Fast and Effective Multi-objective Optimisation of Submerged Wave Energy Converters

Dídac Rodríguez Arbonès¹, Boyin Ding², Nataliia Y. Sergiienko², and Markus Wagner^{3(\boxtimes)}

¹ Datalogisk Institut, University of Copenhagen, Copenhagen, Denmark didac@di.ku.dk

² School of Mechanical Engineering, The University of Adelaide, Adelaide, Australia ³ School of Computer Science, The University of Adelaide, Adelaide, Australia markus.wagner@adelaide.edu.au

Abstract. Despite its considerable potential, wave energy has not yet reached full commercial development. Currently, dozens of wave energy projects are exploring a variety of techniques to produce wave energy efficiently. A common design for a wave energy converter is called a buoy. A buoy typically floats on the surface or just below the surface of the water, and captures energy from the movement of the waves.

In this article, we tackle the multi-objective variant of this problem: we are taking into account the highly complex interactions of the buoys, while optimising the energy yield, the necessary area, and the cable length needed to connect all buoys. We employ caching-techniques and problem-specific variation operators to make this problem computationally feasible. This is the first time the interactions between wave energy resource and array configuration are studied in a multi-objective way.

Keywords: Wave energy \cdot Multi-objective optimisation \cdot Simulation speed-up

1 Introduction

Global energy demand is on the rise, and finite reserves of fossil fuels, renewable forms of energy are playing a more and more important role in our energy supply [11]. Wave energy is a widely available but largely unexploited source of renewable energy with the potential to make a substantial contribution to future energy production [3,9]. There are currently dozens of ongoing wave energy projects at various stages of development, exploring a variety of techniques [10].

A device that captures and converts wave energy to electricity is often referred to as a wave energy device or wave energy converter (WEC). One common WEC design is called a point absorber or buoy. A buoy typically floats on the surface or just below the surface of the water, and it captures energy from the movement of the waves [9]. In our research, we consider three-tether WECs (Fig. 1) as a technological alternative to the common single-tether WECs. While their



Fig. 1. Schematic representation of a three-tether WEC [18].

capital cost are higher than of conventional single-tether heaving buoys, they can extract significantly more energy from the waves [16]. In our case, the buoys are fully submerged and tethered to the seabed in an offshore location. They use the motion of the waves to drive a hermetically sealed hydraulic line to drive hydroelectric turbines to generate electricity, or to power a reverse osmosis desalination plant to create potable water.

A single wave energy converter can only capture a limited amount of energy alone, which is why it is essential to deploy wave energy devices in large numbers. A group of wave energy devices is commonly referred to as a wave energy farm or array [2]. In order to evaluate our arrays, we use a recently developed frequency domain model for arrays of fully submerged three-tether WECs [15]. This model allows us to investigate different parameters, such as number of devices and array layout. In addition to the objective of producing energy, we consider the following two objectives: the cable length needed to connect all buoys as given by the minimum spanning tree, and the area of the convex hull needed to place all buoys. The ideal choice of parameters leads to an optimisation problem: what are the best trade-offs of the buoys' locations, the area needed, and the cable length needed? To the best of our knowledge, this study is the first to investigate this question to reduce costs and to increase efficiency.

We proceed as follows. In Sect. 2, we introduce the multi-objective buoy placement problem and the different objectives that are subject to our investigations. Then, we present in Sect. 3 our speed-ups, the operators, and the constraint handling used. We report on our computational study in Sect. 4 before we conclude with a summary.

2 Preliminaries

In the following, we outline the different objectives and constraints that we consider for the WEC array optimisation.

Let $X = \{x_1, \ldots, x_n\}$ and $Y = \{y_1, \ldots, y_n\}$ be the set of x and y coordinates of n WECs in the plane. The goal is to find a set of coordinates such that the energy output of the whole wave farm is maximised. At the same time, the total length of the cable or pipe necessary to interconnect the buoys, as well as the area necessary for the wave farm, should be minimised.

