Integration of cost modelling within the micro-siting design optimisation of tidal turbine arrays

D.M. Culley ^{a*}, S.W. Funke ^{a,b}, S.C. Kramer ^a & M.D. Piggott ^{a,c}

^a Applied Modelling and Computation Group, Department of Earth Science and Engineering, Imperial College London, London, UK

^b Center for Biomedical Computing, Simula Research Laboratory, Oslo, Norway

^c Grantham Institute for Climate Change, Imperial College London, London, UK

 * Corresponding author. dmc13@imperial.ac.uk

Abstract

The location of individual turbines within a tidal current turbine array – micrositing – 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 paper 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.

Keywords: marine renewable energy, tidal turbines, gradient-based optimisation, adjoint method, shallow water equations, array layout, cable-routing, financial-return

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 – micrositing – 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, 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 locationdependent 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 paper 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 paper is a novel method of optimising turbine micrositing 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 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 3, *OpenTidalFarm*, a code developed by Funke et al. (2014), is presented as a method to optimise the array layout to maximise the power extraction of the array, using the shallow water equations and a gradient-based optimisation approach. In section 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, a model is developed for this application leading to section 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 7.

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 encoded in a 2n-long vector, \boldsymbol{m} , where

$$\boldsymbol{m} = (x_1, y_1, x_2, y_2, \dots, x_n, y_n)^T.$$

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

$$\begin{array}{l} \max \quad R_{\rm fin}(\boldsymbol{m}). \end{array}$$
(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 (1) may be expressed as

$$\max_{\boldsymbol{m}} \quad I_P(P(\boldsymbol{m})) - C_C(L(\boldsymbol{m})), \tag{2}$$

where $P(\boldsymbol{m})$ is the power extracted from the flow by the turbines, $I_P : \mathbb{R} \to \mathbb{R}$ maps power output to financial income, $L(\boldsymbol{m})$ 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 \boldsymbol{m} when it is understood that for a given array configuration (\boldsymbol{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 (2) may be found in section 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 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.

3. OpenTidalFarm

An approach to maximising the power production, $P(\mathbf{m})$, has been developed by Funke et al. (2014). This method is packaged in the open-source code *OpenTidalFarm*, and will be summarised here for completeness. *OpenTidalFarm* solves an optimisation problem constrained by the shallow water equations:

$$\begin{array}{ll} \max & P(\boldsymbol{m}) \\ \text{subject to} & \boldsymbol{b}_l \leq \boldsymbol{m} \leq \boldsymbol{b}_u \\ & g(\boldsymbol{m}) \leq 0. \end{array} \tag{3}$$

The bounds $b_l \leq m \leq b_u$ constrain the turbines to the turbine area (in this case rectangular in shape, for simplicity) and the inequality constraint, $g(m) \leq 0$, 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,

$$P(\boldsymbol{m}) = \int_{\Omega} \rho c_t(\boldsymbol{m}) ||\boldsymbol{u}||^3 \, \mathrm{d}x, \qquad (4)$$

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

$$\boldsymbol{u} \cdot \nabla \boldsymbol{u} - \nu \nabla^2 \boldsymbol{u} + g \nabla \eta + \frac{c_b + c_t(\boldsymbol{m})}{H} ||\boldsymbol{u}|| \boldsymbol{u} = 0,$$

$$\nabla \cdot (H\boldsymbol{u}) = 0.$$
(5)

Here, ν is the viscosity coefficient, g is the acceleration due to gravity, H is the total water depth and c_b represents a constant background bottom friction. Note that $P(\boldsymbol{m})$ is a function of \boldsymbol{m} both directly through $c_t(\boldsymbol{m})$ and through the solution of the shallow water equations, $\boldsymbol{u}(\boldsymbol{m})$. In this paper we consider only the steady-state shallow water model, however this can be generalised to the unsteady case.

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, dP/dm, is computed by solving the adjoint shallow water equations (Funke et al., 2014). Based upon both P and dP/dm, the locations of the turbines, m is 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 m based upon previous functional and gradient evaluations. The *OpenTidalFarm* optimisation loop is outlined in figure 1.

