensitive to the initial guesses made and are by no means guaranteed to be globally optimal.

The analysis codes based upon a perturbation approach which were discussed in the previous sub-section could also be optimised locally – the gradient could be evaluated through finite differencing (perturbing each turbine in each direction in turn and recording the effect on the QoI) or through algorithmic differentiation (note that the discontinuity resulting from the infinitely thin shear layer as assumed in the Jensen model renders it piece-wise differentiable – which is generally sufficient for gradient-based optimisation methods [Culley et al., 2016a]).

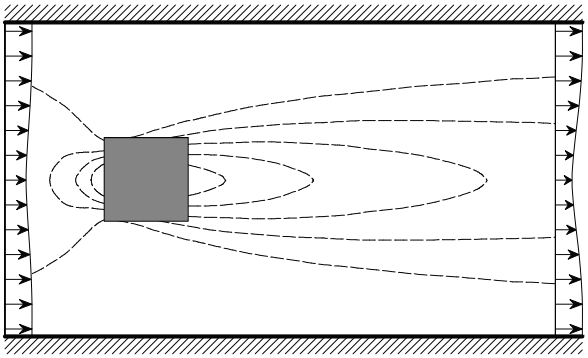


Figure 4.4: Wake behind a 20 m wide turbine in 2 ms^{-1} flow using the shallow water model with the turbine modelled as an area of increased bottom friction.

Shallow water model with parametrised friction

We recall the shallow water equations,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} - \nu\nabla^2\mathbf{u} + g\nabla\eta + \frac{c_b + c_t}{H}\|\mathbf{u}\|\mathbf{u} = 0, \tag{4.11a}$$

$$\frac{\partial \eta}{\partial t} + \nabla\cdot(H\mathbf{u}) = 0, \tag{4.11b}$$

where $\nu \in \mathbb{R}$ and $g \in \mathbb{R}$ are the kinematic water viscosity and acceleration due to gravity respectively, $H : (0,T)\times\Omega \to \mathbb{R}$ is the total water depth defined as the sum of the bottom depth, $h : (0,T)\times\Omega \to \mathbb{R}$ and the free-surface displacement, $\eta : (0,T)\times\Omega \to \mathbb{R}$. The depth-averaged velocity field, $\mathbf{u} : (0,T)\times\Omega \to \mathbb{R}^2$, is the solution to the equations while c_b and c_t are the bottom friction and the augmented bottom friction representation of the turbine(s).

Computing the gradient

One reason that gradient-based optimisation methods are not more prevalent in the coastal design community is the need to compute the gradient of the QoI with respect to the controls in a way that is computationally efficient, even for large control vectores. There are four potential methodologies for the computation of the gradient, depending on the analysis model being used:

Where possible, analytical models may be differentiated using calculus to provide an analytical differential. With large arrays, however, a 'network' of dependency forms where the behaviour of downstream turbines may be contingent on

the behaviour of several upstream turbines who may also be affected by turbines further upstream, etc. Accounting for these dependencies is not trivial and consequently finding the analytical derivative becomes complex and the approach is largely unfeasible if not done pragmatically.

The cost of finite differencing scales linearly with the number of turbines. Once the QoI has been computed for the micro-siting design in question, each control parameter is perturbed individually, and the QoI is recomputed. Division of the change in QoI with the applied change in control parameter provides an estimate of the sensitivity of the QoI with respect to that parameter. Clearly computation of the gradient in this manner requires the analytical code to be solved $2N$ times for the x and y direction for each turbine. While it may be able to improve upon the efficiency of this [Fang et al., 2011], for expensive analytical codes or large arrays the cost of many additional solves may be unfeasible.

For the final two methods, the computational expense is dependent upon whether the forward 'tangent linear' or reverse 'adjoint' mode is used. In situations where the size of the quantity of interest (in terms of degrees of freedom) is greater than the control vector, the tangent linear approach should be employed and will provide the gradient at a cost which scales linearly with the size of the control vector. In contrast, the adjoint mode scales with the size of the QoI and independently of the control vector. In our case, where the QoI is a scalar, the adjoint mode is used and hence the cost of computing the gradient with respect to all controls reduces to the equivalent of one functional evaluation.

In practice, the adjoint implementation is often generated programmatically using algorithmic differentiation (AD). In AD, the analytical code is broken down into a series of basic, fundamental operations (additions, subtractions etc.). The chain rule is then applied repeatedly on these base operations to compute the gradient. The derivatives are as accurate as the original computation, without truncation errors that result from finite differencing. For example, assume our analysis code simply computes $\sin(x)$, 'under the hood' this may be achieved by solving some appropriately truncated version of the Taylor series approximation of sin. All the multiplication, division, subtraction and addition involved in that operation would be recorded in a computational graph when the initial solve is conducted. This graph (sometimes referred to as a 'tape') is then traced backward and the derivative of each of the basic operations is computed and combined using the chain rule to determine the derivative of the function.

With the adjoint approach, a similar computational graph is constructed as the

analytical code is solved from input to output. However, as the 'tape' is reversed, rather than differentiate each base operation, each operation is differentiated symbolically – i.e. given the $\sin(x)$ example above, the code would *symbolically* differentiate sin, then evaluate the result for a given x. The main drawback of this method is that general AD tools are often difficult and intrusive to use – consequently the majority of coastal flow models do not facilitate computation of adjoints.

4.3.4 Hybrid Approach

'Basin-hopping' is a hybrid non-linear optimisation approach developed by Wales and Doye [1997] with the aim of combining the strengths of both stochastic global and gradient-based local optimisation methodologies. The central idea is to randomly perturb the control parameters (as in a global technique) then locally optimise from this initial guess. Once the local optimisation has converged, the optimised control parameters are accepted or rejected based on the Metropolis criterion [Metropolis et al., 1953], rather than greedily as in pure hill-climbing algorithms. The Metropolis criterion is probabilistic in nature – with the ability to accept less optimal configurations dependent upon the QoI of the current state as compared to the initial configuration. This criterion is used as standard in Monte Carlo algorithms and forms the backbone of the simulated annealing optimisation algorithm.

Basin-hopping works because – having been locally optimised – each random perturbation (eventually) leads to a local optimum, so in a way it is a global optimisation of local minima (or maxima as in our case) – thus the user is guaranteed local optimality *plus* assurance that it is the best solution observed from a wide and varied exploration of the search space. Once again, the implementation from the *SciPy* [Jones et al., 2001–] package is used which was developed from Wales and Doye [1997].

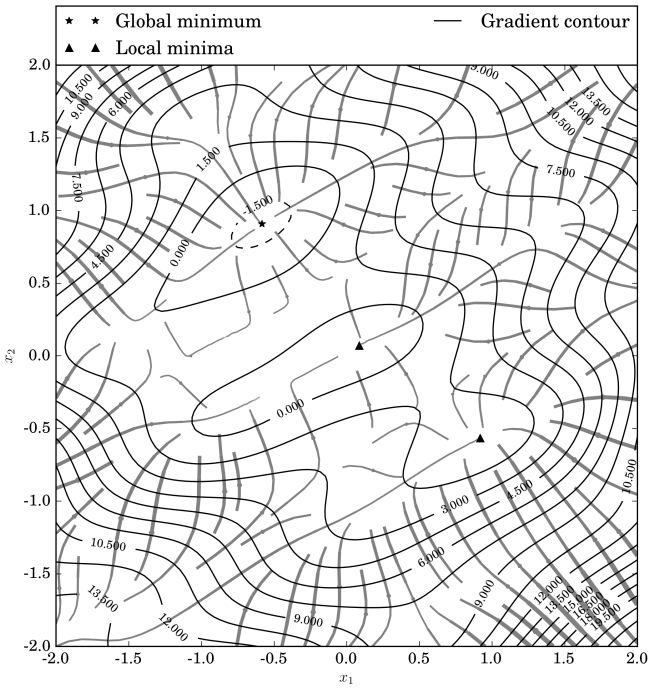


Figure 4.5: Plot of an arbitrary (differentiable) function, $f(x_1, x_2)$, showing direction of gradient (grey), contour lines, local minima and global minimum.

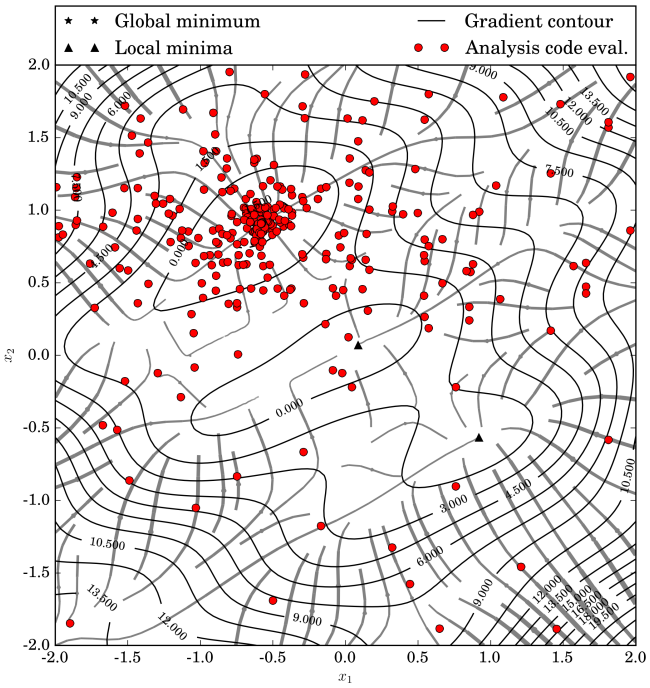


Figure 4.6: Demonstration of global optimisation methodology, points show parameters for which the analysis code was evaluated, global optimum was found to 3 significant figures in 136 evaluations

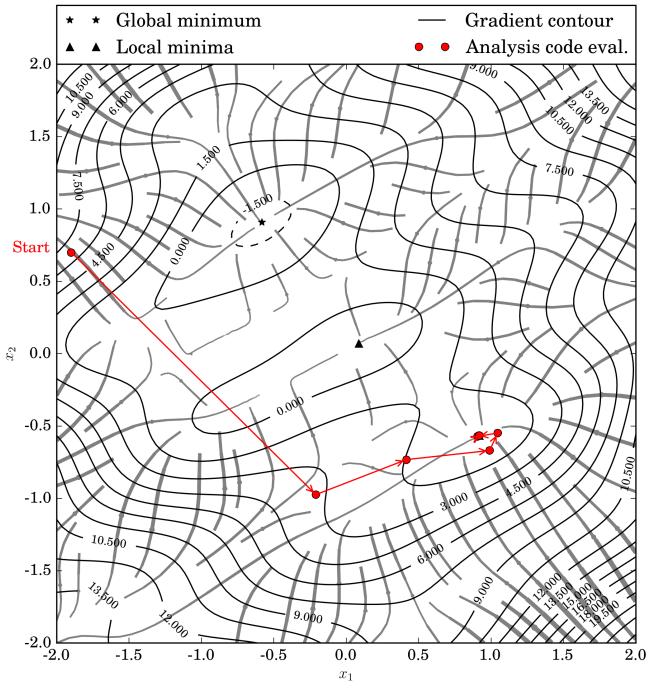


Figure 4.7: Demonstration of local optimisation methodology, function and gradient are evaluated at the points marked on the plot, arrows show the change in parameters from one iteration to the next – these steps are based on the derivative magnitude and direction. Convergence was reached in 6 iterations

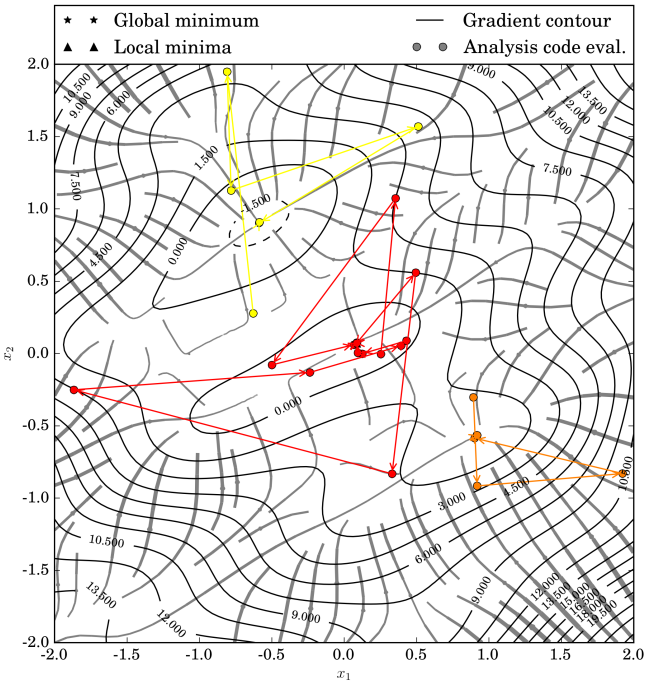


Figure 4.8: Demonstration of basin-hopping optimisation methodology. Parameters are chosen randomly then optimised locally as above (this constitutes one iteration). Once this local optimisation has converged, the parameters are perturbed randomly (shown be a change in colour of the points and arrows on the plot) and locally optimised again. In this case, the algorithm was allowed 3 iterations and used 36 analysis code evaluations.

4.4 Single turbine comparison

As has been seen, design optimisation is a combination of the analysis code and optimisation methodology: Choice of one constrains choice of the other. Consequently, we shall first compare the analysis codes' performance at modelling, first an individual turbine and then an array. Finally, optimisation will be added in so that the relative performances on the array design problem can be assessed.

4.4.1 Calibration

Each of the analytical wake models described in section 4.3.1 have a number of free parameters which may be chosen to adjust the characteristics of the wake deficit. For example, for the Jensen model described by (4.4) and (4.5), the parameters C_t, k and D were optimised. While these parameters refer to real-world measurements (for example D, the turbine diameter), in the interests of fairness all parameters were tuned to minimise the difference between the approximated wake deficit and the wake deficit computed by depth-averaging a high fidelity, three-dimensional CFD simulation of an actuator disk of 18.5 m diameter. The CFD model was based on an actuator disc parametrisation of a turbine and Reynolds-averaged Navier-Stokes implementation with dynamically adaptive mesh optimisation and k-ω-SST turbulence model, implemented in *Fluidity* [AMCG, 2010, Abolghasemi et al., 2016]. This calibration was carried out in a flow of 3 ms^{-1} with a thrust coefficient of 0.15.

For optimisation of the parameters, a genetic algorithm was used where the quantity of interest was the sum of the absolute difference between corresponding points (i.e. the L^1 norm) in the wake of the high-fidelity model and the model under consideration. The goal of the optimisation was to minimise this difference. The difference was computed on a grid of 100 by 100 points equally spaced over 2 turbine diameters either side of the turbine centreline and in three different longitudinal domains – from 0 to 10D, from 3 to 10D and from 5 to 10D (see figure 4.9. Each of these domains puts respectively less emphasis on the importance of matching the near wake region – a place that a second turbine would be unlikely to be placed and so a region of the wake that it is less important to accurately model. Table 4.1 shows the sum of the absolute difference in velocity deficits over those points, for the set of optimal set of parameters which minimised the difference.

It should be noted that the shallow-water model used here (and throughout this

Table 4.1: The percentage velocity difference between wake deficit computed by high-fidelity CFD code and as computed by a wake model, taken over domains of different length.

Wake model	**Average percentage difference** *over longitudinal domain of:*		
	0 to $10D$	3 to $10D$	5 to $10D$
Jensen	2.07	1.84	1.75
Larsen	1.20	0.63	0.46
Ainslie	1.11	0.55	0.44
SW model	2.54	2.59	2.46

chapter) has a fixed viscosity. LES and HLES turbulence models are currently under development and should improve the match. For the remainder of the chapter the models will be used with the optimised parameters from this calibration.

4.4.2 Comparison

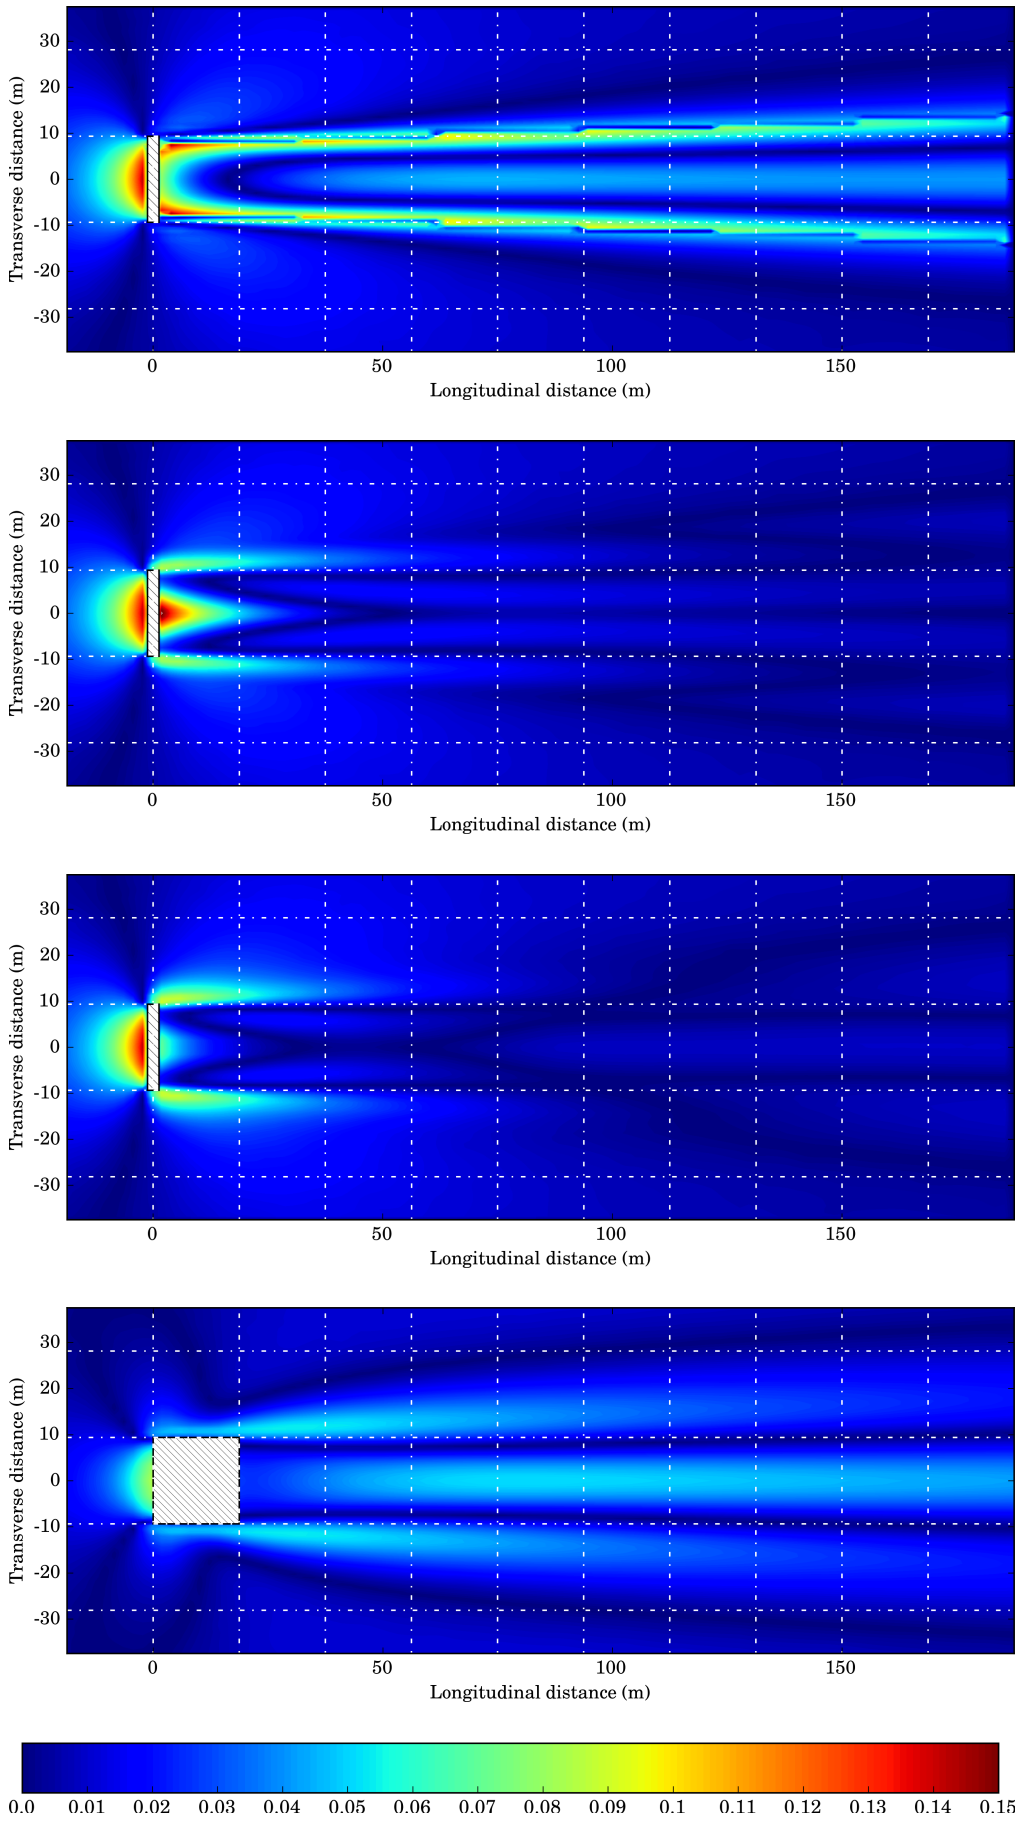


Figure 4.9: The absolute difference (colourbar in ms^{-1}) between (working down from top) Jensen, Larsen, Ainslie and shallow-water models with parametrised friction and a depth-averaged high-fidelity CFD actuator disk computation from Abolghasemi et al. [2016]. The turbine is placed in a 3 ms^{-1} flow. Dotted lines show turbine diameter. Note that these plots interpolate the 100 by 100 sample points into a continuous colour plot – the shear layer in the Jensen model is resolved poorly because of the relatively low sample resolution compared to the shallow angle it transcribes in space and the large discontinuity it represents in the data.

The purpose of the model comparisons in this section is to demonstrate the differences between the analytical perturbation approach and the shallow-water approach with parametrised friction. Validation of models against laboratory data (such as those conducted by Stallard et al. [2013]) has been studied extensively and is not the focus of this work. Rather, the focus is to demonstrate the key differences and investigate how these affect the choice of model for the design optimisation of arrays – and illustrate how the outputs of that model choice may be affected by the modelling simplifications. The model being used for this comparison is itself validated in Abolghasemi et al. [2016].

Figure 4.9 contains a series of plots showing the spatial variation of the absolute difference between the wake deficit computed by the wake models and the depth-averaged wake deficit computed by the high fidelity 3-D CFD code. All models were able to reproduce the reference flow relatively accurately (0.15 ms^{-1} in a 3 ms^{-1} flow). The most important limitation of the Jensen, Larsen and Ainslie wake models is their inability to capture local blockage effects. While the wake deficit is captured relatively well (especially for the Larsen and Ainslie models >3D downstream), the upstream deceleration of the flow caused by the blockage represented by the turbine and the acceleration of the bypass flow around the turbine – a result of mass conservation – are not represented at all. This is a serious limitation of the analytical wake models as the blockage effects of individual turbines tend to combine and augment each other, resulting in array-scale so-called 'global' blockage effects (discussed in more detail in following sections). The absence of local blockage also means that local flow acceleration cannot be exploited – for example through row staggering turbines.

While both the Larsen and Ainslie models reproduce a good likeness of the velocity deficit in the far wake (>3D) region across the width of the domain, the Jensen modelling assumption of an infinitely thin shear layer limits the quality of its reproduction of the deficit in those transition zones. However, even in this region, the absolute difference between the two models is just 0.1 ms^{-1} which – depending on the turbine power curve – is unlikely to cause too great an inaccuracy in the computation of the power extracted by any downstream turbines located in this zone.

In contrast, the shallow-water model captures both the upstream deceleration and the accelerated flow either side of the turbine in the near-wake region implying that, unlike the other models, the blockage effects are being captured to some extent. However, further downstream the performance is generally less good

– though the maximum absolute difference peaks at 0.045 ms^{-1} which, again, is unlikely to result in significant errors in power computation.

4.5 Array comparison

When several turbines are placed in proximity to one another, complex hydrodynamic interactions can develop. While the models performed generally well at recreating the velocity deficit of a single turbine, if they are to be used as a tool for array design and analysis they must also perform well at recreating the array interactions. The results of the previous section have already demonstrated that the shallow-water model deals with *local* blockage more successfully. Once the combined action of local blockage around many turbines translates into *global* blockage around an array the ability of models to capture these effects grows significantly in importance in particular because inaccuracies in the flow velocity has a cubic impact upon the power output. The velocity and thrust coefficient were set as for the individual turnine calibration.

4.5.1 Wake combination approaches

In moving from using the Jensen, Larsen and Ainslie models to represent a single turbine to a full array, there is likely to be circumstances where the wakes of two or more turbines overlap. The way in which the two wakes are combined in these overlapping areas is found using a so-called wake combination formula where j upstream turbines reduce the velocity at point i, for example using

$$\left(1-\frac{u_i}{u_{\text{amb}}}\right)^2=\sum_j\left(1-\frac{u_{ij}}{u_j}\right)^2. \tag{4.12}$$

Here, u_{amb} is the undisturbed flow speed and u_{ij} is the flow speed at point i due to turbine j. This is the 'sum of squares' combination approach, though there are various others. The different approaches are discussed (for wind turbines) in Renkema [2007] where (4.12) was found to perform the best. Consequently this is the method that will be used to combine wakes for this application.

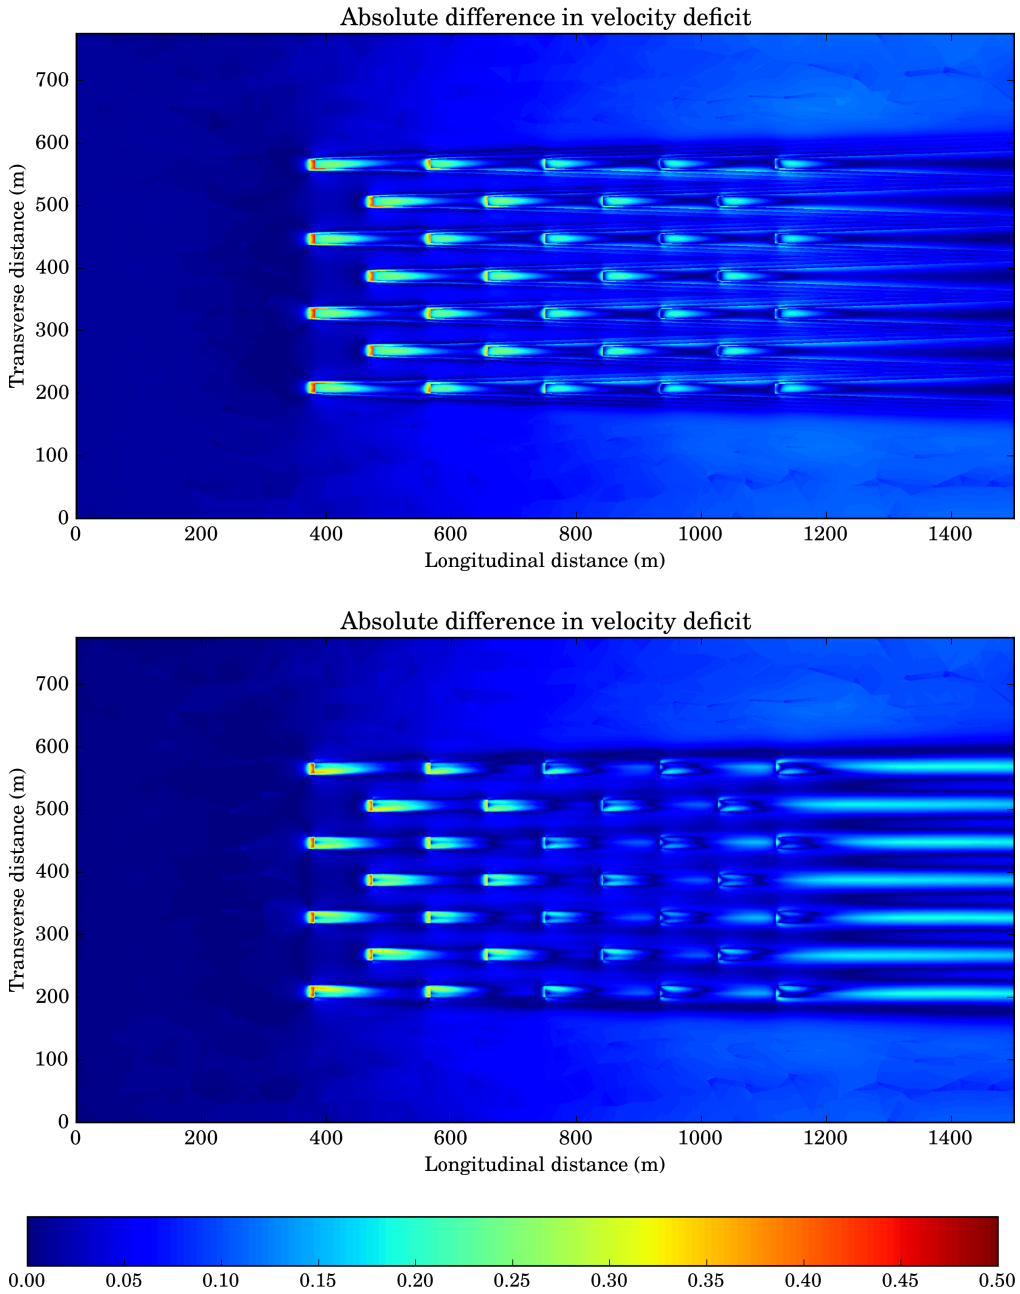


Figure 4.10: The absolute difference (colour-bar in ms^{-1}) between Jensen (top) and Larsen (bottom), models and a depth-averaged high-fidelity CFD actuator disk computation. For a staggered array of turbines placed in a simple channel with 3 ms^{-1} flow.

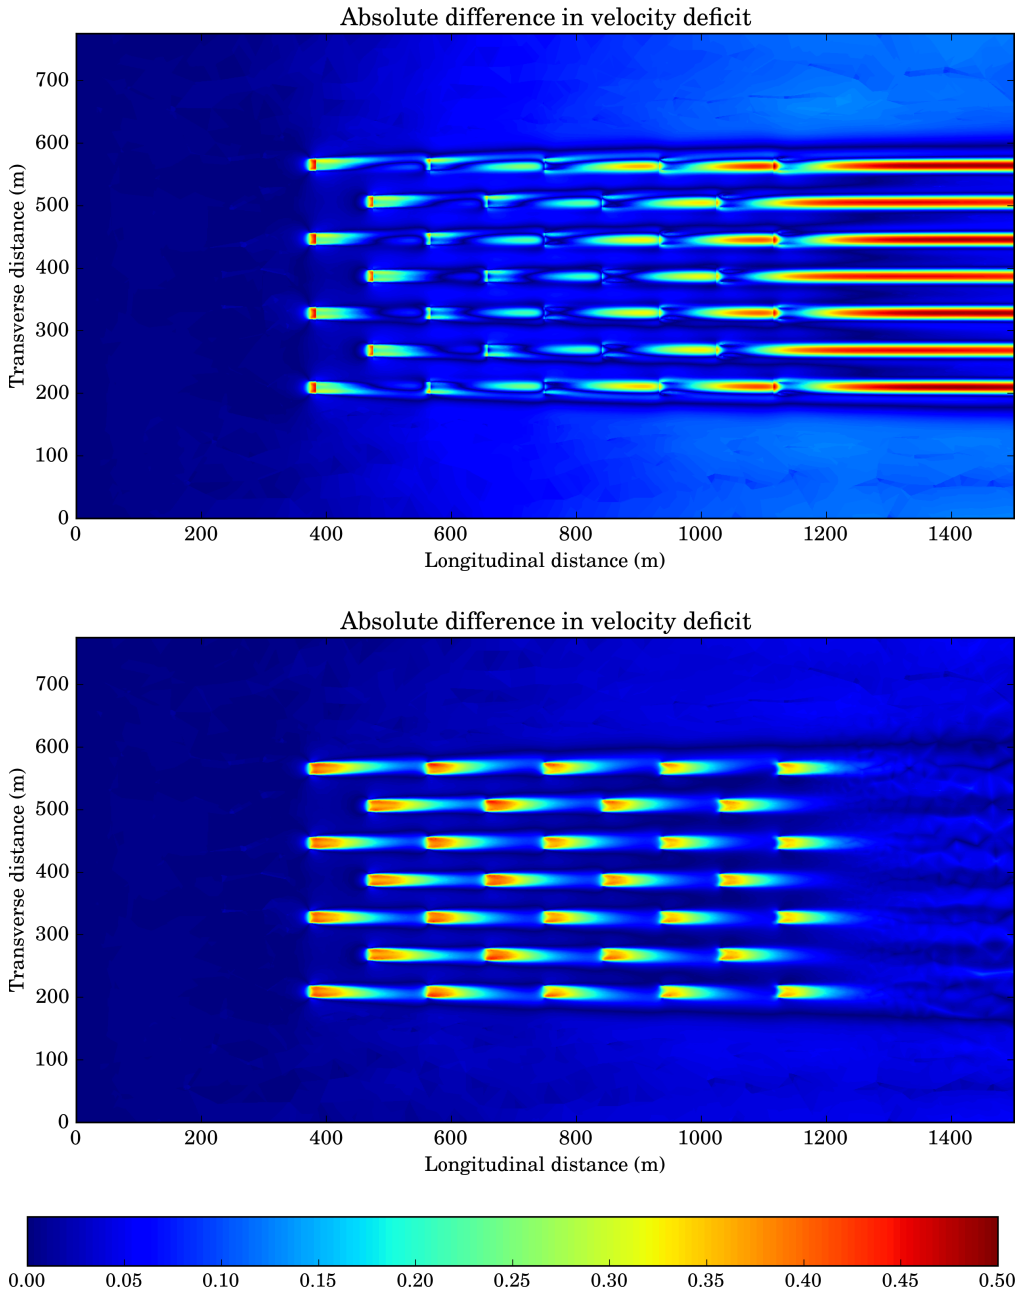


Figure 4.11: The absolute difference (colour-bar in ms^{-1}) between Ainslie and shallow-water models and a depth-averaged high-fidelity CFD actuator disk computation. For a turbine placed in a simple channel with 3 ms^{-1} flow.

