

ON THE VALIDITY OF TIDAL TURBINE ARRAY CONFIGURATIONS OBTAINED FROM STEADY-STATE ADJOINT OPTIMISATION

Christian T. Jacobs¹, Matthew D. Piggott¹, Stephan C. Kramer¹ and Simon W. Funke²

¹Department of Earth Science and Engineering, Imperial College London
London SW7 2AZ, United Kingdom
e-mail: {c.jacobs10,m.d.piggott,s.kramer}@imperial.ac.uk

² Biomedical Computing Department, Simula Research Laboratory
P.O. Box 134, 1325 Lysaker, Norway
e-mail: simon@simula.no

Keywords: Shallow Water Equations, Turbines, Adjoint Optimisation, Large Eddy Simulation.

Abstract. *Extracting the optimal amount of power from an array of tidal turbines requires an intricate understanding of tidal dynamics and the effects of turbine placement on the local and regional scale flow. Numerical models have contributed significantly towards this understanding, and more recently, adjoint-based modelling has been employed to optimise the positioning of the turbines in an array in an automated way and improve on simple man-made configurations (e.g. structured grids of turbines) [1].*

Adjoint-based optimisation of high-resolution and ideally 3D transient models is generally a very computationally expensive problem. Multiple approaches are therefore used in practice to obtain feasible runtimes: using high viscosity values to obtain a steady-state solution, or a sequence of steady-state solutions for “time-varying” setups; limiting the number of adjoint computations; or reformulating the problem to allow for coarser mesh resolution to make it feasible for resources assessment (e.g. [2], [3]). However, such compromises may affect the reliability of the modelled turbines, their wakes and interactions, and thus bring into question the validity of the computed optimal turbine positions.

This work considers a suite of idealised simulations of flow past tidal turbine arrays in a two-dimensional channel. It compares four regular array configurations, detailed by Divett et al. [4], with the configuration found through adjoint optimisation in a steady-state, high-viscosity setup. The optimised configuration produces considerably more power than the other configurations (approximately 40% more than the best man-made configuration). The same configurations are then used to produce a suite of transient simulations that do not use constant high-viscosity, and instead use large eddy simulation (LES) to parameterise the resulting turbulent structures. All simulations are performed using OpenTidalFarm [1].

It is shown that the ‘low background viscosity’/LES simulations produce less power than that predicted by the constant high-viscosity runs. Nevertheless, they still follow the same trends in the power curve throughout time, with optimised layouts continuing to perform significantly better than simplified configurations.

1 INTRODUCTION

Tidal stream turbines are a promising new source of renewable energy. In order for the deployment of turbines to be effective and economically viable, their placement in a turbine array/farm must be carefully considered. Knowledge of the tidal flow conditions, the effect the turbines will have on the flow and any natural habitats, and how any resulting changes will affect turbine power generation downstream (e.g. through blockage effects), all need to be investigated. Experimental studies are often limited to small-scale domains with idealised flow conditions, which may not scale well to realistic marine environments. As more computational resources become available to solve bigger problems, numerical modelling and adjoint-based optimisation of tidal turbine positions is becoming an increasingly popular alternative.

Current work on adjoint-based optimisation has shown very promising results with regards to improving the efficiency of turbine farms over simple man-made configurations, such as a structured grid of turbines or staggering them in some regular way. For example, Funke et al. [1] optimised an array in the Inner Sound of the Pentland Firth, comprising 256 turbines, which resulted in a 33% increase in power extraction. Similar array optimisation studies by Culley et al. [5] and Barnett et al. [6] in idealised scenarios and once again in the Pentland Firth yielded increased power generation of up to 25%. However, despite the power that adjoint-based optimisation offers, it is extremely computationally-demanding. Such optimisation requires both the forward run and the adjoint run (i.e. the simulation run backwards in time) to be performed at every iteration of an optimisation procedure, based on numerical algorithms such as L-BFGS-B (the limited-memory Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm with bound constraints) [7, 8]. High-resolution transient simulations of flow past marine power turbines, particularly those that feature turbulent wakes and thus require even smaller time-steps to capture the physics, are often prohibitively expensive.

