Optimisation of Scheduling and Routing For Offshore Wind Farm Maintenance

Toby Karl Kingsman, M.Sci., M.Res

Submitted for the degree of Doctor of Philosophy at Lancaster University.

May 2020

STOR-i
excellence with impact

Abstract

The growing increase in the size and scope of offshore wind farms motivates the need for industry to have access to mathematical tools that reduce costs by efficiently performing daily operations and maintenance activities. Key offshore activities require the transportation of technicians to and within offshore wind farms to complete corrective and preventive maintenance tasks to keep turbines operating efficiently.

We provide a new deterministic mixed integer linear programming formulation for deciding the optimal vessel routes for transporting technicians around a wind farm and the scheduling of crew transfers, by minimising downtime, travel and technician costs. The model contains sufficient flexibility to account for multiple vessels, shifts and task profiles, whilst being able to prioritise and omit tasks in environments containing limited resources. Computational experiments are performed which quantify and confirm the impact of key instance characteristics such as technician availability, task profiles and weather conditions. We implement and evaluate the impact of a novel industry safety constraint.

The complexity of larger instances motivates a second continuous time formulation, in which preventive maintenance again requires no minimum duration of work before

it can provide benefit. We employ a specific decomposition structure to take advantage of variable preventive maintenance and utilise an adaptive large neighbourhood search procedure to solve instances. We evaluate several distinct acceptance criteria in conjunction with random and adaptive operator selection to determine the best option for our model.

We produce a statistical model of offshore weather conditions to help quantify the likelihood of limited vessel accessibility to offshore wind farms. We model the joint distribution of key meteorological and oceanographic variables over time whilst accounting for seasonal trends using multivariate kernel density estimation. Our method generates alternative metocean realisations from historical data and reproduces the important long term persistence statistics of good and adverse offshore conditions.

Acknowledgements

I would like to take this opportunity to thank the many people who have helped during the arduous process of completing this PhD. Firstly thanks must go to my supervisors, Dr. Burak Boyaci and Professor Jonathan Tawn from Lancaster University, for supporting and guiding me during this process. I am also grateful to JBA Consulting for proposing this project and for all the people who have helped me within the ForeCoast Marine team over the last 3-4 years. Particular thanks must go to Mark and Ed, for their helpful discussions and being my main points of contact with JBA.

My time at Lancaster has been made infinitely more enjoyable by being a part of the STOR-i Centre for Doctoral Training and having access to all of the opportunities they provide. I have been lucky to have studied alongside many fantastic people and friends, who I am sure will all go on to great things. I'd especially like to thank my fellow year group who are a truly great bunch of people. Consistent highlights of the week have been Tuesday afternoons consisting of football and the White Cross pub quiz. Thanks to all the attendees over the years, I hope to get to do "just one more" one day!

I should also thank all the other people I have met during my time at Lancaster

including the scores of Graduate college football teammates and fellow Pokémon Go trainers. Duck sweet!

A very special thanks to my parents and siblings for all their support during the ups and downs of doctoral research. I would not have been able to finish this without any of you. I greatly appreciate you all for helping me to finish off writing this thesis during the COVID-19 lockdown; thank you to Hannah for keeping me "on task".

Finally I am grateful for the chance to learn more about the world in which my grandfather Brian worked, something which I stumbled upon by accident nearly 6 years ago. I would have loved to have been able to discuss the world of Operational Research with you, but I'm sure you'd be proud of where I ended up.

Declaration

I declare that the work in this thesis has been done by myself and has not been submitted elsewhere for the award of any other degree.

Toby Karl Kingsman

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

All across the globe there exists a growing need to produce more electricity than ever before. Both developed and developing countries need new infrastructure to either expand their present levels of production to cope with rising demand or to simply replace existing infrastructure as it reaches the end of its operating lifetime.

Traditionally economic factors are the dominant influence in the long term strategic decisions involved in deciding a countries future energy mix, however there is now also a need to prioritise switching to 'greener' forms of electricity generation in order to reduce carbon dioxide emissions. This is necessary because of the immediate threat posed by a rapidly warming planet which is driven by the presence of excessively high levels of greenhouse gases. As the energy sector in the UK still makes up around 30% of the UK's CO_2 emissions (UK Government, 2015) it is important to find ways of reducing the amounts of these harmful pollutants.

In recent years this has happened through the adoption of medium and long term targets aimed at shifting the UK's energy production in the direction of more renewable forms of energy, examples of which are wind, wave and solar power plants. The UK itself has placed more emphasis on wind energy due to the favourable local conditions found for it around the British Isles. In 2016 wind initially contributed around 6-7% of the UK's total energy needs, (UK Energy Statistics, 2016), whereas by the third quarter of 2019 it was closer to 20%, (Carbon Brief, 2019).

The first generation of wind turbines built in the UK were situated on land and grouped together into nearby locations known as 'wind farms'. The effect from taking this approach was that the average power output across the onshore wind farm became more stable; small fluctuations in wind speed at a specific turbine were balanced out over the entire set of turbines. In conjunction with the development of cheaper and more efficient turbines onshore wind has become one of the most mature and cost effective forms of renewable energy. Despite this it is still subject to several criticisms. These tend to be primarily environmental and quality of living concerns. Many local residents oppose their construction on the grounds of them causing blights on unspoiled landscapes, excessive noise and damage to local wildlife. Although many of these issues have been addressed by industry through new technologies such as quieter blades, focus has shifted in recent years towards building new wind farms out at sea.

This change in policy can be seen as a beneficial switch in terms of both energy generation and environmental considerations. It has been shown that the wind speeds found at positions just a few hundred metres offshore are substantially higher than

those found inland, often reaching the range of being 40% faster, (USA Department of Energy, 2017). This equates to the energy potentially extractable from the wind being nearly three times higher out at sea making these new regions highly desirable. Another benefit of locating these turbines in open waters is the reduced visual impact to any coastal observers and noise pollution becoming negligible. Whilst there are still other issues that need to be accounted for such as fishing areas and marine habitats, the decision has been made that the extra cost associated with offshore wind is worth the price in order to avoid the drawbacks of onshore wind.

Compared to onshore wind, offshore wind farms remain comparatively expensive to develop and operate although these costs have rapidly decreased ahead of projections. In 2011 it cost £136 for each megawatt-hour of offshore wind produced which dropped to £121/MWh by 2014 ahead of a targeted £100/MWh in 2020, (CATAPULT, 2015). In actuality by the third quarter of 2019 this target had already been demolished with some new offshore wind farm contracts containing strike prices around £40/MWh, (offshoreWindBiz, 2019).

These reductions can largely be attributed to the introduction of newer 6MW turbines, but more developments are required to make offshore wind farms profitable without government subsidies. The rapid growth expected in the industry over the next decade means that it is important to not only enhance production but also to lower operating costs.

One way of making wind energy more competitive in a crowded energy marketplace is clamping down on the day to day operating costs. At present these costs make up a large percentage of the total costs over the lifetime of the wind farm, approximately

20-35% (Shafiee, 2015). The demand from industry to tackle this has pushed ideas concerning efficient and effective maintenance scheduling to the forefront.

The key reason for the incursion of such unusually high operating costs is the fact that the costs of travelling to offshore wind sites are significantly higher than those on land. In order to transfer the necessary technicians and equipment to the wind farm, specialised vessels and helicopters are required which are expensive to rent. The money involved with acquiring and utilising these vessels dominates other contributory factors to the total maintenance cost since it is strongly related to losses in power production. Vessels are used to help perform maintenance tasks on wind turbines in order to keep them operational. If said vessels are unable to service faults then the afflicted wind turbines will often be unable to generate electricity thereby missing out on potential revenue. This accessibility of the wind farm is markedly different from onshore activities because of the impact of the weather and sea conditions. Bad weather and heavy seas can limit the number of tasks that can be performed in a day or even completely restrict access to the wind farm when maintenance activities would otherwise be performed.

As faults and breakages in offshore wind farms can potentially have a longer lasting impact than the equivalent breakdown occurring onshore, it is necessary for maintenance operators to decide upon an efficient strategy for dealing with them. Several different strategies have been used that employ different kinds of operations such as reactive, preventive and condition based maintenance to reduce the likelihood of critical failures. Usually a combination approach involving all these types of strategies make for the most sensible and cost effective policy.

1.2 Maintenance Policies

The offshore maintenance problem gets considerable attention from both industrial and academic circles (Hassan, 2013), however before any decisions concerning the maintenance fleet or operations scheduling can be made, a decision concerning the type of maintenance strategy used to cover the wind farm must be taken. Historically there have been several different methods for servicing the wind turbines in order to rectify any breakages and prevent future periods of downtime. These strategies can loosely be classified into reactive maintenance and preventive maintenance, (Shafiee, 2015).

Reactive maintenance is defined as maintenance tasks which are performed in response to any critical or major failures in the offshore wind farms. These are likely to be unpredictable faults that occasionally occur causing the turbine to cease production of electricity for the period of time that the fault persists or alternatively requires it be shut down due to safety considerations. As these stoppages in production are extremely costly to the operator, they are important to resolve as soon as possible which can conflict with the current accessibility of the wind farm. It has been shown by Van Bussel et al. (2001) that using a purely reactive maintenance strategy is a poor choice which produces long periods of downtime and causes excessively high repair costs making the strategy unappealing. These drawbacks are present even for turbines located relatively close to the shore.

Preventive maintenance strategies are rooted in the idea that it can be more economical to spend money to check the condition of components before they break in

order to prevent the need for more costly critical repairs in the long term. These types of inspections are scheduled in advance compared to the on demand nature of reactive maintenance and tend to be periodic in nature. The key issue lies in determining the frequency of these preventive visits; repeated visits will cause costs to quickly build up whilst visits in isolation will be ineffective at detecting likely faults.

The advent of modern monitoring technologies make it possible to improve upon the idea of carrying out simple time based preventive maintenance activities by using sensors that can measure the wear and tear of particularly susceptible components in the wind turbine. These condition based monitoring systems allow for inspections and replacements to be scheduled automatically based on the occurrence of key events. This might be an increase in temperature or corrosion above a certain threshold. The present trend tends to be use several different basic monitoring systems (Mérigaud and Ringwood, 2016) in order to cope with the limited computational power available on site, however this may change in the future.

Most wind farms will eventually employ some kind of hybrid system mixing preventive replacements, inspections and reactive repairs all together. This will require some notion of a priority system where critical repairs take precedence over scheduled inspections in the scheduling protocol. In fact, each individual task type will have different properties that need to be accounted for in terms of the equipment needed, completion time and personnel needed. For example gearbox alterations inside the nacelle (cover housing of key generating components) of the turbine will take longer to complete than inspecting the foundations complicating the scheduling procedure, (Carroll et al., 2015).

1.3 Available Vessels

The variety of tasks present at offshore wind farms has led to a wide selection of vessels suitable for transporting equipment and technicians becoming available for the offshore wind maintenance marketplace. Often vessels have been specifically designed to cope with the demands of offshore wind but some are sourced directly from the existing offshore gas and oil sector. These craft are generally divided into two main categories based on the size of the job that they typically undertake, usually either minor or major repairs (Dalgic et al., 2013).

Minor repairs do not require heavy equipment or the use of specialised lifting vessels such as cranes and jack-up barges. These tasks could include visual inspections or simple repairs and replacements of easy to access components such as sensors or hydraulics. Minor maintenance jobs tend to be performed by smaller and faster boats such as monohulls, catamarans (as shown in Figure 1.3.1a) and SWATH (Small Waterplane-Area Twin Hull) vessels which are collectively known as types of crew transfer vessels. Their main use is for ferrying technicians and equipment to and from wind turbines in order to perform these routine tasks. Due to their size they are limited in the amount of people they can carry. Most monohulls and catamarans have a capacity of around 12 technicians. SWATH vessels can hold more technicians and equipment but at the cost of vessel speed. They travel at a speed of around 15 knots compared to the faster catamarans (20 knots) and monohulls (25 knots). Another important factor to consider are the operating restrictions imposed by the external weather conditions. Crew transfer vessels can only operate in relatively light

seas up to significant wave heights of around 1-1.5m which tend to limit their overall accessibility in comparison with larger vessels. Speed is also an issue as a typical distance from port to wind farm can mean that round trips potentially approach a total of 4 hours each day. Whilst this is all time which could be spent carrying out repairs it is offset by their relatively inexpensive chartering cost.

More serious repairs might need the presence of larger and better equipped vessels that can support heavy lifting operations to replace large components such as wind blades. These vessels tend to be either outfitted with a crane for lifting purposes or are themselves some kind of jack-up vessel. These jack-up vessels, as illustrated in Figure 1.3.1b, form a stable platform for operations by deploying legs that reach down to the surface floor and elevate the ship out of the water. This is a slow process with the ship needing to be present for the entirety of the maintenance activity and their operation being limited to periods of low wave heights and wind speeds, (Sperstad et al., 2016). At present there is a low availability of these vessels causing their rental costs to ascend significantly past smaller crew transfer vessels (Dalgic et al., 2015). Chartering costs per day run into the tens of thousands of pounds making their usage important to optimise. The lead time of ordering these vessels is also sizeable as they often have to be moved from one wind farm to another when hired. Even a relatively short journey from the North Sea to the Irish sea can cause a lead order time of about a week assuming no weather delays.

(a) A catamaran crew transfer vessel (Alicat Workboats, 2016)

(b) A jack-up vessel (A2Sea, 2016)

Figure 1.3.1: Some typical offshore vessels

Recently a larger class of vessels, known as service operations vessels, have been constructed that are capable of staying out at sea for a considerably longer period than smaller crew transfer vessels. As they are bigger ships they are able to undertake both major and minor repair operations by using more advanced facilities such as hydraulic gangways and cranes. Their larger size also means that they can carry substantially more technicians and equipment whilst also being able to operate in heavier seas with a significant wave height limit around 2.5-3m. Although the costs of running such vessels are double those of smaller vessels they have substantially higher accessibility as they can operate on a 24 hour basis out at sea. They have been designed primarily as a response to the industry's need for vessels that can operate in wind farms located in deeper waters, whilst simultaneously avoiding the increasingly costly journey from onshore bases.

Another extension to the idea of longer term support vessels are conceptual mothership designs as seen in Figure 1.3.2b. The key difference between service operation

vessels and motherships is that motherships can stay out of port for an even longer period of time and launch smaller crafts to carry out the maintenance tasks. Generally they accommodate more technicians and resources with current designs looking to accommodate around 200 technicians with space for helicopter landing pads and smaller crew transfer vessels. In a sense this is the evolution of another innovation in the offshore industry: offshore bases. These are permanent structures that are located near to wind farms and permanently manned with technicians and spare parts. In a similar vein to offshore oil and gas platforms they cut down the need to make expensive and frequent trips to the mainland every day just to access the wind farm. It is anticipated that both of these ideas will become more common in the future as construction costs reduce and wind farms are built further away from land.

As wind farms get built further out to sea the use of helicopters in the maintenance fleet also becomes more viable. A helicopters top speed is 4 or 5 times greater than crew transfer vessels, which can help to provide another way of quickly and effectively moving small numbers of technicians to turbines. If this method of transportation is employed tasks serviced via helicopters often end up accumulating significantly amounts of downtime losses compared to if they were accessed by surface vessels. Helicopters are more commonly used in winter because of their ability to operate independently of wave heights and wind speeds. This gives them the highest levels of accessibility of any type of craft. Their downfall comes in the form of their exorbitant operating costs making their practical usage hard to justify. In conjunction with this is that their operation depends on having good levels of visibility, with little fog (Domínguez-Navarro et al., 2014).

(a) A helicopter transfer (Fiberline, 2016)

(b) A mothership design (Focus, 2011)

Figure 1.3.2: Some atypical offshore vessels

1.4 Weather Impact

It is also important to consider influence of weather patterns on both the long term and short term task scheduling and routing problem at offshore wind farms. The long term goal of reliable wind energy generation is chiefly driven by the velocity of the wind, so successful offshore wind farms will need to attempt to maximise the amount of time that turbines are operating for, particularly during periods of windy weather. As the rate of electricity generation depends on the cubic power of wind speed the financial losses sustained from downtime during lower wind speeds are significantly lower than those from periods of high wind speeds. This means that systematic periods of lower windspeed such as the summer months should contain more planned activities in order to produce cost effective schedules.

Individual journeys made by crew and vessels are also affected by the weather even if they can still operate in the prevailing conditions. For example, if the wave heights are higher than normal there will be a repercussive effect on the travel times of vessels. If lots of vessels are scheduled to make trips to the wind farm, and the

turbines within it, this increase in journey time can quickly accumulate into large periods of wasted time.

The daily scheduling of vessel routes can also be impacted by offshore storms meaning that planned routes may not necessarily even be feasible. Short term storms have been known to develop within a days' operations and can cause the need for rapid changes in the schedule with little or notice given. If a day turns out to be too windy, or waves too high, then vessels will be confined to port for an extended period of time, or worse forced to abort their day's activities and return to their home base. This is a worst case scenario as both money and manpower will have been wasted without any restoration in power production or downtime reduction. The overall uncertainty of the weather and sea conditions means that it is wise to have access to a portfolio of vessels that can operate throughout a range of conditions, since poor weather and sea conditions consequently limit the accessibility of the wind farm. This motivates the need for designing vessel routes that take into account the uncertainties in predictions of offshore conditions.

Existing weather forecasts are known to be highly accurate in the short term with very small uncertainty about their predictions. The chaotic nature of weather systems means that uncertainty increases as the projections are made further into the future. Reliable ensemble forecasting methods are usually only applicable up to 7-10 days into the future, after which statistical models of long term metocean (meteorological and oceanographic) conditions are required. These statistical models need to account for the seasonal trends and site-specific details to be applicable for offshore wind farms.

1.5 Routing and Scheduling

In this thesis we focus on the daily operational problem of how best to transport technicians to complete maintenance tasks offshore. We therefore ignore the tactical level decisions associated with how best to acquire the vessels and technicians that we utilise.

When a task is created through an unexpected failure or as a preplanned activity it is classified based upon the urgency that it needs to be dealt with. The more severe tasks such as critical repairs will need to be attended to ahead of less severe tasks which have more flexibility in their completion date. In an ideal world we would schedule a boat and crew to each individual task as soon as they occur, however the costs of doing so are prohibitively large. This is particularly apparent for tasks that involve more expensive jack-up vessels where logic dictates that a more financially sound strategy might be to allow several jobs needing a jack-up vessel to accumulate before one is hired from the spot market.

In order to efficiently utilise a limited set of resources these maintenance activities should not be carried out by individual crews and vessels but rather be grouped together into 'activity bundles'. These activity bundles consist of tasks grouped together into a batch including both preventive inspections and component replacements to be undertaken by a single vessel on its daily trip to a wind farm. Normally because of the length of time needed to complete each task and large transit times to reach the wind farm the bundle will contain no more than 4 or 5 activities in total.

At the start of a shift vessels will depart laden with technicians and equipment

destined to be used at the wind farm. It is assumed that due to the lengthy travel time and fuel costs that each vessel will only access the wind farm once per day, ensuring that every ship will spend its shift working in the vicinity of the offshore wind farm. Once the ship reaches the wind farm it will navigate around to the turbine with the next maintenance task scheduled to be completed as part of its activity bundle on it. After docking so as to transport the assigned workforce and spare parts across the work can begin.

We assume that each activity can be completed by the technicians and equipment transported with the use of a single vessel. These workers then stay at the wind turbine they were deposited at for the tasks duration before waiting to be picked up again by the same vessel. Depending on the scheduled route and the vessel used it is possible for the same technicians to be dropped off elsewhere to complete additional tasks. This is obviously subject to appropriate time constraints and the predicted weather patterns. At the end of a shift the vessels must return back to their base with all of the workers they brought so that no one is abandoned.

It is possible that random faults and failures will be detected during the periods of time when maintenance is being carried out. If this occurs then the easiest method of dealing with these additional tasks is to assign any vessels with spare time, parts and personnel to deal with these as soon as possible, ideally on the same day. This could even be vessels currently located within the wind farm if they match the requirements of the job known better as 'opportunistic maintenance'. If this is not practical then new tasks can simply be added to the bank of activities that still need to be scheduled in the future subject to the usual caveats. Most optimisation based models tend to

take this approach by assuming that any newly occurring tasks are only considered for scheduling from the period after they were detected.

Due to the sheer scale of optimisation models that can capture the key features of the problem it is likely that any models produced will take an appreciably long time to solve to optimality. In order to obtain solutions within a time window useful for industrial purposes we will most likely need to simplify the situation by incorporating a heuristic based approach into the operational model itself. This is known as a 'rolling horizon' approach whereby the planning horizon of the problem is split into two separate periods in order to reduce the total computational complexity.

The first period is a highly detailed system that incorporates all the details of the routing problem to create specific routes for vessels, personnel and equipment that can form the basis of schedules but only for a few initial days. In order to keep the run time of the optimisation relatively low the second period does not attempt to schedule in detail like the first period. Instead the second period tries to ensure that tasks do not get moved around in a manner such that they or others do not become infeasible at a later time. This approach also naturally links in with the fact that weather forecasts are most accurate for one or two days into the future thus limiting the usefulness of any longer term detailed scheduling. Whilst these heuristic simplifications would potentially reduce the quality of the overall solution, we can apply the procedure in an iterative fashion to produce routes for a much longer scenario. We first solve the model to produce a solution for the first period and later periods. Then we try to implement the vessel routes provided by the detailed first period. Once those tasks have been performed, we can rollover to the second day in the scenario by removing

the tasks and work that was actually completed. We note that the work planned and actually completed may differ due to real life uncertainties. The size of the problem to solve starting from the second day is now much smaller and informed by what actually happened in the first period. We note that every activity will thus be performed once it reaches the shorter planning period that is actually implemented.

Heuristics could also be used in addressing cases of the operational problem where a sudden adjustment in the route scheduling is needed. These changes could occur just before the start of a maintenance shift if a vessel required to take part in the days operations became unavailable for some reason. Alternatively the weather could potentially change dramatically before or during operations to the point where the current schedule becomes sub-optimal or infeasible. In these instances it is imperative that a solution to the days activities can be found in a quick manner with the desire for an effective solution being willingly sacrificed to do so.

A statistical model of the weather and sea conditions surrounding an offshore wind farm is important to build as they govern whether the farm is accessible to the vessels or not on a daily basis. As metocean conditions can exhibit substantial variation over time and space it is unwise to simply assume a fixed level of accessibility over a prolonged period of time. Shorter range weather patterns such as storms can occur on a daily basis and often coincide with longer range and more periodic activity such as wind-sea waves. These constituent elements and their interactions need to be understood in order to capture the key elements of the marine conditions. This knowledge can then be utilised to supplement optimisation based routing models for a given shift with a range of possibilities for metocean conditions in subsequent shifts.

We will restrict our focus to the standard variables used in the literature of windspeed and significant wave height that form the core of industry standards for safe operating windows. It has already been shown however that more accurate predictions can be developed by including additional variables Sperstad et al. (2014).

1.6 Research Goals

The main goals of this project was motivated by the needs of our industrial partners, JBA Consulting, with the aim of contributing to the development of their ForeCoast®Marine product. This is a complicated risk management tool that combines forecasting, physical and statistical techniques together to optimise complex marine activities and identify the most favourable time to undertake these operations.

ForeCoast®Marine is currently underpinned by a powerful simulation tool designed to evaluate the impact of the weather and sea conditions on offshore activities. Typically these marine activities require the use of vessels and technicians to complete a set of predefined activities at offshore locations as efficiently as possible. The sequence in which these activities are completed is flexible, with the goal to find a sequence that minimises a given cost function. As real life problems contain ever growing numbers of activities there is a combinatorial explosion in the number of feasible solutions, many of which are not necessarily obvious to a human decision maker. Furthermore, it is challenging to explore all these possibilities in a rule based simulation tool, suggesting the need for a future approach that leverages the power

of mathematical optimisation.

The main objective of our work is to develop an optimisation component that would form a natural extension to ForeCoast®Marine, by considering in greater detail the exact routing of vessels and technicians to offshore tasks. This model should be able to provide knowledge of the exact journeys and crew transfers required to be made by vessels in order to complete the offshore maintenance activities in a cost efficient manner. It should also incorporate some new problem domain knowledge yet to be fully addressed in the literature. The solution methodology will likely need to be heuristic in order to scale well for the larger anticipated problems in the future. This would allow JBA Consulting to both increase the range of problems that can be solved with ForeCoast Marine and increase the quality of routing solutions within them.

One current use of ForeCoast®Marine is the planning of installation campaigns for building new offshore wind farms. These campaigns involve the placement of turbine foundations into the seabed, before constructing the remaining components such as blades on top. Due to their immense size and scale, these operations often taken place over the course of several months and require the chartering of specialised vessels. As these vessels are very expensive and in high demand it is important to not hire a vessel for any longer than needed, and especially important to not overrun a vessel's lease. Doing either of these will incur significant costs to the project. ForeCoast®Marine currently simulates the duration of these installation campaigns under the chief source of uncertainty that is the weather and sea conditions. In order to make the results of the simulation more accurate it needs to be repeated across many different weather

scenarios. Typically the number of these scenarios is limited to the existing historical data available for a particular site. Our goal for the statistical content of this project is to use this historical data to produce a statistical model from which any number of possible weather scenarios can be generated. This will allow for many more simulation replications to be performed and also be potentially able to fill in any missing or erroneous data within the historical data.

1.7 Thesis Outline

We now provide an overview of the content of each chapter in this thesis. In Chapter 2 we introduce our discrete time mixed integer linear programming formulation of the scheduling and routing problem for offshore wind farm maintenance. It incorporates several important features within a shift such as the accumulation of downtime losses at corrective maintenance tasks and operating restrictions implied by weather restrictions. We can confirm and quantify the significant impact that reducing the amount of available vessels and technicians have on the problem. Furthermore, the construction and impact of a novel safety constraint is evaluated.

In Chapter 3 we describe our approach for developing a heuristic method based on adaptive large neighbourhood search to solve larger instances of the offshore wind farm maintenance problem. This is employed to solve a set of tough real world instances that are characterised by the presence of a resource restricted environment.

Chapter 4 is devoted to our statistical model of metocean conditions from which we can simulate many different possible realisations of historical weather conditions.

It is designed to capture the temporal dependence of joint distributions through the use of multivariate kernel density estimation after extracting the seasonal trends of the data. Graphical modelling is introduced to help reduce the number of parameters in our model and improve the accuracy of the accuracy of the remainder. The thesis concludes with Chapter 5 which addresses our conclusions and possible avenues for future research.

Chapter 2

Exact Mathematical Model

2.1 Introduction

In order to meet the defined goals of green energy generation and environmental objectives many countries have focused on growing their renewable energy capacities. The UK alone brought over 2.1GW of offshore wind capacity online in 2018 (RenewableUK, 2018), as a part of its stated target to generate over 30% of all its electricity from renewable sources by 2020 (Energy and Climate Change Committee, 2016).

The rapid growth experienced in the offshore wind sector means there is an increasing opportunity to find savings from conducting operations and maintenance (O&M) activities more efficiently. Presently O&M constitutes up to 35% of a wind farm's total lifetime costs (Shafiee, 2015), providing a clear avenue for efficiency savings. The introduction of turbines with capacities in excess of 8MW has helped to dramatically reduce strike prices making offshore wind comparable with fossil fuels and nuclear power (RenewableUK, 2017). The continual accelerations in the size and

quantity of new offshore wind farms increase the likelihood of large quantities of daily O&M tasks simultaneously occurring, allowing for new economies of scale to be found and exploited in the regular trips made to service and maintain the turbines.

Offshore wind technology is expected to further evolve with novel concepts such as floating turbines designed to operate in deeper waters currently inaccessible with conventional turbine foundations and artificial islands intended to act as inter-connected power hubs between distant wind farms (North Sea Wind Power Hub, 2018). Condition based maintenance strategies that further reduce costs are becoming more common increasing the number of tasks to be performed across wind farms. Mathematical decision support tools will be needed to fully exploit these savings especially in the presence of limited resources and an uncertain offshore environment.

A successful mathematical model of offshore operations must consider both the routing and scheduling aspects of the problem in parallel. A pure scheduling approach cannot harness potential savings from sharing technicians across multiple vessels, whilst a stand alone routing model is unable to cope with discontiguous task completion and prioritisation. Our proposed model utilises properties of both routing and scheduling problems.

The structure of this chapter is as follows. Section 2.2 describes previous literature for short term offshore maintenance routing and outlines some key differences with our approach. A model description is provided in Section 2.3 with the exact mathematical model and its extensions. Section 2.4 details the results of the model on a set of experimental instances with Section 2.5 summarising our conclusions.

2.2 Literature Review

The offshore wind maintenance problem shares some similarities with that of its onshore counterpart as discussed in Kovács et al. (2011). They develop a mixed integer linear programming formulation that minimises lost electricity production and technician transportation costs from scheduled maintenance tasks and unexpected failures. It accounts for the hourly differences in power production on the current day and synthesises this with costs accumulated in later shifts to produce a rolling horizon framework. This model is extended to include a multi-skilled workforce working over several days (Froger et al., 2017). Tasks can also be executed in different modes and be postponed. Gutierrez et al. (2017) consider the onshore problem from a multi-objective viewpoint as there are often several competing stakeholders involved, before solving the resulting mixed integer linear program with an epsilon constraint algorithm. The drawback of applying these models to offshore maintenance is the significantly greater distance between the technician base and the turbines, which leads to the noted pick-up and delivery behaviour of vessels found in offshore wind. Another issue exists with lower offshore accessibility present for offshore wind farms which often prevents technicians visiting offshore locations.

Besnard et al. (2009) develop a model that schedules preventive maintenance operations in offshore wind farms from a cost beneficial viewpoint so they occur in periods of low wind speed and after corrective maintenance tasks are completed. This work is continued in Besnard et al. (2011) to produce a stochastic optimisation model over a longer time horizon. No uncertainty is modelled in the short horizon producing a

rolling horizon model applied to a scenario bush. Camci (2015) looks at the failure rates of turbines and develops a method that uses prognostic information to decide the best time to perform preventive maintenance. Opportunistic maintenance at nearby locations is considered because of the geographical spacing of turbine assets. These works all attempt to determine the best time to perform corrective and preventive maintenance in the medium term, but make no attempt to calculate the optimal routing to complete tasks on their planned day of execution.

Our problem is an extension of the capacitated vehicle routing problem and categorised as an example of a rich vehicle routing problem (RVRP) as defined by Caceres-Cruz et al. (2014), since it contains several features unique to offshore wind farms. A one-to-one pick-up and delivery structure is needed between the wind farm and its onshore base, whilst within the wind farm a many-to-many pick-up and delivery structure is present. The need for wind farms to operate with limited resources means the problem also shares several characteristics with the k-travelling repairman, Nucamendi-Guillén et al. (2016), and team orienteering problems, Chao et al. (1996).

A different offshore routing problem occurs in the construction of supply vessels schedules for offshore oil and gas platforms. It contains similar requirements to offshore wind farm maintenance, as supplies need to be transported by, and loaded/unloaded from, vessels. However the platforms are located significantly further offshore than offshore wind farms, so voyages typically last in excess of a week. They also sequentially visit a few platforms in a row, rather than revisiting multiple close locations at different times in a single voyage. The trade-offs between costs, emissions and robustness to weather conditions is examined by Norlund et al. (2015). They use a

simulation-optimisation methodology to construct weekly schedules of voyages, which will contain periods of waiting or idle time. The accumulation of these idle times leads to slack in the schedules which can either be used to reduce the speed of vessels, or as a buffer against bad weather. Norlund and Gribkovskaia (2017) extends this work by estimating the fuel consumption in different weather conditions, and evaluating the benefit of vessel speed optimisation when the weather uncertainty found in different seasons is included.

The explicit scheduling and routing of offshore wind turbine maintenance on an operational level is first considered by Dai et al. (2015). They develop a mixed integer linear program that finds the optimal sequence of turbines for vessels to visit to complete a set of maintenance tasks over several shifts subject to temporal and personnel constraints. Vessel travel costs and penalty costs for delayed maintenance tasks form the objective function of the model. A similar arc-flow mode for a single shift is presented and reformulated into a path-flow problem through Dantzig-Wolfe decomposition by Stålhane et al. (2015). They further differentiate the downtime costs between corrective and preventive tasks by assuming the latter are only incurred whilst work is being performed.

Irawan et al. (2017) generalise the existing models to include multiple operations and maintenance bases and wind farms. Generic personnel are replaced with a specialist multi-skilled workforce including roles such as electrical and mechanical technicians. Resources such as spare parts, personnel and vessels are treated as scarce and only available in limited quantities. Their approach follows previous solution methods in splitting the problem into a master problem of assigning vessels to routes

and a sub-problem for identifying feasible routes. The creation of feasible routes is an extension of Dai et al. (2015) and is performed with an algorithm incorporating a mixed integer linear program, however the use of an upper limit on the number of turbines a vessel can visit means that optimality is not guaranteed.

Raknes et al. (2017) integrate vessels that can stay offshore for several consecutive days and also tasks which span across multiple shifts into the existing problem. They continue with the distinction between the calculation of downtime costs for preventive and corrective tasks. The difficulty in choosing to perform preventive tasks (which can only be evaluated with a long term view) is overcome in the short term by specifying the number which should be completed in the immediate horizon. Due to the size of the mixed integer program, they propose rolling horizon heuristics to solve bigger instances.

Heuristic methods for short term routing and scheduling of maintenance tasks have also been investigated, chiefly with the aim of solving instances containing higher numbers of offshore tasks. Kennedy et al. (2016) base their genetic algorithm on the model of Besnard et al. (2011) and find savings of around 15% versus unoptimised scheduling. Dawid et al. (2017) takes the approach of several authors and clusters groups of maintenance activities together with compatible vessels to form a list of feasible maintenance plans. They find this cluster-matching approach compares favourably with commercial solvers, particularly for instances containing less than 15-20 tasks. Dawid et al. (2018) extend this to include probabilities of successfully completing tasks.

Task prioritisation is included within Stock-Williams and Swamy (2018) who pro-

duce high quality crew transfer plans through a metaheuristic procedure. The size and scope of technicians assigned to vessel routes is explored through a genetic algorithm procedure that allows for tasks to be scheduled outside of the current day. The fitness of each crew transfer plan is evaluated by a detailed simulation that accounts for both short and long term costs. The inability to complete all the tasks within a shift may be caused by a lack of suitable weather windows or a scarcity of resources. The benefits of sharing limited resources such as skilled technicians across different wind farms are illustrated by Schrotenboer et al. (2018). They develop an adaptive large neighbourhood search (ALNS) heuristic that calculates assignments of technicians to bases and the vessel routes that transport them to turbines on a daily basis. Savings of up to 7% are found from sharing technicians.

Exact models all assume that technicians should be assigned to the same vessel throughout the entire time they remain offshore. We choose to let technicians become disassociated from their original vessel, as swipe on and off access systems permit schedules where technicians may be dropped off and picked up by different vessels. Furthermore, some of the approaches do not consider maintenance tasks that are performed over multiple shifts. This will be necessary for complicated replacement tasks which take several days to complete and motivates the explicit separation of work and technicians in our model. This greater level of detail allows us to model periods of inactivity on the turbines even when technicians are present, which could be theorised to occur when weather conditions prevent certain operations or if turbines are restarted whilst technicians are onboard. Schedules can also be improved from the opportunity to start a task near the end of a shift and complete it at the start

of the subsequent shift. This can be observed from circumstances that involve shifts containing multiple weather windows which vary spatially over an entire wind farm.

Previous works also penalise incomplete tasks after a given deadline on a daily basis in proportion to the lost energy production. This allows for task selection and prioritisation between shifts, but has little effect on the ordering of tasks within a shift. We introduce a variant where tasks do not have completion deadlines, but instead accumulate losses whilst they remain incomplete. These losses might be the physical downtime losses from the performance of turbines or represent some notion of priority relative to the other tasks. We choose to minimise the total value of lost revenue from all of the turbines. This will incentivise the model to fully complete tasks as soon as possible within the horizon whilst considering the impact on the completion of other tasks. This forms a time sensitive framework where vessels will seek to visit and deposit technicians as early as possible at turbines.

