Master's Thesis

Optimisation of Submerged Buoy Arrays for Improved Ocean Wave Energy Production

Prepared by Slava Shekh (a1649566)

Supervised by Dr. Markus Wagner

A thesis submitted in partial fulfilment of the requirements for the degree of Master of Computer Science

June 5, 2016

Abstract

With ever-increasing global energy demand and finite reserves of fossil fuels, renewable forms of energy are becoming increasingly important to consider. Wave energy is a widely available source of renewable energy that is being investigated by dozens of wave energy projects using a variety of techniques. One common design for a wave energy converter (WEC) is the point absorber or buoy, which floats on or below the water surface and captures energy from the movement of the waves. CETO is an example of fully submerged WEC consisting of buoys that are tethered to the seabed in an offshore location for harnessing wave energy.

One of the key aims of wave energy research is to maximise the power absorption of wave energy converters. Since a single WEC can only capture a limited amount of energy, large-scale wave energy production necessitates the deployment of WECs in large numbers called arrays. Despite the current body of research on WEC array optimisation, many of the devices being considered are semi-submerged or floating, while CETO is fully submerged beneath the ocean surface. We are not aware of any research into the optimisation of fully submerged WEC arrays.

In this thesis, we explore different methods for optimising arrays of fully submerged CETO buoys to maximise their energy production. We focus on the problems of finding optimal combinations of buoy radii, exploring the effect of buoy spacing on array performance, and identifying the highest performing buoy layouts. The findings show that larger buoys can result in greater energy yield but will potentially increase the cost of the system, while a mix of small and large buoys can be beneficial for increasing constructive buoy interactions. Furthermore, allowing for sufficient spacing between buoys can minimise destructive interference and increase power output. Certain buoy arrangements can even lead to constructive interference in the array, which results in more power absorption than the sum of using each buoy in isolation.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

With ever-increasing global energy demand and finite reserves of fossil fuels, renewable forms of energy are becoming increasingly important to consider [27]. 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 [13,25]. The idea of harnessing wave energy has been around for at least two centuries, with the first patent for a wave energy device being filed in 1799 by a father and son by the name of Girard [14]. However, it was not until the oil crisis of the 1970s and the publication of Stephen Salter's iconic paper in Nature [36] that interest in wave energy truly began to surge. Since that time, the utilisation of wave energy has continued to be a very active research area. There are currently dozens of ongoing wave energy projects at various stages of development, exploring a variety of techniques [13, 14, 25, 26].

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 the point absorber or buoy, which typically floats on the surface or just below the surface of the water, and captures energy from the movement of the waves [25]. An example of a point absorber is the CETO wave energy converter, developed by Carnegie Wave Energy and named after the Greek sea goddess Ceto [28]. The CETO system consists of one or more fully submerged buoys that are tethered to the seabed in an offshore location, as shown in Figure 1.1. These buoys use the motion of the waves to drive a pump which intakes and pressurises nearby sea water. The sea water is then piped onshore and used to either drive hydroelectric turbines

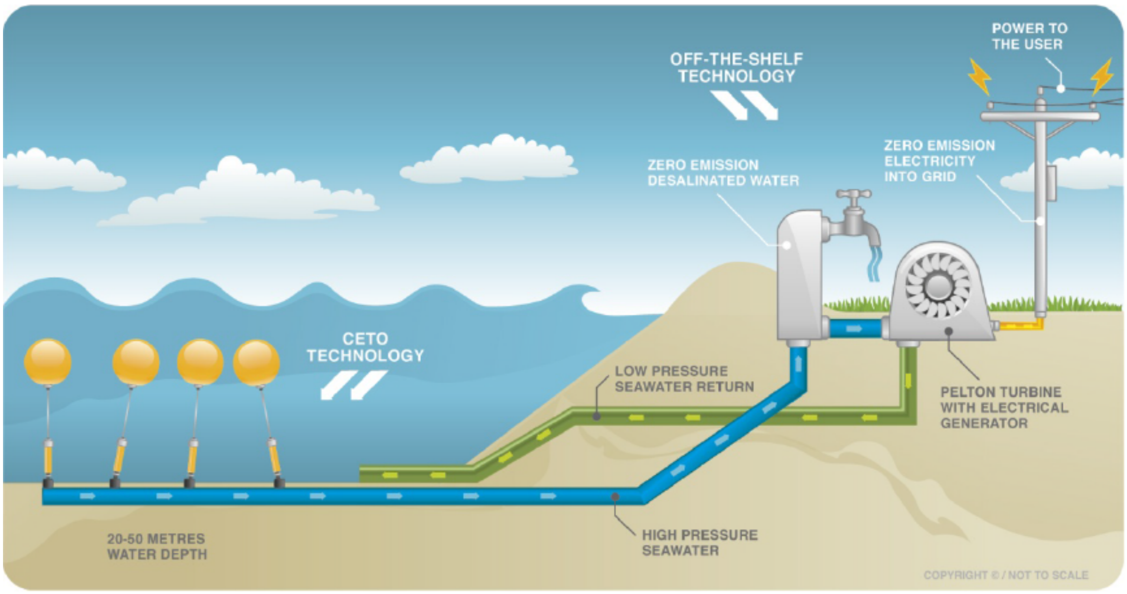

Figure 1.1: Operation of the CETO system [29].

to generate electricity, or to power a reverse osmosis desalination plant to create potable water [29].

One of the central goals in designing and operating a wave energy device is to maximise its overall energy absorption. As a result, the optimisation of various aspects of wave energy converters is an important and active area of research. Two key aspects that are often optimised are geometry and control. Geometric optimisation seeks to improve the shape and/or dimensions of a wave energy converter (or some part of it) with the objective of maximising energy capture [30, 32]. On the other hand, the optimisation of control is concerned with finding good strategies for actively controlling a WEC [34]. A suitable control strategy is needed for achieving high WEC performance in real seas and oceans, due to the presence and abundance of irregular waves [18].

A single wave energy converter can only capture a limited amount of energy alone. For large-scale wave energy production and in order to make any significant contribution to addressing global energy demand, it is essential to deploy wave energy devices in large numbers. A group of wave energy devices working in close proximity to one another is referred to as a wave energy farm or array [12]. Similar to individual wave energy devices, the optimisation of wave energy arrays is also an active area of research. In the case of arrays, the layout and configuration of the array is often optimised [8] as well as active control of individual devices [16].

Despite the current body of research on wave energy converter arrays and their optimisation, many of the devices under consideration are semi-submerged or floating [5, 8, 15]. In contrast, the CETO WEC is fully submerged beneath the ocean surface [29], and the role of the buoy is to pressurise sea water, while the actual power generation occurs onshore. There is very limited research into fully submerged wave energy converters, particularly devices with onshore electricity generation. For instance, there is research into control strategies for the Archimedes Wave Swing (AWS) [43]. Although the AWS is fully submerged below the ocean surface, the operation of the AWS is fundamentally different to CETO, partly because the electricity generation of the AWS occurs offshore. There is also a study that conducts a basic performance analysis of a number of WECs, including a bottom-referenced submerged heave-buoy which is conceptually similar to the CETO buoy [2], yet this study is concerned with the evaluation of WEC performance rather than optimisation. The author is only aware of one study on the optimisation of a fully submerged WEC, which investigates the effectiveness of different control strategies on the power absorption of two buoys, one of which is fully submerged [23].

Although there are limited studies on fully submerged wave energy converters, there is an even bigger gap in research on fully submerged arrays. The author is not aware of any research into optimising the placement or configuration of arrays of fully submerged wave energy converters. Therefore, the aim of this Master's thesis is to explore the most efficient methods for optimising arrays of fully submerged buoys for the production of renewable energy. The results will provide valuable insights into the optimal configuration and placement of buoys in a fully submerged array, which will help to address an important gap in the literature.

The remaining chapters of the thesis are structured as follows. Chapter 2 contains an overview of relevant literature, while Chapter 3 provides some background on the CETO array model and the criteria that can be used to optimise this model. Chapters 4, 5 and 6 then describe the radii, spacing and layout optimisation experiments that we conducted, along with our results and findings. Finally, Chapter 7 presents our conclusions and outlines several directions for future work.

Chapter 2

Literature Review

This chapter provides a review of literature relevant to this thesis. The review begins with a background on wave energy conversion, followed by a discussion of fully submerged wave energy converters, which is a variety of converter that is of particular interest to this thesis. Subsequently, literature on the optimisation of wave energy converters is presented, covering the areas of active control and geometric optimisation. Finally, research on wave energy converter arrays is discussed, particularly the work on array optimisation which is central to this thesis.

2.1 Wave Energy Conversion

The possibility of converting wave energy into electricity has been considered for at least two centuries, since the filing of the first patent for a wave energy device in 1799 [14]. Since that time, a wide variety of methods and techniques for wave energy conversion have emerged. The author of [31] provides an overview of the basic principles that are common to these techniques and groups existing conversion methods into nine categories. In addition, the author discusses relevant aspects of energy transmission, energy storage and the mooring of wave energy devices, which are all important considerations in the design and deployment of a wave energy converter. The existence of different techniques for harnessing wave energy has led to a variety of designs for wave energy converters. In an effort to categorise and compare these designs, a number of review papers have emerged [13, 14, 25, 26].

One review paper [13] discusses the potential benefits and challenges of wave energy, and provides a comparison of many existing wave energy converters with a focus on the United Kingdom (UK). The various WECs are categorised by location, mode of operation and type (either attenuator, terminator or point absorber). Attenuators lie parallel to the wave direction and move with the waves, while terminators lie parallel to the wave front and intercept the waves (see Figure 2.1). In contrast, point absorbers are generally much smaller devices that float on the surface of the water and use heave motion to absorb wave energy. In rare cases, point absorbers can be fully submerged below the surface, such as the CETO WEC that was shown in Figure 1.1. As well as categorising wave energy converters, the authors of [13] discuss two major challenges in current wave energy research and some commonly used approaches for addressing them. The two challenges are identifying the most effective power take-off (PTO) system, and actively controlling WECs to synchronise the natural frequency of the device with the dominant frequency of the wave in order to maximise device efficiency.

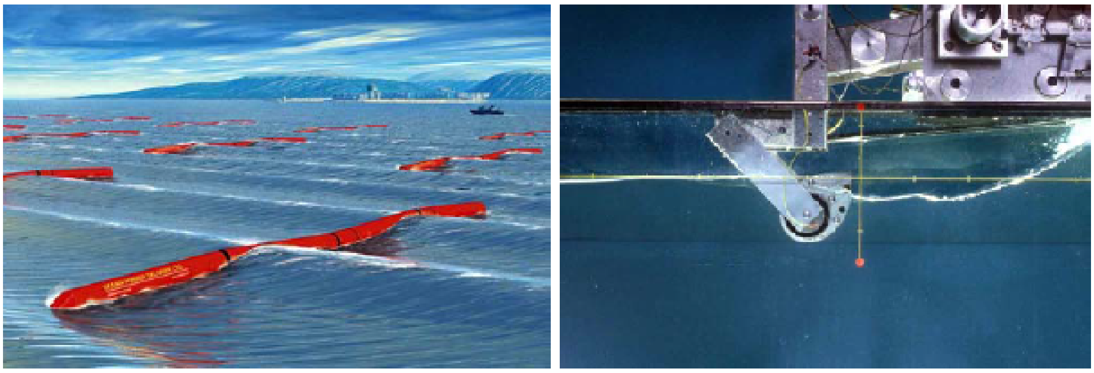

Figure 2.1: Examples of attenuator (left) and terminator (right) WECs [13].

The authors of [25] provide another review of existing wave energy projects, but their review considers projects worldwide, as opposed to the UK-centric focus of [13]. Wave energy converters are categorised by location and as one of six types, including some of the types used in [13]. Unique to this review paper is a discussion of the CETO project [29] and the differences of CETO to traditional point absorbers, including the fact that CETO is fully submerged and utilises onshore electricity generation. In addition to their review of existing wave energy converters, the authors provide a comparison of the power take-off and electricity generation systems used by some of these WECs.

In [14], a further review of wave energy utilisation is provided. This review includes a brief history on the development of wave energy devices from the first wave energy device patent in 1799 through to recent developments. The author estimates that there are around 100 active wave energy projects and describes some of the most significant projects throughout the paper. The WECs are categorised using a hierarchical classification based on their working principle, such as oscillating water column, oscillating body system and overtopping converter. In addition, the author discusses various WEC control strategies (also known as active control or phase control), particularly emphasising the need to control point absorber devices due to their small size. There is also discussion of the different power take-off systems used by WECs, which the author argues is one of the most important aspects of the system and aligns with the discussion in [13].

The review paper in [26] provides the most up-to-date snapshot of wave energy converter development compared to [13, 14, 25]. The authors discuss a wide range of devices and classify them by location, working principle, size and directional wave characteristics. These directional characteristics are similar to the attenuator, point absorber and terminator classification used in [13]. Although there are many systems currently in development, only a small number are at high Technology Readiness Levels (TRL) suggesting that many current developments still lack maturity. The paper also provides a brief analysis of the wave energy resource available worldwide, with the Southern Hemisphere being favourable due to lower seasonal variations, and Australia and New Zealand having access to some of the richest wave energy resources in the world. The later sections of the paper discuss the different stages of energy extraction, conversion and transmission to the grid, providing a detailed explanation of the various methods and technologies used by the current generation of wave energy devices.

2.2 Fully Submerged Wave Energy Converters

Although there are many existing types of wave energy converters [13, 25], very few of them are fully submerged below the surface of the ocean. Two key examples of fully submerged WECs are CETO [29], which is the device being considered in this

Master's thesis, and the Archimedes Wave Swing (AWS) [43], which is submerged like CETO but has fundamentally different operating characteristics. These two wave energy converters are further discussed below.

One of the earliest publications on the CETO wave energy converter was in 2007 in the European Wave and Tidal Energy Conference [28]. This publication describes the design, development and operating principles of this unique wave energy converter. More recently, a book chapter has been published discussing various aspects of the CETO Wave Energy Project [29]. The CETO system uses the motion of the waves to drive a pump which intakes and pressurises nearby sea water. The sea water is then piped onshore and used to either drive hydroelectric turbines to generate electricity, or to power a reverse osmosis desalination plant to create potable water. The author compares CETO to other commercial WEC technologies with one of the key differences being that the CETO device is located offshore and pressurises sea water, rather than generating electricity, which actually occurs onshore in the CETO system. The other key difference of CETO to traditional point absorbers is the full submersion of the device below the ocean's surface. As alluded to earlier, this aspect of CETO is uncommon though not necessarily unique, as the AWS is also a fully submerged point absorber WEC [43].

The authors of [43] describe the Archimedes Wave Swing (shown in Figure 2.2) and compare a number of control strategies for maximising its energy extraction. By modelling and simulating the various control strategies, the authors show that a non-linear control strategy called feedback linearisation control is best for maximising energy absorption. The AWS is particularly relevant to this Master's thesis, because it is one of the few wave energy converters that is fully submerged below the ocean surface. However, the operating principles are very different to CETO, as the AWS compresses air to generate electricity which is then transferred onshore, while CETO compresses sea water offshore and then transfers this sea water onshore for electricity generation [29]. These key design differences mean that any outcomes of AWS modelling and simulation cannot be directly applied to CETO, yet such outcomes can still provide valuable insights.

Another publication relevant to CETO is a study of the performance of eight wave energy devices across five European sites with different availability of wave

Figure 2.2: The Archimedes Wave Swing (AWS) wave energy converter [25].

resource [2]. The devices used in the study are based on real WECs, where each one has a different working principle. In particular, the bottom-referenced submerged heave-buoy (Bref-SHB) was inspired by the CETO WEC [29]. The study involved a numerical assessment of the annual power absorption for each device at each of the five sites. To provide some indication of cost, the assessment also included the annual power absorption per characteristic mass, surface area and PTO force of the eight devices. The findings showed that although the Bref-SHB has a low annual power absorption compared to most of the other devices, all devices including the Bref-SHB had similar energy absorption per mass, surface area and PTO force. This suggests that the estimated cost to power absorption ratio is comparable for all of the WECs in the study.

2.3 Optimisation of Wave Energy Converters

As suggested by the diversity of WEC designs [13, 25], there is a general lack of consensus on the best technique for wave energy extraction. In order for wave energy to become a viable and competitive renewable energy source, there is a need to find the most efficient extraction techniques and maximise their energy capture. This has led to research into the optimisation of wave energy converters, with much of the focus on the optimisation of geometry [30, 32] and active control [18, 34]. Research on the optimisation of these two aspects of wave energy converters is discussed in the following subsections.

It should be noted that optimisation is not only limited to wave energy conversion, and in fact [4] provides a review of over two hundred papers that have applied different single-objective and multi-objective optimisation techniques to solving problems in the field of renewable energy. Although the focus of some of these techniques may be on other forms of renewable energy, many of the techniques themselves, such as genetic algorithms, can be applied to wave energy conversion. In their review, renewable energy consists of wind power, solar energy, hydropower, bioenergy, geothermal energy and hybrid systems that use multiple forms of renewable energy. Ocean wave energy is categorised as a form of hydropower and they provide numerous instances where optimisation techniques have been applied to wave energy. Some examples include using artificial neural networks for predicting water levels, using stochastic optimisation to improve the energy production of a wave to air turbine, using genetic algorithms to optimise the shape of wave energy devices, and using genetic algorithms to optimise the layout of arrays of wave energy converters. The latter two examples have the strongest relevance to this Master's thesis and are further discussed in sections 2.3.2 and 2.4.2.

2.3.1 Active Control

One of the aspects of wave energy converters that is commonly optimised is control. This is also known as optimal control, because the goal is to actively control a WEC in order to optimise its energy absorption. Finding the best control mechanism is an essential part of achieving high WEC performance in irregular waves which are found in real oceans [18].

A relatively early paper on control optimisation [35] considers optimising a partially submerged heaving point absorber. In this context, heaving refers to the fact that the device captures energy from the "up and down" motion of the waves. In particular, the authors focus on optimising the damping element of the power take-off system using a control mechanism called latching. The purpose of latching is to artificially delay the velocity profile of a point absorber in order to synchronise it with the force profile. A genetic algorithm is used to find the damping profile that maximises energy absorption over the period of an incident wave, and the results show that this optimal damping profile is similar to that produced by the latching mechanism. This suggests that latching is an effective way of maximising energy extraction from the damping element, but the authors recommend that their work is extended to consider more realistic device and ocean models.