The WEC design that we consider is a fully submerged spherical body connected to three tethers that are equally distributed around the buoy hull (Fig. 1). Each tether is connected to the individual power generator at the sea floor, which allows to extract power from surge and heave motions simultaneously [14].

2.1 Power Output Prediction

In the following, we briefly outline the model of this kind of WECs arrays as it was derived by Sergiienko et al. [15] and used by Wu et al. [18].

The dynamic equation of the WECs array is derived in the frequency domain using linear wave theory, where a fluid is inviscid, irrotational and incompressible [4]. This model considers three dominant forces that act on the WECs:

- (i) excitation force includes incident and diffracted wave forces;
- (ii) radiation force acts on the oscillating body due to its own motion;
- (iii) power take-off force that exerts on the WEC from machinery through tethers.

The key point in the array performance is the hydrodynamic interaction between buoys that can be constructive or destructive depending on the array size and geometry.

Assuming that the total number of devices in the array is n and p is the body number, then the dynamics of the p-th WEC in time domain is described as:

$$\mathbf{M}_{p}\ddot{\mathbf{x}}_{p}(t) = \mathbf{F}_{exc,p}(t) + \mathbf{F}_{rad,p}(t) + \mathbf{F}_{pto,p}(t), \tag{1}$$

where \mathbf{M}_p is a mass matrix of the *p*-th buoy, $\ddot{\mathbf{x}}_p(t)$ is a body acceleration vector in surge, sway and heave, $\mathbf{F}_{exc,p}(t)$, $\mathbf{F}_{rad,p}(t)$, $\mathbf{F}_{pto,p}(t)$ are excitation, radiation and power take-off (PTO) forces respectively. The power take-off system is modelled as a linear spring and damper for each mooring line with two control parameters, such as stiffness K_{pto} and damping coefficient B_{pto} .

In case of multiple bodies, where $p = 1 \dots n$, Eq. (1) can be extended to include all WECs and expressed in frequency domain:

$$\left(\left(\mathbf{M}_{\Sigma} + \mathbf{A}_{\Sigma}(\omega) \right) j\omega + \mathbf{B}_{\Sigma}(\omega) - \frac{\mathbf{K}_{pto,\Sigma}}{\omega} j + \mathbf{B}_{pto,\Sigma} \right) \hat{\mathbf{x}}_{\Sigma} = \hat{\mathbf{F}}_{exc,\Sigma}, \quad (2)$$

where subscript Σ indicates a generalised vector/matrix for the array of N bodies, $\mathbf{A}_{\Sigma}(\omega)$ and $\mathbf{B}_{\Sigma}(\omega)$ are radiation added mass and damping coefficient matrices that include hydrodynamic interaction between buoys, $\mathbf{K}_{pto,\Sigma}$, $\mathbf{B}_{pto,\Sigma}$ are the stiffness and damping block-matrices of the PTO system.

The total power absorbed by the array of WECs can be calculated as:

$$P_{\Sigma} = \frac{1}{4} (\hat{\mathbf{F}}_{exc,\Sigma}^* \hat{\mathbf{x}}_{\Sigma} + \hat{\mathbf{x}}_{\Sigma}^* \hat{\mathbf{F}}_{exc,\Sigma}) - \frac{1}{2} \hat{\mathbf{x}}_{\Sigma}^* \mathbf{B} \hat{\mathbf{x}}_{\Sigma}, \qquad (3)$$

where * denotes the conjugate transpose.

For more details on the model, we refer the interested reader to [15, 18].

2.2 Constraints and Assumptions

We have the following constraints placed on our optimisation. The first one enforces an upper bound on the area of the farm. This constraint ensures that we can only place a buoy i within a certain area, which is a realistic constraint for most layout problems. For a rectangular wave farm with length l and width w, this constraint is satisfied *iff*

$$0 \le x_i \le l \text{ and } 0 \le y_i \le w, 1 \le i \le n.$$
(4)

Because buoys can damage each other if they get too close, and also maintenance ships need to be able to navigate between them, the second constraint regulates the spacial proximity. It is satisfied iff

$$\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \ge 50 \text{ meters.}$$
 (5)

In addition to the above constraints, we assume that all WECs have the same power take-off characteristics.