A key advantage of this approach is the use of the adjoint equations, since it enables the gradient dP/dm to be determined at a computational cost that is independent of the number of turbines, n. (Technically, the *adjoint* is computed



Figure 1: Schematic of the OpenTidalFarm optimisation procedure.

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 permits 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 $\mathbf{m} \in \mathbb{R}^{2n} \to \mathbb{R}$, 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 paper. Integration of cable cost presents a particular challenge, as determining the length of the cable route required for a given 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.

4. Cable length minimisation

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.

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.

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, 2013b, 2011; Bauer and Lysgaard, 2014; González-Longatt et al., 2012; Lumbreras and Ramos, 2013a).

Lumbreras and Ramos (2013b) 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 capturing the costs of collection points and cabling, and used a genetic algorithm to find a minimum cost routing solution. As with Lumbreras and Ramos (2013b), 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 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 on-shore 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.

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 $n \geq 15$. 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 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 constraints (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.

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 $r_a \cap r_b = \emptyset \quad \forall \ 0 \leq a < b \leq k$.

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

$$L_R(\boldsymbol{m}) = \sum_{(i,j)\in R} l_{ij}(\boldsymbol{m}),\tag{6}$$

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

$$l_{ij}(\boldsymbol{m}) = \sqrt{(\boldsymbol{m}_{xi} - \boldsymbol{m}_{xj})^2 + (\boldsymbol{m}_{yi} - \boldsymbol{m}_{yj})^2},$$
(7)

and m_{xi} corresponds to the x coordinate of the *i*th vertex in m, etc.

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

$$L_{\text{opt}}(\boldsymbol{m}) = \min_{\boldsymbol{R} \in \mathcal{R}} L_{\boldsymbol{R}}(\boldsymbol{m}).$$
(8)

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

The numerical solution of the cable routing problem (8) will be discussed in section 5. Recall however, that L_{opt} is part of the optimisation problem (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 $\frac{dL_{opt}}{dm}$ does not always exist. To illustrate this in one dimension, consider figure 2. 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 3, 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 saw-tooth 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 (6), (7). Next observe that the minimum of two Lipschitz functions 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 $\mathcal{R}(\boldsymbol{m})$ yields that L_{opt} is Lipschitz. The differentiability almost everywhere follows from Rademacher's theorem.



Figure 2: If the turbine layout from figure 4 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 .

5. Tidal Array Routing

Various meta-heuristics were trialled for solving the cable routing problem (8). A genetic algorithm was finally selected as, in initial experimentation, it was found to be the most robust. Furthermore genetic algorithms are the most



Figure 3: Graph showing the routing length for two routings as turbine 11 (T11) in figure 2 traverses in the x direction.

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. As such, much of the terminology associated with this approach is borrowed from genetics. 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.

5.1. 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.



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

Consider figure 4. For this example the two-part chromosome technique would represent this routing as

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

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 $n_{\rm rts}$ individual routes. Compared to other chromosome encodings this has very few redundant solutions and the size of the solution space is shown to be

$$n! \binom{n-1}{n_{\rm rts} - 1}.\tag{9}$$

(Carter and Ragsdale, 2006)

TAR is initialised with a population, \mathcal{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, \mathcal{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.

Figure 5: Schematic of TidalArrayRouting developed for OpenTidalFarm

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 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 $L_{\text{opt}}(\boldsymbol{m})$ and a perturbation vector $\delta \boldsymbol{m}$, then the first-order Taylor remainder should tend to zero linearly with the decreasing perturbation, that is

$$|L_{\text{opt}}(\boldsymbol{m} + h\delta\boldsymbol{m}) - L_{\text{opt}}(\boldsymbol{m})| \to 0 \quad \text{at} \quad \mathcal{O}(h).$$
(10)

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

$$|L_{\text{opt}}(\boldsymbol{m} + h\delta\boldsymbol{m}) - L_{\text{opt}}(\boldsymbol{m}) - h\nabla L_{\text{opt}} \cdot \delta(\boldsymbol{m})| \to 0 \quad \text{at} \quad \mathcal{O}(h^2)$$
(11)

where ∇L_{opt} is the gradient of L_{opt} with respect to \boldsymbol{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.