The authors of [33] extend on the previous work of [35]. Based on the finding that latching is an effective mechanism for optimising the damping element of the power take-off system, the authors develop a way of optimising the latching period, which they argue is the most important control variable when adopting the latching strategy. The outcome of their work is an equation which can be optimised using a standard minimisation routine to find the optimal latching period. This helps to maximise the energy capture of any wave energy devices that use the latching mechanism.

Whereas the works of [35] and [33] considered the optimisation of a point absorber, [34] describe the modelling of an oscillating water column wave energy converter and three control strategies for improving its energy capture. The first strategy involves optimising the characteristics of the Wells turbine, which is a variety of turbine that is often used in the energy conversion process of oscillating water column wave energy converters. This is shown to be effective in improving energy extraction and allowing the device to operate in a broader range of sea states. Another area of control is energy quality and using a turbine-based controller the authors showed that the variability in energy capture over time can be reduced. Finally, the authors find that using phase and amplitude control for the WEC results in a significant improvement to power absorption in regular waves, but their controller is found to be less effective in irregular waves, indicating an area for further research.

Both [35] and [34] identified limitations in their approaches in the context of real ocean waves. [18] addresses this problem by comparing a range of techniques for the control of a semi-submerged heaving point absorber in the presence of irregular waves, as found in real oceans. Most existing control techniques can be categorised into reactive control techniques, such as phase and amplitude control, and resistive bang-bang control techniques, such as latching [35]. The term “bang-bang” refers to a system where a control variable switches between the extremes of its value range, such as the latching mechanism which switches between open and closed. The different control techniques are tested in regular waves, where most techniques achieve close to optimal power absorption, and irregular waves using a total of nine different sea states. Under these irregular conditions, some techniques such as model-predictive control (MPC) are particularly effective for improving power absorption, while other methods such as latching are less effective, but are significantly simpler to implement and do not require any reactive power flow through the system. The authors also consider adaptive control based on changing sea state by using automatic parameter tuning. They find that adaptive control can provide improved power absorption over fixed control parameters.

More recently, [23] investigates the effectiveness of two control strategies on the power absorption of a floating and a submerged cylindrical heaving buoy. The two control strategies considered are real time tuning using wave surface elevation measurements, and non-real time tuning where the device damping is optimised in respect to the dominant wave frequency. Their results show that the real time tuning strategy performs best for both floating and submerged buoys, but particularly for the submerged buoy where it leads to a notable improvement in power absorption. This is one of the few works that give consideration to fully submerged buoys and hence their findings may be particularly relevant to this Master's thesis.

2.3.2 Optimal Geometry

The other aspect of WEC optimisation research involves finding ideal shapes and dimensions for wave energy devices or for specific parts of such devices. One example is the use a genetic algorithm to optimise the shape of a wave energy collector to improve energy capture [30]. The collector uses surge and pitch motion to extract energy, as opposed to heaving motion, such as used in [35]. The shape of the collector is modelled using a series of bi-cubic B-spline surfaces with a small number of control points. In addition, a simple cost function is used to assess the power absorption of a given shape using an incident wave that is weighed according to a wave occurrence distribution. The simplification of the B-spline surfaces and cost function help to ensure that the genetic algorithm can find a good solution in feasible time. The shapes produced by the genetic algorithm are compared to three benchmark shapes and the results show that the very best shapes of the genetic algorithm outperform the benchmarks.

Another example of optimising a specific part of a wave energy device is the optimisation of the airfoil shape of a Wells turbine [32]. Similar to [34], the authors of [32] argue that non-standard airfoil shapes could potentially lead to better performance. To verify their hypothesis, they use multi-objective optimisation to optimise the shape of the turbine in light of two conflicting objectives: turbine efficiency and tangential force (which is directly correlated with the power output of the turbine). The authors use an in-house multi-objective optimisation library called OPAL (OPtimization ALgorithms). The best airfoil shapes found by their optimisation algorithm show both a moderate increase in power output and a slight increase in turbine efficiency when compared to a standard airfoil shape.

Although the previous papers provide insights into some aspects of geometric optimisation, a more detailed study of the modelling, optimisation and control of the geometric aspects of wave energy converters is given in a Doctoral thesis on the topic [24]. In terms of the optimisation, the author uses multi-objective methods similar to [32], except instead of optimising a Wells turbine, the methods are used to assess different geometric configurations of an oscillating WEC. The multi-objective aspect of the problem arises from the need to consider both the power absorption and production cost of the wave energy converter. The results of their experiments show that geometric configurations which maximise power absorption are not always the most cost efficient, which highlights the importance of considering cost in WEC design. The author of [24] also considers varying the actual device geometry over time, which is referred to as geometric control. In this case, geometric control is achieved by adjusting the angle of a flap attached to the oscillating WEC based on the current sea state. Although geometric control is a relatively new area of research, the results are promising, showing that controlling the geometry of a device has the potential to further improve its power absorption.

Comparable to both [32] and [24], the authors of [22] use multi-objective optimisation to find a suitable geometric configuration for a cylindrical buoy. The conflicting objectives of their optimisation problem are maximising the power output of the buoy while minimising production cost by limiting the required volume of sheet plate. Using a technique called multi-objective genetic algorithm (MOGA), the authors find the Pareto-optimal set, which provides a number of alternate buoy configurations that represent different trade-offs of power output and cost. These findings provide useful insights into the optimal design of a resonant buoy, allowing a designer to choose the most suitable configuration based on the relative importance of the trade-off parameters.

2.4 Wave Energy Converter Arrays

The use of a single wave energy converter has limited utility, because there is only so much energy that it can capture alone. For large-scale wave energy production, wave energy devices need to be deployed in groups. A group of wave energy devices working in close proximity to one another is called a wave energy farm or array [12], such as the example shown in Figure 2.3.

Figure 2.3: An array of buoy wave energy converters [13].

The performance of an array is usually measured using the q-factor, as defined in Equation 2.1 [39] for an array of N devices. The power absorption of a device in isolation is represented as P_0, while the power absorption of a device n when placed in the array is represented as P_n. Therefore, the q-factor is the ratio of the power absorption of an array of wave energy converters compared to the power absorption of those same converters in isolation [8, 39]. A q-factor greater than one indicates the presence of constructive interference in the array, as the array of devices is producing more energy than the devices would individually. Conversely, a q-factor less than one is a sign of destructive interference, which may be detrimental to the performance of the array.

$$q = \frac{\sum_{n=1}^{N} P_n}{N P_0} \tag{2.1}$$

One of the first studies on arrays of wave energy devices was conducted in [7] and since then, the study of WEC arrays has become an active area of research. The following subsections explore some of the more recent developments in the field.

2.4.1 Factors Affecting Array Performance

Several studies have considered different factors that influence the power output of a farm of wave energy converters. One example is the study in [10], which considers the effect of device layout and control in the presence of irregular waves, such as those found in real seas. Using a farm consisting of four cylindrical heaving devices and a simple damping control mechanism, the authors find that both the control and layout of an array of devices has a significant impact on the overall energy yield of the farm. Given the simplicity of the scenario considered by the authors, optimisation of the array layout and use of more advanced control strategies could have an even bigger impact on the yield of a large wave energy farm. Additional findings and further discussion of the results can be found in the extended version of the paper [11].

In contrast to the numerical analysis of [10], the authors of [44] use experimental testing to investigate the device interactions in an array of 12 heaving WECs. In their experiments, they subject the array to both regular and irregular waves covering a range of sea states, and use the q-factor to measure device interaction. Their focus is on closely spaced arrays where devices experience the highest levels of interference with one another. The results show that the incident wave frequency and the location of a device in the array can have a significant impact on the performance of that device, including the potential for constructive or destructive interference.

In [42], a phenomenon known as a wake effect is investigated in the context of a wave energy farm consisting of overtopping wave energy converters. In an overtopping converter, the water from incident waves is captured in a basin located above sea level and then releases the water back into the ocean through turbines. Using a farm of nine WECs and two farm configurations (aligned and staggered), the authors investigate the effect of both short and long crested irregular waves on the efficiency of the farm. Overall, their results show that the staggered farm configuration is less affected by wake effects and results in higher energy absorption.

Extending on the work of [10] and [44], the authors of [6] conduct a study to examine the interactions that occur in larger arrays of wave energy converters. The study considers both the heaving cylinder and surging barge variety of WEC, arranging these devices in regular grids of 9–25 devices. Their goal is to investigate the impact of various factors on the annual energy production of the array. Over the course of a year and considering a range of wave periods, they find that the constructive and destructive interference of the devices in the array balances out. When considering such a long period, the positioning and interaction effects of devices in the array have limited significance. Instead, the authors recommend optimising the PTO towards maximising yearly energy production, and using wide-banded WECs in the array to allow for device interactions over a wider range of wave periods.

The author of [1] discusses the Park effect and its application to arrays of wave energy converters. Based on the literature discussed thus far, the term "Park effect" is not widely used. However, the term refers to the modification of the energy output of an array of WECs due to the constructive and destructive interference that occurs between the devices in the array. Based on a review of relevant literature, the authors conclude that the Park effect is negligible for small arrays consisting of less than 10 devices. However, in larger arrays they recommend that the number of rows of devices should be minimised and separating distances maximised in order to mitigate against the Park effect. The author also argues that wake effects are meaningless for WECs, because the wave field is modified in all directions rather than just the downstream region [1]. This is in conflict with other research on the impact of wake effects on wave energy farms [42], but the devices considered in this research were overtopping as opposed to oscillating.

A more recent study on wave energy converter arrays [12], analyses a range of factors that can affect the ideal layout of a wave energy farm, using a farm of partially submerged two-body buoys for their analysis. The assessed factors include the number of WECs (up to a maximum of four), array configuration (line, triangle, rhombus or square), separation distance, and wave directionality. Increasing the number of WECs is found to increase array efficiency, as long as the interference between devices is constructive, while the optimal separation distance is found to be a function of wave length. In terms of array configuration, the triangular and square configurations display the highest efficiency, while the linear configuration appears to be the worst. The effect of different wave climates and the wave directionality in these climates is also assessed, with the findings showing that the square and triangular configurations are best for unidirectional and multi-directional waves respectively.

2.4.2 Optimisation of Array Layout

Analogous to research on the optimisation of individual wave energy devices (Section 2.3), there is a body of research on optimising arrays of such devices. In the case of arrays, an aspect that is frequently optimised is the layout of the devices. As discussed at the start of Section 2.4, the performance of an array is usually measured using the q-factor, which provides some insight into whether the array is encountering constructive or destructive interference and the extent of this interference. In turn, highly constructive interference is indicative of high energy absorption and suggests that arranging the devices in such an array is worthwhile.

A preliminary study of optimal array configurations is conducted in [15] using an array of five semi-submerged spherical WECs. They model the array configuration as a constrained non-linear optimisation problem with the aim of maximising the q-factor. The optimisation problem is solved to find ideal configurations of both symmetric and non-symmetric arrays, which show an improved q-factor over linear arrays. The authors also assess the arrays on a range of wave incidence angles, and show that array configurations with highly constructive interference (i.e. high q-factor) tend to have destructive interference for a wider range of angles. This suggests that the most "optimal" configurations are also the most sensitive to the wave incidence angle, so the wave climate should be carefully considered during the placement of devices in real seas.

Building on the work of [15], the author of [8] explores two different methods for optimising the placement of semi-submerged point absorbers within an array. The two methods are Parabolic Intersection (PI), a heuristic method that constructs arrays based on linear, pentagonal and staggered configurations, and a genetic algorithm (GA) using a custom crossover operator specifically developed for this WEC optimisation problem. The array configurations created by the GA outperform those created by PI in terms of the q-factor, but the GA also requires significantly more computation time. Further details of this research are presented in the Doctoral dissertation [9], including a detailed discussion of the factors that affect WEC arrays and an analysis of array performance in irregular sea states, which proves to be favourable for the array configurations generated using the above methods.

Further optimisation of point absorber WEC arrays is considered in [39]. In alignment with related research [8, 15], the authors optimise performance by maximising constructive interference between devices in the array as measured by the q-factor. However, rather than finding the layout with the highest q-factor, they consider layouts that maximise the q-factor lower bound. Using exhaustive search, they find near optimal layouts that maximise this lower bound and discover that such layouts have certain characteristics, such as being symmetrical to wave direction. The authors then develop a heuristic algorithm for constructing these kind of layouts and show that their technique is able to produce high q-factor layouts over a range of sea states.

The work in [39] is extended in [40] to consider the problem of uncertainty. In realistic ocean environments, the uncertainty and stochasticity associated with the sea state often has a negative impact on the power output of a wave energy farm. To address this problem, the authors propose two models that optimise the layout of a wave energy farm in the presence of uncertainty. The first is a max-min model which maximises the worst case q-factor across all wave directions, while the other model maximises the expected q-factor based on a stochastic distribution of possible wave directions. Both models produce constructive array interaction for a larger range of sea states than a comparable deterministic model, suggesting that their performance is likely to be more robust in uncertain environments.

2.4.3 Other Array Optimisation Research

Besides optimising array layout, other ways to improve power absorption include optimising the power take-off system and actively controlling the devices in the array. A key example of the former is the optimisation in [5] for an array of floating buoy wave energy converters. Rather than optimising the configuration of the array as in [15], they allow each buoy in the array to have a unique damping value, as part of its power take-off system, with the aim of maximising the net power output. The Gauss-Newton algorithm is used to conduct a Least Squares analysis to find suitable damping values for each device such that the power absorption is maximised across a range of wave frequencies. The diversity of damping values across devices gives the array heterogeneous characteristics. An array consisting of five devices is tested in both head and beam seas, and the results show that tuning the damping value of each device individually can result in a significant increase in the net power output of the array, although they recommend further testing in real ocean environments.

Whereas [10] considers the impact of device control on array performance, the authors of [3] optimise two specific control strategies for maximising the power output of an array. The first is a global control strategy where control is centralised and the configuration of the entire WEC array is known to the controller, and the second is an individual control strategy where each device is controlled independently. In general, control strategies assist in maximising the positive interference between neighbouring WECs, which helps to maximise the overall power absorption of the array. This involves optimising the control strategies themselves, and in this case they are optimised by solving the corresponding quadratic programs using the active set algorithm. The authors assess the two optimised control strategies on three array layouts and find that global control provides a substantial performance improvement over individual control, particularly in layouts with the most device interaction, such as the closely packed equilateral triangle array.

Bringing together work on layout and control optimisation, the authors of [16] consider the impact of both layout and control on the positive interference and hence overall power absorption of a wave energy farm. In particular, the two factors are considered simultaneously to determine whether layout and control are interdependent, in contrast to the singular consideration of layout in [8, 15, 39] and control in [3]. Their experiments assess four different array layouts of 2–4 heaving cylindrical WECs using three distinct control strategies: passive, global and independent control. The results show that the choice of control strategy has a direct impact on the optimal array layout, and further that the use of any control strategy can significantly improve array performance in contrast to a complete absence of control. In terms of deploying wave farms in real seas, the sea state and wave direction are additional factors that should be considered when determining the optimal layout. This research shows that both array control strategies and array layout are important factors to consider when optimising a WEC array.

Finally, the authors of [17] introduce approximations into an analytical model to reduce computation cost and allow them to evaluate the power absorption and variance of wave energy "parks" containing up to 1024 wave energy devices. These parks are significantly larger than arrays considered in similar studies, where arrays are often limited to 25 devices or fewer [6,8,15]. In considering a number of different park geometries, their study shows that there are diminishing returns in adding more devices to large parks both in terms of increasing overall power absorption and minimising power variance. Although their model approximations introduce some error into the results, their findings may prove useful for any experiments involving very large WEC arrays.

2.5 Summary

Despite the current body of research, the applicability of this research to the CETO WEC is limited due to its unique design and complete submersion below the surface of the ocean. Although research into the optimisation of semi-submerged buoys such as [18, 33, 35] and buoy arrays such as [8, 15, 39] can provide useful insights, there is very limited research on fully submerged wave energy converters, particularly devices with onshore electricity generation. For example, there is research into control strategies for the fully submerged Archimedes Wave Swing (AWS) [43], but the operation of the AWS is fundamentally different to CETO, in part due to the offshore electricity generation of the AWS. There is also a study that conducts a basic performance analysis of a number of WECs, including a bottom-referenced submerged heave-buoy inspired by the CETO buoy [2], yet this study is concerned with the evaluation of WEC performance rather than optimisation.

In terms of the optimisation of fully submerged wave energy converters, the literature is even sparser. The author is only aware of one study, which investigates the effectiveness of different control strategies on the power absorption of two heaving buoys, one of which is fully submerged [23]. The related problem of finding optimal configurations of wave energy converter arrays has received some attention in recent years [8,15,39], but once again, the focus of this research has been on semi-submerged or floating arrays. To the best of our knowledge, there is a gap in research on optimising arrays of fully submerged buoys.

Chapter 3

Methodology

The literature review has shown that there is currently a lack of research on the optimisation of fully submerged buoy arrays. Therefore, this Master's thesis considers the following research question:

> *What are the most efficient methods for optimising arrays of fully submerged buoys for the production of renewable energy?*

3.1 CETO Array Model

To address the research question, we focus on the CETO WEC, a fully submerged buoy with unique characteristics such as onshore electricity generation (see Section 2.2 for further details). We use a MATLAB model of an array of CETO buoys which has been developed by the School of Mechanical Engineering at the University of Adelaide [37]. In their model, each buoy is attached to the sea floor using three tethers which are evenly spaced around the buoy and correspond to three separate power take-off systems. In order to simplify the optimisation problem, we place the following constraints on the CETO model:

1. A fixed water depth of 30m.

2. A fixed submergence depth for all buoys of 3m.

3. A power take-off with fixed parameters.

4. One direction of wave propagation.

The variable parameters of the CETO model include the size of the array, the radii of individual buoys, the spacing between buoys, and the positions of the buoys relative to one another (i.e. array layout). We optimise each of these parameters in turn through a series of computational experiments which are discussed in Chapters 4, 5 and 6 of this thesis.

