



# Tidal turbine array optimisation using the adjoint approach<sup>☆</sup>



S.W. Funke<sup>a,c,\*</sup>, P.E. Farrell<sup>d</sup>, M.D. Piggott<sup>a,b</sup>

<sup>a</sup> Applied Modelling and Computation Group, Department of Earth Science and Engineering, Imperial College London, London, UK

<sup>b</sup> Grantham Institute for Climate Change, Imperial College London, London, UK

<sup>c</sup> Center for Biomedical Computing, Simula Research Laboratory, Oslo, Norway

<sup>d</sup> Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

## ARTICLE INFO

### Article history:

Received 5 April 2013

Accepted 20 September 2013

Available online

### Keywords:

Marine renewable energy

Tidal turbines

Gradient-based optimisation

Adjoint method

Shallow water equations

Array layout

## ABSTRACT

Oceanic tides have the potential to yield a vast amount of renewable energy. Tidal stream generators are one of the key technologies for extracting and harnessing this potential. In order to extract an economically useful amount of power, hundreds of tidal turbines must typically be deployed in an array. This naturally leads to the question of how these turbines should be configured to extract the maximum possible power: the positioning and the individual tuning of the turbines could significantly influence the extracted power, and hence is of major economic interest. However, manual optimisation is difficult due to legal site constraints, nonlinear interactions of the turbine wakes, and the cubic dependence of the power on the flow speed. The novel contribution of this paper is the formulation of this problem as an optimisation problem constrained by a physical model, which is then solved using an efficient gradient-based optimisation algorithm. In each optimisation iteration, a two-dimensional finite element shallow water model predicts the flow and the performance of the current array configuration. The gradient of the power extracted with respect to the turbine positions and their tuning parameters is then computed in a fraction of the time taken for a flow solution by solving the associated adjoint equations. These equations propagate causality backwards through the computation, from the power extracted back to the turbine positions and the tuning parameters. This yields the gradient at a cost almost independent of the number of turbines, which is crucial for any practical application. The utility of the approach is demonstrated by optimising turbine arrays in four idealised scenarios and a more realistic case with up to 256 turbines in the Inner Sound of the Pentland Firth, Scotland.

© 2013 The Authors. Published by Elsevier Ltd. All rights reserved.

## 1. Introduction

With the increasing cost of energy, tidal turbines are becoming a competitive and promising option for renewable electricity generation. A key advantage of tidal energy is that the power extracted is predictable in advance, which is highly attractive for grid management. In order to amortise the fixed costs of installation and grid connection, arrays consisting of hundreds of tidal turbines must typically be deployed at a particular site. This raises the question of where to place the turbines within the site and how to tune them individually in order to maximise the power output; finding the optimal configuration is of huge importance as it could substantially

change the energy captured and possibly determine whether the project is economically viable. However, the determination of the optimal configuration is difficult because of the complex flow interactions between turbines and the fact that the power output depends sensitively on the flow velocity at the turbine positions.

This problem has heretofore been addressed in two different ways. One approach is to simplify the tidal flow model such that the solutions are either available as explicit analytical expressions, or are extremely fast to compute. This means that the optimum can be analytically derived, or that the whole parameter space of possible configurations can be rapidly explored. For example, Bryden and Couch [5] and Garrett and Cummins [15] optimised simplified models to derive an estimate for the maximum energy that can be extracted from a tidal basin. Vennell [39,40] used simple one-dimensional models to demonstrate the importance of tuning each turbine individually to account for the channel geometry, turbine positions, and the tidal forcing. Thus, optimisation of farms is a crucial step needed to achieve their full potential. However, Vennell [42] observes that this optimisation requires many model runs (if performed naively), thus making it computationally

<sup>☆</sup> This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial-No Derivative Works License, which permits non-commercial use, distribution, and reproduction in any medium, provided the original author and source are credited.

\* Corresponding author. Applied Modelling and Computation Group, Department of Earth Science and Engineering, Imperial College London, London, UK.

E-mail address: [s.funke09@imperial.ac.uk](mailto:s.funke09@imperial.ac.uk) (S.W. Funke).

infeasible to use expensive, physically-accurate flow models for this task. While this approach can provide a coarse estimate for the power potential of a site, these simplified models cannot accurately capture the complex nonlinear flow interactions between turbines.

The second approach is to use more complex flow models to accurately predict the tidal flow, the turbine wakes, and the resulting power output. These models are usually formulated as numerical solutions to partial differential equations (PDEs). The computational expense of these models prohibits the exploration of the whole parameter space [38]. Consequently, typically only a handful of manually identified turbine configurations are investigated in a given scenario [1]. Divett et al. [8] compared the power output of four different layouts in a rectangular channel by solving the two-dimensional nonlinear shallow water equations and was able to improve the power outcome by over 50% compared to a regular layout. Lee et al. [28] used a three-dimensional model to investigate how the distance between adjacent rows in a regular array layout impacts the turbine efficiency and showed an efficiency decay for distances of less than three times the turbine diameter. While these studies show the potential of improving the performance by changing the turbine positions, such manual optimisation guided by intuition and experience becomes difficult in a realistic domain with complex bottom bathymetry, flow dynamics and hundreds of turbines.