A side effect of effectively minimising the time to complete tasks is the presence of slack time at the end of shifts. In practice tasks may require more time than expected, so this slack time can potentially ensure that all the planned activities are still finished before the end of the shift. This approach is limited by the fact that solutions cannot be differentiated after the last task has been completed, since the optimal route for picking up technicians is the same regardless of the time it occurs. We remove this source of symmetry by explicitly incentivising technicians and vessels to return to port as soon as possible. This has the benefit of extending the time sensitive framework to the end of the shift and incorporating a degree of safety into the model by minimising technician time offshore.

We choose to model the problem using a discrete time model for several reasons. Firstly, it allows us to incorporate a novel vessel-technician safety constraint often used in offshore operations. Existing models make no reference to the existence of any safety regulations concerning the presence of technicians and vessels offshore when executing the tasks. The formulation that we introduce later in this paper is flexible enough to integrate a variety of spatial and temporal constraints on crew transfers and vessel movements. Secondly, time discretisation allows us to model the presence of weather windows within shifts in detail. Previous methods neglect the possiblities of multiple non-contiguous weather windows within a single shift, however we can easily model these cases by forbidding vessel or task related activites in specific time periods. Finally, our resulting model allows for the possibility of task preemption so that work can be performed in stages rather than in a single event. This is of particular value given our goal of including weather operating restrictions on tasks.

2.3 Methodology

We introduce a mathematical model for the routing of vessels and technicians around offshore wind farms alongside the scheduling of task completion as a mixed integer linear program in this section.

2.3.1 Problem Description

The problem consists of a set I_W wind turbines and I_p maintenance ports. A set V of heterogeneous vessels are available to transport an offshore maintenance workforce

comprised of P^{Tot} technicians to the offshore tasks. A discretised travel time of T_{vij} periods exists when a vessel v travels from location i to location j. Each vessel v can hold up to P_v^{Max} technicians, travels at a fixed speed v_s and incurs a cost K_v^{Min} per period of travel. A crew transfer can only occur if a vessel v with sufficient technicians or free capacity is present at a location i for a prespecified number of periods τ_{vi}. Crew transfers either drop-off or pick-up technicians, with the former required to occur first if an offshore turbine is maintained. We make no direct connection between technicians and vessels so that technicians can be picked up and dropped off by different vessels.

Maintenance tasks exist on a subset of wind turbines and are unique to that location. Tasks and turbines have a one-to-one relationship: each turbine is limited to a single task and each task is assigned to a single turbine. Tasks are split into two sets: I_W^c corrective and I_W^p preventive maintenance tasks. Corrective tasks reduce the amount of electricity generated by their turbine during every period the task remains incomplete. The degraded output is typically the entire capacity of the wind turbine if it switched off from the start of the shift up until the completion of the task. Alternatively condition based maintenance could be modelled with a cost that varies over time in accordance with the estimated turbine condition. Preventive maintenance tasks only require turbines to be turned off whilst work is performed. Every minute worked on a preventive maintenance task $i \in I_W^p$ yields a monetary benefit of γ_i^{Min}, whilst a corrective maintenance task $i \in I_W^c$ incurs a monetary cost of θ_{is}^t if it is completed in period t of each shift s. This cumulative cost strictly increases in size for more distant time periods in the planning horizon. It is modelled as a

piecewise linear function, made up of the specific downtime losses in each period up to and including period t of shift s. We define Θ_{is}^{t} as the monetary value of downtime losses in period t of shift s if corrective maintenance task i is incomplete, so that $\theta_{is}^{t} = \Theta_{is}^{t} + \theta_{is}^{t-1}$, with $\Theta_{is}^{t} > 0$. This monetary value depends on the turbine, shift and time period involved. This allows it to be informed by the relative priorities of tasks and levels of energy production related to the wind speed in the specific period.

A maintenance task i is considered complete after P_i technicians have worked on it for D_i^{Min} minutes. Task preemption is allowed so that work can be completed in discontinuous stages instead of a single event. Furthermore we allow tasks to be potentially ignored in the planning horizon. This is beneficial when there is insufficient time or resources to complete every task. The planning horizon itself consists of a set of shifts S, forming a short term schedule. Each shift is further split into a set of time periods T_s, each beginning with t_a^s and ending with t_b^s, by which all technicians and vessels must have returned back to port. The length of each time period in shift s is λ_s^{Min} minutes. Technicians and vessels are incentivised to remain onshore in shift s at a rate of $\omega_{is}^{Min}, \Omega_{vis}^{Min}$ per minute respectively. Unnecessary personnel are therefore kept away from the wind farm reducing the likelihood of seasickness and its contagion within the transported technicians.

Vessel routes should be constructed so that each vessel starts and ends each shift in the port and has access to a shared pool of technicians. The vessels need to transport technicians to and from the offshore wind farm, whilst also moving them around to different turbines in order to complete more work. The number of technicians onboard a vessel is not allowed to exceed its technician capacity at any point. Technicians can

only be dropped off and picked up at a turbine once in each shift, however different vessels are allowed to perform these crew transfers. Furthermore, a vessel can drop-off technicians at a turbine and later pick them up again without ever having left the turbine. Once a drop-off has completed technicians can potentially perform work until a vessel starts to pick them up.

The vessel routes and timing of crew transfers should be chosen so that the most valuable combination of work on corrective and preventive maintenance tasks are performed. The cost for completing corrective tasks increases the later they are finished in the planning horizon, however the model may choose to intentionally not complete certain tasks if it leads to a better overall solution or if there are insufficient resources to complete them. Visits to preventive maintenance tasks should also be scheduled if they are found to be beneficial. The length of time that technicians and vessels are kept offshore should also be minimised.

We use the following decision variables in our mathematical model:

- $x_{vis}^t = 1$, if the vessel v is at location i at the beginning of time period t in shift s, after any relocations in or out of i. 0 otherwise.
- $r_{vijs}^t = 1$, if the vessel v starts a relocation from location i to location j before the start of time period t in shift s and completes it before the start of time period $t + T_{vij}$ in shift s. 0 otherwise.
- $s_{vis}^t = 1$, if a drop-off occurs at location i with vessel v before time period t in shift s. 0 otherwise.
- $\bar{s}_{vis}^t = 1$, if a pick-up occurs at location i with vessel v before time period t in shift s. 0 otherwise.
- l_{vis}^t, Number of technicians dropped-off by vessel v at port i before time period t and arrive the turbine before the start of time period $t + \tau_{iv}$ in shift s.
- $\bar{l}_{vis}^t$, Number of technicians picked-up by vessel v at port i before time period t and arrive onboard before the start of time period $t + \tau_{iv}$ in shift s.
- y_{vs}^t, Number of technicians on vessel v at the beginning of time period t in shift s, after any crew transfers have occurred.
- $p_{is}^t = 1$, if P_i technicians are present on turbine i at the beginning of time period t in shift s, after any crew transfers have occurred. 0 otherwise.
- $\tilde{p}_{is}^t$, Number of technicians at port i at the beginning of time period t in shift s, after any crew transfers have occurred.
- $W_{is}^t = 1$, if a period of work at location i is performed in time period t in shift s. 0 otherwise.
- $a_i = 1$, if the corrective task at location i is completed within the planning horizon. 0 otherwise.
- c_i, Total downtime cost accumulated from the corrective task at location i.

2.3.2 Mathematical Model

$$\min \overbrace{\sum_{i\in I_W^c} c_i}^{\text{corrective costs}} - \overbrace{\sum_s \sum_{t\in T_s} \sum_{i\in I_W^p} \lambda_s^{Min}\gamma_i^{Min} W_{is}^t}^{\text{preventive benefits}} + \overbrace{\sum_s \sum_{v,i,j} \lambda_s^{Min} K_v^{Min} T_{vij}^{Min} \sum_{t\in T_s} r_{vijs}^t}^{\text{travel costs}}$$
$$- \underbrace{\sum_s \sum_{t\in T_s} \sum_{i\in I_P} \lambda_s^{Min}\omega_{is}^{Min}\tilde{p}_{is}^t}_{\text{technician onshore incentive}} - \underbrace{\sum_s \sum_{t\in T_s} \sum_{i\in I_P} \sum_v \lambda_s^{Min}\Omega_{vis}^{Min} x_{vis}^t}_{\text{vessel onshore incentive}} \quad (2.3.1)$$

$$\text{s.t. } x_{vis}^{t+1} = x_{vis}^t + \sum_j r_{vjis}^{t+1-T_{vji}} - \sum_j r_{vijs}^{t+1}, \quad \forall v\in V, \forall i\in I, \forall s\in S, \forall t\in T_s^- \quad (2.3.2)$$

$$\sum_j r_{vijs}^{t+1} \le x_{vis}^t, \quad \forall v\in V, \forall i\in I, \forall s\in S, \forall t\in T_s^- \quad (2.3.3)$$

$$\sum_v (x_{vis}^t + \sum_j r_{vijs}^t) \le 1, \quad \forall i\in I_W, \forall s\in S, \forall t\in T_s \quad (2.3.4)$$

$$\sum_i (x_{vis}^t + \sum_j \sum_{w=0}^{T_{vji}-1} r_{vjis}^{t-w}) = 1, \quad \forall v\in V, \forall s\in S, \forall t\in T_s \quad (2.3.5)$$

$$s_{vis}^t \le x_{vis}^u, \quad \bar{s}_{vis}^t \le x_{vis}^u, \quad \forall v\in V, \forall i\in I, \forall s\in S, \forall t\in T_s, \forall u\in[t, t+\tau_{vi}-1] \quad (2.3.6)$$

$$l_{vis}^t \le \min(P^{Tot}, P_v^{Max}) s_{vis}^t, \quad \forall v\in V, \forall i\in I_P, \forall s\in S, \forall t\in T_s \quad (2.3.7)$$

$$\bar{l}_{vis}^t \le \min(P^{Tot}, P_v^{Max}) \bar{s}_{vis}^t, \quad \forall v\in V, \forall i\in I_P, \forall s\in S, \forall t\in T_s \quad (2.3.8)$$

$$\sum_v \sum_{t\in T_s} s_{vis}^t = \sum_v \sum_{t\in T_s} \bar{s}_{vis}^t \le 1, \quad \forall s\in S, \forall i\in I \quad (2.3.9)$$

$$\sum_v \sum_{t\in T_s} t\bar{s}_{vis}^t \ge \sum_v \sum_{t\in T_s} (t+1) s_{vis}^t, \quad \forall s\in S, \forall i\in I_W \quad (2.3.10)$$

$$y_{vs}^t \le P_v^{Max}, \quad \forall v\in V, \forall s\in S, \forall t\in T_s \quad (2.3.11)$$

$$y_{vs}^{t+1} = y_{vs}^t + \sum_{i\in I_P} \bar{l}_{vis}^{t+1-\tau_{vi}} - \sum_{i\in I_P} l_{vis}^{t+1} + \sum_{i\in I_W} P_i \bar{s}_{vis}^{t+1-\tau_{vi}} - \sum_{i\in I_W} P_i s_{vis}^{t+1}, \quad \forall v\in V, \forall s\in S, \forall t\in T_s^- \quad (2.3.12)$$

$$p_{is}^{t+1} = p_{is}^{t} + \sum_{v} s_{vis}^{t+1-\tau_{vi}} - \sum_{v} \bar{s}_{vis}^{t+1}, \qquad \forall i \in I_W, \forall s \in S, \forall t \in T_s^- \tag{2.3.13}$$

$$\tilde{p}_{is}^{t+1} = \tilde{p}_{is}^{t} + \sum_{v} l_{vis}^{t+1-\tau_{vi}} - \sum_{v} \bar{l}_{vis}^{t+1}, \qquad \forall i \in I_P, \forall s \in S, \forall t \in T_s^- \tag{2.3.14}$$

$$\begin{aligned} &\sum_{v} y_{vs}^{t} + P_i \sum_{i \in I_W} p_{is}^{t} + \sum_{i \in I_P} \tilde{p}_{is}^{t} \\ &\quad + \sum_{v \in V} \sum_{i \in I_W} \sum_{u=0}^{\tau_{vi}-1} P_i (s_{vis}^{t-u} + \bar{s}_{vis}^{t-u}) \\ &\quad + \sum_{v \in V} \sum_{i \in I_P} \sum_{u=0}^{\tau_{vi}-1} (l_{vis}^{t-u} + \bar{l}_{vis}^{t-u}) = P^{Tot}, \qquad \forall s \in S, \forall t \in T_s \end{aligned} \tag{2.3.15}$$

$$\sum_{(v,i) \in \phi_1^{kst}} x_{vis}^{t} + \sum_{(v,i,j,u) \in \phi_2^{kst}} r_{vijs}^{t-u} \geq p_{ks}^{t}, \qquad \forall k \in I_W, \forall t \in T_s, \forall s \in S \tag{2.3.16}$$

$$p_{is}^{t} \geq W_{is}^{t}, \qquad \forall i \in I_W, \forall s \in S, \forall t \in T_s \tag{2.3.17}$$

$$\sum_{s} \sum_{t \in T_s} \lambda_s^{Min} W_{is}^{t} \geq D_i^{Min} a_i, \qquad \forall i \in I_W^c \tag{2.3.18}$$

$$\sum_{s} \sum_{t \in T_s} \lambda_s^{Min} W_{is}^{t} \leq D_i^{Min}, \qquad \forall i \in I_W^p \tag{2.3.19}$$

$$c_i \geq \theta_{is}^{t} W_{is}^{t}, \qquad \forall i \in I_W^c, \forall s \in S, \forall t \in T_s \tag{2.3.20}$$

$$c_i \geq \theta_{is}^{t_b^s} (1 - a_i), \qquad \forall i \in I_W^c, \forall s \in S \tag{2.3.21}$$

$$\sum_{i \in I_P} (x_{vis}^{t_a^s} + \sum_{j} r_{vijs}^{t_a^s}) = 1, \qquad \forall v \in V, \forall s \in S \tag{2.3.22}$$

$$\sum_{i \in I_P} (x_{vis}^{t_b^s} + \sum_{j} r_{vjis}^{t_b^s + 1 - T_{vji}}) = 1, \qquad \forall v \in V, \forall s \in S \tag{2.3.23}$$

$$\sum_{i \in I_P} (\tilde{p}_{is}^{t_a^s} + \sum_{v} \bar{l}_{vis}^{t_a^s}) = P^{Tot}, \qquad \forall s \in S \tag{2.3.24}$$

$$\sum_{i \in I_P} (\tilde{p}_{is}^{t_b^s} + \sum_{v} l_{vis}^{t_b^s + 1 - \tau_{vi}}) = P^{Tot}, \qquad \forall s \in S \tag{2.3.25}$$

$$x_{vis}^{t} \in \{0, 1\}, \qquad \forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s \tag{2.3.26}$$

$$r^t_{vijs} \in \{0,1\}, \qquad \forall v \in V, \forall i \in I, \forall j \in I, \forall s \in S, \forall t \in T_s \tag{2.3.27}$$

$$s^t_{vis}, \bar{s}^t_{vis} \in \{0,1\}, \qquad \forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s \tag{2.3.28}$$

$$l^t_{vis}, \bar{l}^t_{vis} \in \mathbb{Z}, \qquad \forall v \in V, \forall i \in I_P, \forall s \in S, \forall t \in T_s \tag{2.3.29}$$

$$y^t_{vs} \in \mathbb{Z}, \qquad \forall v \in V, \forall s \in S, \forall t \in T_s \tag{2.3.30}$$

$$p^t_{is} \in \{0,1\}, \qquad \forall i \in I_W, \forall s \in S, \forall t \in T_s \tag{2.3.31}$$

$$\tilde{p}^t_{is} \in \mathbb{Z}, \qquad \forall i \in I_P, \forall s \in S, \forall t \in T_s \tag{2.3.32}$$

$$W^t_{is} \in \{0,1\}, \qquad \forall i \in I_W, \forall s \in S, \forall t \in T_s \tag{2.3.33}$$

$$a_i \in \{0,1\}, \qquad \forall i \in I^c_W \tag{2.3.34}$$

$$c_i \in \mathbb{R} \qquad \forall i \in I^c_W \tag{2.3.35}$$

The objective function (2.3.1) involves the minimisation of five components. The first is the accumulation of downtime losses from corrective O&M tasks over the entire horizon. The second represents the benefit gained by performing preventive maintenance tasks - it includes a negative sign since we are minimising the objective (minimising the additive inverse of a number is the same as maximising the original number). The third assigns vessel specific travel costs per minute travelled. The fourth and fifth components incentivise technicians and vessels to spend time in port respectively. Penalising this in the objective function allows us to account for solutions where work is finished early or when vessels return to port without any technicians onboard. We set the values of ω^{Min}_{is} and Ω^{Min}_{vis} to be very small in relation to other objective function coefficients, so it is always beneficial to complete additional tasks.

For each location $i \in I$, shift s and period $t \in T_s$, their combination (i, s, t) represents a time-space node. In addition to the time-space nodes, there are two different types of arcs in the network: dwelling arcs and relocation arcs. These correspond to the earlier outlined variables. The presence of a vessel v at a time space node is measured by the dwelling arc (v, i, s, t). These connect with relocation arcs that can be represented as a tuple (v, i, j, s, t), where (i, s, t) represents the tail time-space node and $(j, s, t - T_{vij})$ indicate the head time-space node. Both types of arcs do not cross the boundary between different shifts.

Constraints (2.3.2) are illustrated in Figure 2.3.1. Vessels are constrained to balance their arrivals via a relocation $\sum_j r_{vjis}^{t+1-T_{vji}}$, plus those present the period before x_{vis}^t with their departures. These are the vessels that remain at the location for the next period x_{vis}^{t+1} and those relocating away from the location $\sum_j r_{vijs}^{t+1}$. This is enforced for every time-space period (i, s, t) except for the first and final time periods in the shift, t_a^s and t_b^s. The boundary conditions ensure that vessels can only leave before the start of t_a^s and arrive at the end of t_b^s.

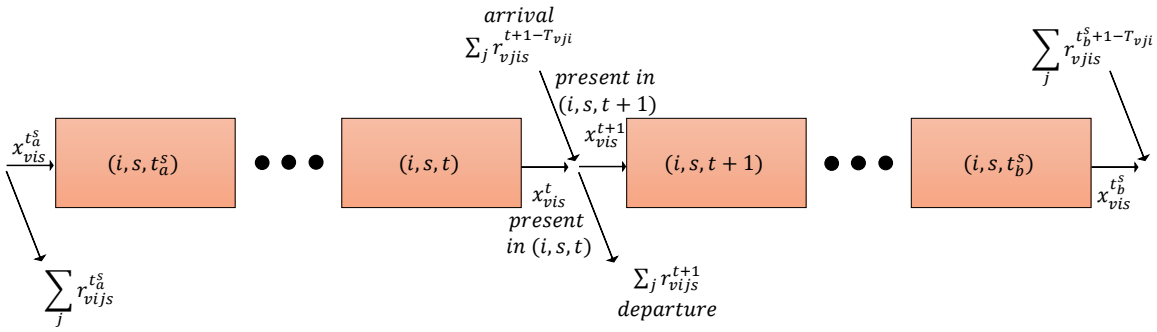

Figure 2.3.1: Individual vessel flow including boundary conditions.

Constraints (2.3.3) ensure that vessels cannot arrive to and then instantly depart from a location. Constraints (2.3.4) restrict each turbine to a maximum of one vessel at any time on the grounds of safety. Conservation of vessels is guaranteed by constraints (2.3.5) since each vessel is present at a location or involved in a relocation.

Constraints (2.3.6) require a vessel to be present for the entirety of any crew transfer. Constraints (2.3.7-2.3.8) decide the number of technicians involved if a crew transfer occurs in port. We assume that P_i technicians are transferred at turbines thus excluding partial transfers. In practice vessels do not perform multiple crew transfers at locations within a shift (constraints (2.3.9)), whilst constraints (2.3.10) force drop-offs to occur before pick-ups at turbines. Constraints (2.3.11) determine the technician capacity of each vessel. Technician flow on and off vessels is controlled by constraints (2.3.12). It states that technicians either arrive onboard by completing a crew transfer, depart by initiating a crew transfer or remain onboard the vessel. An equivalent form for the port and turbine viewpoints is given by constraints (2.3.13-2.3.14). Constraints (2.3.15) conserve the total number of technicians in the system at any time by tracking them at locations, on vessels and during crew transfers. Constraints (2.3.16) refer to our technician-vessel maximum safety range constraint, which we describe in detail in Section 2.3.3.

Technician idle time and working time is differentiated with constraints (2.3.17). This allows for technicians to be present on a turbine for a period, but not necessarily perform work in it. This gives flexibility to model the effects of weather on task completion. Specifically, we can model the effect of weather forcing technicians to stop working on turbines during periods of bad weather without having to leave it.

This can be achieved since $p_{is}^t = 1$ does not necessarily imply that $W_{is}^t = 1$, which we will make use of in Section 2.3.5. Constraints (2.3.18) provide for discontiguous work on corrective maintenance tasks across multiple shifts and is only considered complete after D_i^{Min} minutes of work across the planning horizon. An upper limit on the length of preventive maintenance is given by constraints (2.3.19).

Corrective cost accumulation is modelled through constraints (2.3.20-2.3.21) which we now explain in more detail. Analysing the constraints shows that we have to consider two cases which are related to whether the corrective task at turbine i is completed or not, i.e. if $a_i = 1$. If $a_i = 1$ then the right hand side of constraints (2.3.21) evaluates to zero, causing constraints (2.3.20) to become binding in the form of $c_i \geq \theta_{is}^t W_{is}^t$ for every time period t and shift s. We notice that the total downtime cost from accumulated from the corrective maintenance task at i is greater than or equal to the total losses incurred from completing the task in each period that it is worked on. Therefore as θ_{is}^t strictly increases for later time periods and shifts, according to its definition, the actual binding constraint amongst all these constraints will be $c_i \geq \theta_{is^*}^{t^*}$ where s^* and t^* are the last shift and last period the task is worked on and thus is completed.

Alternatively if $a_i = 0$, then constraints (2.3.21) will become binding in the form of $c_i \geq \theta_{is}^{t_b^s}$, since $\theta_{is}^{t_b^s}$ is the downtime cost from finishing the work in the last period of the last shift of the planning horizon. This is cannot actually occur if the work is completed, since the vessels and technicians are required to be back in port at this time, so it is the cost of not actually completing the task. Conceptually this cost is greater than every other θ value, so that we heavily penalise incomplete corrective

maintenance tasks, which follows naturally from the definition of $\theta_{is}^{t_b^s}$.

Finally as the objective function minimises c_i, the optimal solution will select the minimum value of c_i. This minimum value clearly occurs for the outlined binding constraints, when the left hand side equals the right hand side. Furthermore as shown later in the thesis, the downtime losses typically represent over 90% of the optimal value in practice. This means that even in slightly sub-optimal solutions we are likely to determine the correct downtime costs, since they are by far the most important quantity to minimise.

The use of W_{is}^t variables in constraints (2.3.20) implies that turbines with corrective tasks are switched on after the task is completed. An alternative assumption where they are only switched on after the technicians who completed the task are picked up, would require using p_{is}^t variables instead. We prefer the first option since it implicitly limits the amount of time technicians spent on turbines and has the advantage of structuring vessel pick-up routes. In practice the turbines may simply be switched on after the shift has concluded rendering this distinction moot. Constraints (2.3.22-2.3.25) state that vessels and technicians must start and end each shift in port. Model complexity can be further reduced by noticing that each turbine has an earliest arrival time (EAT) and latest departure time (LDT). The assumption that locations can only be visited for pick-up and drop-off once means that vessels and technicians can only be present at the wind farm between the EAT and LDT. We therefore fix associated variables outside of this window to zero during pre-processing.

2.3.3 Technician-Vessel Maximum Safety Range Constraint

Discussions with our industry partners have highlighted the need to include several real world technician safety constraints. One of these refers to the relative positions of vessels and technicians onboard turbines and is included in our model as constraints (2.3.16). We now describe our process of deriving said constraints.

Safety regulations stipulate that in an emergency any technician should be able to be picked up from their working location by a vessel within a set amount of time, α. This is to ensure that in the event of serious injury, technicians can always be quickly reached by a vessel, and if necessary transported to hospital in an expedited manner. A typical value of α is 30 minutes which is not a significant restriction for smaller wind farms, but can become more restrictive for larger wind farms that use slower vessels. We will also ignore any technician capacity considerations, since if the vessel is already fully loaded with technicians some can be dropped off to make room for the injured technicians.

$$\sum_{(v,i)\in\phi_1^{kst}} x_{vis}^t + \sum_{(v,i,j,u)\in\phi_2^{kst}} r_{vijs}^{t-u} \geq p_{ks}^t, \qquad \forall k \in I_W, \forall t \in T_s, \forall s \in S \tag{2.3.16}$$

Constraints 2.3.16 describe the constraint for ensuring at least one vessel is within the safety range of a turbine with technicians onboard and involves two sets ϕ_1^{kst} and ϕ_2^{kst}. ϕ_1^{kst} denotes the set of locations for vessels that are within α minutes of turbine k before time period t in shift s. ϕ_2^{kst} represents the set of vessel relocations in shift s that lie inside an α minute radius of turbine k in time period t. This includes relocations which intersect the radius and only temporarily lie within the safety coverage region. For convenience we convert the safety time radius α to a distance r by multiplying with

the vessel speed v_s, so that $r = \alpha v_s$. We also convert a location z in the mathematical model to its physical location in Cartesian co-ordinates, $\vec{z} = (z_x, z_y)$.

We first define $\vec{r}_{vij}$ to be the vector defined by the possible travel of a vessel v from location $\vec{i} = (i_x, i_y)$ to location $\vec{j} = (j_x, j_y)$. We consider $\vec{e}_{vij}$ as the vector of one period of travel of a vessel v from location $\vec{i}$ to location $\vec{j}$, so that $\vec{e}_{vij} = \vec{r}_{vij}/T_{vij}$. The start point of the wth period of travel within the relocation is written as $\vec{i} + w\vec{e}_{vij}$, where $w \in [0, T_{vij} - 1]$. In order for a period of a relocation to provide valid coverage within the safety range, both the start and end point of that period must lie within the safety radius. This radius r emanates from a circle centered at turbine $\vec{k} = (k_x, k_y)$.

We can determine whether a point $\vec{a}$ lies within the safety radius of a turbine k by checking whether the Euclidean distance from location $\vec{a}$ to location $\vec{k}$ is smaller than r. The Euclidean distance between two locations $\vec{a} = (a_x, a_y)$ and $\vec{k} = (k_x, k_y)$ is defined as $dist(\vec{a}, \vec{k}) = \sqrt{(a_x - k_x)^2 + (a_y - k_y)^2}$. We can therefore determine if a relocation that starts at period $t - w$ provides valid coverage for turbine k at period t by checking that the conditions $\text{dist}(\vec{i} + w\vec{e}_{vij}, \vec{k}) \leq r$ and $\text{dist}(\vec{i} + (w+1)\vec{e}_{vij}, \vec{k}) \leq r$ are true. This utilises the fact that if a vessel v lies within the interval $[o_1, o_1 + 1]$ at time period t, it must have started its relocation from i to j at time period $t - o_1$.

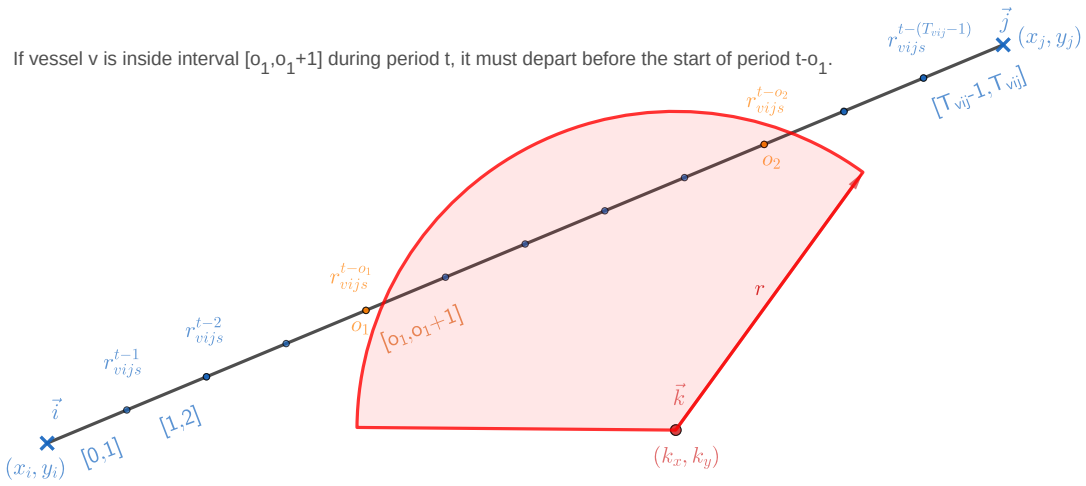

Figure 2.3.2: The vessel v travelling from i to j can only provide safety coverage for turbine k in period t of shift s, if it departs i before period $t-5, t-6, t-7,$ or $t-8$.

Figure 2.3.2 illustrates an example vessel relocation which partially covers the turbine k. Specifically the vessel can only provide coverage if it departs at a specific set of time periods. Other cases are possible where the entire relocation is either completely exterior or interior to the safety region. We also need to check whether a vessel is dwelling at a location i within the safety radius r in period t, which requires that $\text{dist}(\vec{i}, \vec{k}) \leq r$.

Using our outlined notation, the sets $\phi_1^{kst}, \phi_2^{kst}$ can be described as,

$$\phi_1^{kst} = \{(v,i) : \text{dist}(\vec{i},\vec{k}) \leq r\}, \tag{2.3.36}$$

$$\phi_2^{kst} = \left\{(v,i,j,w) : \left[\text{dist}(\vec{i} + w\vec{e}_{vij},\vec{k}) \leq r\right] \,\&\, \left[\text{dist}(\vec{i} + (w+1)\vec{e}_{vij},\vec{k}) \leq r\right] \,\&\, [w \in [0, T_{vij} - 1]]\right\}. \tag{2.3.37}$$

For simplicity we have assumed that a relocation time period must be completely enclosed within the safety region in order to provide valid coverage. This excludes points o_1, o_2 in Figure 2.3.2. An alternative approach would require a minimum fraction of a period to lie within the safety region to be designated as providing coverage. A fraction of 0.8 would include o_1, but not o_2 in Figure 2.3.2. A potential extension could allow deviations to straight vessel paths to form curved arcs to increase the vessel time spent within the safety radius. Further optimisation should then be performed if several turbines are covered by a vessel relocation.

2.3.4 Model Simplification - Transfer Time Approximation

One method of reducing the number of variables in our model lies in using a larger time period discretisation to reduce the overall number of time periods. We propose that combining vessel travel times and crew transfers times together can allow us to increase the period length and reduce the size of the model. The combined travel and transfer activity is typically the event with shortest duration in the model allowing us to increase the period length to around 20-30 minutes.

Existing models tend to append transfer times into vessel travel times as drop-off and pick-up turbines are modelled as separate locations, (Christiansen et al., 2013).

In our model this would relate to setting the transfer times to zero and excluding them, whilst the new vessel travel times are increased to include the old transfer times, $T'_{vij} = T_{vij} + \tau_{vi}$. Our model requires an additional constraint to concatenate these activities since the outlined transformation can omit the time needed to pick-up technicians if a vessel remains present whilst the task is worked on. Crew transfer times are required to be symmetric with respect to pick-up and drop-off to make the approach valid. This is illustrated in the first line of Figure 2.3.3. Our proposed method is shown in the second line, where we extend the time technicians spend working on the turbine by τ_{vi} periods.

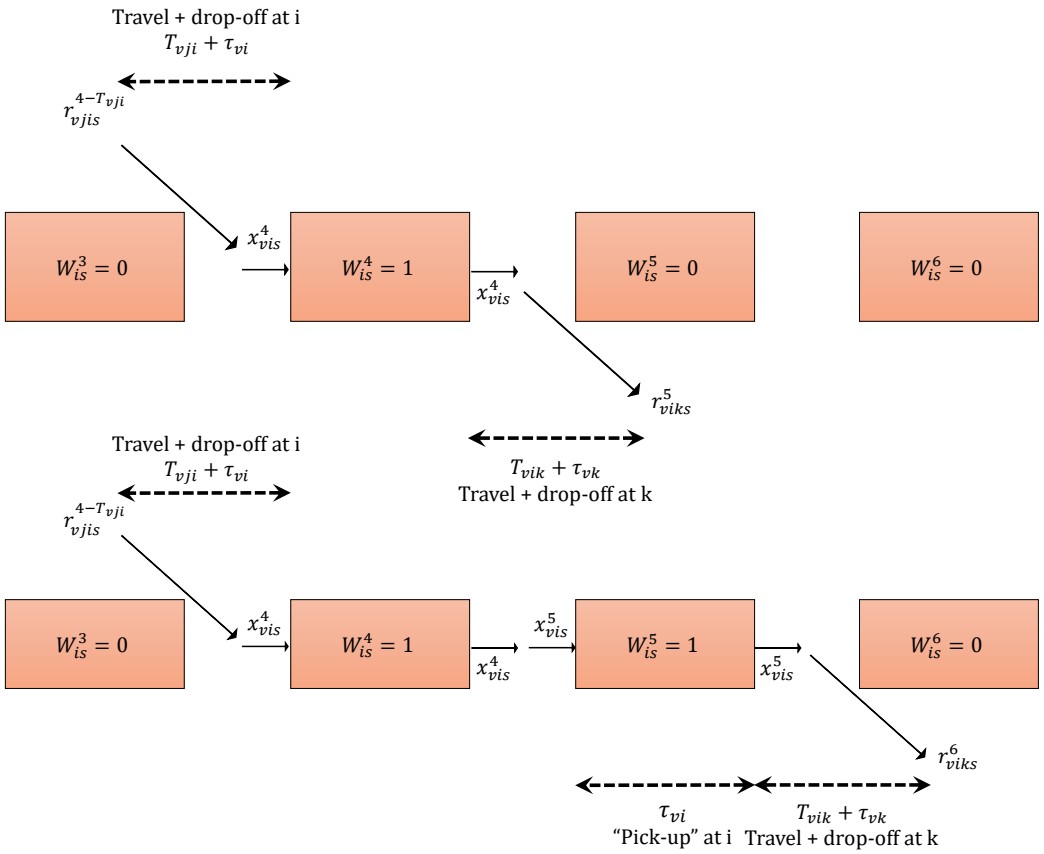

Figure 2.3.3: Above: Pick-up time at location i is omitted when a vessel performs a single visit if the standard transformation is used. Below: The missing τ_{vi} periods of pick-up are modelled as dummy periods of work with the vessel waiting at the turbine.