3.2 Optimisation Criteria

In order to optimise any aspect of a CETO array, we need to define one or more optimisation criteria. The CETO model provides four output parameters which are suitable for this purpose:

1. q-factor
2. Relative capture width (RCW)
3. Average power absorption
4. Tether force

The simplest of these is the average power absorption of the buoys in the array, as shown in Equation 3.1. Given an array of N devices, P_n is the power absorption of a device n when operating in that array. The power absorption values are calculated across a range of wave frequencies and possible sea states to ensure that the results are robust to the kind of variability that occurs in real ocean environments. Naturally, the higher the average power absorption of the array, the better the array configuration. Yet, if we only consider the power characteristics of the array itself, we do not gain an understanding of the effectiveness of that array compared to the individual buoys. This forms the motivation for the q-factor as an optimisation criterion.

$$P = \frac{\sum_{n=1}^{N} P_n}{N} \tag{3.1}$$

The q-factor is defined as the ratio of the power absorption of an array of buoys compared to the power absorption of the same buoys in isolation. If an array configuration has a q-factor less than one, then there is destructive interference within the array and placing the buoys in such a configuration may be detrimental to the performance of the array. On the other hand, if an array configuration has a q-factor greater than one, then the array is potentially absorbing more wave energy than the individual buoys could absorb in isolation. These properties of the q-factor have led to it being a widely used metric for measuring the performance of WEC arrays [8, 15, 39] and it is also one of the criteria that we consider for our array optimisation problem.

Related work on wave energy converter arrays typically defines q-factor as per Equation 2.1 [39]. Since two given CETO buoys can have different radii, they can also differ in power absorption and therefore the CETO model uses a modified definition of q-factor (Equation 3.2). Similar to Equation 2.1, P_n is the power absorption of a device n when operating in an array of N devices. However, the original equation denominator has been modified to include the individual power absorption of each buoy, where P_i is the power absorption of device i in isolation. Once again, the q-factor is calculated over a range of wave frequencies and possible sea states.

$$q = \frac{\sum_{n=1}^{N} P_n}{\sum_{i=1}^{N} P_i} \tag{3.2}$$

Although q-factor is a useful measure of the positive interference in an array, it does not provide any indication of the absolute power being absorbed by that array. In some situations, an array configuration with a very small power absorption may have the most positive interference and therefore the highest q-factor compared to the other configurations. However, in practical settings it is not cost effective to deploy array configurations with low power output, even if they have very high q-factors. This limitation provides the motivation for the remaining optimisation criteria: relative capture width (RCW) and tether force.

RCW is a measure of the power extracted by each buoy in the array with respect to its size, as defined in Equation 3.3. Similar to previous optimisation criteria, P_n is the power absorption of each buoy weighted over a range of wave frequencies and possible sea states. However, the denominator is replaced with the product of P_w, the power absorption per unit frontage of a device, and the sum of device radii, where a_p is radius of the p-th device. RCW attempts to provide some balance between q-factor and absolute power absorption, by giving more preference to arrays consisting of larger devices, particularly those arrays which may have been overlooked in the q-factor calculation.

$$RCW = \frac{\sum_{n=1}^{N} P_n}{P_w \left(2 \sum_{p=1}^{N} a_p \right)} \tag{3.3}$$

In addition to the optimisation criteria outlined above, another possible metric is the capital cost of a wave energy converter. Imagine that there are two options for the design of a CETO buoy, and that one design has slightly higher power absorption and q-factor than the other, but that it is also significantly more expensive to construct. If we had infinite capital, we would simply select the design with the highest power or q-factor, regardless of cost. In reality however, cost can be a significant factor and it is worth considering how much benefit one design provides over another in light of the additional cost. Another term for this concept is "efficiency", which in the context of wave energy converters is a measure of how much energy a WEC captures with respect to the cost of construction, installation, operation, maintenance, insurance etc. The efficiency of wave energy converters has previously been explored in [32] for the Wells turbine used in various WECs, and in [38] for the layout optimisation of wave farms.

In the CETO model, the optimisation criterion related to cost is tether force. Also known as the power take-off (PTO) force, it represents the combined forces exhibited by the power take-off system [41], as shown in Equation 3.4. The tension applied on a tether t is expressed as the sum of the spring rate k and the damper rate d multiplied by the extension of that tether δ_t. The tether force of a buoy n is then the maximum tension across its three tethers, and the tether force of the array is simply the average tether force of all N buoys in that array. This calculation is again weighted over a range of wave frequencies and possible sea states.

$$T = \frac{\sum_{n=1}^{N} \left(\max_{t=1}^{3} \left(k\delta_t + d\delta_t \right) \right)}{N} \tag{3.4}$$

Generally speaking, the higher the tether force, the higher the cost of running the PTO system and operating the buoy [2]. Therefore, by minimising the tether force criterion we attempt to reduce the capital cost of the CETO system. Although outside the scope of this thesis, future work could combine tether force with average power absorption to provide some notion of CETO buoy efficiency in a similar way to the related works described above.

Chapter 4

Radii Optimisation

Using the model described in Section 3.1, we investigate the optimal choice of radii for buoys in differently sized CETO arrays. The following subsections define the radii optimisation problem, describe some suitable solution techniques, and then present our experimental results.

4.1 Problem Definition

A solution to the radii optimisation problem is represented as $\begin{pmatrix} r_1 & \dots & r_n \end{pmatrix}$, where r_k is the radius of buoy k and n is the total number of buoys in the array. Each buoy can have a radius of either 2m, 2.5m, 3.2m, 4m or 5m, and the radius of each buoy in the array can be different. Buoys are arranged in a staggered layout as shown in Figure 4.1. The radius of each buoy is stored based on its position in this layout, ordered from left-to-right, bottom-to-top. In the array shown in Figure 4.1, the 2m buoy appears first because it is located at the bottom of the left column of the array, and is then followed by the 2.5m, 4m and 5m buoys from left-to-right, bottom-to-top.

The goal is to optimise the radii of the buoys in the array with respect to the four optimisation criteria defined in Section 3.2. For q-factor, RCW and average power absorption, larger values indicate better array configurations, while smaller values of tether force are more desirable as these are expected to reduce the cost of the power take-off system.

Irrespective of the solution technique used, one or more starting points are also needed to initialise the search for solutions. A simple option is to start with a same-buoy configuration, such as $\begin{pmatrix} 2 & 2 & \dots & 2 & 2 \end{pmatrix}$ or $\begin{pmatrix} 5 & 5 & \dots & 5 & 5 \end{pmatrix}$. Since there are only five such configurations corresponding to each of the five possible buoy radii, it is even feasible to evaluate all same-buoy configurations and choose the best one as the starting point. Another option is to start with a random configuration, though depending on the algorithm, there may be some deviation in the results over multiple runs.

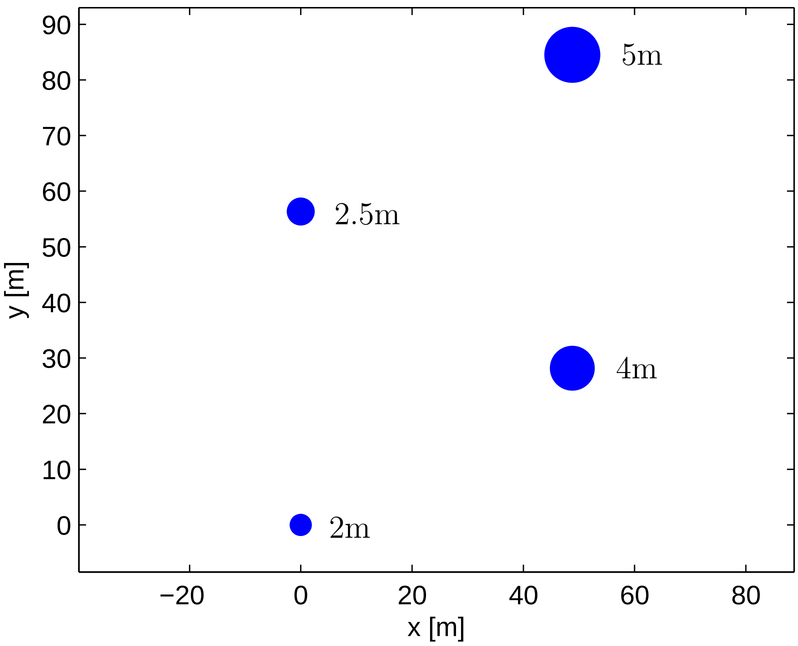

Figure 4.1: The $\begin{pmatrix} 2 & 2.5 & 4 & 5 \end{pmatrix}$ radii configuration of a 2x2 array using a staggered layout. The direction of wave propagation is from right to left.

4.2 Solution Techniques

Two algorithms have been investigated for solving the radii optimisation problem: brute force search (BFS) and random local search (RLS). Each approach is explained in more detail below, along with its benefits and drawbacks.

4.2.1 Brute Force Search

In the simplest case, a brute force search can be used to evaluate all possible radii combinations. This method provides a clear picture of the entire fitness landscape, including the best and worst radii combinations, as well as the characteristics that define these configurations. Since it is exhaustive in nature, using brute force search is only computationally feasible for small array sizes (supporting results will be presented in Section 4.3). Nonetheless, being able to examine the whole fitness landscape for smaller array sizes can provide valuable insights into the kind of characteristics that are needed to create high performing arrays of larger dimensions.

4.2.2 Random Local Search

In order to address the computational performance issues of brute force search, we have developed a random local search algorithm (Algorithm 1). The random local search starts with an initial array configuration and then changes (mutates) the radius of one random buoy in the configuration. If the fitness of the mutated configuration (offspring) is better than the fitness of the existing configuration (parent), then the offspring replaces the parent. This process is then repeated for a specified number of generations (maxGens). The fitness can be determined using any of the optimisation criteria outlined in Section 3.2.

Algorithm 1 Random Local Search

```
function RLS(initial, arraySize, maxGens)
    parent ← initial;
    for gen ← 1 to maxGens do
        buoy ← random(1 to arraySize);
        offspring ← mutate(parent, buoy);
        if fitness(offspring) > fitness(parent) then
            parent ← offspring;
        end if
    end for
end function
```

The mutation aspect of random local search can be implemented in a number of ways including an up/down mutation (UDM) or a fully random mutation (FRM). UDM increases or decreases the buoy radius by a single increment and cycles the values when required. For example, a 2.5m buoy could be mutated to a 3.2m or 2m buoy, while a 2m buoy could potentially become a 2.5m or 5m buoy. On the other hand, FRM randomly changes the buoy radius to a different value. Since this provides FRM with four options for changing buoy radius while UDM only has two, FRM can be considered a more exploratory mutation operator.

As the name suggests, random local search is an algorithm that only evaluates solutions in its local neighbourhood and hence there is no guarantee of finding the optimal solution. However, the non-exhaustive nature of the search also means that it is computationally feasible to find solutions for larger arrays, even if those solutions are not necessarily optimal.

4.3 Results and Discussion

The results of the radii experiments are presented in the following sections, grouped by array size, starting with trivial cases and continuing towards the 5x5 array consisting of 25 CETO buoys. All experiments have been run in a shared cluster computing environment based on a Lenovo NeXtScale System. Due to the nature of computing environments with shared resources, the computation time can fluctuate between runs, so any computation time values in the experimental results are only indicative. One of the most computationally expensive parts of the experiment is the fitness evaluation of different array configurations using the CETO model. Therefore, we have developed a caching mechanism to re-use some of the calculations and speed-up the evaluation process. This provides an approximately 350-fold performance improvement [45] which is very beneficial to our experiments, particularly those involving larger arrays where the execution time of the CETO model is substantially higher.

4.3.1 Simple Arrays

Initially, we consider a number of trivial array configurations consisting of 1–2 CETO buoys, including the 1x1, 1x2 and 2x1 arrays. Since these are all very simple cases, a brute force search has been used to evaluate all possible radii configurations. Less than a minute of computation time is required to run a brute force search of the 1x1 array, while the 1x2 and 2x1 can be computed exhaustively in approximately 5 minutes each.

The optimal radii configurations for each array size are listed in Table 4.1 in terms of the four optimisation criteria: q-factor, RCW, average power absorption and tether force. All criteria are weighted based on the expected probabilities of different sea states. The values shown in bold are the optimal criteria values found for the corresponding array size. For example, a single 5m buoy is an optimal 1x1 configuration in terms of q-factor, RCW and power absorption, but not tether force since we are attempting to minimise this criterion.

Array Size	Radii [m]	q-factor	RCW	Power [W]	Tether Force
1x1	$\begin{pmatrix} 5 \end{pmatrix}$	**1**	**1.283**	**412,816**	1,022,498
	$\begin{pmatrix} 2 \end{pmatrix}$	**1**	0.465	54,303	**64,437**
1x2	$\begin{pmatrix} 5 & 5 \end{pmatrix}$	**1.010**	**1.309**	**839,413**	1,024,996
	$\begin{pmatrix} 2 & 2 \end{pmatrix}$	1.001	0.465	108,682	**64,838**
2x1	$\begin{pmatrix} 5 & 5 \end{pmatrix}$	0.919	**1.146**	**748,799**	960,069
	$\begin{pmatrix} 2 & 2 \end{pmatrix}$	**0.989**	0.454	106,473	**64,046**

Table 4.1: The optimal radii configurations of small arrays for different optimisation criteria. The values shown in bold are the optimal criteria values found for the corresponding array size.

The optimal 1x2 and 2x1 configurations are also visualised in Figure 4.2 and Figure 4.3 respectively. The colours represent the average power absorption of each buoy. There are some slight variations in the positioning of buoys of different radii, and this is due to the way in which the staggered array layout is constructed by the CETO model. In order to maintain the same tether angle and water depth for buoys of different sizes, the positions of those buoys need to be slightly adjusted in order to accommodate their difference in size.

The results of the 1x1 radii configurations show that the 5m buoy generates significantly more power than the 2m buoy. At the same time, they have a larger RCW value, meaning that their power output more than compensates for their larger size. However, the 2m buoy also generates a much lower tether force, suggesting that it could also be much cheaper to operate.

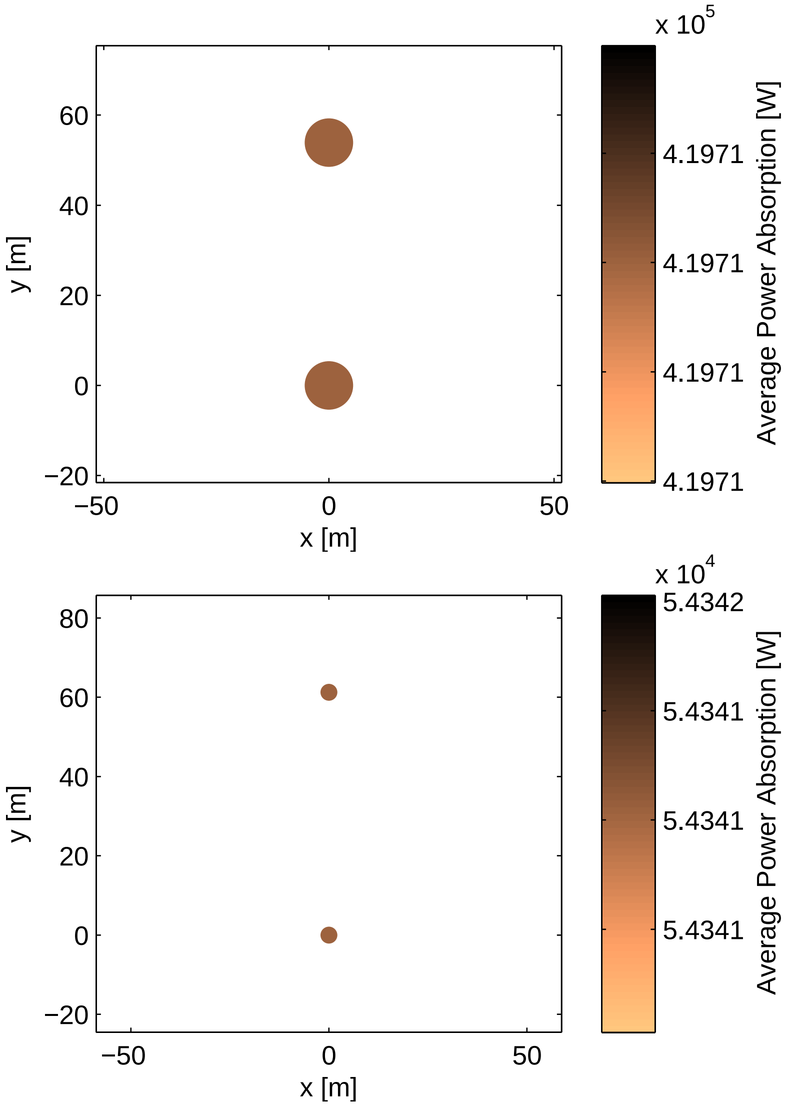

Figure 4.2: The optimal radii configurations of the 1x2 array. The top configuration is optimal in terms of q-factor, RCW and average power absorption. The bottom configuration is optimal in terms of tether force.