Table 4.2: Showing the average of the percentage difference between wake deficit computed by high-fidelity CFD code.

Wake model	**Avg. percentage difference**
Jensen	2.74
Larsen	5.13
Ainslie	8.55
SW model	1.30

The comparison of arrays of turbines modelled by the different methods to the high-fidelity actuator disk run was conducted in much the same way as in the previous section for individual turbines. An array of 32 turbines was modelled in a staggered formation, placed in a simple channel measuring 1500 metres in length, 775 metres in width and 33 metres in depth. The turbine parameters were defined as in the previous section. The turbine area spans from 375 to 1125 metres from the inflow and occupies the central 375 metres of the channel width, this gave a longitudinal spacing of just over $5D$ and a lateral spacing of circa $6.75D$ (i.e. adjacent rows of turbines are offset laterally by circa $3.38D$).

Once again, the model from Abolghasemi et al. [2016] was used to provide the high-fidelity comparison, and the spatial absolute error in velocity deficit is plotted in figure 4.10 for the Jensen (top) and Larsen models (bottom) and in figure 4.11 for the Ainslie (top) and shallow-water (bottom) models. These comparisons were computed using discrete sampling on a grid of 1000 by 1000 points over the whole domain and linearly interpolated to provide the plots. The average percentage difference over the domain for each model is shown in table 4.2.

Overall the models are able to reproduce the flow profile well, however in the near-wake region of the turbines the difference can be large, and errors tend to be compounded in downstream rows so that the total error grows – this can be seen in the results for the Ainslie model. Another difference between the perturbation-approach models (Jensen, Larsne and Ainslie) and the shallow-water model is in the flow bypassing the array. For the perturbation-approach models, the difference from reference velocity of this bypass flow grows from the leading edge of the array to circa 15% by the downstream end of the domain, while for the shallow-water model the discrepancy is much less; reaching circa 5% by the end of the domain. This is because blockage effects are being captured by the shallow-water model so the acceleration of this bypass flow is being modelled, unlike the perturbation-approach wake models. While this omission was not overly prob-

lematic for a single turbine (as demonstrated in the previous section), one of the effects of having an array of turbines is that blockage is compounded and thus so too is the error of neglecting to model it.

4.6 Optimisation

The manner in which the modelling approaches have been presented in this work is as motivated by the requirements of the optimisation methodology – for example global optimisation requires a computationally cheap analysis model which motivated discussion on perturbation methods and the Jensen, Larsen and Ainslie models: There is an intrinsic relationship between choice of optimisation method and choice of analysis model. Rather than compromise on the fidelity of the analysis model, oftentimes in previous work (for example Malki et al. [2014]) the optimisation problem itself has instead been simplified: Rather than allowing turbines to move independently they are confined to (usually staggered) rows, the optimal longitudinal and transverse spacing between which is then investigated by brute force or trial and error. For the purposes of straightforward reference this will be termed the 'spacing-adjustment' strategy.

The purpose of this section is to evaluate such an approach and see what it can illuminate about the differences between the perturbation and the shallow-water approaches. Since the Larsen model performed consistently well in both the individual turbine comparison and the array comparison it alone of the wake models will be used from here forward.

The discussion in this section will be based around two test scenarios – designed to highlight certain interesting features about both the spacing-adjustment optimisation strategy and the differences in the underlying analysis approaches. These two scenarios are a flat-bottomed simple channel and a channel with an island in the middle to promote ambient flow non-uniformity (figures 4.12 and 4.13).

4.6.1 Test scenario 1

The domain for this test scenario measures 1300 by 600 metres, allowing a large centrally located turbine area measuring 700 by 500 metres. The lateral boundaries of the domain were modelled with a free-slip condition and the right hand outlet boundary with a $\eta = 0$ Dirichlet condition. The inflow is driven by a con-

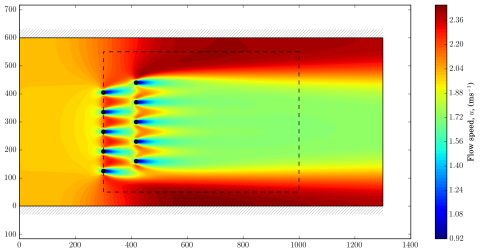


Figure 4.12: Flow through test domain of scenario 1 with two staggered rows of turbines (turbines represented with the shallow-water model).

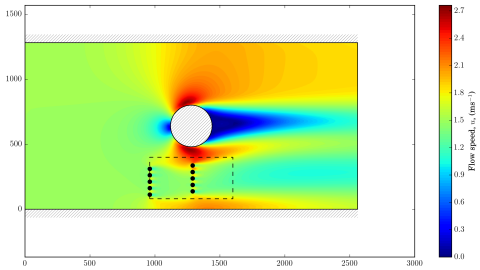


Figure 4.13: Flow through test domain of scenario 2 with two staggered rows of turbines (turbines represented with the shallow-water model).

stant velocity Dirichlet condition of 2 ms^{-1}. Uni-directional flow has been chosen in this simple example as array dynamics are much easier to interpret. Transient scenarios will be used in later sections. The other parameters are as follows:

Parameter	**Value**
Water depth	50 m
Viscosity coefficient	$\nu = 3\ m^2 s^{-1}$
Acceleration due to gravity	$g = 9.81\ m\ s^{-2}$
Water density	$\rho = 1000\ kg\ m^{-3}$
Bottom friction coefficient	$c_b = 0.0025$
Turbine radii	$r = 10$ m

An array of 10 turbines placed in two rows was tested, with the power that the array extracts sampled for a range of different longitudinal and transverse spacings. The downstream row of turbines is staggered – offset in the transverse direction by half the transverse spacing. The upstream row's longitudinal position remains static – at the upstream end of the turbine area – and the downstream row moves relative to it.

Local blockage

The graphs in figure 4.14 and 4.15 was produced from a parameter study of the longitudinal and transverse spacing of the two rows.

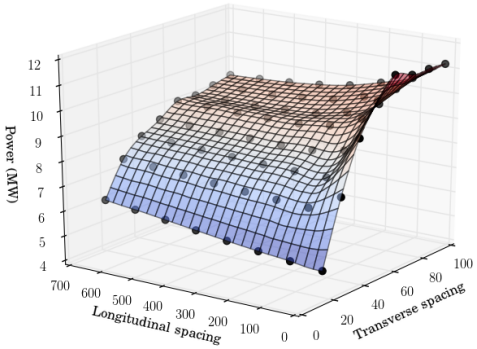


Figure 4.14: Power extracted – as computed using the shallow water model – by an array of 10 turbines as the transverse and longitudinal spacing between the two rows is adjusted.

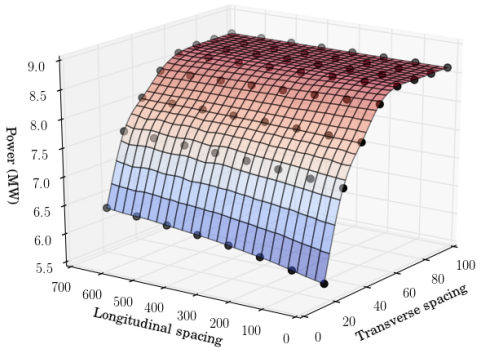


Figure 4.15: Power extracted – as computed by the Larsen model – to the same parameter adjustments as in figure 4.14.

Discussion

This set up is designed to highlight the potential impact of local acceleration on array design, and the potential danger of using a wake model which does not naturally account for it. There are some quantitative differences between the runs conducted with the shallow water model and the analytical model, for example the peak power achieved (circa 11 MW for the shallow-water model versus 8.5 for the wake model). However, such differences are arguably not as important as the qualitative similarities between the two plots. Recall that the purpose of these models is to find the optimal array design – so scaling differences in the objective are less important than the form of the surfaces being broadly similar. This similarity should indicate that both models would agree on the array layout that maximises the power output of the array, and it is determining this layout that is the goal of this work.

As such, there is a strong similarity between the surfaces in 4.14 and 4.15 except for one key difference; the ridge that runs along the line of 60 m transverse spacing. This ridge is created because at this spacing, the downstream row of turbines is ideally placed off the shoulders of the upstream turbines and is most benefiting from the bypass flow acceleration around the sides of them. This acceleration, caused by the local blockage resulting from the presence of the turbine is not captured by the wake model. Thus in situations where local flow acceleration is likely to be an important mechanism within the array hydrodynamics, it is important that coupled modelling approaches – like the shallow water model – be used.

4.6.2 Test scenario 2

The domain for test scenario 2 (figure 4.13) measures 2560 by 1280 with a centrally located circular island of 320 metre diameter and a 640 by 320 metre turbine area located longitudinally central and in the middle of the channel to the 'south' of the island. The parameters and boundary conditions were set as in the previous test scenario, although the inflow was reduced to 1.5 ms^{-1}, so the faster flowing water coming of the island would be in the turbines' operating range.

The presence of the island allows the impact on optimisation methodology of two important processes to be observed: First is local flow acceleration – caused by the zone of fast-flowing water that is projected into the turbine area due to the island. The second is global blockage – the presence of turbines to the 'south' of the island causing bulk flow redirection through the 'northern' passage.

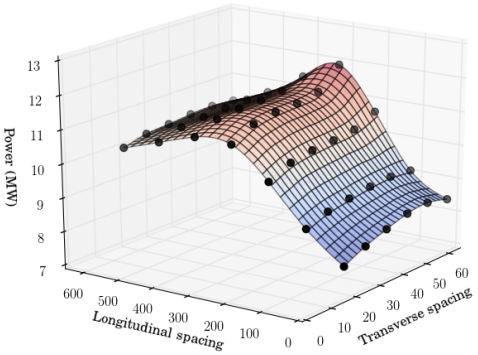


Figure 4.16: Plot showing the impact on power of array spacing in test scenario 2, using the shallow water model.

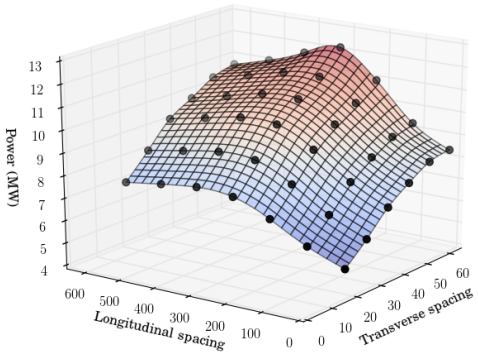


Figure 4.17: Plot showing the impact of row spacing on power using the Larsen wake model.

Discussion

As before, there are some qualitative similarities between the two plots, figures 4.16 and 4.17. The ridge running along the longitudinal spacing of 300 metres shows where the downstream row of turbines moves into the local acceleration of the 'jet' coming off the island. Note that the front row is, as before, fixed so cannot benefit from this faster flow. As before, with small longitudinal spacing, the shallow-water model shows a 'hump' where the back row benefits from local acceleration of bypass flow from the front row, while this is not evident in figure 4.17.

A key shortcoming of the spacing-adjustment optimisation strategy – while it

works reasonably well for flat-bottomed simple channels – is that when the natural localised flow variation is important (as in this case with the 'jet'), it may be that the rows should be located to exploit that variation rather than mitigate array interaction. Local flow variation can be capitalised upon to a greater extent by allowing turbines to be adjusted freely – rather than requiring them to conform to rows. A gradient-based optimisation with such freedom performed by *OpenTidalFarm* yields the array layout in figure 4.18. The 'barrage' that has been formed by the optimisation makes sense from a global blockage sense, and is located to maximally exploit the local flow variation. This motivates the use of the more comprehensive optimisation methods tested in the next section.

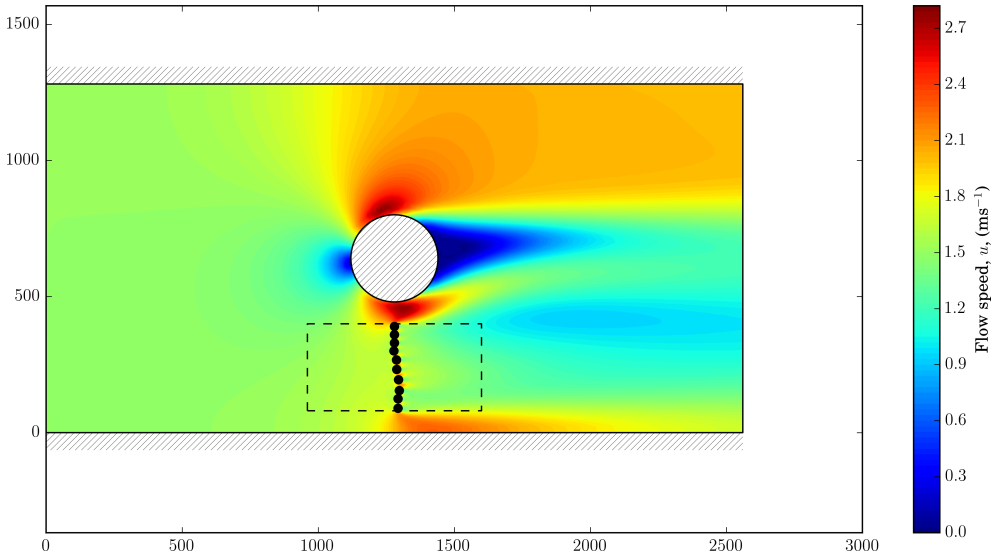


Figure 4.18: Micro-siting location as found by gradient-based optimisation with shallow water model, showing the benefits of controlling each turbine individually – with no designer intuition, a single barrage automatically forms at the location of highest local flow speed. This optimisation took 62 iterations, a total run time of 1hr58 on a single CPU and achieved an array which extracts 12.43 MW.

Temporarily setting aside the optimisation approach, in terms of modelling strategy the differences between figures 4.16 and 4.17 highlight the impact of another hydrodynamic mechanism; global blockage. Since the wake modelling approach uses superposition of wake deficit onto the ambient velocity – i.e. turbines are uncoupled from the flow – the redirection of flow around the unobstructed side of the island due to the presence of an array is not captured. Thus no matter what array is placed in the channel, the same amount of energy is flowing through that channel, and placing more turbines captures a greater and greater proportion of that. In reality, however, as the channel is obstructed and flow is diverted, the

total energy flowing through the channel reduces – and while adding more turbines may capture a great portion of the total energy flowing through, it also reduces that total. This means that as more turbines are added, the power extracted by the array flattens off at a higher rate than a superposition approach will tend to predict (and then starts to reduce, which the superposition approach will not capture). This is highlighted in figure 4.19 which shows the power extracted by optimised arrays comprising different numbers of turbines as modelled with the shallow water method and the perturbation (analytical wake model) approach (to be clear, for each data point gradient-based optimisation with the shallow-water model was used to find an optimised array formation and that array formation was then evaluated with the shallow-water model and the Larsen model).

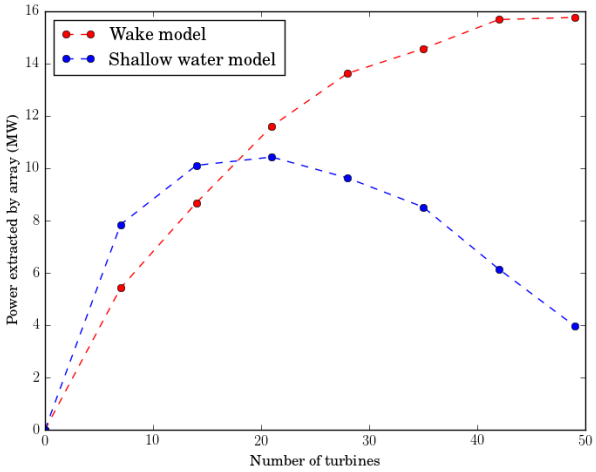


Figure 4.19: Plot showing the power extracted by optimised arrays comprising different numbers of turbines as modelled by the shallow water model and Larsen wake model. Since the wake modelling approach uses the ambient flow, the flow redirection resulting from the blockage caused by the turbines is not modelled, so the power extracted by the array is over-estimated for higher turbine numbers.

4.7 Optimisation comparison

While limitation of the degrees of freedom of the control parameter greatly simplifies the optimisation (with two free parameters, a brute force optimisation of the longitudinal and transverse spacing is trivial) it also limits the quality of the solution that may be obtained. As has been seen, local flow variations caused by bathymetric features or coastal geometry makes the problem far more complex; the n'th row of turbines may produce more power if moved to exploit some local flow feature rather than being naively placed at the 'optimal' downstream spac-

ing from the $n-1$'th row. While increasing the degrees of freedom of the control parameter increases the size of the search space (realistically beyond around five degrees of freedom the problem is no longer solvable by brute force), the potential quality of the eventual solution clearly increases as the restrictions on turbine placement are removed. The logical goal is to allow the placement of each turbine to be fine-tuned individually, without restricting the array geometry to rows. Such an approach leads to a control parameters with $2N$ degrees of freedom; one for each of the two dimensional Cartesian co-ordinates of the N turbines.

In this section several combinations of optimisation approaches and models will be tested on three idealised example scenarios. Gradient-based approaches will be used with the shallow water model and Larsen model, and both global optimisation and the hybrid basin-hopping approach will be used with the Larsen model. Finally, a global-then-gradient-based optimisation approach will be used in which the turbine locations returned from the global optimisation with the Larsen model are used as a starting guess and improved by gradient-based optimisation with the shallow water model.

The following three test scenarios are chosen to exploit the differences between global and local optimisation approaches. The first scenario is a simple channel with manufactured smoothly undulating bathymetry (see figure 4.20) – this should benefit the gradient-based approaches. The second scenario is a simple channel with flat bathymetry except for a large step in the middle of the turbine area – this divides the turbine area into a 'fast' and a 'slow' flow area, with a sudden transition between. This scenario should favour the global and hybrid approaches, with the local approach not 'seeing' the improvement possible from 'jumping' turbines across the transition. Both these first two scenarios will be demonstrated with arrays of 8 turbines. The final scenario is a larger simple channel with semi-realistic bathymetry, but the array comprises 72 turbines. Optimisation of such a large array should favour the gradient-based approach.

4.7.1 Test 1

The domain for this test scenario measures 1000 by 500 metres and tests the capability of the optimisation approach in situations with localised flow variation.

The lateral boundaries are modelled with a free slip condition and the 'outlet' boundary with a $\eta = 0$ condition. The inflow is velocity driven by a sinusoidal function $u_{\text{in}} = u_{\text{max}} \sin(2\pi t/T)$, where $u_{\text{max}} = 2.0$ ms^{-1}, and $T = 500$ s. This, therefore,

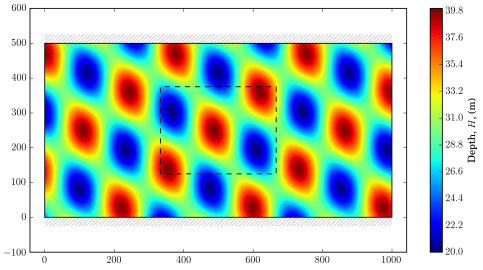


Figure 4.20: Bathymetry for test scenario 1

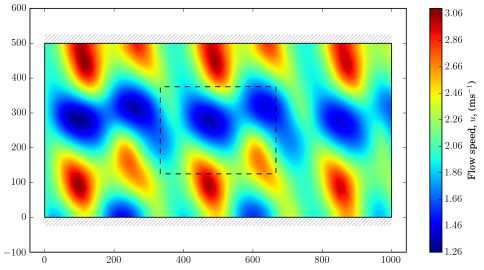


Figure 4.21: Peak positive ambient flow through the test scenario 1 domain.

simulates a (rapid) idealised reversing tide through the channel.

Parameter	**Value**
Water depth	$20 < H < 40$ m
Viscosity coefficient	$\nu = 5$ m^{2}s^{-1}
Acceleration due to gravity	$g = 9.81$ m s^{-2}
Water density	$\rho = 1000$ kg m^{-3}
Bottom friction coefficient	$c_b = 0.0025$
Turbine radii	$r = 10$ m

The ambient flows used for the wake models were snapshots taken at $t = pT/9$ for $0 \leq p \leq 9$. The QoI was evaluated by linearly interpolating between the power evaluated at those times and thus inferring an estimate of the energy extracted over the period.

Results

The optimised layouts for this test scenario are shown in figure 4.24.

The below table shows the power extracted by the optimised array configuration (for all methods the power extracted was evaluated using the shallow water model so as to ensure consistency in the comparison), the number of optimisation iterations and the total optimisation time (as run on a single CPU).

Opt Method	Model	Opt. Energy	Iterations	Time
Genetic	Larsen	21.306 GJ	8686	00:53:22
Gradient-based	Larsen	19.082 GJ	43	0:05:54
Basin-hopping	Larsen	22.674 GJ	287	01:29:14
Gradient-based	S.W.	27.086 GJ	52	06:33:06
Genetic then gradient	Larsen then S.W.	24.084 GJ	8686 then 35	05:44:33

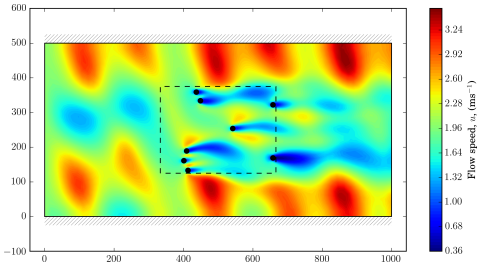


Figure 4.22: Best optimal array layout for scenario 1 in peak 'flood' flow, in this case the best configuration was found using gradient-based optimisation and the shallow water model.

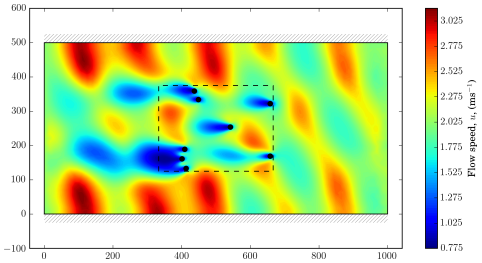
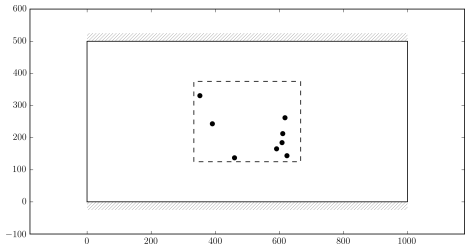


Figure 4.23: Best optimal array layout for scenario 1 in peak 'ebb' flow.

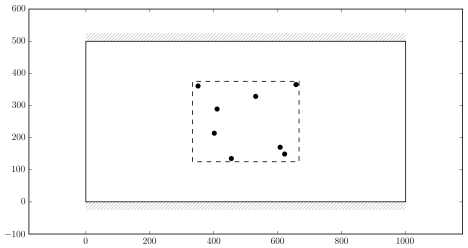
Discussion

As expected, the optimisation methodologies employing the gradient in general fared better than the genetic approach in terms of optimal power production. The (lack of) effect of the turbines on the local flow variations means that the gradient optimisation of the wake model was the least successful – while the layout exploited these variations in the ambient flow, the power was computed using the shallow-water model (for fair comparison between layouts).

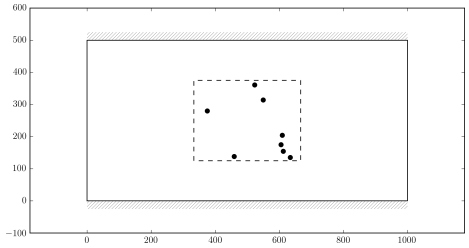
The difference between the genetic-then-gradient approach and the gradient-based shallow water approach highlight the impact of the starting 'guess' on the optimised result – in this particular case the starting guess of 4 rows of two turbines resulted in a better performing optimised layout than the starting guess of the genetic optimised wake model.



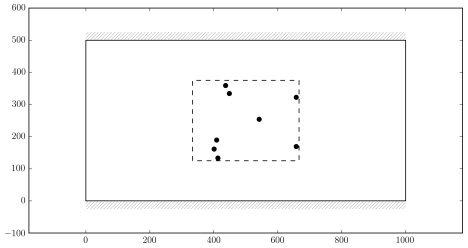
(a) Genetic optimisation of Larsen analytical wake model



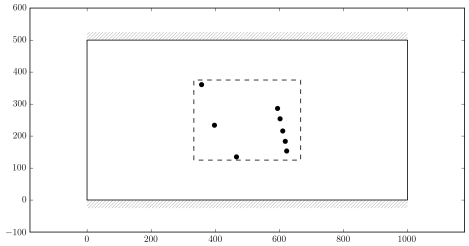
(b) Gradient based optimisation of Larsen analytical wake model



(c) Basin-hopping optimisation of Larsen analytical wake model



(d) Gradient based optimisation of shallow water model



(e) Genetic & Larsen start then gradient & shallow water

Figure 4.24: Comparison of optimised layouts for test scenario 1.

4.7.2 Test 2

The domain, boundary conditions and parameters are the same for test scenario 2 as they were for test scenario 1, except for the bathymetry. Rather than the smooth undulations of the previous example, the bathymetry varies in a step function located in the centre of the domain (and turbine area) giving a shallow western half of the domain at 30 m depth and a deep eastern part at 60 m depth.

Results

The optimised layouts for this test scenario are shown in figure 4.29.

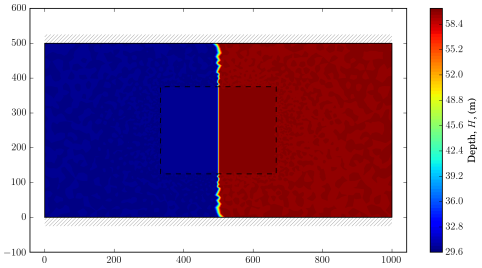


Figure 4.25: Bathymetry for test scenario 2

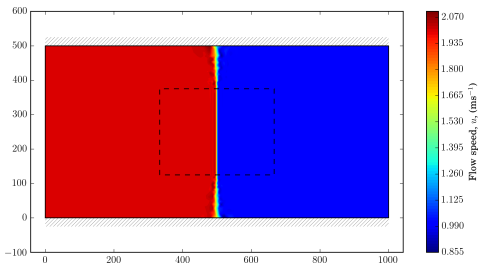


Figure 4.26: Peak positive ambient flow through the test scenario 2 domain.

Opt Method	Model	Opt. Energy	Iterations	Time
Genetic	Larsen	8.989 GJ	4207	00:26:34
Gradient-based	Larsen	4.615 GJ	56	00:06:25
Basin-hopping	Larsen	9.798 GJ	643	01:22:14
Gradient based	S.W.	6.547 GJ	52	05:47:03
Genetic then gradient	Larsen then S.W.	11.142 GJ	4207 then 55	06:48:53

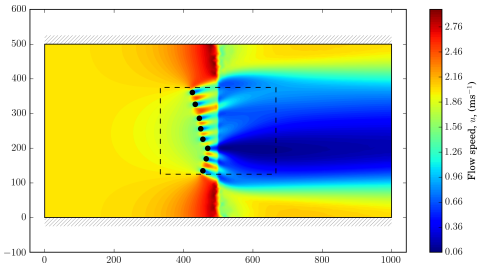


Figure 4.27: Best optimal array layout for scenario 2 in peak 'flood' flow.

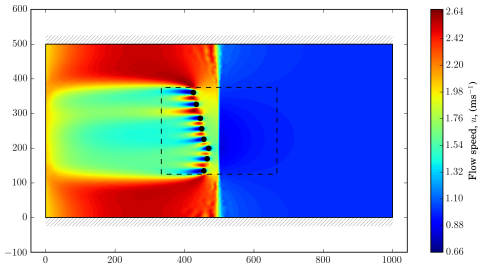
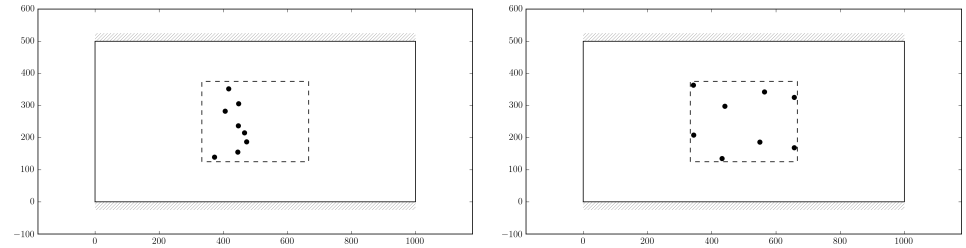


Figure 4.28: Best optimal array layout for scenario 2 in peak 'ebb' flow.

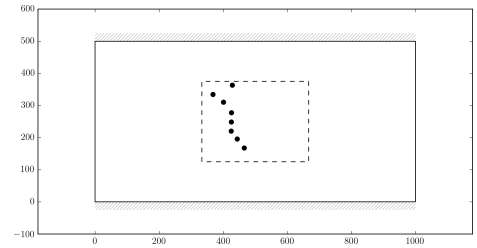
Discussion

This test scenario was chosen to highlight the limitations of gradient-based optimisation – that in non-convex domains, a poor starting guess (starting the optimisation in one 'basin' of the objective) will lead to a poor optimised solution (the solution never escapes that basin). For the gradient-based optimisations, the starting guess was, again, 4 rows of 2 turbines – so the 4 upstream turbines are located in fast flow, while the 4 downstream turbines are located in slow flow. While the two purely gradient-based approaches fail to move all turbines into the fast flowing portion of the domain, all other methods (which have some global element) do achieve this. By far the best performing approach in this case is the genetic-

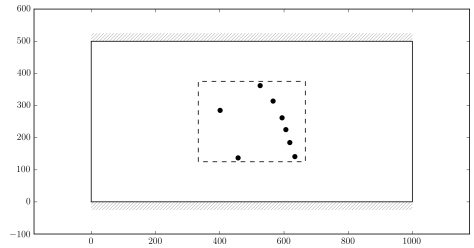


(a) Genetic optimisation of Larsen analytical wake model

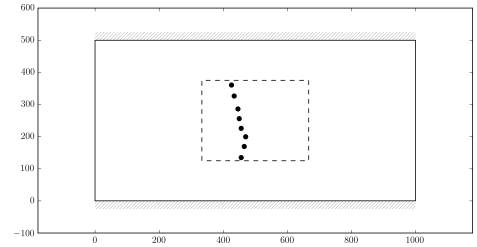
(b) Gradient based optimisation of Larsen analytical wake model



(c) Basin-hopping optimisation of Larsen analytical wake model



(d) Gradient based optimisation of shallow water model



(e) Genetic & Larsen start then gradient & shallow water

Figure 4.29: Comparison of optimised layouts for test scenario 2.

then-gradient approach, the genetic optimisation provides a starting guess with all turbines in the fast flow, for the gradient optimisation which tweaks the turbine locations into an optimal barrage.

In both test cases so far, the basin-hopping approach has performed well, though at substantial additional cost (as compared to the other approaches using solely the Larsen model). The additional per-iteration cost for the gradient-based Larsen and basin-hopping come from the gradient evaluations which are performed by perturbing each turbine in each direction and re-evaluating the flow. This method could be finessed to improve performance [Fang et al., 2011].

4.7.3 Test 3

The final test example only compares shallow-water model gradient-based optimisation with Larsen model global optimisation so as to demonstrate the different scaling abilities (with size of control parameter) of the two approaches. The domain is a scaled-up version of test 1, two times larger in length and width (the test 1 domain was discretised with 9752 vertices and 19 352 cells, compared to 20 115 vertices and 40 070 cells for the larger domain). The bathymetry is the same as test 1 – smoothly undulating between 20 and 40 metres in depth. The other parameters and boundary conditions are as before. In this example the array comprises 72 turbines.

Results

The optimised layouts for this test scenario are shown in figures 4.30 and 4.31.