6. Integrating the micro-siting and cable-routing optimisations

Integrating the cable-routing kernel into OpenTidalFarm (see figure 1) 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 (dL/dm) 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 (2) that the financial return, R_{fin} is defined as

$$R_{\text{fin}} = I_P(P(\boldsymbol{m})) - C_C(L(\boldsymbol{m})), \qquad (12)$$

and consequently,

$$\frac{dR_{\rm fin}}{d\boldsymbol{m}} = I_P \left(\frac{dP}{d\boldsymbol{m}}\right) - C_C \left(\frac{dL}{d\boldsymbol{m}}\right). \tag{13}$$

Consequently, 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,

$$NPV(d, N) = \sum_{t=0}^{N} \frac{R_t}{(1+d)^t}.$$
 (14)

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, 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%.

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

$$5.16 \times 10^6 + 254.51L \le C_C(L) \le 6.22 \times 10^6 + 862.48L.$$
(15)

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 $[LCOE] = \pounds MWh^{-1}$). Which is the present value of all the project costs divided by the discounted lifetime energy production of the plant

$$LCOE = \frac{\sum_{t=0}^{N} \frac{C_{ann}^{t}}{(1+d)^{t}}}{\sum_{t=0}^{N} \frac{O_{ann}}{(1+d)^{t}}},$$
(16)

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 $\pounds 61 \text{ MWh}^{-1}$ to $\pounds 106 \text{ MWh}^{-1}$ for between 15 and 20 years

$$61 \times P \times \sum_{t=0}^{15} \frac{2920}{(1.06)^t} \le I_P(P) \le 106 \times P \times \sum_{t=0}^{20} \frac{2920}{(1.15)^t}$$
$$1.91 \times 10^6 P \le I_P(P) \le 2.24 \times 10^6 P \tag{17}$$

where P is the installed capacity of the array – as returned by OpenTidalFarm.

7. Test cases

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 6 shows the schematic of the test layout overlaid on the Meygen site (or figure 7 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 bed roughness set to a constant value of 0.0025. The viscosity was set to 3 m²s⁻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) gradientbased 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 the starting guess of the turbine locations 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



Figure 6: 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*.

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, 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 m^0 .

As such, 8 turbines are initialised in a grid formation (see figure 7(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 7(c)) and cable length is minimised. The optimisation is then restarted from this layout, with the objective now defined as in (1), to maximise the *return* provided by the array, assuming a 'worst case' scenario – i.e. using the upper bound of C_C from (16) along with the lower bound of I_P from (17). The model converges as shown in figure 7(d) with a total power extraction of 32.19 MW for 3135 m of cabling. In contrast, figure 7(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 7: Eight turbines are optimised in the domain from figure 6 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).

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 8 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 8: 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 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 9(b) resulting in 35.00 MW power extracted for 2387 m of cabling. For comparison the power-only optimisation is shown in figure 9(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).

7.3. Pentland Firth

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 10, 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 $90 \text{ m}^2 \text{s}^{-1}$ 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



Figure 9: 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).

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 10 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 10). Maximisation of power output yielded a micrositing design as shown in figure 11 (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 11 (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 11 (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.



Figure 10: Domain, mesh and inflow boundary of the idealised Orkney example.

8. Conclusions

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.

The main extensions of this work are to increase the realism of the model. On the power optimisation side, elements such as bathymetry and realistic timedependent tidal forcing are simple to include, but come at an increased (sometimes significantly) computational cost. Likewise the inclusion of turbulence models will likely be an important step to fully capturing the structure and shape of the turbine wake, which is obviously a vital part of the interaction between turbines. On the cabling side, 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 11: 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.

9. Acknowledgements

The authors would like to acknowledge Imperial College London, the Engineering and Physical Sciences Research Council (projects: EP/J010065/1 and EP/L000407/1) and a Center of Excellence grant from the Research Council of Norway to the Center for Biomedical Computing at Simula Research Laboratory for funding which supported this work. Finally, the authors would like to acknowledge helpful discussions with B. Jones and B. Amor, and to thank M. Nordaas for his feedback on section 4.5.