The 1x2 and 2x1 results show similar trends for RCW, power absorption and tether force. In order to maximise RCW and power absorption, larger buoys are required, while tether force is naturally minimise with smaller buoys. However, optimising q-factor requires different configurations for the two arrays, with only the 1x2 array favouring larger buoys for maximising q-factor. This may be due to the different orientation of the arrays in relation to the incident wave, which proceeds from right to left in the figures. In the 1x2 array, both buoys are facing the incident wave, while in the 2x1 array, the left buoy is slightly shadowed by the right buoy. This shadowing reduces the power output of the left buoy and may also

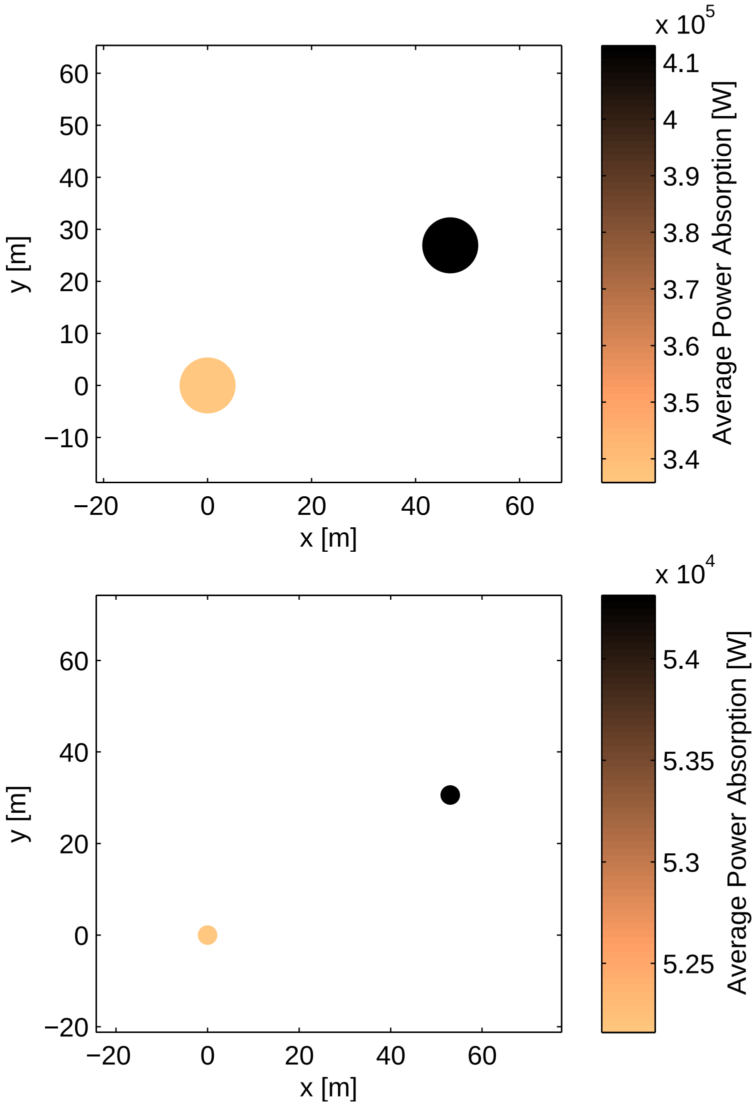

Figure 4.3: The optimal radii configurations of the 2x1 array. The top configuration is optimal in terms of RCW and average power absorption. The bottom configuration is optimal in terms of q-factor and tether force. The direction of wave propagation is from right to left.

cause some destructive interference between the buoys. By choosing buoys with the smallest radius, it is possible to minimise the magnitude of these effects and therefore maximise the q-factor of the array.

4.3.2 2x2 Array

The number of possible radii combinations for the 2x2 array is substantially larger than the smaller arrays, but is still computationally feasible to evaluate using brute force search. The 2x2 array consists of 4 buoys and each buoy can have one of 5 possible radii. This means that there are $5^4 = 625$ possible combinations to evaluate. The CETO model takes approximately 1 minute to evaluate the fitness of a single 2x2 array configuration, so the entire search space takes roughly 10 hours to compute.

The optimal configurations of the 2x2 array are listed in Table 4.2 and visualised in Figure 4.4, where each configuration is optimal for at least one of the four optimisation criteria. Notably, all of the optimal 2x2 configurations exclusively contain buoys with a radius of either 2m or 5m, regardless of the criteria being optimised. This is consistent with the optimal 1x1, 1x2 and 2x1 configurations, and may prove to be true for larger arrays. The ability to limit the search to 2m and 5m buoys would significantly reduce the search space for larger radii optimisation problems.

Radii [m]	q-factor	RCW	Power [W]	Tether Force
$\begin{pmatrix} 5 & 5 & 5 & 5 \end{pmatrix}$	0.897	**1.122**	**1,465,869**	934,463
$\begin{pmatrix} 5 & 2 & 2 & 5 \end{pmatrix}$	**0.986**	1.034	923,586	540,089
$\begin{pmatrix} 2 & 2 & 2 & 2 \end{pmatrix}$	0.984	0.449	211,183	**63,976**

Table 4.2: The optimal radii configurations of the 2x2 array for different optimisation criteria. The values shown in bold are the optimal criteria values found across all 2x2 configurations.

Overall, the $\begin{pmatrix} 5 & 2 & 2 & 5 \end{pmatrix}$ configuration provides a well-balanced solution, offering a q-factor near 1 with minimal destructive interference. At the same time, the RCW value is not significantly lower than the all 5m buoy solution, and the power absorption and tether force are competitive but balanced between the two extremes.

The other two optimal configurations use exactly the same radius for all buoys in the array. If it proves to be infeasible to evaluate all radii solutions for larger arrays, then using all 5m or all 2m buoy configurations may prove to be a valuable starting point, especially if the primary criterion of interest is anything other than q-factor. Furthermore, the left buoys in these arrays are shadowed by the right buoys, which has a clear impact on their power absorption (see bottom two configurations in Figure 4.4). This shadowing effect could potentially be reduced through layout optimisation which is further explored in Chapter 6.

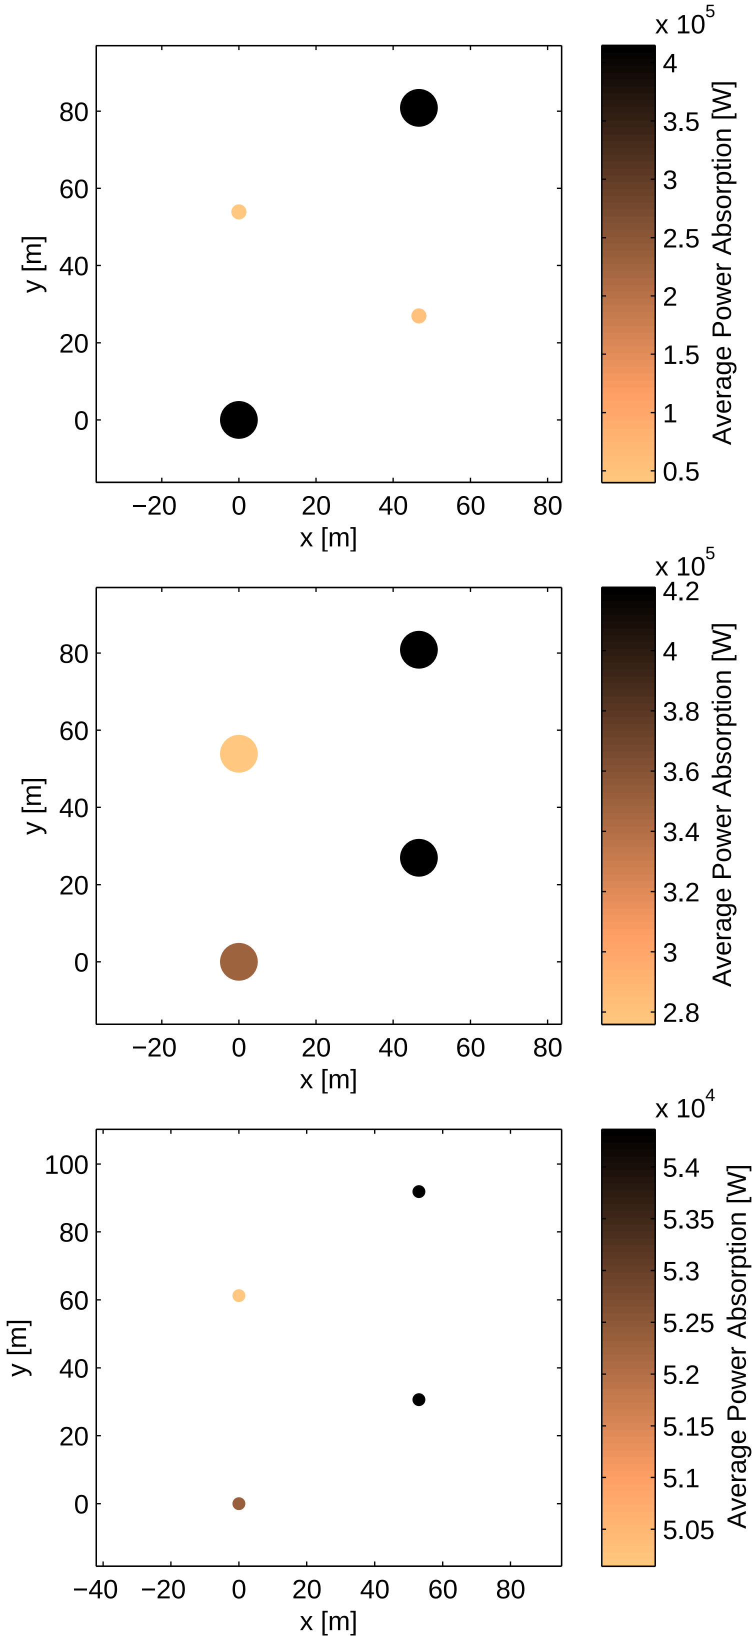

Figure 4.4: The optimal radii configurations of the 2x2 array. The top configuration is optimal for q-factor, the middle configuration is optimal for RCW and average power absorption, while the bottom configuration is optimal for tether force.

4.3.3 3x3 Array

With increasing array size, the CETO model takes longer to evaluate a single radii configuration and the number of possible configurations rapidly increases. In the 3x3 array, there are 9 buoys and therefore $5^9 = 1,953,125$ possible combinations. Each configuration takes around 2.5 minutes to evaluate, so it would take over 9 years to evaluate all combinations, and therefore a complete brute force search is not feasible for the 3x3 array.

Nonetheless, using insights from the 2x2 array, we have successfully run a partial brute force search. For the 2x2 array, we observed that the optimal configurations only consisted of buoys with a radius of either 2m or 5m. Since there are only $2^9 = 512$ combinations of 2m and 5m buoys for the 3x3 array, it is computationally feasible to evaluate all such combinations. With a computation time of 2.5 minutes per configuration, this partial brute force search takes around 21 hours to run.

Once again, we consider all four optimisation criteria and the best found configurations are listed in Table 4.3 and visualised in Figure 4.5, where each configuration is optimal for at least one criterion. The results again show that the all 5m buoy configuration is ideal for maximising power absorption, while the all 2m buoy configuration is best for minimising tether force. However, this all 2m buoy configuration is now also ideal for maximising q-factor, which suggests that the increase in the number of buoys from the 2x2 to the 3x3 array has also caused an increase in destructive interference between the buoys. If this is true, then the all 2m buoy configuration should also exhibit high q-factor for larger array sizes.

Radii [m]	q-factor	RCW	Power [W]	Tether Force
$\begin{pmatrix}5 & 5 & 5 & 5 & 5 & 5 & 5 & 5 & 5\end{pmatrix}$	0.777	1.065	**3,167,337**	808,966
$\begin{pmatrix}5 & 5 & 5 & 2 & 5 & 5 & 5 & 2 & 2\end{pmatrix}$	0.881	**1.103**	2,572,726	617,030
$\begin{pmatrix}2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2\end{pmatrix}$	**0.973**	0.402	433,952	**68,752**

Table 4.3: The optimal radii configurations of the 3x3 array for different optimisation criteria. The values shown in bold are the optimal criteria values found across all 3x3 configurations consisting of 2m and 5m buoys.

Another interesting finding is that a suitable mix of 2m and 5m buoys can result in an RCW value which exceeds the all 5m buoy configuration. One possible expla-

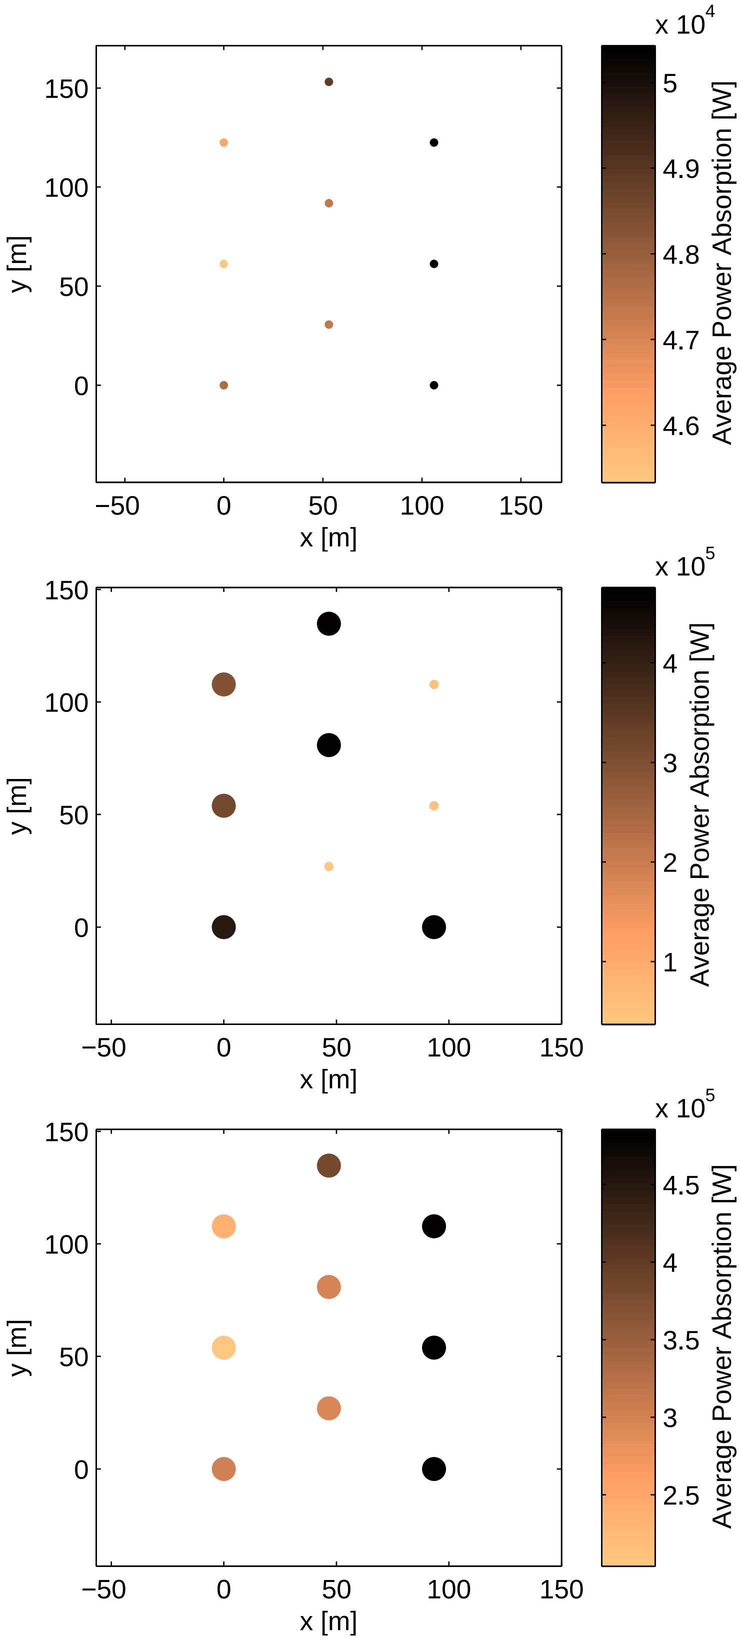

Figure 4.5: The optimal radii configurations of the 3x3 array. The top configuration is optimal for q-factor and tether force, the middle configuration is optimal for RCW, while the bottom configuration is optimal for average power absorption.

nation is that the leftmost buoys in the all 5m array have a substantial reduction in power absorption due to potential shadowing effects from the rightmost buoys, which absorb most of the wave energy propagating from right to left (Figure 4.5). This would cause a reduction in the relative capture width (RCW) of the buoys compared to the mixed 2m / 5m solution which is not as strongly affected by the shadowing effects due to the placement of 2m buoys in suitable positions.

4.3.4 4x4 Array

The radii configurations of the 4x4 array are too numerous to evaluate using either brute force search or the partial brute force search used for the 3x3 array. Even if we only consider buoys of 2m and 5m radius, there are $2^{16} = 65,536$ possible combinations. The CETO model takes approximately 7 minutes to evaluate a 4x4 configuration, which means that almost a year would be required to run a partial brute force search. This has led us to consider non-exhaustive methods, such as the random local search (RLS) described in Section 4.2.2.

Radii experiments of smaller arrays have generally shown that all 5m buoy and all 2m buoy configurations can be optimal for some criteria. Therefore, as a starting point, we conducted a local optima check for these two configurations to confirm whether it would even be possible for RLS to make any improvement with the up/down mutation (UDM). This was achieved by iterating over the buoys in the array configuration, and increasing or decreasing each buoy radius by a single increment from the starting configuration. The local optima check revealed that RLS with UDM would be unable to improve on the all 2m buoy configuration in terms of q-factor and tether force, and it would also be unable to improve on the all 5m buoy configuration in terms of power absorption. This suggests the 2m and 5m configurations of the 4x4 array are likely to be good solutions with respect to these optimisation criteria.

In spite of the local optima for these criteria, the local optima check was successful in finding an improvement to the all 5m buoy solution with respect to RCW. This improvement replaced one of the 5m buoys with a 2m buoy, which is reminiscent of the best RCW-based configuration found in the 3x3 radii experiment. With this knowledge, we have used RLS to perform RCW-based radii optimisation of the

4x4 array, starting with an all 5m buoy solution with an RCW value of 0.865. The algorithm was run for 24 hours and the best result was recorded at the end of that period. This process was repeated 30 times. The RCW of the best solution found in each repetition has been visualised using a box plot in Figure 4.6. The radii configuration of the best solution found across all repetitions is shown in Figure 4.7.

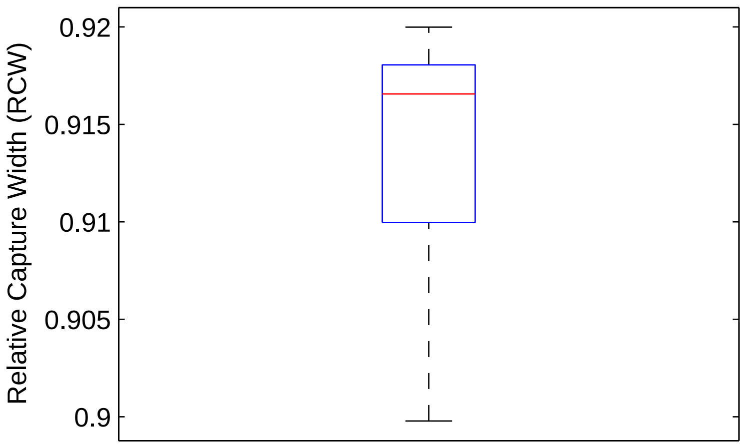

Figure 4.6: The RCW values of the best solutions found in each repetition of RLS for the 4x4 array.