Opt Method	Model	Opt. Power	Iterations	Time
Genetic	Larsen	116.6 GJ	22434	40:12:03
Gradient based	S.-water	168.1 GJ	102	19:29:37

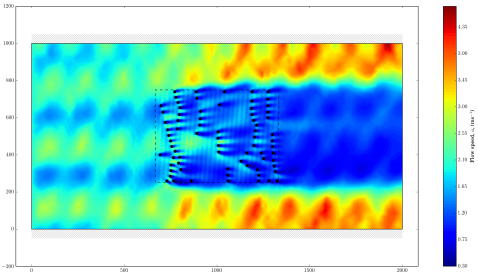


Figure 4.30: Gradient-based optimisation of array represented by the shallow water model in scenario 3 showing array in peak 'flood' flow.

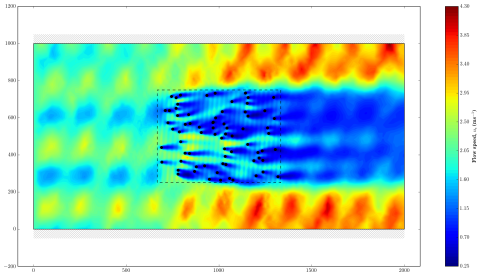


Figure 4.31: Genetic optimisation of array represented with Larsen analytical wake model in scenario 3 showing array in peak 'flood' flow.

4.7.4 Discussion

The purpose of this scenario was to illustrate the poor scaling qualities of genetic optimisation approach as compared to the gradient-based. This is due to two reansons – firstly the search domain increased in size, that is to say there are very many more potential solutions. In addition to this, even though since the ambient

flow has been computed the cost of analysis at each individual iteration is dependent largely on the number of turbines. Thus the large increase in time required to perform the genetic optimisation of the array is almost entirely due to the poor scaling properties of the approach with respect to the number of turbines in the array.

In contrast, for the gradient-based shallow water model optimisation, the larger domain means a higher computational cost at each iteration, and while there has been a 9-fold increase in the number of turbines, the number of iterations has only doubled, and the total time to complete the optimisation has increased less than 4-fold (explained in part by the increased mesh size).

As such, for large arrays the value of a global optimisation strategy must be weighed against the much increased computational cost.

4.8 Conclusions

This work has looked at different approaches to first model, then optimise the design of tidal stream turbine arrays. The goal was to find an efficient method for determining the array micro-siting design which maximises the power extracted by the turbines.

Much of the body of previous work on tidal turbine array optimisation has limited the control parameters, e.g. to the longitudinal and transverse spacing between rows of (sometimes staggered) turbines. However this is inherently inflexible to exploiting the local variations in flow that occur as a result of non-flat bathymetry, or irregular coastal geometry. Consequently for this work, array micro-siting optimisation refers to methods based on control vectors containing the two-dimensional Cartesian coordinates of each turbine comprising the array.

Two different optimisation methodologies were considered alongside two hybrid approaches. Global optimisation has the benefit of exploring the optimisation space well, however it requires large numbers of iterations and the product of the optimisation has no guarantees of optimality. In contrast gradient-based methods require fewer iterations and the product of the optimisation is guaranteed to be (locally) optimal. However the optimum that is found is sensitive to the starting parameters guess. Basin-hopping is a hybrid approach which randomly perturbs the design vector once it has converged from gradient based optimisation – thus testing several local optima. Finally, global optimisation used to provide a starting

guess to a gradient-based optimisation was the other hybrid approach employed.

The large numbers of iterations required by the global optimisation algorithms requires that the analysis code used to evaluate sets of design parameters be computationally inexpensive so that the optimisation problem does not become intractable. Consequently a perturbation approach was presented which used analytically calculated wake deficits superposed onto the ambient flow through the domain as pre-computed by a suitable flow solver. Three such analytical models were presented, calibrated, analysed and the best one selected for use in optimisation.

Conversely the fewer number of iterations required by gradient-based methods meant that the analysis code used with such methodologies could be proportionally more computationally expensive. This is beneficial on the basis that more expensive models should be of higher fidelity and generally more realistic. Indeed, this approach enabled the control parameter at each iteration to be evaluated using a shallow-water model which meant that the hydrodynamic coupling between turbines and flow was not neglected as in the perturbation approach. This meant that local blockage effects (such as acceleration of flow around the sides of the turbine) and global blockage effects (such as bulk flow redirection) where captured.

A suite of three heavily idealised test scenarios were devised to test the strengths and weaknesses of the optimisation approaches under consideration. These approaches were:

- Global, gradient-based and basin-hopping optimisation used with the Larsen analytical wake model.
- Gradient-based optimisation used with the shallow-water model
- An approach which used gradient-based optimisation with the shallow-water model started from the product of a global optimisation with the Larsen analytical wake model.

These demonstrated the following:

- Where there is significant local flow variation, row spacing adjustment is unlikely to produce the best design.
- Where blockage is likely to be a dominant mechanism, a coupled model is required to capture the flow redirection and correctly estimate power.

- Gradient-based optimisation approaches perform well in generally convex domains without sharp bathymetric features.
- A computationally cheap genetic optimisation using an analytical wake model can be used to provide a good starting guess for gradient-based optimisation in non-convex domains. This helps to address the limitations arising from the local nature of the gradient-based optimisation
- For large arrays of turbines, genetic optimisation becomes impractical – and takes longer than gradient-based optimisation with a fully coupled shallow-water model.

A clear direction for further work would be to more comprehensively compare the other analytical wake models and combination models within the optimisation framework. Furthermore, background turbulence intensity can be parametrised within the Ainslie model framework – which may have an important impact on the placement of downstream turbines and therefore potentially alter the designs produced by optimisation approaches using this as an analysis code.

Chapter 5

Integration of cost modelling within the micro-siting design optimisation

Contents

5.1 Introduction 112
5.2 Problem Formulation 113
5.3 OpenTidalFarm 115
5.4 Cable length minimisation 116
5.5 Tidal Array Routing 123
5.6 Integrating the micro-siting and cable-routing optimisations . 127
5.7 Test cases 130
5.8 Conclusions and summary 136

Abstract

This chapter is largely based upon Culley et al. [2016a], D.M. Culley, S.W. Funke, S.C. Kramer and M.D. Piggott (2016). Integration of cost modelling within the micro-siting design optimisation of tidal turbine arrays. *Renewable Energy*, 85, 215-227.

The location of individual turbines within a tidal current turbine array – micro-siting – can have a significant impact on the power that the array may extract

from the flow. Due to the infancy of the industry and the challenges of exploiting the resource, the economic costs of realising industrial scale tidal current energy projects are significant and should be considered as one of the key drivers of array design. This chapter proposes a framework for the automated design of tidal current turbine arrays in which costs over the lifespan of the array may be modelled and considered as part of the design optimisation process. To demonstrate this approach, the cost of sub-sea cabling is incorporated by implementing a cable-routing algorithm alongside an existing gradient-based array optimisation algorithm. Three idealised test scenarios are used to demonstrate the effects of a financial-return optimising design approach as contrasted with a power maximisation approach.

5.1 Introduction

Tidal current turbines are devices which convert the momentum of tidally induced ocean currents into electricity. Much as with wind power, several individual turbines may be formed into an array to yield power on an industrial scale. Determining the optimum location of turbines within an array – micro-siting – is an issue of critical importance, on which the viability of the project may hinge. Rapid spatial variations in current flow speed can be caused by complex bathymetry and the presence of the turbines themselves. Since the power extraction of a turbine is dependent upon the cube of that flow speed (i.e. is highly sensitive to it), optimisation of the micro-siting design is a complex and challenging problem. Funke et al. [2014] have demonstrated in test scenarios that such micro-siting optimisation has the potential to increase the power extracted by an array by 33 % as compared to an array of the same number of turbines arranged in a rectangular grid.

Just as power production is dependent on the turbine micro-siting design, so too may certain *costs* be dependent on turbine location. These location-dependent costs, such as the cost of cabling, water depth or difficulty of installation, may have just as great an impact on the project viability and must be considered in the design process. As noted by Thomson et al. [2011], there has been a general focus across the renewables industry on design optimisation based solely on energy yield. The goal of this chapter is to develop a more holistic approach to array design which balances both energy yield and cost, and thus enables array developers to maximise their overall return on investment.

For this work, sub-marine cabling costs have been chosen as an example location-based cost. In offshore wind projects, connection costs typically represent 18 – 20% of the capital cost of a project [Lumbreras and Ramos, 2011, Quinonez-Varela et al., 2007]. Research into the cost of elements of tidal current power installations has suggested that these costs may be similarly significant, and the largest cost affected by micro-siting design [Allan et al., 2011]. Negative wake interactions between turbines may motivate the spacing of turbines – where lease area allows – far apart from each other and conditions may result in higher flow velocities out in the main channel away from land. However, as spacing between turbines and distance from shore increases, so too does the cost of the cabling required to connect the turbines to the power grid. Thus the requirement to minimise the cost associated with increased cable length will likely be a competing factor to maximising power output.

The contribution of this chapter is a novel method of optimising turbine micro-siting design, which integrates evaluation and consideration of both costs and benefits. Cable length minimisation and power-extraction maximisation are integrated and optimised using an efficient gradient-based algorithm. This approach requires fewer iterations than alternate (for example genetic algorithm) approaches meaning that a more computationally expensive (and therefore more accurate / realistic) prediction model may be used.

In the following section, the array micro-siting design optimisation problem is formulated mathematically, and it is shown how the financial return is defined as a function of incomes – such as power extracted by the array, and costs – such as the length of cabling required to connect the turbines. In section 5.4 the cable routing problem is outlined, and previous work in tackling it is explored. The problem is mathematically formulated and it is shown that the proposed integration into a gradient-based framework is valid. In section 5.5, a model is developed for this application leading to section 5.6 in which the challenges of integrating the *OpenTidalFarm* and cable routing models are examined. Finally the approach is demonstrated on idealised test-cases in section 5.7.

5.2 Problem Formulation

The overall goal of this work is to optimise the financial return over the array lifespan, R_{fin}, for the developer of an array of N tidal turbines located within a bounded site – the 'turbine area'.

The domain is modelled in two-dimensional Cartesian space and the coordinates of the turbine locations are, as before, encoded in a $2N$-long vector, $\mathbf{m}$, where

$$\mathbf{m} = (x_1, y_1, x_2, y_2, \ldots, x_N, y_N)^T.$$

R_{fin} is considered to be a function of $\mathbf{m}$ and is therefore maximised through adjusting the turbine locations,

$$\max_{\mathbf{m}} \quad R_{\text{fin}}(\mathbf{m}). \tag{5.1}$$

In order to obtain a framework through which cost and income models can be integrated in a modular fashion, R_{fin} is expressed as the sum of income functions and cost functions. For example (5.1) may be expressed as

$$\max_{\mathbf{m}} \quad I_P(P(\mathbf{m})) - C_C(L(\mathbf{m})), \tag{5.2}$$

where $P(\mathbf{m})$ is the power extracted from the flow by the turbines, $I_P : \mathbb{R} \to \mathbb{R}$ maps power output to financial income, $L(\mathbf{m})$ is the total length of sub-sea cabling required to connect the turbines to base stations on the shore, and $C_C : \mathbb{R} \to \mathbb{R}$ maps cable length to financial cost. Both P and L can implicitly be written as functions of $\mathbf{m}$ when it is understood that for a given array configuration ($\mathbf{m}$), the evaluation of P will involve the solution of a problem describing the tidal dynamics and evaluation of L involves the solution of a routing optimisation problem. Further details on the functions I_P and C_C used in (5.2) may be found in section 5.6.

Models which evaluate a physical quantity, such as power extracted by the array, or the length of sub-sea cabling required to service it, can thus be included in the composition of (5.1) if functions can be identified which map those physical quantities into a financial dimension.

In a general setting, there might be additional location-based cost functionals, such as installation depth, which could be incorporated in a similar fashion as the cable cost, but these are not considered here. For this work, these (and all other costs) are assumed to be constant.

5.3 OpenTidalFarm

We recall from (3.15) the default optimisation problem which *OpenTidalFarm* is built to tackle;

$$\begin{aligned} \max_{\mathbf{m}} \quad & P(\mathbf{m}) \\ \text{subject to} \quad & \mathbf{b}_l \leq \mathbf{m} \leq \mathbf{b}_u \\ & g(\mathbf{m}) \leq 0. \end{aligned} \tag{5.3}$$

Once again, the bounds $\mathbf{b}_l \leq \mathbf{m} \leq \mathbf{b}_u$ constrain the turbines to the turbine area (in this case rectangular in shape, for simplicity) and the inequality constraint, $g(\mathbf{m}) \leq 0$, implements a minimum distance spacing constraint between adjacent turbines. By default, the performance of the turbine layout is evaluated as the power extracted by the turbines,

$$P(\mathbf{m}) = \int_\Omega \rho c_t(\mathbf{m}) ||\mathbf{u}||^3 \, \mathrm{d}x. \tag{5.4}$$

Note in this chapter the use of $c_t(\mathbf{m})$ rather than the more realistically parametrised $c_t(\mathbf{u}, \mathbf{m})$ which I developed after this work was completed.

In this chapter we consider only the steady-state shallow water model, however this can be generalised to the unsteady case. As such, energy yields quoted in this chapter are instantaneous values. This is clearly a weakness in the methodology, however it is useful in two respects. Firstly, using the steady-state shallow water model, proof of concept of integration with the cable routing is demonstrated at a much reduced computational cost (as compared to using the unsteady model. Secondly, with flow coming from just one direction, it is much more easy to interpret the arrangement of the turbines and intuitively grasp how the optimal turbine arrangement has developed in response to the simplified flow.

A key advantage of this approach is the use of the adjoint equations, since it enables the gradient $\mathrm{d}P/\mathrm{d}\mathbf{m}$ to be determined at a computational cost that is independent of the number of turbines, N. (Technically, the *adjoint* is computed at an independent cost, while the gradient requires matrix multiplication of the adjoint with another matrix, both of whose size varies with N – however, this cost is negligible for the sizes of N that are realistically under consideration). The accessibility of the gradient facilitates the use of a gradient-based optimisation algorithm. The iteration numbers for these algorithms are also observed to be independent

of the number of turbines [Funke et al., 2014], and very small compared to genetic algorithms [Barnett et al., 2014]. Therefore optimisation of even a large number of turbines is feasible. Clearly the benefit of this approach is the scalability with number of turbines. In addition, since so few iterations are required, each iteration can be more computationally expensive. This means that a more realistic flow model may be used to capture the complex hydrodynamics more accurately.

The computational economy of the gradient-based optimisation is at the heart of the *OpenTidalFarm* approach. Therefore any additional (cost) model to be integrated into this framework must be able to interface with a gradient based approach and be reasonably efficient to evaluate. More precisely it must be a differentiable map $\mathbf{m} \in \mathbb{R}^{2n} \rightarrow \mathbb{R}$, which characterises the cost of the array with regard to the turbine positions in terms of a single scalar value, *and* have the gradient efficiently computable. To demonstrate how this may be achieved even for complicated cost models, sub-sea cabling costs have been selected for this chapter. Integration of cable cost presents a particular challenge, as determining the length of the cable route required for a given $\mathbf{m}$ is, itself, an optimisation problem as described in the following section. Unlike micro-siting design, which is a continuous optimisation, the cable routing is an integer problem. This presents an interesting challenge in integrating an integer problem within a gradient-based optimisation framework.

5.4 Cable length minimisation

5.4.1 Problem Formulation

The electrical power generated in each turbine must be gathered together at a collection point where it can be properly conditioned and delivered to the national grid. Power is carried by submarine cables which run along, or are trenched into, the sea bed. Turbines may be connected in series, however the cables have a finite power capacity and as such there is a limit to the number of turbines that may be linked by a single cable.

In early deployments, it is likely that individual turbines will be connected directly to the collection point. However, as array sizes grow, this method of connection would become too expensive and tidal array designers will likely follow the wind industry in connecting turbines in more efficient ways – for example in series as described above. Direct connection of each turbine to the collection point

is a special case of the methodology described above, where the cable capacity is simply 1 turbine. So this formulation is robust both for small, early turbine arrays and the larger more intricately connected arrays to come.

The cable routing problem is summarised as: Given the position of N turbines, find a cable routing with minimal cable cost which connects each turbine to a collection point, under the conditions that each turbine supports a maximum of two connections (one cable in and one cable out) and that each cable may connect a limited number of turbines as defined by the cable capacity.

This problem may be graphed as a set of vertices (turbines) and edges (cables) and formulated as a combinatorial problem consisting of two coupled optimisations to minimise the length of the cable network. The first optimisation is akin to a bin-packing problem, in which the complete set of vertices must be partitioned into subsets of the turbines which will be connected in series by a single cable. As such, the size of each subset is limited by the cable capacity. The second optimisation is to determine the order in which the vertices within the subset are linked so as to minimise the route length. The overall objective is to minimise the length of the union of the subset routes.

5.4.2 Cable routing optimisation in literature

A few papers have considered the optimisation of cable routing design in offshore *wind* turbine arrays which shares many design concerns with cable routing for tidal arrays [Lumbreras and Ramos, 2013a, 2011, Bauer and Lysgaard, 2014, González-Longatt et al., 2012, Lumbreras and Ramos, 2013b].

Lumbreras and Ramos [2013a] developed a model employing mixed integer programming in order to capture stochasticity in both wind resource and component failure rates in an optimisation of electrical layout. The model takes the position and characteristics of the wind turbines (as found by a separate micro-siting model) and determines the cables (routing and type), transformers and converter stations. In a case study comparison, existing offshore wind farm electrical layout designs were compared to those produced by the model. This demonstrated that cable routing optimisation can result in significant capital savings on a project, together with reduced transmission losses and increased redundancy against component failure.

In a similar vein, González-Longatt et al. [2012] developed a cost model cap-

turing the costs of collection points and cabling, and used a genetic algorithm to find a minimum cost routing solution. As with Lumbreras and Ramos [2013a], the locations of the turbines were determined from a prior micro-siting assessment.

Bauer and Lysgaard [2014] identify the utility of coupling micro-siting and cable routing design. They consider offshore wind farms and identify TopFarm [Larsen and Réthoré, 2013] as the most sophisticated array optimisation tool and indicative of many such tools – in that it requires in the order of thousands of array configuration evaluations. The authors propose a strategy in which a quick cable routing estimation is made at each micro-siting optimisation, with a more developed optimisation carried out once the turbine sites are determined.

One of the key differences between the electrical layout design for offshore wind farms as opposed to tidal farms is that the power collection point (at which the cable routes terminate) may be located onshore, rather than within the array (as with offshore wind). This is because for nearshore installations it may be cheaper to run the extra cabling and have the electrical collection equipment on dry land. This, however, means that the first connection in every route (from the shore-side collection point to the first turbine in the series) is the most expensive – as it has to run from shore to the array. As such there is an inherent 'start-up' cost for each route in a routing, and the optimal solution must exploit the cable capacity to have the minimum number of routes possible.

Having reproduced the algorithm proposed by Bauer and Lysgaard [2014] and tested it upon typical array formations produced by *OpenTidalFarm* with an onshore collection point it was found to produce routings with more than the minimum number of routes and which therefore were clearly sub-optimal. This is because the heuristic was designed to work very rapidly – which it achieves by working greedily. As noted, the cable routing problem is a coupled set-partitioning and vertex-ordering problem and while heuristics can work well on either one aspect, they are generally poor at dealing with both aspects simultaneously.

5.4.3 Approaches to solving the cable-routing problem

The routing problem is related to the travelling salesman problem (TSP) to find the shortest route for a salesman to tour a number of cities (N), visiting each exactly once and returning to where they started. Here, the 'cities' represent the turbines and the salesman's route represents the cable. This is an optimisation problem which has been intensively studied in mathematical literature since its

conception in 1930. The multiple TSP (mTSP) is a version of the problem in which several salesmen must tour their own individual set of cities. This variant is closest to the cable routing problem since the capacity constraints of the cable require that overall array routings be composed of several individual routes. Solution approaches have been proposed from a wide range of operational research problems, such as scheduling and mission planning [Bektas, 2006], and for each application the method is tailored to suit the purpose.

Hence, there are a wide variety of approaches adopted in literature to solving the TSP and its variants. A 'brute force' attempt to find a solution (i.e. trying every permutation) would have a run time falling within a polynomial factor of $\mathcal{O}(N!)$ [Bellman, 1962]. As such, this becomes infeasible for $N \gtrsim 15$. Practical approaches to finding a solution can be split into heuristics and meta-heuristics. Heuristics tend to be specifically designed for a certain problem, often work greedily and have little or no ability to revise solutions. Meta-heuristics go a step further; they tend to be more generalised and use simple heuristics to produce multiple solutions and often have the ability to accept poorer solutions as part of a mechanism to better explore the whole search space.

The majority of classic heuristic methods suffer from the limitation that capacity constraints do not play a central decision role in the algorithm, but rather are enforced as a check and limitation at each iteration. While this greatly simplifies the formulation of the problem, it generally results in the set of vertices being partitioned poorly, or into more routes than the minimum necessary. While this is acceptable when the collection point is located within or close to the array, as discussed previously in the context of the Bauer and Lysgaard [2014] algorithm, if the distance from the collection point increases, the first edge in the route becomes more expensive than other edges connecting turbines – so there is an inherent 'start-up' cost for a route. As such, minimising the number of routes (the 'set partitioning' side of the optimisation) grows in importance compared to optimising each individual route (the TSP side of the optimisation). In near-shore tidal array developments it will likely be more economical to build the collection point on-shore, even at the expense of the additional cabling. As such it is important that any algorithm employed for the tidal scenario will still produce a good cable routing in this situation.

Meta-heuristic methods come in many forms, for example ant-colony, simulated annealing and genetic algorithms. The main benefits of these approaches is the ability to comprehensively search the solution space and have design con-

straints (such as cable capacity) codified as an integral part of the solution generation – rather than them being enforced as a limiting check with each iteration. In the case of the genetic algorithm, only solutions that comply with the cable capacity are proposed, so this constraint is incorporated from the start – as opposed to heuristic methods in which a check has to be carried out with each iteration. This means that, as problems become more constrained, the solution space shrinks and the algorithm works more efficiently.

5.4.4 Notation

We formulate the cable routing problem on a graph of vertices and edges. The vertices, V represent the individual turbines and the power collection point, and are referred to by the indices i and j ($i, j \in V$). The sub-sea cables are represented by undirected edges where $E = \{(i, j) \mid i, j \in V,\ i < j\}$ is the set of all possible edges. One cable connecting a set of turbines is called a route, $r \subset E$. The number of cables connected along an individual route is limited by the cable capacity, C_{ap}, so $|r| \leq C_{ap}$. Here all the turbines are assumed to have the same power rating, although this could be generalised. A routing $R = r_1 \cup r_2 \cup \cdots \cup r_k$ is a collection of individual routes (here k) and defines a possible cable network. Since each cable may only be used once, we require that $r_a \cap r_b = \emptyset \ \ \forall\ 0 \leq a < b \leq k$.

The cable length L of a routing R with turbine locations m is

$$L_R(\mathbf{m}) = \sum_{(i,j)\in R} l_{ij}(\mathbf{m}), \tag{5.5}$$

where l_{ij} is the Euclidean distance between vertices i and j

$$l_{ij}(\mathbf{m}) = \sqrt{(\mathbf{m}_{xi} - \mathbf{m}_{xj})^2 + (\mathbf{m}_{yi} - \mathbf{m}_{yj})^2}, \tag{5.6}$$

and $\mathbf{m}_{xi}$ corresponds to the x coordinate of the ith vertex in $\mathbf{m}$, etc.

Given fixed turbine locations, there exist a finite, but possibly large number of routings which satisfy the above constraints. Define $\mathscr{R}$ to be the complete set of these. The solution to the cable routing problem is thus

$$L_{\text{opt}}(\mathbf{m}) = \min_{R\in\mathscr{R}} L_R(\mathbf{m}). \tag{5.7}$$

5.4.5 Use of a gradient-based optimisation approach with the cable-routing

The numerical solution of the cable routing problem (5.7) will be discussed in section 5.5. Recall however, that L_{opt} is part of the optimisation problem (5.2), which is to be solved by a gradient-based optimisation algorithm. Most gradient-based optimisation algorithms assume that the functional is at least once differentiable. So the natural question arises whether the minimal cable length is a continuous and differentiable function of the turbine locations.

Unfortunately, it can easily be shown that the gradient $\frac{dL_{\text{opt}}}{d\mathbf{m}}$ does not always exist. To illustrate this in one dimension, consider figure 5.1. If turbine number 11 (T11) is moved in the x direction, at some point, the optimal routing will be described with turbine 11 in route 3 (r_3) rather than route 2 (r_2). The effect of this movement of turbine 11 on L is shown in figure 5.2, both for the turbine being held in route 3 and route 2. The intersection shows the point at which the optimum routing will change to maintain an optimum route length. Here the minimal cable length becomes non-smooth.

Instead, we will only be able to show that L_{opt} is Lipschitz continuous. Rademacher's theorem then proves that L_{opt} is differentiable almost everywhere. Further it can be shown that the differentiable pieces have a bounded differential. Under the assumption that the optimisation algorithm never hits a non-differentiable point, we can apply a standard quasi-Newton algorithm to solve this non-smooth problem. Some work has been done looking at the effects of applying a gradient-based optimisation approach to a non-smooth function [Zhang et al., 2000, Lewis and Overton, 2009, Skajaa, 2010]. Zhang et al. [2000] use the L-BFGS algorithm (the algorithm at the heart of SLSQP) to find the minimum of a discontinuous sawtooth function, with some success – dependent upon the size of the discontinuities. Lewis and Overton [2009] and Skajaa [2010] observe that the inexact line search performed by optimisation algorithms will almost certainly miss points at which the function is non-differentiable and hence demonstrated successful implementations of gradient-based optimisation of non-smooth functions.

Theorem 1. *The minimal cable length $L_{opt}: \mathbb{R}^{2N} \to \mathbb{R}$ is Lipschitz continuous and differentiable almost everywhere.*

Proof. First, note that $L_R(\mathbf{m})$ is Lipschitz for a fixed R. This follows directly from the definitions (5.5), (5.6). Next observe that the minimum of two Lipschitz func-

tions f and g is also Lipschitz. This can be seen from the relationship

$$\min(f,g) = \frac{f+g}{2} - \frac{|f-g|}{2}.$$

Iteratively applying this rule over the finite set $\mathscr{R}(\mathbf{m})$ yields that L_{opt} is Lipschitz. The differentiability almost everywhere follows from Rademacher's theorem. □

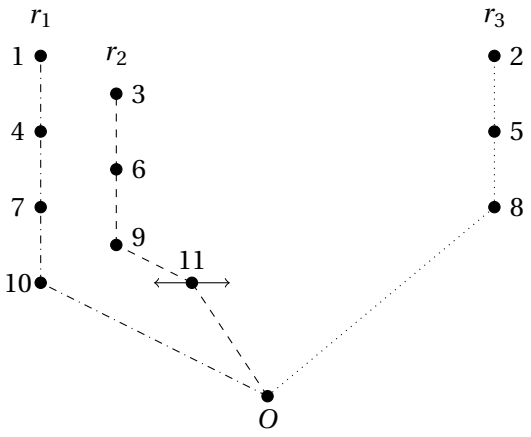


Figure 5.1: If the turbine layout from figure 5.3 is taken, and turbine 11 is moved in the x direction, as shown by the arrow, when the turbine is positioned as shown, the optimal routing will be described with turbine 11 in r_2.

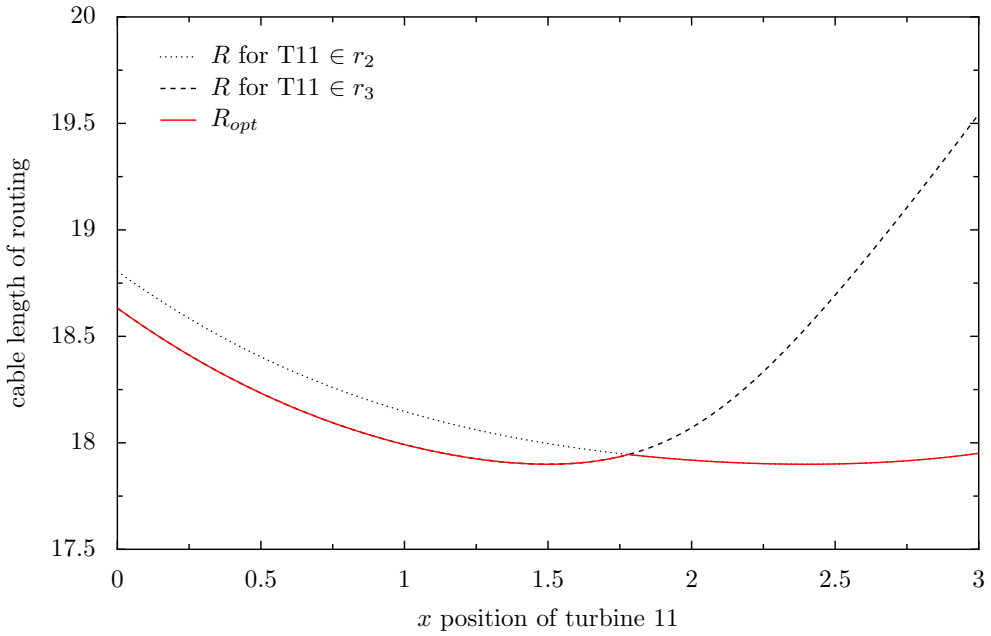


Figure 5.2: Plot showing the routing length for two routings as turbine 11 (T11) in figure 5.1 traverses in the x direction.

5.5 Tidal Array Routing

Various meta-heuristics were trialled for solving the cable routing problem (5.7). A genetic algorithm was finally selected as, in initial experimentation, it was found to be the most robust. Furthermore genetic algorithms are the most popular approach used in operational research problems and therefore the most developed. Genetic algorithms (GAs) are search meta-heuristics which simulate the process of evolution. Recall from previous chapters that possible solutions are encoded as numeric vectors termed 'chromosomes'. An initial population of chromosomes is generated, either randomly or using a pseudo-random heuristic. This population is evaluated using a fitness function and the fittest chromosomes are selected and mutated to populate the next generation. The pseudo-random nature of the mutation ensures that the search space is effectively explored, while a good chromosome representation minimises this search space through eliminating repeated solutions and thereby reducing redundancy [Carter and Ragsdale, 2006]. GAs are used on a wide variety of problems and are adapted to purpose by designing suitable chromosomes and mutations. The particular GA being developed for the application considered in this work shall be referred to as *TidalArrayRouting* (TAR) to distinguish design choices specific to this algorithm for this application from GAs in general.

5.5.1 Chromosome selection

The goal in selecting a chromosome representation is to use a numeric vector which minimises the potential for the same solution to be represented in different ways (thus undesirably increasing the size of the search space) but is also easy and computationally efficient to manipulate in order to mutate. A comprehensive review of the available approaches for the mTSP is given in [Carter and Ragsdale, 2006] which compares the one chromosome technique [Tang et al., 2000] and the two chromosome technique [Malmborg, 1996] with a novel two-part chromosome technique.

Consider figure 5.3. For this example the two-part chromosome technique would represent this routing as

$$[\underbrace{10\,,\,7\,,\,4\,,\,1\,,\,9\,,\,6\,,\,3\,,\,11\,,\,8\,,\,5\,,\,2}_{\text{order of turbines visited}}\,,\,\underbrace{4\,,\,7}_{\text{breaks}}].$$

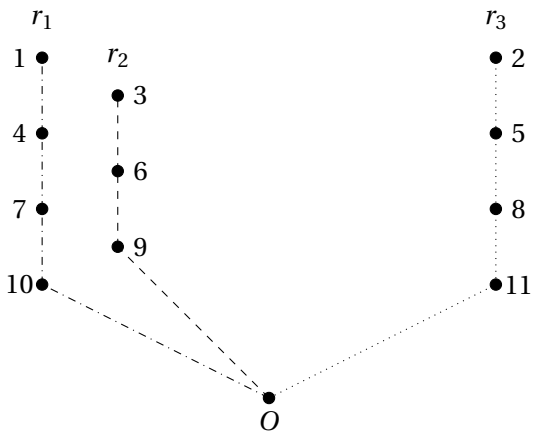


Figure 5.3: Schematic of an idealised routing, $R = r_1 \cup r_2 \cup r_3$ with capacity constraint $C_{ap} = 4$

The first part of the chromosome are the N turbines in the array, the second part are the breaks – which denote the index of the first turbine in a new route and thus split the routing into n_{rts} individual routes. Compared to other chromosome encodings this has very few redundant solutions and the size of the solution space is shown to be

$$N! \binom{N-1}{n_{\text{rts}}-1}. \tag{5.8}$$

[Carter and Ragsdale, 2006]

TAR is initialised with a population, $\mathscr{P}_0$, comprising p routings, each represented as two-part chromosomes, where p is the population size and is a multiple of eight (one plus the number of mutation operators; of which TAR has seven). The population is then randomly subdivided into groups of eight and the length of each route is evaluated. The route with the shortest length in each group is preserved and the remaining seven routes are replaced by mutated versions of the survivor. A new population, $\mathscr{P}_i$, is assembled from the survivors and the mutant chromosomes. By design, all mutations produce valid routings (i.e. routings which satisfies the constraints). If the shortest route does not change for a predefined number of consecutive iterations, it is returned as the optimal routing.