2.3 Euclidean Minimum Spanning Tree

We use the Euclidean minimum spanning tree (MST) to calculate the minimum length of cable or pipe required to connect all buoys in a particular array configuration. It is computed by first constructing the complete graph on the set of points that represent the buoys and edge costs given by the Euclidean distance between any pair of buoys. Then, the minimum spanning tree for this graph is computed and used as an objective representing the costs of the cable or pipe length. Figure 2 displays an array layout as, as well as the minimum spanning tree, represented by lines joining each buoy.



Fig. 2. An example WEC array. The circles visualise the safety distance.

2.4 (Cost of the) Convex Hull

In our study, the cost of the convex hull is defined as the area contained by the set of points forming the convex hull. This value is the minimum land area that is required for a wave farm layout. Figure 2 displays a buoy layout, as well as the area (cost) of the convex hull shaded in grey.

3 Computational Speed-Up, Operators, and Constraint Handling

3.1 Speed-Up of Simulation

In order to make the simulations computationally feasible, and to make the best use of the available hardware, we reimplemented the PTO system in C++ and

Table 1. Runtime per evaluation in seconds (median of 20 runs). Ω is the set of frequencies ω used. In each cell, single-thread results are on the left and multi-thread ones on the right. *Laptop* denotes a computer with a Intel Core i7-4910MQ CPU (up to 3.9 GHz, used with 4 threads) and 32 GB RAM. *Server* is a compute server with four AMD Opteron 6348 CPUs (up to 3.4 GHz, used with 48 threads) and 128 GB RAM.

n	$ \Omega $	MATLAB		C++	
		Laptop	Server	Laptop	Server
4	50	29.28/11.64	56.15/3.54	2.83/0.98	5.48/0.38
9	25	80.96/31.14	153.75/9.48	8.32/3.05	16.28/0.98
16	25	262.86/97.58	508.63/29.88	29.31/9.94	55.42/3.23
25	25	658.21/239.37	1265.97/72.46	71.92/26.20	141.16/8.95

parallelised it with OpenMP [13]. Because the system is defined as a series, it is inherently parallelisable and a linear speed-up possible. Furthermore, OpenMP's framework allows for nested paralellisation, which further decreases the overall running time. The integrals are calculated with the GNU Scientific Library [5], which has support for integrals with singularities. After comparing the performance of complex linear system solving in C++ and MATLAB [12], the result from Eq. (2) is obtained from MATLAB. Note that the evaluation of Eq. (3) can only be parallelised for each of the frequencies ω considered.

The evaluation of a WEC array is time consuming even with parallelisation. The integral calculation is the bottleneck consuming upwards of 95 % of the running time. Wu et al. [18] used caching of integral computations because a large portion of these integrals are repeated, achieving a factor 7 speed-up in running time. We use the caching approach even more comprehensively, by caching results not only within a single layout evaluation like Wu et al., but by reusing them across multiple evaluations. This additional improvement can help if an optimisation algorithm modifies only part of a solution at each iteration, therefore reusing integrals computed in previous iterations.

In Table 1 we list the achieved time needed to compute the intra-buoy interactions. As we can see, the speed-ups (up to 142-fold) allow us to run significantly more evaluations if the overall available time is limited.

3.2 Problem-Specific Operators

As the problem is highly constrained due to a large number of buoys and the given safety margin around each buoy, the operators have to ensure that feasible placements are produced. We investigate the benefit of the two variation operators MOVEMENTMUTATION and BLOCKSWAPCROSSOVER by Tran et al. [17] over the commonly used Polynomial Mutation and Simulated Binary Crossover. The former pair was designed for wind turbine placement optimisation, where safety distance constraints and area constraints also need to be considered.

MOVEMENTMUTATION is an operator that does a local change to the current solution. For a randomly picked WEC, MOVEMENTMUTATION moves it to a randomly selected spot along a selected direction to a feasible location.