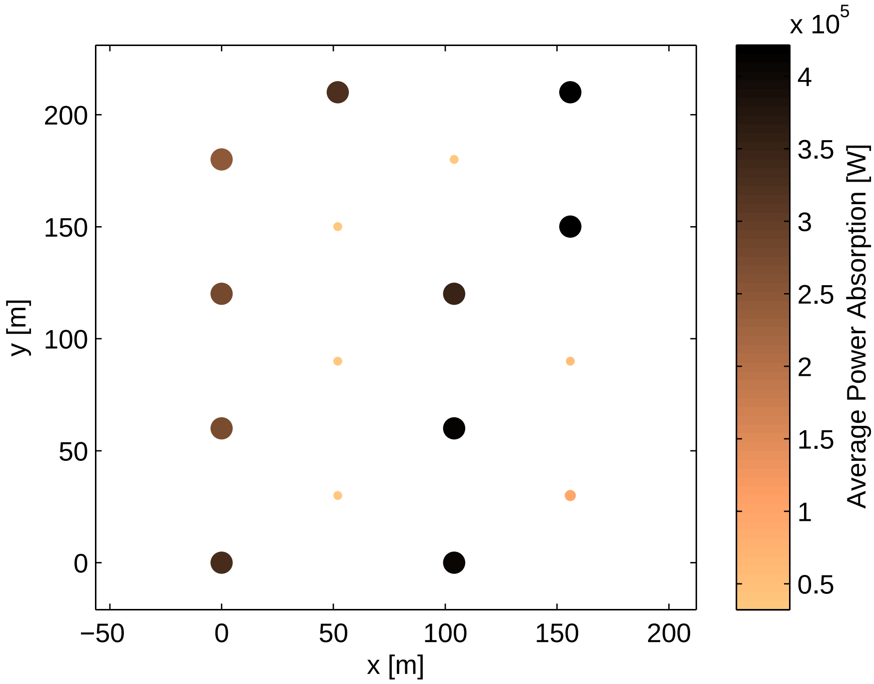

Figure 4.7: The radii configuration of the best RCW-based solution found by RLS for the 4x4 array. RCW = 0.920; q-factor = 0.853; Power = 3,751,603; Tether Force = 590,251.

The results show that random local search can consistently find 4x4 configurations with an RCW value equal to or exceeding 0.9, which is a notable improvement on the RCW value of the all 5m buoy starting configuration (0.865). This demonstrates that RLS is a useful for technique for finding improvements to radii configurations in larger arrays, where brute force search is no longer feasible. Furthermore, the best solution found by RLS for the 4x4 array is reminiscent of the best RCW-based solution for the 3x3 array. There appears to be a pattern of alternating columns of larger and smaller buoys, which may help to reduce the shadowing effects on the larger buoys, and would ultimately lead to higher power absorption and an increase in relative capture width.

4.3.5 5x5 Array

Given that radii optimisation of the 4x4 array was not feasible to complete exhaustively, the 5x5 optimisation problem is even larger and is therefore much better suited towards approximate methods like random local search. Consequently, we conducted the same local optima checks on the all 5m buoy and all 2m buoy configurations of the 5x5 array and found very similar results to the 4x4. The 2m buoy configuration is locally optimal in terms of q-factor and tether force, while the 5m buoy configuration is locally optimal in terms of power absorption but can be improved upon in relation to RCW. As a result, we have applied RLS to the 5x5 radii optimisation problem in a very similar way. Starting with an all 5m buoy solution (RCW = 0.771), RLS was run for 24 hours with a focus on optimising RCW, and the best result was recorded at the end of that period. This process was repeated 30 times. The RCW of the best solution found in each repetition has been visualised using a box plot in Figure 4.8. The radii configuration of the best solution found across all repetitions is shown in Figure 4.9.

Based on these results, random local search appears to be effective for consistently finding solutions that exceed the all 5m buoy configuration in terms of RCW (0.771). The best solution also shows visual similarities to the best 3x3 and 4x4 RCW-based configurations. The differences may be partly attributable to the increased computation time required for the CETO model to evaluate 5x5 configuration which ultimately reduces the number of 5x5 configurations that can be

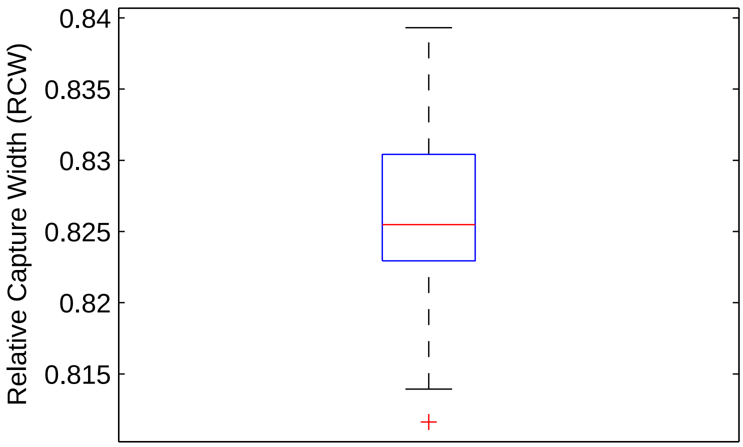

Figure 4.8: The RCW values of the best solutions found in each repetition of RLS for the 5x5 array.

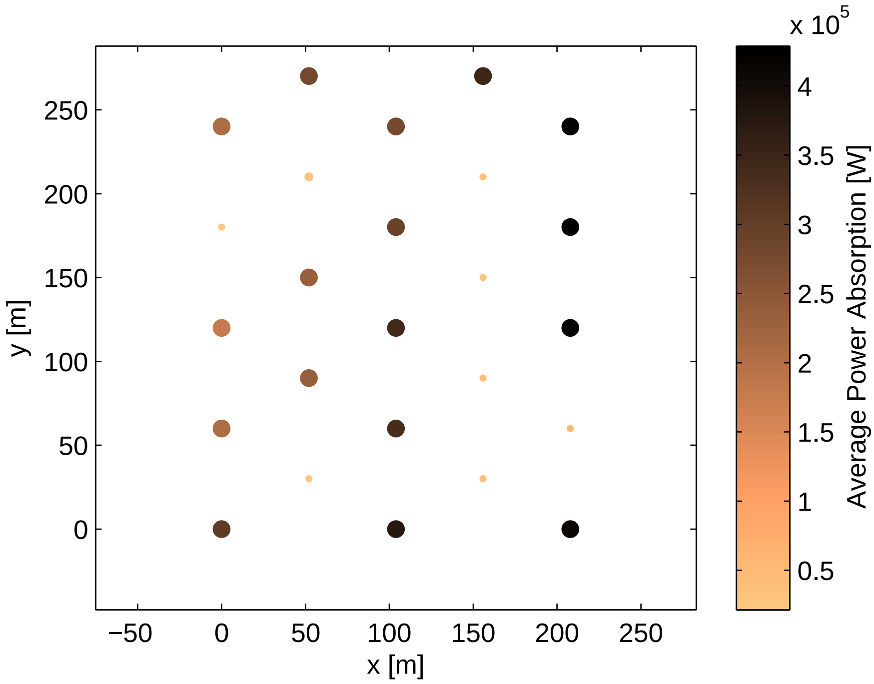

Figure 4.9: The radii configuration of the best RCW-based solution found by RLS for the 5x5 array. RCW = 0.839; q-factor = 0776; Power = 5,622,107; Tether Force = 592,966.

considered in a 24-hour period. Given more computation time, it is expected that the best 5x5 configuration will show similar trends to the smaller arrays.

Finally, throughout the radii experiments we have observed a general decrease in array q-factor as the size of the arrays increases. For example, the all 5m 2x2 configuration has a q-factor of 0.897, while the all 5m 5x5 configuration has a q-factor of only 0.66. This supports an earlier observation that an increase in the number of buoys seems to increase the destructive interference in the CETO array. This is consistent with the best q-factor solutions found for larger arrays, which tend to favour the 2m buoys. These smaller buoys absorb less power and exhibit smaller forces, but as a consequence they also diminish the effect of any destructive interference in the array.

Chapter 5

Spacing Optimisation

For our next set of experiments, we consider the variation of spacing between the buoys in a CETO array. We investigate configurations where buoys are almost touching to those where the buoys are four times further apart than in the default staggered configuration (see Figure 4.1). The following subsections define the spacing optimisation problem, describe an effective solution technique, and then present our experimental results.

5.1 Problem Definition

A solution to the spacing optimisation problem is represented as a real scalar value m_s, which reflects the spacing multiplier that is applied to the buoys in a 2x2 array. For example, if m_s is set to 2 then buoys are spaced twice as far apart as in the default configuration. In addition, we experiment with both the default staggered layout (Figure 4.1), as well as an aligned grid as shown in Figure 5.1.

The aim is to identify the most promising spacing arrangements and explore the trends in optimisation criteria as the spacing between the buoys increases. In alignment with the radii optimisation experiments, we consider the same four optimisation criteria: q-factor, average power absorption, RCW and tether force. However, the spacing optimisation is intentionally much simpler than the radii optimisation, as we are interested in understanding the effect of the spacing parameter in isolation. This will help to inform the more advanced layout optimisation experiments that will be covered in Chapter 6.

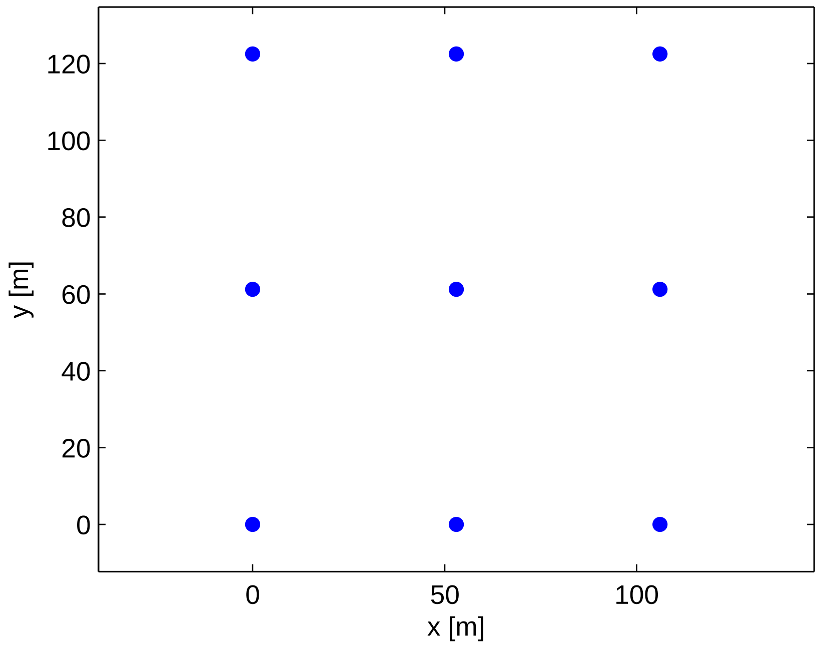

Figure 5.1: The aligned grid layout of a 3x3 array in the CETO model.

5.2 Solution Technique: Iterative Search

The brute force search that was used for some of the radii optimisation experiments in Chapter 4 is not possible for the spacing experiments, as the spacing multiplier is represented as a real number. Therefore, we use a form of iterative search to increase the spacing multiplier by a fixed quantity (spacing increment) and thereby evaluate increasingly larger array configurations. By carefully selecting this spacing increment, we can choose the granularity at which to explore the search space given the computational resources available and the sensitivity of buoy spacing on the performance of the array. For example, a smaller spacing increment will explore the search space more thoroughly at the expense of increased computation time. For our experiments, we use a value of 0.01 for the spacing increment, which translates to an average spacing increase of about 0.5m between the buoys, and provides a suitable level of granularity for our purposes.

5.3 Results and Discussion

A total of four spacing experiments have been run using different configurations of a 2x2 buoy array:

1. Staggered layout consisting of 5m buoys

2. Aligned layout consisting of 5m buoys

3. Staggered layout consisting of 2m buoys

4. Aligned layout consisting of 2m buoys

The experiments have been run in the same cluster computing environment described in Section 4.3. The results are shown in Figures 5.2 and 5.3 where the first figure shows the effect of buoy spacing on q-factor and tether force, while the second figure shows the effect of buoy spacing on RCW and average power absorption. Within each plot, the four lines show the results of the four experiments listed above with respect to the relevant optimisation criterion.

The results provide several insights into the effects of spacing on the dynamics of the buoy array. Firstly, arrays consisting of 2m buoys are not very sensitive to spacing between buoys, regardless of whether the array is staggered or aligned. The only notable exception is when the buoys are very close together some criteria, such as q-factor, can be negatively influenced. On the other hand, spacing can have a significant impact on many aspects of 5m buoy arrays, and their configuration (aligned or staggered) can also affect the power absorption and other criteria.

There is a clear trend in the spacing of the 5m staggered array, where all four optimisation criteria consistently increase with further spacing of the buoys. There is a similar trend with the 5m aligned array, although the criteria values seem to plateau once an average spacing around 100m is reached. This means that with sufficient spacing between buoys, the 5m staggered array generally outperforms the 2m staggered array, at least for the spacing values that we have investigated. The only exception is tether force, which we are aiming to minimise but which also happens to be higher in the 5m staggered configuration.

Comparing the 2m and 5m buoy arrays, the 5m variants are clearly superior in terms of power absorption and RCW, and this is consistent with the findings of many

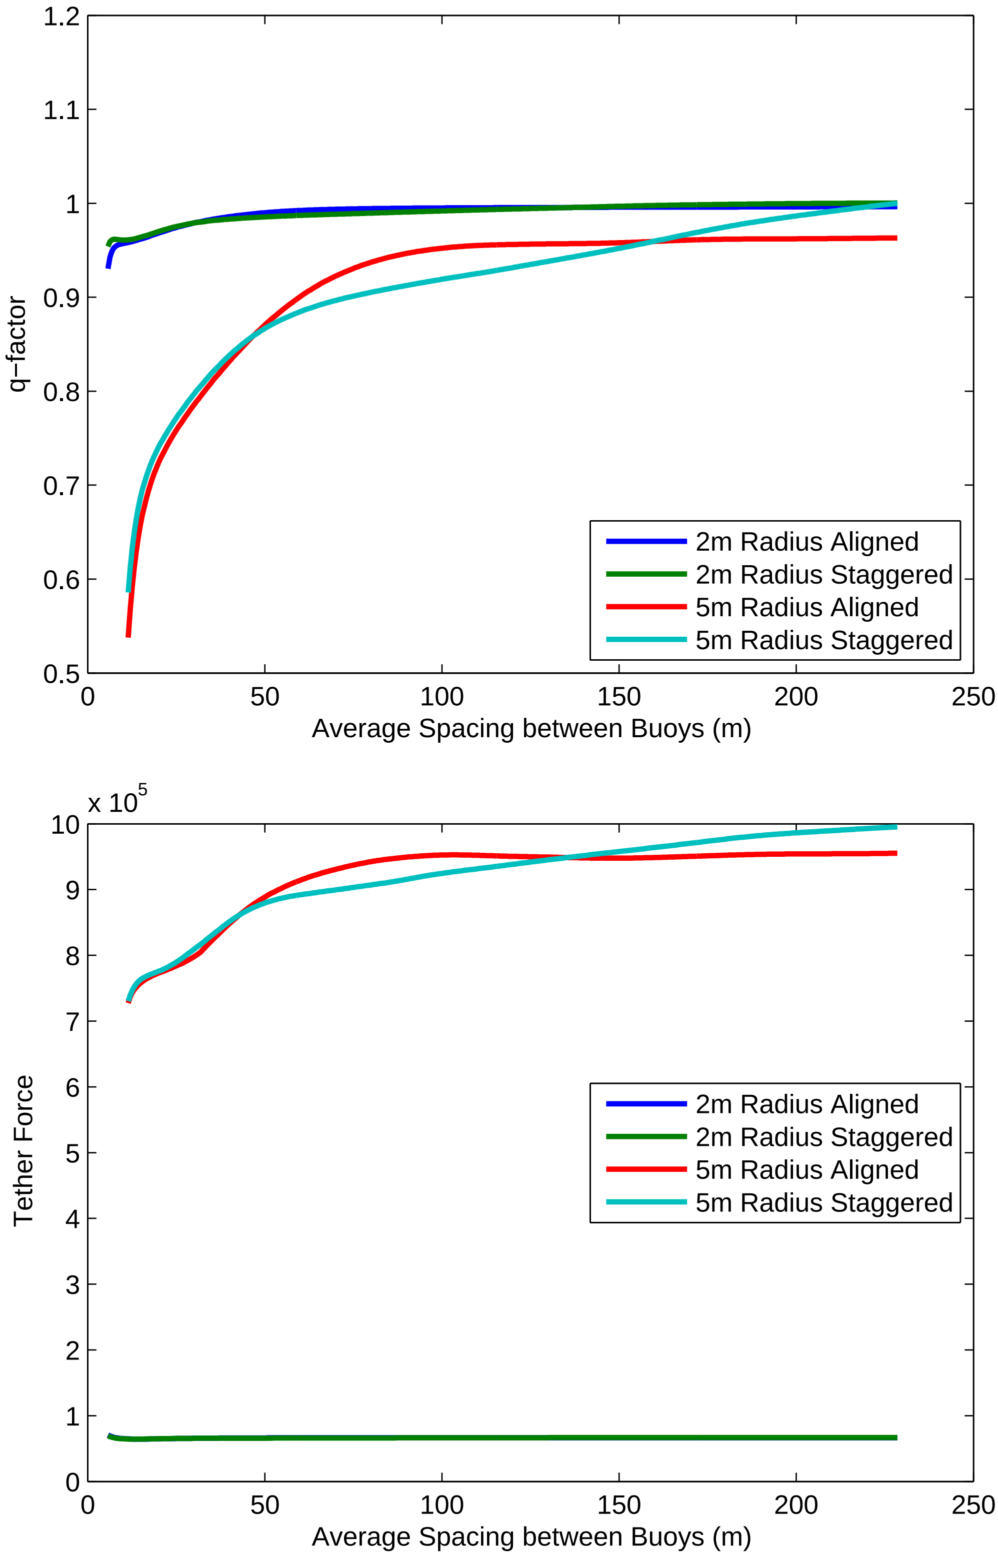

Figure 5.2: The effect of buoy spacing on the q-factor and tether force of a 2x2 array.