The mutation operations involve shuffling either the first part, the second part or both parts of the chromosome, in the hope of improving its performance. This is done naively so as to keep the computational expense of this operation minimal. For example, the turbines at randomly chosen indices i and j may be swapped

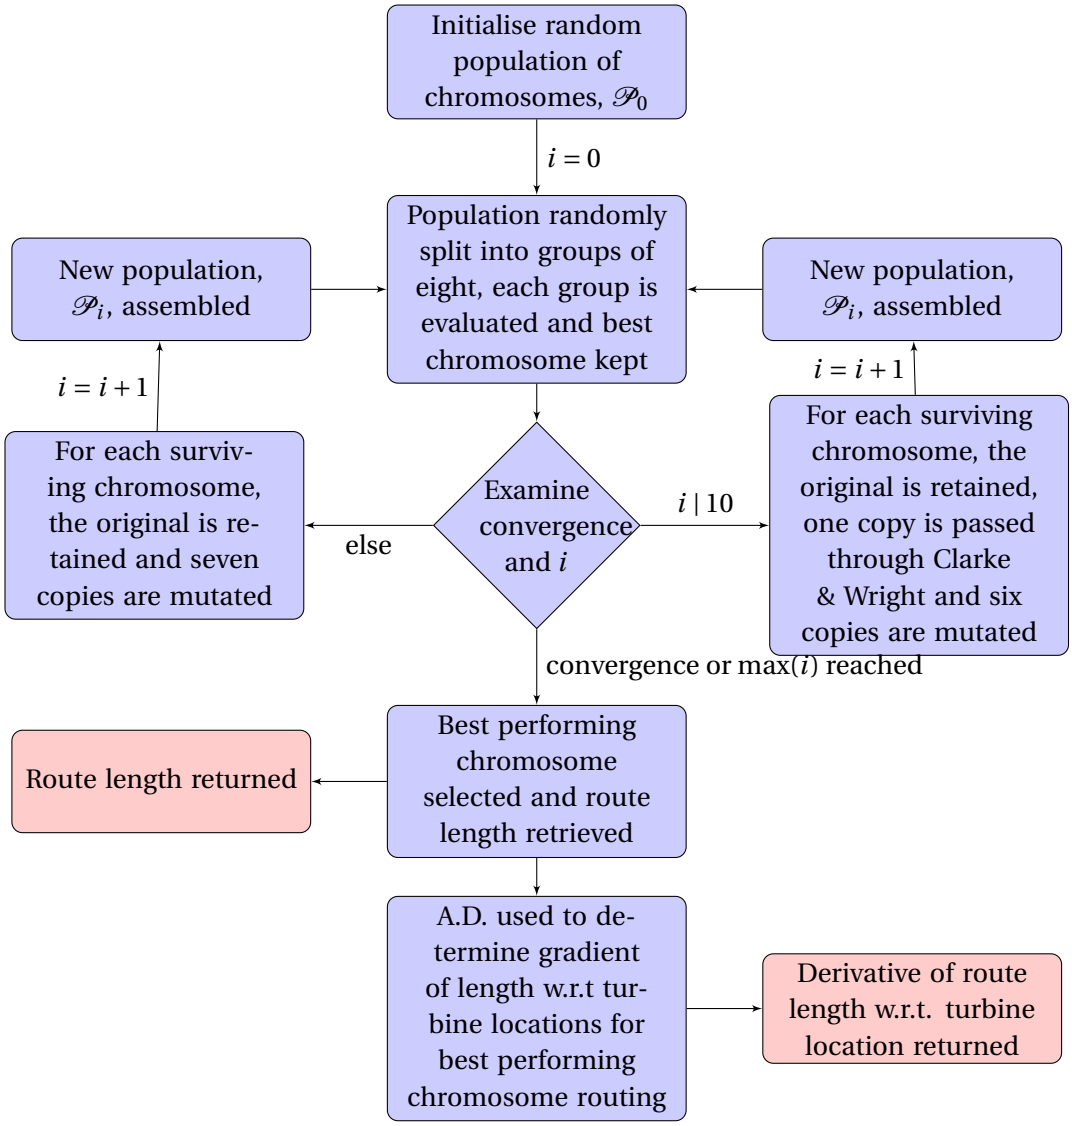


Figure 5.4: Schematic of *TidalArrayRouting* developed for *OpenTidalFarm*

with each other, or the section of the chromosome between i and j may be reversed. To improve convergence, a more complex heuristic is occasionally used in place of one of the naive mutations. This heuristic is computationally more expensive than the naive mutation and experimentation has informed the decision to use it only on every tenth iteration of TAR. This gives a good balance for a wide range of N and population sizes between convergence at fewer iterations and the computational expense of those iterations.

The heuristic used is a modification of the Clarke and Wright greedy algorithm [Clarke and Wright, 1964]. This connects every turbine directly to the collector,

then calculates the saving of removing a given turbine's connection to the collector and instead connecting it to a neighbour. It creates a list of the saving for each turbine being connected to each other turbine rather than the collector, and arranges this list in order of decreasing saving. Starting at the top of the list, the algorithm determines whether the new connection is valid (e.g. will not exceed the cable capacity), and if so it makes the change. This algorithm was the basis of the approach used by Bauer and Lysgaard [2014] and as has been discussed, such an approach is limited in itself as it does not perform the set partitioning function well. Consequently this responsibility is taken away, and the heuristic is applied separately to each route within the chromosome which is passed to it. Since it does not act naively like the mutation operators, it often makes an improvement, aiding convergence of the overall TAR model and justifying its additional computational expense.

Once the algorithm has found a routing (either because it has converged to the specified tolerance or because it has reached the maximum number of iterations), the gradient of the length of that routing with respect to the turbine positions is computed. Since the length is simply a sum of Euclidean norms, finding this differential is trivial.

The gradient implementation of L_{opt} can be rigorously verified by employing a Taylor remainder test. Given $L_{\text{opt}}(\mathbf{m})$ and a perturbation vector $\delta\mathbf{m}$, then the first-order Taylor remainder should tend to zero linearly with the decreasing perturbation, that is

$$|L_{\text{opt}}(\mathbf{m} + h\delta\mathbf{m}) - L_{\text{opt}}(\mathbf{m})| \to 0 \quad \text{at} \quad \mathcal{O}(h). \tag{5.9}$$

Similarly, the second-order remainder should converge quadratically in h,

$$|L_{\text{opt}}(\mathbf{m} + h\delta\mathbf{m}) - L_{\text{opt}}(\mathbf{m}) - h\nabla L_{\text{opt}} \cdot \delta(\mathbf{m})| \to 0 \quad \text{at} \quad \mathcal{O}(h^2) \tag{5.10}$$

where ∇L_{opt} is the gradient of L_{opt} with respect to $\mathbf{m}$. In other words, if h is halved, the second-order remainder should reduce by a factor of 4. These properties were found to hold for the cable length.

5.6 Integrating the micro-siting and cable-routing optimisations

Integrating the cable-routing kernel into *OpenTidalFarm* (see figure 3.6) requires that for every functional evaluation, the cable-routing is optimised and the shortest route length found. Additionally, for every gradient evaluation, the effect of the turbine positions on the route length ($\mathrm{d}L/\mathrm{d}\mathbf{m}$) must also be calculated. Thus we have a nested optimisation with an inner cable-routing optimisation for each iteration of the micro-siting optimisation.

In order to achieve dimensional parity, a financial approach seems natural and as such we recall from (5.2) that the financial return, R_{fin} is defined as

$$R_{\text{fin}} = I_P(P(\mathbf{m})) - C_C(L(\mathbf{m})), \tag{5.11}$$

and consequently,

$$\frac{\mathrm{d}R_{\text{fin}}}{\mathrm{d}\mathbf{m}} = I_P\left(\frac{\mathrm{d}P}{\mathrm{d}\mathbf{m}}\right) - C_C\left(\frac{\mathrm{d}L}{\mathrm{d}\mathbf{m}}\right). \tag{5.12}$$

As such, we require a definition of an income function $I_P : \mathbb{R} \to \mathbb{R}$ determining the 'income' from the power extracted from the flow and a cost function $C_C : \mathbb{R} \to \mathbb{R}$ determining the 'cost' of cable. A commercial array designer would have their own cost information based upon the cost and performance of the equipment to be installed. Unfortunately the proprietorial nature of such information means that it is unobtainable for academic publication. Consequently, for this work, I_P and C_C will be estimated based upon available information and assumption.

For the purposes of this work, a present value (PV) approach will be used to adjust for the time value of money. The essence of this approach is to recognise that there is an opportunity cost to investing cash in a project with a payoff at a later date. Therefore the value of future earnings are discounted by the rate an investor could receive in alternate investments. This approach enables the cost of an investment made in the present to be effectively weighed against the benefits of that investment repaid in the future, to arrive at the net present value of the project. This is achieved through a *discount rate*, d, whereby n future cash flows (each covering the same period of time – usually a year), R_t, earned at year t, are discounted to give the *net* present value, NPV, of an investment,

$$\mathrm{NPV}(d,n) = \sum_{t=0}^{n} \frac{R_t}{(1+d)^t}. \tag{5.13}$$

It is assumed that construction is instantaneous and occurs at $t = 0$ and is therefore included.

Clearly the NPV is highly dependent upon the choice of the discount factor, and therein is the first assumption that must be made in defining the financial models. Allan et al. [2011] cite a potential range of values for d in tidal-stream power investment, from between 6% to 15%. Compounded over an assumed turbine lifetime of 20 years (based on manufacturer's specified design life) and assuming a static price per unit of energy, this variation in d is equivalent to a variation in lifetime income of circa 35%.

5.6.1 Cost model

It is assumed that construction of the array is instantaneous, and therefore the cost of the cabling does not require discounting to attain its present value. Green et al. [2007] found purchase prices between two companies and for different types of cables ranged from £104.16 m^{-1} to £500.87 m^{-1} (converted to pounds sterling and adjusted for inflation to 2014 prices), with installation prices ranging from £64.39 m^{-1} to £70.71 m^{-1} and a vessel mobilisation fee of £3.42 M to £4.12 M (adjusted as before). Problematically, there has been a significant increase in the cost of energy in general and in renewable energy in particular. Although specific data is unavailable for tidal current power, Heptonstall et al. [2012] determine that the main drivers behind the increase in UK offshore *wind* energy are linked to:

- Increasing costs of materials and workforce;
- Constraints in material and production capacity throughout the supply chain;
- Detrimental trends in currency exchange rates;
- Constraints on capacity of installation infrastructure – i.e. shipping and ports;
- Planning and consenting bottlenecks.

Heptonstall et al. [2012] determined there had been a 51% rise in the cost of generating wind energy from 2006 to 2010. Given that all the above points (with the exception of planning and consenting) apply as much to sub-sea cabling costs,

and in the absence of data to the contrary, it is assumed that the 51 % cost rise can also be applied to the cable cost in tidal power installation.

Therefore the total cost of cabling is estimated at between £254.51 m^{-1} with a fixed £5.16 million vessel mobilisation fee and £862.48 m^{-1} with a fixed £6.22 million vessel mobilisation fee

$$5.16 \times 10^6 + 254.51L \leq C_C(L) \leq 6.22 \times 10^6 + 862.48L. \tag{5.14}$$

5.6.2 Income model

The price per unit of energy delivered to the grid is very difficult to estimate, especially over the anticipated lifetime of 15 to 20 years. The industry standard is to use a levelised cost of energy approach (LCOE, where [LCOE]= £ MWh^{-1}). Which is the present value of all the project costs divided by the discounted lifetime energy production of the plant

$$\text{LCOE} = \frac{\sum_{t=0}^{n} \frac{C_{\text{ann}}^t}{(1+d)^t}}{\sum_{t=0}^{n} \frac{O_{\text{ann}}}{(1+d)^t}}, \tag{5.15}$$

where O_{ann} is the annual energy output of the array (assumed constant) and C_{ann}^t is the annual cost of the array in year t (including capital and operational expenditure). It is assumed that the cost of capital and the developer's profit is encapsulated by the discount factor, d. Project planners will look at the price that will be achieved through sale of the electricity (including relevant carbon credits) and subtract from it the LCOE, if this difference is greater than or equal to zero then the project is viable.

A full exploration of the intricacies of discount factor choice is beyond the scope of this work, for which we simply seek a realistic value to demonstrate the OTF+TAR algorithm. Consequently, the interested reader is directed towards Allan et al. [2011] from where an upper and lower LCOE of £61 MWh^{-1} and £106 MWh^{-1} have been taken, based on a discount rate of 6% and 15% respectively.

The annual output of each MW of installed capacity (assuming one third of the rated plant capacity is delivered to the grid each hour) is 2920 MWh, discounted at between 6% and 15% and sold at a present value of £61 MWh^{-1} to £106 MWh^{-1} for between 15 and 20 years

$$61 \times P \times \sum_{t=0}^{15} \frac{2920}{(1.06)^t} \leq I_P(P) \leq 106 \times P \times \sum_{t=0}^{20} \frac{2920}{(1.15)^t}$$

$$1.91 \times 10^6 P \leq I_P(P) \leq 2.24 \times 10^6 P \tag{5.16}$$

where P is the installed capacity of the array – as returned by the instantaneous energy extraction from *OpenTidalFarm* using the steady-state shallow-water model.

5.7 Test cases

5.7.1 Idealised channel

The first test instance is an idealisation of a phased deployment in a simple channel. It is likely that in early schemes, a small number of turbines (perhaps eight) will be installed on a site that will end up accommodating a much larger number (perhaps hundreds). The initial deployment should provide a good business case to motivate developing out the rest of the site.

On this basis a scenario has been devised loosely scaled on a site in the Inner Sound of the Pentland Firth which has been leased by the company *MeyGen Ltd* (*www.meygen.com*) for tidal stream energy deployment. Figure 5.5 shows the schematic of the test layout overlaid on the Meygen site (or figure 5.6 shows the layout by itself). The site is a 1000 by 2200 m simple flat-bottomed channel with a 333 by 600 m leased array area located as shown in the figure. Flow enters at the left hand boundary at a constant 3 ms^{-1} and exits at the right. The top and bottom boundaries are frictionless. The domain is modelled at a constant depth of 30 m with (manning's) bed roughness set to a constant value of 0.0025. The viscosity was set to 3 m^2s^{-1}. The power collection point is located at the south-west corner and the cable capacity, C_{ap} is set to 6 (i.e. a maximum of 6 turbines may be connected in series).

Turbines were modelled as 20 m square with a minimum distance constraint of 20 m edge to edge between turbines.

With complex, nonlinear optimisation problems (such as in this case) gradient-based optimisation can be prone to finding local, rather than global optima [Barnett et al., 2014]. This means that the optimised layout can be dependent upon

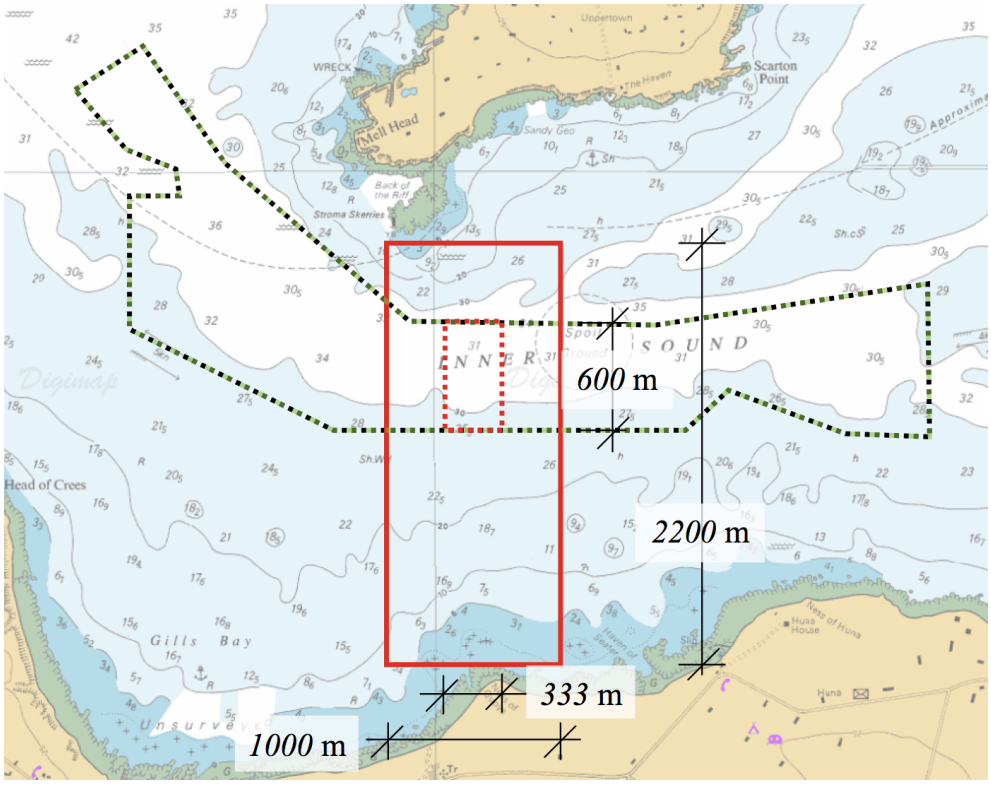


Figure 5.5: Map of idealised channel set up. The MeyGen Ltd (*www.meygen.com*) site in the Inner Sound of the Pentland Firth is shown by the black dashed line, the dimensions of the test domain (solid red line) and turbine area (dashed red line) are based upon a section of this real-world development site. Map from *edina.ac.uk/digimap*.

the starting guess of the turbine locations $\mathbf{m}^0$. To combat this issue and provide a robust procedure for array optimisation, the turbine locations were initially optimised solely to minimise cable length and once this converged (at a 'minimum cost' turbine layout) the optimisation was then restarted for both income and cost. The benefit of this approach is that, no matter the initial guess, $\mathbf{m}^0$, the resulting minimum cost layout will be broadly the same, so that the *combined* optimisation step will have the same starting layout independent of $\mathbf{m}^0$.

As such, 8 turbines are initialised in a grid formation (see figure 5.6(a)) within the lease area and the TAR optimisation is run. Within 8 iterations, the turbines have been relocated into the lower left corner of the turbine site (see figure 5.6(c)) and cable length is minimised. The optimisation is then restarted from this layout, with the objective now defined as in (5.1), to maximise the *return* provided by the array, assuming a 'worst case' scenario – i.e. using the upper bound of C_C from (5.15) along with the lower bound of I_P from (5.16). The model converges as shown

in figure 5.6(d) with a total power extraction of 32.19 MW for 3135 m of cabling.

In contrast, figure 5.6(b) shows the optimisation as run from the grid starting layout solely maximising the power. The turbines move far enough so as to no longer sit in each other's wake, achieving a power extraction of 29.45 MW – however this is only a local minimum – the algorithm has not found the solution in which the turbines are formed up into a single fence so as to exploit blockage effects. Consequently, in this instance, the coupled OTF+TAR algorithm actually provides a better solution.

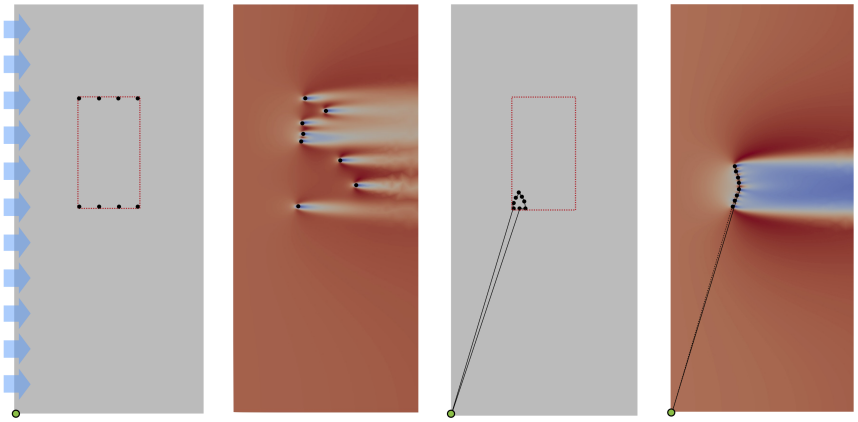


Figure 5.6: Eight turbines are optimised in the domain from figure 5.5 with the collection point located at the south-west corner (green dot). (a) shows domain layout and initial turbine locations for both runs (power = 11.55 MW, cable length would be 6424 m), (b) shows the layout after optimisation to maximise power (power = 29.45 MW, cable length would be 6500 m), (c) shows the layout after optimisation to minimise cable length (power = 12.92 MW, cable length = 2523 m), this is used as the starting point for the combined maximisation of financial return which results in the layout shown in (d) (power = 32.19 MW, cable length = 3135 m). The velocity magnitude field is shown in figures (b) and (d).

It may be more common to present power values in, for example, MWh delivered over a year. However, given the quoted MW values are based on an instantaneous power extraction computed from a steady-state shallow-water model solve, it seems most appropriate to report these values directly, rather than extrapolating the yield to represent a year's production – with all the assumption that this would entail. In addition, given the high viscosity and the highly idealised nature of the domain, reporting figures representing yield over a year implies a realism that is not reflected in the problem set-up and belies the idealisations that have been made.

Inflow conditions for the more realistic Pentland Firth scenario are designed so this value is broadly comparable to peak spring tide power.

5.7.2 Idealised channel bend

In more densely populated sites, there is less potential for the power-only optimisation to become 'stuck' in local minima – as there is less room for turbines to be positioned entirely out of each others way (when the ambient flow is monotonic and the turbines are out of the way of each other's wake, the gradient of the solution with respect to turbine position is flat – therefore the optimisation is deemed to have converged). Figure 5.7 shows a scenario taken from Funke et al. [2014] which is 30 m deep, flat bottomed simple channel shaped in a right-angled bend. The turbine site shown in the figure is initialised with a grid of turbines as before however this time there are a total of 16 turbines on a site that is 320 by 160 metres (that is a density of circa 312 turbines per square kilometer rather than 40 as before). C_C, I_P, C_{ap} and the boundary conditions were set as before.

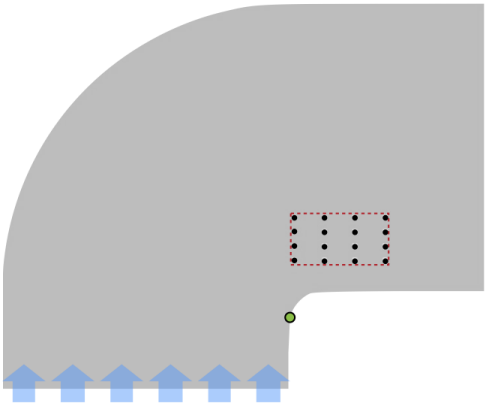
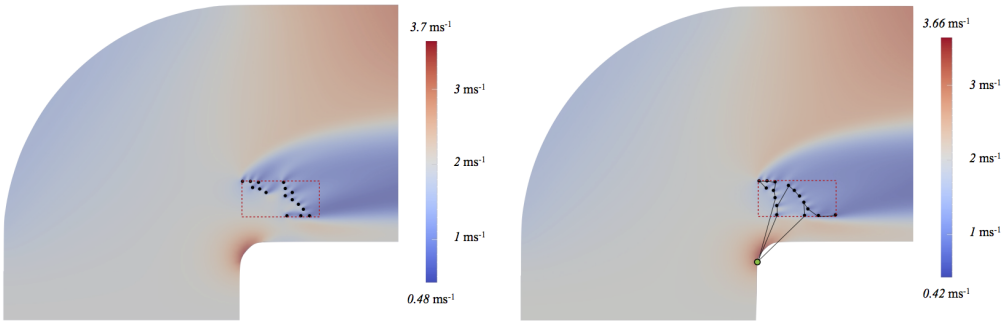


Figure 5.7: Domain of right angle simple channel example. The location of the power collection point is shown by the green dot, the turbine area by the red dashed line, and the initial turbine layout by the black dots. Inflow is shown by the arrows, outflow is across the vertical edge, the curved sides are frictionless boundaries. Scenario taken from Funke et al. [2014] (for $\mathbf{m}_0$ (shown), power = 28.86 MW, cable length would be 2587 m).

Once again, the OTF+TAR algorithm was first run to find the minimum cost solution and then run from that starting point to find the solution shown in figure 5.8(b) resulting in 35.00 MW power extracted for 2387 m of cabling. For comparison the power-only optimisation is shown in figure 5.8(a) and achieves 36.11 MW. The overall number of *OpenTidalFarm* iterations was broadly the same for both optimisations (60 iterations with cabling, 62 iterations without).



(*a*) (*b*)

Figure 5.8: Optimised micro-siting for (a) maximisation of power only (power = 36.11 MW, cable length would be 2413 m) and (b) maximisation of financial return (power = 35.00 MW, cable length = 2387 m).

5.7.3 Pentland Firth

While the previous scenarios were relatively academic, a more realistic scenario, which is also based upon an example in Funke et al. [2014], is modelled more recognisably on the geometry of the Inner Sound of the Pentland Firth. Once again, the bathymetry is fixed at a constant 30 m depth. The mesh and inflow direction are shown in figure 5.9, the inflow velocity was fixed to 3 ms^{-1}. In comparison to the previous examples, the domain is much larger, and thus to ensure the problem remains tractable on a desktop computer, the smallest mesh element size (located within the turbine area) was set to 10 m compared to 2 m in the previous two examples.

The relative increase in mesh element size meant that the viscosity also had to be increased to 90 m^2s^{-1} in order to ensure the convergence of the flow solve. This somewhat limits the quality with which hydrodynamic structures – such as wakes – may be represented. The potential inclusion of turbulence models is active work for the *OpenTidalFarm* project, this will help to balance the computational cost of a flow solve (which must be repeated for each optimisation iteration) against an assurance that the physical processes involved are being well captured. Funke et al. [2014] use a much finer mesh run on 64 cores in order to reduce this problem.

Despite these limitations, the mesh from figure 5.9 still serves to demonstrate the combined OTF+TAR optimisation. 32 turbines were initialised as before from a regular 8 by 4 grid within the turbine area of dimensions 1000 m by 500 m (shown dashed in red in figure 5.9). Maximisation of power output yielded a micrositing

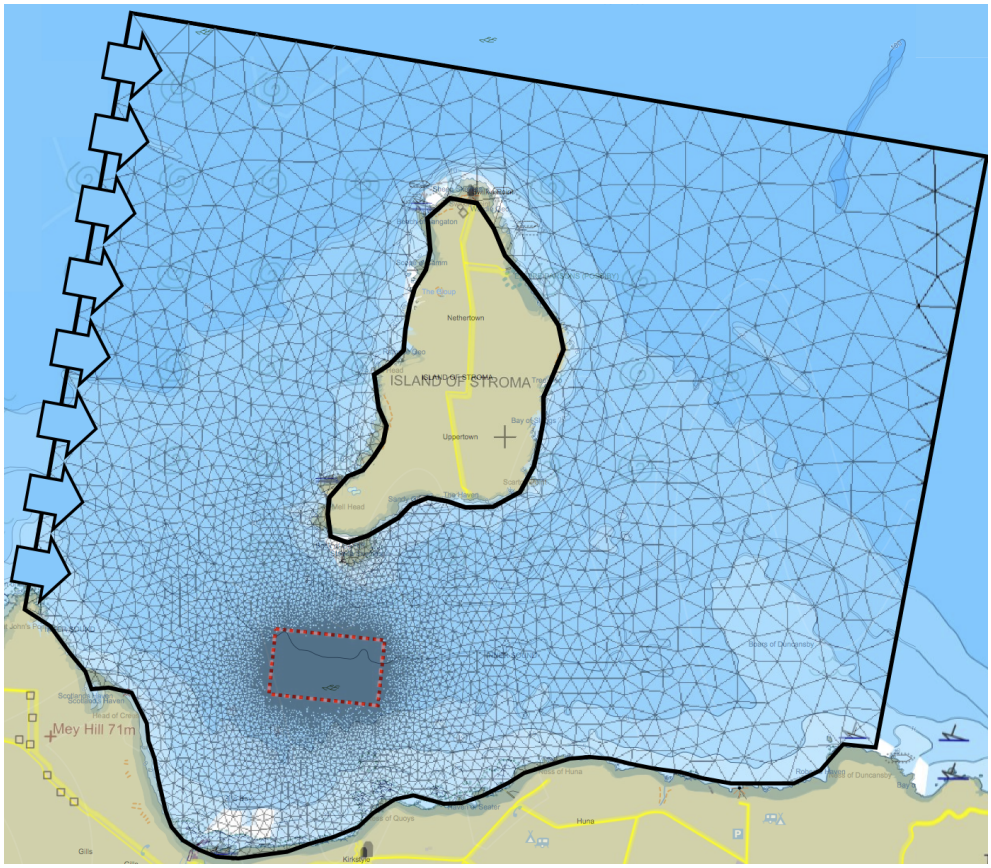


Figure 5.9: Domain, mesh and inflow boundary of the idealised Orkney example.

design as shown in figure 5.10 (a) which is broadly as expected from the work of Funke et al. [2014]. A barrage of turbines forms across the flow, with smaller 'spurs' acting as smaller ancillary barrages. As in Funke et al. [2014] there are, perhaps, too many turbines for the site; some of the turbines which cannot fit into the main barrage form up along the top boundary of the site and help to funnel flow in, the remaining turbines lie behind the main barrage.

When cable costs are included and the micrositing is optimised for the financial return, the main structure of the array, as described above is retained, implying that power output should not be radically different. As can be seen in figure 5.10 (b) the turbines have been drawn in towards the collection point. This resulted in an optimised array formation producing 69.47 MW and using 9.23 km of cabling, versus 70.57 MW produced by the array optimised for power only.

As a further investigation into the interplay between array design and financial return, the Orkney case was rerun for a scenario in which the the cost of the cabling

was doubled (perhaps due to supply bottlenecks or high demand for installation equipment). This resulted in an array formation as shown in figure 5.10 (c). As expected, the turbines are, in general, pulled further toward the collection point. However, the secondary 'spur' has grown significantly and now more closely resembles a full barrage. This reduced the cable length to 8.91 km for a power output of 67.53 MW.

5.8 Conclusions and summary

In this work, a general framework has been presented for the design of tidal turbine array micro-siting design with the inclusion of costs. This framework has been explored by considering the specific example of how sub-marine cable costs may be determined and balanced with the income from power generation. This work has given a practical example of how a cost functional with a discrete nature may be incorporated into a gradient-based optimisation approach. Thus showing that even relatively complex cost functions can be integrated into the *OpenTidalFarm* framework.

The benefit of this framework is that the industrial designer is furnished with a more holistic approach to micro-siting optimisation. The benefits of certain arrangements of turbines are automatically balanced with their associated costs. This inclusive approach is distinct from the more traditional methodologies in which the array design is determined largely based upon maximisation of power, while costs are considered at a later stage in the design process.

The combined optimisation algorithm was applied to three idealised scenarios, one under-populated simple channel, one simple-channel with a right-angle bend and one more realistic scenario based upon the Inner Sound of the Pentland Firth. In the first case, inclusion of the cable-routing cost minimisation facilitated an arrangement in which the power was actually increased as compared to a power-only optimisation run. This was because the power-only approach found, and could not improve upon, a localised optimum solution – while inclusion of the cable-cost enabled the turbines to be grouped together more closely thus minimising cable cost but also maximising blockage effects and thus the power extracted.

In the second case, with a right angle bend, the turbine area was much more highly populated and inclusion of the cable cost resulted in a decreased cable cost at the expense of a slight decrease in power production. Thus the balance of the

cost and income functionals was demonstrated.

Finally, an idealised model of the Inner Sound of the Pentland Firth was run with 32 turbines. Here the site had a good population of turbines for its size. The characteristic shape of the array, when optimised with consideration of cabling cost, was similar to the power-only optimisation but the turbines were arranged more tightly and towards the back of the site, close to where the collection point for the cables was located. The power extracted by this array was broadly the same as the array optimised for power alone while the cable cost could be reduced significantly. This demonstrated that the move to a more holistic design approach does not always mean a compromise in the power output of the array.

5.8.1 Future work

It is likely that cable contractors will aim to route cables to follow bathymetric features – rather than necessarily use straight runs as assumed here. An algorithm which could identify such features and direct cable routings to exploit them would be of value to developers. There is also scope to improve the parallelism of the TAR code, to improve computational run time. Apart from the methodology, the functions used to run the test cases, which convert the cable length and power extracted into a financial costs and incomes could be much improved. As has been discussed the proprietorial nature of such information makes it difficult to reliably source. Finally, cable cost is only one of several location-based costs, there is an ongoing effort to construct a library of such costs which, when used together, can provide a full accounting of the costs which feed into a micro-siting design.