As a result, current work on the adjoint optimisation of tidal turbine placement in the literature has been mostly limited to steady-state simulations in which very high, uniform and constant, non-physical values of the background viscosity are required to ensure that a steady-state solution exists (e.g. $0.5 \text{ m}^2\text{s}^{-1}$ [2], $3 \text{ m}^2\text{s}^{-1}$ [9, 6], $30 \text{ m}^2\text{s}^{-1}$ [5, 3, 1], $90 \text{ m}^2\text{s}^{-1}$ [10]). Such simulations only require one solve of the governing forward equations (with the time derivatives removed) and one solve of the adjoint equations¹, per optimisation iteration, making the problem tractable with available resources. However, the wake downstream of each turbine is likely to be diffused out in such simulations, and no turbulence is modelled. The validity of the results from the optimisation process is therefore questionable. The size of the wake and the presence of any turbulent eddies can potentially affect the power generated by the turbines further downstream, and any blockage effects from neighbouring turbines may not be realistic. When the optimised turbine positions from the steady-state computations are applied to a realistic, transient flow comprising a much lower background viscosity and higher Reynolds number, the amount of power generated by the flow may be substantially different and the benefits of the optimisation process may not be as pronounced compared to the steady-state, high-viscosity case.

The work presented in this paper first considers four simple man-made turbine array configurations by Divett et al. [4]. Steady-state simulations are run with each of these configurations, and the configuration that provides the highest amount of generated power is then optimised using adjoint-based optimisation. The resulting optimised configuration, and also the four man-

¹The simulation can also be divided up into n steady-states, in which case $2n$ solves are required per optimisation iteration.

made configurations, are then each considered in a transient setup in which a more realistic background viscosity is used and Large Eddy Simulation (LES) is applied to parameterise the turbulence. Section 2 presents the shallow water equations that are solved by the numerical model, along with details of the Smagorinsky LES turbulence parameterisation that is enabled in the transient flow simulations. Setup details are provided in Section 3, followed by an analysis of the steady-state and transient simulation results in Section 4. The paper closes with some concluding remarks in Section 5.

2 NUMERICAL MODEL

This work uses the OpenTidalFarm numerical modelling package [1] for solving the shallow water equations. It uses FEniCS [11] to discretise the equations with the finite element method, and the Dofin-Adjoint framework [12, 13] to automatically annotate the forward model system and compute the adjoint model based on that.

2.1 Momentum equation

The equation governing conservation of momentum is given by

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -g \nabla \eta + \nabla \cdot \mathbb{T} - \frac{c_b + c_t}{H} \|\mathbf{u}\| \mathbf{u}, \quad (1)$$

where \mathbf{u} is the depth-averaged velocity, g is the acceleration due to gravity, η is the free-surface displacement, h is the free-surface height at rest, H is the total free-surface height (i.e. $h + \eta$), c_b is the (dimensionless) bottom drag coefficient, and c_t is the (dimensionless) tidal turbine drag coefficient. The Euclidean norm $\|\mathbf{u}\| = \sqrt{\mathbf{u} \cdot \mathbf{u}}$ is used here. The stress tensor \mathbb{T} is given by

$$\mathbb{T} = \nu (\nabla \mathbf{u} + \nabla \mathbf{u}^T), \quad (2)$$

where ν is the kinematic viscosity.

Each of the N turbines is represented by a two-dimensional Gaussian profile, such that for a turbine i of radius r_i centred at (x_i, y_i) ,

$$c_t = \sum_{i=1}^N K_i \phi_{x_i, r_i}(x) \phi_{y_i, r_i}(y), \quad (3)$$

where K_i is a dimensionless friction coefficient and

$$\phi_{p,r}(x) = \begin{cases} \exp(1 - 1/(1 - |\frac{x-p}{r}|^2)) & \text{if } |\frac{x-p}{r}| < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

2.2 Continuity equation

The continuity equation is given by

$$\frac{\partial \eta}{\partial t} + \nabla \cdot (H \mathbf{u}) = 0. \quad (5)$$