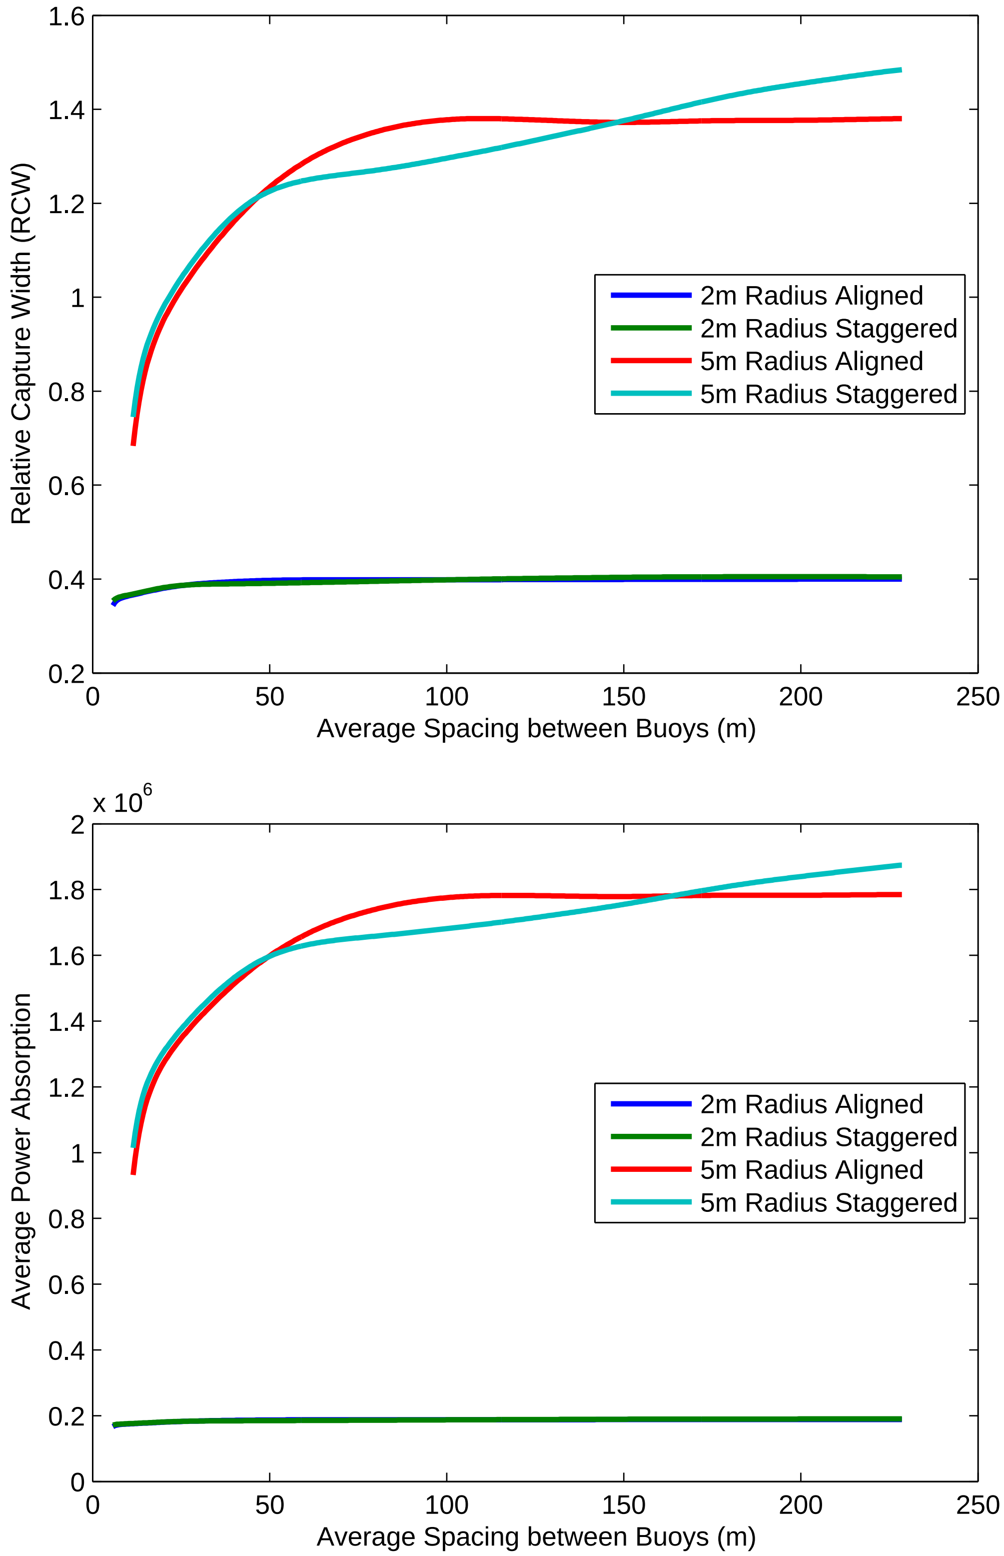

Figure 5.3: The effect of buoy spacing on the RCW and average power absorption of a 2x2 array.

of the radii optimisation experiments. In terms of q-factor, the 2m buoys result in a higher value in closely spaced arrays, but this advantaged is quickly diminished with increased spacing. In particular, the 5m staggered array has a comparable q-factor once spacing reaches 200m or more.

From this spacing study, the main disadvantage of using 5m buoys appears to be the increased tether force, which would theoretically increase the operational cost of the array. If this is not a major concern, then using 5m buoys would be a better option for wave energy capture, particularly in terms of the raw power absorption and relative capture width of the buoys. In terms of staggered and aligned configurations, the two are quite comparable, although the 5m staggered variant seems to be superior once the buoys are spaced sufficiently far apart.

Chapter 6

Layout Optimisation

Using the insights gained from the radii and spacing experiments, we now consider the problem of finding effective layouts for differently sized CETO arrays. More specifically, we investigate the best combination of buoy positions in an array given proximity and boundary constraints. The following subsections define the layout optimisation problem, describe an effective solution technique, and then present our experimental results.

6.1 Problem Definition

A solution to the layout optimisation problem is represented as a list of coordinates, $\begin{pmatrix} x_1 & \dots & x_n & y_1 & \dots & y_n \end{pmatrix}$, where x_k and y_k are the x and y coordinates of buoy k and n is the total number of buoys in the array. For example, the layout visualised in Figure 6.1 is represented as $\begin{pmatrix} 0 & 30 & 50 & 80 & 0 & 50 & 30 & 80 \end{pmatrix}$.

Layouts are also subject to proximity and boundary constraints, which restrict the number of valid layouts that are possible. The proximity constraint ensures that no two buoys are placed within 50m of one another to allow sufficient space between the buoys for the movement of certain ships and watercraft. The boundary constraint ensures that all buoys are placed within a fixed area to ensure that search algorithms do not seek impractical solutions. The placement area is calculated based on array size, as listed in Table 6.1, allowing approximately $20,000m^2$ per buoy within a square area. The placement areas are intentionally large, because small areas have previously been explored through the staggered layout in the radii and

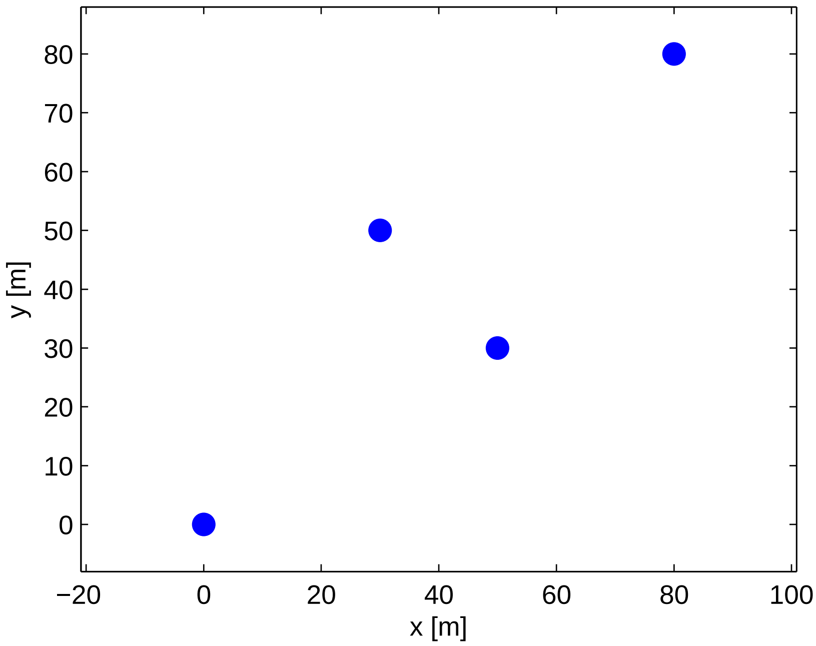

Figure 6.1: An example layout of a 2x2 array.

spacing experiments. In addition, the spacing experiments showed that spacing the buoys further apart resulted in higher RCW, power and q-factor (see results in Section 5.3), so using a large boundary constraint will provide the necessary flexibility to explore these broader spacing arrangements.

Array Size	Number of Buoys	Placement Area
2x2	4	283m x 283m
3x3	9	424m x 424m
4x4	16	566m x 566m
5x5	25	707m x 707m

Table 6.1: The placement areas of different arrays based on $20,000m^2$ per buoy.

The goal of layout optimisation is to identify the best layouts for arrays of different sizes, while ensuring that neither the proximity constraint nor the boundary constraint is violated. For this optimisation problem, we focus on using relative capture width (RCW) and q-factor to measure the quality of a given layout. These criteria have provided particularly useful insights in earlier experiments, and in some radii experiments have acted as conflicting objectives, offering multiple viewpoints on the same optimisation problem.

In terms of searching for solutions, there are several options for initialisation. One option is to start with a fixed layout, such as the staggered or aligned layout used in the array spacing experiments, while another option is to start with a completely random layout. The fixed layout can improve the consistency of results over multiple runs, while the random layout can help some optimisation algorithms to escape from local optima by initialising the algorithm in different parts of the search space.

6.2 Solution Technique: CMA-ES

Random local search was effectively used in larger radii optimisation problems to find local improvements to simple solutions consisting of same sized buoys (see Section 4.3.4). However, due to the more complex nature of the layout optimisation problem and the fact that solutions to this problem are represented in the real number domain, we now explore a more advanced algorithm called Covariance Matrix Adaptation Evolution Strategy (CMA-ES). An evolution strategy is an optimisation algorithm that explores the search space by sampling new solutions from a multivariate normal distribution, where each problem variable corresponds to a dimension of this distribution. CMA-ES is a particular type of evolution strategy that uses covariance matrix adaptation to modify the covariance matrix of this distribution during execution [21]. This allows CMA-ES to significantly reduce the number of generations and computation time needed to converge to an optimum solution, making it a powerful search algorithm.

For our experiments, we use the CMA-ES implementation provided by the authors of the algorithm [19]. This implementation is based on minimising an objective function, while our goal is to maximise the optimisation criteria, so the objective function value is inverted before and after being passed to the CMA-ES algorithm. Due to the expensive computational costs of running the CETO model to evaluate solutions, we use a small (2,2)-CMA-ES variant, which maintains a population of 2 and uses 2 parents to generate new solutions each generation. Despite this small population, we provide the algorithm with the ability to explore the search space by repeating each CMA-ES run multiple times and where possible allowing sufficient time for convergence.

In terms of constraint handling, boundary or box constraints can be handled in CMA-ES in a number of ways. A simple approach is to repair the solution to avoid violating the constraint, or alternatively a penalty term can be added to the objective function based on the extent to which the constraint is violated [20]. For the layout optimisation problem, we specify the boundary where buoys can be placed using the 'LBounds' and 'UBounds' parameters of the CMA-ES algorithm. These parameters use a coordinate wise boundary handling technique that evaluates invalid solutions as though they were on the boundary of the feasible space, but applies a penalty term based on their distance from the feasible space to push the search away from the boundaries.

On the other hand, the proximity constraint is unique to our optimisation problem and is handled through solution resampling. If a generated solution contains any buoys that are within 50m of one another, a 'NaN' value is returned by the objective function, which signals CMA-ES to resample for a new solution. This process is repeated until a valid solution is found.

6.3 Results and Discussion

The results of the layout experiments are presented in the following two sections, focusing on different optimisation criteria. Section 6.3.1 presents the results when optimising for the RCW criterion, while Section 6.3.2 presents the results when optimising for q-factor. All experiments have been run in the same cluster computing environment described in Section 4.3.

6.3.1 Optimising for RCW

For these experiments, we use CMA-ES to optimise the layout of the 2x2, 3x3, 4x4 and 5x5 arrays with respect to RCW. Starting with a random initial layout, the CMA-ES algorithm is given a 7-day period to find the layout with the highest possible RCW. The exception is the 2x2 array, which can be evaluated much faster by the CETO model and is therefore only given a 24-hour period, which provides sufficient time for CMA-ES to converge to a local optima. The 3x3, 4x4 and 5x5 arrays were also initially optimised over 24 hours, but then subsequently extended to 7-day experiments which ultimately led to a 1-2% improvement in RCW.

The CMA-ES experiments are repeated 30 times for each array size. The RCW of the best solution found in each repetition has been visualised using a box plot in Figure 6.2. The best overall layouts for each array size are also listed in Table 6.2 and visualised in Figures 6.3 and 6.4.

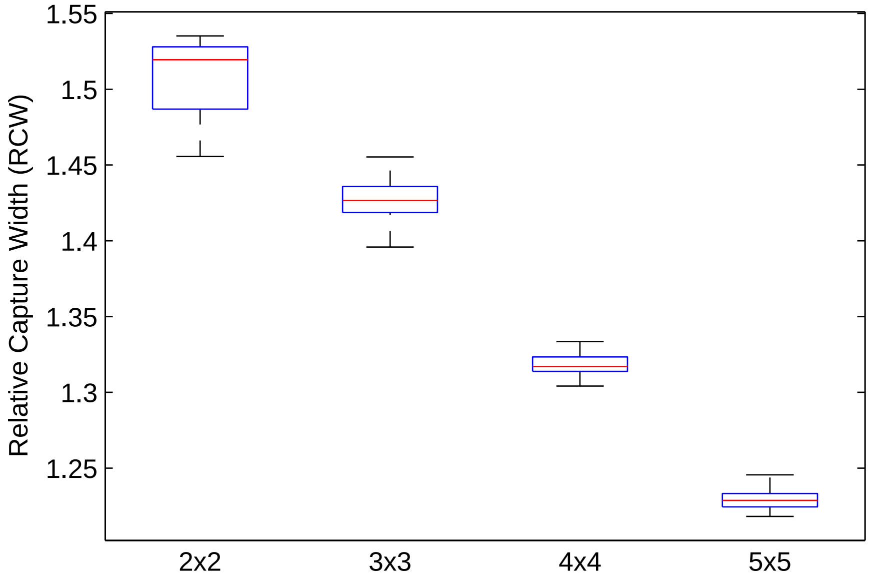

Figure 6.2: The RCW of the best solutions found in each repetition of CMA-ES layout optimisation. The results are grouped by array size.

Array Size	q-factor	RCW	Power [W]	Tether Force
2x2	1.022	1.535	1,938,912	1,004,544
3x3	0.982	1.455	4,157,851	973,753
4x4	0.923	1.333	6,853,572	928,819
5x5	0.876	1.246	10,075,302	894,265

Table 6.2: The criteria values of the best overall layout found for each array size when optimising for RCW.