BLOCKSWAPCROSSOVER is designed to implant a randomly selected rectangular "block" of WECs from each of the two parents to produce two children, each with a varying degree of information from each parent. A repair operator is applied in case the number of WECs does not match to the target number.

Note that the fundamental difference between both the WEC positioning and the wind turbine positioning is that "shading" is the primary inter-turbine effect, while the primary inter-buoy effect is "phase shifting". However, as the operators do not consider these effects directly, we can apply them to our problem as well.

3.3 Constraint Handling

As described in Sect. 2.2, we consider area constraints and safety distance constraints in this study.

The area (box) constraint is enforced by applying a sinusoidal-shaped function that maps any value to a closed range. The function used has the form [6]:

$$x = a + (b - a) * (1 + \cos(\pi * x/(b - a) - \pi))/2$$
(6)

The advantage of this function is twofold. First, the boundaries of the region are automatically enforced without the need for a check on each iteration. Second, it provides a smooth transition of the movements of buoys close to the boundaries, contributing to the performance of the optimisation algorithm.

The inter-buoy distance is enforced by applying a penalty to the objectives. This penalty is proportional to the distance that the buoys lie outside the safety margin. The resulting objectives O' are used in the optimisation process:

$$O' = O\left(1 - K \sum_{i=1}^{n} \sum_{j \neq i}^{n} \max(M - d(i, j), 0)\right)$$
(7)

where n is the number of buoys, M is the safety distance to keep between buoys, d(i, j) is the Euclidean distance between the buoys, and $K \in \mathbb{R}^+$ is the penalty regularisation parameter. This parameter is meant to control the slope of the penalty applied, acting as a trade-off between discouraging solutions that lie far into the infeasible region, and allowing the exploration of boundary regions.

4 Experimental Study

In this section, we describe our experimental setup and report on the results of different multi-objective evolutionary algorithms using our speed-ups and variation operators for the multi-objective buoy placement problem.

4.1 Experimental Setup

For the basis of our study, we utilise the algorithms SMS-EMOA [1] and MO-CMA-ES [7], as implemented in the optimisation framework Shark 3.0 [8]. We use SMS-EMOA in two variants: (i) the default SMS-EMOA with Polynomial Mutation and Simulated Binary Crossover, and (ii) the problem-specific SMS-EMOA^{*} with MOVEMENTMUTATION and BLOCKSWAPCROSSOVER (see Sect. 3.2).

We use a population of size $\mu = 50$ for all experiments, and the evaluation budget for each run is 6000 evaluations. All other parameters are used with their default values in the Shark library. Unless stated otherwise, we report the results of 20 independent runs.

For all runs, we initialise the first population with regular grids that are scaled from the tightest grid to the most generous one where buoys are placed on the boundaries as well. In Fig. 3 we show an example, which also shows the non-linear effect that arises from the constraint handling of the box constraints (Eq. 4). The side-effect of this initialisation is that we already achieve right from the beginning a population of solutions that is guaranteed to be diverse in the size of the convex hull and in the length of the minimum spanning tree. In preliminary experiments, we observed that this approach performed better than one with random initial layouts.



Fig. 3. Initial population in the 25-buoy scenario. The $\mu = 50$ layouts are shown in different colours. (Color figure online)

The scenarios are defined as follows. The goal is to place 4, 9, 16, and 25 buoys subject to the three objectives in a quadratic area. We scale the area available with the number (considering an area of $20,000 \text{ m}^2$ per buoy), which results in squares with sides of length 283 m, 424 m, 566 m, and 707 m respectively.

To compare the performance of the different setups, we inspect the sets of trade-offs visually, and we employ the hypervolume indicator. To compute the latter, we rescale the final solution set into the unit cube that is defined by the extreme values (of feasible layouts) observed for each scenario.

In each scenario, we use the sea state, i.e. the wave frequency distribution, and the features of the buoys as defined in the single-objective investigations in [18], which allows us to compare results for 25 buoys. The WEC radius is a = 5 meters, and their power take-off characteristics are kept static. The mass of each buoy is equal to 0.85 times the mass of the displaced water. Ocean depth is chosen to be 30 m and all WECs are submerged 6 m to the centre of buoy.