2.3 Turbulence parameterisation

Large eddy simulation (LES) is used to parameterise the turbulence generated in the transient flow simulations with a low background viscosity. The crux of LES is to resolve the large-scale eddies by using a fine enough mesh spacing, whilst modelling any smaller, sub-grid scale eddies

with an eddy viscosity term which acts as a diffusive agent; it essentially describes the energy transfer from the larger (resolved) eddies down to the smaller eddies. The particular LES model employed in this work is the standard Smagorinsky model [14, 15]. The eddy viscosity term is given by

$$\nu' = (c_s \Delta)^2 |\mathbb{S}|, \quad (6)$$

where c_s is the Smagorinsky coefficient, set to 0.2 in this work (within the range of typical c_s values [16]). The filter width Δ is an estimate of the local element size which is defined here as the square root of the area of each element. The strain rate tensor \mathbb{S} is defined as

$$\mathbb{S} = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^T), \quad (7)$$

The modulus of \mathbb{S} is given by

$$|\mathbb{S}| = \sqrt{2 \sum_i \sum_j \mathbb{S}_{ij} \mathbb{S}_{ij}}. \quad (8)$$

where \mathbb{S}_{ij} is the (i,j) -th component of \mathbb{S} . The eddy viscosity ν' is added to the background viscosity ν in (2) and is updated at the beginning of each time-step.

Turbulence is a truly three-dimensional phenomenon. Furthermore, LES requires a substantial amount of resolution to sufficiently resolve the large-scale eddies. The validity of applying an LES model in this work to a two-dimensional domain is therefore questionable. Nevertheless, 2D turbulence modelling in the context of tidal turbine simulation has shown good agreement with experimental data [17], and similar LES schemes such as MILES have generated promising results in resolving turbine wakes [18, 19].

3 SIMULATION SETUP

All the simulations presented here are two-dimensional and were performed in a $3 \times 1 \text{ km}^2$ rectangular domain (in the x - y plane), following a similar setup to the work by Divett et al. [4]. A total of 15 turbines were deployed in the four man-made configurations as described in [4], and a suite of steady-state shallow water simulations was first performed for each of the configurations. The physical parameters are provided in Table 1.

3.1 Mesh

The computational mesh that discretised the $3 \times 1 \text{ km}^2$ domain, comprising triangular elements, was generated using Gmsh [20]. The inner region defined by $500 \leq x \leq 2,500 \text{ m}$ and $125 \leq y \leq 875 \text{ m}$, where all the turbines are situated, contained a structured grid of solution nodes with a characteristic element length of $\Delta x = 2 \text{ m}$. The outer section of the domain's mesh, which is outside the region of interest and therefore does not require high numerical resolution, was unstructured with $\Delta x = 100 \text{ m}$.

3.2 Initial and boundary conditions

The initial conditions $\mathbf{u}(x, t = 0) = [0, 0]^T \text{ ms}^{-1}$ and $\eta(x, t = 0) = 0 \text{ m}$ were applied at $t = 0 \text{ s}$. Throughout the simulations a Dirichlet velocity boundary condition of 2 ms^{-1} was strongly imposed to allow inflow from the left. For the steady-state simulations, outflow was

| Parameter | Value |
|------------------------------------|---|
| Free-surface height at rest | $h = 50$ m |
| Kinematic viscosity (steady-state) | $\eta = 1$ m ² s ⁻¹ |
| Kinematic viscosity (transient) | $\eta = 10^{-6}$ m ² s ⁻¹ |
| Smagorinsky coefficient | $c_s = 0.2$ |
| Gravitational acceleration | $g = 9.81$ ms ⁻² |
| Water density | $\rho = 1,000$ kgm ⁻³ |
| Bottom drag coefficient | $c_b = 0.0025$ |
| Number of turbines | $N = 15$ |
| Turbine radii | $r_i = 10$ m $\forall i = 1, \dots, N$ |
| Turbine friction coefficient | $K_i = 12$ $\forall i = 1, \dots, N$ |

Table 1: The common parameter values used in all simulations. The Smagorinsky coefficient c_s is only required in the transient simulations where the Smagorinsky LES model is used. Note also the difference in the background kinematic viscosity between the steady-state and transient simulations.

modelled by a zero-value Dirichlet free-surface condition strongly imposed along the right-hand boundary. For the transient simulations, a Flather [21] condition was used to minimise spurious reflections at the outlet. No-normal flow conditions were weakly imposed along the lateral walls.

3.3 Spatial and temporal discretisation

The Galerkin finite element method was used to discretise the model equations in space. Continuous piecewise quadratic Lagrange basis functions were used to represent the velocity solution field, whereas continuous piecewise linear Lagrange basis functions were used for the free-surface η (thereby forming the P2-P1, or Taylor-Hood, element pair [22]) as well as all other fields (including ν and ν').