10. Bibliography

- Allan, G., Gilmartin, M., McGregor, P., Swales, K., Jan. 2011. 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.
- Barnett, G. L., Funke, S. W., Piggott, M. D., 2014. Hybrid global-local optimisation algorithms for the layout design of tidal turbine arrays. Submitted to Renewable Energy.
- Bauer, J., Lysgaard, J., 2014. The offshore wind farm array cable layout problem: a planar open vehicle routing problem. Journal of the Operational Research Society.
- Bektas, T., Jun. 2006. The multiple traveling salesman problem: an overview of formulations and solution procedures. Omega 34 (3), 209–219.
- Bellman, R., 1962. Dynamic programming treatment of the travelling salesman problem. Journal of the ACM (JACM) 9 (1), 61–63.
- Carter, A., Ragsdale, C., 2006. A new approach to solving the multiple traveling salesperson problem using genetic algorithms. European Journal of Operational Research 175 (1), 246–257.
- Clarke, G., Wright, J., 1964. Scheduling of vehicles from a central depot to a number of delivery points. Operations research 12 (4), 568–581.
- Funke, S., Farrell, P., Piggott, M., 2014. Tidal turbine array optimisation using the adjoint approach. Renewable Energy 63 (0), 658 – 673.
- González-Longatt, F., Wall, P., Regulski, P., Terzija, V., 2012. Optimal electric network design for a large offshore wind farm based on a modified genetic algorithm approach. Systems Journal, IEEE 6 (1), 164–172.
- Green, J., Bowen, A., Fingersh, L., Wan, Y., 2007. Electrical collection and transmission systems for offshore wind power. In: Offshore technology conference. Vol. 30. Offshore Technology Conference.

- Heptonstall, P., Gross, R., Greenacre, P., Cockerill, T., 2012. The cost of offshore wind: Understanding the past and projecting the future. Energy Policy 41, 815–821.
- Jones, E., Oliphant, T., Peterson, P., 2001. SciPy: Open source scientific tools for Python.
- Kraft, D., 1988. A software package for sequential quadratic programming. Forschungsbericht / Deutsche Forschungs- und Versuchsanstalt f
 ür Luft- und Raumfahrt 88 (28).
- Larsen, G., Réthoré, P.-E., 2013. Topfarm–a tool for wind farm optimization. Energy Procedia 35, 317–324.
- Lewis, A., Overton, M., 2009. Nonsmooth optimization via BFGS. To appear; SIAM Journal of Optimization.
- Lumbreras, S., Ramos, A., 2011. A Benders' decomposition approach for optimizing the electric system of offshore wind farms. In: PowerTech, 2011 IEEE Trondheim. pp. 1–8.
- Lumbreras, S., Ramos, A., 2013a. Offshore wind farm electrical design: a review. Wind Energy 16 (3), 459–473.
- Lumbreras, S., Ramos, A., 2013b. Optimal design of the electrical layout of an offshore wind farm applying decomposition strategies. IEEE Transactions on Power Systems 28 (2), 1434–1441.
- Malmborg, C., 1996. A genetic algorithm for service level based vehicle scheduling. European Journal of Operational Research 93 (1), 121–134.
- Quinonez-Varela, G., Ault, G., Anaya-Lara, O., McDonald, J., 2007. Electrical collector system options for large offshore wind farms. Renewable Power Generation, IET 1 (2), 107–114.
- Skajaa, A., 2010. Limited memory BFGS for nonsmooth optimization. Master's thesis, New York University, New York, US.
- Tang, L., Liu, J., Rong, A., Yang, Z., Jul. 2000. 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.
- Taylor, C., Hood, P., 1973. A numerical solution of the Navier-Stokes equations using the finite element technique. Computers & Fluids 1 (1), 73–100.
- Thomson, M., Whelan, J., Gill, L., 2011. 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.
- Zhang, S., Zou, X., Ahlquist, J., Navon, I., Sela, J., 2000. Use of differentiable and nondifferentiable optimization algorithms for variational data assimilation with discontinuous cost functions. Monthly Weather Review 128 (12), 4031–4044.