In this paper, we present a novel technique for maximising the power extraction of array configurations that combines the physical fidelity of PDE-based flow models with advanced automated optimisation techniques. This approach allows the identification of optimal solutions in a computationally feasible number of iterations, circumventing the computational limitations noted in Ref. [42]. The turbine configuration problem is formulated as a PDE-constrained optimisation problem, which is a major topic of research in applied mathematics [20,21]. The resulting maximisation problem is solved using a gradient-based optimisation algorithm that takes orders of magnitude fewer iterations than genetic algorithms or simulated annealing approaches (see e.g. Ref. [3]). In this paper, the power extracted by an array configuration is predicted using a two-dimensional nonlinear shallow water model, which captures the interactions between the geometry, the turbines, and the flow. The gradient of the power is efficiently computed using the *adjoint technique* of variational calculus, which solves an auxiliary system that propagates causality backwards through the physical system. This yields the gradient at a cost almost independent of the number of turbines to be optimised, which is crucial for the method to be applied to large arrays. This gradient is used by the optimisation algorithm to automatically reposition the turbines and to adjust their tuning parameters. The flow solution is re-evaluated, and the algorithm iterated until an optimum is found.

This approach has several key advantages. Firstly, it closes the optimisation loop, by accounting for the effects of the turbines on the flow field itself. This is necessary to find the actual optimum of the nonlinear optimisation problem. Secondly, unlike gradient-free methods, the approach requires a relatively small number of model evaluations and scales to large numbers of turbines, which is necessary for the optimisation of industrial arrays. For example, in Section 6, an array of 256 turbines is optimised in a realistic domain at an approximate cost of 200 flow solutions. Thirdly, the optimisation algorithm can incorporate complex constraints such as minimum separation distances, bathymetry gradient constraints, and legal site restrictions. Finally, the same mathematical framework extends naturally to more realistic flow models such as the Reynolds-averaged Navier–Stokes equations, and to other functionals such as profit or environmental impact.

The approach is implemented in an open-source software framework called OpenTidalFarm; all code and examples from this paper are available at <http://opentidalfarm.org>.

## 1.1. Optimisation algorithms

Optimisation algorithms can be divided into two categories: gradient-free and gradient-based algorithms. Gradient-free optimisation algorithms use the functional of interest (in this case, power extracted by the array) as a black box. They proceed by evaluating the functional at many points in parameter space and use these values to decide which areas merit further exploration. While these methods tend to be robust and can, under certain smoothness conditions, provably find globally optimal solutions [33], they typically require a very large number of functional evaluations that scales linearly or superlinearly with the number of parameters to be optimised. For example, Bilbao and Alba [3] used a genetic algorithm that mimics the process of natural evolution to optimise the location of 8 wind turbines. The algorithm was able to improve the power output by about 70% compared to the initial layout after 17,300 functional evaluations. This large number of evaluations clearly introduces a practical upper limit for the number of turbines that can be optimised. This difficulty is compounded if a more realistic (and hence more expensive) model is used.

By contrast, gradient-based optimisation algorithms use additional information to update the position in parameter space at each iteration: the first or higher derivatives of the functional of interest with respect to the parameters. Depending on the problem, this can lead to a significant reduction in the number of iterations required compared to gradient-free algorithms, making these the only feasible choice for large scale optimisation problems [20]. One caveat of applying gradient-based optimisation algorithms is that they find only local optima. This issue can be circumvented by using hybrid approaches [22]. The main difficulty of applying gradient-based methods is that the implementation of the gradient computation can be difficult for complex models, as it involves differentiating through the solution of a partial differential equation.

One way to obtain the derivative information is to approximate the gradient using finite differences. However, a major disadvantage of this approach is that a single gradient evaluation requires a large number of functional evaluations that scales linearly with the number of optimisation parameters. This sets a practical upper bound on the number of turbines to be optimised, and discards the main advantage of gradient-based optimisation algorithms. Alternatively, the tangent linearisation of the model (i.e. the derivative of the model evaluated at a particular solution) can efficiently compute the derivative of all outputs with respect to a single input, while the adjoint linearisation can efficiently compute the derivative of a single output with respect to all inputs [19]. For the turbine optimisation problem, we wish to maximise a single output (the power extracted) with respect to many input parameters (the positions and tuning parameters of the turbines); this means that the adjoint approach is the natural choice, as the required gradient information can be computed in a number of equation solves that is independent of the number of turbines.

The development of adjoint models is generally considered as very complicated [17,30]. However, this problem has been solved in recent work for the case where the forward model is discretised using finite elements, in the high-level FEniCS framework [11]. This allows for the extremely rapid development of optimally efficient adjoint models, which significantly reduces the development effort required to implement gradient-based optimisation algorithms for PDE-constrained optimisation problems [13].

To the best of our knowledge, this paper presents the first application of the adjoint method to the optimisation of turbine arrays. While the examples are shown in the marine context, it is expected that the presented techniques can also be applied to the optimisation of wind farms. As the wind turbine layout problem is

Download English Version:

<https://daneshyari.com/en/article/6769106>

Download Persian Version:

<https://daneshyari.com/article/6769106>

[Daneshyari.com](https://daneshyari.com)