Unlike the steady-state simulations, the transient simulations required temporal discretisation due to the presence of time derivatives. The implicit backward Euler method was used for this purpose so that larger time-steps could be taken. All transient simulations were run until $T = 2,000$ s with a time-step of $\Delta t = 4$ s, which gave ample opportunity for the turbulent flow to become fully developed and for any spurious reflections from the outflow to dissipate. Newton iteration dealt with the non-linearity introduced through the advection term in (1); for each time-step, k iterations were performed such that $\|\mathbf{u}^{k+1} - \mathbf{u}^k\| \leq 10^{-9}$, where k is typically between 3 and 10. Once both (1) and (5) are discretised and assembled to form a fully block-coupled system, the system was solved directly with LU decomposition.

3.4 Man-made turbine configurations

3.4.1 Centred

As per the work in [4], the centred configuration illustrated in Figure 1 considered a regular 5×3 grid of turbines positioned in the centre of the channel. The turbines in each row are $7.5d$ away from each other, and each row is $10d$ apart, where $d=2r$ is the turbine diameter.




Figure 1: The centred turbine configuration. The whole $3 \times 1 \text{ km}^2$ channel is shown.




Figure 2: The offset turbine configuration. The whole $3 \times 1 \text{ km}^2$ channel is shown.

3.4.2 Offset

The offset case (see Figure 2) considers the same regular grid of turbines as the centred case, but instead of being placed in the centre of the domain, the grid is positioned closer to one of the side walls. Such a configuration may be required to, for example, allow cargo to be transported via ship down one side of the channel [4].

3.4.3 Staggered

This is similar to the centred configuration, but each row is staggered such that the turbines in one row are aligned between the turbines of the adjacent row. This is illustrated in Figure 3.

3.4.4 Larger spacing

In the larger spacing case (see Figure 4), the spacing between each turbine in a given row remains the same as the centred case, but the spacing between each row is now $20d$ instead of $10d$.




Figure 3: The staggered turbine configuration. The whole $3 \times 1 \text{ km}^2$ channel is shown.




Figure 4: The larger spacing turbine configuration. The whole $3 \times 1 \text{ km}^2$ channel is shown.

3.5 Adjoint optimisation

Adjoint-based optimisation was employed to maximise the amount of power extracted by the turbines with respect to the individual turbine positions. Such optimisation of high-resolution transient simulations is computationally expensive. Therefore, in order for the adjoint problem to be feasibly solved with available resources, only the steady-state problem was optimised in this work.

The Dofin-Adjoint library [12, 13] was used to annotate the governing equations and automatically differentiate the forward model with respect to the optimisation variables, chosen to be the individual turbine positions here. The adjoint system was solved to obtain gradient information, which was then passed to the L-BFGS-B algorithm [7, 8] to optimise the functional describing the total power extracted by the turbine array, defined by

$$P(\mathbf{u}, \mathbf{m}) = \int_{\Omega} \rho c_t \|\mathbf{u}\|^3 dx, \quad (9)$$

where ρ is the density of water ($1,000 \text{ kgm}^{-3}$), \mathbf{m} is the vector of optimisation variables (in this case, the coordinates of each turbine), and Ω denotes the whole domain [1].

The staggered array configuration was used as the initial condition for the optimisation process, since this was found to be the best man-made configuration (see Section 4). The optimisation was constrained such that turbines must be at least $1.5d$ apart, where d is the turbine diameter, and can only be placed in the region defined by $500 \leq x \leq 2,500 \text{ m}$ and $125 \leq y \leq 875 \text{ m}$ (i.e. the area of highest mesh resolution). The optimisation was terminated after 100 iterations, resulting in the configuration shown in Figure 5.