The results in Figure 6.2 and Table 6.2 indicate that there is a general decrease in RCW, q-factor and tether force as array size increases. This supports our observations in the radii optimisation experiments, where increasing the number of buoys seems to increase the destructive interference in the CETO array, which has a negative impact on criteria such as q-factor and RCW. As expected, there is also some deviation in results across runs due to the stochastic nature of CMA-ES. The

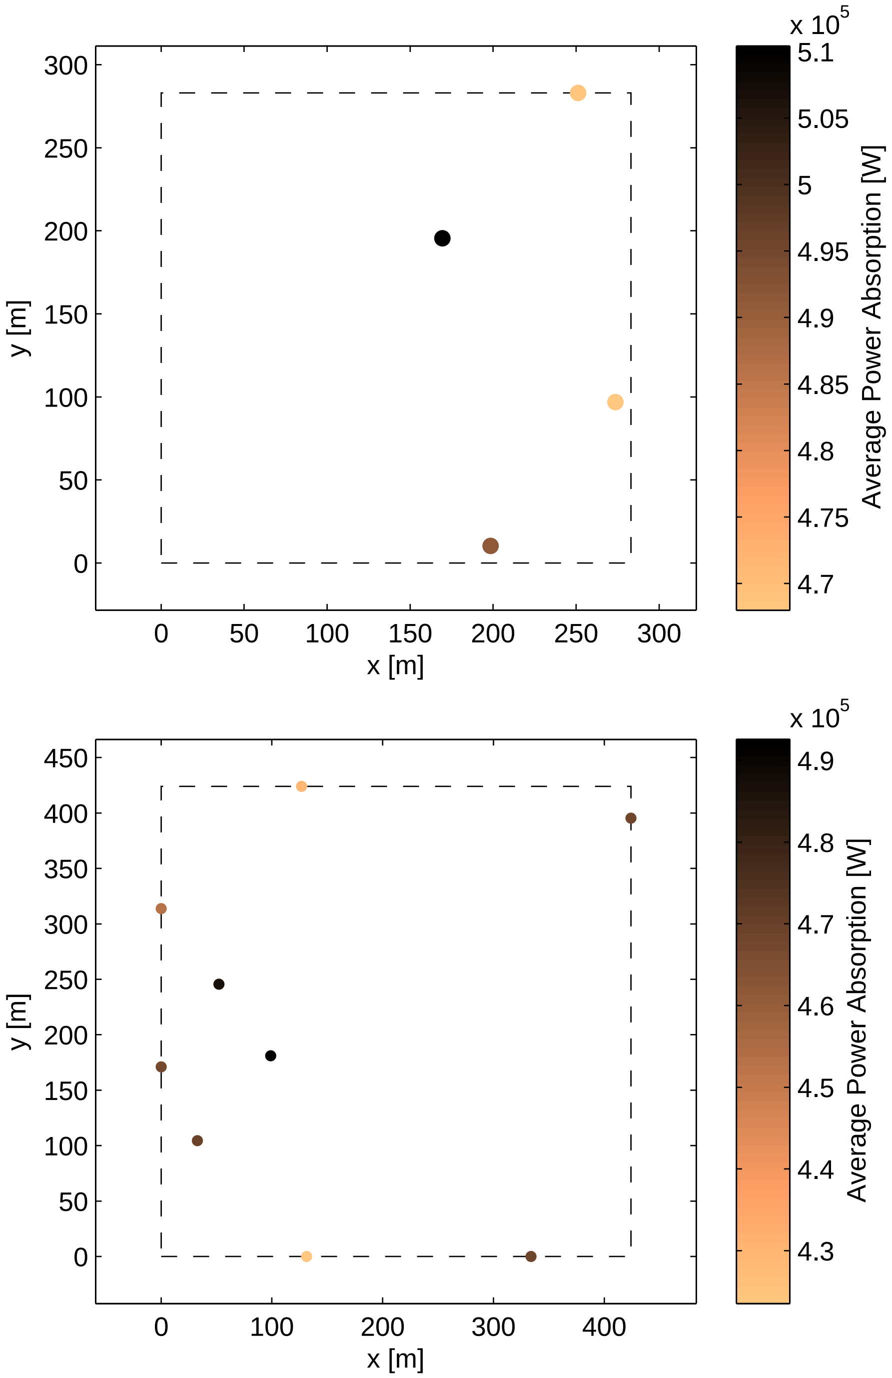

Figure 6.3: The best layouts found for the 2x2 and 3x3 arrays when optimising for RCW. The optimisation criteria values of these layouts are given in Table 6.2.

fact that this deviation is greater in the 2x2 and 3x3 layout experiments may be due to the nature of our experimental setup. Since the CETO model takes progressively longer to evaluate larger arrays, the 2x2 and 3x3 experiments naturally had time to complete a larger number of generations. The spread of these solutions may be due to the algorithm getting stuck in various local optima after many generations, while

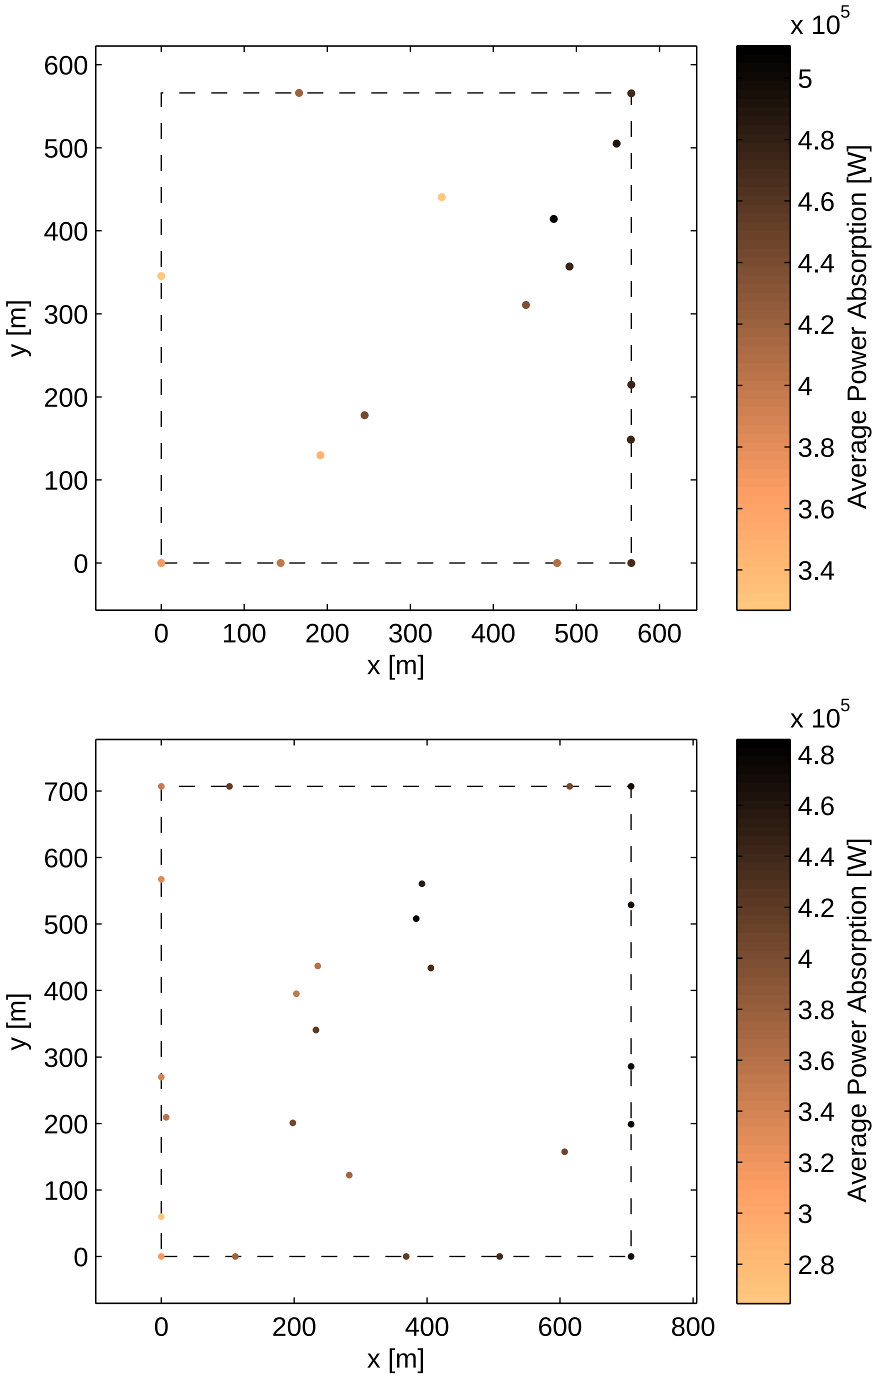

Figure 6.4: The best layouts found for the 4x4 and 5x5 arrays when optimising for RCW. The optimisation criteria values of these layouts are given in Table 6.2.

arrays larger were not given as much opportunity to explore the fitness landscape so the final solutions are much closer together in terms of RCW.

There are also several trends in terms of buoy placement, which may seek to take advantage of the wave direction that propagates from right to left in Figures 6.3 and 6.4. In the 2x2 case, the buoys seem to be grouped in pairs and arranged in such a way that the buoys on the left, which are further from the wave front, absorb more of the wave energy. Although this is an unusual arrangement, the RCW is above 1.5 and the q-factor is higher than 1 suggesting that this particular layout is taking advantage of constructive interference between the buoys. This is a very promising result, which was not possible to achieve by only optimising the radii of the buoys in earlier experiments, where the optimal 2x2 radii configurations were only able to reach an RCW of 1.122 and a q-factor of 0.986.

In the 3x3, 4x4 and 5x5 arrays, the trend is slightly different. In the case of the 3x3 array, there are still several buoy pairs on the left side of the placement area, but the remaining buoys are placed along the perimeter. Similar arrangements can be seen for the 4x4 and 5x5 arrays, although the trend is not as clear. Extrapolating from the 2x2 layout result, one explanation is that these buoy pairs are resulting in constructive interference in the array. Since the q-factor of these larger arrays is still less than 1, perhaps the non-paired buoys are still causing destructive interference. However, by placing them along the perimeter of the placement area and therefore increasing the separation distance between them, this destructive effect may be reduced.

Although the q-factor of the larger arrays is still less than 1, these layout-optimised arrays are able to achieve significantly higher RCW and q-factor than the best solutions found in the radii optimisation experiments. For example, the best solution found in the 5x5 radii experiments had an RCW value of 0.839, while the 5x5 layout experiments resulted in a solution with an RCW value of 1.246. Similar improvements can be seen in the results of the 3x3 and 4x4 layout optimisation. Nonetheless, since the overall interference within the array is still slightly negative, there may be scope in future work for further optimising the configurations of these larger arrays, possibly by finding specific arrangements of buoy pairs that result in constructive interference for the entire array, if these arrangements exist.

6.3.2 Optimising for q-factor

For these experiments, we use CMA-ES to optimise the layout of the 2x2, 3x3, 4x4 and 5x5 arrays, but this time with respect to q-factor. Starting with a random initial layout, the CMA-ES algorithm is given a 7-day period to find the layout with the highest possible q-factor. The exception is the 2x2 array, which can be evaluated much faster by the CETO model and is therefore only given a 24-hour period, which provides sufficient time for CMA-ES to converge to a local optima. The 3x3, 4x4 and 5x5 arrays were also initially optimised over 24 hours, but then subsequently extended to 7-day experiments which ultimately led to a 1-2% improvement in q-factor.

The CMA-ES experiments are repeated 30 times for each array size. The q-factor of the best solution found in each repetition has been visualised using a box plot in Figure 6.5. The best overall layouts for each array size are also listed in Table 6.3 and visualised in Figures 6.6 and 6.7.

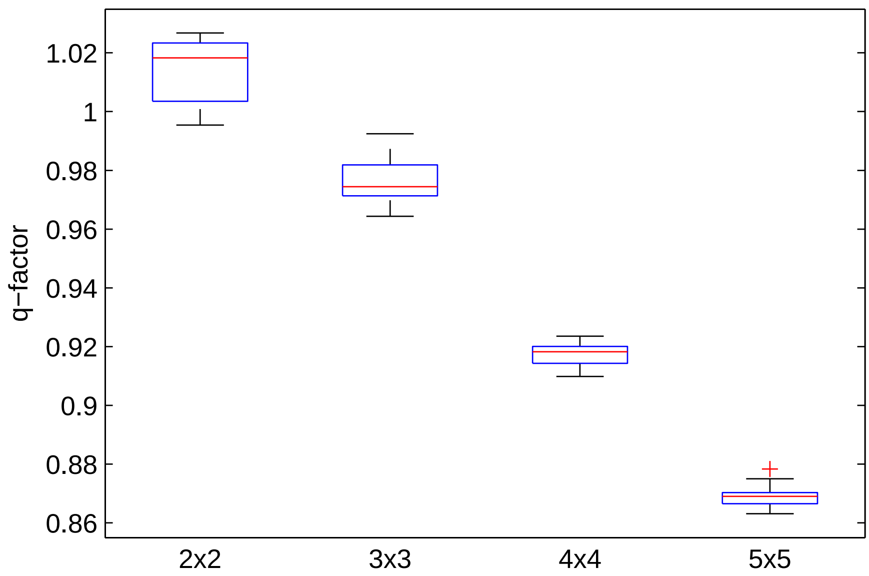

Figure 6.5: The q-factor of the best solutions found in each repetition of CMA-ES layout optimisation. The results are grouped by array size.

Array Size	q-factor	RCW	Power [W]	Tether Force
2x2	1.027	1.524	1,940,530	1,001,031
3x3	0.992	1.469	4,196,407	978,378
4x4	0.924	1.330	6,856,921	928,666
5x5	0.878	1.243	10,092,368	892,075

Table 6.3: The criteria values of the best overall layout found for each array size when optimising for q-factor.

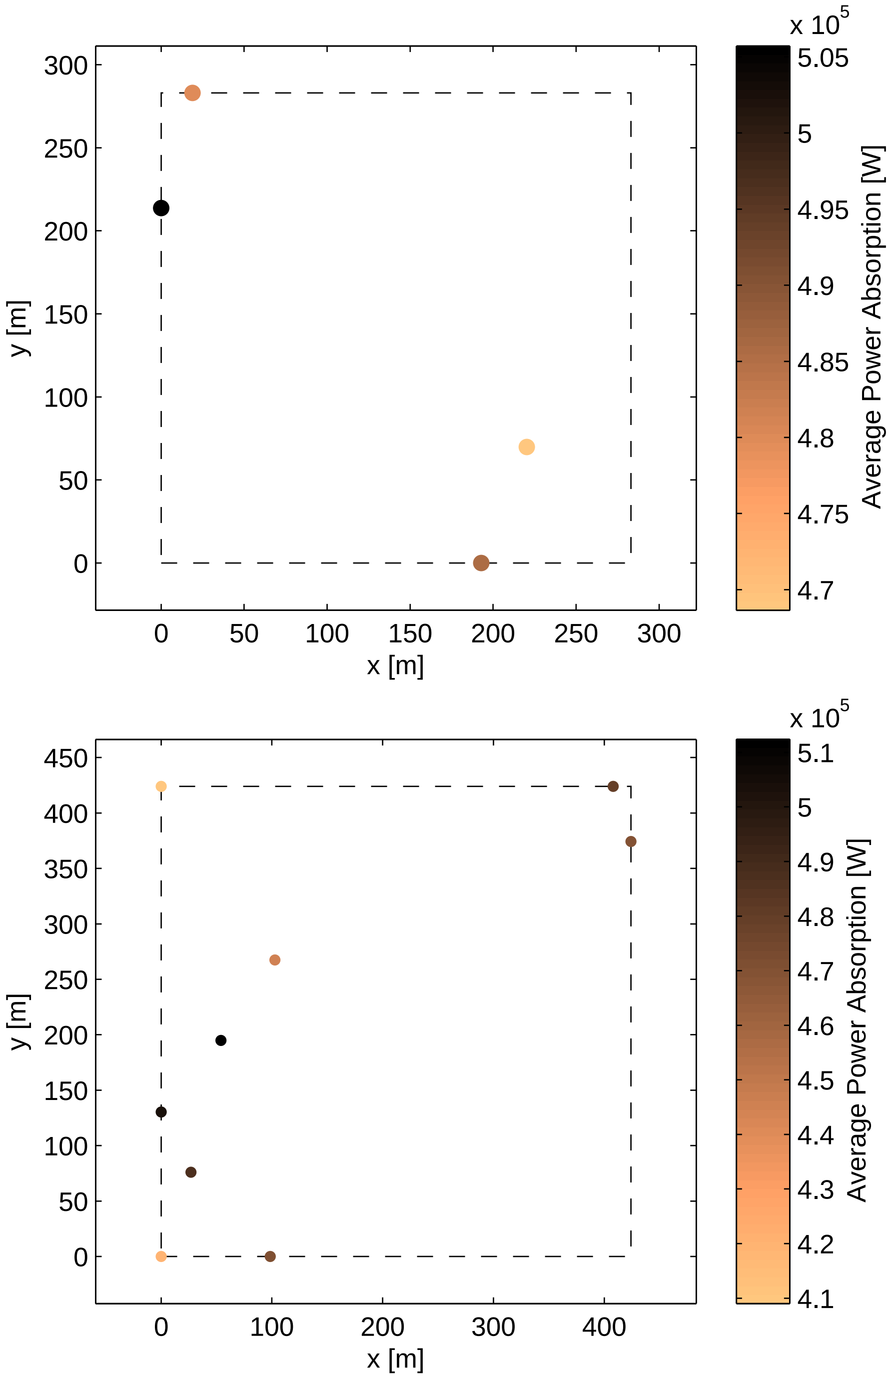

Figure 6.6: The best layouts found for the 2x2 and 3x3 arrays when optimising for q-factor. The optimisation criteria values of these layouts are given in Table 6.3.

Similar to the RCW-based layout optimisation, the results of the q-factor-based experiments show a general decrease in q-factor, RCW and tether force with increasing array size, presumably due to the presence of more destructive interference in larger arrays. For all array sizes, improved layouts in terms of q-factor have successfully been found with varying degrees of improvement. Interestingly, the RCW

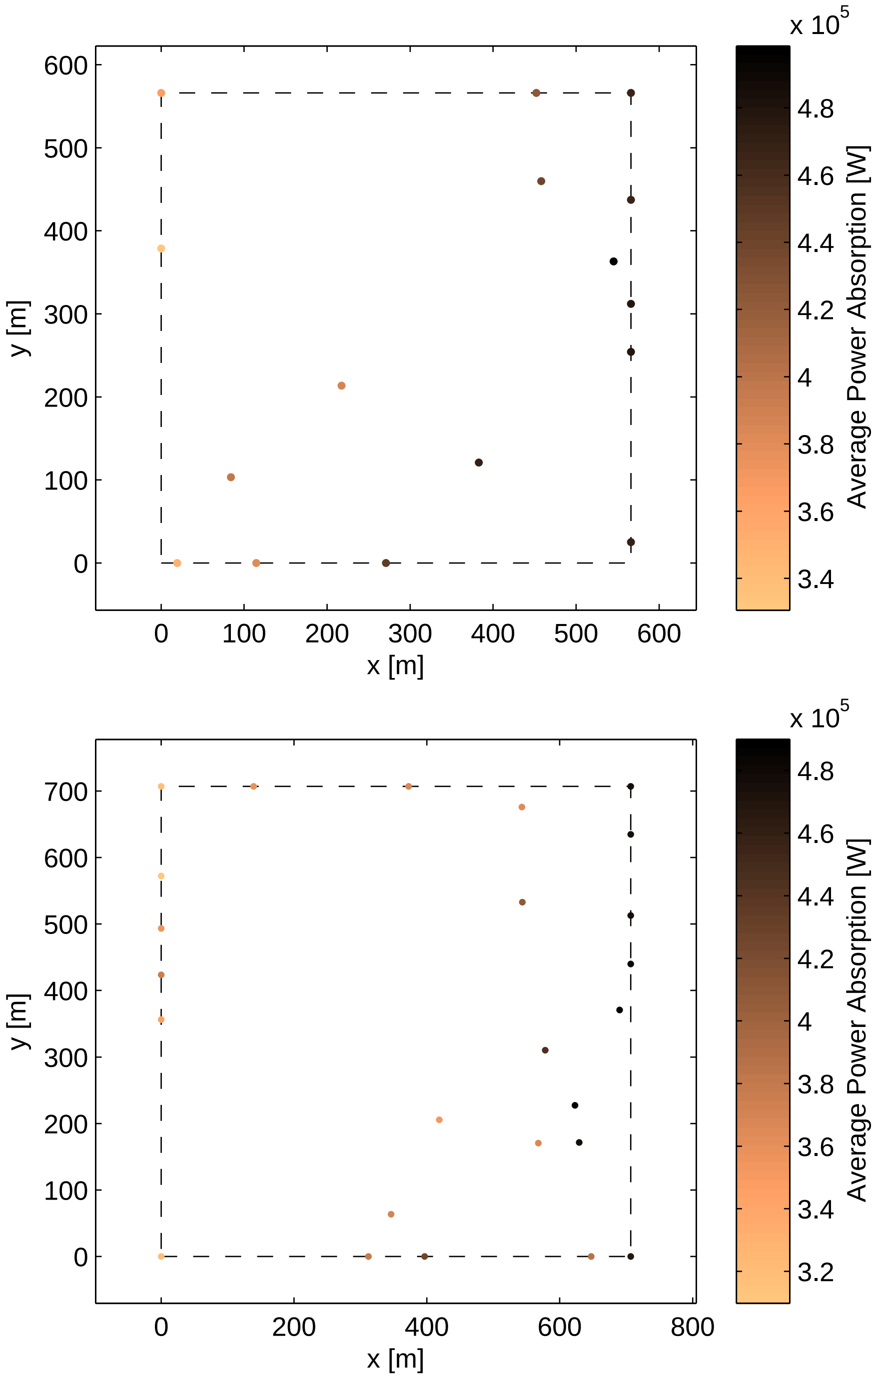

Figure 6.7: The best layouts found for the 4x4 and 5x5 arrays when optimising for q-factor. The optimisation criteria values of these layouts are given in Table 6.3.

of the best layout found for the 3x3 array is actually higher than the layout found in the earlier experiments which focused on optimising RCW. Yet, the difference between the RCW and q-factor values from the two sets of experiments are quite small, which suggests that the ideal layouts for maximising these two criteria may ultimately be very similar.