We represent the missing pick-up time as additional dummy work time since we lack a threshold for the time spent on a turbine. This modification is underpinned by the assumption that it is sub-optimal to leave technicians on a turbine after the work is completed if the vessel is present. We correctly adjust the timings by first differentiating between tasks that have both crew transfers performed within a single vessel visit and those performed across two vessel visits. If there is only a single departure from a turbine in a shift, then both pick-up and drop-off must be performed

in the same visit. If there are two departures from a turbine, pick-up and drop-off are performed across separate visits. We introduce new binary variables to track this: $\chi_{is} = 1$ if turbine i is visited exactly once in shift s and $\psi_{is} = 1$ if turbine i is visited exactly twice in shift s.

$$\chi_{is} + 2\psi_{is} = \sum_{v,j,t \in T_s} r^t_{vijs}, \qquad \forall i \in I_W, \forall s \in S \tag{2.3.38}$$

$$\chi_{is} + \psi_{is} \leq 1, \qquad \forall i \in I_W, \forall s \in S \tag{2.3.39}$$

$$\chi_{is}, \psi_{is} \in \{0, 1\} \qquad \forall i \in I_W, \forall s \in S \tag{2.3.40}$$

We introduce w_{is} to track the minutes spent working on turbine i in shift s to ensure dummy work is performed in its correct shift. The calculation is performed in minutes to potentially fit the pick-up transfer into any excess work time arising from the conversion to discrete time, thus avoiding the need for dummy work. The assumption that θ^t_{is} is a piecewise linear function $\theta^t_{is} = \Theta^t_{is} + \theta^{t-1}_{is}$, where Θ^t_{is} is the losses in period t of shift s for turbine i allows us to adjust downtime costs for dummy work.

$$\sum_{t \in T_s} \lambda^{Min}_s W^t_{is} - \tau^{Min}_{vi} \chi_{is} = w_{is} \qquad \forall i \in I_W, \forall s \in S \tag{2.3.41}$$

$$\sum_s w_{is} \geq D^{Min}_i a_i, \qquad \forall i \in I^c_W \tag{2.3.42}$$

$$\sum_s w_{is} \leq D^{Min}_i \qquad \forall i \in I^p_W \tag{2.3.43}$$

$$c_i \geq \theta^t_{is} W^t_{is} - \tau^{Min}_{vi} \Theta^{Min}_{is} \chi_{is} \qquad \forall i \in I^c_W, \forall s \in S, \forall t \in T_s \tag{2.3.44}$$

$$c_i \geq \theta^{t^b_s}_{is}(1 - a_i), \qquad \forall i \in I^c_W, \forall s \in S \tag{2.3.45}$$

$$w_{is} \in \mathbb{R} \qquad \forall i \in I_W, \forall s \in S \tag{2.3.46}$$

Experiments have shown that the model will try to actively avoid solutions which contain dummy work by scheduling extra visits to the tasks to reduce their cost. The additional visits occur before or after task completions and do not trigger crew transfers. We forbid these artificial solutions by introducing two further constraints. Constraints (2.3.47) state that if a vessel arrives to a location, the location must either already have technicians onboard it or a drop-off must occur when it arrives. Constraints (2.3.48) are the symmetric counterpart of constraints (2.3.47), stating that if a vessel leaves a location technicians should remain on the turbine after it leaves or the vessel must execute a pick-up on it's departure.

$$\sum_j r_{vjis}^{t+1-T_{vji}-\tau_{vi}} \leq p_{is}^t + s_{vis}^{t+1}, \qquad \forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s^- \tag{2.3.47}$$

$$\sum_j r_{vijs}^t \leq p_{is}^t + \bar{s}_{vis}^t, \qquad \forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s \tag{2.3.48}$$

We now restate the mathematical model with the transfer time approximation and technician-vessel maximum safety range constraint for completeness. Any relocation variables involving T_{vij} are rescaled to $T_{vij} + \tau_{vj}$, whilst τ_{vi} is removed from any crew transfer terms. Constraints (2.3.6) are omitted since the transfer time approximation removes the need to explicitly force vessels to wait at locations during crew transfers.

$$\min \sum_{i \in I_W^c} c_i - \sum_s \sum_{i \in I_W^p} \gamma_i^{Min} w_{is} + \sum_s \lambda_s^{Min} \left(\sum_{v,i,j} K_v^{Min} T_{vij}^{Min} \sum_{t \in T_s} r_{vijs}^t - \sum_{t \in T_s} \sum_{i \in I_P} \left(\omega_{is}^{Min} p_{is}^t + \sum_v \Omega_{vs}^{Min} x_{vis}^t \right) \right)$$

s.t. (2.3.2), (2.3.4) – (2.3.5), (2.3.7) – (2.3.11), (2.3.12) – (2.3.15), (2.3.16), (2.3.47) – (2.3.48), Flow constraints

$s_{vis}^t \leq \sum_j r_{vjis}^{t-T_{vji}-\tau_{vi}}, \quad \bar{s}_{vis}^t \leq \sum_j r_{vijs}^t$ $\forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s$

(2.3.38) – (2.3.39), (2.3.17), (2.3.41) – (2.3.45), (2.3.21) Task execution constraints

(2.3.22) – (2.3.25) Return to base constraints

(2.3.26) – (2.3.35), (2.3.46), (2.3.40) Variable definitions

(2.3.49)

2.3.5 Metocean Restrictions On Offshore Activities

The limitations imposed by meteorological and oceanographic (metocean) conditions differentiate the execution of offshore O&M activities from their onshore counterparts. The modelling choice of a discrete time space network provides the flexibility to incorporate a variety of restrictions on both turbine accessibility and task execution at any location and time.

The metocean conditions are created by introducing a wind speed and wave height for every time period into our model. A single wind speed across the entire windfarm is allowed to vary throughout the planning horizon. This is incorporated into the objective function costs so that greater downtime losses are accrued by corrective tasks during periods of higher wind speed. Conversely preventive maintenance tasks receive a greater benefit when performed in periods of low wind speed since the lost energy production is smaller. If the wind speed exceeds a threshold $Wind^{Lim} = 10m/s$, work cannot be performed although technicians can remain idle on the turbine. Therefore solutions in which work is paused and restarted within shifts may become optimal.

Significant wave height controls whether crew transfers or vessel relocations can be performed. A crew transfer can only be performed if the significant wave height stays below a threshold $Wave^T$ for the entire transfer, whereas a vessel relocation requires it to be below $Wave^R$ at both the start and end of the relocation. For simplicitly we make an assumption that the wave heights will either slowly increase or decrease throughout a shift at a constant rate. Specifically we assume that if the wave height is below $Wave^R$ at both the start and end of the relocation, it must

have stayed below the threshold at all times during the relocation. This means we do not need to check the wave height in the middle of a relocation. Crew transfer vessels can safely operate in significant wave heights up to $1.5m$ which in conjunction with our previous assumption that every relocation will result in a crew transfer leads us to set $Wave^R = Wave^T = 1.5m$. These thresholds allow us to model multiple non-contiguous weather windows within a shift. More complicated weather windows across parts of the wind farm could be modelled if weather data exists for each wind turbine.

The wave height at turbine i in period t of shift s, $Wave_{ist}$ and the wind speed in period t of shift s, $Wind_{st}$ are known parameters that restrict variables in our mathematical model as follows,

$$Wave_{ist} r^t_{vijs} \leq Wave^T \qquad \forall v \in V, \forall i \in I, \forall j \in I, \forall s \in S, \forall t \in T_s \qquad (2.3.50)$$

$$Wave_{jst+T_{vij}} r^t_{vijs} \leq Wave^T \qquad \forall v \in V, \forall i \in I, \forall j \in I, \forall s \in S, \forall t \in T_s \qquad (2.3.51)$$

$$Wave_{ist} s^t_{vis} \leq Wave^T, \qquad \forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s \qquad (2.3.52)$$

$$Wave_{ist} \bar{s}^t_{vis} \leq Wave^T, \qquad \forall v \in V, \forall i \in I, \forall s \in S, \forall t \in T_s \qquad (2.3.53)$$

$$Wind_{st} W^t_{is} \leq Wind^{Lim}, \qquad \forall i \in I_W, \forall s \in S, \forall t \in T_s. \qquad (2.3.54)$$

We choose to allow vessels to dwell at offshore turbines at any point in time regardless of the weather conditions, so constraints (2.3.50)-(2.3.54) do not include any x variables. Constraints (2.3.50) and (2.3.51) check that the wave weights for a potential relocation from time-space node (i, s, t) to $(j, s, t + T_{vij})$ are below $Wave^T$. Constraints (2.3.52) and (2.3.53) only allow crew transfers to occur when the signifi-

cant wave height is below the threshold. Constraints (2.3.54) forbid work from being performed in a period of high wind speed.

2.3.6 Stochastic Extension For Uncertain Weather Conditions

As the metocean conditions, $Wave_{ist}, Wind_{st}$, cannot be predicted with a 100 percent confidence, there is a degree of uncertainty associated with them. It is important to capture this uncertainty, particularly if it could cause an operational threshold $Wind^{Lim}, Wave^{R}, Wave^{T}$ to be exceeded. The optimal plan of action for vessel and personnel can vary significantly dependent on whether access to a turbine is allowed at a particular time point. This motivates the need for a stochastic extension to our model. The aim of this is to account for the uncertainty in whether vessels and personnel can access turbines, as defined through the metocean conditions.

A stochastic version of our model to account for uncertainty in offshore weather conditions can be developed in which the aim would be to minimise the overall expected cost across several shifts. In general this would form a multi-stage stochastic model using several different weather forecasts with different probabilities of occurrence. This would optimise the cost of the routing and scheduling in a first time period, plus the expected costs in the subsequent later shifts. The metocean conditions for the first shift are assumed to be known exactly, whereas the conditions in later shifts are uncertain. This idea is motivated by the fact that weather forecasts are normally very accurate in the present, but tend to decrease in accuracy when predicting futher into the future. In the context of our routing and scheduling problem we consider the metocean conditions in the first shift to be common to each scenario

as the latest weather forecast is assumed to be an accurate prediction of conditions. Predictions for subsequent shifts are subject to greater uncertainty and need to be modelled with distinct metocean conditions with fixed probabilities of occurring. The differing possibilities of metocean conditions can lead to operational thresholds being exceeded in different scenarios, all of which should be considered when building the plan for the first shift.

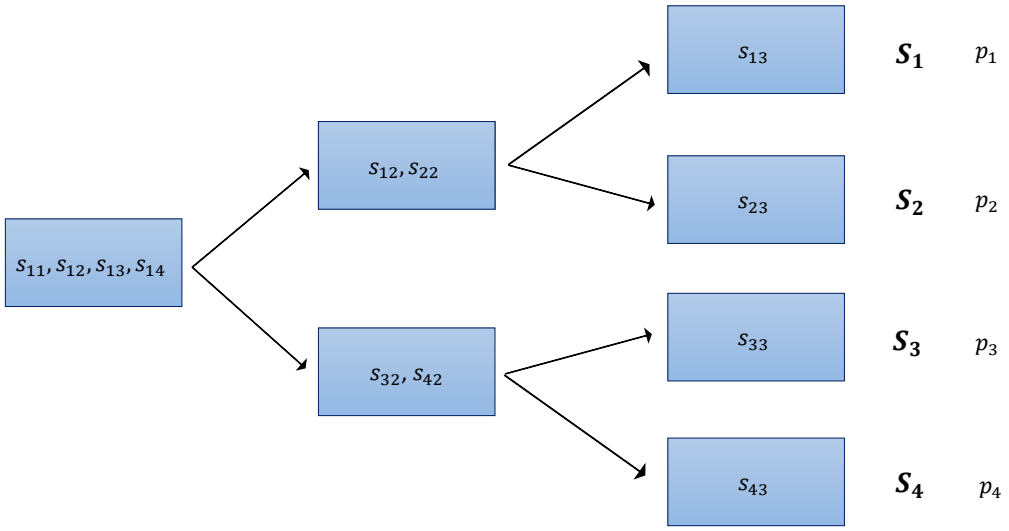

Figure 2.3.4: A scenario tree for a three stage stochastic optimisation. s_{xy} represents the metocean conditions in the yth day of scenario x. There are three days in the planning horizon, with the second and third each having two distinct possibilities for metocean conditions. Several scenarios have common metocean conditions.

An example of a three stage stochastic optimisation is shown in Figure 2.3.4. There are 4 different scenarios, S_1, S_2, S_3 and S_4, each of which share the same metocean conditions in the first shift, but contain different conditions in the second and third shifts. Instances using several vessels and shifts with size comparable to those

described in Section 2.4.1 were found to be too large to be solved exactly using a standard mixed integer programming solver. We therefore restrict our focus to a smaller two stage model.

We choose to focus on a two stage model with two scenarios for the metocean conditions in the second stage. The first reason for this is to ensure that we can solve problems of a reasonable size. The second reason is that we wish to focus on the case where metocean conditions are either slightly above or below operational thresholds at different locations. This is the most interesting example where a small change in offshore conditions could have a significant impact on the optimal solution. Otherwise the effect of the weather is negligible, as the wind farm is either entirely accessible or inaccessible most of the time.

If a larger multi-stage problem does need to be solved, one method could be to use a rolling horizon approach. A rolling horizon approach aims to solve smaller sub-problems that incorporates less information from future shifts, in a sequence that moves through the shifts. This will then provide a good approximation to the full multi-stage problem. The first step is to optimise the first shifts activities in detail whilst taking into account its impact on the work planned for the second shift, rather than all the successive shifts. The plan created for the first shift is then implemented and the work available for the next shifts is updated based on the solution of the first shift. In order to solve for the next shift we roll the time horizon forward, and in the case of later shifts update the metocean conditions with their newly revealed information. This is repeated to form an iterative process that covers the entire planning horizon.

The stochastic programming model we now present is an enhanced version of our previously stated deterministic model. For this reason, constraints that remain unchanged from its deterministic counterpart will be ignored to avoid unnecessary repetitions. The proposed modifications will now be described.

We can convert our existing model to a deterministic equivalent of a two stage stochastic model using two scenarios, $K = \{1, 2\}$ from a set of shifts $S = \{1, 2A, 2B\}$. Scenario 1 can be represented as a set of shifts $S_1 = \{1, 2A\}$, whilst scenario 2 is represented as $S_2 = \{1, 2B\}$. This ensures that both scenarios have a common first shift, but a different second shift. A probability p_k is assigned to each scenario under the conditions $\sum_{k \in K} p_k = 1$. Then we can minimize the expected value of the objective, subject to the constraints from both scenarios.

We first need to extend the definition of the variables associated with work on a corrective maintenance. They should now operate on the basis of scenarios, rather than all the possible shifts.

$a_i^k = 1$, if the corrective task at location i is completed within the planning horizon in scenario k. 0 otherwise. (2.3.55)

c_i^k, Total downtime cost accumulated from the corrective task at location i in scenario k. (2.3.56)

The objective of our mathematical model can now be rewritten to minimise the expected costs of scenario 1 and scenario 2, by including the probability of that scenario occurring. Constraints (2.3.42-2.3.45) are also altered to account for multiple scenar-

ios as shown below,

$$\min p_1 \left(\sum_{i \in I_W^c} c_i^1 - \sum_{s \in S_1} \sum_{i \in I_W^p} \gamma_i^{Min} w_{is} + \sum_{s \in S_1} \lambda_s^{Min} \left(\sum_{v,i,j} K_v^{Min} T_{vij}^{Min} \sum_{t \in T_s} r_{vijs}^t - \sum_{t \in T_s} \sum_{i \in I_P} \left(\omega_{is}^{Min} p_{is}^t + \sum_v \Omega_{vs}^{Min} x_{vis}^t \right) \right) \right) \tag{2.3.57}$$

$$+ p_2 \left(\sum_{i \in I_W^c} c_i^2 - \sum_{s \in S_2} \sum_{i \in I_W^p} \gamma_i^{Min} w_{is} + \sum_{s \in S_2} \lambda_s^{Min} \left(\sum_{v,i,j} K_v^{Min} T_{vij}^{Min} \sum_{t \in T_s} r_{vijs}^t - \sum_{t \in T_s} \sum_{i \in I_P} \left(\omega_{is}^{Min} p_{is}^t + \sum_v \Omega_{vs}^{Min} x_{vis}^t \right) \right) \right) \tag{2.3.58}$$

$$\sum_{s \in S_k} w_{is} \geq D_i^{Min} a_i^k, \qquad \forall i \in I_W^c, \forall k \in K \tag{2.3.59}$$

$$\sum_{s \in S_k} w_{is} \leq D_i^{Min} \qquad \forall i \in I_W^p, \forall k \in K \tag{2.3.60}$$

$$c_i^k \geq \theta_{is}^t W_{is}^t - \tau_{vi}^{Min} C_{is}^{Min} \chi_{is}, \qquad \forall i \in I_W^c, \forall k \in K, \forall s \in S_k, \forall t \in T_s \tag{2.3.61}$$

$$c_i^k \geq \theta_{is}^{t_s^b} (1 - a_i^k), \qquad \forall i \in I_W^c, \forall k \in K, \forall s \in S_k \tag{2.3.62}$$

2.4 Experimental Results

In this section we outline the results of three experiments which were designed to illustrate the breadth of our model. Our main experiment solves our deterministic model on a modern wind farm layout with the goal of finding the maximum instance size we can solve within a reasonable computation time. We also evaluate the impact of two different safety ranges on the same set of instances, so that the cost of implementing a safety range can be quantified. The final experiment is designed to assess the value of a stochastic solution for uncertain metocean conditions.

The geographic layout of an offshore wind farm can have a significant impact on the design of routes for conducting O&M activities. Initial offshore wind farms arranged

turbines in a traditional square lattices with a typical inter-turbine separation of around 1km. Existing analysis tends to be performed on wind farms of this shape (Dawid et al., 2017), or commonly seen variations such as trapezoidal layouts. Newly emerging wind farms optimise turbine placement based on predicted wind speeds and foundation costs (Fischetti and Pisinger, 2018), which leads to irregular spatial structures such as that found in Figure 2.4.1. This structure has been provided by our industrial partners as a realistic example of a future offshore wind farm, which we use as the basis for our experiments. Some turbines have been deliberately omitted and distances rescaled to ensure anonymity. The port is located 45 km south east of the centre of the wind farm.

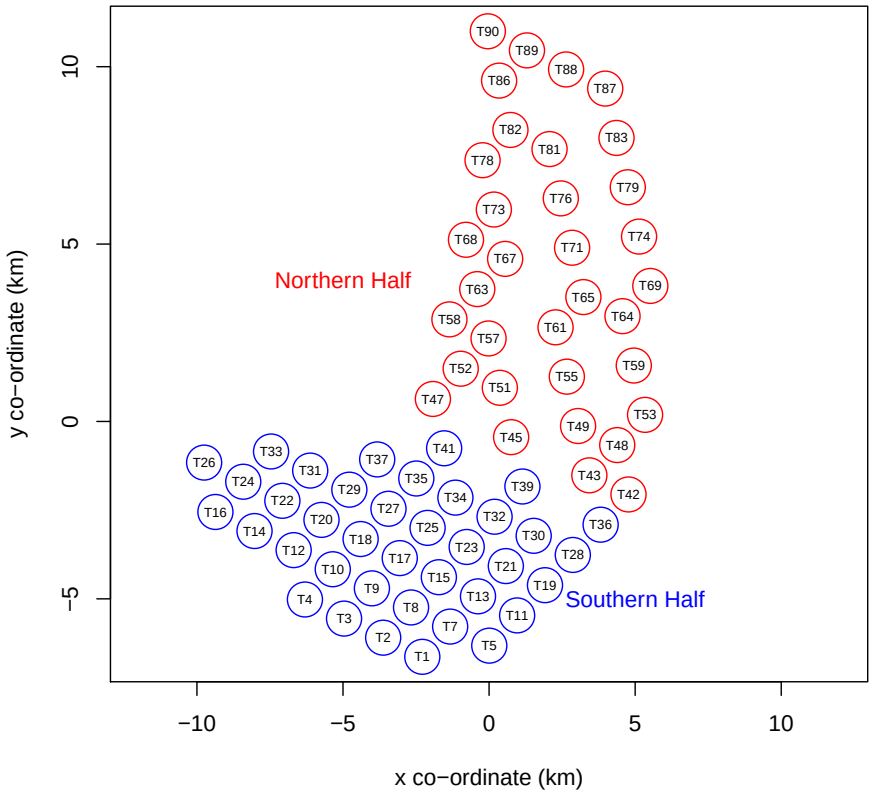

Figure 2.4.1: Example of a future offshore windfarm layout.

Task lists were adapted from Dawid et al. (2018) and include fixed ratios of corrective and preventive maintenance tasks. Preventive maintenance requires 2 technicians and provides the smallest benefit of 20 per minute worked, but possess durations in excess of 8 hours. This allows opportunistic maintenance to be performed. Corrective maintenance tasks are split into sub-categories based on their priority. The most severe tasks accrues losses of 80 per minute whilst they are incomplete and require 4 technicians. Typical corrective tasks require 2-4 hours of work by 2-3 technicians to be considered complete. Task locations were randomly assigned to existing turbines

within the wind farm. A maximum of three vessels are available, the first two being smaller crew transfer vessels with capacity for 12 technicians moving at a speed of 36km/h. The third vessel can carry 16 technicians but moves at a slower 30km/h. The technicians available in port totalled 10, 15 or 24. We follow the approach of Dawid et al. (2018) and reduce vessel speed within the wind farm by a third to account for turning and acceleration/deceleration. Wave and sea conditions are ignored in these scenarios, which consisted of a single shift of 11 hours split into 15 minute time intervals. Instances were solved using IBM CPLEX Optimizer using default settings on a Intel(R) Xeon(R) CPU E5-2640 v3 @ 2.60GHz with 32.0GB of RAM with a time limit of 2hrs.

Name	Type	Cost/min	P_i	D_i^{Min}	τ_{vi}^{Min}
Annual Service (Short)	Prev	20	2	480	20
Annual Service (Long)	Prev	20	2	600	20
Minor Reset	Corr	40	2	120	20
Minor Repair	Corr	50	3	240	20
Medium Repair	Corr	60	3	300	20
Major Repair	Corr	80	4	360	20

Table 2.4.1: Overview of task profiles.

2.4.1 Combining Preventive and Corrective Maintenance (Without A Safety Range)

$(\lvert I_W^c \rvert), \lvert I_W^p \rvert)$	P^{Tot}	$\lvert V \rvert = 1$			$\lvert V \rvert = 2$			$\lvert V \rvert = 3$		
		Time(s)	#Opt.	Gap(%)	Time(s)	#Opt.	Gap(%)	Time(s)	#Opt.	Gap(%)
(3,1)	10	1.12	10	0						
(3,1)	15	1.00	10	0						
(3,1)	24	1.01	10	0						
(5,1)	10	8.34	10	0	70.55	10	0			
(5,1)	15	10.49	10	0	37.60	10	0			
(5,1)	24	9.98	10	0	18.49	10	0			
(6,2)	10	33.40	10	0	353.9	10	0	1131	10	0
(6,2)	15	58.22	10	0	256.0	10	0	751.3	10	0
(6,2)	24	50.26	10	0	86.15	10	0	128.4	10	0
(8,2)	10				2271	10	0	3846	6	4.99
(8,2)	15				2555	10	0	5281	3	2.51
(8,2)	24				532.7	10	0	1961	9	0.6
(9,3)	10							N/A	0	3.35
(9,3)	15							7067	2	6.45
(9,3)	24							N/A	0	3.46
(3,1)	Avg	1.05	30	0						
(5,1)	Avg	9.06	30	0	42.21	30	0			
(6,2)	Avg	47.30	30	0	232.0	30	0	670.4	30	0
(8,2)	Avg				1786	30	0	3142	18	3.18
(9,3)	Avg							7067	2	4.28

Table 2.4.2: Results from preventive and corrective maintenance instances.

As previous studies have indicated that vessels tend to visit no more than 4 or 5 turbines in optimal solutions (Irawan et al., 2017), we investigate similar levels of tasks per vessel. Single vessel instances are comprised of 4, 6 and 8 tasks, 2 vessel instances contain 6, 8 and 10 tasks, whilst 3 vessel instances consist of 8, 10 and 12 tasks. Table 2.4.2 shows the effects of changing the number of tasks, vessels and technicians within this set-up using the parameters outlined earlier in this section. The first column shows the number of corrective and preventive maintenance tasks in the instance set. The second column gives the number of technicians available. For each instance set we created 10 different instances and report the average time to solve the instances, the number of instances (out of 10) solved to optimality within 2 hours and the average optimality gap.

The solution time scales dramatically with the number of tasks to the point that larger instances with 3 vessels and 12 tasks cannot be solved to optimality within 2 hours. The number of vessels available similarly affects the solution time as there are more ways of starting tasks earlier within the shift. The optimal routing and transfer plans for multiple vessel instances tends to follow a line drop-off and pick-up approach, as there is no need to reuse technicians across multiple tasks regardless of the number available. An example of line drop-off and pick-up for three tasks, A,B and C would be: drop-off at A, drop-off at B, drop-off at C followed later by pick-up at A, pick-up at B and pick-up at C.

Table 2.4.2 proves this is not always the case since the average solution time often decreases upon adding more technicians into the system. This is particularly visible for the 3 vessel, 8 task instances. The reason behind this is attributable to the change

in task completions. Optimal solutions to instances with fewer technicians make use of the ability to ignore certain tasks which adds additional complexity to the model, whereas the presence of large quantities of technicians often allow all the tasks to be completed. In the 3 vessel, 8 task instances, the majority of the 15 technician instances ignore at least 1 task whilst the 24 technician instances complete all the tasks. Further evidence comes from the 15 technician instances utilising all the available technicians at the wind farm, whereas the 24 technician instances often leave a few excess technicians behind in port. This highlights a class of tougher resource restricted instances which our model can solve without any prior specifications on the subset of tasks to perform. Discussions with our industrial partners emphasized the requirement of coping with variable numbers of technicians and vessels. In severe cases wind farms can accrue numerous tasks to the point where they cannot all be completed within a single shift even with maximum vessel and technician availability. Moreover these availability levels are often significantly lower producing resource restricted instances where there is an abundance of work to perform with few technicians and vessels.

2.4.2 Effect Of Adding A Vessel Safety Range

We quantify the financial cost of potential safety regulations by solving instances with and without a vessel safety radius. The inclusion of this as an additional constraint will cause an increase in the solution costs as certain schedules will be forbidden. Using a small safety radius can produce solutions where locations are visited purely for their geographic location without completing tasks. This provides circumstances where corrective tasks are only temporarily worked on. The average percentage increase in

the objective function for a safety radius of k minutes is measured as

$$\Delta\text{Obj}_k = 100 * \frac{1}{n} \sum_{i=1}^{n} \left(\frac{Z_i^k - Z_i^*}{Z_i^*} \right), \tag{2.4.1}$$

where Z_i^* is the optimal cost of task list i if no safety radius is imposed and Z_i^k the optimal cost if a safety radius of k minutes is imposed.

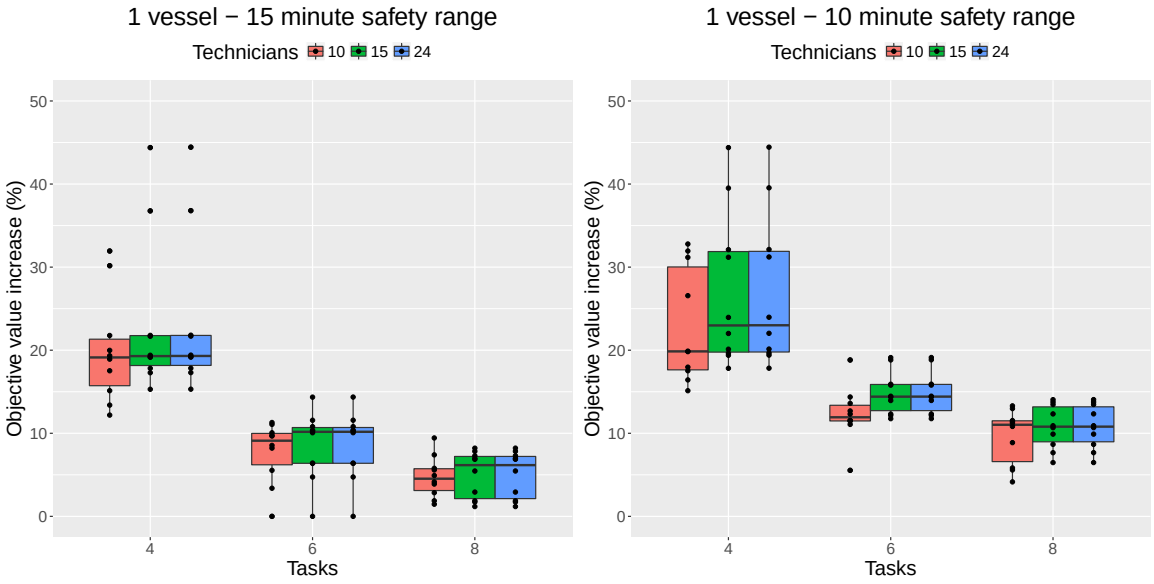

(a) Single vessel instances. Left: 15 minute safety range. Right: 10 minute safety range.

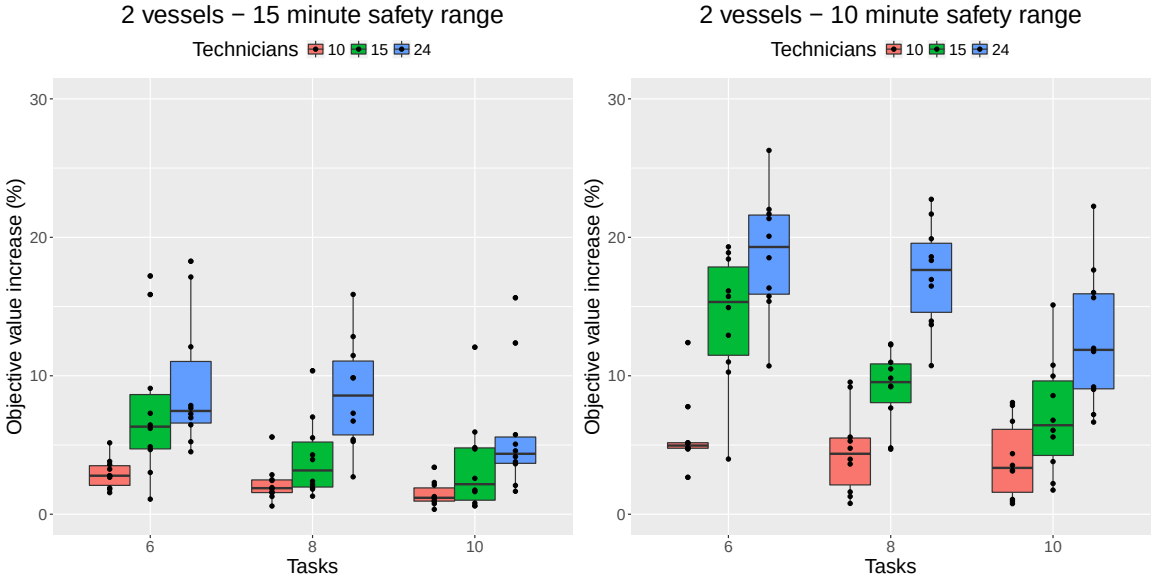

(b) Two vessel instances. Left: 15 minute safety range. Right: 10 minute safety range.

Figure 2.4.2: Quantifying the effect of a safety range.

Figure 2.4.2 shows the results of applying both a 15 and 10 minute safety radius on a subset of instances from Table 2.4.2 in box-plot form. A black dot represents an individual task list in each of the instance sets. Although the total cost is used, it is

heavily dominated by the corrective maintenance tasks. Any significant increases will therefore correspond to fewer tasks being completed or existing tasks being finished at a later time. The expected increase in objective function costs as the safety radius tightens occurs in all the presented instances. The average increase in the optimal cost of single vessel instances which contain 6 tasks is $\Delta\text{Obj}_{15} = 8.25\%$ and $\Delta\text{Obj}_{10} = 14.01\%$. This effect is less pronounced when 2 vessels are available as there is more flexibility to deal with the tightened safety restrictions resulting in increases of 6.62% and 12.90% when 15 and 10 minute safety radii are used. The computational time for solving instances that involve a safety radius is significantly larger than those without. The largest 2 vessel, 10 task instances contained some task lists that were unable to be solved to optimality within 2 hours, so calculations were performed with the best solution found to that point. A decrease in the safety radius appears to have a bigger effect on solutions containing larger quantities of technicians and tasks, however this is not always guaranteed.

Both of these aspects can be explained by the model's need to consider omitting tasks once all available resources have been fully utilised. At this stage newly introduced corrective tasks cannot be completed (or should not replace an existing task) and substantially increase the proportion of downtime costs within the total objective cost, thus reducing the percentage increase due to the safety radius. E.g. the downtime costs could increase from $60,000$ to $80,000$, but the extra costs due to the safety range h may remain stable. Then the percentage that h represents out of the total costs will naturally reduce, as the downtime costs are the main component of the objective function. Solutions to instances involving additional technicians are re-

stricted more heavily for the same safety radius since fewer transfer plans will remain feasible. This causes a growth in the percentage increase in the objective function due to the safety radius. This effect is capped by the ability to transfer technicians to the wind farm and illustrated in the single vessel instances containing 6 tasks. The inconsequential difference in the objective value from using 24 over 15 technicians can be attributed to the model receiving a reward for keeping the surplus technicians in port, since a maximum of 12 can be transported on the vessel. Significant reductions in downtime costs from utilising more technicians is only possible when there is both spare vessel capacity and sufficient uncompleted tasks to make use of them, as observed in the two vessel instances.

2.4.3 Uncertain Offshore Accessibiilty

In this experiment we consider a simple 2 stage stochastic optimisation with a look of a single shift to produce the scenario tree in Figure 2.4.3. We test our model in situations where task premption within and across shifts is necessary by considering periods of time when $Wind^{Lim}, Wave^{R}$ and $Wave^{T}$ are all exceeded. In the first shift every location in the wind farm is accessible by vessels at any time, however there is a significant period of time in the middle of the shift where $Wind^{Lim}$ is exceeded. Work cannot be performed on turbines during this interval meaning technicians will become idle. The weather conditions diverge in the second shift as the significant wave height at a turbine depends on its geographical location. In scenario 1 all the turbines in the northern half of the wind farm are inaccessible due to the wave heights, whilst in scenario 2 the southern half is inaccessible. This means that certain activities can only

be completed within the first shift for each scenario. We set $p_1 = 0.6$ and $p_2 = 0.4$ to represent a stronger belief in the first forecast. The downtime losses for corrective tasks and benefits for preventive maintenance are linked to the wind speed at each turbine for every period.

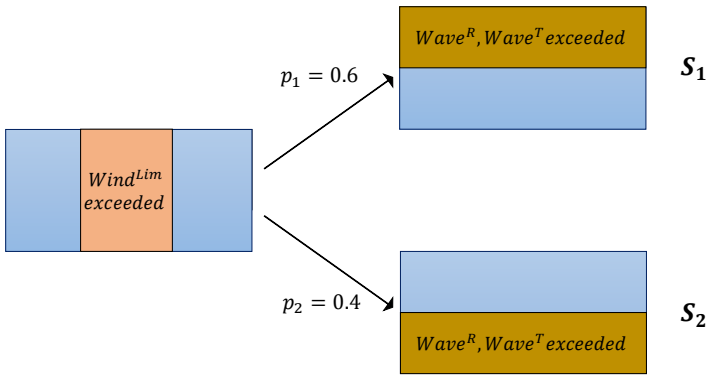

Figure 2.4.3: A simple scenario tree with metocean conditions.

A direct comparison between the vessel routes and transfer plans for the first shift is visualised in Figure 2.4.4 using the task list provided by Table 3.5.1. Directed arcs denote the direction of travel with the time periods spent travelling and the number of technicians onboard illustrated through arc labels. Tasks with longer durations were selected to find optimal solutions containing tasks that were performed discontiguously, whilst the number of tasks was balanced across both the northern and southern halves of the wind farm. A single vessel with capacity for 12 technicians was available.

The deterministic solution for scenario 1 shown in Figure 2.4.4a, exhibits the traditional planning technique of line drop-off and pick-up as the vessel visits the

Name	Place	Cost	Type	P_i	D_i^{Min}	τ_{vi}^{Min}
Medium Repair	T45	240	Corr	3	300	20
Minor Repair	T74	220	Corr	3	240	20
Medium Repair	T19	240	Corr	3	300	20
Annual Service	T8	100	Prev	2	480	20
Major Repair	T78	400	Corr	4	600	20
Minor Repair	T13	220	Corr	3	240	20
Medium Repair	T63	240	Corr	3	300	20
Annual Service	T21	100	Prev	2	480	20

Table 2.4.3: Task profiles for Figure 2.4.4.

tasks at T74, T19, T13 and T21. Work completion is paused on these tasks during the interval of high wind speed, after which the vessel starts picking up technicians from the turbines in the reverse order having waited at T21. The corrective tasks at T13 and T19 have 7 and 11 periods of works performed on them respectively. These tasks can be completed in second shift as the southern half of the wind farm remains accessible. The task completed at T74 is an example of prioritisation as it is completed earlier than the more urgent task at T19, because it cannot be accessed in the second shift. As the task is only considered finished after its technicians have been picked up, 3 extra periods of work can potentially be performed. The optimal solution only transports 11 technicians to the wind farm despite having the availability of an extra technician.