4 RESULTS

4.1 Steady-state simulations

The results from the steady-state simulations in Figure 6 show that the centred array was the poorest choice of configuration for this setup. The power extracted by such a setup was limited because, apart from the turbines in the first row, all the turbines downstream were situated in the wakes/separated flow regions of the ones upstream (as shown in Figure 7). By definition there is a velocity deficit in these regions and thus, from equation (9), the power extracted was reduced. When a larger turbine spacing was used, the wake was able to recover before interacting with the downstream turbines (as shown in Figure 8), and thus 23% more power could be extracted from the flow. By offsetting the centred array configuration such that the turbines are closer to a side wall, the power extraction was increased, albeit fairly insignificantly, relative to the centred configuration. The staggered array configuration yielded the best power extraction potential,




Figure 5: The optimised turbine configuration from the steady-state simulation after 100 optimisation iterations. The whole $3 \times 1 \text{ km}^2$ channel is shown. The optimisation process was constrained such that turbines were not allowed to be placed outside of the dotted region shaded in blue. For the steady channel flow considered here a symmetric configuration would be expected, but this is not the case since the optimisation process was terminated after 100 iterations and had not fully converged (see Figure 10). Asymmetry may also be expected if the optimisation gets stuck in a local minimum.

generating 75% more power than the centred configuration. As Figure 9 shows, the accelerated region of flow around the first row of turbines interacts with the second (staggered) row behind it, resulting in more energy being removed from the flow and thus more power extraction. Each of these findings agrees qualitatively with the transient simulations of [4] which showed that, relative to the centred array, only a 4% increase in power extraction could be obtained with the offset array, whereas a 31% and 54% increase could be obtained with the larger-spaced and staggered configurations, respectively.

Since the staggered array yielded the most power, it was used as the initial condition and improved on throughout the adjoint optimisation process. This yielded a configuration that could extract 48.12 MW of power — a 156% increase relative to the centred array, as shown in Figure 10. This is greater than any difference in power extraction between the simple man-made configurations. Relative to the staggered array (i.e. the optimisation’s initial condition and the best man-made configuration considered here), 46% more power could be extracted.

Flow around a turbine results in acceleration around its edges, but also in a velocity deficit immediately behind the turbine (i.e. in its wake). It would be unwise to place a turbine in the wake of another since the power is related to the velocity. It therefore makes sense that the optimisation process re-positioned the turbines such that the effect of wakes from other turbines is minimised. Furthermore, regions of accelerated flow can be used to harness more power which is why some turbines are packed closer together in the cross-flow direction (e.g. the row of four turbines on the right-hand side of Figure 5). This is highlighted in Figure 11.

4.2 Transient simulations

Unlike the steady-state simulations which featured highly dissipated wakes due to the presence of a uniformly high background viscosity, the transient simulations exhibited longer, sharper regions of separation with some downstream turbulent eddy shedding as Figure 12 demonstrates. From the power profiles in Figure 6, the oscillations caused by turbulent eddies were relatively small. However, considerably less power was extracted compared to the steady-state setups, after the initial spin-up period; the initial peak was caused by the inflow condition of 2 ms^{-1} superimposing itself with a region of 2 ms^{-1} outflow generated by the Flather boundary condition, which eventually exited the domain.

With respect to the differences in power between the configurations, qualitatively similar behaviour was observed. For example, in the centred configuration the turbine rows were still