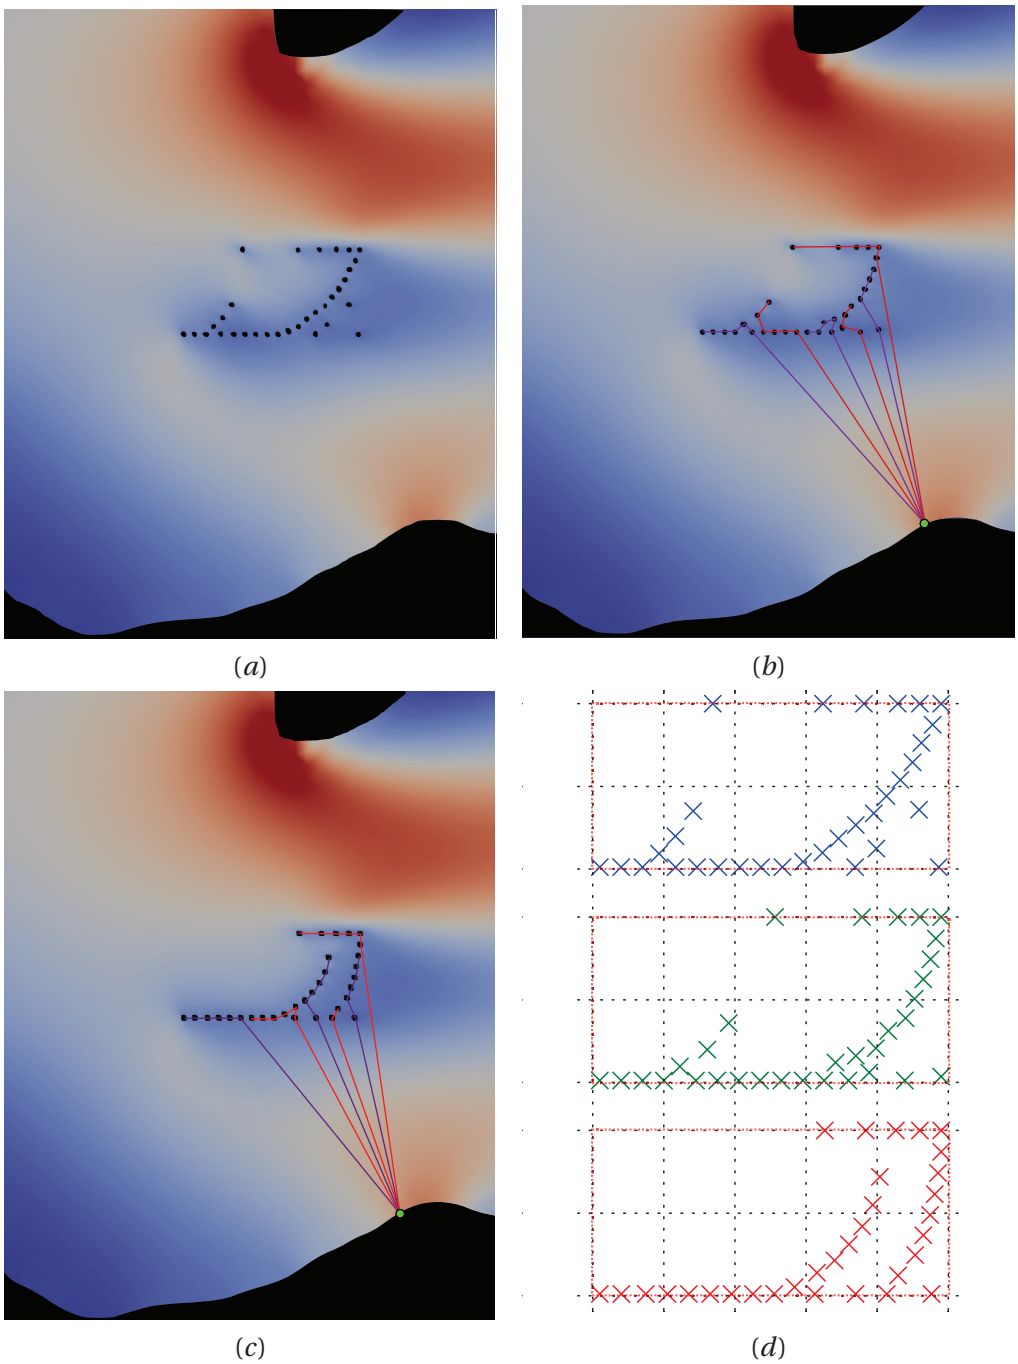


Figure 5.10: Close up of Orkney case with optimised micro-siting for (a) maximisation of power only (power = 70.57 MW, cable length would be 9.70 km) and (b) maximisation of financial return (showing cable routing) (power = 69.47 MW, cable length = 9.23 km). Figure (c) shows a scenario where the cost of cabling doubles (power = 67.53 MW, cable length = 8.91 km) and (d) shows a comparison of the three scenarios, with (a) in blue, (b) in green and (c) in red.

Chapter 6

Sizing a tidal turbine array

Contents

6.1 Introduction 140
6.2 Problem Formulation 142
6.3 Surrogate modelling 143
6.4 Combined Model 148
6.5 Idealised example 152
6.6 Pentland Firth example 156
6.7 Conclusions 160

This chapter is based upon Culley et al. [2016b]

Abstract

The new and costly nature of tidal stream energy extraction technologies can lead to narrow margins of success for a project. The design process is thus a delicate balancing act – to maximise the energy extracted, while minimising cost and risk. Scenario specific factors, such as site characteristics, technological constraints and practical engineering considerations greatly impact upon both the appropriate number of turbines to include within a tidal current turbine array (array size), and the individual locations of those turbines (turbine micro-siting). Both have been shown to significantly impact upon the energy yield and profitability of an array.

The micro-siting arrangement for a given number of turbines can significantly influence the power extraction of a tidal farm. Until the layout has been optimised (a process which may incorporate turbine parameters, local bathymetry and a host of other practical, physical, legal, financial or environmental constraints) an accurate forecast of the yield of that array cannot be determined. This process can be thought of as 'tuning' an array to the proposed site to maximise desirable outcomes and mitigate undesirable effects.

The influence of micro-siting on the farm performance means that determining the optimal array size needs to be coupled to the micro-siting process. In particular, the micro-siting needs to be repeated for any new trial array size in order to be able to compare the performance of the different farm sizes. Considering the large number of design variables in the micro-siting problem (which includes at least the positions of each turbine) it becomes clear that algorithmic optimisation is a key tool to rigorously determine the optimal array size and layout.

This chapter proposes a nested optimisation approach for solving the array size and layout problem. The core of this approach consists of two nested optimisation procedures. The 'outer' optimisation determines the array size. At each 'outer' iteration the power extracted by N turbines is found via a separate optimisation of their micro-siting on the proposed site. The 'inner' optimisation is treated as a computationally expensive black-box solver, mapping array size to power (and additionally returning the optimal micro-siting design). This forms the basis of a practical approach to the array sizing problem based on Bayesian optimisation, in which a surrogate model is built and used to maximise the utility of each evaluation made by the solver.

This chapter introduces and reports on the implementation of this novel surrogate-assisted array design approach which, coupled with a simple financial model is demonstrated through optimisation of array size for two test scenarios to maximise the financial return on the array for the developer.

6.1 Introduction

While the micro-siting optimisation problem has received some limited attention [Funke et al., 2014, Roc et al., 2013, Sullivan and McCombie, 2013, Thomson et al., 2011], the more fundamental question of how many turbines there should be in the first place, has been given very little practical consideration. Resource

assessment methodologies often make unrealistic assumptions as to how the resource may be harnessed – for example assuming that turbine barrages stretch across tidal straights from shore to shore, not leaving space for shipping or acknowledging that nearer shore there is insufficient depth to install turbines. Meanwhile, theoretical work – such as Vennell [2011b] – has framed the question in terms of blockage ratio and turbine density – a useful abstraction which facilitates general analysis of sites of various size (for example, two turbines in a narrow channel may have the same blockage as one hundred turbines in a wider channel). These models serve a useful purpose in identifying the upper bound of power available in a given region, and in helping us to understand how the hydrodynamic characteristics of the site/array change as the array is tuned. However, for a developer – who has a clearly defined turbine deployment area leased from the landowner (in the UK this is the Crown Estate) – these approaches are less helpful, and a more practical methodology is called for.

Since the flow dynamics are coupled to the locations of the turbines, and the power produced by those turbines is highly sensitive to the flow velocity, accurately forecasting the potential power output of an array of a given number of turbines, N, is both site and array-design specific and requires a micro-siting optimisation. Therefore in the course of determining the array size for a given site through iterative optimisation, significant computational effort is expended each time a different N is sampled. Consequently the problem is a computationally expensive integer optimisation whereby each time the performance of a given array size is evaluated an 'inner' micro-siting optimisation for that array is required.

The novel contribution of this chapter is the proposal of a practical methodology to study this relationship. The goal is that this approach may be adopted to scope projects on potentially complex sites and to specify the optimal design (both in terms of turbine numbers and their positions) of potentially large arrays. In this work, use of developer's return upon investment is proposed as the metric of array design quality, and a simple financial model is implemented to demonstrate the utility of this approach.

In the following section, the array sizing optimisation problem is formulated mathematically and split into an 'inner' continuous and 'outer' integer optimisation problem. The 'inner' micro-siting design optimisation approach is briefly reviewed in section 6.2 and the implementation of a surrogate model for the 'outer' optimisation is explored in section 6.3. The combined model is detailed in section 6.4 and applied to a simple, idealised, test case in section 6.5. In section 6.6 the

approach is demonstrated on a more realistic case based upon the initial deployment plans for a project which is currently under development in the Inner Sound of the Pentland Firth.

6.2 Problem Formulation

The power of a tidal current turbine array is computed as the sum of the power extracted from the flow due to the presence of the turbines. Encoding the Cartesian coordinates of each turbine as before in the vector, $\mathbf{m}$, recall that the power extracted by the turbine may be expressed as

$$\text{Power}(\mathbf{m}) = \int_{\Omega} \rho c_t(\mathbf{m}, \mathbf{u}) ||\mathbf{u}||^3 \, \mathrm{d}x. \tag{6.1}$$

Assuming the number of turbines, N, is fixed, then the maximum power that may be extracted from the flow is found to be

$$\begin{aligned} \tilde{P}(N) = \max_{\mathbf{m}} \quad & \text{Power}(\mathbf{m}) \\ \text{subject to} \quad & b_l \leq \mathbf{m} \leq b_u \\ & ||\mathbf{p}_i - \mathbf{p}_j|| \leq 0 \quad \forall\ i \neq j. \end{aligned} \tag{6.2}$$

This is a *continuous* optimisation problem; given $N \in \mathbb{Z}^+$, the turbine locations which provide the maximum power, $\text{Power}(\mathbf{m}) \in \mathbb{R}$, are returned as the vector $\mathbf{m} \in \mathbb{R}^{2N}$. Since N resides in the integer space, finding the optimum number of turbines in an array is found as an *integer* optimisation problem

$$\max_{N} \quad \tilde{P}(N) \qquad N \in \mathbb{Z}^+ \mid N \leq N_{\max}, \tag{6.3}$$

where $N_{\max}$ is the maximum number of turbines that the site is capable of supporting, which can be calculated from the minimum turbine distance constraints and the turbine lease area (though in reality may instead be limited by practical factors such as the number of turbines the developer can afford).

Using the notation $|_N$ to indicate the length of the preceding vector, the combined problem may be described as

$$\max_{N,\mathbf{m}} \quad P(\mathbf{m}|_N). \tag{6.4}$$

6.3 Surrogate modelling

The concept of surrogate modelling is to replace the (complex or expensive) model based on the *physical processes* being simulated – the *simulation* model – with a (simpler or cheaper) model which largely disregards the physical processes and is rather a mathematically constructed input-output map based on the observations made so far. For example, the shallow-water equations are based upon mathematical representation of conservation of properties. We may run an experiment twice with differing parameters using the shallow-water model, and then fit a function (which is the surrogate model) to the resulting data-set. The output at a different set of inputs can then be found by either re-running the experiment (requiring a shallow water solve), or else by evaluating the surrogate model at that point. The surrogate model is constructed via a fundamentally bottom-up approach, starting with data, using this to mathematically formulate an input-output mapping and then to evaluate it at unknown points rather than finding those values via the analytical model.

In this work, we use *OpenTidalFarm* to evaluate the maximum power that may be extracted from the site by a given number of turbines, which it does by optimising their locations – the array's micro-siting design. Taking a step back, we can treat *OpenTidalFarm* as a 'black-box' – an unknown, non-linear solver – which, given N, solves (6.2) and returns $\tilde{P}$. While *OpenTidalFarm* is efficient at the performing this task, it is still computationally expensive in the context of running a great number of separate times, evaluating different N's. This motivates the use of surrogate modelling to help minimise the number of calls to *OpenTidalFarm* in the course of optimising N.

Using surrogate models in this way – to replace an expensive black-box functional – is very common in computational approaches to aerodynamic design [Keane and Nair, 2005, Jones et al., 1998]. The surrogate chosen for this application is a stochastic process, and the overall approach is known as Bayesian optimisation. Stochastic processes and Bayesian optimisation are two large intersecting areas of a group of techniques often known as 'machine learning' and therefore represent large subject areas in their own right. A brief overview of the theory of being employed will be presented in the remainder of this section. Once again, the in-

terested reader is directed to more complete resources such as Rasmussen [2006], Zhou et al. [2004], Keane and Nair [2005] for more details.

6.3.1 Gaussian process regression

Gaussian process regression (GPR) was chosen to construct the surrogate model because it accounts for and quantifies uncertainty in the predictions it produces, supports multi-fidelity sampling (which will be explained in due course) and is fairly well established since it is one of the earliest developed stochastic process regressions. GPR was developed in the 1950's for geo-statistical applications – namely to estimate the level below the surface of strata formed by minerals of interest, based on borehole measurements taken elsewhere. GPR is part of the field of statistical learning theory which has its own distinct notation – consequently we have made the choice to keep discussion to a high level.

The key idea is that variables are no longer treated deterministically but are instead treated as *random variables*. This means the variable can carry several values each with an associated probability – i.e. can be represented by a probability density function (pdf). We begin by making assumptions about our underlying problem by defining, based on judgement, what this pdf may look like, this is called a *prior*. When samples are generated by running the analysis code, the model is 'trained' and combining the prior with these observations the *posterior* distribution is obtained. It is this which we use as a surrogate model. These observations are called the *training set*, $\mathscr{D} = \{(\mathbf{x}_i, f(\mathbf{x}_i))|i = 1,\ldots,n\}$.

A Gaussian process is a stochastic process X_t indexed by a continuous set t, where for any finite subset of t, say $\{t_1, t_2, t_3, \ldots, t_n\}$, $X_{t_1,\ldots,t_n} = (X_{t_1}, \ldots, X_{t_n})$ is a multivariate normal random variable. For clarity, begin by considering the bivariate Gaussian in figure 6.1. This plot shows the combined pdf for random variables x_1 and x_2. By projecting the bivariate Gaussian onto either the x_1y-plane or the x_2y plane the *marginal* probability distributions $P(f(x_1))$ & $P(f(x_2))$ are obtained – that is the probability distribution of one variable without reference to the other. By transecting the bivariate Gaussian along an x_1y-plane or x_2y plane, the *conditional* probability distributions $P(f(x_1|x_2))$ & $P(f(x_2|x_1))$ are obtained, which are contingent on the location of that transection. So, given the value of one variable, the probability distribution of the other can be obtained as this conditional pdf. The major advantage of using a Gaussian distribution in this context is that both the conditional and marginal probabilities can be written in closed form and are,

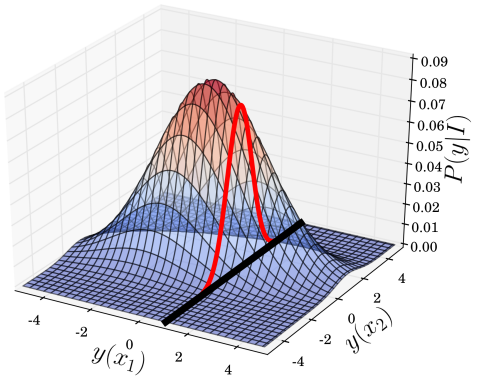


Figure 6.1: Linking two points using a bivariate Gaussian distribution enables the conditional pdf (red) for the function at the unknown point x_2 to be calculated based upon the known point x_1.

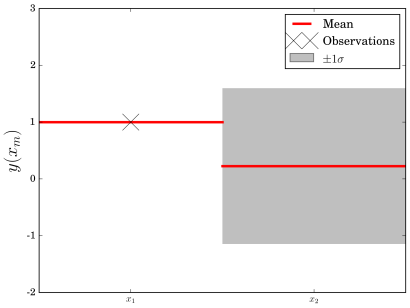


Figure 6.2: The posterior distribution of the relationship in figure 6.1 showing the distribution of $y(x_2)$ given $y(x_1) = 1$.

themselves, Gaussian.

In figure 6.1 the bivariate Gaussian represents the *prior*. If the 'experiment' were then run and a value $f(x_1) = 1$ obtained (shown by the black line), we can evaluate the conditional probability distribution for $f(x_2)$ which can (like any Gaussian distribution) be parametrised by a mean and covariance. This can be plotted in an alternate but equivalent way as in figure 6.2. This figure illustrates the *posterior* model – the result of combining the prior with observations and this is what is used as the surrogate. Since $f(x_1)$ was measured (from the analysis code), this can be plotted directly and without error bars (we're certain of its value), for $f(x_2)$ the posterior mean is plotted – that is the mean of conditional probability distribution of $f(x_2)$ given $f(x_1)$, $P(f(x_2|x_1))$. However error-bars can also be plotted or, more accurately, confidence intervals, based upon the covariance of that conditional probability.

This concept can be generalised to multivariate Gaussians of m dimensions or indeed an infinite-dimensional Gaussian process. To generalise to a multivariate Gaussian, the input space is discretised into m segments and the prior is defined by an m-long mean vector, $\boldsymbol{\mu}$, and an $m \times m$ covariance matrix, $\boldsymbol{K}$. Frequently a zero vector is used for $\boldsymbol{\mu}$ and the covariance is set using a kernel function such as

$$\kappa(x, x') = \sigma_f^2 \exp\left(-\frac{1}{2l^2}(x - x')^2\right). \tag{6.5}$$

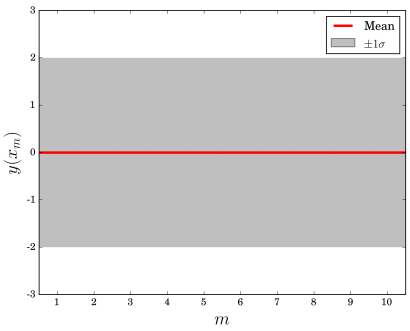


Figure 6.3: Prior distribution for multivariate Gaussian distribution over 10 variables – here a zero-mean has been assumed.

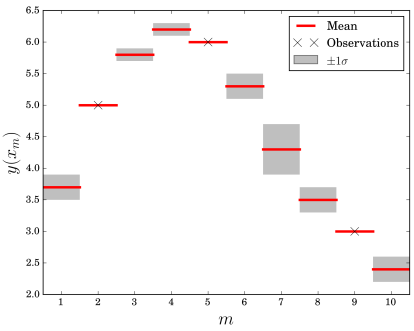


Figure 6.4: Posterior distribution for multivariate Gaussian of figure 6.3 given a training set, $\mathscr{D}$ of observed values.

which represents the prior assumption that similar inputs will beget similar outputs, and the similarity of the outputs can be estimated by the exponential of the squared difference of the inputs (x and x' simply represent two different inputs). This squared exponential is one particular choice of the covariance function (probably the most popular, though various others may be used). The 'hyperparameters', l and σ_f^2 represent the (horizontal) length scale and maximum allowable covariance (i.e. related with y-axis scale) respectively and allow the function to be 'tuned' to the particular problem – i.e. to modify the assumptions upon which we build the prior.

Generalising figure 6.2 to 10 dimensions, we might define a prior as in figure 6.3 which, with observations may become the posterior in figure 6.4.

The major strength of using a Gaussian process regression to construct the surrogate, rather than a parametric model, is that the former much better reflects what is actually known based on the observations and incorporates uncertainty away from those observations. For example, a parametric model may be fitted to some arbitrary function of interest, $y(x)$, based on some samples of x taken in the range $0 < x < 100$. If this is used as a surrogate to evaluate $y(\alpha)$ for $\alpha >> 100$ a *value* of y will be directly returned – with no indication that the model has been extrapolated. Conversely if GPR was used to construct the surrogate, evaluating $y(\alpha)$ would return a *distribution* with a large covariance (due to $(x-\alpha)^2$ in (6.3.1) being large) reflecting low confidence in the expected value.

A second major strength of using GPR to construct a surrogate model is that the observations which comprise the training set are also random variables. This means that observations can be taken using variable fidelity analytical models,

with observations from high fidelity models being passed in with a pdf closer to a delta function (i.e. very low covariance) and observations from lower fidelity models being passed in with a larger covariance to reflect the uncertainty attached to use of a less accurate model. This is very useful because the fidelity of the *OpenTidalFarm* 'inner' optimisation can be controlled by adjusting the convergence criteria of the gradient-based optimisation algorithm. To explain; a feature of gradient-based (or 'first-order') optimisation routines is that they generally operate greedily – i.e. an iteration only concludes if the functional has improved. In fact a graph of functional by iteration generally takes a form similar to logarithmic growth, so while a fully converged micro-siting optimisation may require 80 iterations (treating *OpenTidalFarm* as a black-box solver this would represent a *high* fidelity solution) if the optimisation were only allowed 10 iterations the functional value will be lower than the fully converged value, but will be in the vicinity of it. This unconverged value represents the *low* fidelity solution.

6.3.2 Bayesian optimisation

The general procedure of Bayesian optimisation is that each time the experiment is run, a new data-point is sampled and this is added to the training set. The posterior is recomputed based upon the larger training set and the next data-point to sample is chosen based upon an *acquisition function* constructed from this surrogate model. The acquisition function chosen for this application is a function of the posterior mean and variance

$$\mathrm{acq}(N) = \boldsymbol{\mu}(N) + \kappa\boldsymbol{\sigma}(N). \tag{6.6}$$

where κ is a tuning parameter which can be chosen to either 'explore' or 'exploit' N. 'Exploring' means taking samples away from observations that have already been made – i.e. where the variance is largest – hence larger values of κ will induce this behaviour in the optimisation. 'Exploitation' means testing around the area of the expected optimum in greater detail and thus identifying it with tighter tolerance – in this mode $\kappa = 0$.

The strategy for this array design application is to start the optimisation in an exploration mode so as to identify the general shape of $\tilde{P}(N)$ – this will be done using low fidelity samples (i.e. limiting *OpenTidalFarm* to fewer iterations than

required for full convergence). Once the the region of N that is likely to contain the optimum has been identified, κ is reduced and new sample points will be chosen so as to exploit this region and close in on the optimal solution.

One additional benefit of having access to the confidence bounds on estimated values is that a 'probabilistic branch and bound' method – as outlined in de Freitas et al. [2012] – can be used to quickly identify feasible regions for the optimal solution. Recall that the purpose of the surrogate is not to recreate $\tilde{P}(N)$ accurately over the whole range N, but rather to represent $\tilde{P}(N)$ sufficiently well that the region containing the optimal solution is correctly identified (through the exploration phase) and then in this region ensure the representation of $\tilde{P}(N)$ is accurate. Using confidence bounds provided by the variance enables infeasible regions of the function to be quickly eliminated: Figure 6.7 has the 95 % confidence bounds plotted (as computed by 1.96σ from the mean). These can be used to discard regions in which the upper confidence bound is less than the maximum of the lower confidence bound, $\max(\mu - 1.96\sigma)$, as these regions are highly unlikely to contain the optimum [de Freitas et al., 2012].

6.4 Combined Model

To make it easier to refer to, the Bayesian optimisation routine as implemented is called *TidalArraySizer*. *TidalArraySizer* interfaces with *OpenTidalFarm* to optimise array size through a Bayesian optimisation approach with variable fidelity sampling and a probabilistic branch and bound method which identifies feasible regions for the optimal solution.

The optimisation begins in exploration mode ($\kappa = 1.96$) taking low fidelity samples at the bounds and midpoint of N. The geometry of the turbine area and the spacing constraints (for example minimum distance constraints) on the turbines determine the maximum number of turbines the site can support. The micro-siting is optimised for $N_{\max}$ to give an evaluation of the power output of the most populous array possible for the site, likewise the $N = 0$, $P = 0$ data point fully evaluates the training set bounds. The next data point is taken at $N_{\max}/2$ which completes initialisation of $\mathscr{D}$, the training set for the GPR.

The construction of the GPR model is handled by the Python package *SKLearn* [Pedregosa et al., 2011] which optimises the hyper-parameters to minimise the variance in the predictions given the observations. When a new data point is

added to the training set, the model is reconstructed and those hyper-parameters re-optimised. For an in-depth discussion on the optimisation of hyper-parameters see Rasmussen [2006] and on implementation see Pedregosa et al. [2011].

A surrogate is built up based upon this initial training set, with low fidelity samples. Each time an additional N is sampled, the data is added to the training set and the surrogate is rebuilt and the next point for testing identified using the acquisition function, (6.3.2). Efficiency in the methodology is achieved by using low fidelity samples to identify the general trend of the function over the range of $[0, N_{\max}]$ during the initial exploration phase.

As the size of the training set increases, and the design space becomes more populated with (low fidelity) samples, there reaches a point where the range of array sizes for which the upper confidence bound exceeds the maximum lower confidence bound ceases to shrink with additional sampling. The rate of this shrinkage is used as a criterion for switching to exploitation mode, with high fidelity sampling.

When this stage is reached, *TidalArraySizer* switches to collecting high-fidelity samples - i.e. tightening the convergence criteria of the 'inner' *OpenTidalFarm* optimisation. If necessary, low fidelity samples can be restarted and allowed to converge and then are updated to high fidelity samples in $\mathscr{D}$. This is the 'exploitation' phase wherein the function has been sufficiently well explored to be certain the maximum is in the region being tested at high fidelity. At this point, the acquisition function κ weighting is set to zero and samples are taken at high fidelity.

At the conclusion of the optimisation, the optimum number of turbines is returned, along with their optimal micro-siting design on the site and a plot of the variable-fidelity surrogate of $\tilde{P}(N)$. Various detailed diagnostic and prognostic output data is also available (for example the flow velocity plots at each timestep or power output per turbine / per timestep, among many others). A flow chart of the combined model is shown in figure 6.7.

6.4.1 Financial model

In realistic turbine areas, limited by constraints such as allowing shipping lanes, channel blockage can often not reach a point where a clear maximum of $\tilde{P}(N)$ is found – rather the function tends to plateau as $N_{\max}$ is approached. In general, tidal array developers are interested in maximising the value of their investment,

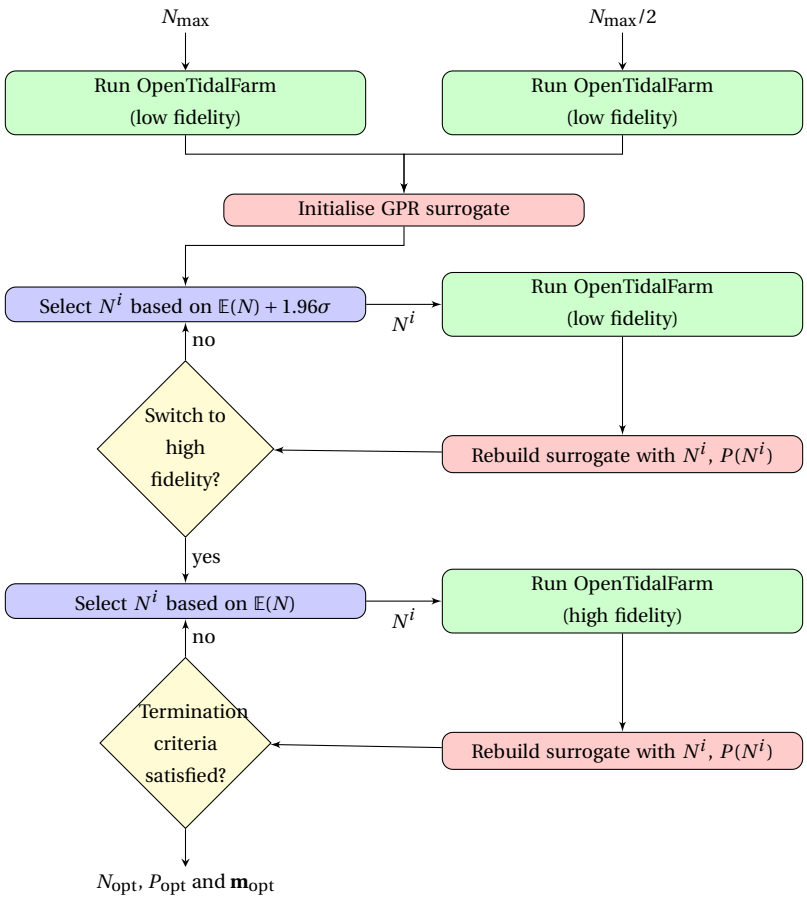


Figure 6.5: Flow chart showing the ensemble *TidalArraySizer* model.

and diminishing returns on investment in additional turbines for an array runs contrary to this goal. While numerous measures of project finance exist (measures such as levelised cost of energy, LCOE, and net present value, NPV, are commonly used) the data upon which these models are based is generally proprietorial and the assumptions required to model a project over, say, a 25 year lifespan give the models a complexity which is beyond the scope of this work. Instead, a simplified financial model has been implemented which demonstrates the utility of the methodology. The modular, object-based design of the *TidalArraySizer/OpenTidalFarm* code allows users to substitute their own financial models to run their own analyses. A similar idealised approach to financial modelling was used in Culley et al. [2016a] to demonstrate optimisation of micro-siting design to maximise profit as computed by the income from the power extracted by the turbines minus the cost of sub-sea electrical cabling. For that work, detailed

financial data was likewise unavailable – so simplifying assumptions were made to demonstrate the utility of the methodology.

For this work, a notional measure of profit will be used which is computed by subtracting the lifetime cost of the array from the lifetime income generated by the power production of the array. It is assumed that the cost of the array is solely a function of the number of turbines multiplied by a constant cost per turbine. A cost per turbine, of £4 million has been chosen to reflect the capital expenditure on the turbine itself. In addition to a per-turbine contribution to support-infrastructure (cabling, electrical conditioning infrastructure etc.), installation and maintenance of an additional £4 million has been added giving a total cost of £8 million per turbine.

Modelling the income generated from power production over the array lifetime is much more difficult as it depends not only on the performance of the array but also on the energy market and governmental support schemes which can vary enormously over 25 years. Consequently the income will be computed as the idealised income per unit of energy multiplied by the power extracted by the array over the lifetime of the project. The idealised income per unit extracted, $I^{\text{idealised}}_{\text{per unit}}$ is computed from the price per unit of energy that is necessary for an ideal turbine to break even over its life. For this, we measure the energy extracted by a single 1 MW turbine in a wide simple channel with a sinusoidally head-driven inflow which varies between $\pm$ 3 ms^{-1} over a 12.5 hour tidal cycle. Multiplying by the design life, $N^{\text{design}}_{\text{tidal cycles}} = 43,800$, gives the energy extracted by that turbine over its life and dividing the cost of the turbine C_{turbine} = £8 million by this value gives the idealised income per unit of energy, $I^{\text{idealised}}_{\text{per unit}}$. The energy extracted by the array over its lifetime is found by choosing a representative 12.5 hour cycle in which the average power extraction corresponds to the average over the longer period (running the optimisation based on a longer model time would become too computationally expensive). By measuring the energy extracted over this period and multiplying by the number of tidal cycles in the array lifetime and the idealised income per unit energy we arrive at the lifetime income of the turbine array.

$$\text{Profit}(N) = \text{Income} - \text{Cost}, \tag{6.7}$$

where

$$\text{Cost} = N \times C_{\text{turbine}}, \tag{6.8}$$

$$\text{Income} = I^{\text{idealised}}_{\text{per unit}} \times E^{\text{cycle}}_{\text{array}} \times N^{\text{design}}_{\text{tidal cycles}}, \tag{6.9}$$

and

$$I^{\text{idealised}}_{\text{per unit}} = \frac{C_{\text{turbine}}}{E^{\text{ideal cycle}}_{\text{turbine}} \times N^{\text{design}}_{\text{tidal cycles}}}. \tag{6.10}$$

Clearly this is a heavily simplified financial model whose suitability only extends to proof of concept. When realistic, publishable data is available to improve upon this, future work will look at the impact of a less idealised model. In addition, greater accuracy could be achieved by modelling for longer than a single tidal cycle. Recall that this model is designed to produce an optimal array size and micro-siting design and an idea of the power output for resource assessment purposes. Once an optimised design has been returned more accurate, higher fidelity hydrodynamic and financial models may be used to produce more reliable power and financial return estimates for the array, modelled over a longer period of time.

6.5 Idealised example

The framework will first be demonstrated on a highly idealised example problem. The domain is a 2560 by 1280 m simple channel with a circular island, 320 metres in diameter located at its centre. The turbine area measures 640 by 320 metres and is centred longitudinally in the channel and in the middle of the sub-channel to the south of the island. The element size for the (unstructured) finite-element mesh varies from 5 metres in the turbine area to 25 metres furthest from it (figure 6.6). The depth is a constant 50 metres throughout the domain.

Flow is driven by a sinusoidal velocity-forcing ranging from ± 2 ms^{-1} at the in-flow boundary over a period of 12.5 hours, divided into 15 time-steps. The viscosity was set to 2 m^{2}s^{-1} and the turbines were parametrised as per (3.13) and figure 3.3.