The deterministic solution for scenario 2 in Figure 2.4.4b services 5 turbines and reuses technicians across multiple tasks to improve efficiency. The task at T13 is visited first to allow it to be completed within the current shift as it will not be accessible in the subsequent shift. The vessel then transports technicians to T63 and T74, before immediately returning to T63 to pick-up its technicians after 6 periods of work are performed. During the interval of high wind speed these technicians are moved to the corrective task at T45 and the preventive task at T21 before the vessel returns for its second visit to T74. This action allows 4 tasks to be immediately restarted after the pause in work but exacerbates the imbalance in work performed by different technicians. The return transit to port picks up all 12 technicians from T74, T45, T21 and T13 respectively. The asymmetric travel time between P1 and T13 is due to the rounding differences in the travel and transfer time approximations, in conjunction with the outgoing transit including time for the initial loading in port.

Figure 2.4.4c illustrates the stochastic solution which is informed about the relative likelihoods of wind farm accessibility in shift 2 and correspondingly performs maintenance on corrective tasks located in both the northern and southern halves of the wind farm. This contrasts with the deterministic solutions where work is never performed on turbines that will become inaccessible in the second shift. The vessel initially transports technicians to T74, T45, T21 and T13 via a line drop-off. The task at T21 is a preventive task so the vessel returns to collect its technicians at the start of the interval of high wind speed. These technicians are then directly moved to T63 along with the final spare technician onboard before the vessel travels to and waits at T13. After the interval of high wind speed concludes the vessel picks up the

remaining technicians from T13, T63, T74 and T45. This solution differs from the second deterministic scenario as the work performed is more heavily concentrated on the two corrective tasks at T74 and T45.

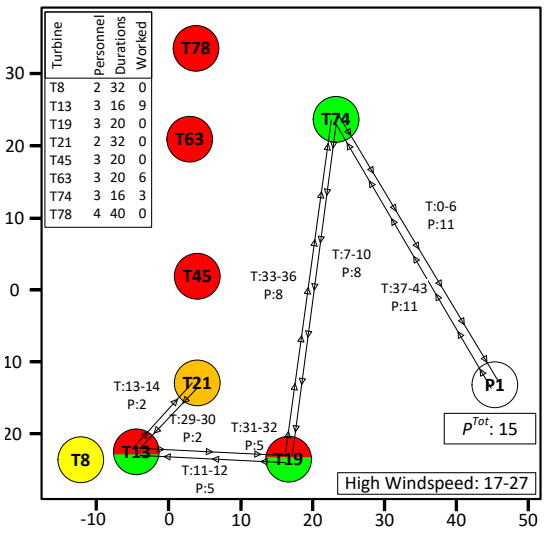

(a) Deterministic solution for scenario 1.

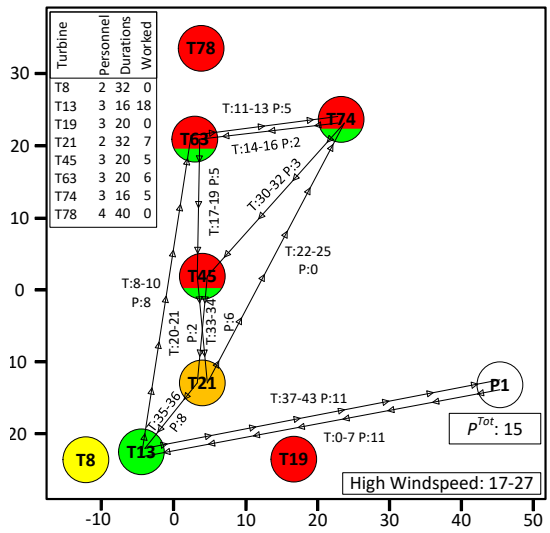

(b) Deterministic solution for scenario 2.

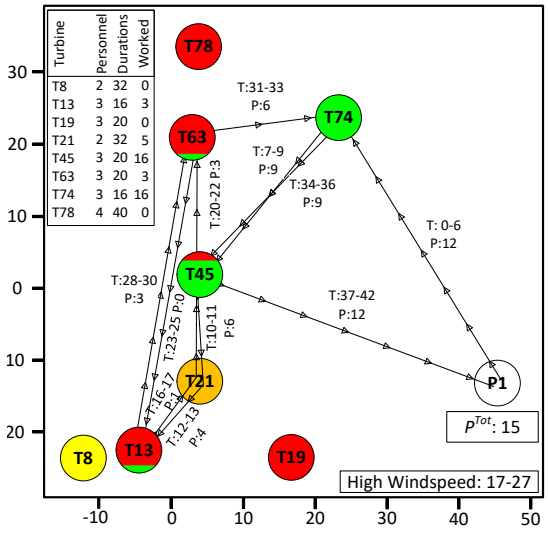

(c) Stochastic solution for scenario 1 and scenario 2.

Figure 2.4.4: Comparison of vessel routes and task completions for the common first shift. Preventive tasks not worked on are yellow, those worked on are orange. Corrective tasks are coloured red. Green is added to corrective tasks in proportion to the percentage of total work performed on them. Arc label T states the periods of travel, whilst P represents the number of technicians onboard.

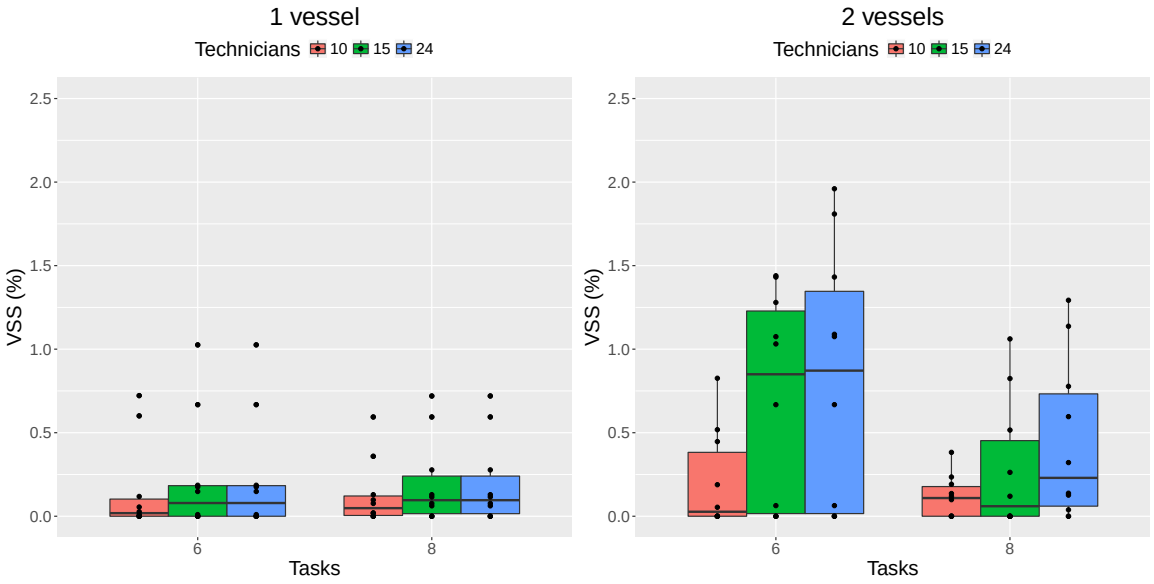

Figure 2.4.5: Quantifying the value of stochastic solutions for the single look ahead approach.

To measure the improvement that can be achieved by solving the stochastic model instead of its deterministic counterpart, we compute the value of a stochastic solution (VSS). The VSS and the relative VSS_r show the impact of the uncertainty on the solution in the first shift,

$$VSS = z^*_{det} - z^*_{stoc} \tag{2.4.2}$$

$$VSS_r = (z^*_{det} - z^*_{stoc})/z^*_{det} \tag{2.4.3}$$

where z^*_{stoc} is the optimal solution of the stochastic problem and z^*_{det} is the optimal solution of the stochastic problem with the actions performed in the first stage fixed to those found in the optimal solution of the deterministic problem. A large value of VSS_r indicates that there is a value in solving a stochastic model, compared to its deterministic equivalent.

The results of an initial implementation of this stochastic approach are given in Figure 2.4.5, which shows the relative value of a stochastic solution (VSS_r). Certain instances have optimal costs which are identical across the deterministic and stochastic solutions as the majority of tasks can be completed before the second shift. This became less common when the second vessel was added as it can be leveraged to further improve solutions. A single task list's best stochastic solution found within 2 hours was greater than its corresponding optimal deterministic solution due to it not being solved to optimality, so we set its percentage saving to zero. We observe that for our extreme example, containing contradictory expectations of the metocean conditions in the second shift, there is little benefit to a stochastic solution. This is because most of the extra work completed is optional preventive maintenance.The implementation of this iterative framework can be extended and improved upon in further work. A multi-period look ahead method using more than 2 shifts could be used to better account for future uncertainty. The effect of introducing new tasks into the system could also be investigated to validate the value of stochastic solution in a dynamic setting.

2.5 Conclusions

We have presented a new mathematical formulation of the offshore maintenance routing problem that addresses several limitations of the existing literature. The existing approach to task completion is extended to a fully time sensitive framework both within and across shifts. We have illustrated the effect of omitting or postponing

O&M tasks, which helped to define a new class of resource restricted instances that are tougher to solve to optimality. A novel technician-vessel maximum safety range constraint is introduced with its effect quantified on smaller instances. Stochastic solutions show the value of performing opportunistic maintenance in the presence of a variety of weather and resource restrictions.

This work could be developed further to incorporate multiple concurrent tasks at offshore turbines and fully differentiate technician skill types. Our model could be extended to include larger service operations vessels used primarily at distant far-shore wind farms by introducing an extra offshore location where these vessels wait overnight. It would then be possible to determine optimal routes and schedules involving a mixed vessel portfolio. An additional target would be to quantify and fully integrate the long term benefits of completing preventive maintenance into a short term horizon which is mainly focused on corrective maintenance. Condition based maintenance strategies that monitor individual component degradations provide significant motivation for further research in this direction. It is hoped that this work can eventually be applied to a real life instance.

Chapter 3

Heuristic Method

3.1 Introduction

Offshore wind power is currently one of the fastest growing methods of generating renewable electricity on the planet. The global offshore wind industry installed a record 6.1GW of new capacity in 2019, bringing the total capacity to 29GW, according to Council (2019). It is predicted that the total global capacity could exceed 90GW within the next five years.

The breakneck expansion of offshore wind farms has been powered by environmental subsidies and the development of larger turbines. Opportunities to reduce the high operating costs incurred by offshore wind farms still exist in the form of daily operations and maintenance activities. O&M activities are believed to contribute up to 35% of a wind farm's total lifetime costs (Shafiee, 2015), providing a clear avenue for efficiency savings. One key direction involves the day to day scheduling of maintenance activities and the design of optimal routes for vessels and technicians to

complete said tasks.

Two common types of maintenance actions required at offshore wind farms are known as preventive and corrective maintenance respectively. Preventive maintenance requires turbines to be temporarily shutdown in order to maintain and restore components before they fail, ideally at the least disruptive times. Corrective maintenance is a responsive action to an unexpected failure which ideally merits a rapid repair response to avoid a costly accumulation of lost potential revenue whilst the turbine is inactive. This can conflict with the notion of minimising travel costs, but may harmonize with opportunistic preventive maintenance once a vessel is offshore. A successful mathematical model of offshore operations must consider both the routing and scheduling aspects of the problem in parallel. A pure scheduling approach cannot harness potential savings from sharing technicians across multiple vessels, whilst a stand alone routing model is unable to cope with discontiguous task completion and prioritisation.

These challenges will be particularly apparent inside newer offshore wind farms that contain more turbines over a larger area than ever before. The Hornsea Project One wind farm is set to be the world's biggest wind farm containing over 170 turbines and is planned to expand further in several stages, Orsted (2018). Routing and scheduling problems of this size will require mathematical decision support tools and likely need heuristic methods given the nuances associated with offshore work and transportation.

The structure of this chapter is as follows. Section 3.2 presents a brief review on existing literature associated with the offshore wind farm maintenance and routing

problem. Section 3.3 describes the optimisation model for the problem. Section 3.4 motivates and describes the metaheuristic framework for solving the model, before Section 3.5 presents the computational experiments. Section 3.6 states our conclusions.

3.2 Literature Review

Dai et al. (2015) introduce the first mixed integer linear programming (MILP) model of offshore operations that integrates the vessel routing and task scheduling aspects. The model aims to complete a set of offshore tasks so that travel costs and downtime losses are minimised over a short term planning horizon. Tasks completed after a set period are penalised with the model solved through a commercial solver. This model is extended by Irawan et al. (2017) to consider multiple offshore windfarms being serviced by several ports and introduces different technician skill types.

Different problem aspects are examined by subsequent works such as Raknes et al. (2017) who investigate longer tasks which often span across several working days. Their optimisation model also incorporates larger service vessels that can stay offshore for several consecutive shifts. Schrotenboer et al. (2018) consider the impact of sharing technicians across multiple windfarms instead of remaining localised to a single windfarm. They quantify the estimated saving as around 7% by developing an adaptive large neighbourhood search (ALNS) heuristic to assign technicians to ports and vessel routes.

Heuristic methods for short term routing and scheduling of maintenance tasks have

also been investigated, chiefly with the aim of solving instances containing higher numbers of offshore tasks. Kennedy et al. (2016) base their genetic algorithm on the model of Besnard et al. (2011) and find savings of around 15% versus unoptimised scheduling. Dawid et al. (2017) takes the approach of several authors and clusters groups of maintenance activities together with compatible vessels to form a list of feasible maintenance plans. They find this cluster-matching approach compares favourably with commercial solvers, particularly for instances containing less than 15-20 tasks. Dawid et al. (2018) extend this to include probabilities of successfully completing tasks. Stock-Williams and Swamy (2018) utilise a genetic algorithm metaheuristic to create vessel routes that allow for a degree of task prioritisation inside vessel routes. The fitness of each vessel route is later evaluated through a simulator.

We believe a more realistic model of offshore operations lies in the form of a cumulative capacitated vehicle routing problem (CCVRP). This problem differs from the standard CVRP since the objective is to minimise the sum of arrival times to locations rather than the total cost of the routes. It often occurs in disaster management settings such as earthquakes, where it is important to reach as many victims as soon as possible Campbell et al. (2008), Rivera et al. (2015).

Multiple works employ an adaptive large neighbourhood search (ALNS) procedure to solve larger instances and variations of the classical CCVRP. Ribeiro and Laporte (2012) illustrate the benefits of ALNS compared to other methods on a set of classical benchmark instances. Li et al. (2016) use an ALNS procedure to solve a pick-up and delivery problem with profits and reserved requests in a collaborative logistics setting. There are a set number of reserved requests which must be completed within the time

window and an additional number of selective requests which can be carried out if determined to be profitable.

The CCVRP forms the basis of the more accurate downtime losses used by Stålhane et al. (2015) wherein corrective tasks are penalised until their completion, yet preventive tasks are only charged whilst being worked on. Their MILP models a single shift and is reformulated into a path-flow model via Dantzig-Wolfe decomposition. The effects of this model under uncertainty is considered by Irawan et al. (2019) who use it as the basis of a simulation-optimisation algorithm. The deterministic problem is first solved with a large neighbourhood search to efficiently create high quality routes before a simulation tool is used to evaluate the underlying parameters such as vessel travel and personnel transfer time, previously assumed to be deterministic.

We utilise the adaptive large neighbourhood search (ALNS) framework first proposed by Ropke and Pisinger (2006a) for vehicle routing problems. ALNS extends the large neighbourhood search heuristic of Shaw by incorporating multiple destroy and repair operators within the same search space. During each iteration a pair of operators are probabilistically selected based on their historical performance relative to the other combinations. It has been successfully applied to a variety of rich vehicle routing problems (Aksen et al. (2014),Grangier et al. (2016). The direct application of ALNS in our context is more challenging because of the desire to include optional preventive maintenance of variable duration and the separation of work completion from arrival times. This results in the breakdown of the assumption of a fixed time between a drop-off and pick-up activity. We overcome this difficulty by decoupling the timing aspect of the problem from the route sequencing. The binary variables of

the mathematical model corresponding to the turbine visits and the vessel routes are managed by the removal and repair operators of the ALNS. The continuous variables denoting the arrival times to the given locations and amount of work performed are determined by solving the mathematical model with the binary variables fixed for a given route or through a simple heuristic.

3.3 Mathematical Model

We now introduce our mathematical model of the offshore wind farm maintenance routing (OWFMRP) problem in order to state the OWFMRP in a precise fashion. The main differences between the model described in this section and the model described in Chapter 2 are as follows. Firstly, this model is a continuous time model instead of a discrete time period model. Secondly, the objective of the model no longer includes the minimisation of technician and vessel time offshore. Thirdly, this model partitions the node set into two groups: one for pick-up and one for drop-off. Vessel routes must drop-off technicians at the drop-off node and pick-up personnel from the pick-up node. Finally, we longer explicitly model technicians when they are present on a turbine. If they are on a turbine, they are assumed to be performing work.

3.3.1 Notation

The OWFMRP is defined on a complete undirected graph $G = (N, E)$. The set of nodes $N_d = \{1, \ldots, n\}$ and $N_p = \{n + 1, \ldots, 2n\}$ represent the drop-off and pick-up nodes of all tasks respectively. The port at node $\{0\}$ hosts a set of heterogeneous

vessels $V = \{1, \ldots, v\}$ with variable travel costs of K_v which end their respective routes in a dummy node $\{2n + v\}$ respectively. The technician capacity of each vessel is Q_v^{Cap}. Each edge has a vessel specific travel time T_{vij} and each node has a transfer time of τ_{vi} before pick-up or drop-off. These transfer times are the same for pick-up and drop-off. Travel times are symmetric and respect the triangle inequality. This means that the direct travel time of a vessel v between two locations a and b is at least as fast as the travel time via a third location c, so that $T_{vab} \leq T_{vac} + T_{vcb}$. The planning horizon is comprised of a set of distinct shifts $S = \{1, \ldots, s\}$ with the length of shift s limited to T_s^{Max}. The task at turbine $i, i = 1, \ldots, n$ has a drop-off node i and a pick-up node $n + i$, a duration D_i, a technician demand P_i and a classification as either corrective or preventive maintenance. For the pick-up node we define $P_{i+n} = -P_i$. The set of turbines with corrective maintenance tasks $i \in N^{Corr}$ accrue losses of $v_i > 0$ from the start of the shift until their completion, whereas preventive maintenance tasks $i \in N^{Prev}$ provide a benefit v_i only whilst being worked on. A scaling term δ is included to penalise corrective tasks not completed within the planning horizon more severely. Both corrective and preventive maintenance tasks can be ignored if it leads to a more profitable solution.

The aim of the OWFMRP is to define a set of vessel routes starting and ending at port, which collectively visit the drop-off and pick-up nodes of a subset tasks in order to complete the most valuable combination of maintenance. Each route must not require more technicians than the vessel capacity and finish within the duration of the shift. We seek to minimise the lost revenue accrued before corrective tasks are completed, in conjunction with the benefit of working on variable preventive

maintenance.

We use the following decision variables in the model:

- x^s_{vij}, equals to 1, if vessel v travels directly along arc (i, j) in shift s; 0, otherwise.
- y^s_{vi}, equals to 1, if vessel v visits location i in shift s; 0, otherwise.
- c^s_i, time that a vessel arrives to location i in shift s.
- q^s_{vi}, number of technicians onboard vessel v after visiting location i in shift s.
- t_i, completion time of corrective task at turbine i.
- w_i, amount of time spent working on preventive task i.
- $\tilde{a}^s_i$, equals to 1, if the turbine i which contains a corrective task is visited in shift s; 0, otherwise.
- a_i, equals to 1, if the turbine i which contains a corrective task is completed; 0, otherwise.

3.3.2 Formulation

$$\min \sum_{i \in N_p^{Corr}} v_i t_i - \sum_{i \in N_p^{Prev}} v_i w_i + \sum_{v \in V} \sum_{i \in N} \sum_{j \in N} \sum_{s \in S} K_v T_{vij} x^s_{vij} \tag{3.3.1}$$

$$\sum_{\substack{i,j \in N_d \cup N_p \\ \cup \{0, 2n+v\} \\ i \neq j}} T_{vij} x^s_{vij} + \sum_{i \in N_d \cup N_p} \tau_{vi} y^s_{vi} \leq T^{Max}_s, \quad \forall v \in V, \forall s \in S \tag{3.3.2}$$

$$\sum_{j \in N_d \cup N_p \cup \{2n+v\}} x^s_{v0j} = 1, \quad \forall v \in V, \forall s \in S \tag{3.3.3}$$

$$\sum_{j \in N_d \cup N_p \cup \{0\}} x^s_{vj,2n+v} = 1, \qquad \forall v \in V, \forall s \in S \tag{3.3.4}$$

$$y_{vis} = \sum_{\substack{j \in N_d \cup N_p \cup \{2n+v\} \\ j \neq i}} x^s_{vij}, \qquad \forall v \in V, \forall i \in N_d \cup N_p, \forall s \in S \tag{3.3.5}$$

$$y_{vis} = \sum_{\substack{j \in N_d \cup N_p \cup \{0\} \\ j \neq i}} x^s_{vji}, \qquad \forall v \in V, \forall i \in N_d \cup N_p, \forall s \in S \tag{3.3.6}$$

$$c^s_i \geq c^s_j + T_{vji} + \tau_{vi} - M(1 - x^s_{vji}), \qquad \forall v \in V, \forall i \in N_d \cup N_p \cup \{2n+v\}, \tag{3.3.7}$$

$$\forall j \in N_d \cup N_p \cup \{0\}, i \neq j, \forall s \in S$$

$$c^s_0 = 0, \qquad \forall s \in S \tag{3.3.8}$$

$$\sum_{v \in V} y^s_{vi} = \tilde{a}^s_i, \qquad \forall i \in N_d \cup N_p, \forall s \in S \tag{3.3.9}$$

$$a_i \leq \tilde{a}^s_i, \qquad \forall i \in N^{Corr}_p \tag{3.3.10}$$

$$c^s_i \leq T^{Max}_s \tilde{a}^s_i, \qquad \forall i \in N_d \cup N_p, \forall s \in S \tag{3.3.11}$$

$$\sum_{s \in S} c^s_i - \sum_{s \in S} c^s_{i-n} \geq D_i a_i, \qquad \forall i \in N^{Corr}_p \tag{3.3.12}$$

$$\sum_{s \in S} c^s_i - \sum_{s \in S} c^s_{i-n} \geq w_i, \qquad \forall i \in N^{Prev}_p \tag{3.3.13}$$

$$t_i \geq \sum_{u=1}^{s-1} T^{Max}_u \tilde{a}^s_i + c^s_i, \qquad \forall i \in N^{Corr}_p, \forall s \in S \tag{3.3.14}$$

$$t_i \geq \delta \sum_{u=1}^{s} T^{Max}_u (1 - a_i), \qquad \forall i \in N^{Corr}_p, \forall s \in S \tag{3.3.15}$$

$$q^s_{vi} \geq q^s_{vj} + P_j - (Q^{Cap}_v + P_j)(1 - x^s_{vji}), \qquad \forall v \in V, \forall i \in N_d \cup N_p \cup \{2n+v\}, \tag{3.3.16}$$

$$\forall j \in N_d \cup N_p \cup \{0\}, i \neq j, \forall s \in S$$

$$x^s_{vij} \in \{0,1\}, \qquad \forall v \in V, \forall i,j \in N_d \cup N_p \cup \{0, 2n+v\}, \forall s \in S \tag{3.3.17}$$

$$y^s_{vi} \in \{0,1\}, \qquad \forall v \in V, \forall i \in N_d \cup N_p, \forall s \in S \tag{3.3.18}$$

$$c^s_i \geq 0, \qquad \forall i,j \in N_d \cup N_p \cup \{0, 2n+v\}, \forall s \in S \tag{3.3.19}$$

$$0 \leq q^s_{vi} \leq Q^{Cap}_v, \qquad \forall v \in V, \forall i, j \in N_d \cup N_p \cup \{0, 2n+v\}, \forall s \in S \tag{3.3.20}$$

$$t_i \geq 0, \qquad \forall i \in N^{Corr}_p \tag{3.3.21}$$

$$D_i \geq w_i \geq 0, \qquad \forall i \in N^{Prev}_p \tag{3.3.22}$$

$$\tilde{a}^s_i \in \{0,1\} \qquad \forall i \in N_d \cup N_p, \forall s \in S \tag{3.3.23}$$

$$a_i \in \{0,1\} \qquad \forall i \in N^{Corr}_p \tag{3.3.24}$$

The objective function (3.3.1) minimises the net cost of the solution from the sum of the lost revenue accrued before corrective tasks are completed subtracted by the benefit from working on preventive tasks and the total vessel transportation costs. Constraints (3.3.2) limit the duration of a vessel route to the length of its shift. Constraints (3.3.3) and (3.3.4) ensure that each vessel must leave the port and return back to it at the end of the route. Constraints (3.3.5) establish that if a vessel visits a drop-off or pick-up node it must depart onwards to another location. Similarly constraints (3.3.6) provide an immediate predecessor to a visited node. Constraints (3.3.7) specify that if location i is visited directly after location j by vessel v, then the difference between their arrival times must be at least $T_{vji} + \tau_{vi}$. M is a sufficiently large positive constant. Constraints (3.3.8) set the departure time of all vessels in port to zero, whilst constraints (3.3.9) ensures that at most one vessel can visit each location in a shift. Constraints (3.3.10) allows work to be performed on corrective tasks in a shift without necessarily completing the task and constraints (3.3.11) force the arrival to time to unvisited locations to zero. Constraints (3.3.12) and (3.3.13)

determine the amount of time spent on corrective and preventive maintenance tasks respectively. Constraints (3.3.14) calculate the completion time for finished corrective tasks, whereas constraints (3.3.15) ensure extra penalised downtime losses accrue for the entire planning horizon for incomplete tasks. Constraints (3.3.16) deal with the technician load of the vessels. Finally constraints (3.3.17-3.3.24) define the nature of the decision variables.

The formulation is a less compact version of the models derived from Stålhane et al. (2015) versions since we consider variable preventive maintenance durations and include explicit variables for if tasks are executed. We choose this form to allow us to model vessels visiting turbines without necessarily completing tasks within the same shift. For example we might complete the majority of a preventive task in the first shift and the remainder in the subsequent shift. Additionally we no longer impose the task duration as a hard constraint for time spent on a turbine or between visits, we have a softer form which penalises costs until technicians are picked up.

The outlined formulation contains a large number of binary, integer and continuous variables and is unable to be solved efficiently using optimisation software. After thirty minutes of computational time Gurobi reported an optimality gap of 51.17% when dealing with just fourteen corrective maintenance tasks and three vessels. Our industrial partners indicated their desire for a solution to be produced in less than thirty minutes, which ruled out the use of an exact formulation for all but very small instances. The difficulty in solving instances with multiple vessels and more than 12 tasks matches with our findings from Chapter 2, so we conclude that simply solving the mathematical model is impractical for realistic sized instances containing more

than twenty tasks.

We note that 3.3.7 and3.3.16 are the classic Miller,Tucker and Zemlin constraints for the temporal and capacitated aspects of the OWFMRP respectively (Miller et al., 1960). These can be lifted to form tighter constraints as described by Desrochers and Laporte (1991) for a variety of routing problems. We experimented with including or excluding the lifted versions of the MTZ constraints, however we found them their inclusion did not allow us to solve any larger instances of our model to optimality. The reasons for this could be investigated further, however we chose not to focus on this as our goal was to provide a heuristic for large problem instances.

3.4 Solution Framework

In this section we define the decomposition strategy to solve the problem. The motivation for this is to reduce the complexity of the overall OWFMRP by decomposing the problem into two smaller sub-problems that are each easier to solve. The binary variables x^s_{vij}, y^s_{vi} and a^s_i are handled by the removal and repair operators of an ALNS procedure, whereas the continuous variables c^s_i, t_i and w_i are determined from a timing sub-problem.

We choose this decomposition structure for three reasons. Firstly it provides for an efficient calculation of the amount of preventive maintenance work performed, including cases where it might be beneficial to delay a pick-up from a completed corrective maintenance task. Secondly it caters for the potential inclusion of drop-off and pick-up nodes of the same task being visited by different vessels. This is an

emerging development in offshore systems, as new vessels contain swipe-on and swipe-off technician access systems. These systems can allow for technicians' locations to be tracked throughout the day across different vessels. Finally we can embed feasibility checks into the sequencing sub-problem as equipment loads, technicians, and skill types can be generated from the sequence of visits a priori to the timing sub-problem.

In a similar vein to Adulyasak et al. (2012) we refer to the algorithm as an optimisation-based adapative large neighbourhood search (Op-ALNS), since the decomposed timing sub-problem can be handled with exact methods. The pseudocode for our ALNS procedure is given in Algorithm 1 with more details outlined later in Section 3.4.2.

Data: Instance data; Set of operators O

$S \leftarrow$ simple construction heuristic;

$S_{best} \leftarrow S$;

$\mathcal{T} \leftarrow \mathcal{T}_{start}$;

Initialise weights and scores of operators;

for $seg \leftarrow 1$ **to** N_{segs} **do**

  Select the policy for operators;

  for $iter \leftarrow 1$ **to** N_{iters} **do**

    Select the destroy, repair operators O^-, O^+ via roulette wheel selection of operator weights;

    Apply operators O^-, O^+, $S \rightarrow S'$;

    if S' *feasible* **then**

      if $f(S') < f(S)$ **or** $Unif(0,1) < e^{\frac{f(S)-f(S')}{\mathcal{T}}}$ **then**

        $S \leftarrow S'$;

        if $f(S) < f(S_{best})$ **then** $S_{best} \leftarrow S$;

      end

    end

    Update scores of O^-, O^+;

    $\mathcal{T} \leftarrow \mathcal{T} * c$;

  end

  Update the weights and reset scores of the operators;

  Call local search procedure on S;

  if $f(S) < f(S_{best})$ **then** $S_{best} \leftarrow S$;

end

return $S_{best}, f(S_{best})$

Algorithm 1: ALNS procedure

3.4.1 Timing Sub-Problem

The output of the removal and repair operators can be thought of as creating a sequence of locations to be visited for each vessel. Note that the sequence of tasks proposed by the removal and repair operators will be a feasible assignment of tasks to vessels, but not necessarily a feasible sequence of visits for each vessel route. In the timing sub-problem we clarify this and determine the optimal arrival times and hence task completions for the planned visits. We fix x^s_{vij}, y^s_{vi} to their value from the sequencing sub-problem in the full mathematical model (3.3.1)-(3.3.24) to form the timing sub-problem. The work completion variables a_i can be dealt with in two ways. We can either assume their values from the proposed routes, so that if a location is included in a route it must be visited and worked to completion. However it is possible that a better solution could be obtained in some circumstances by potentially omitting work at certain locations from the routes. The cumulative aspect to the problem means that any travel time directed towards wasteful tasks is heavily penalised in the corrective downtime losses. This can be avoided by assuming that work is performed, but not necessarily completed, if a turbine is visited in a shift. This is equivalent to allowing only the a_i variables to remain unfixed in the timing sub-problem. The resulting model becomes mixed integer albeit with only a few binary variables and the proposed routes trimmed afterwards if tasks are ignored.

Timing Sub-Problem Simplification

Whilst our results have shown that this sub-model remains tractable for the instance sizes we wish to tackle, it may be harder to directly incorporate within the ALNS

procedure. Our practical implementation revealed that the frequent creation and solving of the sub-problem took longer than expected. Deeper analysis indicated that the despite its simplicity, repeatedly updating and fixing binary variables formed a temporal bottleneck on the process. This lag was an order of magnitude larger than the solving time of the MILP vastly reducing the number of iterations that could be performed in a fixed time. We therefore deemed this direct approach unbeneficial because of its associated computational complexity.

For practical use it may be sufficient to solve a restricted form of the sub-model with additional assumptions. In this case the aim would be to complete as many, if not all of the corrective tasks, before squeezing in additional preventive maintenance. This requires the assumptions that the losses per unit time of corrective tasks is strictly bigger than the benefit per unit time of any preventive tasks, which in themselves are strictly bigger than the routing costs and that all corrective tasks in a route are completed. As we illustrate later these are natural assumption for realistic instances which allow us to determine the total cost of vessel routes directly.

The variability in preventive maintenance makes it natural to model its inclusion as a two step process. The first step requires the calculation of arrival times to tasks assuming corrective tasks are always completed and zero work is performed on preventive tasks. Any slack time found at the end of a vessel route should then be repurposed into time spent working on preventive tasks in the subsequent step. As the inclusion of preventive work impacts subsequent tasks, we first examine the marginal effects of delaying existing activities in a vessel route.

The marginal benefit of moving implicit slack past a given location depends on

its type. Table 3.4.1 lists the possible combinations which depend on several factors. For example delaying a corrective drop-off causing additional downtime costs at the rate of ω per minute. We do not consider moving the preventive maintenance slack from T_p^* to locations preceding T_d^* as it would imply the corrective maintenance task itself becomes incomplete. An example case is illustrated in Figure 3.4.1. Each of the positions ($\uparrow$) is a candidate for having the implicit slack reassigned to. Every position the slack is moved past causes the accumulation of marginal benefits according to the outlined rules. Thus the total benefit is determined from the summation of all the marginal benefits of locations past which the implicit slack is moved. For example the total benefit at $posn_4$ is $0 - \omega_2 - \omega_9 + \omega_1$, since the implicit slack would need to be moved prior to a preventive pick-up, a preventive drop-off and a corresponding corrective drop-off and pick-up. Notice the correction to T9d in order to avoid double penalising delay to corrective tasks. We note that this is essentially the same situation as the slack presenting itself at the end of the shift. In that case the T_d^* and T_p^* would be replaced by the port.

Type	Partner	Marginal Benefit
Corrective drop-off	Its pick-up after T_p^*	-ω
Corrective pick-up	Its drop-off before T_d^*	-ω
Preventive drop-off	Its pick-up after T_p^*	-ω
Preventive pick-up	Its drop-off before T_d^*	+ω
Corrective drop-off	Its pick-up before T_p^*	0
Preventive drop-off	Its pick-up before T_p^*	+ω

Table 3.4.1: Marginal benefits for delaying maintenance activities due to inserting slack.

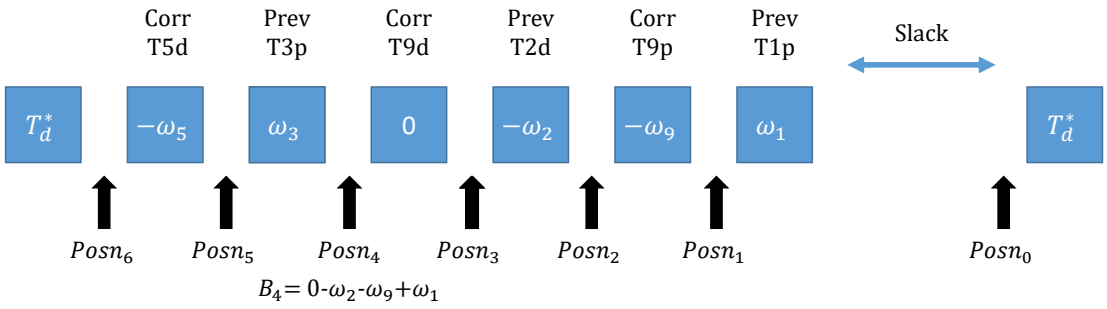

Figure 3.4.1: Analysis of marginal benefit of moving slack around a route.