To compute the power output of a solution, we need to choose a number of discrete wave frequencies from the wave spectrum. While Wu et al. [18] observed that a single frequency from the entire spectrum of waves can be used with reasonable accuracy during buoy placement optimisation, we prefer to use a significantly more time-consuming approach with 25 or 50 wave frequencies. This provides us with very accurate power output predictions. Also, this greatly

reduces the risk of unrealistic exploitation of local optima due to ill-conditioned scenarios, which we have observed in the one-frequency case.

4.2 Experimental Results

In the following, we compare the performance of the different multi-objective approaches. Figure 4 summarises the hypervolumes achieved by the final populations for the different scenarios. While the standard version of MO-CMA-ES outperforms the standard version of SMS-EMOA, both are easily outperformed (in terms of achieved hypervolume) by SMS-EMOA^{*}. It appears that even though the latter employs operators previously used in wind turbine placement optimisation, they are also beneficial in our case.

Note that MO-CMA-ES hardly benefits from the caching of simulation results, as it tends to sample new coordinates for all buoys every time. While this behaviour is typically an advantage, it was infeasible for us to apply MO-CMA-ES to the optimisation of the larger scenarios.

Exemplarily, we show in Fig. 5 how the average objective scores across the populations as they evolve over time. Interestingly, SMS-EMOA and SMS-EMOA^{*} behave quite differently, even though they differ only in their variation operators. For example, the standard SMS-EMOA performs best in terms of convex hull and length of the minimum spanning tree, but it produces on average



Fig. 4. Hypervolumes covered by the final populations. MO-CMA-ES results are missing for 16 and 25 buoys due to the unacceptable run-times.



Fig. 5. Evolution of objective scores over time in case of the 9-buoy scenario. Shown are the averages and the 95% confidence intervals. Note that we are minimising the negative power output, meaning that smaller values are better.



Fig. 6. Population with the highest power output layout for 25 buoys (left, computed by SMS-EMOA^{*}), with the three-dimensional space being projected twice into the two-dimensional space. The layout with the highest power output (3.590 MW) is shown in the middle, and it uses more than twice the area of the layout with the lowest output (3.270 MW) in this population. Note that the waves arrive from the top of the layouts. The circles in the middle and in the right show the 5 m-buoys (to scale).

the layouts with the worst power output. It appears that all three approaches explore quite different parts of the objective space across all runs. This is not only reflected by the final mean values, but also by the spread across different runs by the same algorithm. In summary, this shows to us that a decision maker should not blindly trust multi-objective performance indicators, but inspect the solution sets as well. In practice, a single trade-off layout has to be chosen for implementation, and even though the nature of the problem is multi-objective, the decision eventually boils down to hidden preferences or economic factors.

Lastly, we briefly compare our results for 25 buoys with the ones by [18]. We take their best result and recompute the power prediction more accurately using 25 frequencies instead of the single frequency they used. As a result, their best layout has a predicted power output of 3.459 MW. Compared to this, our best layout (Fig. 6) has an output of 3.590 MW, which is 3.8% better.

To conclude, we can see that the multi-objective optimisation of arrays of wave energy converters is feasible, if software engineering tricks are employed to speed-up the simulations, and if problem-specific variation operators are used. Also, it is important for engineers to explore different algorithms as these explore the objective spaces with different biases—and the consequence of this bias might matter to the decision maker.

5 Conclusions

Wave energy plays an increasing role in the energy supply world-wide. We have investigated the problem of placing wave energy converters on a given offshore area using different conflicting objective functions.

In a first step, we speeded up and parallelised the computations of the buoys' interactions, which resulted in a speed-up by a factor of up to 142 for 25 buoys. In order to improve the actual optimisation, we employed variation operators from the loosely related wind turbine problem. Interestingly, these problem-specific

operators proved to be effective in our setting as well, despite the intra-device interactions being fundamentally different.