The surrogate is initialised with $N = 0$, $P = 0$ and low fidelity optimisation runs of $N_{\max}$ and $N_{\max}/2$, where $N_{\max}$ is the maximum numbers of turbines that can fit on the site (assuming a minimum centre-to-centre spacing of 30 metres $N_{\max} = 21 \times 10 = 210$). For this example, the low fidelity standard was set by an optimisation tolerance of 1% – which for most samples represented an optimisation requiring circa 15 iterations (that is 15 adjoint solves and typically a similar number of shallow water solves). The initialised surrogate model is shown in figure 6.7.

In the low fidelity phase of the optimisation, the acquisition function is defined as (6.3.2) with $\kappa = 1.96$. Low fidelity samples continued to be taken until the range of the region in which the upper confidence bound is greater than the maximum

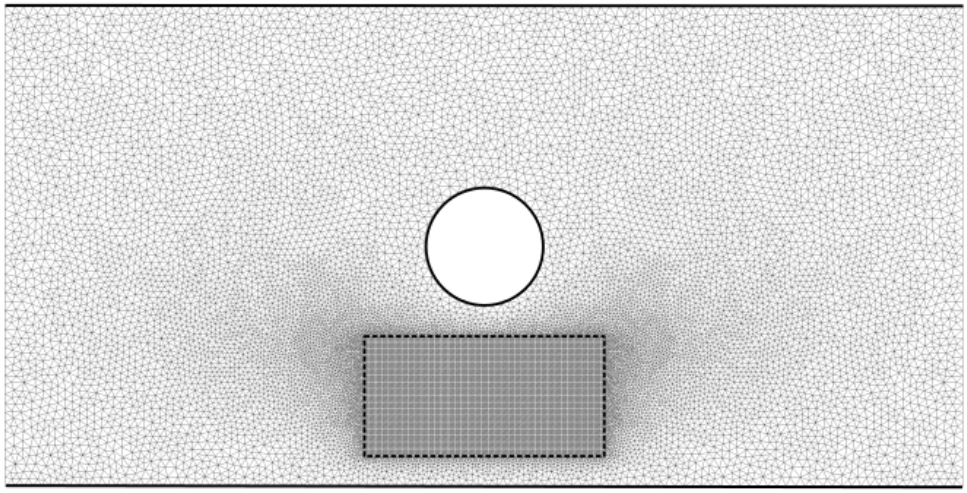


Figure 6.6: Unstructured computational mesh of the simple channel with circular island test example

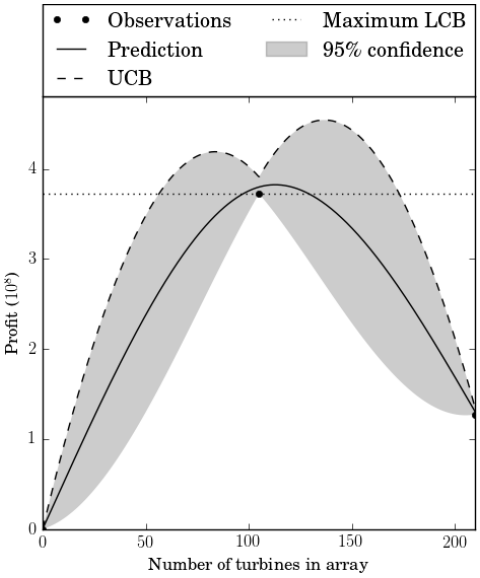


Figure 6.7: Plot of the Profit(N) surrogate initialised with training set data points at Profit($N = 0$), Profit($N_{\max}$), Profit($N_{\max}/2$).

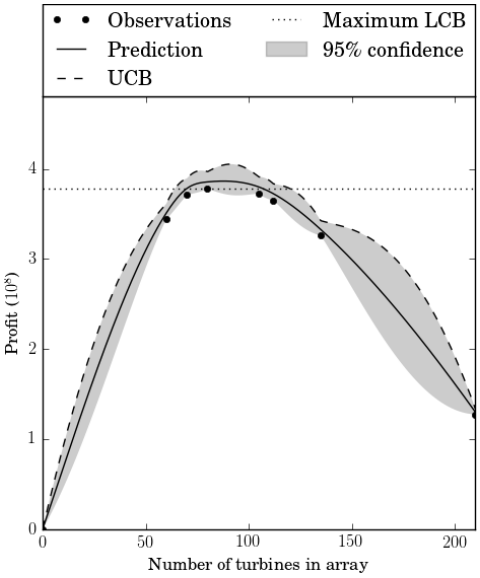


Figure 6.8: Plot of the Profit(N) surrogate before switching to high fidelity sampling.

lower confidence bound (shown where the dashed upper confidence bound exceeds the dotted maximum lower confidence bound line in figures 6.7 and 6.8) stopped decreasing. Figure 6.8 shows the surrogate model of the objective function once the algorithm has completed collecting low fidelity samples.

Now the $\kappa = 0$ acquisition function selects sample points based on the maxi-

mum of the surrogate prediction. Low fidelity runs can be restarted and allowed to converge to high fidelity. Low fidelity sample points can be clearly identified in the finished surrogate shown in figure 6.9 where the upper confidence bound does not fully 'pinch' in at a data point.

The finished surrogate can be seen in figure 6.9, along with the final turbine layout in figures 6.10 and 6.11. It can be seen that the optimisation was achieved in four low fidelity and five high fidelity *OpenTidalFarm* calls. In this instance an optimum number of 80 turbines (±5 turbines) is returned as the optimum array size, this is less than half of the maximum number the site can support given the minimum distance constraint imposed between turbines.

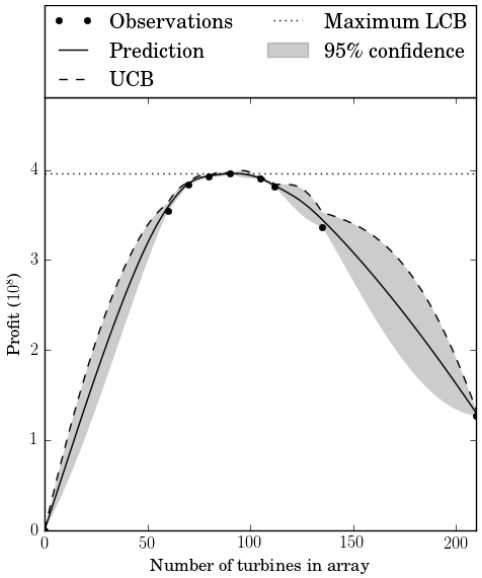


Figure 6.9: GPR surrogate at end of optimisation.

The turbine layout is significantly different from Funke et al. [2014] which shows a similar problem, though on a smaller domain and without any power curve (in Funke et al. [2014] turbines are not realistically parametrised, rather Power $\propto \mathbf{u}^3$ and is uncapped). At the up and downstream ends of the turbine area, the turbines are lined up in a clearly recognisable barrage formation. Between these two barrages the layout seems less coherent. Rather than large scale barrages, turbines are grouped in sets of 3 to 5 turbines which are staggered in diagonal formations, allowing downstream turbines to capitalise on the local flow acceleration around the sides of the upstream turbine. Turbines along the north and south turbine area boundaries capture energy as the streamlines diverge while the flow moves through the turbine area, but also as flow within the turbine area is re-energised

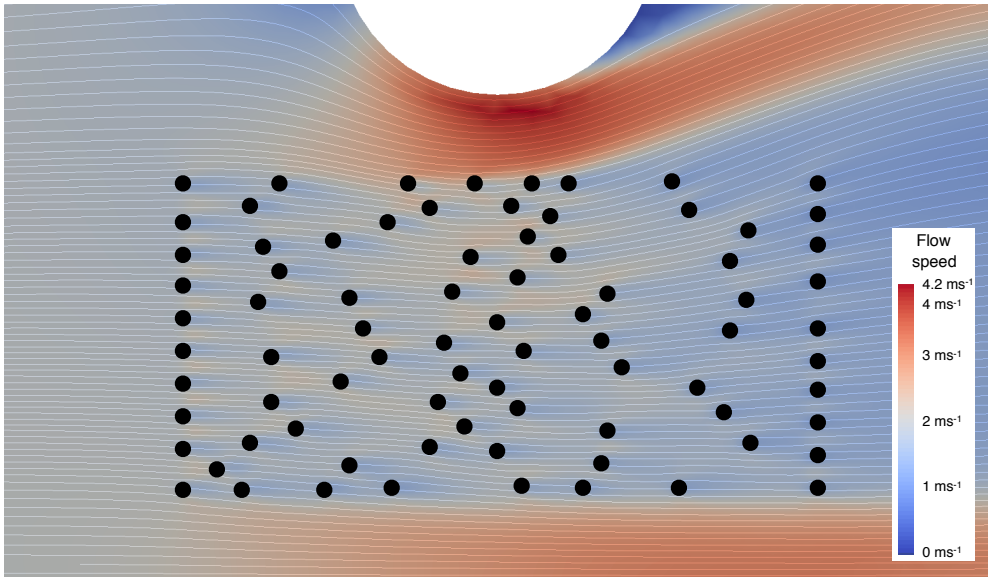


Figure 6.10: Optimised location of the optimal number of turbines for the site showing flow velocity magnitude and streamlines for 'flood' tide (flow left to right).

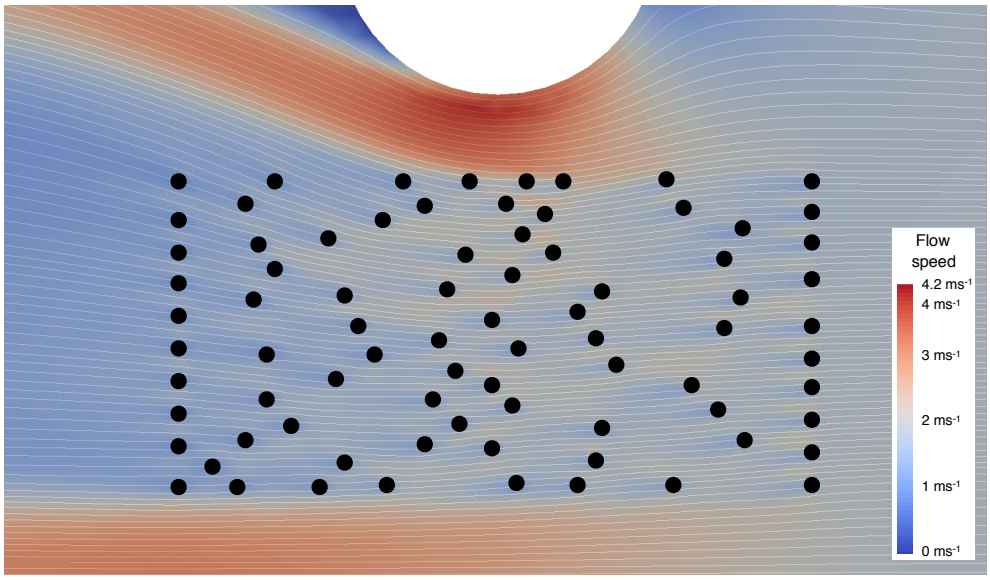


Figure 6.11: Optimised location of the optimal number of turbines for the site showing flow velocity magnitude and streamlines for 'ebb' tide (flow right to left).

through shear effects from the bypass flow. This layout represents an 11 % improvement on the starting layout of a regular grid formation, the energy extracted over the tidal cycle for the optimised array was 2.458 TJ.

6.6 Pentland Firth example

This example uses a domain based on the Pentland Firth, a site between the Scottish mainland and the Orkney Islands. This is one of the prime sites being considered for deployment of tidal stream turbine arrays and the turbine area used in this example is based on the phase 1 deployment site for which the developer *MeyGen* has obtained consent to develop an 86 MW array. A simplified version of this example on a smaller, flat-bottomed analogue of this domain was considered in Culley et al. [2015].

The unstructured mesh is shown in figure 6.12, the element size ranges from 5 metres in the turbine area, to 40 metres along the coastline and islands to a maximum of 400 metres away from these features. The mesh was generated by the open-source meshing software Gmsh [Geuzaine and Remacle, 2009]. The domain is bounded to the south by a portion of the Scottish mainland coast from Dunnet Head to just round Duncansby Head. This is artificially extended out to the eastern open boundary, which is situated a few kilometres east of the Skerry Islands. The western open boundary stretches from Dunnet Head to the nearest point on the island Hoy. The northern boundary is modelled (as are the coastlines and island boundaries) with a zero velocity Dirichlet condition – i.e. no-slip and no-normal flow. Clearly this is a simplification as it closes off the Pentland Firth's connection to Scapa Flow through the waterways between Hoy, Flotta and South Ronaldsay. This simplification was made to keep the problem size tractable on a single HPC node of 12 cores.

The domain is based on a simplification of the validated set-up used by Funke et al. [2015a]. Bathymetry information is provided by combining data from four sources (as outlined in Funke et al. [2015a]). Flow is head-driven with the head difference across the eastern and western boundary calculated using the OSU Tidal Prediction Software [Egbert and Erofeeva, 2010]. The model simulates one tidal cycle starting at 1040 on 18-09-2001 using the Q1, O1, P1, K1, N2, M2, S2 and K2 tidal constituents. The start of the simulation time was chosen as the head difference is close to being zero using these constituents at this time.

A period of 12 hours is used, subdivided into 40 time steps. The viscosity varies over the domain with proximity to the open boundaries to aid stability. At the turbine site, the viscosity is 10 m^2s^{-1}, a higher viscosity of 1000 m^2s^{-1} was used in a narrow region (1 km) near the boundaries. Once again, the turbines were parametrised as per (3.13) and figure 3.3. The turbine area is 1000 by 600 metres in

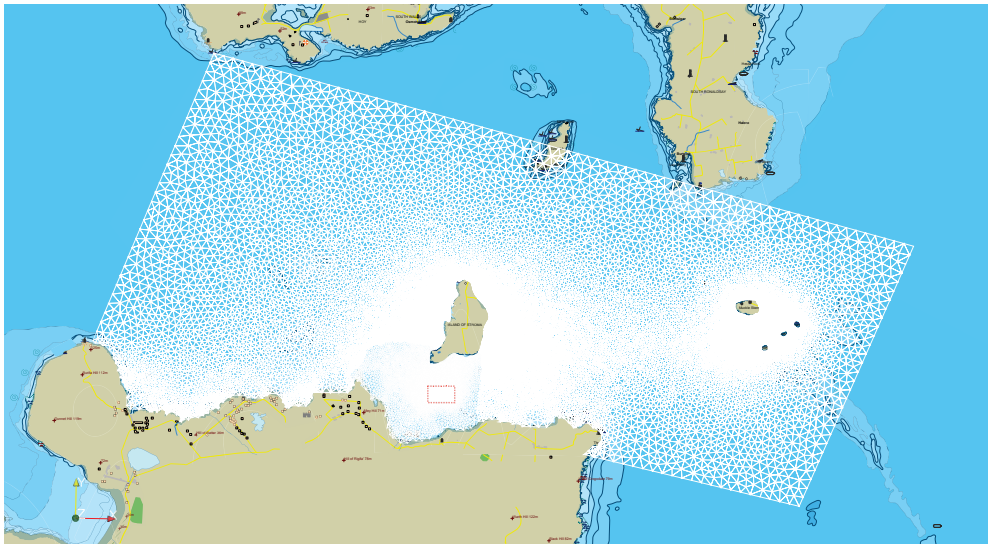


Figure 6.12: Map of Pentland Firth with unstructured mesh superimposed in white

size and located to the south of Stroma in the Inner Sound of the Pentland Firth, as shown in figure 6.12 by a dashed red box.

As before, the surrogate was initialised with $N = 0$, $P = 0$ and low fidelity optimisation runs of $N_{\max}$ and $N_{\max}/2$, where again $N_{\max}$ is the maximum numbers of turbines that can fit on the site (assuming the same minimum centre-to-centre spacing as before $N_{\max} = 33 \times 20 = 660$).

Once again, low fidelity samples were taken over the range of $0 \leq N \leq N_{\max}$ from which the surrogate shown in figures 6.13 and 6.14 was constructed.

In contrast to the results for the example in section 6.5, in this case, the optimised turbine micro-siting design for the optimal number of turbines is much more ordered (see figure 6.17), exhibiting several clearly defined rows of turbines. To the east of the turbine area, two well ordered rows of turbines have emerged, within these rows there is some offsetting of one turbine to the next. This staggering is, once again, likely to exploit local blockage induced flow accelerations around the sides of turbines. One reason for this may be that the previous example was driven by a constant velocity in a comparatively small domain, while this example is driven by a head difference over the domain. This means that (in the previous example) flow is forced through the site and bulk flow redirection to the north of the island is a less important mechanism. The result of this is that the site can support more turbines and the site is much more densely populated than in this Pentland Firth example. With more densely packed sites, the interactions

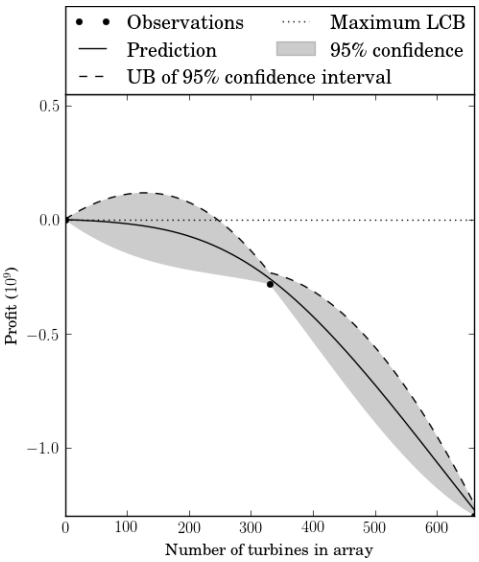


Figure 6.13: Surrogate model initialised with starting training data.

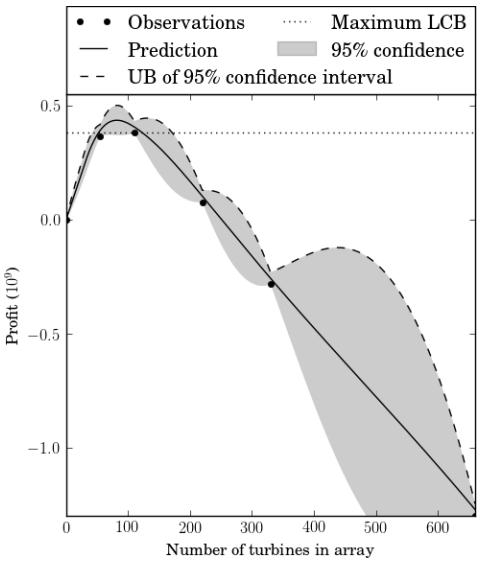


Figure 6.14: Surrogate model at conclusion of low fidelity runs.

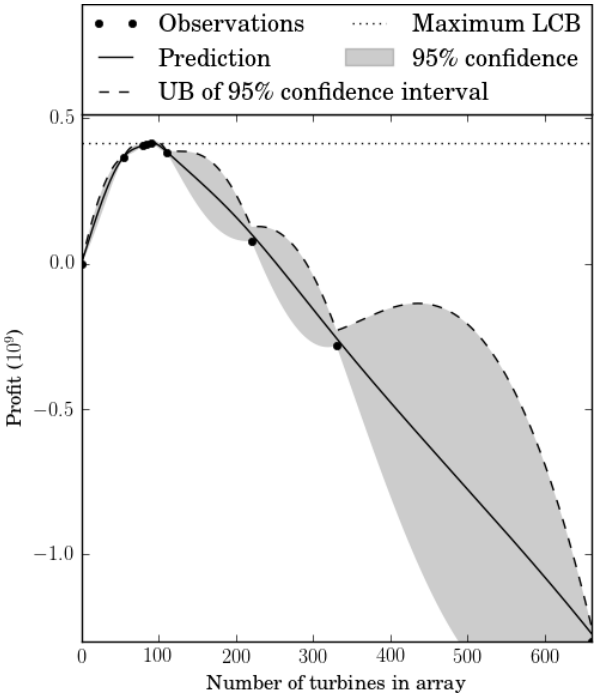


Figure 6.15: Final surrogate model of N to P mapping for the Inner Sound example.

between the turbines over the tidal cycle are that much more intricate and chaotic – leading to a less well ordered arrangement of turbines.

To the west of the turbine area, several smaller barriers have formed, which lie

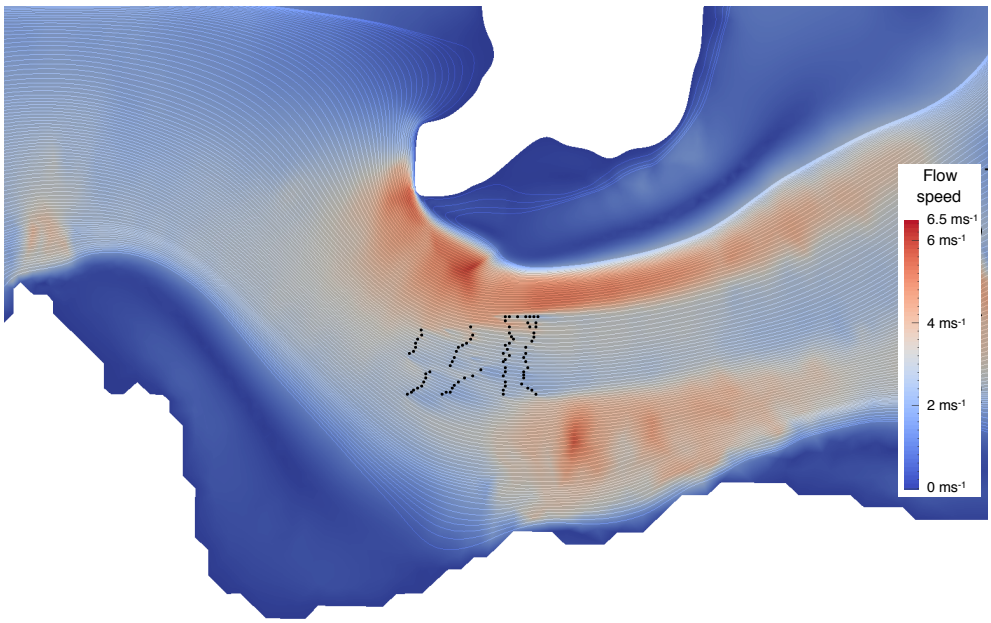


Figure 6.16: Optimised location of the optimal number of turbines for the site showing flow velocity magnitude and streamlines for 'flood' tide (flow left to right).

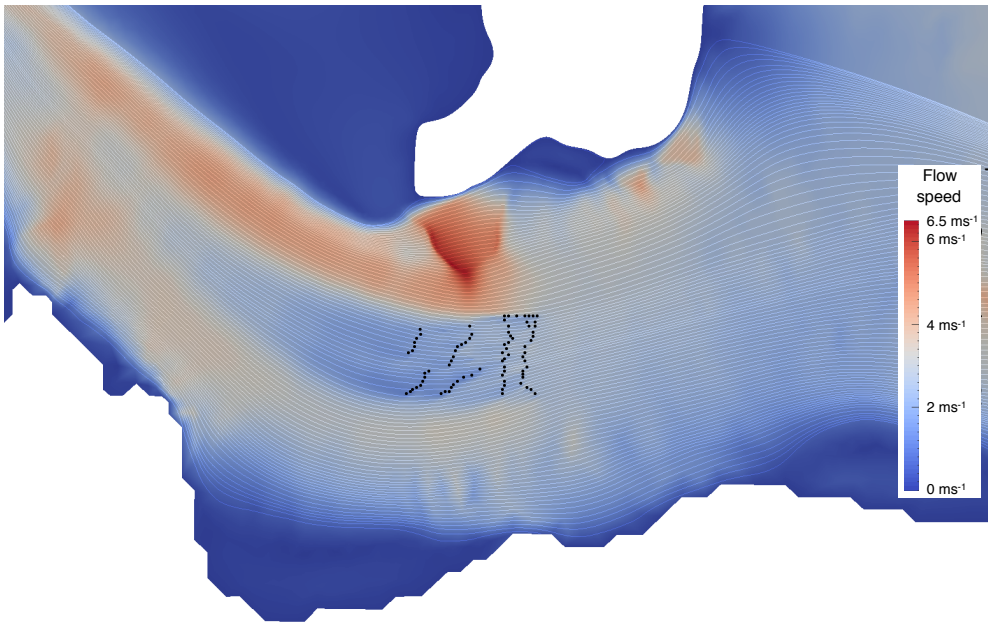


Figure 6.17: Optimised location of the optimal number of turbines for the site showing flow velocity magnitude and streamlines for 'ebb' tide (flow right to left).

broadly perpendicular to flow entering the turbine area during the flood phase, but also perpendicular to flow leaving the turbine area during the ebb phase. Acceleration of the flow around the sides of these mini-barrages is exploited (where

the position of turbine area boundaries allows) by downstream barrages.

The proposed array consists of 90 turbines, which in the optimised layout extract 32.59 TJ over the tidal cycle. This number of turbines is broadly consistent with the array size proposed by MeyGen for phase one of their development of the Inner Sound – on which this example is based. Clearly this example is based upon simplified modelling assumptions (for example the idealised financial model), however these assumptions are grounded in reality and the similarity is encouraging.

6.7 Conclusions

In this work, a methodology for the optimisation of tidal turbine array design has been presented which considers both the number of turbines in the array (array size) and the locations of those turbines relative to one another and the domain (micro-siting design). Both of these design parameters are strongly coupled to the regional and local hydrodynamics to which the power extracted by the array is highly sensitive (and thus are any prognostic variables based on it).

Previous work has demonstrated the strong impact turbine micro-siting design can have on the performance of the array, and has introduced *OpenTidalFarm* as a tool to maximise some desirable quantity through 'tuning' the array to the site through manipulation of turbine locations. This approach has shown utility whether the array performance is measured in terms of the energy extracted by the array [Funke et al., 2014], the financial performance of the array [Culley et al., 2016a], or indeed by some other barometer. However, until a given array has been tuned to the site, the designer does not know how that array will perform (because of the strong coupling of hydrodynamics to design parameters and the high sensitivity of power to them), i.e. for a given array size the performance of that array is unknown until its micro-siting design has been optimised. Consequently, for a given site and conditions, finding the array size which maximises the performance indicator is an 'outer' optimisation for which each sample point investigated requires an 'inner' micro-siting optimisation provided by *OpenTidalFarm*.

Due to the computational cost of performing a full micro-siting optimisation at each array size sample point taken, a methodology suited to the optimisation of expensive black-box functions is required – where the micro-siting optimisation performed by *OpenTidalFarm* is treated as the black-box – an unknown non-linear solver which for a given array size determines the optimal performance of that ar-

ray and the micro-siting design which achieves this. This work demonstrates the utility of a surrogate-based approach which uses a Bayesian methodology combining prior-assumed knowledge of the function of interest with a training set of recorded observations, to achieve a posterior expectation of how unobserved array sizes may be capable of performing. This approach minimises the number of observations which are required by using Gaussian process regression to maximise the utility of the observations.

This approach was used to optimise the array size and micro-siting design for two test examples; the first being a simple channel with a centrally located circular island and the second being loosely based on the Inner Sound of the Pentland Firth, in the Scottish Orkney Islands. In both cases, turbines were parametrised with realistic thrust and power curves and optimised for a full flood-ebb cycle. In the first case, the array size was optimised in four low fidelity and five high fidelity *OpenTidalFarm* runs, an array size of 80 turbines was returned as optimal to maximise the financial profitability of the project, and an array design of those turbines was proposed to maximise the power extracted by the array. In the Inner Sound example, the array size was likewise optimised for financial profitability and in four low and five high fidelity *OpenTidalFarm* runs an optimal array size of 90 turbines was found.

Chapter 7

Uncertainty estimation

Contents

7.1 Introduction 163
7.2 Motivation 164
7.3 Studying output sensitivities to parameters using the adjoint . . 166
7.4 Bayesian methods 171
7.5 Complexity and uncertainty quantification 174
7.6 Optimisation under uncertainty 177

Abstract

In tidal stream array modelling, the quantification of uncertainty is infrequently discussed beyond the requirements of model validation. The bulk of this thesis has dealt with the optimisation of different aspects of tidal stream turbine array design – and integrating those approaches such that they may be used cohesively together. There is, however, significant cross-over between optimisation and the quantification of uncertainty – in that both are concerned with mapping system response to parameter variability. In this chapter we discuss and explore the cross-over. Some new work is presented – looking at the sensitivity of objectives with respect to various model parameters (such as bathymetry and bed friction). However, much of the chapter comprises high-level discussion of an important direction for future work; considering uncertainty quantification techniques which may be employed to better understand how uncertainty in modelling assumptions and variability in parameters propagates through to the objective of interest.

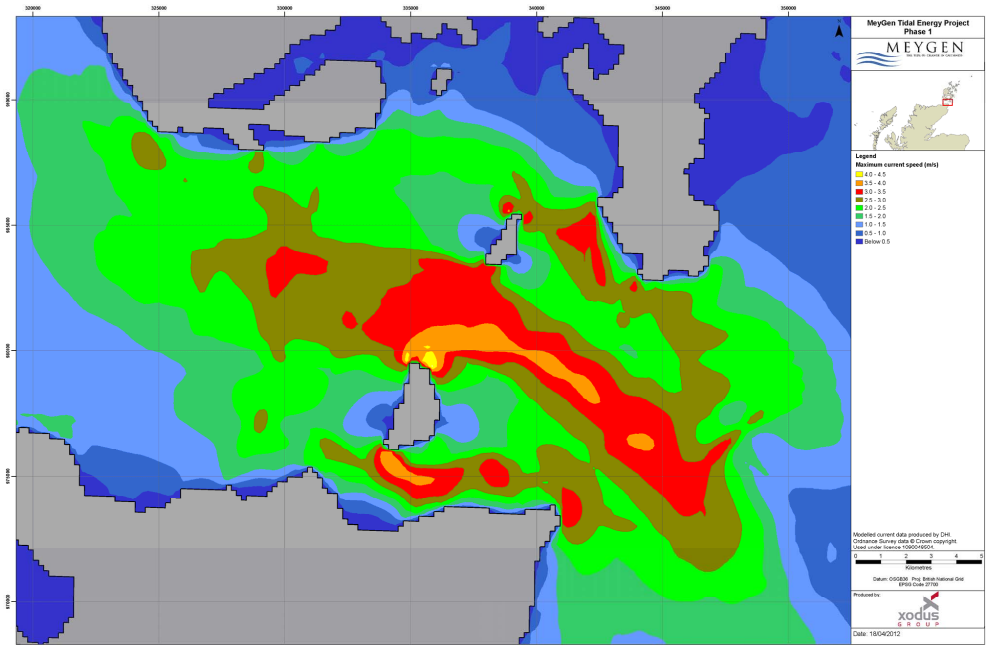


Figure 7.1: Plot showing mean peak flow speed in the Inner Sound of the Pentland Firth. [MeyGen, 2014].

7.1 Introduction

In complex and dynamical systems, in which non-linear and multi-scale processes interact, uncertainty can exist in many forms. Often uncertainty is introduced through the presence of randomness (for example turbulent eddies or storm-wave induced currents) that is inherent to the physical processes being modelled, and is impossible to describe in a deterministic way. While sometimes it arises because there is an intrinsic lack of knowledge relating to either the process or the parameters being modelled.

One of the shortcomings of computational modelling for the tidal stream energy industry is that, while the outputs that are being provided to decision-makers are fairly simple (for example a colour plot showing spring tidal velocities over the domain, figure 7.1) there may be great uncertainty in the information and data from which they were generated. Compounding this, computational modelling is often an intricate procedure, and those who develop and run the models are generally not the same as those who use the results to make project decisions. Uncertainty in the inputs, randomness in the physical processes and limitations of the models (either from wrong / incomplete equations or discretisation errors) are often not communicated alongside the results, meaning that it is all too easy

for this uncertainty to be (inadvertently) neglected and for it to go unaccounted for. Simply put, it is imperative that results which are to be used to determine the viability of – and make design decisions on – a multi-million pound industrial project are presented with appropriate 'error-bars'.

Fortunately, the parallels between optimisation and uncertainty quantification are strong – meaning that study of the former leads naturally on to consideration of the latter: The process of optimisation is concerned with exploring the influence of variance in a *design* parameter on a chosen objective (albeit with the goal of maxi-/minimising that objective) whereas uncertainty quantification is concerned with exploring the impact of the variance in an uncertain *input* parameter on the objective. In the parlance of the previous chapter, the former is about exploitation of the function, while the latter is about exploration of the function.

Three key concepts which have been explored thus far in this thesis have direct application to the quantification of uncertainty – use of Bayesian approaches to produce non-parametric surrogate models, the concept of 'balancing complexity' in fidelity of model choice and finally use of the adjoint method to provide sensitivity information. While sensitivity analysis is not technically a part of uncertainty quantification, it can feed into it and be an invaluable tool in identifying where uncertainty might exist, or where it may impact the performance of the design in question. Consequently, for the purposes of this discussion we shall consider it to be under the umbrella of uncertainty quantification.

In the following section, the importance of uncertainty quantification will be presented – motivating the discussions of the rest of the chapter. These discussions will revolve around application of the concepts detailed above and in the previous chapter to the uncertainty quantification problem. Towards the end of the chapter will be some discussion of the more advanced, concept of optimisation under uncertainty (OUU). This latter section represents aspirations for future development of this work.