Assuming the slack originates in $Posn_0$ with $value_{Posn_0} = 0$ then we can recursively

compute the value of the preceding positions $Posn_1$, $Posn_2$, ..., $Posn_n$.

$$\forall n \geq 1, \quad Benefit_{Posn_n} = Benefit_{Posn_{n-1}} + \begin{cases} +\omega, & \text{if } Posn_n \text{ is a preventive pick-up} \\ 0, & \text{if } Posn_n \text{ is a corrective drop-off} \\ & \quad \& \text{ corresponding pick-up already delayed} \\ -\omega, & \text{otherwise} \end{cases} \tag{3.4.1}$$

Under the assumption that we want to place the slack into a position containing as large a cumulative benefit as possible, Equation 3.4.1 shows us that we should never choose a $-\omega$ location. This corresponds to only ever placing the slack in front of either a corrective drop-off with the corresponding pick-up also delayed before T_p^* or directly before a preventive pick-up. We choose to only consider ending on a $+\omega$ position rather than a 0 valued position. In this case the surplus slack will be associated with the travelling from a corrective drop-off to a preventive pickup. The option to delay the departure time gives more freedom to complete the corrective maintenance task in practice. The mathematical formula is equivalent to the logical case described earlier as it only relies on the assumption that every $\omega > 0$.

$$\begin{array}{l}\text{Net benefit of} \\ \text{preventive pick-up} \\ \text{in } Posn_i\end{array} = \left\{\begin{array}{l}\text{Preventive} \\ \text{benefit} \\ \text{of task } i\end{array}\right\} - \left\{\begin{array}{l}\text{Costly delays} \\ \text{to subsequent} \\ \text{corrective pick-ups}\end{array}\right\} + \left\{\begin{array}{l}\text{Beneficial delays to} \\ \text{later preventive pick-ups} \\ \text{that started before } i\end{array}\right\} \tag{3.4.2}$$

We formalise this approach mathematically in Equation 3.4.3, where we define B_i as the net benefit of delaying the pick-up activity of preventive maintenance activity

i.

$$B_i = \omega_i - \sum_{\substack{j \text{ after } i \\ j \in N_p^{Corr}}} \omega_j + \sum_{\substack{k \text{ after } i \\ k \in N_p^{Prev} \\ k-n \text{ before } i}} \omega_k \tag{3.4.3}$$

The benefit calculation is comprised of three terms. The first term is the pure benefit from performing work on a preventive task, whereas the remaining terms quantify the effects on subsequent corrective and preventive tasks respectively. Additional time assigned to preventive task i will delay the arrival to any subsequent corrective pick-up nodes in the same route. Costs will continue to accrue for these tasks despite the work being finished. Preventive tasks which overlap the task, defined by their drop-off node being visited before i and their pick-up node after, provide benefit as scheduled work at i will similarly occur on them.

Our heuristic for assigning preventive work applies the available slack time to tasks with the highest benefit preference to the tasks with the highest benefit. This process continues until either all the slack time is applied or there are no more locations with positive benefit available. The assumption that corrective tasks must be completed often means that surplus time separates consecutive locations in routes. A second update procedure allows for surplus time associated with preventive tasks to be shifted around locally to yield greater benefits. For simplicity and computational speed we only consider moving the surplus time to preceding tasks. Both of these steps are described in Algorithm 2 for a given vessel route.

Data: Route sequence $R_{1:n}$, slack, net benefits B

Sort B such that $i < j \implies B_i > B_j$;

for $q \leftarrow 1$ **to** $|B|$ **do**

  if $slack > 0$ **and** $B_q > 0$ **then**

    for $l \in R_{route.index(q):n}$ **do**

      $C_l^s \leftarrow C_l^s + \min(slack, D_q)$;

    end

    $slack \leftarrow slack - \min(slack, D_q)$;

  end

end

for $q_0, q_1 \in R_{n-1:1} \odot R_{n:2}$ **do**

  if $q_0 \in N_p^{Prev}$ **then**

    $availableTime = C_{q_1}^s - C_{q_0}^s$;

    if $availableTime > T_{vq_0q_1} + \tau_{vq_1}$ **then**

      $C_{q_0}^s \leftarrow C_{q_0}^s + availableTime - T_{vq_0q_1} - \tau_{vq_1}$;

    end

  end

end

Algorithm 2: Preventive work assignment heuristic to update the values of timing sub-problem variables.

Whilst this approach will likely determine the optimum amount of preventive maintenance in a given route, we note that it is a heuristic method and as such may not always provide the optimal solution.

3.4.2 ALNS Framework

ALNS was first introduced by Ropke and Pisinger (2006a) to generalise the large neighbourhood search heuristic by utilising multiple removal and repair operators within the same search procedure. Some of these heuristics attempt to focus the search around previously discovered high quality solutions whilst other heuristics help to diversify the search. It can be seen as an extension of large neighbourhood search (LNS) to include a reinforcement learning layer for choosing which operators to apply.

During each iteration a removal and repair operator combination is selected from a list according to an adaptive and probabilistic mechanism. These are applied in succession to destroy a significant percentage of the current solution before repairing it. This new solution is then accepted or rejected according to an acceptance criterion such as simulated annealing. The likelihood of reselecting the previous removal-repair combination is then updated based on its performance. The repeated application of these steps cause the chance of choosing to be linked with their historical performance. We now describe the key elements of our ALNS heuristic for the problem.

Neighbourhood size

During each iteration a solution containing L tasks will have a total of $\lceil \rho L \rceil$ tasks removed from the solution, where $\rho \in [0, 1]$ is the removal fraction. These tasks will then be gradually reinserted into the routes with only tasks that improve upon the incumbent solution without violating constraints being kept. Tasks deemed to be unprofitable and not inserted in a particular iteration are retained in the task pool. This pool is merged with the tasks removed from subsequent destroy operators allowing

the possible reinsertion of tasks later in the algorithm. Every solution produced at the end of an iteration will be feasible since every corrective and preventive maintenance task is considered optional.

If the algorithm cannot find an improved feasible solution in $\Gamma = 100$ iterations, we double the removal fraction ρ. This helps to avoid repeatedly searching the same solution space. Additionally, after the first 10 iterations we apply each previously unselected operator combination to ensure all combinations are applied before the initial weight calculations. The size of ρ is of crucial importance. If only a small percentage of tasks in the solution are removed then the benefits of large neighbourhood search will be lost. If a large percentage of the solution is destroyed then the application of repair operators can be very time consuming and inefficient at finding improvements. We initially set $\rho = 0.2$ in our experiments in accordance with previous works.

Adaptive search engine

The selection of removal and repair operators should be informed by their previous performance but not at the total expense of diversification. Previously unsuccessful operators should maintain a small chance of selection in the future, requiring a probabilistic approach. A weight ω_i is associated with each removal-repair operator combination, which is a guide to the historical performance of combination i in previous iterations. Assuming a set of n operator combinations, the probability p_j of choosing combination j is controlled by roulette-wheel selection of all the weights, $p_j = w_j / \sum_{i=1}^{n} w_i$.

Adaptive weight adjustment

Sets of consecutive iterations are defined as segments during which the weights remain fixed. At the beginning of the ALNS each operator combination weight is set to one, but these are updated at the end of a segment based on the scores collected during it. The scores reflect the success of each operator at finding new or improved solutions within the segment. If the operator finds a feasible solution we update the operator score by either $\sigma_1, \sigma_2, \sigma_3$ or σ_4. A reward of $\sigma_1 = 10$ is given should a new globally best solution be discovered, $\sigma_2 = 5$ for an undiscovered and $\sigma_3 = 3$ for a rediscovered solution that improves upon the incumbent respectively. $\sigma_4 = 1$ is awarded if a worse solution is accepted through a stochastic acceptance criteria. We assume $\sigma_1 > \sigma_2 > \sigma_3 > \sigma_4$ as there is a clear hierarchy in the properties of new solutions. This follows the approach of Li et al. (2016).

Once a segment is completed the next set of adaptive weights are calculated from the existing scores before they are reset to zero for the next segment. Let $w_{i,seg}$ be the weight of operator combination i in segment seg, $time_{i,seg}$ be the total time combination i was used and $score_{i,seg}$ its resulting score. Then

$$w_{i,seg+1} = \begin{cases} w_{i,seg}, & \text{if score}_{i,seg} = 0. \\ (1-\eta)w_{i,seg} + \eta\frac{\text{score}_{i,seg}}{\text{time}_{i,seg}}, & \text{otherwise,} \end{cases} \tag{3.4.4}$$

where $\eta \in [0, 1]$ is a reaction factor which controls the speed of blending between the historical operator performance and its most recent scores. $\eta = 1$ implies that weights are purely controlled by the score to computational time ratio in the previous segment, whereas $\eta = 0$ simply maintains the previous weight. As our experiments have a

limited computational time we scale scores by the total running time rather than their frequency. This helps to incentivise quick and efficient operator combinations. We set $\eta = 0.7$ in our experiments.

Acceptance and stopping criteria

Acceptance criteria within meta-heuristics can be classified as either deterministic or stochastic. Deterministic methods produce the same decision regardless of any exterior factors. Stochastic methods produce results that depend on the state of the search and often involve probabilistic actions. We include simulated annealing, time dependent simulated annealing and great deluge as stochastic acceptance criteria in our study. The deterministic acceptance criteria of only allowing improving solutions is also considered.

The most common stochastic acceptance method often employed within ALNS procedures is simulated annealing. It specifies that improving solutions should always be accepted, whilst worse solutions are to be accepted with a given probability. This probability depends on the cost of the proposed move and the state of the search, which helps to reduce the chance of the method getting trapped in a local optima. The process is underpinned by the temperature $\mathcal{T}$ which is gradually reduced during the search by multiplication with a cooling constant $0 < c < 1$, so that the probability of accepting a worse solution correspondingly lowers. The formula for this probability is

$$p = e^{-\frac{\Delta f}{\mathcal{T}}}, \tag{3.4.5}$$

where $\Delta f = f(S) - f(S')$ is the change in cost from the proposed move. $\mathcal{T}$ is initially

set to $\mathcal{T}_{start}$, which is computed from the choice to accept a solution 30% worse than the initial solution with a probability $\tau = 20\%$. When $\mathcal{T}$ eventually reaches zero the process reduces to the deterministic approach of only accepting improving solutions.

Simulated annealing has also been extended to a variety of settings. Bilgin et al. (2006) adjusted the formula for limited computational time to accept non-improving solutions with a probability at time t governed by the following equation

$$p_t = e^{\frac{-\Delta f}{\Delta F\left(1-\frac{t}{T}\right)}}, \tag{3.4.6}$$

where Δf is the change in the cost from the proposed move at time t, ΔF is the range for the maximum change in the objective function, T is the time limit and t is the current time within the search.

The great deluge algorithm was first proposed by Dueck (1993) and is based on a stochastic framework which allows improving moves by default. The decision of whether to accept a worsening move depends on if its quality is better than an expected cost. The common analogy with nature describes this expected cost as a gradually rising water level in which a person climbing a hill will try to move in any direction in order to avoid getting their feet wet. For our minimisation problem with limited computational time, the water level is set equal to the cost of the initial solution. It decreases linearly in proportion with the remaining time

$$\tau_t = f_0 + \Delta F\left(1 - \frac{t}{T}\right), \tag{3.4.7}$$

where ΔF is the expected maximum change in the objective function, f_0 is the expected final cost, T is the time limit and t is the current time within the search. As f_0 is unknown it can be replaced with a suitable lower bound. Great deluge has an

advantage over simulated annealing in the sense it requires fewer parameters to be fixed.

3.4.3 Removal (O^-) and Repair (O^+) Heuristics

This section describes seven removal heuristics and four repair heuristics. Given the set of tasks currently contained in the vessel routes L and a removal fraction $\rho \in [0, 1]$, each removal operator removes a total of $\lceil \rho |L| \rceil$ tasks from the routes. A degree of randomisation is associated with the removal operators, using the approach of Ropke and Pisinger (2006a) $p \geq 1$ in order to help diversification. Our randomisation scheme picks the $\lfloor U(0,1)^3 |L| \rfloor$th item from a ranked list, where U(0,1) is a random number generated from the standard uniform distribution.

Random removal heuristic

This simple removal heuristic removes a total of $\lceil \rho |L| \rceil$ tasks from the current solution S. The advantage of including this operator is to enhance the diversification of the explored solutions.

Least expensive removal heuristic

This heuristic emphasises the fact that it is more beneficial to perform tasks with greater downtime costs suggesting that the least expensive tasks are better candidates for removal.

Price similarity removal heuristic

This heuristic attempts to remove tasks that are most similar in price to an initial seed removal request subject to randomisation. A price similarity measure can be defined between two tasks i and j according to their price difference, $P(i,j) = |p_i - p_j|$. This makes the heuristic biased towards removing tasks that can be easily interchanged, given the expectation that reshuffling them is more likely to produce better solutions.

Shaw removal heuristic

The Shaw removal heuristic employed by Ropke and Pisinger (2006a) and Shaw (1997) expands the notion of similar tasks to examine more features. A seed task i is randomly selected from the tasks contained in the incumbent solution, which is used to define the similarity $\mathcal{S}(i,j)$ between itself and another task j. The lower the $\mathcal{S}(i,j)$ the more similar the two tasks are. We use a similarity measure consisting of three terms: a distance term, an arrival time term and a technician term. These terms are weighted using α, β and γ respectively. Each term is scaled to be inside $[0,1]$ as α, β and γ correspond to the maximum possible inter-turbine distances, arrival times and technician differences respectively.

$$\mathcal{S}(i,j) = \frac{\sqrt{d_{i,j}^2 + d_{i+n,j+n}^2}}{\alpha} + \frac{|C_i - C_j| + |C_{i+n} - C_{j+n}|}{\beta} + \frac{|P_{i+n} - P_{j+n}|}{\gamma} \tag{3.4.8}$$

Neighbour removal heuristic

Ropke and Pisinger (2006b) and Ribeiro and Laporte (2012) suggest the use of another class of removal heuristics that make use of the existing historical information

when removing tasks. This may be a beneficial inclusion but can counteract the diversification benefits from the overall ALNS procedure. The neighbour removal heuristic stores information in a complete, directed and weighted graph called the 'neighbour graph', an example of which is shown in Fig 3.4.2a. Each node represents a location that has been been visited in the problem. Initially all edge weights are set to infinity (∞). The weight of edge (i, j) in the graph represent the best objective value found among solutions for which i is visited directly before j. These weights are updated, if necessary, whenever a new ALNS solution is discovered. The choice of which tasks to consider for removal is informed by calculating a score for each task in the current solution by examining the state of the neighbour graph. We determine the score of a given task by summing the weights of arcs leaving both its drop-off and pick-up locations. Tasks with high scores are more likely to be misplaced and can be removed with the goal of reinserting them into superior positions later. Some randomisation is included to ensure we don't only remove tasks with the highest score.

Request historical removal heuristic

This heuristic also uses historical information to remove tasks from the routes, but focuses on vessel routes rather than predecessor-successor relations. The information is stored in a complete, weighted and undirected graph called the 'request graph' where each node represents a task's drop-off or pick-up location. All the edge weights are initially set to zero. The cost of the edge (i, j) is defined as the number of times i and j have been visited by the same vessel in the best B solutions found so far in the search. When a solution better than an existing solution among the top-B

is discovered it replaces it causing the request graph to be updated accordingly. In our computational experiments we use B=20. Task i and task j are considered more similar if they have higher edge weights connecting them. Specifically the similarity of task i and j is constructed from the pairwise connections, $S_{ij} = E_{i+d,j+p} + E_{i+d,j+d} + E_{i+p,j+p} + E_{i+p,j+d}$. More similar tasks can then be removed in the same method as outlined for the Shaw removal heuristic.

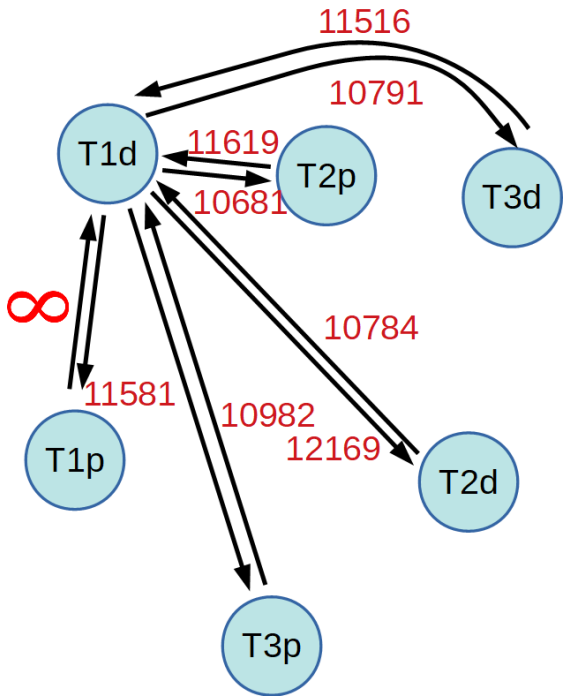

(a) Neighbour removal heuristic. Part of the neighbour graph involving the task at turbine 1. The best solution found so far with a location directly succeeding T1d is T2p with a cost of 10681. (Other edges have been omitted for clarity.)

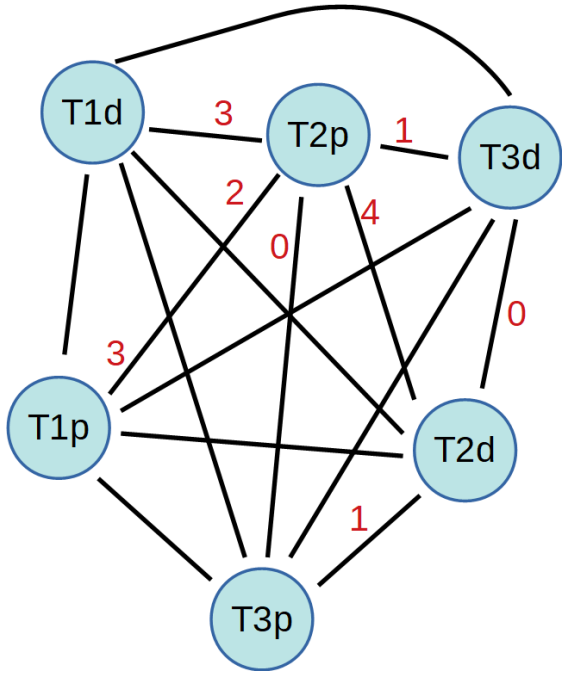

(b) Request historical removal heuristic. Part of the request graph involving the task at turbine 2. The similarity between task 1 and task 2 is (3+2+3+1)=9. (Some edge weights have been omitted for clarity.)

Figure 3.4.2: History based removal heuristics

Smallest objective value increase removal heuristic

This heuristic is designed to remove tasks which provide the lowest benefit in the current solution S so that they can be potentially inserted later in better positions. This often corresponds to removing corrective tasks with the lowest effect on total downtime costs accumulation or preventive tasks with little or no work performed. The cost of removing a task from solution S is defined as $f(S') - f(S)$, given $f(S')$ is the value of the timing sub-problem with task i omitted. This is calculated for every task contained in the routes and requires invoking the timing sub-problem to evaluate every proposed intermediate solution S'. The cumulative aspect of the problem makes it possible that removing a task from the routes lowers the overall cost of the solution. The heuristic then randomly removes a less detrimental task based on the ranked costs of removing each task. The entire process is then repeated until the necessary number are removed, as shown by Algorithm 3.

Data: Solution S containing initial L^* tasks, $\rho \geq 1$

D is the set of removed tasks;

while $|D| < \lceil \rho L^* \rceil$ **do**

$L \rightarrow$ tasks in S;

for $i \in L$ **do**

$Cost^-(L_i, S) = f(S \setminus \{L_i\}) - f(S')$;

end

Sort L such that $i < j \implies Cost^-(L_i, S) < Cost^-(L_j, S)$;

$k \leftarrow \lfloor U(0,1)^3 |L| \rfloor$;

$D \leftarrow D \cup \{L_k\}$;

$S \leftarrow S \setminus \{L_k\}$;

end

return S, D;

Algorithm 3: Smallest objective value increase removal heuristic

Random repair heuristic

We include the naive random repair heuristic primarily as a check that the adaptive operator selection is working correctly. It randomly chooses two locations in a vessel route and inserts the drop-off and pick-up of a given task in an order respecting their precedence. It is not expected to provide good solutions, but can help with diversification and benefits from its speed in being an $\mathcal{O}(1)$ operation.

Sequential greedy repair heuristic

This heuristic is a classic sequential insertion heuristic that seeks to minimise the cost of repairing solutions by inserting tasks into their immediate best position. We define Δz_{vi} to be the cost reduction from inserting the pick-up and drop-off nodes i and $i+n$ of task i into the positions of vessel route v that produce the largest possible decrease in total cost. The maximum cost reduction among the vessels $\max_{v \in V}\{\Delta z_{vi}\}$ is then selected to determine the location for insertions. If a task cannot be inserted in a route or $\max_{v \in V}\{\Delta z_{vi}\} < 0$, it is not added and is returned to the task pool. The complexity of this heuristic is reduced by processing each task in $\mathcal{P}$ in a sequential order and therefore only requires $\mathcal{O}(\mathcal{P})$ operations.

Greedy repair heuristic

This heuristic extends the sequential repair to the classic greedy repair framework. Rather than evaluating tasks in a particular order, the task i with the globally most profitable insertion is selected. This process is repeated until $\mathcal{P}$ contains no tasks with profitable potential insertions and takes $\mathcal{O}(\mathcal{P}^2)$ operations.

Regret repair heuristic

The regret repair heuristic has the advantage of recognising the limited possibilities for task insertions earlier than greedy based heuristics and tries to account for this. We define z^1_{vwi} as a variable that indicates the cost of placing a task in its best position i in its best vessel route and z^2_{vwi} its second best vessel route. A regret value is then calculated for each task i in $\mathcal{P}$ as the difference in cost savings between inserting i in

its best and second best route. The task with maximum regret value is then inserted at its best position. This can be extended to a general regret-κ heuristic by including the cost differences up to the κth best route combination, with the task R_i chosen for insertion.

$$R_i = \arg\max \sum_{j=2}^{\kappa} (\Delta z_{vwi}^1 - \Delta z_{vwi}^{\kappa}) \tag{3.4.9}$$

As a result, requests with a high regret value will be inserted first in an attempt to preserve valuable positions for later tasks. The regret-2 repair heuristic we utilise is computationally expensive relative to the greedy heuristics since it requires $\mathcal{O}(\mathcal{P}^3)$ operations.

3.4.4 Local Search

Following the approach of Li et al. (2016) we include a local search component within the ALNS metaheuristic. This is applied at the end of each segment to steadily improve solutions and keep running times to an acceptable length. A total of $n_{moves} = 3$ moves are randomly selected from the list and applied in the order listed below. Each move explores all the feasible possibilities before choosing the best option. This is repeated until no more improvements can be found. At this point the local search moves to the next type of neighbourhood. The moves outlined below assume that drop-off and pick-up nodes of the same task are visited by the same vessel, but can be generalised if the nodes are split.

- *Intraroute relocate:* the drop-off and pick-up nodes of a task are moved to different positions in the same vessel route.

- *Interroute relocate:* a task is removed from a vessel and reinserted in another.
- *Intraroute exchange:* a pair of tasks within the same vessel route are exchanged.
- *Interroute exchange:* a pair of tasks contained on separate vessel routes are exchanged.
- *Profitable task insertion:* insert tasks from the pool following a greedy procedure if it improves the objective.
- *Profitable task removal:* remove tasks from the solution following a greedy procedure if it improves the objective.

3.5 Computational Results

The geographic layout of an offshore wind farm can have a significant impact on the design of routes for conducting O&M activities. Initial offshore wind farms arranged turbines in a traditional square lattices with a typical inter-turbine separation of around 1km. Existing analysis tends to be performed on wind farms of this shape Dawid et al. (2017), or commonly seen variations such as trapezoidal layouts. Newly emerging wind farms optimise turbine placement based on predicted wind speeds and foundation costs Fischetti and Pisinger (2018), which leads to irregular spatial structures such as that found in Fig. 3.5.1. This structure has been provided by our industrial partners as a realistic example of a future offshore wind farm, which we use as the basis for our experiments. Some turbines have been deliberately omitted and distances rescaled to ensure anonymity. The port is located 45 km south east of the centre of the wind farm.

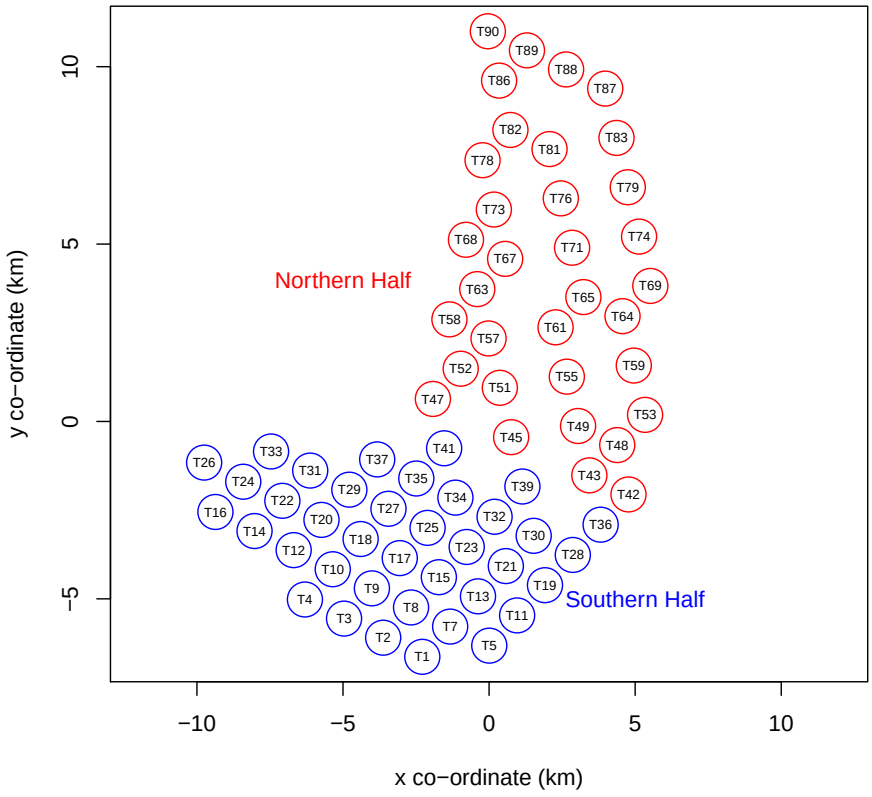

Figure 3.5.1: Example of a future offshore windfarm layout.

We perform computational experiments on an instance set with characteristics defined by Table 3.5.1. Individual task profiles were designed in accordance with previous works and validated with data from our industry partner. We consider instances involving 3-4 vessels, 20-26 tasks and 20-40 technicians. This will allow us to model instances that are resource restricted in terms of having either insufficient numbers of technicians or vessels. We later refer to specific instances as a-b-c, where a is the number of vessels, b the number of tasks and c the number of technicians available.

Corrective maintenance tasks require 2-4 technicians and have typical durations of 180-450 minutes to ensure they can be completed in a single shift. This variety is included to provide a more interesting scenario as all the associated turbines are assumed to be inactive at the start of shift losing potential revenue of £10.8/min. These values are based on a strike price of £60/MWh on a windy day where power output is on average 10MW. Small adjustments to the downtime cost have been made to ensure the model fully prioritises corrective maintenance over preventive tasks. The same analysis could be repeated for summer work days featuring slower winds which would result in a much smaller difference between preventive and corrective task downtime costs.

We can also model corrective tasks which have not caused the turbine to be completely shutdown. The task at T23 is an example where a minor repair has caused the turbine to be generating electricity at a reduced capacity of 72%. A mixture of preventive maintenance tasks are modelled with different prioritisations. These levels are scaled relative to the corrective downtime losses and can be informed by either stakeholders' preferences or wind conditions since revenue is lost only during task execution. The fuel consumption of the crew transfer vessels per minute travelled is approximated to £5/min. Each vessel can transport a maximum of twelve technicians and travel at a speed of 35km/h. Minor alterations are made to the fuel cost of each vessel to differentiate between them and avoid multiple optimal solutions.

Operations are considered during a single shift of twelve consecutive hours, so that all vessels must have returned back to port before the time is elapsed. Time to transfer technicians between vessels and turbines is fixed at 11 minutes independent of

each vessel and task type. Incomplete corrective maintenance activities are penalised with an entire day's downtime costs to reflect the loss of revenue during overnight conditions between shifts. This is achieved by setting $\delta = 2$.

It is important to note that we have deliberately constructed far more challenging instances than in previous works or typically found in practice. The reason for this is twofold. Firstly we believe that as offshore wind farms start to grow in size larger routing and maintenance problems will begin to occur and need to be solved. Secondly larger problems allow us to demonstrate the power of our model to determine the best times to perform a mixture of corrective and preventive tasks in situations without an obvious solution. Problems of this scale will make better use of the ability of the model to not only route vessels and technicians to tasks, but also determine the best combination of tasks to perform within the shift. Typical workloads might allow for 3-4 tasks to be completed on average by a vessel per day. Our system is significantly overloaded with more tasks per vessel than this, in order to guarantee that the resulting solutions use a large degree of task prioritisation. Despite this we do not overload the system only with corrective maintenance tasks in order to ensure a blend of corrective and preventive maintenances tasks are performed.

Turbine	Type	Cost	Technicians	Duration
T11	Prev	3.0	4	420
T86	Prev	3.0	2	420
T3	Prev	5.2	2	420
T20	Prev	5.2	2	420
T32	Prev	3.0	4	420
T67	Prev	3.0	2	420
T24	Prev	3.0	3	420
T78	Corr	10.8	4	180
T55	Corr	10.8	4	240
T73	Corr	10.8	3	450
T45	Corr	10.8	2	180
T19	Corr	10.8	3	180
T81	Prev	5.2	2	420
T59	Prev	3.0	2	1200
T8	Corr	10.8	3	180
T13	Corr	10.8	4	240
T28	Prev	3.0	2	1200
T42	Prev	3.0	3	650
T33	Prev	3.0	2	1200
T9	Corr	10.8	4	240
T1	Corr	10.8	3	180
T23	Corr	7.8	4	240
T68	Prev	3.0	2	1200
T52	Prev	3.0	3	650
T63	Prev	3.0	2	1200
T29	Prev	3.0	4	240

Table 3.5.1: Example task profiles.

A sensible idea would be to benchmark the results of our ALNS based heuristic against the optimal solution to determine the true quality of the heuristic. A simple approach is to solve the mathematical model described earlier in Chapter 3 to optimality using a mathematical solver, and compare the optimal solution with the best objective found by our heuristic. However, as described earlier, we can only compare the results of our heuristic with the optimal solution found from the formulation for very small instances. As a result we decided to not spend much time benchmarking our heuristic's performance compared to the optimal solution. Table 3.5.2 shows a few of the largest instances that we could solve in acceptable computational time.

Tasks	#Vessels	Optimal soln.	Time (s)	ALNS soln.
T11,T81,T73,T45,T8,T55,T78	1	23 333.550	265.1	23 333.550
T11,T81,T73,T45,T55,T78,	2	13 979.291	586.2	13 979.291
T11,T81,T73,T45,T8,T13,T55,T78	1	25 121.691	3154.0	25 121.691

Table 3.5.2: Sample benchmarking instances.

Each row lists the tasks included and the number of vessels for the given instance; all instances included 24 technicians. The last three columns provide the optimal solution, the best solution found by the heuristic after 30 minutes of run time and the time taken to solve the model to optimality using Gurobi in Python 2.7. We see that that for these smaller cases our heuristic successfully reproduces the optimal solution. This is to be expected as the problem instances are fairly small, but does not give

any real insight into how well our method will perform as well for more challenging problem instances. Regardless, our analysis now proceeds on the assumption that our heuristic will produce high-quality solutions for the larger instances that we are actually interested in. Rather than benchmarking against the optimal solution, we will focus on comparing results between different operator selection methods and acceptance criteria.

We compare four different acceptance criteria on our instance set in conjunction with either random operator selection or adaptive operator selection. We evaluate the performance of several variations of the ALNS procedure constructed from alternative pairings of operator selection methods and acceptance criteria. We pair adaptive operator selection and random operator selection with the move acceptance criteria of only allowing improvements, simulated annealing, great deluge acceptance and time dependent simulated annealing. To ensure fair comparisons, all heuristics were run thirty times on the same computer after being implemented in Python 2.7.

$\|V\|$	$\|I\|$	P^{Tot}	Acceptance	Random				Adaptive			
				Best (#)	Avg.	Sd.	Dev(%)	Best (#)	Avg.	Sd.	Dev(%)
3	20	20	classicSA	**29512.87**	30 507.21	451.93	1.31	29 812.89	30 204.37	278.72	2.34
3	20	20	greatDeluge	**29606.16**	30 403.11	422.45	1.64	29 649.15	30 359.85	354.05	1.78
3	20	20	onlyImprove	**29129.83**	30 455.98	519.21	0.00	29 300.78	30 404.61	302.00	0.59
3	20	20	timeDepSA	29 743.64	30 465.87	386.14	2.11	**29674.78**	30 229.56	333.75	1.87
Avg				29 498.13	30 458.04	444.93	1.26	29 609.40	30 299.60	317.13	1.65
3	20	40	classicSA	17 808.99	17 967.08	125.47	0.01	**17807.36**	17 943.16	132.67	0.00
3	20	40	greatDeluge	**17811.22**	18 012.66	363.73	0.02	**17811.22**	18 121.51	412.71	0.02
3	20	40	onlyImprove	**17811.22**	18 068.81	377.33	0.02	**17811.22**	18 007.54	280.98	0.02
3	20	40	timeDepSA	**17807.36**	17 883.45	92.95	0.00	**17807.36**	17 862.19	104.22	0.00
Avg				17 809.70	17 983.00	239.87	0.01	17 809.29	17 983.60	232.65	0.01
3	26	20	classicSA	43 802.75	46 878.57	1765.70	1.21	**43657.44**	45 852.60	1722.70	0.88
3	26	20	greatDeluge	43 918.71	46 636.52	1747.90	1.48	**43278.34**	45 457.06	1845.38	0.00
3	26	20	onlyImprove	43 497.42	47 561.71	2984.81	0.51	**43428.07**	46 935.50	2173.33	0.35
3	26	20	timeDepSA	**43999.88**	48 087.38	1692.91	1.67	44 260.71	47 201.02	1756.23	2.27
Avg				43 761.58	47 282.71	2058.50	1.22	43 656.14	46 361.55	1874.41	0.87
3	26	40	classicSA	**26846.68**	28 286.96	542.76	0.98	27 152.06	28 243.02	456.98	2.13
3	26	40	greatDeluge	26 736.27	27 402.02	440.71	0.57	**26730.78**	27 687.12	570.76	0.55
3	26	40	onlyImprove	26 695.90	27 722.65	782.52	0.41	**26693.59**	27 840.65	671.50	0.41
3	26	40	timeDepSA	26 950.86	27 724.17	460.28	1.37	**26585.69**	27 858.58	581.15	0.00
Avg				26 807.43	27 783.95	556.57	0.83	26 790.53	27 907.34	570.10	0.77

Table 3.5.3: Computational results from three vessel instances using a thirty minute time limit.