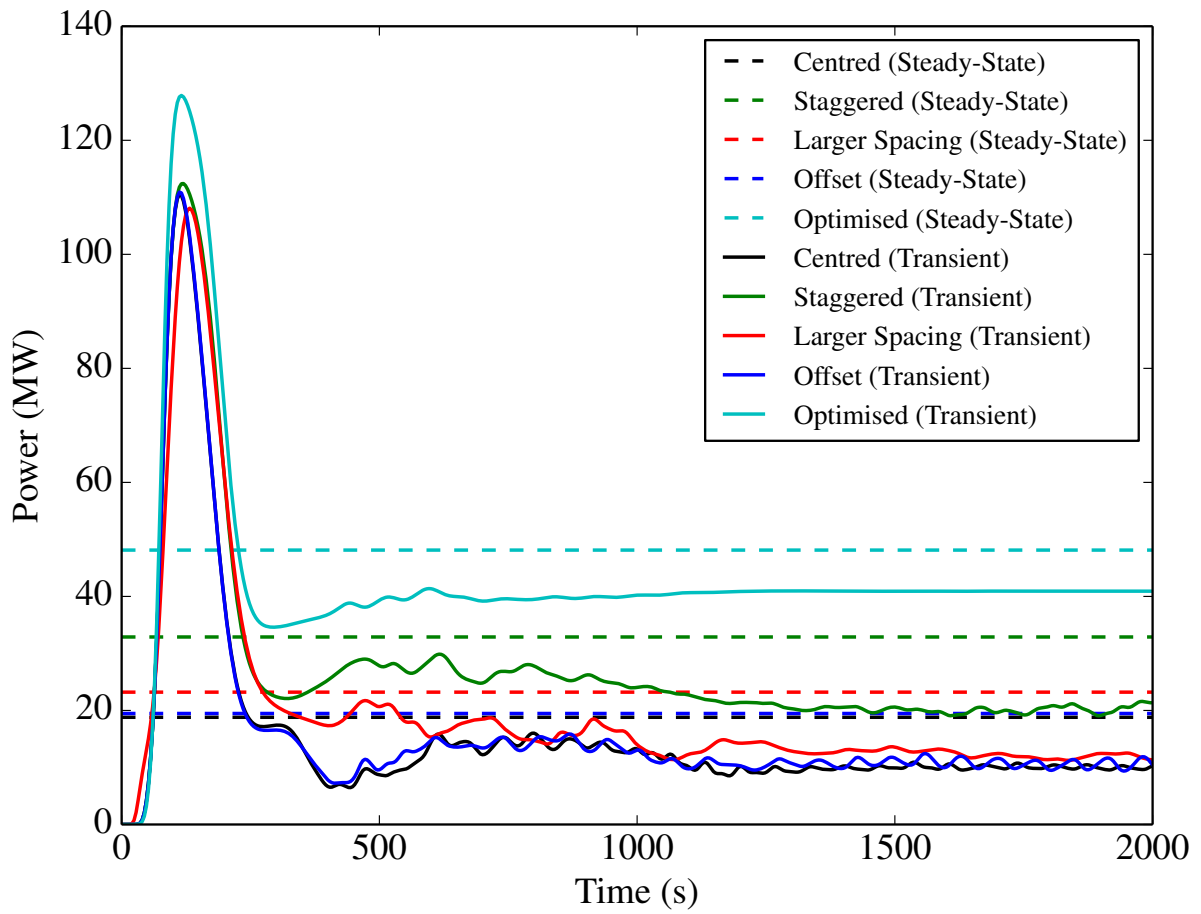


Figure 6: The total power generated against time, for each of the steady-state and transient simulations.




Figure 7: The flow speed around the centred turbine array in the steady-state simulation.




Figure 8: The flow speed around the array with larger turbine spacing in the steady-state simulation.




Figure 9: The flow speed around the staggered turbine array in the steady-state simulation.




Figure 10: Total power generated at each adjoint optimisation iteration. Iteration 0 represents the power from the initial staggered configuration.




Figure 11: The flow speed around the optimised turbine array in the steady-state simulation. Note the increased upper limit of the colour bar.




Figure 12: The flow speed around the centred turbine array in the transient simulation. Note the decreased lower limit of the colour bar.




Figure 13: The flow speed around the optimised turbine array in the transient simulation. Note the increased upper limit of the colour bar as a result of increased flow speed around the turbines on the right-hand side that are situated very close together.

sufficiently close to be in the wakes of the upstream turbines. However, in the transient runs the velocity deficit was much stronger, dropping to a near-zero flow speed immediately behind the turbines, resulting in decreased power generation. In the simulation with larger turbine spacing, the improvement in power extraction relative to the centred array was not as significant compared to the steady-state setup, because the wakes were longer and sharper. The turbines should therefore be placed even further apart to allow the wake to sufficiently recover.

Once again, the optimised array configuration out-performed all four man-made configurations under consideration. Interestingly, many of the wakes were entrained into the regions of acceleration and flowed around the downstream turbines as Figure 13 shows, instead of immersing them in the region of low flow speed and limiting power generation. Unlike the other configurations, turbulent eddies were only shed downstream of all the turbines, which explains why the power in Figure 6 remained constant once the flow became fully established.

5 CONCLUSIONS

- Adjoint optimisation can be used to greatly improve the power extraction potential of tidal turbine arrays/farms.
- The use of a lower background viscosity and turbulence modelling shows that the velocity deficit is much stronger in reality, resulting in decreased power generation relative to the high viscosity steady-state runs. Furthermore, turbulent eddies have the potential to affect the amount of power extracted, to a small degree in the configurations considered here.
- The power curves of the transient, turbulent runs are qualitatively similar to the steady-state runs, in that the best (staggered) man-made setup is still out-performed by the optimised configuration from adjoint modelling.
- Horizontal LES (HLES) [23] or 3D LES turbulence models, coupled with more realistic turbine parameterisations (e.g. rectangular blocks of drag as opposed to smooth Gaussian functions), should be adopted in future work.