7.2 Motivation

As best practice in approaches to modelling and design for tidal energy extraction applications becomes more developed, the quantification and mitigation of uncertainty is becoming an area of increasing interest. In IEC [2015] – a recent technical standards report on tidal energy resource assessment – there is an em-

phasis on uncertainty quantification (UQ) that has not been given such prominence in past specifications of best practice. In Catapult [2015] – a report reviewing current research on tidal energy yield uncertainty – the uncertainties were subdivided into four categories: Site measurement, temporal variation, spatial variation and plant performance. However, the report and the research it reviewed focused more on adoption of best practice to mitigate the impact of uncertainty – rather than the development of methodologies for its explicit quantification.

This limitation has a lot to do with "a lack of field and model data for a statistically significant number of sites" [IEC, 2015], but also (from a computational viewpoint) the additional complexity of adding uncertainty quantification on to an already computationally taxing modelling problem. The notion of the complexity of a problem modelled with UQ can be formalised in a manner similar to that presented in chapter 2 where problem complexity with optimisation was considered. Extending the problem to design optimisation with reference to this uncertainty adds an additional dimension to the complexity.

The purpose of computational models in engineering design is to provide virtual prototyping for problems which cannot, physically, economically or feasibly, be prototyped in reality. However, in all engineering design problems, there is uncertainty in the characterisation, measurement and modelling of input and design parameters and inevitable variability in these parameters once the design has been deployed in the real world. This variability and uncertainty may cause significant changes in the performance and reliability of the finished design and so, if good account is not made of them, the efficacy of the model is reduced and the fundamental motivation for virtual prototyping undermined.

Traditional approaches to parameter optimisation for system design utilise deterministic models and characterise uncertainty through the use of mean / worst case values, or the application of safety factors [Schuëller and Jensen, 2008]. While this may have the result that finished system designs are inoculated against variability and uncertainty, such approaches are inherently conservative and almost inevitably over-design the system, which has cost implications.

One of the key driving forces behind the quantification of uncertainty is to improve investment into the industry. With high uncertainty comes greater financial risk for investment, which is passed on to the industry through increased cost of borrowing. The Offshore Renewable Energy Catapult report quoted investors' concerns about unquantified (which is therefore treated as high) risk in investment in the sector as a major barrier to main-stream deployment [Catapult, 2015]. Since

debt, fundamentally, is a claim against future wealth, then immunising that wealth against variability, risk and uncertainty serves to reduce to cost of borrowing which will promote development and growth in this nascent industry. Reducing (or at least identifying) uncertainty surrounding that future wealth increases the accuracy of financial modelling and the predictability of the future financial returns.

Perhaps less nebulous than the *fiscal* ramifications of uncertainty is the *engineering* impact of uncertainty. Design for durability is fundamentally costly and uncertainties in loading and operating conditions are reflected through the specification of design safety factors. These are used to scale up the design loads to ensure that the product is robust to varying conditions. High safety factors lead to over-engineering which requires more material, more resilient components and generally higher specification at all levels (from turbine foundation construction to lifetime array-maintenance strategy).

7.3 Studying output sensitivities to parameters using the adjoint

Tools such as sensitivity analysis are invaluable to determine which model parameters most strongly impact upon the performance metrics of the finished design. Computationally economical access to gradient data means that *OpenTidalFarm* already offers this capability [Culley et al., 2015]. In the case of the array design problem, differentiating with respect to bathymetry provides critical insight into which areas of the domain the depth (or uncertainty over the depth) may be affecting the value of the objective, and therefore allows for potential additional, targeted physical surveying of the proposed site. This facilitates 'adaptive observation' as parameters to which the performance metrics are most sensitive are identified and may be re-surveyed at higher fidelity.

As has been seen throughout this thesis, a key benefit of the adjoint method is that the gradient of the objective is computed at a computational cost which is essentially independent of the size of the control variable. In the application of optimising turbine positions, this means that the expense of computing the gradient is independent of the number of turbines. However, the gradient of the objective may also be computed with respect to control variables with many more degrees of freedom, for example the bathymetry, H, bed roughness, c_b, or viscosity, ν, which vary over the whole domain.

Recall from (3.22) the definition of the adjoint variable as

$$\lambda^T = J_{\mathbf{u}} F_{\mathbf{u}}^{-1} \qquad or \qquad \lambda = F_{\mathbf{u}}^{-T} J_{\mathbf{u}}^T, \tag{7.1}$$

so the gradient with respect to, for example the viscosity, bottom friction or depth can be expressed as

$$\frac{dJ}{d(\cdot)} = -\lambda^T F_{(\cdot)} + J_{(\cdot)}, \tag{7.2}$$

where $(\cdot)$ is substituted for ν, c_b or H respectively.

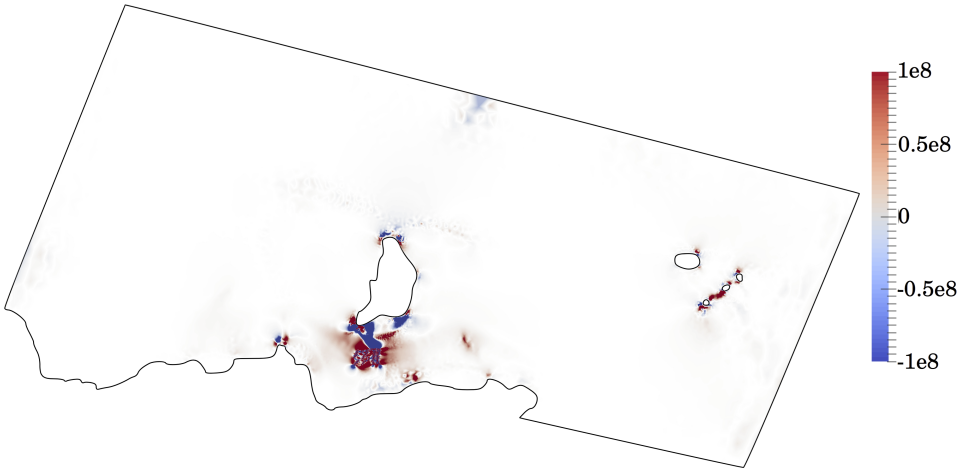


Figure 7.2: Sensitivity of array power with respect to bathymetry over the domain. Positive values indicate that if the depth in that area were deeper than modelled, the power extracted would be larger than modelled, conversely negative values imply that a decreased depth would result in increased the power

Recall the example from the previous chapter which found the optimal array design for a turbine area to the south of the island of Stroma in the Pentland Firth. Having obtained the array size and micro-siting design, we analyse the sensitivity of the performance of that array to the bathymetry and bottom friction over the domain.

In figure 7.2 the sensitivity of power extracted by the array is shown with respect to the water depth, H. Positive values indicate that increased depth would increase the power extracted by the array. The general trend is, as one might expect, showing that increased depth may be used to the east and west of the site to 'funnel' flow towards the turbine area, and that reduced depth to the north and south of the site adds to this effect and makes it more difficult for flow to escape

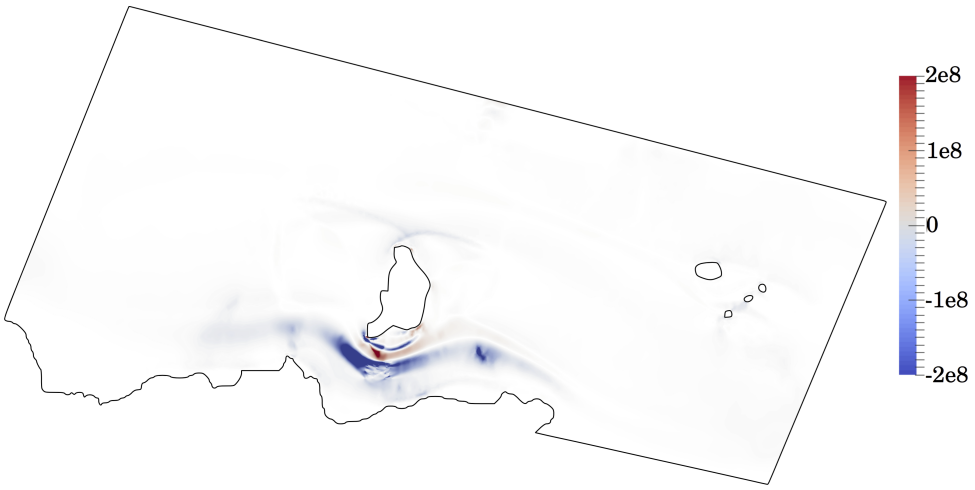


Figure 7.3: Plot showing sensitivity of the power extracted by the array with respect to the bottom friction over the domain. Positive values of $\mathrm{d}P/\mathrm{d}c_b$ indicate that increasing the bottom friction will increase the power, negative values that decreasing the bottom friction will increase the power.

from the turbine area as the stream-lines diverge. Additionally, to the north of the island, decreased depth would encourage flow through the Inner Sound where the turbines are located, thus increasing the power extracted by the array.

An analogous trend is shown in figure 7.3 which shows the sensitivity of power with respect to bottom friction. Once again, positive values indicate that increased friction at an area would result in increased power, and negative values that decreased friction will increase power. Decreasing friction in a channel running in the dominant ebb/flow direction clearly increases the flow velocity through the site, increasing power.

It should be noted that sensitivities only provide a 'snapshot' of the underlying function – The linearisation is only true for small perturbations around the current state: If the underlying function were highly non-linear it is likely that small changes to the input parameters (for example through updating bathymetry data from a higher resolution source) may drastically change both the objective and the sensitivity map in an unexpected way. Thus the development of a suite of more comprehensive uncertainty quantification (UQ) techniques is important.

7.3.1 Model boundary conditions

The capability to directly differentiate the objective with respect to the boundary conditions would be an invaluable tool to determine the impact of forcing as-

sumptions – and give a clearer idea as to how cyclical variation in tidal current magnitudes may affect the project – by computing the array's sensitivity to inflow conditions.

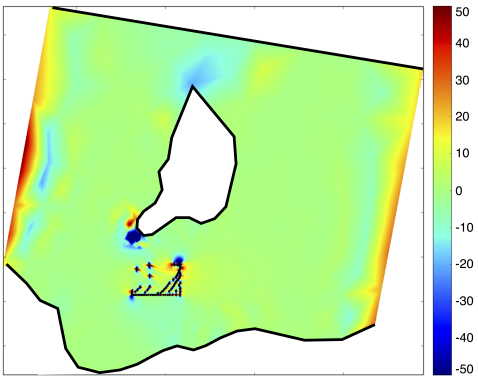


Figure 7.4: Plot showing $\mathrm{d}P/\mathrm{d}H$ for simplified Inner Sound domain. Note the high sensitivities at the inflow/outflows – implying these boundaries are placed too close to the turbine area.

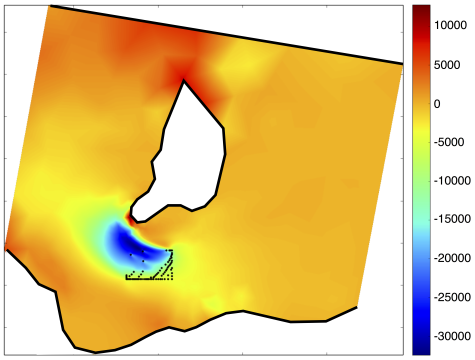


Figure 7.5: Plot showing $\mathrm{d}P/\mathrm{d}c_b$ for simplified Inner Sound domain.

While this capability is not yet fully developed, some idea of the suitability of boundary location can be inferred with the current tools. Boundary location is critical to ensuring model accuracy [IEC, 2015]. A clear example of this is given in figures 7.4 and 7.5 from Culley et al. [2015]. These figures shows the sensitivity of the power extracted by an optimised (both size and turbine location) array located in the Inner Sound of the Pentland Firth (UK) to the water depth and bottom friction. This can be contrasted with figures 7.2 and 7.3 on an extended domain which have already been discussed. While these models were run with different problem parameters (for example viscosity) the point being illustrated is the high sensitivity of the power in the smaller domain to the depth immediately next to the inlet and outlet boundaries. For both boundaries, the sensitivity map implies that an increase in depth immediately inside the boundaries would increase the power. The physical reason for this seems clear – the smaller model is forced by a sinusoidal *velocity* profile; a greater depth at the inflow would mean a greater volumetric flow rate through the domain; meaning more water is forced through the Inner Sound in a given time requiring a high flow velocity through the turbine site.

This – albeit crude – measure of the impact of the proximity of boundaries on the objective is important. Boundary location is often judged more by the modeller's intuition informed by experience, the location of key geographic features (like a continental shelf) or the limits of the available computational budget (as

it was for the Culley et al. [2015] example). Tools like this help to provide a more formal validation of boundary choice, but should be used with caution given their crudeness.

7.3.2 Data assimilation

Another key use of the adjoint, is for data assimilation. This is a technique that is used in the meteorological modelling and forecasting industry [Courtier et al., 1993], but (to the author's knowledge) has not been incorporated into a coastal scale models used for tidal stream energy extraction modelling and design applications. Models are generally calibrated using a hindcasting approach – the model is run for a period in history for which data (either from velocity profilers or surface height gauges) is available. The modelled surface height or flow velocity at the location of the data is then compared to the data itself, and model inputs such as boundary conditions, bottom friction, finite element grid and turbulence parameters are adjusted in an effort to improve the agreement.

The idea behind variational data assimilation is much the same – but automated. The difference between observed and modelled parameters is used as the objective functional, and parameters (such as bottom friction) are used as control variables. A (gradient-based) optimisation approach may then be used to minimise the difference by adjusting the control parameter.

Not only does this have the benefit of improving the accuracy of the model to match the observed data, but also – potentially – to 'improve' patchy input parameters when real-world data could not be obtained. For example, in locations where it is not known in high resolution, if the bathymetry were optimised to calibrate the model and this bathymetry then tested in *another* model – to see if this model's accuracy is also improved, it may be an economical and pragmatic method for improving input data when expensive surveying of the unknown areas is not possible. Of course in reality one should realise that some of this calibration will simply be accounting for biases in the models.

While much of the technical infrastructure to achieve this is in place, one of the main limitations to deployment of this methodology is the poor quality of the data: ADCP deployments are often not concurrent; one or two devices are left in locations for a month at a time – then moved to a new location – meaning that the velocity data for the site is strung out over a period of many months. Another technique for deploying ADCPs is to mount them to boats and measure from the

surface downwards as the boat navigates transects of the flow. This, of course, means that the data's spatial location is temporally varying.

Problems with data quality notwithstanding, the key technical issue is that of optimising a control parameter with many degrees of freedom (e.g. the bathymetry over the domain will have a value at each mesh node) constrained by data from a very small number of points (e.g. a handful of profiler deployments) is always going to be problematic due to the existence of numerous potential solutions to this under-constrained problem. For example, by manipulating the bathymetry in immediate proximity to the measured data points, the modelled error can be minimised – mathematically this is a valid solution, but clearly this isn't a meaningful result in engineering terms. Thus *regularisation* is required which, for example, enforces limits on the spatial gradient of the control parameter to curb this tendency for extreme local adjustment. This is probably the most challenging part of variational data assimilation in practice.

7.4 Bayesian methods

In chapter 6 Bayesian methods were used to optimise the number of turbines comprising a tidal array on a given site. Recall that this optimisation proceeded in two modes; 'exploration' and 'exploitation'. In the former, the general trend of the objective over the range of the number of turbines, N, was established. In the latter, the vicinity of the optimum was sampled in higher fidelity to establish the optimal value of N. The first 'exploration' mode can be thought of in similar terms to uncertainty quantification – where N is an uncertain input parameter to which we wished to judge the response of the objective. As compared to sensitivity analysis which provides a snapshot of the derivative at a particular state, uncertainty quantification (in this case using Bayesian methods) provides a more complete picture – the response of the objective over the range of the input uncertainty.

Such an approach to UQ carries all the benefits of Bayesian optimisation laid out in chapter 5 – most notably that the models can be constructed from variable fidelity samples and are non-parametric. However, just as optimisation requires re-evaluation of the analysis code as we iterate towards maximum exploitation of the control variable, so too does uncertainty quantification require re-evaluation of the analysis code to iterate through different model parameter values towards exploration of that parameter (as varying the dimension of each parameter requires additional samples to be taken). Unlike in the adjoint method, where the sensitiv-

ity of the objective is returned at a cost that is essentially independent of the size of the parameter with respect to which it is being computed, Bayesian methods – for example Gaussian process regression (GPR) scale super-linearly. Consequently, one significant limitation is poor scaling properties with the dimension of the uncertain parameter.

One reason for this is that there are computational bottlenecks to traditional implementation of GPR – for example storage of the training set ($\mathcal{O}(n^2)$ memory, where n is the size of the training set) Keane and Nair [2005]. Of course, the return from this computational investment is much higher than when using the adjoint – rather than a just a linearisation of the sensitivity for a given input-parameter value, Bayesian methods provide a full view of the variation of the function to the whole range of values that the variable parameter may take.

To illustrate the effectiveness of using GPR for UQ, consider the domain from the example problem used in chapter 5 and the previous section. Here two zones have been selected within the domain which represent areas in which the bathymetry is uncertain (figure 7.6) – and may vary from the previously modelled depth by $\pm 10\%$. This can be modelled by multiplying the bathymetry function ($H : \Omega \to \mathbb{R}$) by a function which is equal to 1 over the domain but steps to a uniformly higher or lower factor in the region of interest. This step is modelled using a hyperbolic tangent function to smooth the transition over 0.5 km to prevent sudden changes in bathymetry which would result in changes to the hydrodynamics due to the step rather than due to the adjusted bathymetry.

Nine equally spaced samples were taken (as shown in figure 7.7 and GPR was used to construct the surface showing how the energy extracted by the array responded to the variation of the depth in the two zones. This is shown in three-dimensions in figure 7.7 with the 95 % confidence intervals shown as translucent surfaces, or in two dimensions in figure 7.8. The results are as expected – that increasing the depth in areas in which the objective is sensitive to it, will result in the power extracted by the array being affected as indicated by the sensitivity study.

Figure 7.8 shows that the worst-case scenario is for the bathymetry in the positive sensitivity zone to be 10 % shallower than previously modelled while the bathymetry in the negative sensitivity zone is 10 % deeper. Such a scenario would result in a circa 30 % reduction in the power extracted by the array.

If data is known on the probability distribution of the bathymetry in those locations (i.e. if the probability density function of the bathymetry is known – rather

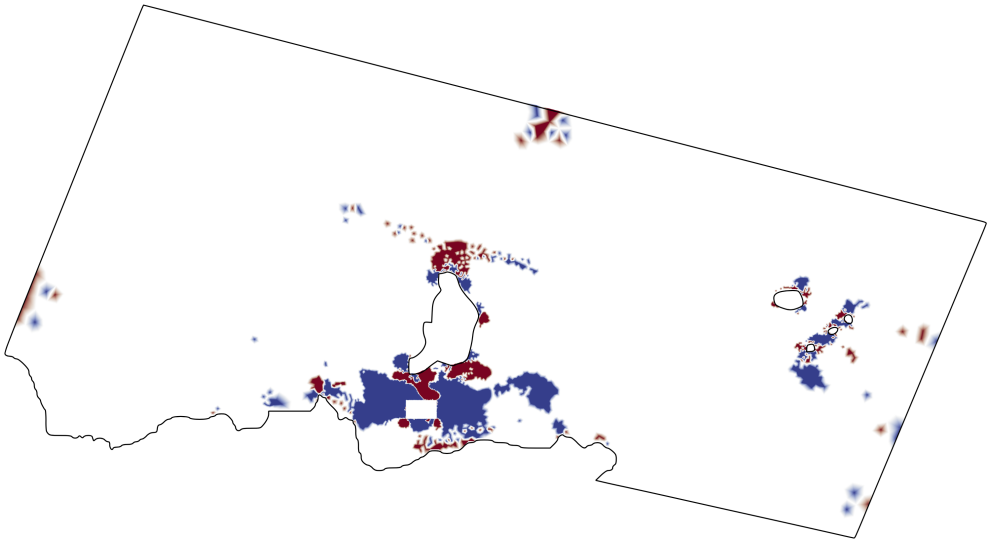


Figure 7.6: Pentland Firth problem domain showing the regions to whose bathymetry the power extracted by the turbine array is most sensitive. Red zones are those in which the gradient implies that a 1 m increase in bathymetry would result in at least a 10 kJ per tidal cycle increase in energy extracted, the blue zones a decrease.

than just a range of values as assumed in this example) then the liklihood of the reduction in power could be computed and decisions made on the basis of whether such a scenario warrants mitigation.

7.4.1 Other approximation techniques

The surrogate modelling approach discussed here and in the preceding chapter was based on creating an input-output mapping using training data but without any reference to the underlying physics. Other, physics-based, approximation techniques exist and play an important part in design optimisation and UQ in other (for example aerospace) industries. While these will not be discussed in detail or employed in earnest in this thesis they warrant acknowledgement for completeness.

The biggest family of such physics-based approximation technique are *reduced basis methods*. The aim is to reduce the number of unknowns in the original physical equations so as to produce a model which is correspondingly simpler and computationally cheaper to solve. Various flavours of reduced basis methods exist depending on the problem, *Proper Orthogonal Decomposition* (POD) and *Reduced Order Modelling* (ROM) are two relatively well known examples. The interested reader is directed to Keane and Nair [2005] for a more complete exposition on

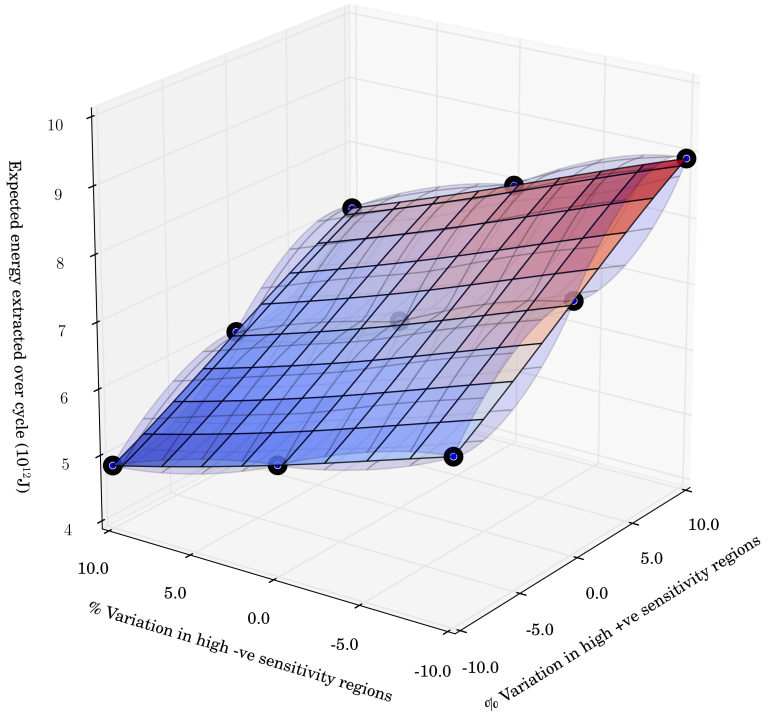


Figure 7.7: Plot of functional response to uncertainty in depth in the Pentland Firth. The points show sample locations and translucent surfaces the bound of the 90% confidence interval.

physics based approximations.

7.5 Complexity and uncertainty quantification

The previous two sections have demonstrated how uncertainties in data (that used for intialisation, parametrisation and evaluation of models) can be handled using adjoint and Bayesian methodologies. Some of the more difficult uncertainties to quantify are those related to the structure of the model itself and its adequacy for representing the processes in which the modeller is interested.

The problem complexity of modelling with uncertainty can be visualised as a

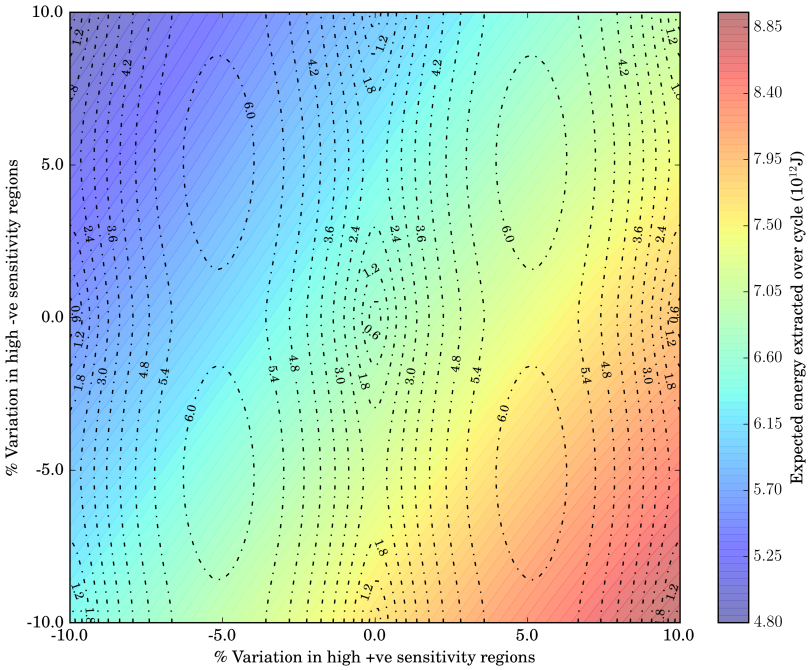


Figure 7.8: Plot of functional response to uncertainty in depth in the Pentland Firth. The dot-dashed lines show the percentage difference of the bound of the 90% confidence interval.

balance between 3 variables; the complexity of the model, the complexity of the scenario and the complexity of the uncertainty quantification methodology. As with optimisation the goal is to achieve 'balanced complexity'. For example, higher spatial resolution allows use of better turbulence parametrisation and thus for turbulent structures to, more accurately, be modelled. However, to maintain computational tractability this likely comes at the cost of reducing the size of the domain in question – so the condition of the flow as it approaches the turbine, for example, is likely to be less realistic than when modelling a larger domain (where it has had more opportunity to develop from inlet). In this scenario, which model is more 'accurate'? The answer, of course, is that it depends on the hypothesis the modeller is testing, i.e. the question the modeller is trying to answer.

Similarly to section 2, this trade-off can be represented graphically as in figure 7.9 – where the front shown in red once again represents the limits imposed by the available computational budget. Organising models and methods in this way highlights the importance of inter-comparison of similar models, which is vi-

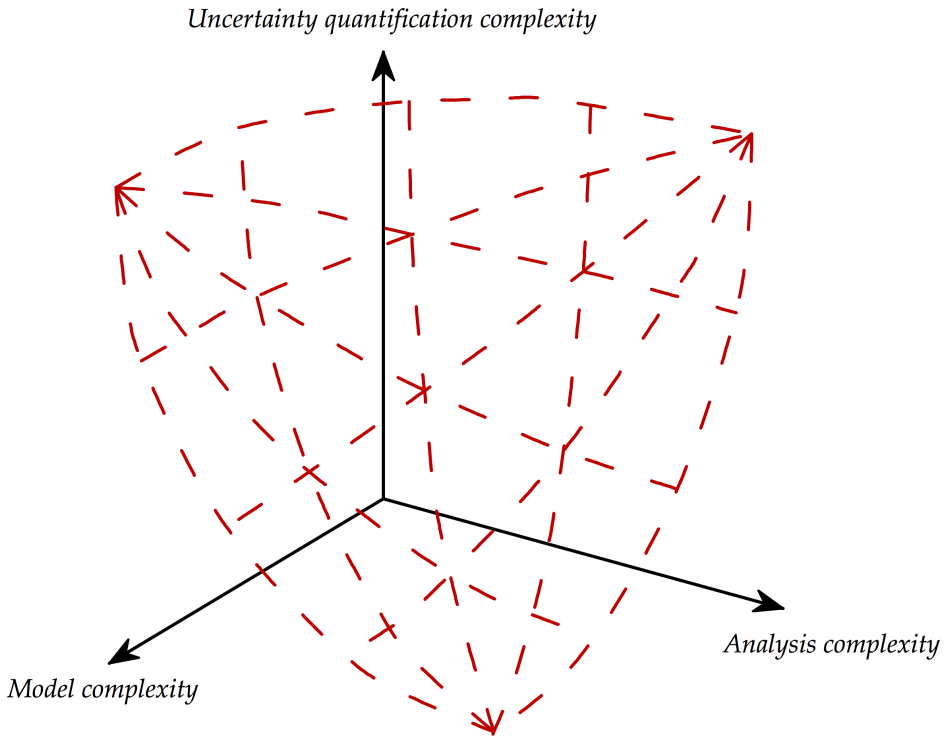


Figure 7.9: Illustrative graph of problem complexity as a function of model, analysis and uncertainty quantification complexity with limits of computational budget shown by red surface

tal in understanding the model's impact on uncertainty. This is especially true when (as in the tidal array case) real-world validation data does not exist, or is not available. By answering questions like whether a finite-volume representation of a given phenomenon performs differently to a finite-element discretisation, the underlying structural uncertainty inherent to a modelling approach may, to some extent, be understood.

With inter-model comparisons, it should be recognised that the variation in results may be biased by the shared heredity or shared formulation of the models being compared [Smith et al., 2014]. The goal, however, is to determine the prominence in variation due to hardwear and software differences as compared to that arising from changing parameters.

In some respects, using a model based on the *FEniCS* framework as is *OpenTidalFarm*, is ideal for these kind of comparison studies. *FEniCS* supports a wide range of finite element discretisations and the light-weight design of *OpenTidal-*

Farm means that it is not overly involved to implement and compare alternate schemes – for example $\mathscr{P}1_{\text{bubble}}\mathscr{P}1$ of even discontinuous Galerkin schemes. Even more straightforward is testing performance on different hardware architectures. As the low-level code (which is generated at run-time) is optimised for the hardware at hand, obtaining fair comparisons between machines is as straightforward as installing the software and running the model. Perhaps more interesting than testing different discretisations of the same model, the *FEniCS* framework allows relatively straightforward implementation of PDEs so is an ideal platform for testing other models (e.g. 2D versus 3D, hydrostatic versus non-hydrostatic etc.).

The ultimate goal of this kind of work is that by more robustly modelling the uncertainty in a system it is easier to more fully understand where that uncertainty arises. In the case of parameter uncertainty this means that surveying effort can be directed to the locations in the domain that are most important. But in general this means that scientific and developmental effort is properly targeted at reducing unknowns; is there true benefit to developing more sophisticated turbulence modelling capabilities if the accuracy of the overall model is limited by some other factor?

7.6 Optimisation under uncertainty

As has been discussed, the quantification of uncertainty naturally leads on to optimisation under uncertainty (OUU). This can fundamentally change the approach of developers to project specification and design at all levels (from array placement, to array design, to device design). While so far in this thesis, the concern has been to find the *optimal* design point, in OUU the goal is to find a *robust* optimal design point – one that is inherently immune to variability in the parameters and modelling choices. In site-selection this may demonstrate that a less energetic site with lower risk may represent a more desirable investment choice than a highly energetic but highly risky one. Without quantification of risk, in those scenarios the designer will inexorably be drawn to the highly energetic site and may ultimately be disappointed.

The concept of optimisation under uncertainty is one that is familiar to sectors involved in the design of safety critical systems (such as nuclear reactors) but only recently is being applied more broadly. As the name suggests, this is a design paradigm in which the uncertainty is incorporated into the design optimisation (rather than merely audited afterwards), and therefore the finished design will be

inherently immune to it.

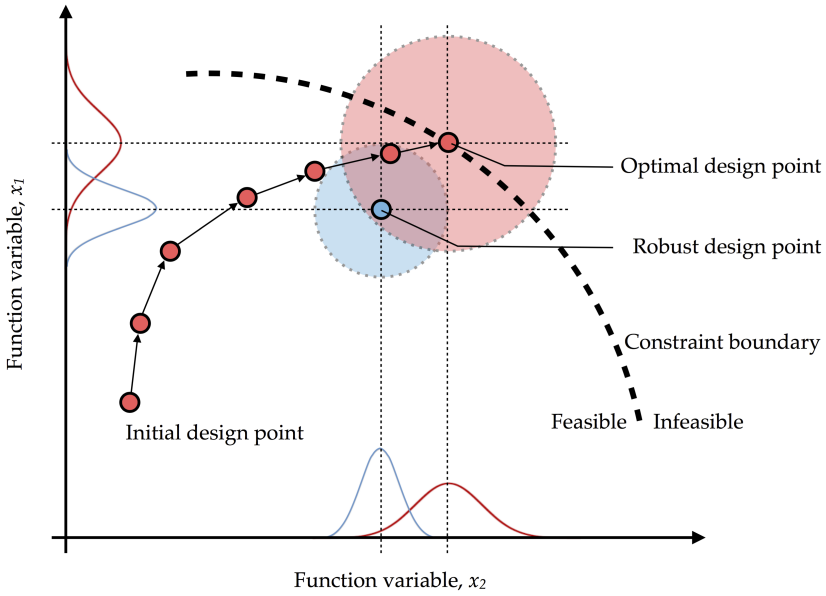


Figure 7.10: Sketch plot illustrating the difference between traditional optimisation (red) and robust optimisation (blue) from an initial design point.