$\|V\|$	$\|I\|$	P^{Tot}	Acceptance	Random				Adaptive			
				Best (#)	Avg.	Sd.	Dev(%)	Best (#)	Avg.	Sd.	Dev(%)
4	20	20	classicSA	**29732.66**	30 346.14	346.16	1.66	29 752.03	30 340.84	320.38	1.72
4	20	20	greatDeluge	29 743.98	30 403.70	343.20	1.70	**29534.04**	30 104.66	367.49	0.98
4	20	20	onlyImprove	**29247.94**	30 281.92	339.40	0.00	29 798.63	30 315.96	251.65	1.88
4	20	20	timeDepSA	**29456.99**	30 243.68	437.87	0.71	29 729.29	30 121.70	224.37	1.65
Avg				29 545.39	30 358.77	535.42	1.02	29 703.50	30 220.79	290.97	1.56
4	20	40	classicSA	16 316.83	16 416.51	60.86	0.06	**16308.46**	16 436.93	72.12	0.01
4	20	40	greatDeluge	**16307.27**	16 319.53	22.92	0.00	**16307.27**	16 335.34	43.44	0.00
4	20	40	onlyImprove	**16307.27**	16 332.35	49.53	0.00	**16307.27**	16 333.50	54.27	0.00
4	20	40	timeDepSA	**16306.90**	16 332.21	33.26	0.00	16 309.94	16 346.14	42.72	0.02
Avg				16 309.57	16 350.15	41.64	0.02	16 308.24	16 362.98	53.14	0.01
4	26	20	classicSA	44 259.71	46 409.89	1584.69	3.80	**43343.52**	44 980.78	1449.21	1.65
4	26	20	greatDeluge	43 532.46	46 432.35	1548.89	2.10	**42638.37**	45 802.84	1753.88	0.00
4	26	20	onlyImprove	43 540.90	46 745.03	2363.30	2.12	**43407.47**	46 836.29	2808.77	1.80
4	26	20	timeDepSA	44 186.40	47 401.72	1706.02	3.63	**43902.57**	46 520.60	1788.50	2.96
Avg				43 879.87	46 747.20	1800.72	2.91	43 322.98	46 035.13	1950.09	1.61
4	26	40	classicSA	**23526.64**	24 126.81	302.21	1.11	23 737.29	24 155.35	254.21	2.02
4	26	40	greatDeluge	**23267.88**	23 729.39	365.48	0.00	23 292.04	23 588.54	323.12	0.10
4	26	40	onlyImprove	23 304.13	23 875.39	430.36	0.16	**23267.76**	23 827.74	430.68	0.00
4	26	40	timeDepSA	23 334.47	23 685.78	196.04	0.29	**23318.83**	23 648.48	200.81	0.22
Avg				23 358.28	23 854.34	323.52	0.39	23 403.98	23 805.03	302.20	0.59

Table 3.5.4: Computational results from four vessel instances using a thirty minute time limit.

The results of our experiments are summarised in Tables 3.5.3 and 3.5.4. The

initial column headings refer to the number of vessels available, the total number of maintenance tasks, the number of technicians present and the acceptance criteria used. Results are defined for both the random operator selection and the adaptive operator selection methods. We specify the best solution found in thirty runs of the instance (column Best), the average solution found (column Avg.), the standard deviation among the results (column Sd.) and the deviation (column Dev %). This measures the relative quality of results without the knowledge of the true optima and is calculated as Dev(%)=100(Best - $\underline{\text{Best}}$)/$\underline{\text{Best}}$, where $\underline{\text{Best}}$ is the best known solution obtained by any of the acceptance criteria operator selection combinations for a given instance. The best values found across selection methods are highlighted in boldface. We also visualise the spread of the objective values for each combination in the histograms presented in Figures 3.5.2 and 3.5.3.

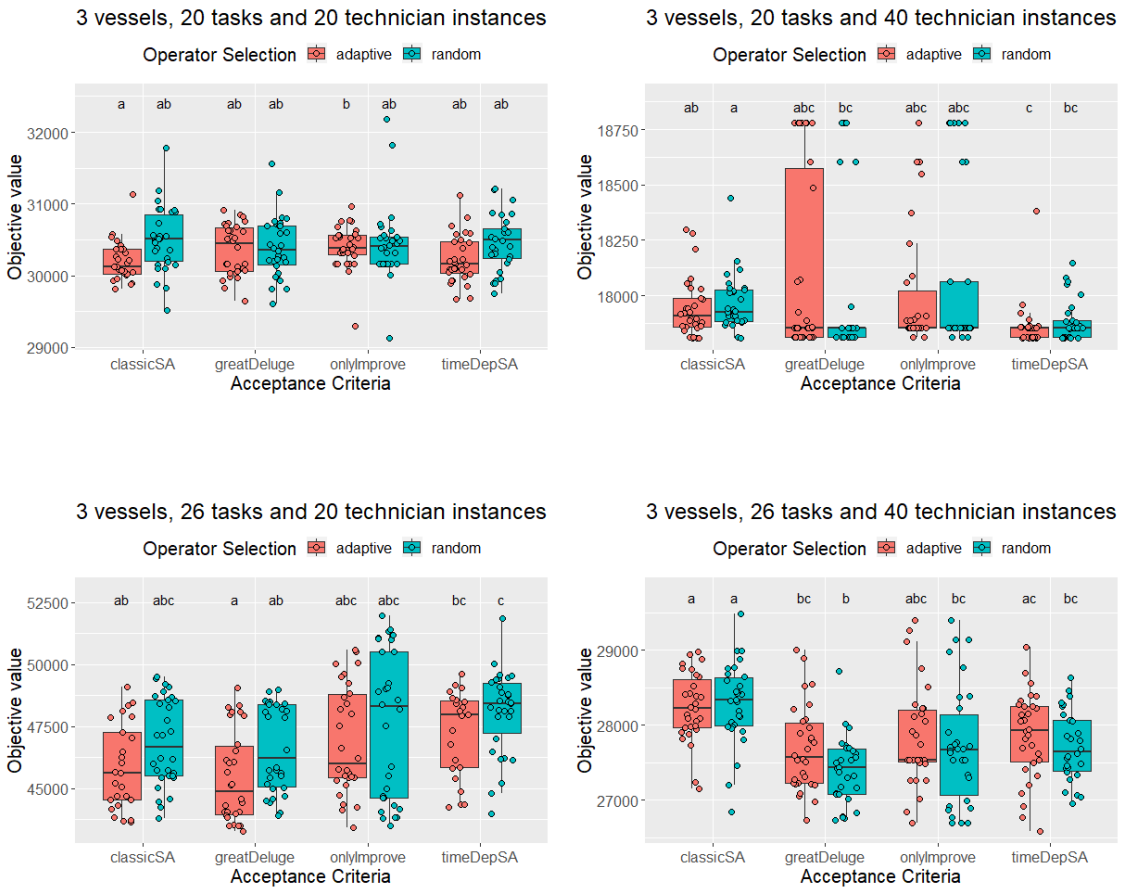

Figure 3.5.2: Histograms of experimental results for three vessel instances.

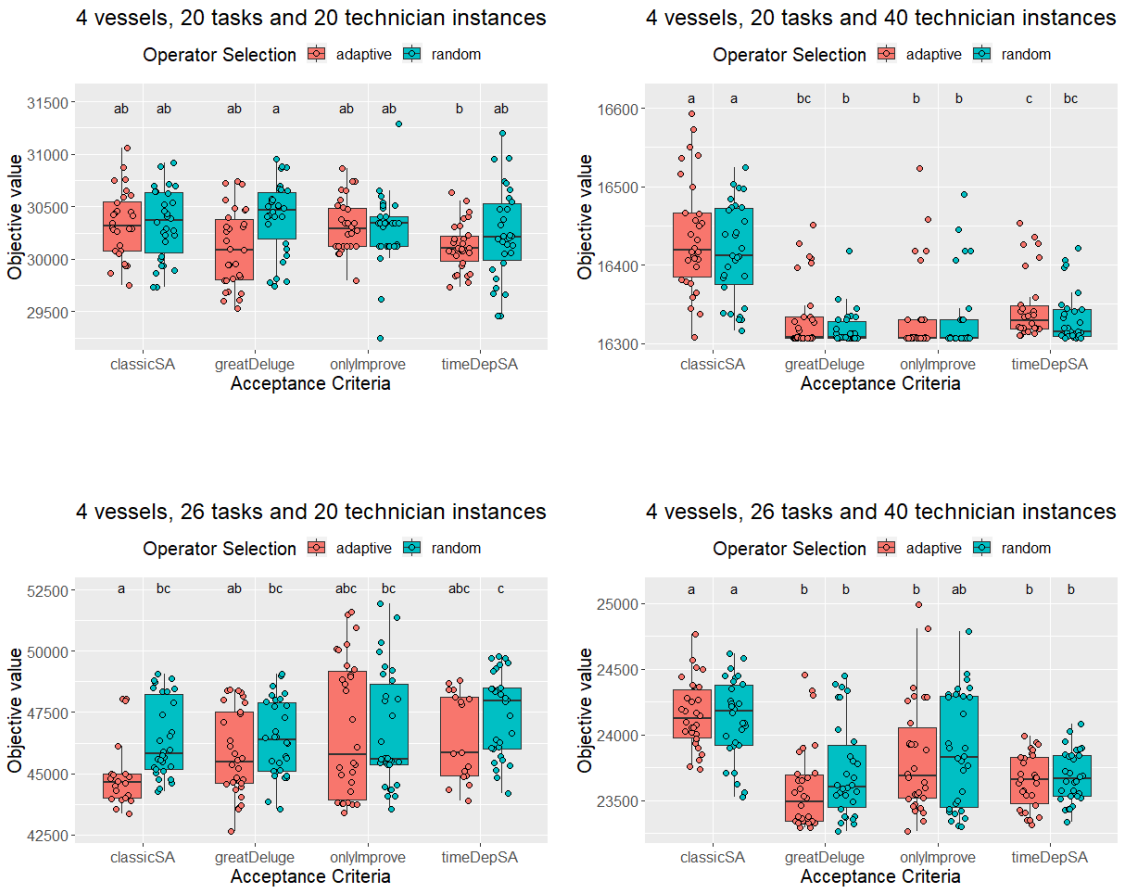

Figure 3.5.3: Histograms of experimental results for four vessel instances.

For both of the instances containing 20 tasks and 40 technicians all the tested combinations were able to produce very similar results. Indeed the best results found for the 3-20-40 instance are all within four units of each other with many combinations all reaching the same objective value. We believe that these instances are less complex than the others since they contain the lowest ratio of work volume to resources. This effectively reduces the complexity of the problem since the model no longer has to solve the knapsack portion of the problem as the task prioritisation needed is clear. The fact that the average deviations never exceed 0.02% adds further weight to the idea that the

methods all converge to a similar high quality solution. Interestingly it appears that classic simulated annealing performs significantly worse than the other combinations on the 4-20-40 instance in terms of its average performance. Notice however that since the problem is relatively simple the absolute difference in objective values is only 300. We therefore conclude that there is little benefit to using an adaptive selection method in these scenarios.

The adaptive operator selection - great deluge acceptance criteria combinations appears to perform the best for the 3-26-20 task instance. It provides both the best solution found (43278.34) and the lowest average objective (45457.06) for the three vessel version. The best results of other combinations have noticeable deviations ranging from 0.35% for only accepting improving solutions and 2.27% for time dependent simulated annealing. Standard deviations are roughly similar with the not unexpected exception of the random selection method and only accepting improving solutions combination. We also observe that the standard deviations for all the combinations are around 3.5 times bigger than in the corresponding case with 40 technicians indicating that this is a more difficult instance. Furthermore the greatest difference in average objective values across the random and adaptive selection methods (921.17) occurs for this instance suggesting that adaptive selection holds the highest benefits for our most challenging resource restricted instance.

These findings also hold for the four vessel extension with great deluge acceptance again providing the best ultimate solution. Adaptive operator selection performs the best for every acceptance criteria and therefore has a lower average deviance of 1.61% versus 2.91%. We note that the gaps between average performance and best

solution are larger in this instance when compared with the 3-26-20 instance despite the overall objectives being similar. Interestingly the best average performance occurs for adaptive selection within classic simulated annealing rather than the great deluge acceptance.

Surprisingly, random operator selection with only accepting improving solutions appears to hold its own against the other methods on the 3-20-20 and 4-20-20 instances. It provides the best solution found for both instances, albeit with only one or two runs performing significantly better than the other methods. This may be due to the presence of the local search component potentially dominating the behaviour of the removal and repair heuristics.

Results for the 3-26-40 and the 4-26-40 instances are less conclusive and hence are harder to draw meaningful conclusions from. In the 3 vessel instance the average performance is mixed and often contradictory to the best solution. Some combinations exhibit better performance with random operator selection and others with adaptive operator selection. For example, the best solution comes from adaptive selection with time dependent simulated annealing, but it has a worse average performance than five other combinations. This ambiguity is highlighted in the average performance being very similar between adaptive and random operator selection. We note that classic simulated annealing acceptance criteria appears to perform worst across all the tested combinations. Adaptive selection method with great deluge has the best average performance (23588.54) for the 4-26-40 instance, but is beaten to the best ultimate performance by both random selection with great deluge acceptance and adaptive selection with only accepting improving solutions. We again observe contradictory

results between most operator selection acceptance criteria combinations. However the performance of simulated annealing is again worse than the other combinations in accordance with the 4-20-40 instance with the additional caveat that its best solution fails to come close to the other combinations.

We introduce statistical hypothesis tests as an additional step in our analysis as a measure for detecting whether any relative performance difference can be considered significant. We first utilise the Kruskal-Wallis H test (Kruskal and Wallis, 1952) as a non-parametric test to detect whether there is least one significant difference between selection method-acceptance criteria combinations. In other words, it will check if the choice of selection method-acceptance criteria actually has an impact on the objective value found at the end of the ALNS. A p-value below 0.05 indicates that there is evidence to reject the null-hypothesis: that the objectives found by all selection method-acceptance criteria originate from the same underlying distribution. Our findings listed in Table 3.5.5 indicate significance for every instance at a 95% confidence level. This confirms the need for further analysis, since the selection method-acceptance criteria does impact the quality of the ALNS solution.

Instance	p-value	Reject?
3-20-20	0.005 21	✓
3-20-40	9.467×10^{-5}	✓
3-26-20	1.663×10^{-5}	✓
3-26-40	6.637×10^{-9}	✓
4-20-20	0.002 355	✓
4-20-40	$<2.2 \times 10^{-16}$	✓
4-26-20	0.000 454 4	✓
4-26-40	1.563×10^{-12}	✓

Table 3.5.5: Krusal-Wallis H test results for detecting a significant difference between combinations.

We next examine our results to determine if adaptive operator selection is superior to random operator selection by pooling the results for each instance and performing a Mann-Whitney U test for significance. We use the Mann-Whitney U test (Fagerland and Sandvik, 2009), (Kruskal, 1957) as a non-parametric statistical test to detect whether two independent samples are drawn from the same underlying distribution. Our results are summarised in Table 3.5.6 based on a 95% confidence interval. It indicates that the performance of adaptive operator selection is statistically superior to the random operator selection for the 20 technician instances. The null hypothesis

is confirmed for the 3-20-40 instance with the largest p-value of 0.9978 which validates the idea that both selection methods have converged to a similar local or near optimal minima, since it is conceptually the simplest instance. Results are less clear for the remaining instances.

Instance	p-value	Reject?	Comments
3-20-20	0.003 179	✓	Adaptive is superior; median of 30268 vs 30443.
3-20-40	0.9978	✗	-
3-26-20	0.000 839	✓	Adaptive is superior; median of 45847 vs 47902.
3-26-40	0.1045	✗	-
4-20-20	0.011 55	✓	Adaptive is superior; median of 30190 vs 30347.
4-20-40	0.1481	✗	-
4-26-20	0.001 422	✓	Adaptive is superior; median of 45264 vs 46289.
4-26-40	0.2539	✗	-

Table 3.5.6: Mann-Whitney U test results for comparison between adaptive and random operator selection.

We also investigate whether a particular acceptance criteria is dominant for each instance by performing pairwise Wilcoxon-Rank Sum tests between all the combinations in order to detect which specific operator selection-acceptance criteria results are significantly different from each other. As multiple comparisons are needed an adjustment to the p-value threshold for significance is required to maintain a low

overall type-I (false positives) error. We use the 'Bonferroni correction' which divides the alpha value α for each comparison by the number of comparisons $n = 28$. Thus our adjusted p-values will be approximately 0.001786.

Furthermore we classify combinations into a graphical structure denoted by letters according to the significant differences detected as described by Piepho (2004). Structures containing the same letter are not considered significantly different, whilst groups not containing any mutual letters are. Note that transitivity does not hold, namely if a is significantly different to b and b significantly different from c it does not imply that a is also significantly different from c. The results of this graphical method of representing pairwise significant differences is found at the top of each graph in Figures 3.5.2 and 3.5.3. We choose to interpret these significance/non-significant relations as quasi-groups, which can be ranked or further split in conjunction with the performance metrics outlined earlier.

We broadly observe few statistical differences between operator selection-acceptance criteria combinations on most of the instances. For example the 3-20-20 instance only has a single significant difference between adaptive operator selection with classic simulated annealing and only accepting improving solutions acceptance criteria. The statistical tests unsurprisingly indicate that adaptive great deluge and the only accept improving combinations perform similarly given that their upper quartiles and tails are the largest in the 3-20-40 instance. We observe that classic simulated annealing performs significantly different to the majority of other combinations on the 3-26-40 and 4-26-40 instance so we conclude that its performance is definitely worse.

An example where the pairwise tests add value is the 3-26-20 instance. This in-

stance presents several significant differences using structures represented by three letters a,b and c. It is clear that the time dependent simulated annealing acceptance methods (bc, c) are statistically worse than the great deluge acceptance with adaptive operator selection (a). If we incorporate additional metrics such as average performance then we can rank quasi-groups which perform better but below statistical significance. This creates a pseudo-ranking of $a > ab > abc > bc > c$ which implies that a is the most valuable structure. This matches our earlier belief that adaptive operator selection with great deluge acceptance performs the best on this instance.

Furthermore, adaptive operator selection beats random selection on every performance criteria (Mann-Whitney U test, best solution, lowest average, smallest deviation for each acceptance criteria) in the 4-26-20 instance, whilst the three vessel version is only denied this by two solutions. We therefore conclude that the benefits of adaptive operator selections are greatest on the most complicated instances. It should be noted that the inclusion of a local search mechanism within our ALNS procedure may be responsible or interfere with the analysis of these results.

3.5.1 Analysis of ALNS Procedure

The key advantage of adaptive operator selection over random selection lies in its ability to self adjust to use the more successful heuristics repeatedly within the ALNS framework. We determine the average utilisation percentage of the operator combinations from the relative frequency of combinations used during the search procedure. Random operator selection will produce utilisation percentages that are static and equal to 1/n, where n is the number of combinations. Our adaptive operator selection

allows these percentages to change in favour of more successful and efficient heuristics. The degree of success will vary in accordance with the score obtained and the operators' duration which together act as a proxy for information learnt per unit time.

We present this information in Figure 3.5.4 for an arbitrarily chosen run of the 3-26-20 instance. The inner ring of Figure 3.5.4a represents the utilisation percentage of the removal operator and the outer ring the utilisation percentage of the repair operator. Operators are listed in clockwise descending order of their utilisation percentage. Figure 3.5.4b shows the same information but with the contents of the rings swapped. This allows us to display the conditional percentages of operator combinations with the absolute percentages listed on the larger fractions for clarity.

We observe that some combinations are clearly boosted or reduced relative to random operator selection as we would not expect all of the removal and repair operator combinations to necessarily generate high scores per unit time. Smallest objective value increase removal is used in 25% of the iterations with request historical removal, neighbour removal and price similarity also having utilisation percentages larger than 14%. The heuristic still makes use of the remaining removal operators as a diversification technique with random removal employed 6% of the time, but effectively ignores the expected worst combination of random removal and random repair. We note that request historical removal was paired with random repair 6% of the time perhaps indicating that some of this score was achieved by chance near the start of the search procedure.

Figure 3.5.4b shows that the most heavily utilised repair operator was the sequential randomised greedy repair heuristic which appears to achieve a large score. This

score is effectively increased in our time dependent setting since it consumes significantly less computational time than full greedy repair or regret repair. We observe that regret repair has a higher utilisation percentage (28%) than greedy repair (20%) indicating that it manages to accumulate a large enough score to offset its comparatively long running time. The random repair heuristic is used 9% less often than would be expected under random operator selection. Even though the heuristic is extremely quick it becomes less favoured since it does not discover new solutions very frequently. Interestingly the smallest objective increase removal heuristic was most often paired with the three key repair heuristics with relative utilisation percentages of 27.8%, 28.6% and 25.0% respectively. These are approximately double the rate of their usage under random operator selection. The next set of operator combinations with boosted utilisation percentages relative to random selection vary dependent on repair heuristic but include neighbour removal, price similarity removal, Shaw removal and request historical removal.

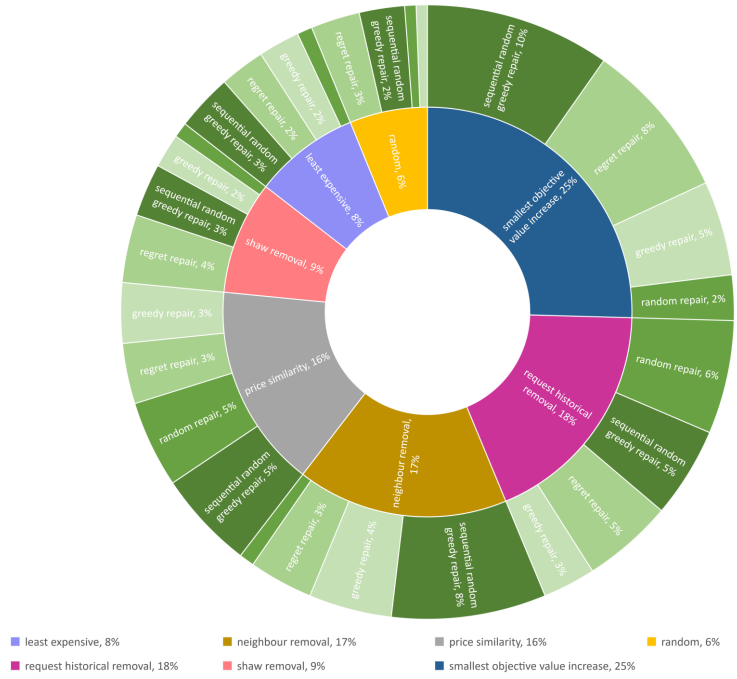

(a) Utilisation percentage of removal-repair combinations for adaptive operator selection.

(b) Utilisation percentage of repair-removal combinations for adaptive operator selection.

Figure 3.5.4: Utilisation percentage of removal and repair operators during adaptive large neighbourhood search.

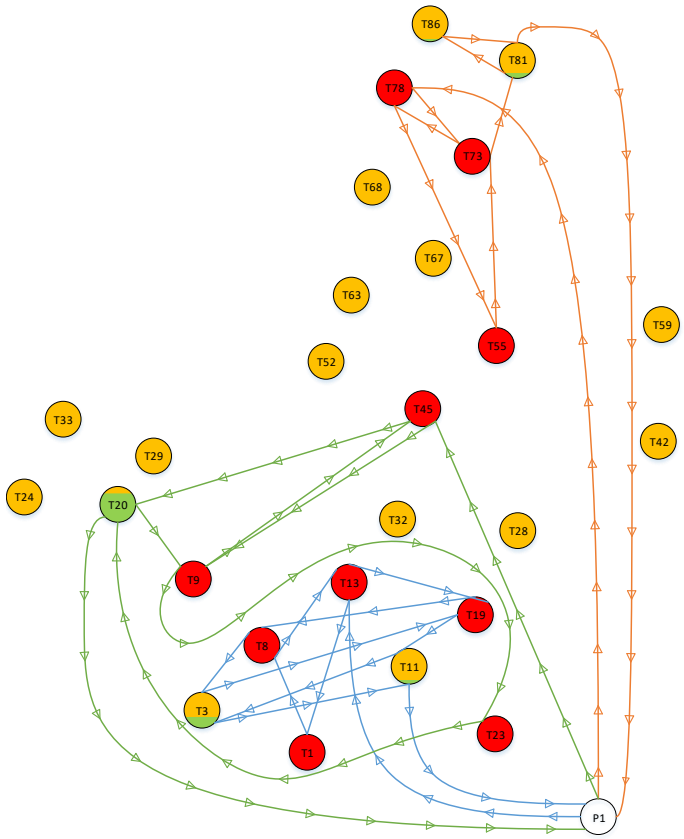

Figure 3.5.5: Illustration of an example vessel route produced for a 3-26-20 instance. Corrective maintenance tasks are coloured red. Preventive tasks are coloured orange with green added in proportion to the percentage of work completed.

V1	Nodes	P1	T13d	T1d	T1p	T8d	T13p	T19d	T8p	T3d	T19p	T11d	T3p	T11p	P1
	Time	0	90	105	285	299	330	345	479	494	585	598	622	643	720
	#Techs	7	3	0	3	0	4	0	3	1	5	1	3	7	7
V2	Nodes	P1	T78d	T73d	T78p	T55d	T55p	T73p	T81d	T86d	T86p	T81p	P1		
	Time	0	96	110	276	299	539	560	575	590	623	638	720		
	#Techs	7	3	0	4	0	4	7	5	3	5	7	7		
V3	Nodes	P1	T45d	T9d	T45p	T20d	T9p	T23d	T23p	T20p	P1				
	Time	0	90	112	270	293	352	369	609	631	720				
	#Techs	6	4	0	2	0	4	0	4	6	6				

Table 3.5.7: Details of the vessel routes for Figure 3.5.5 including the arrival time to each location and the number of technicians onboard the vessel when departing the location. Arrival times are rounded to the nearest minute.

Figure 3.5.5 shows the set of vessel routes generated from a run of the ALNS procedure with adaptive operator selection and great deluge acceptance on the 3-26-20 instance. The specific properties of the routes are listed in Table 3.5.7 including the locations visited, departure times and number of technicians onboard the vessel when departing the location. We see that the first vessel modelled with the lowest fuel costs is required to visit six turbines, the most of any vessel. The route prioritises corrective maintenance tasks ahead of the preventive maintenance tasks with technicians being required to work on multiple turbines during the shift. A small degree of opportunistic maintenance is performed on the preventive tasks at T3 and T11 with the limited time remaining in the shift. Meanwhile the second vessel takes seven technicians to deal

with a group of tasks in the more distant northern half of the offshore wind farm. The solution suggests to complete the task at T78 and start the task at T73 with the transported technicians. The vessel is instructed to wait alongside T55 until its corrective maintenance is completed before moving onto opportunistic preventive maintenance. This behaviour is a natural consequence of the costs found in our instances since the corrective downtime losses dominate both preventive maintenance benefits and transportation costs. Once preventive task benefits increase we start to observe solutions that deviate from this behaviour. This can be seen in the third vessels route which drops technicians at the preventive maintenance task at T20 before relocating the technicians from T9 to complete the corrective task at T23.

We notice that the model has chosen to omit many of the preventive maintenance activities included within the model. This prioritisation allows for a better utilisation of resources with technicians spending 61.9% of the shift working on tasks. The downtime losses from corrective maintenance tasks were calculated as 43272.972 with the smaller benefit of preventive maintenance tasks being 2989.676. The transportation costs are in line with the results obtained by Irawan et al. (2019) with a cost of 3212.86 representing 6.5% of the total costs compared to 5%. We cannot compare the corrective and preventive maintenance percentages since we include significantly more corrective maintenance tasks in our dataset. We note that a saving of 1% in the overall objective value approximately corresponds to either a 14% drop in the effective routing costs or a 15% increase in preventive maintenance benefits assuming the same corrective downtime losses are accrued.

Table 3.5.8 describes the completion times for corrective maintenance tasks and

the time spent working on preventive maintenance tasks for the twenty technician instance and a less restricted forty technician instance. We observe that almost all of the omitted preventive tasks have large percentages of work performed on them often in excess of 50%. Interestingly the corrective maintenance tasks also tend to be completed earlier suggesting that the benefits of the additional technicians is not limited to performing extra preventive maintenance after all the corrective tasks. Note that strict improvements in completion time or working time are not always guaranteed. Sometimes it is beneficial to slightly delay the pick-up activity of a corrective task compared to the twenty technician instance. T73 is an example of this myopic behaviour with its completion time occurring 76.5 minutes later.

Task	Completion Time		Task	Duration Worked		Task	Duration Worked	
N^{Corr}	$P^{Tot}=20$	$P^{Tot}=40$	N^{Prev}	$P^{Tot}=20$	$P^{Tot}=40$	N^{Prev}	$P^{Tot}=20$	$P^{Tot}=40$
T19	585.0	326.5	T28	0.0	327.1	T52	0.0	0.0
T23	608.8	541.7	T67	0.0	11.0	T3	128.4	420.0
T73	559.6	636.5	T59	0.0	327.6	T32	0.0	75.2
T13	330.0	365.2	T33	0.0	208.4	T24	0.0	182.0
T1	285.4	272.8	T86	32.4	367.9	T20	338.5	420.0
T45	269.9	284.7	T63	0.0	353.9	T29	0.0	0.0
T8	478.9	286.3	T11	45.2	234.8			
T9	351.9	359.8	T81	63.3	420.0			
T55	538.7	327.6	T68	0.0	353.9			
T78	276.1	290.2	T42	0.0	304.4			

Table 3.5.8: A comparison of corrective task completion times and preventive maintenance task work durations between a 20 and a 40 technician instance.

3.5.2 Impact of Resource Availability

We now examine the impact of adding additional resources such as technicians and vessels into the OWFMRP. Figure 3.5.6 illustrates the relative saving in the optimal cost in the different instances normalised by the average value found from the instance containing 3 vessels, 26 tasks and 20 technicians. This was chosen as it is the most expensive and complex instance on average. We note that some runs lie below 0%

since there are obviously more expensive solutions than the average case. As our system contains uncompleted corrective maintenance tasks in the 26 task instance there is an approximate reduction of 35% in costs when six of the additional tasks are omitted in the 20 task instance version. The bulk of this saving comes from the lowered downtime costs from finishing corrective maintenance tasks earlier within the shift. We observe that the savings from doubling the number of technicians present is significantly larger than those from adding an extra vessel, an expected result since our instances are effectively limited by the number of technicians rather than vessels. Each vessel has a technician capacity of 12, so the 3 vessel instances can use up to 36 technicians with the additional technicians. The 80% increase in technicians can be fully utilised as potential tasks remain in the system. Furthermore adding technicians in conjunction with a fourth vessel into the system allows us to make use of the 4 idle technicians in the 3 vessel instances to harness further savings in optimal cost of achieve a 49% saving in the 26 task instance.

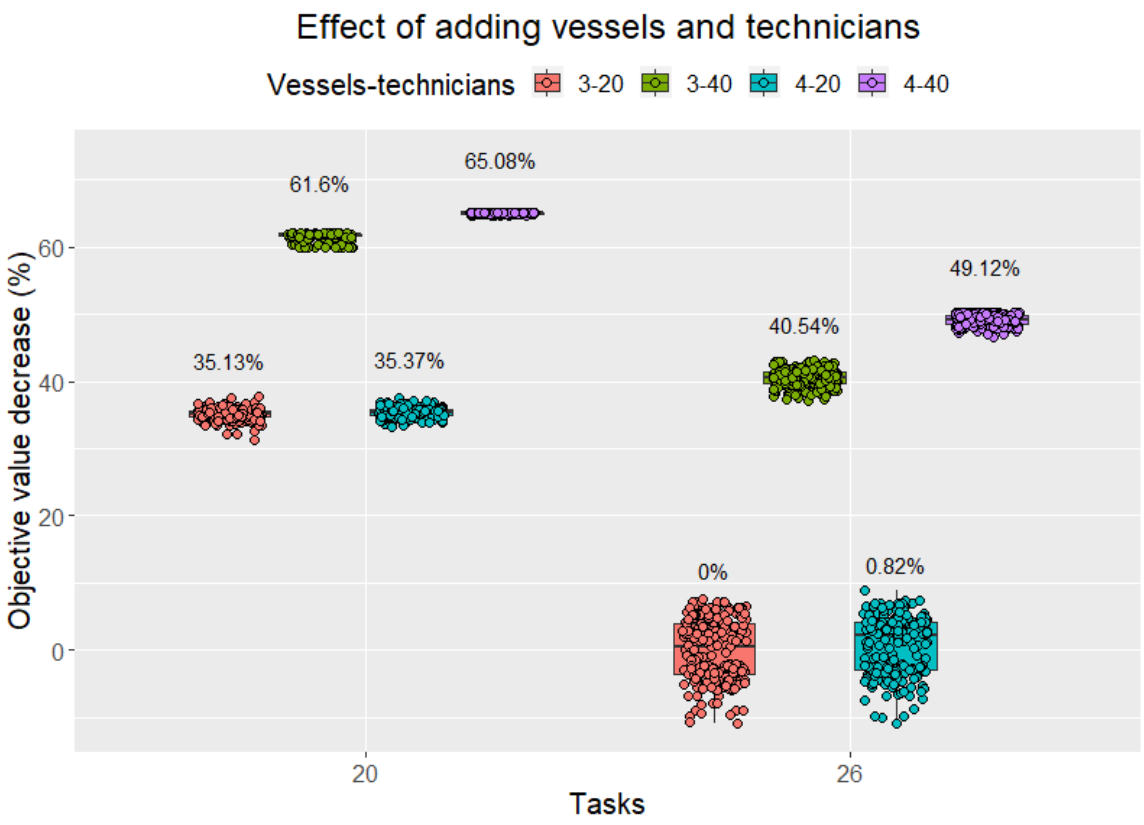

Figure 3.5.6: Effect of adding extra resources into task instances.

We note that the analysis and discussion in this section is operating on the assumption that the heuristic yields solution close to optimality. Also that the optimality gaps of the solutions are roughly the same.

3.6 Conclusions and Future Work

We have presented a mathematical formulation for the offshore wind farm maintenance routing problem involving both corrective and optional preventive maintenance activities. We propose a novel decomposition structure to decouple the model into a timing sub-problem and a master routing problem to reflect the intricacies of downtime cost

calculations. An adaptive large neighbourhood search procedure is developed to solve a set of tough real world instances which reflect undertaking operations in the midst of possible resource restrictions. A variety of performance measures and hypothesis tests were employed to determine the relative performance of a variety of acceptance criteria in conjunction with either random or adaptive operator selection methods. Our results indicate that there is no clear winner for the simpler instances but that adaptive operator selection performs better on the more challenging personnel limited instances. Evidence exists to suggest that adaptive operator selection combined with great deluge acceptance was the best choice for several instances such as the 3-26-20 case.

The full power of our decomposition approach could be extracted with a deeper examination of the potential for pick-up and drop-off nodes to be visited by different vessels. The inclusion of multiple concurrent tasks with technician skill types fully differentiated would further expand this framework. Another goal would be to quantify and incorporate the long term benefits of preventive maintenance into a short term horizon which is mainly focused on corrective maintenance. Condition based maintenance strategies that monitor individual component degradations provide significant motivation for further research in this direction. Our model can be extended to consider the challenge of routing larger Service Operation Vessels (SOV) which can remain offshore for several shifts.

Chapter 4

Developing a Statistical Model of Metocean Conditions

4.1 Introduction

4.1.1 Motivation

Key offshore activities such as offshore wind turbine installation or complex repair tasks require vessels and technicians to spend significant periods of time working in uncertain offshore environments. In order to avoid wasteful offshore trips and potentially unsafe working conditions detailed planning and scheduling is needed to minimise the likelihood of unnecessary costs being incurred. Furthermore, the offshore conditions also directly impact on the revenue of the offshore wind farm given that the turbines are powered by the wind which heavily contribute to its overall economic performance.

Future technological developments will likely seek to hybridise multiple data sources such as turbine degradation and wave conditions to plan and automate operations and maintenance tasks. This will increase the need for decision making based on a wide array of possible metocean scenarios. Accurate judgements can only be made when uncertainties are incorporated from a large number of metocean time series, particularly given the existence of key operational limits on parameters such as offshore windspeed W_s, and significant wave height H_s. The knowledge of the possible evolutions of offshore conditions can then be used in conjunction with improved vessel and technician routing to fully extract potential savings from efficiently performing operations and maintenance activities. As a result methods are needed for generating alternative realisations from a relatively small pool of historical data and hindcasting.

Our overall goal is to create a statistical model that is capable of generating scenarios, lasting up to a fortnight in length starting from any future time point. This model should be able to synthesise several inputs together to produce a prediction of future conditions that includes uncertainty. If the time point of interest is within two weeks it is important to utilise the recent metocean conditions and current seasonal trends, as they both play a role in generating future metocean patterns. In contrast for time periods further into the future only statistical methods are of use. Here we focus on the latter case but our discussion in Chapter 5 outlines how our approach can be adapted for the former case. A statistical approach requires modelling the joint distribution of the metocean variables from a temporal viewpoint. This joint distribution could vary with seasonal trends with possible realisations created by simulating from it used to aid decision making.