The computational study shows that multi-objective evolutionary algorithms can be used for the multi-objective buoy placement problem; in particular, our best performing configuration even improves the power output upon a previous single-objective results by 3.8%. This improvement can translate into millions of additional dollars of income per year during the lifetime of the wave farm.

In the future, we will extend our research to optimising the individual power take-off characteristics of the buoys in addition to their position, as the effective sea state within the WEC arrays differs significantly from the state outside.

We made the code and results available online: http://cs.adelaide.edu.au/ ~optlog/research/energy.php

Acknowledgments. This work has been supported by the ARC Discovery Early Career Researcher Award DE160100850.

References

- Beume, N., Naujoks, B., Emmerich, M.: SMS-EMOA: multiobjective selection based on dominated hypervolume. Eur. J. Oper. Res. 181(3), 1653–1669 (2007)
- de Andrés, A., Guanche, R., Meneses, L., Vidal, C., Losada, I.: Factors that influence array layout on wave energy farms. Ocean Eng. 82, 32–41 (2014)
- Drew, B., Plummer, A., Sahinkaya, M.N.: A review of wave energy converter technology. Proc. Inst. Mech. Eng. Part A: J. Power Energ. 223(8), 887–902 (2009)
- Falnes, J.: Ocean Waves, Oscillating Systems: Linear Interactions Including Wave-Energy Extraction. Cambridge University Press, Cambridge (2002)
- GNU Scientific Library: Version 1.16 (2013). http://www.gnu.org/software/gsl/. Accessed 7 Apr 2016
- Hansen, N.: CMA-ES Source Code: Practical Hints. https://www.lri.fr/~hansen/ cmaes_inmatlab.html. Accessed 7 Apr 2016
- Igel, C., Hansen, N., Roth, S.: Covariance matrix adaptation for multi-objective optimization. Evol. Comput. 15(1), 1–28 (2007)
- Igel, C., Heidrich-Meisner, V., Glasmachers, T.: Shark. J. Mach. Learn. Res. 9, 993–996 (2008)
- Lagoun, M., Benalia, A., Benbouzid, M.: Ocean wave converters: state of the art and current status. In: IEEE International Energy Conference, pp. 636–641 (2010)
- López, I., Andreu, J., Ceballos, S., de Alegría, I.M., Kortabarria, I.: Review of wave energy technologies and the necessary power-equipment. Renew. Sustain. Energ. Rev. 27, 413–434 (2013)
- Lynn, P.A.: Electricity from Wave and Tide: An Introduction to Marine Energy. Wiley, Hoboken (2013)
- 12. MATLAB: Version 8.6.0 (R2015b). The MathWorks Inc., Natick (2015)
- OpenMP Architecture Review Board: OpenMP Application Program Interface, Version 3.0, May 2008
- Scruggs, J.T., Lattanzio, S.M., Taflanidis, A.A., Cassidy, I.L.: Optimal causal control of a wave energy converter in a random sea. Appl. Ocean Res. 42(2013), 1–15 (2013)

- Sergiienko, N.Y., Cazzolato, B.S., Ding, B., Arjomandi, M.: Frequency Domain Model of the Three-Tether WECs Array (2016). https://www.researchgate. net/publication/291972368_Frequency_domain_model_of_the_three-tether_WECs_ array. Accessed 7 Apr 2016
- Srokosz, M.A.: The submerged sphere as an absorber of wave power. Fluid Mech. 95(4), 717–741 (1979)
- Tran, R., Wu, J., Denison, C., Ackling, T., Wagner, M., Neumann, F.: Fast and effective multi-objective optimisation of wind turbine placement. In: Genetic and Evolutionary Computation, pp. 1381–1388. ACM (2013)
- Wu, J., Shekh, S., Sergiienko, N., Cazzolato, B., Ding, B., Neumann, F., Wagner, M.: Fast and effective optimisation of arrays of submerged wave energy converters. In: Genetic and Evolutionary Computation Conference (2016). Accepted for publication