Put simply, optimisation under uncertainty (OUU) is concerned with determining how the design variables can be adjusted to mitigate the impact of potentially unfavourable configurations of model variables. This is shown in figure 7.10 which sketches out the process of optimising a function $f(x_1, x_2)$. The traditional 'optimal design point' lies on the constraint boundary but the variance in the two parameters means that there is a high probability that in application it will turn out to be infeasible. Conversely there is high confidence that the *robust design point* will be feasible and the reduced variance in the probability distribution of this point means designers and investors can have high confidence that the implemented design will perform closely to how expected.

True optimisation under uncertainty approaches are significantly more involved than their deterministic analogues. Different types of uncertainty must be treated in different ways, for example depending whether they can be treated probabilistically or not, and whether they are uncertainties in input parameters (such as the bathymetry) or design parameters (such as installation tolerances). The result is that, as with optimisation, a library of efficient optimisation techniques and approximation concepts are required to measure the uncertainty in question and propagate it through the optimisation while accounting for and minimising its effect on the performance metric. These problems often appear as a multi-objective

decision problem – for example finding a trade-off between maximal performance and minimal variance, perhaps with an added constraint that the solution lies within the feasible design space to a certain degree of certainty.

Chapter 8

Summary and future work

Contents

8.1 Summary 180

8.2 Future work 182

8.1 Summary

The objectives of this thesis were to develop tools and processes which promoted a practical and more holistic approach to the design of tidal turbine arrays. In the survey of models and methods, it was identified that while the sector is rich in techniques for the *analysis* of tidal stream power scenarios at various scales, there is a significant deficit in design / optimisation tools. This was formalised in a three-dimensional plot of existing models which demonstrated the large open space that the research in this thesis aims to fill.

OpenTidalFarm was introduced in chapter 3 as a practical method for determining the optimal micro-siting design of tidal turbine arrays. The method by which turbines were parametrised within the shallow-water model so as to reflect the variable thrust characteristics of the real devices was laid out in detail.

In chapter 4, *OpenTidalFarm* was compared against an alternate array optimisation methodology based on the approach used by the commercial code *TidalFarmer*. This uses global optimisation algorithms to improve turbine placement. These require a large number of iterations meaning that to ensure the overall problem complexity remains tractable, the fidelity of the evaluation at each iteration

must be reduced. In contrast, *OpenTidalFarm*'s access to gradient information via the adjoint model facilitates gradient-based optimisation which requires fewer iterations to converge, enabling each iteration to be computed with a higher fidelity (and hence more realistic) analysis code. The global-optimisation approach is limited by the low fidelity of the analysis code – which in this case neglects coupling between turbine and flow. The *OpenTidalFarm* approach is limited in that it only finds local optima and is potentially sensitive to the starting 'guess' of turbine locations. The chapter concluded that using the approaches in concert was an effective way to mitigate the weaknesses inherent to each approach: Using the global optimisation as a fast method to find a near-optimal starting guess that the gradient-based *OpenTidalFarm* could build upon and improve.

The remainder of the thesis focused on improvements to the *OpenTidalFarm* framework, with the goal of making it a more practical tool for designers and project decision makers.

In chapter 5, a methodology for the inclusion of financial costs to the array design problem was considered. Subsea cable costs were chosen as they represent a significant proportion of project cost, but are also an interesting case for modelling (finding the optimal cable routing for a particular array layout is, itself, an optimisation problem). The multi-objective optimisation framework developed for this work has been adopted and extended in du Feu et al. [2016], which considered the impact of environmental constraints on turbine micro-siting design.

Chapter 6 looked at the more fundamental – but more computationally demanding – question of the optimal number of turbines that should comprise an array. This question of array size has received some analytical scrutiny but very little attention in terms of developing practical approaches: For an engineer with a clearly defined site with a specified turbine, the work in this chapter outlined a practical approach to determining the array size by using Bayesian optimisation techniques to minimise the number of calls to *OpenTidalFarm*.

In the previous chapter some preliminary work and some high-level thoughts were presented on the subject of sensitivity analysis and uncertainty quantification in tidal applications. It is well understood that optimisation and uncertainty quantification techniques share a common theoretical basis, and as such many of the themes and approaches from preceding chapters were re-examined in the context of uncertainty quantification. The concept of optimisation under uncertainty was presented as an interesting potential direction for future work.

8.2 Future work

8.2.1 More realistic test cases

The work in this thesis has focused on developing processes for array optimisation – and has had much less emphasis on producing analyses of particular situations. In explanation; when examples have been discussed they have been presented in the context of proof-of-concept of the underlying methodology rather than to draw hard conclusions on the situations under analysis – for example that "the Inner Sound array should comprise x number of turbines". There is a simple practical reason for this – that the requisite engineering and business data on which such assertions would be fundamentally grounded are highly proprietorial and not available for academic publication. Certainly the examples have been run under realistic assumptions, but the results could change significantly on the basis of those assumptions.

This lack of certain information fed into the design of the code – construction of which was the main effort in this thesis. Object orientated code design ensures that within the frameworks outlined in chapters 3, 4, 5 and 6 the implementations can be replaced or upgraded by the user to reflect more realistic assumptions or models. For example the thrust curve implementation outlined in chapter 2 allows the curve itself to be user specified and easily changed to reflect that of the turbine which is to be installed on site. The lack of access to proprietorial information has begot a more flexible code; necessity is the mother of invention.

A key direction for future work, however, is to work more closely with industry to improve the underlying assumptions regarding the turbines' performance characteristics and the array developers business strategy. This may lead to some interesting new (but related problems): One example may be phased deployment of an array – how many turbines should be installed in each phase and where should they be placed so that initial phases perform optimally (but without compromising latter phases) to ensure future investment.

8.2.2 Turbulence modelling

One of the major strengths of the *OpenTidalFarm* methodology is the intrinsic coupling between the turbines and the hydrodynamics. This means that, as compared to perturbation approaches (discussed in chapters 2 and 4) the hydro-

dynamics are modelled in a fundamentally more realistic fashion, and processes such as global and local blockage caused by turbines are represented rather than neglected.

There is, however, much that can be done to improve the hydrodynamic solver at the heart of *OpenTidalFarm* so that it better captures the important physical processes and thus produce better solutions to the array design problem. Most important is the inclusion of more sophisticated approaches to modelling turbulence. Turbine sites are chosen for their highly energetic (and therefore advection-dominated) flow regimes. As a consequence, turbulence is an intrinsic characteristic of the tidal turbine array's environment. The addition of the turbines themselves only make this environment more turbulent.

In the models and methods section it was demonstrated that the world of computational models for tidal power applications is wide and varied, with many methodologies including sophisticated approaches to turbulence modelling. The barrier to such models being included in the *OpenTidalFarm* framework is currently one more related to computational economy rather than computational formulation.

One promising avenue is the inclusion of an LES solver, towards which much progress has already been made. This is because, in a quasi two-dimensional setting, the energy cascade is inverted – i.e. energy tends to coalesce into larger structures. In comparison, in three-dimensions, where additional mesh resolution is required to resolve turbulence, the demands on the spatial discretisation of the two-dimensional case are much less. This may make inclusion in an optimisation framework a tractable proposition.

Two options stand out for implementation in *OpenTidalFarm*. The first is 2D LES – which has been implemented but not exhaustively tested and integrated into the framework. Recognising that in well mixed coastal systems a lot of the turbulent structures are large-scale depth-averaged, in 2D LES, 3D effects – such as backscatter – on the 2D turbulence are neglected. I.e. the vertical structure does not affect the horizontal turbulence (and therefore mixing). The second option is HLES, which assumes a vertical velocity profile and treats the vertical mixing via a Reynolds-Averaged approach and allows its influence to contribute to the plan-view turbulence modelled via LES.

HLES has received surprisingly little attention within the tidal array modelling community, with the exception of Mungar [2014], who demonstrated its clear applicability to the tidal turbine modelling problem using the implementation in the

two-dimensional version of Delft3D.

8.2.3 Operational optimisation

For a fixed array size, this thesis has focused on the optimisation of turbine location alone. However, the *OpenTidalFarm* framework is not limited to this. The turbine diameter or amount of friction by which it is represented could also be used as control parameters for an optimisation. Here, since amount of friction controls the thrust that the turbine exerts on the flow, there is the opportunity to optimise the thrust characteristics of each individual turbine over the course of the tidal period to maximise the yield of the whole array. In a physical senses, blade feathering and generator gearing can change a turbine's hydrodynamic characteristics. Reduction of the thrust applied by an upstream turbine may benefit those downstream and promote higher energy extraction throughout the array.

Rather than the control vector representing the location coordinates of each turbine, it would instead contain the temporal thrust value for each turbine over the period. This vector would then be optimised in the same manner as the turbine locations (or even in concert with the turbine locations) to maximise power (or indeed any other implemented functional).

Bibliography

M.A. Abolghasemi, M.D. Piggott, A.V. Spinneken, C.J. Cotter, and S. Crammond. Simulating tidal turbines with multi-scale mesh optimisation techniques. *Submitted to Journal of Fluids and Structures*, 2016.

T. A. A. Adcock, S. Draper, and T. Nishino. Tidal power generation – a review of hydrodynamic modelling. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 2015.

T.A.A. Adcock, A.G.L. Borthwick, and G.T. Houlsby. The open boundary problem in tidal basin modelling with energy extraction. In *Proceedings of the 9th European Wave and Tidal Energy Conference. Southampton, UK*, 2011.

T.A.A. Adcock, S. Draper, G.T. Houlsby, A.G.L. Borthwick, and S. Serhadlıoğlu. The available power from tidal stream turbines in the Pentland Firth. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 469(2157), 2013.

R. Ahmadian and R.A. Falconer. Assessment of array shape of tidal stream turbines on hydro-environmental impacts and power output. *Renewable Energy*, 44:318–327, August 2012.

G. Allan, M. Gilmartin, P. McGregor, and K. Swales. Levelised costs of Wave and Tidal energy in the UK: Cost competitiveness and the importance of "banded" Renewables Obligation Certificates. *Energy Policy*, 39(1):23–39, January 2011.

AMCG. Fluidity manual. *Imperial College London*, November 2010.

M. Anderson and P. Stuart. Simplified solution to the eddy-viscosity wake model. *Renewable Energy Systems report*, July 2009.

U. Aytun Ozturk and B. A. Norman. Heuristic methods for wind energy conversion system positioning. *Electric Power Systems Research*, 70(3):179–185, August 2004.

P. Bachant, M. Wosnik, et al. Reynolds number dependence of cross-flow turbine performance and near-wake characteristics. In *Proceedings of the Marine Energy Technology Symposium, 2014, Seattle, USA*, 2014.

S. Balay, S. Abhyankar, M.F. Adams, J. Brown, Brune P., K. Buschelman, L. Dalcin, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L. McInnes, K. Rupp, B.F. Smith, S. Zampini, and H. Zhang. PETSc Users Manual. Technical Report ANL-95/11 - Revision 3.6, Argonne National Laboratory, 2015.

G.L. Barnett, S.W. Funke, and M.D. Piggott. Hybrid global-local optimisation algorithms for the layout design of tidal turbine arrays. *Submitted to Renewable Energy*, 2014.

W.M. Batten, M.E. Harrison, and A.S. Bahaj. Accuracy of the actuator disc-RANS approach for predicting the performance and wake of tidal turbines. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1985), 2013.

J. Bauer and J. Lysgaard. The offshore wind farm array cable layout problem: a planar open vehicle routing problem. *Journal of the Operational Research Society*, 2014.

T. Bektas. The multiple traveling salesman problem: an overview of formulations and solution procedures. *Omega*, 34(3):209–219, June 2006.

R. Bellman. Dynamic programming treatment of the travelling salesman problem. *Journal of the ACM (JACM)*, 9(1):61–63, 1962.

J. Blanchfield, C. Garrett, A. Rowe, and P. Wild. Tidal stream power resource assessment for masset sound, haida gwaii. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 222(5):485–492, 2008.

I.G. Bryden and S.J. Couch. ME1: marine energy extraction: tidal resource analysis. *Renewable Energy*, 31:133–139, February 2006.

I.G. Bryden, T. Grinsted, and G.T. Melville. Assessing the potential of a simple tidal channel to deliver useful energy. *Applied Ocean Research*, 26(5):198–204, 2004.

Z. Cai, J. Douglas, and M. Park. Development and analysis of higher order finite volume methods over rectangles for elliptic equations. *Advanced Computational Mathematics*, 19(1-3):3–33, 2003.

Carbon Trust. Accelerating marine energy: The potential for cost reduction - insights from the Carbon Trust Marine Energy Accelerator. *Carbon Trust*, 2011.

A.E. Carter and C.T. Ragsdale. A new approach to solving the multiple traveling salesperson problem using genetic algorithms. *European Journal of Operational Research*, 175(1):246–257, November 2006.

O.R.E. Catapult. Wave and Tidal Energy Yield Uncertainty. Technical report, Offshore Renewable Energy Catapult, 2015.

R.H. Charlier. A sleeper awakes: tidal current power. *Renewable and Sustainable Energy Reviews*, 7(6):515–529, December 2003.

G.U. Clarke and J.W. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581, 1964.

M. Claussen, L. Mysak, A. Weaver, M. Crucifix, T. Fichefet, M.F. Loutre, S. Weber, J. Alcamo, V. Alexeev, A. Berger, et al. Earth system models of intermediate complexity: closing the gap in the spectrum of climate system models. *Climate Dynamics*, 18(7):579–586, 2002.

P. Coker, J. Barlow, T. Cockerill, and D. Shipworth. Measuring significant variability characteristics: An assessment of three UK renewables. *Renewable Energy*, 53: 111–120, May 2013.

G. Cornelis van Kooten. Wind power: the economic impact of intermittency. *Letters in Spatial and Resource Sciences*, 3(1):1–17, November 2009.

P. Courtier, J. Derber, R. Errico, J.-F. Louis, and T. Vukićević. Important literature on the use of adjoint, variational methods and the kalman filter in meteorology. *Tellus A*, 45(5), 1993.

D.M. Culley, S.W. Funke, S.C. Kramer, and M.D. Piggott. A hierarchy of approaches for the design of tidal turbine arrays. *5th International Conference on Ocean Energy, ICOE, Halifax*, November 2014.

D.M. Culley, S.W. Funke, S.C. Kramer, and M.D. Piggott. Tidal stream resource assessment through the optimisation of array design with quantification of uncertainty. *Proceedings of the 11th European Wave and Tidal Energy Conference, EWTEC, Nantes*, 2015.

D.M. Culley, S.W. Funke, S.C. Kramer, and M.D. Piggott. Integration of cost modelling within the micro-siting design optimisation of tidal turbine arrays. *Renewable Energy*, 85(0):215 – 227, 2016a. doi: 10.1016/j.renene.2015.06.13.

D.M. Culley, S.W. Funke, S.C. Kramer, and M.D. Piggott. A surrogate-assisted approach for optimising the size of tidal turbine arrays. *Submitted to Journal of Marine Renewable Energy*, 2016b.

D.M. Culley, S.W. Funke, S.C. Kramer, and M.D. Piggott. A comparison of approaches to tidal stream turbine array optimisation. *Submitted to Journal of Marine Renewable Energy*, 2016c.

N. de Freitas, A. Smola, and M. Zoghi. Regret bounds for deterministic gaussian process bandits. *arXiv preprint arXiv:1203.2177*, 2012.

DHI. MIKE 21 – 2D modelling coast and sea. *http://www.mikebydhi.com*, -2016a.

DHI. MIKE 3 – 3D modelling coast and sea. *http://www.mikebydhi.com*, -2016b.

T. Divett. *Optimising design of large tidal energy arrays in channels: layout and turbine tuning for maximum power capture using large eddy simulations with adaptive mesh.* PhD thesis, University of Otago, 2014.

S. Draper and T. Nishino. Centred and staggered arrangements of tidal turbines. *Journal of Fluid Mechanics*, 739:72–93, 2014.

S. Draper, G.T. Houlsby, M.L. Oldfield, and A.G. Borthwick. Modelling tidal energy extraction in a depth-averaged coastal domain. *Proceedings of the European wave and tidal energy conference*, 2009.

DTOcean. Deliverable 2.1 – Assessment of capabilities of available tools. *DTOcean Project, www.dtocean.eu*, 2014.

R.J. du Feu, S.W. Funke, Kramer S.W., D.M. Culley, J. Hill, B. Halpern, and Piggott M.D. The trade off between tidal-turbine array yield and environmental impact: a multi-objective optimisation problem. *Submitted to Renewable Energy*, 2016.

G. Egbert and L. Erofeeva. OSU Tidal Prediction Software (OTPS). *http://volov.oce.orst.edu/tides/otps.html*, 2010.

EnergyTransition.de. German power prices negative over weekend . http://energytransition.de/2014/05/german-power-prices-negative-over-weekend/, May 2014. Accessed: 2016-05-06.

EPRI. Methodology for estimating tidal current energy resources and power production by tidal in-stream energy conversion devices. *North American in stream power feasibility demonstration project*, 2006.

Y. Eroğlu and S. U. Seçkiner. Design of wind farm layout using ant colony algorithm. *Renewable Energy*, 44:53–62, August 2012.

R.M. Errico. What Is an Adjoint Model ? *Bulletin of the American Meteorological Society*, 78(11):2577–2591, 1997.

I. Fairley, I. Masters, and H. Karunarathna. The cumulative impact of tidal stream turbine arrays on sediment transport in the pentland firth. *Renewable Energy*, 80:755–769, 2015.

F. Fang, C.C. Pain, I.M. Navon, G.J. Gorman, M.D. Piggott, and P.A. Allison. The independent set perturbation adjoint method: A new method of differentiating mesh-based fluids models. *International Journal for Numerical Methods in Fluids*, 66(8):976–999, 2011.

P.E. Farrell, D.A. Ham, S.F. Funke, and M.E. Rognes. Automated derivation of the adjoint of high-level transient finite element programs. *SIAM Journal on Scientific Computing*, 35(4):C369 – C393, 2013.

S.W. Funke. *The automation of PDE-constrained optimisation and its applications*. PhD thesis, Imperial College London, Department of Earth Science and Engineering, October 2012.

S.W. Funke and P.E. Farrell. A framework for automated PDE-constrained optimisation. *arXiv:1302.3894*, 2013.

S.W. Funke, P.E. Farrell, and M.D. Piggott. Tidal turbine array optimisation using the adjoint approach. *Renewable Energy*, 63:657 – 673, 2014.

S.W. Funke, S.C. Kramer, and M.D. Piggott. Design optimisation and resource assessment for tidal-stream renewable energy farms using a new continuous turbine approach. *submitted to Renewable Energy*, 2015a. doi: http://arxiv.org/abs/1507.05795.

S.W. Funke, S.C. Kramer, and M.D. Piggott. Design optimisation and resource assessment for tidal-stream renewable energy farms using a new continuous turbine approach. *arXiv preprint arXiv:1507.05795*, 2015b.

S.W. Funke, D.M. Culley, S.C. Kramer, and M.D. Piggott. OpenTidalFarm: A software package for the optimisation of tidal turbine arrays. *Submitted to Geosci. Model Dev.*, 2016.

C. Garrett and P. Cummins. Generating power from tidal currents. *Journal of waterway, port, coastal, and ocean engineering*, 130(3):114–118, 2004.

C. Garrett and P. Cummins. The power potential of tidal currents in channels. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2060):2563–2572, August 2005.

C. Garrett and P. Cummins. The efficiency of a turbine in a tidal channel. *Journal of fluid mechanics*, 588:243–251, 2007.

C. Geuzaine and J.F. Remacle. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.

J. S. González, A. G. Gonzalez Rodriguez, J. C. Mora, J.R. Santos, and M.B. Payan. Optimization of wind farm turbines layout using an evolutive algorithm. *Renewable Energy*, 35(8):1671–1681, August 2010.

F.M. González-Longatt, P. Wall, P. Regulski, and V. Terzija. Optimal electric network design for a large offshore wind farm based on a modified genetic algorithm approach. *Systems Journal, IEEE*, 6(1):164–172, 2012.

S.A. Grady, M.Y. Hussaini, and M.M. Abdullah. Placement of wind turbines using genetic algorithms. *Renewable Energy*, 30(2):259–270, February 2005.

J. Green, A. Bowen, L.J. Fingersh, Y. Wan, et al. Electrical collection and transmission systems for offshore wind power. In *Offshore technology conference*, volume 30. Offshore Technology Conference, 2007.

G. Hagerman and R. Bedard. Massachusetts tidal in-stream energy conversion (TISEC): survey and characterization of potential project sites. Technical report, Report for Electrical Power Research Institute, 2006.

C. Hee Jo, J. Young Yim, K. Hee Lee, and Y. Ho Rho. Performance of horizontal axis tidal current turbine by blade configuration. *Renewable Energy*, 42:195–206, 2012.

P. Heptonstall, R. Gross, P. Greenacre, and T. Cockerill. The cost of offshore wind: Understanding the past and projecting the future. *Energy Policy*, 41:815–821, February 2012.

L. Hirth. The market value of variable renewables; the effect of solar and wind power variability on their relative price. *Energy Economics*, 38:218–236, July 2013.

M. Hofmann. A review of decision support models for offshore wind farms with an emphasis on operation and maintenance strategies. *Wind Engineering*, 35(1): 1–16, 2011.

IEC. Marine energy - Wave, tidal and other water current converters – Part 201: Tidal energy resource assessment and characterization. Technical Report 1, International Electrotechnical Commission, 4 2015.

C.T. Jacobs, M.D. Piggott, S.C. Kramer, and S.W. Funke. On the validity of tidal turbine array configurations obtained from steady-state adjoint optimisation. *VII European Conference on Computational Methods in Applied Sciences and Engineering, ECCOMAS, Crete.*, 6 2016.

N.O. Jensen. *A note on wind generator interaction.* Risø-M. 1983. ISBN 87-550-0971-9.

C.M. Johnstone, D. Pratt, J.A. Clarke, and A.D. Grant. A techno-economic analysis of tidal energy technology. *Renewable Energy*, 49:101–106, January 2013.

D.R. Jones, M. Schonlau, and W.J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL `http://www.scipy.org/`. [Online; accessed 2015-12-01].

R. Karsten, A. Swan, and J. Culina. Assessment of arrays of in-stream tidal turbines in the bay of fundy. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1985):20120189, 2013.

R.H. Karsten, J.M. McMillan, M.J. Lickley, and R.D. Haynes. Assessment of tidal current energy in the minas passage, bay of fundy. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 222(5):493–507, 2008.

A. Keane and P. Nair. *Computational approaches for aerospace design: the pursuit of excellence.* John Wiley & Sons, 2005.

D. Kraft. A software package for sequential quadratic programming. *Forschungsbericht / Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*, 88 (28), 1988.

S.C. Kramer and M.D. Piggott. A correction to the enhanced bottom drag parameterisation of tidal turbines. *arXiv preprint arXiv:1506.03611*, 2015.

A. Kusiak and Z. Song. Design of wind farm layout for maximum wind energy capture. *Renewable Energy*, 35(3):685–694, 2010.

W.Y Kwong, P.Y. Zhang, D. Romero, J. Moran, M. Morgenroth, and C. Amon. Multi-objective wind farm layout optimization considering energy generation and noise propagation with NSGA-II. *Journal of Mechanical Design*, 2013.

M. Lalander, E. Leijon. Numerical modeling of a river site for in-stream energy converters. In *Proceedings of the 8th European Wave and Tidal Energy Conference, EWTEC09, Uppsala, Sweden*, 2009.

G.C. Larsen and P.-E. Réthoré. Topfarm–a tool for wind farm optimization. *Energy Procedia*, 35:317–324, 2013.

G.C. Larsen, J. Højstrup, and A.H. Madsen. Wind fields in wakes. In *1996 European Union wind energy conference*, pages 764–768. H.S. Stephens & Associates, 1996.

LCICG. Technology Innovation Needs Assessment; Marine Energy, Summary Report, 2012. *Low Carbon Innovation Coordination Group*, 2012.

R.J. LeVeque. *Finite volume methods for hyperbolic problems*, volume 31. Cambridge university press, 2002.

A.S. Lewis and M.L. Overton. Nonsmooth optimization via bfgs. *Submitted to SIAM J. Optimiz*, 2009.

A. Logg, K. A. Mardal, and G. Wells. *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84. Springer Science & Business Media, 2012.

S. Lumbreras and A. Ramos. A benders' decomposition approach for optimizing the electric system of offshore wind farms. In *PowerTech, 2011 IEEE Trondheim*, pages 1–8. IEEE, 2011.

S. Lumbreras and A. Ramos. Optimal design of the electrical layout of an offshore wind farm applying decomposition strategies. *Power Systems, IEEE Transactions on*, 28(2):1434–1441, 2013a.

S. Lumbreras and A. Ramos. Offshore wind farm electrical design: a review. *Wind Energy*, 16(3):459–473, 2013b.

R. Malki, I. Masters, A.J. Williams, and T.N. Croft. Planning tidal stream turbine array layouts using a coupled blade element momentum–computational fluid dynamics model. *Renewable Energy*, 63:46–54, 2014.

C.J. Malmborg. A genetic algorithm for service level based vehicle scheduling. *European Journal of Operational Research*, 93(1):121–134, 1996.

R. Martin-Short, J. Hill, S.C. Kramer, A. Avdis, P.A. Allison, and M.D. Piggott. Tidal resource extraction in the Pentland Firth, UK: Potential impacts on flow regime and sediment transport in the Inner Sound of Stroma. *Renewable Energy*, 76(0): 596 – 607, 2015. doi: 10.1016/j.renene.2014.11.079.

J. McNaughton, I. Afgan, D.D. Apsley, S. Rolfo, T. Stallard, and P.K. Stansby. A simple sliding-mesh interface procedure and its application to the cfd simulation of a tidal-stream turbine. *International Journal for Numerical Methods in Fluids*, 74 (4):250–269, 2014.

N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.

MeyGen. MeyGen Tidal Energy Project Phase 1 Environmental Statement. Technical report, Project consents documentation, 2014.

G. Mosetti, C. Poloni, and B. Diviacco. Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm. *Journal of Wind Engineering and Industrial Aerodynamics*, 51:105–116, 1994.

S. Mungar. Hydrodynamics of horizontal-axis tidal current turbines: A modelling approach based on Delft3D. Master's thesis, Delft University of Technology, 2014.

T. Nishino and R. Willden. Low-order modelling of blade-induced turbulence for rans actuator disk computations of wind and tidal turbines. In *Wind Energy-Impact of Turbulence*, pages 153–158. Springer, 2014.

T. Nishino and R.H.J. Willden. Effects of 3-D channel blockage and turbulent wake mixing on the limit of power extraction by tidal turbines. *International Journal of Heat and Fluid Flow*, 37:123–135, 2012.

F. O'Rourke, F. Boyle, and A. Reynolds. Tidal energy update 2009. *Applied Energy*, 87(2):398–409, February 2010.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

B. Polagye and R. Bedard. Tidal in-stream energy resource assessment for Southeast Alaska. Technical report, Report to Electric Power Research Institute, 2006.

B. Polagye, M. Previsic, and R. Bedard. System Level Design, Performance, Cost and Economic Assessment: Knik Arm Alaska Tidal In-stream Power Plant. Technical report, Reprot to EPRI, 2006.

B. Polagye, M. Kawase, and P. Malte. In-stream tidal energy potential of puget sound, washington. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 223(5):571–587, 2009.

G. Quinonez-Varela, G.W. Ault, O. Anaya-Lara, and J.R. McDonald. Electrical collector system options for large offshore wind farms. *Renewable Power Generation, IET*, 1(2):107–114, 2007.

C.E. Rasmussen. *Gaussian processes for machine learning*. Citeseer, 2006.

D.J. Renkema. Validation of wind turbine wake models. Master's thesis, TU Delft, 2007.

T. Roc, D.C. Conley, and D. Greaves. Methodology for tidal turbine representation in ocean circulation model. *Renewable Energy*, 51:448–464, March 2013.

D. Sale, J. Jonkman, and W. Musial. *Hydrodynamic optimization method and design code for stall-regulated hydrokinetic turbine rotors*. National Renewable Energy Laboratory, 2009.

G.I. Schuëller and H.A. Jensen. Computational methods in optimization considering uncertainties–an overview. *Computer Methods in Applied Mechanics and Engineering*, 198(1):2–13, 2008.

W. Shi, D. Wang, M. Atlar, and K.C. Seo. Flow separation impacts on the hydrodynamic performance analysis of a marine current turbine using cfd. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 227(8):833–846, 2013.

A. Skajaa. Limited memory BFGS for nonsmooth optimization. Master's thesis, New York University, New York, US, 2010.

M.J. Smith, P.I. Palmer, D.W. Purves, M.C. Vanderwel, V. Lyutsarev, B. Calderhead, L.N. Joppa, C.M. Bishop, and S. Emmott. Changing how earth system modeling is done to provide more useful information for decision making, science, and society. *Bulletin of the American Meteorological Society*, 95(9):1453–1464, 2014.

Thomas Sørensen, Morten Lybech Thøgersen, Per Nielsen, and Niels Jernesvej. Adapting and calibration of existing wake models to meet the conditions inside offshore wind farms. *EMD International A/S. Aalborg*, 2008.

T. Stallard, R. Collings, T. Feng, and J. Whelan. Interactions between tidal turbine wakes: experimental study of a group of three-bladed rotors. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1985):20120159, 2013.

P. Stansby and T. Stallard. Fast optimisation of tidal stream turbine positions for power generation in small arrays with low blockage based on superposition of self-similar far-wake velocity deficit profiles. *Renewable Energy*, 92:366–375, 2016.

P. Sullivan and P. McCombie. Optimisation of tidal power arrays using a genetic algorithm. *Proceedings of the Institution of Civil Engineers-Energy*, 166(1):19–28, 2013.

G. Sutherland, M. Foreman, and C. Garrett. Tidal current energy assessment for johnstone strait, vancouver island. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 221(2):147–157, 2007.

L. Tang, J. Liu, A. Rong, and Z. Yang. A multiple traveling salesman problem model for hot rolling scheduling in Shanghai Baoshan Iron & Steel Complex. *European Journal of Operational Research*, 124(2):267–282, July 2000.

M. Tarbotton and M. Larson. Canada ocean energy atlas (phase 1) potential tidal current energy resources analysis background. Technical report, Report to Canadian Hydraulics Centre, 2006.

C. Taylor and P. Hood. A numerical solution of the Navier-Stokes equations using the finite element technique. *Computers & Fluids*, 1(1):73–100, 1973.

M.D. Thomson, J.I. Whelan, and L. Gill. The development of a tool for the design and optimisation of tidal stream turbine arrays. In *Proceedings of the European Wave and Tidal Energy Conference*, 2011.

W. Tong, S. Chowdhury, A. Mehmani, and A. Messac. Multi-objective wind farm design: Exploring the trade-off between capacity factor and land use. In *Proceedings of the 10th World Congress on Structural and Multidisciplinary Optimization, Orlando, FL, USA*, pages 19–24, 2013.

H. Torrens-Spence, P. Schmitt, P. Mackinnon, and B. Elsaesser. Current and turbulence measurement with collocated ADP and turbulence profiler data. In *Current, Waves and Turbulence Measurement (CWTM), 2015 IEEE/OES Eleventh*, pages 1–8. IEEE, 2015.

Triton Consultants. Tidal Current: BC Hydro, BC Green Energy Study - Phase 2. Technical report, Report to BC Hydro, 2002.

S. Venkataraman and R.T. Haftka. Structural optimization complexity: what has Moore's law done for us? *Structural and Multidisciplinary Optimization*, 28(6): 375–387, 2004.

R. Vennell. Tuning turbines in a tidal channel. *Journal of Fluid Mechanics*, 663:253 – 267, 2010. doi: 10.1017/S0022112010003502.

R. Vennell. Estimating the power potential of tidal currents and the impact of power extraction on flow speeds. *Renewable Energy*, 36(12):3558–3565, 2011a.

R. Vennell. Tuning tidal turbines in-concert to maximise farm efficiency. *Journal of Fluid Mechanics*, 671:587–604, 2011b.

R. Vennell, S.W. Funke, S. Draper, C. Stevens, and T. Divett. Designing large arrays of tidal turbines: A synthesis and review. *Renewable and Sustainable Energy Reviews*, 41:454–472, 2015a.

R. Vennell, A. Harang, M. Smeaton, and M. Gerritsen. Giga watt arrays: How many tidal turbines will it take? *Presentation at the European Wave and Tidal Energy Conference in Nantes*, 2015b.

D. J. Wales and J. P. Doye. Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28):5111–5116, 1997.

J.I. Whelan, J.M.R. Graham, and J. Peiro. A free-surface and blockage correction for tidal turbines. *Journal of Fluid Mechanics*, 624:281–291, 2009.

S. Zhang, X. Zou, J. Ahlquist, I.M. Navon, and J.G. Sela. Use of differentiable and nondifferentiable optimization algorithms for variational data assimilation with discontinuous cost functions. *Monthly Weather Review*, 128(12):4031–4044, 2000.

Z. Zhou, Y.S. Ong, and P.B. Nair. Hierarchical surrogate-assisted evolutionary optimization framework. In *Evolutionary Computation, 2004. CEC2004. Congress on*, volume 2, pages 1586–1593. IEEE, 2004.