4.1.2 Weather Windows

The most important factors in offshore activities are the operational limit of tasks and of the vessels used to transport personnel to turbines. If weather conditions exceed the operational thresholds, then vessels cannot travel safely and repair activities cannot be executed thus leading to the accretion of significant downtime costs. An operation can only be undertaken if the conditions remain calm for at least the duration of the repair or a specified amount of time, often referred to as a weather window, Anastasiou and Tsekos (1996). A weather window is technically defined as a number of successive multivariate observations (e.g. wind speed, significant wave height etc.) all of which concurrently lie below their respective thresholds, with its persistence defined as the length of time it remains below the threshold. Weather windows are always defined to be as large as possible. This means that twelve consecutive hours of conditions always below a threshold will be considered a single weather window of 12 hours, rather two consecutive windows of 6 hours or 3 windows of 4 hours.

We also examine the frequency and persistence of the 'waiting time' before a weather window occurs. It is defined as the number of successive multivariate observations which have at least one of their marginal variables taking a value above the threshold. This is the converse of a weather window since it requires at least one of the environmental parameters to constantly exceed its threshold, rather than every observation lying below it. As a result every data point will be classified into the state of either being in a weather window or waiting for a weather window. As an offshore wind farm will only be considered accessible during a weather window of

sufficient duration it will important to accurately reproduce their distribution in any model, and hence by any synthetic dataset generated under that model. Accurate planning of operations and maintenance activities will require a good representation of the distribution of waiting times until subsequent weather windows as this directly impacts lost revenue if tasks cannot be completed beforehand.

Leontaris et al. (2016) attempt to determine windows of time during which offshore operations can be completed. In the context of offshore wind farm installation this requires modelling wind speeds and significant wave heights. They follow a copula approach that constructs the joint distribution of metocean variables with knowledge of the dependence structure and autocorrelation built in. Seasonal trends are accounted for by analysing the data on a monthly basis. A simple vine copula is used to model the dependence structure. The autocorrelation within the wind speed is captured through a Gaussian copula, whilst the dependence between wind speed and significant wave height is modelled with a Gumbel copula for the majority of months. As a result they generate continuous values of metocean parameters which can be evaluated for multiple potential weather thresholds.

Bruijn et al. (2019) take an alternative approach to modelling the weather windows for a series of sequential offshore operations. They introduce a discrete time 2-state Markov chain model representing whether an operation can be worked on or not in a period of time. The transition rates can either be assumed to be constant within predetermined periods such as months or seasons or modelled as time-dependent. In this case kernel density estimation is employed to place more weight on observed transitions close to the time of interest and correspondingly less weight on more distant

transitions. A first-order Markov chain is incorporated to replicate the persistency of metocean conditions. De Masi et al. (2015) builds a first order multivariate Markov chain model of offshore wave heights using a triple state of significant wave height, peak wave period and mean wave direction. Their work is calibrated on the autocorrelation of wave heights and the persistence of weather windows and the waiting time before a weather window.

A drawback of the model provided by Leontaris et al. (2016) is that data is split into monthly blocks with the environment parameters allowed to have different statistical parameters from block to block. This leads to unrealistic discontinuous jumps between the end and start of successive months. We will utilise a rolling window of local observations to ensure smooth transitions in statistical parameters over time. Furthermore whilst the empirical cumulative distribution function does capture non-trivial patterns, it will not be able to give detail between the ranks of data. As this transformation is eventually employed in reverse to reincorporate seasonality, it will be unable to accurately generate observations at higher precision than those found in the original data. Our use of kernel density estimation will better account for seasonal effects and allow for new previously unobserved observations to be generated. We will also show that their assumed graphical dependence structure in which the wave height is driven by the current windspeed, which itself is driven by the previous windspeed, is inappropriate for our dataset. We find that it is important to condition on multiple previous lags as opposed to a single time step.

Our work differs from the Markov chain based approaches as described by the likes of Bruijn et al. (2019) as it generates a continuous model of metocean conditions rather

than a discrete set of states. This means that an equivalent workability distribution based on a threshold u_2 instead of u_1 can be automatically generated from a single synthetic dataset as opposed to needing to rerun and calibrate a discrete Markov chain model. It is also likely that a first order Markov chain is too simplistic a model to fully capture all the features of metocean conditions. This is illustrated in Fig. 4.1.1. Another advantage is the potential for sensitivity analysis for observations close to the threshold, given in practice there is likely to be some uncertainty surrounding measurements.

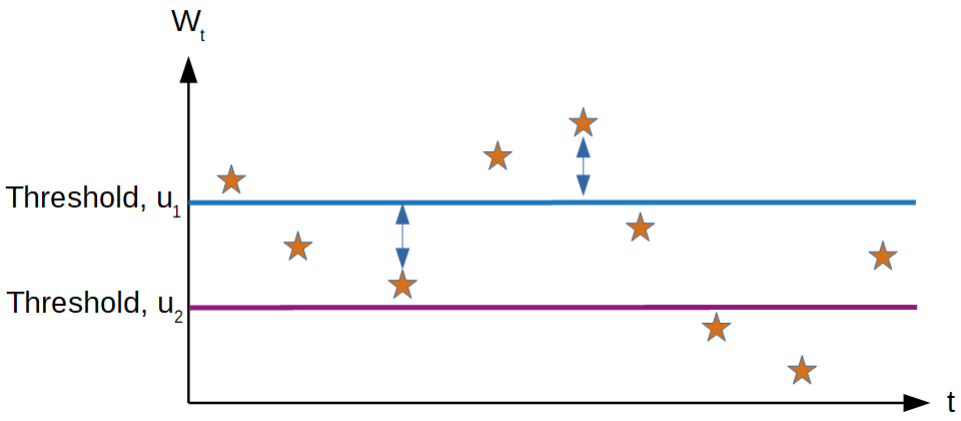

Figure 4.1.1: Weather windows based on different thresholds are easily deduced from a continuous model of metocean conditions. Sensitivity analysis can also be explored to determine how close weather windows are to being shortened or extended.

4.1.3 Model Overview

The methodological approach described in this section utilises several statistical concepts to simulate stochastic offshore time series that preserve key properties of the original data. The foundation of our statistical model is a multivariate Markov model.

We will assume that the underlying structure can be reproduced with a Markov chain model; a sequence of random vectors $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_N$ which follow a kth order Markov property, where $\underset{\sim}{\mathbf{X}}_t = (X_{1,t}, \ldots, X_{d,t})$ is the vector of d environmental parameters. This requires that the probability of moving to a future state $\mathbf{X}_{t+1}$ in any set A conditioned on all the previously visited states is exactly the same as the probability conditioned on just the previous k observations. In other words, the choice of the next state is influenced only by the last k states. This is expressed probabilistically by

$$P(\mathbf{X}_{t+1} \in A | \mathbf{X}_t = \mathbf{x}_t, \mathbf{X}_{t-1} = \mathbf{x}_{t-1}, \ldots, \mathbf{X}_1 = \mathbf{x}_1) = P_t(\mathbf{X}_{t+1} \in A | \mathbf{X}_t = \mathbf{x}_t, \ldots, \mathbf{X}_{t-k+1} = \mathbf{x}_{t-k+1}), \tag{4.1.1}$$

for all $t > k$ and all sets A.

The current Markov model described in Eq. 4.1.1 is a time dependent process, namely that the transition probabilities between states vary depending on the time of year. This is emphasised through the probability P_t on the right hand side being a function of t. The presence of temporal non-stationarity is a reflection of the long term seasonality in the metocean conditions with winter conditions typically being more severe than those in summer. In order to remove the dependence of P_t on t we seek to transform the data to make them stationary. Our reason for doing this is that data from any t will be then informative about any other time t', which will help give efficient inference. In particular, a stationary time series ensures that all of its statistical properties such as mean, variance and other higher order moments can be considered constant over time. Ensuring that our time series have the stationarity

property is important since it allows the assumption that any two sub-blocks of the series follow the same underlying model. Transforming the data to be stationary effectively increases the sample size for any analysis leading to greater accuracy in results. Furthermore any block of deseasonalised and stationary data can be resampled without restrictions and used to create a new synthetic time series by adding in the necessary seasonality.

As the rate and scale of seasonality is unique to each metocean variable, we first adjust the individual marginal distributions to account for their seasonal effects. One option could be to use a parametric family of distributions to model each metocean variable individually, however we employ a non-parametric approach. This involves constructing a kernel density estimate of the cdf for the ith environmental parameter, denoted $\hat{F}_{i,t}$, at each time point t in a year to deseasonalise the data. The transformed marginal variables $\hat{F}_{i,t}(X_{i,t})$ will therefore all follow a common (standard uniform) distribution.

In order to generate the necessary joint and conditional distributions we transform the deseasonalised marginals to the standard normal scale by use of the probability integral transform, $Y_{i,t} = \Phi^{-1}(\hat{F}_{i,t}(X_{i,t}))$, to give the transformed vector $\underset{\sim}{\mathbf{Y}}_t = (Y_{1,t}, \ldots, Y_{d,t})$. This choice is made to exploit the properties of the multivariate normal distribution, namely that its conditionals are also distributed as multivariate normal. Furthermore correlation analysis is best employed with the use of normal margins. We observe that a Gaussian multivariate copula fails to capture the true complexity of the joint distribution of $\mathbf{Y}_t$ in Section 4.3, but because of the Gaussian margins we feel it is appropriate to use a Gaussian kernel, see Section 4.2.1. We therefore

simulate from the full multivariate kernel density estimate instead.

We are not limited to examining purely kth order Markov chains as we can condition on a non-consecutive set of time lags and variables. For example a second order Markov chain could be supplemented with higher order dependencies for significant wave height H_s and wave period T_p such as the example in Eq. 4.1.2. This could be viewed alternatively as a fourth order Markov chain with some higher order terms removed. Specifically for all $t > 4$ we could have

$$P(\mathbf{Y}_{t+1} \in A|\mathbf{Y}_t = \mathbf{y}_t, \ldots, \mathbf{Y}_1 = \mathbf{y}_1) = P(\mathbf{Y}_{t+1} \in A|\mathbf{Y}_t = \mathbf{y}_t, \mathbf{Y}_{t-1} = \mathbf{y}_{t-1}, H_{s,t-2} = h_{s,t-2}, T_{p,t-4} = t_{p,t-4}). \tag{4.1.2}$$

We observe that Eq. 4.1.2 no longer contains explicit dependence on t. This graphical model may be a more appropriate form than a fully saturated fourth order Markov model since it may match better with the underlying physics of the system. In addition, models containing fewer parameters will have improved accuracy due to the improved ratio between number of parameters and sample size. The relative benefits of simpler graphical structures are determined by the calibration of auto- and cross-correlation of data. New observations can then be simulated on the stationary scale $\mathbf{Y}_t$ from any multivariate Markov structure with the desired structure. The completed synthetic dataset is obtained by back-transforming the marginals of the simulated data through their respective inverse kernel density estimates to restore both the observed marginal structure and seasonality to the time series, $\mathbf{X}_t = \hat{F}_t^{-1}(\Phi(\mathbf{Y}_t)) = (\hat{F}_{1,t}^{-1}(\Phi(Y_{1,t})), \ldots, \hat{F}_{d,t}^{-1}(\Phi(Y_{d,t})))$, whilst preserving the observed dependence.

4.1.4 Data Information

The dataset analysed in this chapter was obtained from our industry partners JBA Consulting and consists of 10 years of data (2000-2009) sampled at 3 hourly intervals. We choose to focus our research on three key meteorological and oceanographic variables: wind speed (m/s), the significant wave height (m) and wave period (s).

The wind speed W_s, is required as it has the largest impact on the safety of offshore operations. The significant wave height H_s, is defined as the average of the highest third of all the waves measured within a 20 minute period. This convention has been adopted as larger waves tend to be more destructive and impactful than smaller waves. Significant wave height measurements are therefore a more meaningful tool than simple averages when considering the sea's impact on offshore activities. Wave period T_p, is taken to be the mean zero-up crossing wave period which provides information about the frequency of wave crests. The calculation of this mean is achieved by taking the average at the points of the wave where its vertical displacement crosses the z-axis, rather than at its peaks or troughs.

4.2 Methodology

We now introduce the statistical techniques that we use to develop our statistical model of offshore metocean conditions. Kernel density estimation is introduced with the goal of converting the original data to a stationary scale upon which a statistical dependence model (both across variables and over time) can be built. In Section 4.2.1 we outline the principles associated with kernel density estimation, before adapting

the standard methods to help us deseasonalise our data series in Section 4.2.2. Section 4.2.3 defines the conditional property of multivariate normal distributions and Section 4.2.4 describes the process we use to generate new observations from a conditional multivariate kernel density estimate.

4.2.1 Kernel Density Estimation

Kernel density estimation is a widely used non-parametric technique for estimating the distribution of a random variable (Silverman, 2018). Compared to parametric techniques, which involve selecting a parametric family of distributions and estimating the parameters of this family via maximum likelihood estimation, it offers greater flexibility in modelling and is capable of discerning multiple modes, skewness and other structures present without requiring the selection of a parametric family of models. The goal of kernel density estimation (KDE) is to determine a smoothed estimate of an underlying probability distribution from a limited set of data points. We first consider a set of independent and identically distributed (i.i.d.) univariate data points $(x_1, x_2, \ldots, x_n)$ which are assumed to follow an unknown probability density function f. The kernel density estimate of these points is

$$\hat{f}_h(x) = \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{x - x_i}{h}\right), \tag{4.2.1}$$

where K is a kernel (a symmetric probability density function centered at zero which integrates to 1) and $h > 0$ is the bandwidth. We employ a Gaussian kernel henceforth so $K = \phi$, the standard normal density function. The bandwidth determines the scale of smoothing in the estimate. Mathematically speaking the kernel density estimate

applies a kernel function at each data point x_i, so that the estimate at a point x is the sum of the contributions from each kernel function at that point. A kernel estimate of the cumulative distribution function is obtained by integrating $\hat{f}_h(x)$, i.e.,

$$\hat{F}_h(x) = \frac{1}{n}\sum_{i=1}^{n} \Phi\left(\frac{x - x_i}{h}\right), \tag{4.2.2}$$

where Φ is the cumulative distribution function of the standard normal distribution. As described in Silverman (2018) there is a necessary trade-off between the bias and variance of the KDE dependent on the size of h. A smaller bandwidth with less smoothing produces a less biased estimate at the cost of a larger variance. Estimates with an over-smoothed bandwidth yield the opposite: a larger bias but smaller variance. The standard approach is to choose the bandwidth that minimises the asymptotic mean integrated squared error (MISE) in order to best reduce both the bias and variance of the estimator. However if the underlying density is approximately Gaussian and Gaussian kernels are used then a good estimate for the optimal bandwidth h is 'Silverman's rule of thumb' Silverman (2018),

$$h = \frac{0.9\min(\hat{\sigma}, \frac{IQR}{1.34})}{n^{\frac{1}{5}}}. \tag{4.2.3}$$

It involves the standard deviation $\hat{\sigma}$ of the data sample of size n and their interquartile range IQR. Note that as n increases h decreases, so as more data is obtained less smoothing is required. The advantage of this rule is its simplicity to compute.

We also employ boundary reflection on our kernel density estimates to cope with the domain boundaries found in our data. These are important to consider as negative windspeeds and negative significant wave heights make no physical sense. Standard kernel density estimation suffers from boundary effects from a lack of data below a

lower boundary. This can cause an underestimation in the density estimate compared to the true density near the boundary. A simple adjustment can be made to incorporate weights from below the boundary by reflecting data points about the boundary and including them in the estimate. The set of data points $(x_1, x_2, \ldots, x_n)$ is augmented with $(-x_1+2b, -x_2+2b, \ldots, -x_n+2b)$, where b is the lower bound of the data. The natural lower bounds of windspeed and significant wave heights are 0m/s and 0.1m respectively, so the zero probability is assigned to values below these boundaries in the reflected kernel density estimate of the form

$$\hat{f}_h(x) = \frac{1}{nh} \sum_{i=1}^{n} \left\{ K\left(\frac{x - x_i}{h}\right) + K\left(\frac{x + x_i - 2b}{h}\right) \right\}, \tag{4.2.4}$$

for $x \geq b$ and $\hat{f}_h(x) = 0$ for $x \leq b$. This estimate will always have zero derivative at the boundary assuming the kernel is symmetric and differentiable, which is the case when $K = \phi$. In practice we only need to reflect observations close to the boundary since they are the only points that will contribute to the density estimate near it. As our approach uses $K = \phi$ there is no practical need to reflect points with $x_i > 4h + b$, since ϕ decays very quickly from its mode.

In higher dimensions the d-dimensional kernel density estimate $\hat{f}_{\mathbf{H}}$ is constructed from a set of i.i.d d-variate vectors $(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ as

$$\hat{f}_{\mathbf{H}}(\mathbf{x}) = \frac{1}{n|\mathbf{H}|^{\frac{1}{2}}} \sum_{i=1}^{n} K\left(\frac{\mathbf{x} - \mathbf{x}_i}{\mathbf{H}^{\frac{1}{2}}}\right), \tag{4.2.5}$$

where $|\mathbf{H}|$ is the determinant of a positive definite bandwidth matrix $\mathbf{H}$ and K is a symmetric multivariate kernel. We will employ the standard multivariate normal kernel of d dimensions, $K = \phi_d$ in our subsequent kernel density estimation process,

so

$$\hat{f}_{\mathbf{H}}(\mathbf{x}) = \frac{1}{n\sqrt{|\mathbf{H}|(2\pi)^d}} \sum_{i=1}^{n} e^{-\frac{1}{2}(\mathbf{x}-\mathbf{x}_i)^T \mathbf{H}^{-1}(\mathbf{x}-\mathbf{x}_i)}, \tag{4.2.6}$$

where the bandwidth matrix $\mathbf{H}$ is analogous to the covariance matrix.

4.2.2 Deseasonalising Data

The first step of our approach is designed to reduce the prevalence of long term seasonality in our dataset. This requires us to estimate $\hat{F}_{i,t}$ as described in Section 4.1.3 for each of our series ($i = 1, \ldots, d$). We consider historical time series of windspeeds, significant wave heights and wave periods measured at 3 hourly intervals over the course of 10 years. In order to apply our statistical methods we seek to transform these time-series into stationary series.

Recall that the probability integral transform (PIT) allows us to alter the distribution of a set of univariate data to that of a standard uniform distribution by constructing their cumulative distribution function. A continuous univariate random variable X with an unknown distribution can be transformed to U(0,1) by passing it through its cdf, F_X, so that $Z = F_X(X)$ has a $U(0,1)$ distribution. Furthermore, if we wish to transform X to Y, where Y follows a standard normal distribution then $Y = \Phi^{-1}(F_X(X))$.

The classic kernel density methods introduced in Section 4.2.1 require that the data be identically distributed, so the methods are only applicable to specific windows during which an assumption of an identical distribution is deemed reasonable. Given our initial dataset starts with unknown probability distributions we need to

construct a kernel density estimate for the cdf at various points in time. This choice is made since the true distribution will vary depending on the time year, for example windspeeds are significantly higher in winter months compared to summer months. We create a kernel density estimate of each variable at time t to fully capture the overall distribution's time dependency. We utilise a local window E_t around a central time point x_t containing κ observations before and after x_t, i.e. a window of $2k+1$ values. In addition we add the equivalent data points in subsequent years into the window since they also follow the same distribution if we neglect the effects of climate change. These occur with a periodicity of $w = 8 * 365 = 2920$, which is the number of observations in a year. The mathematical definition of such a window is thus,

$$E_t = \{s : |s \mod w - t| \leq \kappa\}. \tag{4.2.7}$$

We choose to set $\kappa = 200$ under the belief that only short term local patterns up to the range of a few weeks worth of observations can be considered to follow the same underlying distribution as x_t. Specifically the window incorporates 25 days worth of observations from both before and after the central time point (recall that observations occur every 3 hours). The complete reflected kernel distribution estimate about a lower bound b using data from the local window E_t for the cumulative distribution function at time t is given as,

$$\hat{F}_t(x) = \frac{1}{|E_t|} \sum_{i \in E_t} \left\{ \Phi\left(\frac{x - x_i}{h}\right) + \Phi\left(\frac{x + x_i - 2b}{h}\right) \right\}. \tag{4.2.8}$$

Long term seasonality trends can be reincorporated by transforming data back onto its original scales. A new set of observations generated on a different scale can be altered to follow the original scale by reversing the PIT. A set of observations

Z following a $U(0,1)$ distribution can be transformed via the inverse of a given cdf F_X, so that $F_X^{-1}(Z)$ follows the same distribution as X. Additionally if we wish to transform Y which follows a standard normal distribution to X with a cdf F_X, then $X = F_X^{-1}(\Phi(Y))$.

4.2.3 Conditional Distribution of a Multivariate Normal Distribution

Our previous steps were aimed at extracting the long term seasonality from the data set, however consideration needs to be made for the shorter term dependence structures that remain present. These effects tend to be driven by environmental processes such as offshore storms which persist across several days. A simple pass through the deseasonalised data highlights the significant degree of autocorrelation within each metocean variable and the large amount of cross-correlation between variables within this range. Here a statistical model is developed to capture and represent these features.

Markov models provide a general framework for modelling continuous stochastic processes. First order Markov models rely on the Markov property: that the conditional probability distribution of the future depends only on the present state. This is also referred to as the "memoryless" property. Higher order Markov models in contrast, depend on more than just the previous value thus incorporating a degree of "memory" into the model. Higher order processes therefore utilise more historical information than lower order processes. We utilise a higher order Markov model to

better capture the features within our deseasonalised data. A single model can drive the entire process since the data were transformed to be marginally stationary. Formally the joint distribution at time t of $\mathbf{Y}_t$, is assumed to be driven by the a set of k previous values (lags) $\mathbf{Y}_{t-1}, \ldots, \mathbf{Y}_{t-k}$.

A natural approach for modelling this Markov process comes from transforming the deseasonalised data to the normal scale via the PIT in order to use the conditional properties of multivariate normal distributions. We describe the equations involved for completeness. Assume that $\mathbf{Z} \sim MVN(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is an N-dimensional set of variables with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. These variables can be also be split into a set of variables $\mathbf{Z}^a$ and a set $\mathbf{Z}^b$. In our context $\mathbf{Z}^a = \mathbf{Y}_t$ and $\mathbf{Z}^b = (\mathbf{Y}_{t-1}, \ldots, \mathbf{Y}_{t-k})$, so that $\mathbf{Z}^a$ contains variables at the current time step t and $\mathbf{Z}^b$ environmental variables at any of the previous k lags. Furthermore, we can partition $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ as follows

$$\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_a \\ \\ \mathbf{Z}_b \end{bmatrix}, \boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_a \\ \\ \boldsymbol{\mu}_b \end{bmatrix}, \boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{a,a} & \boldsymbol{\Sigma}_{a,b} \\ \\ \boldsymbol{\Sigma}_{b,a} & \boldsymbol{\Sigma}_{b,b} \end{bmatrix}. \tag{4.2.9}$$

An analytical formula for the conditional distribution exists which states that the conditional distribution of a set of variables $\mathbf{Z}^a$ given another set $\mathbf{Z}^b$ taking the value $\mathbf{c}$ is again multivariate normal. Formally $\left(\mathbf{Z}^a | \mathbf{Z}^b = \mathbf{c}\right) \sim MVN\left(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}\right)$, where

$$\hat{\boldsymbol{\mu}} = \boldsymbol{\mu}_a + \boldsymbol{\Sigma}_{a,b}\boldsymbol{\Sigma}_{b,b}^{-1}\left(\mathbf{c} - \boldsymbol{\mu}_b\right) \tag{4.2.10}$$

$$\hat{\boldsymbol{\Sigma}} = \boldsymbol{\Sigma}_{a,a} - \boldsymbol{\Sigma}_{a,b}\boldsymbol{\Sigma}_{b,b}^{-1}\boldsymbol{\Sigma}_{b,a}. \tag{4.2.11}$$

We can condition on as many lags or subsets of variables as is deemed necessary in order to accurately model future observations.

Our empirical modelling found this approach to produce poor replications of the

original data. We tried a variety of graphical structures and values for k, but none yielded a sufficiently good level of fit. Some variables such as the wave height and wave period did not appear to follow a bivariate Gaussian distribution post deseasonalisation, with a different regime for swell and wind-sea waves. We also found that the distributions of weather windows and waiting time before weather windows appeared to be significantly underestimating the empirical distributions observed in our dataset. Therefore we instead chose to simulate data from a multivariate kernel density estimation as described in Section 4.2.4.

4.2.4 Simulating From a Multivariate Kernel Density Estimate

We now seek to instead construct a kernel density estimate of the conditional distribution of $\mathbf{Z}^a|\mathbf{Z}^b$ from the joint kernel distribution of $\mathbf{Z}^a$ and $\mathbf{Z}^b$. Sampling from a conditional density estimate constructed from kernel density estimates is a relatively straightforward process.

We begin by illustrating this for the bivariate case containing data $(x_1, y_1), \ldots, (x_n, y_n)$ in which we want to simulate from $Y|X = x$. Instead of randomly sampling from the existing data points (x_i, y_i) we are required to use weighted sampling in proportion to each data points contribution to the overall density given a observed value for X . This can be understood from the formula in Eq. (4.2.15), which has been derived from a two-dimensional kernel density estimate. We utilise the standard bivariate normal density function $\phi(x, y; \rho)$ as the kernel, with ρ the correlation coefficient and a band-

width matrix $\mathbf{H} = \begin{bmatrix} h_x^2 & h_x h_y \rho \\ h_y h_x \rho & h_y^2 \end{bmatrix}$. Then the conditional kernel density estimate of Y given X is

$$\hat{f}(y|x) = \frac{\hat{f}(x,y)}{\hat{f}(x)} \tag{4.2.12}$$

$$= \frac{\frac{1}{nh_xh_y}\sum_{i=1}^{n}\phi\left(\frac{x-x_i}{h_x}, \frac{y-y_i}{h_y}; \rho\right)}{\frac{1}{nh_x}\sum_{i=1}^{n}\phi\left(\frac{x-x_i}{h_x}\right)} \tag{4.2.13}$$

$$= \frac{\frac{1}{h_y}\sum_{i=1}^{n}\left(\phi\left(\frac{x-x_i}{h_x}\right)\phi\left(\frac{y-(y_i+\rho\frac{h_y}{h_x}(x-x_i))}{h_y(1-\rho^2)^{\frac{1}{2}}}\right)\right)}{\sum_{i=1}^{n}\phi\left(\frac{x-x_i}{h_x}\right)} \tag{4.2.14}$$

$$= \frac{1}{h_y}\sum_{i=1}^{n}\omega_i\phi\left(\frac{y-(y_i+\rho\frac{h_y}{h_x}(x-x_i))}{h_y(1-\rho^2)^{\frac{1}{2}}}\right), \tag{4.2.15}$$

where $\omega_i = \phi\left(\frac{x-x_i}{h_x}\right) / \sum_{j=1}^{n}\phi\left(\frac{x-x_j}{h_x}\right)$ is the fraction that the kernel centered at x_i contributes to the overall kernel density estimate for x, with $0 < w_i < 1$ for $i = 1, \ldots, n$ and $\sum_{i=1}^{n} w_i = 1$. Thus with probability ω_i the y realisation comes from a $N\left(y_i + \rho\frac{h_y}{h_x}(x-x_i), h_y(1-\rho^2)^{\frac{1}{2}}\right)$ random variable.

We now consider the multivariate extension of the kernel density estimation as shown in Eq. 4.2.5. We continue to use the sets $\mathbf{Z}^a$ and $\mathbf{Z}^b$ as outlined in Section 4.2.3 but they now include a time index t. This allows us to build a time series in an iterative manner. Joint observations of $\mathbf{Z}$ at time i, $\mathbf{z}_i$, can be partitioned into two sets $[\mathbf{z}_i^a, \mathbf{z}_i^b]$ according to whether they relate to the current or preceding time steps. The bandwidth matrix is partitioned in the same manner, $\mathbf{H} = \begin{bmatrix} \mathbf{H}_{a,a} & \mathbf{H}_{a,b} \\ \mathbf{H}_{b,a} & \mathbf{H}_{b,b} \end{bmatrix}$. We first generate an initial realisation of $\mathbf{Z}_0 = (\mathbf{Y}_0, \mathbf{Y}_{-1}, \ldots, \mathbf{Y}_{1-k})$ by randomly sampling from the historical data points $(\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_n)$ and then simulating from the

marginalised kernel associated with it. Then we can simulate future observations, $\mathbf{Z}^a_{t+1} = \mathbf{Y}_{t+1}$, from past observations, $\mathbf{Z}^b_{t+1} = (\mathbf{Y}_t, \ldots, \mathbf{Y}_{t-k+1})$, (for $t = 0, \ldots, m$) by probabilistically selecting a kernel i at centered at $\mathbf{z}_i = [\mathbf{z}^a_i, \mathbf{z}^b_i]$ with probability ω_i defined by Eq. (4.2.16) and simulating from the kernel's conditional distribution, E.q. (4.2.17). This process is repeated to construct a time series of length m as shown below.

1. Generate $\mathbf{Z}_1$ by selecting a kernel i with probability $1/n$ and simulating from its marginalised kernel, $\mathbf{Z}_1 = \mathbf{z}^b_i + MVN(\mathbf{0}, \mathbf{H}_{b,b})$.

2. For $t = 0, \ldots, m$

   (a) Normalise the weights $w^{(t)}_i$ that the kernel applied to each data point $\mathbf{z}_i$ contributes to the KDE at $\mathbf{Z}^b_{t+1}$ to give

   $$w^{(t)}_i = \phi\left(\frac{\mathbf{Z}^b_{t+1} - \mathbf{z}^b_i}{\mathbf{H}^{\frac{1}{2}}}\right) \Big/ \sum_{j=1}^{n} \phi\left(\frac{\mathbf{Z}^b_{t+1} - \mathbf{z}^b_j}{\mathbf{H}^{\frac{1}{2}}}\right). \tag{4.2.16}$$

   (b) Select kernel i (centered at $\mathbf{z}_i = [\mathbf{z}^a_i, \mathbf{z}^b_i]$) to simulate from with probability $w^{(t)}_i$.

   (c) Generate the next observations $\mathbf{Z}^a_{t+1}$ from the conditional multivariate normal distribution.

   $$\mathbf{Z}^a_{t+1} | \mathbf{Z}^b_{t+1} \sim MVN\left(\mathbf{z}^a_i + \boldsymbol{H}_{a,b}\boldsymbol{H}^{-1}_{b,b}\left(\mathbf{Z}^b_{t+1} - \mathbf{z}^b_i\right), \boldsymbol{H}_{a,a} - \boldsymbol{H}_{a,b}\boldsymbol{H}^{-1}_{b,b}\boldsymbol{H}_{b,a}\right) \tag{4.2.17}$$

4.2.5 Selection of Bandwidth Smoothing Parameter

Our approach for determining the correct set of variables and time lags to condition on relies upon a set of statistical diagnostics. These diagnostics include direct compar-

isons between the deseasonalised and simulated data via the autocorrelation within variables, the cross correlation between variables and marginal QQ-plots. Another key validation tool are the weather window persistence distributions which would underpin any practical use of the synthetic time series.

Our results indicate that the quality of the fit depends highly on the bandwidth scale. We first utilise a standard normal scale plug-in bandwidth selector given by Chacón et al. (2011) denoted as,

$$\mathbf{H}_{ns} = \mathbf{S}\left(\frac{4}{n+2}\right)^{\frac{2}{d+4}}, \tag{4.2.18}$$

where $\mathbf{S}$ is the sample covariance matrix, d is the dimensionality of the data and n is the number of observations. An example bandwidth matrix $\mathbf{H}_{ns}$ is shown below and conditions on multiple variables and time steps. In particular it has a graphical structure in which the set of variables $\mathbf{Z}_{t+1}$ depend on $\mathbf{Z}_t, \mathbf{Z}_{t-1}, \mathbf{Z}_{t-2}, \mathbf{Z}_{t-3}, \mathbf{Z}_{t-7}$ and $\mathbf{Z}_{t-10}$. We note that our earlier use of the probability integral transform ensures that all of the marginal variables have the same variance so the bandwidth matrix is symmetric with its diagonal elements being identical. It appears that classical bandwidth estimators such as this oversmooth the kernel density estimate causing the autocorrelation functions in the marginals to become a poor fit to those of the observed data. We have found through empirical analysis that the elements of the bandwidth matrix $\mathbf{H}$ in our conditional simulations should be at least an order of magnitude smaller than those implied by $\mathbf{H}_{ns}$. This result captures the acf at larger lags significantly better allowing us to reduce the amount of variables to be conditioned on.

	W_s^t	H_s^t	T_p^t	W_s^{t-1}	H_s^{t-1}	T_p^{t-1}	W_s^{t-2}	H_s^{t-2}	T_p^{t-2}	W_s^{t-3}	H_s^{t-3}	T_p^{t-3}	W_s^{t-7}	H_s^{t-7}	T_p^{t-7}	W_s^{t-10}	H_s^{t-10}	T_p^{t-10}
W_s^t	0.339	0.263	0.172	0.289	0.225	0.152	0.244	0.189	0.131	0.204	0.158	0.112	0.109	0.082	0.06	0.072	0.054	0.040
H_s^t	0.263	0.339	0.269	0.271	0.305	0.246	0.253	0.273	0.22	0.228	0.238	0.188	0.137	0.138	0.113	0.095	0.094	0.081
T_p^t	0.172	0.269	0.339	0.175	0.243	0.269	0.166	0.225	0.231	0.153	0.201	0.199	0.102	0.124	0.115	0.079	0.090	0.082
W_s^{t-1}	0.289	0.271	0.175	0.339	0.263	0.172	0.289	0.225	0.152	0.244	0.189	0.131	0.125	0.096	0.070	0.082	0.061	0.045
H_s^{t-1}	0.225	0.305	0.243	0.263	0.339	0.269	0.271	0.305	0.246	0.253	0.273	0.220	0.155	0.159	0.134	0.108	0.108	0.093
T_p^{t-1}	0.152	0.246	0.269	0.172	0.269	0.339	0.175	0.243	0.269	0.166	0.225	0.231	0.112	0.139	0.134	0.086	0.100	0.096
W_s^{t-2}	0.244	0.253	0.166	0.289	0.271	0.175	0.339	0.263	0.172	0.289	0.225	0.152	0.145	0.112	0.082	0.096	0.070	0.053
H_s^{t-2}	0.189	0.273	0.225	0.225	0.305	0.243	0.263	0.339	0.269	0.271	0.305	0.246	0.177	0.182	0.153	0.122	0.128	0.108
T_p^{t-2}	0.131	0.220	0.231	0.152	0.246	0.269	0.172	0.269	0.339	0.175	0.243	0.269	0.124	0.154	0.156	0.094	0.120	0.121
W_s^{t-3}	0.204	0.228	0.153	0.244	0.253	0.166	0.289	0.271	0.175	0.339	0.263	0.172	0.173	0.133	0.096	0.109	0.082	0.060
H_s^{t-3}	0.158	0.238	0.201	0.189	0.273	0.225	0.225	0.305	0.243	0.263	0.339	0.269	0.201	0.217	0.177	0.137	0.138	0.113
T_p^{t-3}	0.112	0.188	0.199	0.131	0.220	0.231	0.152	0.246	0.269	0.172	0.269	0.339	0.139	0.187	0.200	0.102	0.124	0.115
W_s^{t-7}	0.109	0.137	0.102	0.125	0.155	0.112	0.145	0.177	0.124	0.173	0.201	0.139	0.339	0.263	0.172	0.204	0.158	0.112
H_s^{t-7}	0.082	0.138	0.124	0.096	0.159	0.139	0.112	0.182	0.154	0.133	0.217	0.187	0.263	0.339	0.269	0.228	0.238	0.188
T_p^{t-7}	0.060	0.113	0.115	0.070	0.134	0.134	0.082	0.153	0.156	0.096	0.177	0.200	0.172	0.269	0.339	0.153	0.201	0.199
W_s^{t-10}	0.072	0.095	0.079	0.082	0.108	0.086	0.096	0.122	0.094	0.109	0.137	0.102	0.204	0.228	0.153	0.339	0.263	0.172
H_s^{t-10}	0.054	0.094	0.090	0.061	0.108	0.100	0.070	0.128	0.120	0.082	0.138	0.124	0.158	0.238	0.201	0.263	0.339	0.270
T_p^{t-10}	0.040	0.081	0.082	0.045	0.093	0.096	0.053	0.108	0.121	0.060	0.113	0.115	0.112	0.188	0.199	0.172	0.270	0.339

Figure 4.2.1: Estimated bandwidth matrix $\mathbf{H}_{ns}$ before scaling reduction.