In terms of buoy placement, we can see similar trends to the RCW-based optimisation. In the 2x2 array, the buoys are once again grouped in pairs, with the left buoy of each pair harnessing more of the wave energy presumably through strong positive interference. The buoy pairs are spaced even further apart than in the best RCW-based layout, suggesting that appropriate array spacing helps to further increase q-factor, which is consistent with the results of our spacing optimisation experiments (see Section 5.3).

In the 3x3, 4x4 and 5x5 arrays, the best layouts again consist of placing many buoys along the perimeter of the placement area, while some are placed in the middle or right of the placement area in various arrangements including buoy pairs. Since the results of the q-factor and RCW-based layout experiments are very comparable in terms of both q-factor and RCW, the layouts found in either experiment are likely to be effective arrangements for maximising energy capture. The general trend across all layout experiments seems to involve specific buoy arrangements, such as buoy pairs, that maximise constructive interference within the array. The remaining buoys are then arranged along the perimeter of the placement area to increase buoy separation distance and therefore decrease any destructive interference that may be created by these additional buoys.

Chapter 7

Conclusion

In this Master's thesis, we have explored different methods for optimising arrays of fully submerged CETO buoys to maximise their production of renewable energy. Using a MATLAB model of a CETO array, we investigated three specific optimisation problems: finding optimal combinations of buoy radii, exploring the effect of buoy spacing on array performance, and identifying the highest performing buoy layouts.

The effectiveness of different array configurations was measured using four optimisation criteria: average power absorption, q-factor, relative capture width (RCW) and tether force. While ultimately we were seeking to maximise power absorption, the other three criteria provided deeper and potentially more valuable insights into ideal array configurations. More specifically, the q-factor is a measure of the effectiveness of the array compared to using the same buoys individually, RCW represents the power extracted by each buoy with respect to its size, while tether force provides some indication of the cost of different configurations by measuring the combined forces of the power take-off system.

The radii optimisation experiments involved using brute force search and random local search to explore different selections of buoy radii using any combination of 2m, 2.5m, 3.2m, 4m or 5m as the possible radii. Across different array sizes from 1x1 through to 5x5, the results consistently showed that 5m buoys were best for maximising power absorption, while 2m buoys were ideal for minimising tether force which can subsequently reduce system running costs. For some smaller arrays, q-factor could be maximised using specific arrangements of 2m and 5m buoys, while larger arrays favoured the 2m buoys for optimising q-factor due to the increase of destructive interference in the array as additional buoys were included in the configuration. RCW could also be optimised by using specific arrangements of 2m and 5m buoys, particularly for larger arrays where the random local search technique proved to be valuable in identifying these arrangements.

For spacing optimisation, we explored trends in the four optimisation criteria as the spacing between the buoys was varied using iterative search. Arrays consisting of the larger 5m buoys, were more influenced by changes to buoy spacing, but also showed more favourable results in terms of power absorption and RCW. The overall trend was that increasing the spacing resulted in larger values of all optimisation criteria, albeit with diminishing returns after 200m for some configurations. Therefore, the ideal choice appears to be the 5m buoy array with as much spacing as possible within practical constraints, where the only real disadvantage is the increase to tether force which may lead to higher operational costs.

Using the insights from the spacing experiments, we conducted layout optimisation of arrays ranging from 2x2 to 5x5 using large placement areas to allow exploration of the search space. The results showed that although larger arrays produce more power, there is also a general decrease in RCW, q-factor and tether force due to an increased number of buoys creating more destructive interference. In general, the layout experiments showed that specific arrangements of buoys can create constructive interference in the array, such as placing the buoys in certain paired arrangements in the 2x2 array to achieve a q-factor greater than 1. The remaining buoys should be placed along the perimeter of the placement area to increase buoy separation distance, which should minimise the effect of any destructive interference and ultimately increase the power absorption of the array.

Overall, we hope that our findings will help to inform the development and deployment of fully submerged wave energy farms, such as CETO, by offering valuable insights and approaches for maximising wave energy capture. By contributing to the field of wave energy research, we hope to further the use of wave energy as a viable and competitive form of renewable energy moving into the future. To this end, a selection of our results have been published as part of a conference paper in the 2016 Genetic and Evolutionary Computation Conference [45].

7.1 Future Work

This research has a number of directions for future work:

1. Optimising other aspects of the CETO array, such as the power-take off system, or optimising both radius and layout simultaneously to potentially find even more interesting and novel solutions.

2. Considering efficiency as an optimisation criterion, similar to efficiency measures explored in related work [32,38]. In the CETO array, efficiency could be measured by combining tether force with power absorption to provide an idea of the power per unit cost.

3. Optimising multiple conflicting criteria simultaneously by using multi-objective algorithms to find the set of Pareto optimal solutions that provide different criteria trade-offs. An example is trying to optimise power absorption and system cost at the same time. There may be multiple optimal solutions depending on the relative importance of power and cost to an end user, so it may be appropriate to offer these different solutions to the user and allow them to exercise their judgement in selecting the best solution.

4. Exploring other optimisation techniques, other than random local search and CMA-ES. Developing optimisation algorithms that can exploit knowledge of the problem domain could lead to more efficient exploration of the search space and ultimately the ability to find even higher performing array configurations.

5. Exploring mixed-initiative optimisation that combines the knowledge and experience of a human expert with the computational capabilities of a machine. Preliminary work on a mixed-initiative system has been conducted and the result is shown in Figure 7.1. A human expert is able to manually drag the buoys to different positions in the placement area and then request the machine to either evaluate the performance of the given array or attempt to improve the array further by using the buoy positions as inputs to an optimisation algorithm.

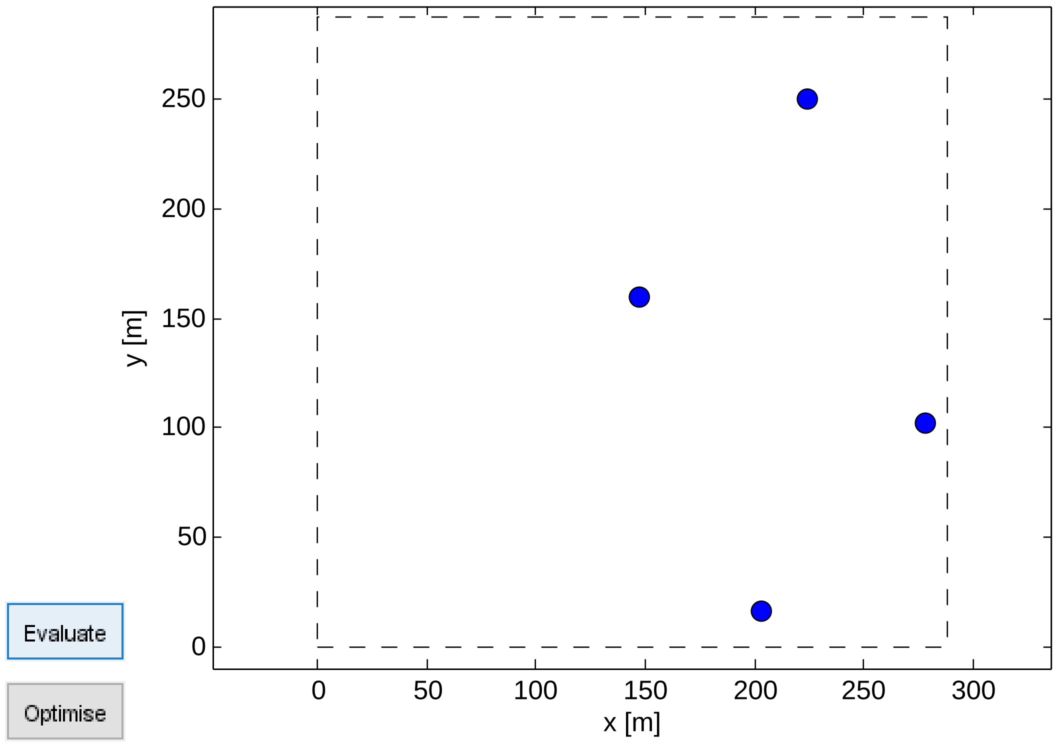

Figure 7.1: A screenshot of a mixed-initiative optimisation system that allows a human expert to manually evaluate different array layouts or refine them further using optimisation.

Bibliography

[1] A. Babarit. On the park effect in arrays of oscillating wave energy converters. *Renewable Energy*, 58:68–78, 2013.

[2] A. Babarit, J. Hals, M. Muliawan, A. Kurniawan, T. Moan, and J. Krokstad. Numerical benchmarking study of a selection of wave energy converters. *Renewable Energy*, 41:44–63, 2012.

[3] G. Bacelli and J. Ringwood. Constrained control of arrays of wave energy devices. *International Journal of Marine Energy*, 3(4):e53–e69, 2013.

[4] R. Banos, F. Manzano-Agugliaro, F. Montoya, C. Gil, A. Alcayde, and J. Gómez. Optimization methods applied to renewable and sustainable energy: A review. *Renewable and Sustainable Energy Reviews*, 15(4):1753–1766, 2011.

[5] S. Bellew, T. Stallard, and P. Stansby. Optimisation of a heterogeneous array of heaving bodies. In *Proceedings of the 8th European Wave and Tidal Energy Conference*, pages 519–527, 2009.

[6] B. Borgarino, A. Babarit, and P. Ferrant. Impact of wave interactions effects on energy absorption in large arrays of wave energy converters. *Ocean Engineering*, 41:79–88, 2012.

[7] K. Budal. Theory for absorption of wave power by a system of interacting bodies. *Journal of Ship Research*, 21(4), 1977.

[8] B. Child and V. Venugopal. Optimal configurations of wave energy device arrays. *Ocean Engineering*, 37(16):1402–1417, 2010.

[9] B. F. M. Child. On the configuration of arrays of floating wave energy converters. 2011.

[10] J. Cruz, R. Sykes, P. Siddorn, and R. E. Taylor. Wave farm design: preliminary studies on the influences of wave climate, array layout and farm control. In *Proceedings of the 8th European Wave and Tidal Energy Conference*, 2009.

[11] J. Cruz, R. Sykes, P. Siddorn, and R. E. Taylor. Estimating the loads and energy yield of arrays of wave energy converters under realistic seas. *Renewable Power Generation, IET*, 4(6):488–497, 2010.

[12] A. de Andrés, R. Guanche, L. Meneses, C. Vidal, and I. Losada. Factors that influence array layout on wave energy farms. *Ocean Engineering*, 82:32–41, 2014.

[13] B. Drew, A. Plummer, and M. N. Sahinkaya. A review of wave energy converter technology. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 223(8):887–902, 2009.

[14] A. F. O. Falcão. Wave energy utilization: A review of the technologies. *Renewable and sustainable energy reviews*, 14(3):899–918, 2010.

[15] C. Fitzgerald and G. Thomas. A preliminary study on the optimal formation of an array of wave power devices. In *Proceedings of the 7th European Wave and Tidal Energy Conference*, 2007.

[16] P. B. Garcia-Rosa, G. Bacelli, and J. V. Ringwood. Control-informed optimal array layout for wave farms. *IEEE Transactions on Sustainable Energy*, 6(2):575–582, 2015.

[17] M. Göteman, J. Engström, M. Eriksson, and J. Isberg. Optimizing wave energy parks with over 1000 interacting point-absorbers using an approximate analytical method. *Int. J. Mar. Energy*, 10:113–126, 2015.

[18] J. Hals, J. Falnes, and T. Moan. A comparison of selected strategies for adaptive control of wave energy converters. *Journal of Offshore Mechanics and Arctic Engineering*, 133(3):031101, 2011.

[19] N. Hansen. CMA-ES source code, 2016. URL: `https://www.lri.fr/~hansen/cmaes_inmatlab.html`.

[20] N. Hansen. The CMA evolution strategy: A tutorial. *ArXiv e-prints*, 2016. `arXiv:1604.00772`.

[21] N. Hansen, S. D. Müller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18, 2003.

[22] H.-J. Koh, W.-S. Ruy, I.-H. Cho, and H.-M. Kweon. Multi-objective optimum design of a buoy for the resonant-type wave energy converter. *Journal of Marine Science and Technology*, 20(1):53–63, 2015.

[23] U. A. Korde and R. C. Ertekin. Wave energy conversion by controlled floating and submerged cylindrical buoys. *Journal of Ocean Engineering and Marine Energy*, pages 1–18, 2015.

[24] A. Kurniawan. Modelling and geometry optimisation of wave energy converters. 2013.

[25] M. Lagoun, A. Benalia, and M. Benbouzid. Ocean wave converters: State of the art and current status. In *2010 IEEE International Energy Conference*, pages 636–641. IEEE, 2010.

[26] I. López, J. Andreu, S. Ceballos, I. M. de Alegría, and I. Kortabarria. Review of wave energy technologies and the necessary power-equipment. *Renewable and Sustainable Energy Reviews*, 27:413–434, 2013.

[27] P. A. Lynn. *Electricity from Wave and Tide: An Introduction to Marine Energy.* John Wiley & Sons, 2013.

[28] L. Mann, A. Burns, and M. Ottaviano. CETO, a carbon free wave power energy provider of the future. In *Proceedings of the 7th European Wave and Tidal Energy Conference*, 2007.

[29] L. D. Mann. Application of ocean observations & analysis: The CETO wave energy project. In *Operational Oceanography in the 21st Century*, pages 721–729. Springer, 2011.

[30] A. McCabe, G. Aggidis, and M. Widden. Optimizing the shape of a surge-and-pitch wave energy collector using a genetic algorithm. *Renewable Energy*, 35(12):2767–2775, 2010.

[31] M. E. McCormick. *Ocean wave energy conversion*. Courier Corporation, 2013.

[32] M. Mohamed, G. Janiga, E. Pap, and D. Thévenin. Multi-objective optimization of the airfoil shape of Wells turbine used for wave energy conversion. *Energy*, 36(1):438–446, 2011.

[33] G. Nolan, J. Ringwood, W. Leithead, and S. Butler. Optimal damping profiles for a heaving buoy wave-energy converter. In *Proceedings of the Fifteenth International Offshore and Polar Engineering Conference (ISOPE)*, volume 1, pages 477–485, 2005.

[34] G. Nunes, D. Valério, P. Beirao, and J. S. Da Costa. Modelling and control of a wave energy converter. *Renewable Energy*, 36(7):1913–1921, 2011.

[35] J. Ringwood and S. Butler. Optimisation of a wave energy converter. In *IFAC Conference on Control Applications in Marine Systems*, pages 7–9. Maynooth University, 2004.

[36] S. H. Salter. Wave power. *Nature*, 249(5459):720–724, 1974.

[37] N. Sergiienko, B. Cazzolato, B. Ding, and M. Arjomandi. An optimal arrangement of mooring lines for the three-tether submerged point-absorbing wave energy converter. *Renewable Energy*, 93:27–37, 2016.

[38] C. J. Sharp. Wave energy converter array optimization: Array economic analysis and preliminary results of a genetic algorithm approach introducing cost factors. 2015.

[39] L. V. Snyder and M. M. Moarefdoost. Layouts for ocean wave energy farms: Models, properties, and heuristic. In *Proceedings of the 2nd Marine Energy Technology Symposium (METS2014)*, 2014.

[40] L. V. Snyder and M. M. Moarefdoost. Optimizing wave farm layouts under uncertainty. In *Proceedings of the 3rd Marine Energy Technology Symposium (METS2015)*, 2015.

[41] M. Srokosz. The submerged sphere as an absorber of wave power. *Journal of Fluid Mechanics*, 95(04):717–741, 1979.

[42] P. Troch, C. Beels, J. De Rouck, and G. De Backer. Wake effects behind a farm of wave energy converters for irregular long-crested and short-crested waves. *Coastal Engineering Proceedings*, 1(32):53, 2011.

[43] D. Valério, P. Beirão, and J. S. da Costa. Optimisation of wave energy extraction with the Archimedes Wave Swing. *Ocean Engineering*, 34(17):2330–2344, 2007.

[44] S. Weller, T. J. Stallard, and P. K. Stansby. Experimental measurements of irregular wave interaction factors in closely spaced arrays. *IET renewable power generation*, 4(6):628–637, 2010.

[45] J. Wu, S. Shekh, N. Sergiienko, B. Cazzolato, B. Ding, F. Neumann, and M. Wagner. Fast and effective optimisation of arrays of submerged wave energy converters. In *Proceedings of the 2016 Genetic and Evolutionary Computation Conference*. ACM, 2016.