6 ACKNOWLEDGEMENTS

The authors acknowledge the support and use of the Imperial College High Performance Computing Service and the UK National Supercomputing Service (ARCHER). CTJ was supported by ARCHER eCSE grant 03-7. SWF is supported by The Research Council of Norway through a Centres of Excellence grant to the Center for Biomedical Computing at Simula Research Laboratory, project number 179578. The authors would also like to acknowledge the support of the UK Engineering and Physical Sciences Research Council under projects EP/J010065/1, EP/L000407/1 and EP/M011054/1.

REFERENCES

- [1] S. W. Funke, P. E. Farrell, and M. D. Piggott. Tidal turbine array optimisation using the adjoint approach. *Renewable Energy*, 63:658–673, 2014.
- [2] 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.

- [3] D. M. Culley, S. W. Funke, S. C. Kramer, and M. D. Piggott. Tidal stream resource assessment through optimisation of array design with quantification of uncertainty. In *Proceedings of the European Wave and Tidal Energy Conference (EWTEC) 2015*, 2015.
- [4] T. Divett, R. Vennell, and C. Stevens. Optimization of multiple turbine arrays in a channel with tidally reversing flow by numerical modelling with adaptive mesh. *Philosophical Transactions of the Royal Society A*, 371(1985):1471–2962, 2013.
- [5] D. M. Culley, S. W. Funke, S. C. Kramer, and M. D. Piggott. A hierarchy of approaches for the optimal design of tidal turbine farms. In *Proceedings of the 5th International Conference on Ocean Energy*, 2014.
- [6] G. L. Barnett, S. W. Funke, and M. D. Piggott. Hybrid global-local optimisation algorithms for the layout design of tidal turbine arrays. *Renewable Energy*, Submitted.
- [7] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [8] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)*, 23(4):550–560, 1997.
- [9] S. W. Funke. *The automation of PDE-constrained optimisation and its applications*. PhD thesis, Imperial College London, 2012.
- [10] 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*, 85:215–227, 2016.
- [11] A. Logg, K.-A. Mardal, G. N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012.
- [12] P. E. Farrell, D. A. Ham, S. W. 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.
- [13] S. W. Funke and P. E. Farrell. A framework for automated PDE-constrained optimisation. *ACM Transactions on Mathematical Software*, Submitted.
- [14] J. Smagorinsky. General Circulation Experiments with the Primitive Equations. *Monthly Weather Review*, 91(3):99–164, 1963.
- [15] J. Deardorff. A numerical study of three-dimensional turbulent channel flow at large Reynolds numbers. *Journal of Fluid Mechanics*, 41(2):453–480, 1970.
- [16] J. W. Deardorff. On the magnitude of the subgrid scale eddy coefficient. *Journal of Computational Physics*, 7(1):120–133, 1971.
- [17] S. Nash, A. I. Olbert, and M. Hartnett. Towards a Low-Cost Modelling System for Optimising the Layout of Tidal Turbine Arrays. *Energies*, 8:13521–13539, 2015.

- [18] T. Divett, R. Vennell, and C. Stevens. Channel scale optimisation of large tidal turbine arrays in packed rows using large eddy simulations with adaptive mesh. In *Proceedings of the 2nd Asian wave tidal energy conference, Tokyo*, 2014.
- [19] T. Divett, R. Vennell, and C. Stevens. Channel-scale optimisation and tuning of large tidal turbine arrays using LES with adaptive mesh. *Renewable Energy*, 86:1394–1405, 2016.
- [20] C. Geuzaine and J.-F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.
- [21] R. A. Flather. A tidal model of the northwest European continental shelf. *Memoires de la Société Royale des Sciences de Liège*, 10(6):141–164, 1976.
- [22] M. G. Larson and F. Bengzon. *The Finite Element Method: Theory, Implementation, and Applications*. Springer, 2013.
- [23] G. A. Liek. Horizontal Large Eddy Simulation using Delft2D-MOR: the influence of large horizontal eddies on the shape, depth and extent of scour holes. Master’s thesis, TU Delft, 2000.