Figure 4.2.2 demonstrates that scaling the bandwidth matrix $\mathbf{H}_{ns}$ to become $\mathbf{H}_{ns}/10$ and $\mathbf{H}_{ns}/20$ produces a superior match to the autocorrelation in the deseasonalised data. This was performed for a graphical model with the following structure, where the only change was the scaling in the covariance matrix.

$$P(\mathbf{Z}_{t+1} \in A|\mathbf{Z}_t = \mathbf{z}_t, \ldots, \mathbf{Z}_1 = \mathbf{z}_1) = P(\mathbf{Z}_{t+1} \in A|\mathbf{Z}_t = \mathbf{z}_t, \mathbf{Z}_{t-1} = \mathbf{z}_{t-1}, \mathbf{Z}_{t-2} = \mathbf{z}_{t-2}, \mathbf{Z}_{t-7} = \mathbf{z}_{t-7}, \mathbf{Z}_{t-10} = \mathbf{z}_{t-10}) \tag{4.2.19}$$

A similar pattern is observed for the remaining marginal variables. Whilst this may appear dramatic, note that this reduction in covariance effectively only reduces the standard deviation in the kernel by a factor of $\sqrt{20} \approx 4.47$ from a small number to

begin with.

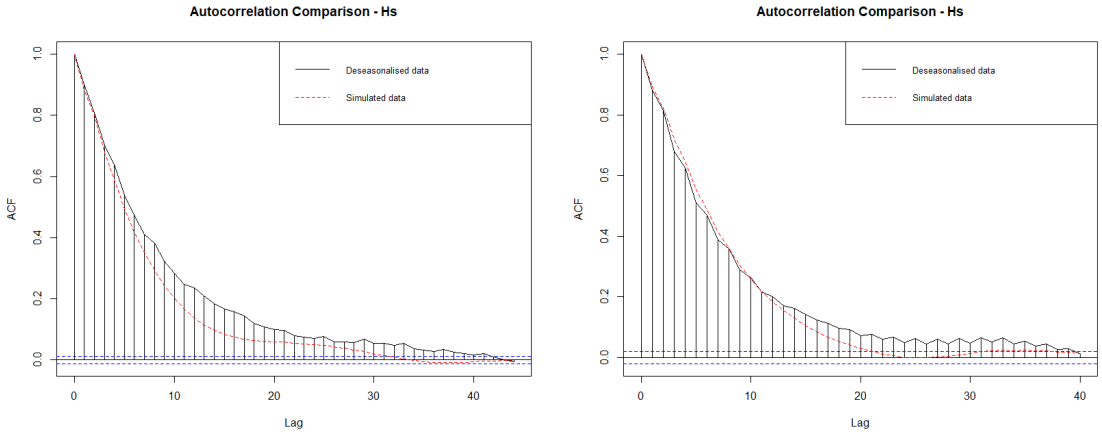

(a) $\mathbf{H}_{ns}$ (b) $\mathbf{H}_{ns}/10$

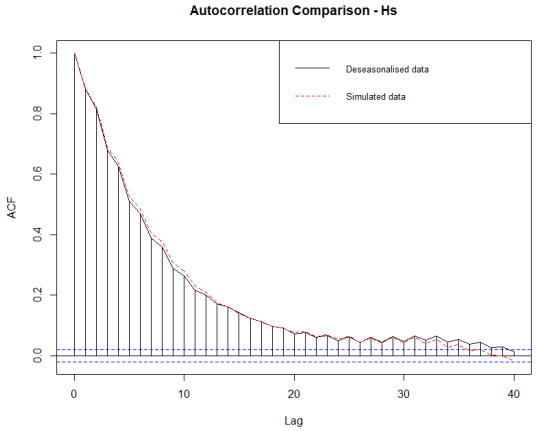

(c) $\mathbf{H}_{ns}/20$

Figure 4.2.2: The use of a scaling reduction on the $\mathbf{H}_{ns}$ improves the autocorrelation fit. Illustrated for deseasonalised significant wave heights.

4.2.6 Graphical Modelling

In order to make our statistical model more reliable we seek to reduce the number of parameters in the model by investigating whether all of the links between variables in

a kth order Markov chain are strictly necessary. A complicated model which conditions on multiple variables across a multitude of previous values may perform equally as well as a reduced, simpler version. This reduction will also make our estimates of the remaining variables more accurate as estimates will be less variable given fewer parameters are estimated. The correlation matrix C and the inverse correlation matrix $D = C^{-1}$ provide the necessary information. The elements of the inverse correlation matrix D inform us of the correlation between pairs of variables conditioned on the remaining variables, (Whittaker, 1990). These are better known as the partial correlations. Under our assumption that all the variables follow a multivariate Gaussian distribution, a partial correlation element $D_{ij} = 0$ implies that variables i and j are conditionally independent given the other variables. This property can therefore be used to reduce a dense graphical model to a simpler form such as that shown in Fig 4.2.3 by replacing small values in D with zero. Formally two random variables X and Y are conditionally independent given an extra variable Z if and only if joint cumulative distribution function given Z can be decomposed into the product of the independent cdf given Z for both X and Y.

$$(X \perp\!\!\!\perp Y)|Z \iff F_{X,Y|Z=z}(x,y) = F_{X|Z=z}(x).F_{Y|Z=z}(y) \qquad \forall x, y, z. \tag{4.2.20}$$

A drawback to the partial correlation analysis is that the assumption of a partial correlation being zero has an effect on the remaining partial correlation structure. Making one such approximation requires solving an equation to find how the maximum likelihood estimates change, but multiple zero approximations have to be solved simultaneously and not iteratively. This makes the situation significantly more com-

plicated.

A pairwise comparison can be made between two graphical models using the following log-likelihood based approach. The log likelihood of a q-variate Gaussian graphical model with a mean μ, covariance matrix Σ, based on a random sample of N observations is given as,

$$2l(\mu, \Sigma) = const - N\log(|\Sigma|) - Ntr(S\Sigma^{-1}) - N(\bar{x} - \mu)^T\Sigma^{-1}(\bar{x} - \mu), \qquad (4.2.21)$$

where $\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i$ is the sample mean, $S = \frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})(x_i - \bar{x})^T$ is the sample covariance and tr is the trace of a matrix (sum of all the elements along its main diagonal). We observe that this log-likelihood is maximised when the maximum likelihood estimators take the value $\hat{\mu} = \bar{x}$ and $\hat{\Sigma} = S$. Using these estimators causes the final term in Eq. (4.2.21) to evaluate to zero. The term $Ntr(S\hat{\Sigma}^{-1})$ evaluates to $Ntr(SS^{-1}) = Ntr(I_q) = Nq$, where q is the number of variables in the model. Ignoring the constant term, the maximised log-likelihood under the saturated model M_s simplifies to

$$2l_s(\hat{\mu} = \bar{x}, \hat{\Sigma} = S) = -N\log(|S|) - Nq. \qquad (4.2.22)$$

In order to compare graphical models we need to construct the deviance, which is denoted as twice the difference between the maximised log-likelihood under a fully saturated model M_s and a candidate model M_c. This simplified model M_c has a different covariance matrix S_c where we have approximated several entries in the inverse matrix ${S_c}^{-1}$ to be zero. Whittaker (1990) shows that the approximation of elements to zero in S_c^{-1} only effects the corresponding elements in S. The remaining

elements of S_c are the same as those in S, i.e. $S_{ij} = (S_c)_{ij}$ for all i, j not approximated to zero in $(S_c)_{ij}^{-1}$. This means the term $Ntr(S\hat{\Sigma}^{-1})$ also evaluates to $Ntr(SS_c^{-1}) = Ntr(I_q) = Nq$, where q is the number of variables in the model, so

$$2l_c(\hat{\mu} = \bar{x}, \hat{\Sigma} = S_c) = -N\log(|S_c|) - Nq. \tag{4.2.23}$$

This means that the deviance difference between the two models can be written as,

$$\text{Deviance difference for } M_c \subseteq M_s \ = 2(\hat{l}_s - \hat{l}_c) \tag{4.2.24}$$

$$= -N\log(|S|) + N\log(|S_c|) \tag{4.2.25}$$

$$= N\log(|S_c|/|S|). \tag{4.2.26}$$

In order to test whether a model M_2 is a better fit than model M_1, we can perform a likelihood ratio test based on E.q. (4.2.24) which is χ_p^2 distributed where p is the difference in the number of parameters. The testing of these models for significance is normally performed in a sequential manner.

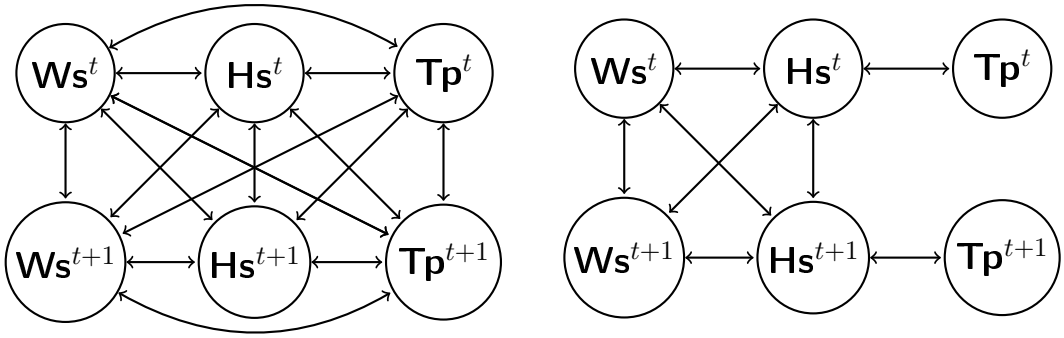

Figure 4.2.3: Two graphical model candidates. One is more saturated than the other.

The two main selection methods are forwards selection and backwards elimination. Forwards selection starts from the independence model, which assumes that all of the

variables are completely independent, and adds in edges based on deviance difference rules. Backwards elimination removes edges from the model since it begins from the fully saturated model that contains every possible link between the variables. As the underlying physics of the problem is well understood, we hope that any statistical selection methods will converge with our understanding of the physical processes involved. For example it would make little sense for the wave period at time t to affect the wind speed at $t + 20$.

4.3 Time Series Validation

We chose to employ a graphical model of the following form with a rescaled bandwidth matrix of $\mathbf{H}_{ns}/20$,

$$P(\mathbf{Z}_{t+1} \in A|\mathbf{Z}_t = \mathbf{y}_t, \ldots, \mathbf{Z}_1 = \mathbf{z}_1) = P(\mathbf{Z}_{t+1} \in A|\mathbf{Z}_t = \mathbf{z}_t, \mathbf{Z}_{t-1} = \mathbf{z}_{t-1}, \mathbf{Z}_{t-2} = \mathbf{z}_{t-2}, \mathbf{Z}_{t-5} = \mathbf{z}_{t-5}, \mathbf{Z}_{t-9} = \mathbf{z}_{t-9}). \tag{4.3.1}$$

This graphical model conditions the $\mathbf{Z}_{t+1}$ future observations using a tenth order Markov chain with many of the higher order terms removed, specifically keeping only the $\mathbf{Z}_t, \mathbf{Z}_{t-1}, \mathbf{Z}_{t-2}, \mathbf{Z}_{t-5}, \mathbf{Z}_{t-9}$ lagged terms. We believe that the inclusion of the both the lower order terms and higher order terms provide benefit to the resulting model. The first three lags help to provide a sense of immediate direction and the connected rate of change of the metocean variables, whilst the higher order terms balance this with information from over 24 hours in the past to reflect the dependence between different weather systems. Whilst this model may still appear to be overparametrised compared to other works and the underlying physics of the problem, it produces good results for our dataset. We note that although our modelling framework is flexible enough to condition on individual variables at higher order terms, every time step conditioned on in our model involves all three variables (windspeed, significant wave height and wave period).

This choice of graphical model was made after experimentation with several graphical structures. Analysis was performed on the stationary scale before the diagnostic process was repeated on the synthetic data with seasonality back transformed in. The synthetic data was generated so as to be the same length as the observed data and

for the same time periods. The results for our chosen model are now described.

4.3.1 Marginal Distribution Analysis On The Stationary Scale

We introduce the autocorrelation plot as a visual measure of comparison between two time series. These are shown for the marginal variables in Figures 4.3.1a, 4.3.1c and 4.3.1e. The autocorrelation within a time series is the correlation between values of the process at different times and for a stationary series y_t at lag k is denoted as $\rho_k = \text{Corr}(y_t, y_{t-k})$. The autocorrelation at lag zero is always 1, since this represents the autocorrelation between each term and itself. The value of the autocorrelation function at each lag is marked with a vertical line in the plots. Only acf values that are greater (or less) than the 95% tolerance bounds for independence (blue dashed line) are considered statistically significant with our plots omitting insignificant lags greater than 50 (over 6 days in the past), as all autocorrelations are essentially zero then. We have added a curve connecting each autocorrelation value in order to aid with visual comparison. The autocorrelation for the original deseasonalised dataset is marked in the solid black line, whilst the autocorrelation of the simulated dataset is presented with the red dashed line.

Figures 4.3.1b, 4.3.1d and 4.3.1f accompany their respective marginal autocorrelation plot and display their quantile-quantile (q-q) plot. This is a more powerful technique for assessing whether two distributions follow the same underlying distribution than simply comparing histograms of the two samples. We include the line $y = x$ in red as a reference line along which points will lie if the two distributions are identical. Points that significantly deviate from this line can indicate different levels

of skewness between the distributions or that one contains heavier tails.

For both types of figures we observe a good match between the deseasonalised data and the simulated data indicating that we reproduce the key features. In fact for all three variables the autocorrelations become statistically insignificant beyond lag forty highlighting the removal of long term seasonal patterns beyond five days in the past. The argument for some marginal variables is more successful than others, which we will now describe in more detail.

Windspeed autocorrelation plot

For example, the autocorrelation plot of windspeed matches very well for the first seven lags in particular and is even successful at picking up the minor piecewise deviations to the underlying curve. After this the acfs begin to slightly diverge with weaker autocorrelations observed in the simulated data than in the deseasonalised data. The difference in acfs is no more than 0.05 for any given lag.

Significant wave height autocorrelation plot

The significant wave height autocorrelation function also fits well for significant wave height at lower lags. At higher lags the acfs again exhibit a mild divergence at higher order lags albeit in a less noticeable fashion.

Wave period autocorrelation plot

The autocorrelation plot for the wave period exhibits a comparative poorer fit at lower order lags, but a closer fit at higher lags. Figure 4.3.1e shows that the jagged structures present are captured more accurately by the simulated data

from lags 26 onwards than they are in lags 1-8. We note the possible presence of a cyclical pattern where the simulated data contains an additional spike at lags 2 and 6 compared to the deseasonalised data.

Q-Q plots

The q-q plots for all three of the metocean variables indicate an excellent fit of the simulated data quantiles to the deseasonalised data quantiles.

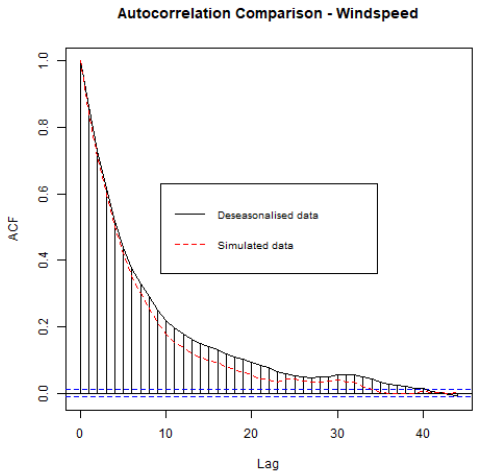

(a) Windspeed acf

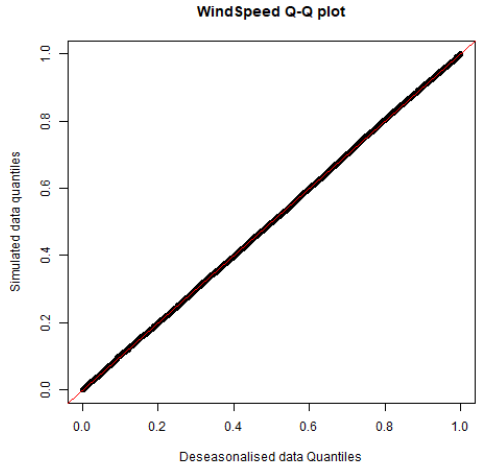

(b) Windspeed q-q plot

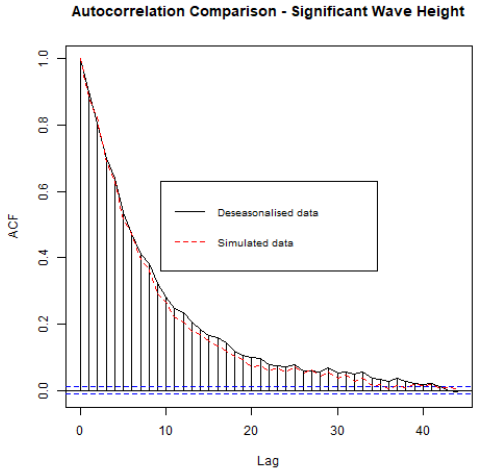

(c) Significant wave height acf

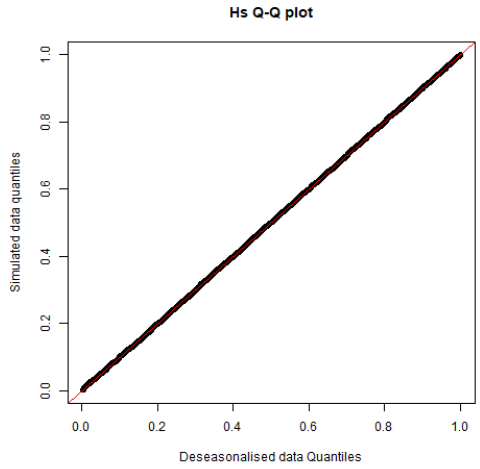

(d) Significant wave height q-q plot

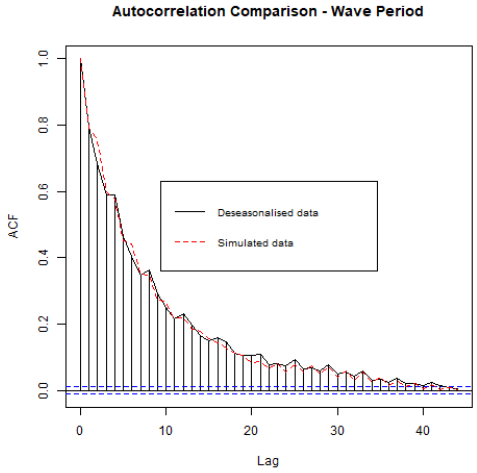

(e) Wave period acf

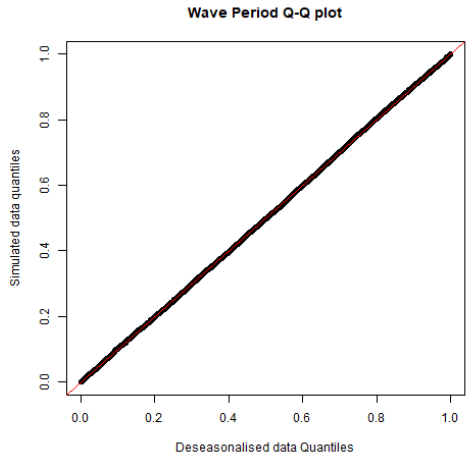

(f) Wave period q-q plot

Figure 4.3.1: Autocorrelation structure and Q-Q plots for the deseasonalised and simulated data.

4.3.2 Joint Distribution Analysis On The Stationary Scale

It is possible to generalise the concept of autocorrelation between distinct time points in a marginal variable to consider cross-correlation between different marginal variables at different time steps. The cross-correlation function will therefore reveal deeper information about the interactions between the metocean variables which are important features to capture within our joint statistical model. Mathematically the cross-correlation function between two stationary time series x_t and y_t at lag k is denoted as $R_k = \text{Corr}(x_t, y_{t-k})$. An examination of the cross-correlation function (ccf) between the windspeed, significant wave heights and wave periods is provided in Figure 4.3.2. The graphs presented follow the same approach as Figure 4.3.1 and highlight the difference between the deseasonalised and simulated data.

Similar patterns emerge in the cross-correlation structure as those found in the autocorrelation structure.

Windspeed and significant wave height cross-correlation plot

The windspeed and significant wave height ccf displays a good fit for lags close to lag zero. The simulated data starts to deviate away from the deseasonalised data for lags outside $[-8, 8]$ with the underestimation of the ccf more noticeable on the left hand side of the plot creating an asymmetry in Figure 4.3.2a. The maximum ccf value occurs for both the deseasonalised and simulated data occurs at lag -1, indicating that the strongest correlation is $\text{Corr}(Ws_t, Hs_{t+1})$. This suggests that the prevailing wind conditions precede the arrival of waves, which matches with the physical generation of waves. As one lag represents three

hours it also appears to capture the local nature of this effect.

Windspeed and wave period cross-correlation plot

The next plot comparing windspeed and wave period displays the simulated data accurately reproducing the ccf values for the first twenty lags save for the initial value at lag zero. The negative lags appear to match for far fewer lags and again underestimate the cross-correlation function. Figure 4.3.2b shows that the maximum ccf value for the deseasonalised data occurs at lag -1, whilst for the simulated data it occurs at lag 0. Physically we would expect the wave period to lag the windspeed by the same amount as the significant wave height does (1 lag), given that they occur together. It appears that our model has both shifted the peak and increased it compared to the deseasonalised data, which could be addressed in the future.

Significant wave height and wave period cross-correlation plot

The best match of all the cross-correlation functions occurs for significant wave height and wave period since the simulated and deseasonalised data agree well at both low and high lags, albeit not always picking up the jagged underlying structure at greater lags. Surprisingly the cross-correlation function for the deseasonalised data becomes statistically insignificant around lag forty but remains significant at lag minus forty. The simulated data appears to broadly capture this feature, which is present in each comparison. The maximum ccf value for both the deseasonalised and simulated data occurs at lag 0. This matches with the underlying physics that waves and wave periods occur simultaneously.

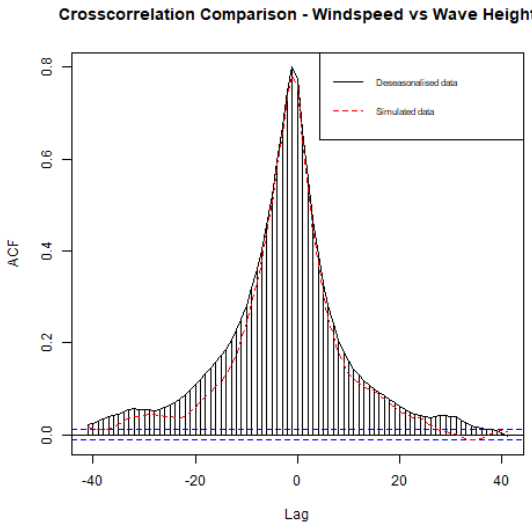

(a) Ws vs Hs cross-correlation function

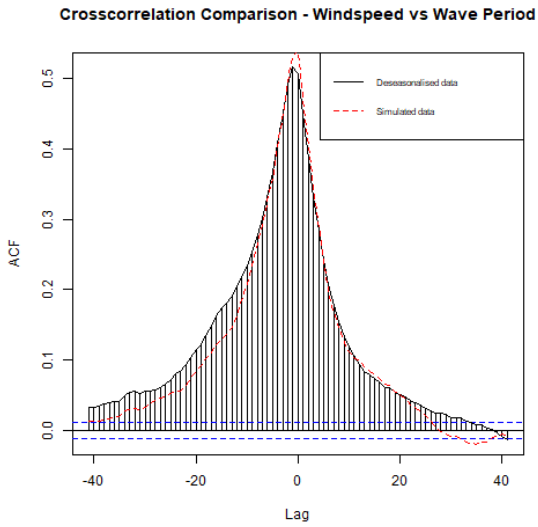

(b) Ws vs Tp cross-correlation function

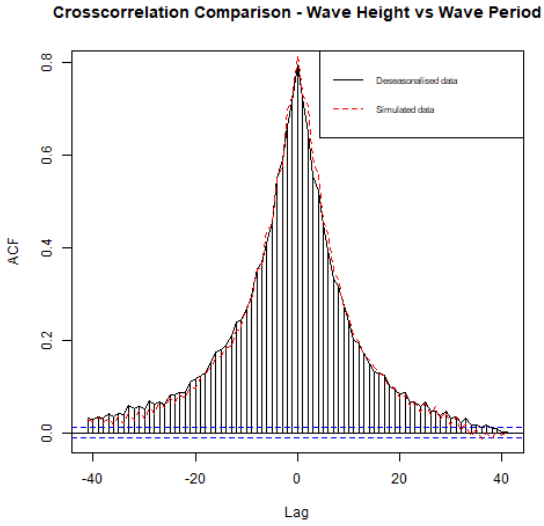

(c) Hs vs Tp cross-correlation function

Figure 4.3.2: Cross-correlation structure within the deseasonalised and simulated data.

4.3.3 Marginal Distribution Analysis On The Original Scale

As we have shown the multivariate Markov model appropriately represents the stationary data we now repeat the analysis with the seasonality added back into the data. Figure 4.3.3 again displays the marginal autocorrelation and q-q plots of the windspeed, significant wave height and wave period.

We observe that the decaying structure of the autocorrelation function is broadly reproduced as the plots capture the decaying structure of the acf found in the original data. All three plots highlight that longer term seasonality is successfully reincorporated as there are statistically significant acf values around 0.2 at lag forty (5 days in the past).

Windspeed autocorrelation plot

The windspeed acf for the synthetic data begins a gradual divergence from its value in the original dataset starting from the first lag and reaches a maximum difference from the observed data of 0.08 past lag ten.

Significant wave height autocorrelation plot

A similar result occurs for the significant wave height acf, however the overall divergence between the acfs for the original and synthetic data is smaller.

Wave period autocorrelation plot

The wave period acf could be considered the closest match to the original dataset as only the tail of the synthetic data acf appears to be an underestimation. Furthermore, some of the spikes in the function are captured within the synthetic results. It is noted that the underestimation of autocorrelation in the synthetic

data was expected as it also occurred albeit to a lesser degree on the stationary scale.

Q-Q plots

The q-q plots have steps in them since the original data was provided to us on a non-continuous scale, which leads to the original data having ties. This is a features that doesn't occur for the simulated data. For example, the windspeed data was given to the nearest metre/second. This creates the vertical banding effect seen in the q-q plots where there are noticeable gaps in the original data quantiles. The plots themselves suggest a good matching between the original and synthetic datasets. A closer inspection shows a small degree of curvature in the plots which is particularly visible for significant wave height. This occurs in the central quantiles where the red reference line forms a tangent to the vertical banding rather than an intersection.

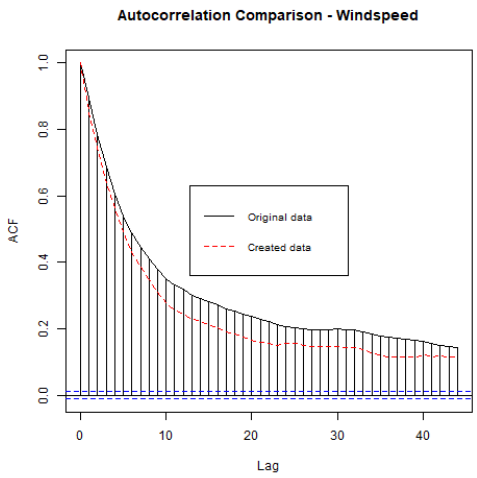

(a) Windspeed acf

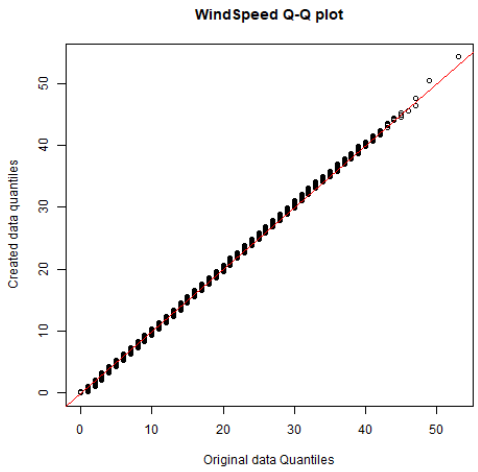

(b) Windspeed q-q plot

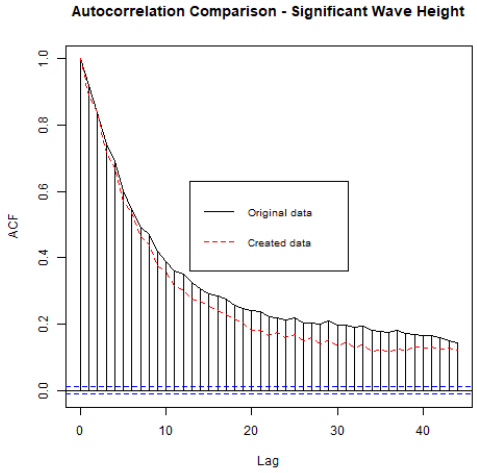

(c) Significant wave height acf

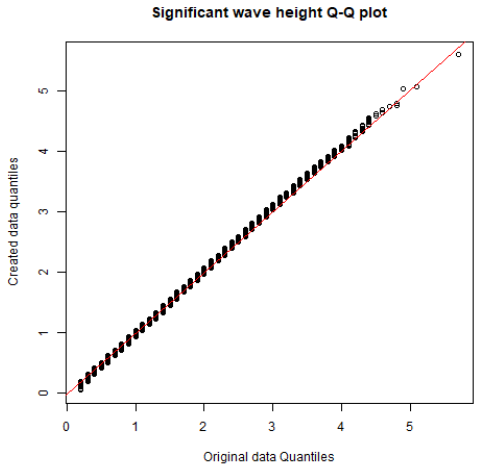

(d) Significant wave height q-q plot

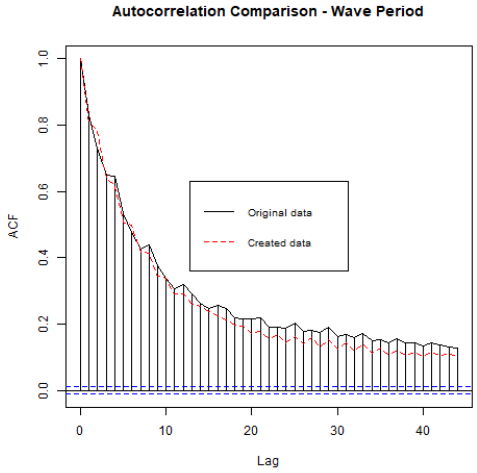

(e) Wave period acf

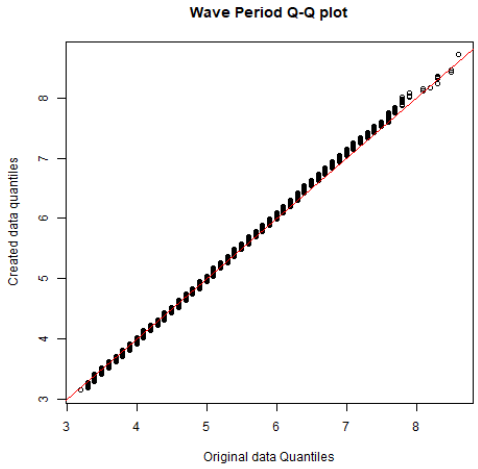

(f) Wave period q-q plot

Figure 4.3.3: Autocorrelation structure within the original and synthetic data.

4.3.4 Joint Distribution Analysis On The Original Scale

An examination of the cross-correlation between the windspeed, significant wave heights and wave periods is provided in Figure 4.3.4. Our approach appears to provide a good fit for the cross-correlation structure between the marginal variables. This is particularly evident for shorter range lags near zero which have very small errors that are typically less than 0.03. We note that all three cross-correlations contain slightly higher errors at longer lags, as shown in Figures 4.3.4a, 4.3.4b and 4.3.4b once seasonality is included.

Windspeed and significant wave height cross-correlation plot

The ccf between windspeed and significant wave height in the synthetic data appears to be a slight, but consistent, underestimation of the ccf in the original data. This does not exceed 0.08 at any point.

Windspeed and wave period cross-correlation plot

The ccf is again an underestimation of the original value at low lags with the windspeed vs wave period plot even exhibiting some asymmetry in this divergence between the positive and negative lags. This may be because of we are unable to distinguish effectively between the swell and wind-sea waves in our dataset.

Significant wave height and wave period cross-correlation plot

The third graph comparing significant wave height and wave period performs well as the acfs only significantly diverge outside of the lag [-10,+10] interval.

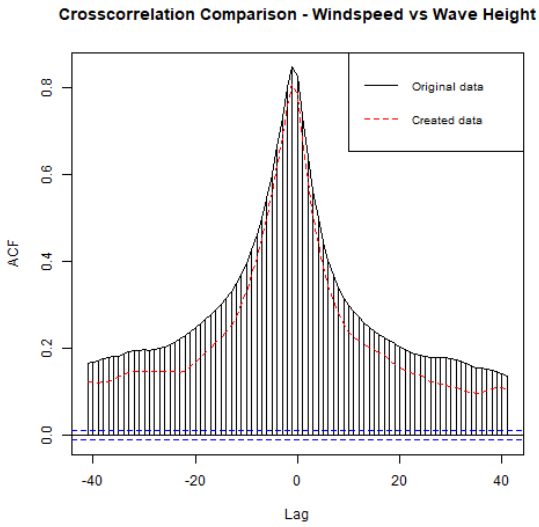

(a) Ws vs Hs cross-correlation function

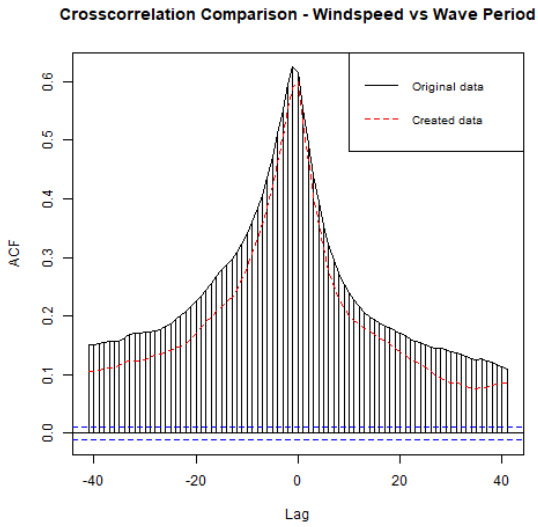

(b) Ws vs Tp cross-correlation function

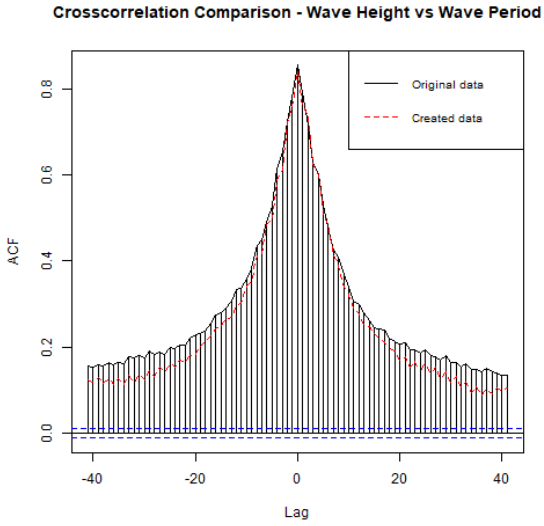

(c) Hs vs Tp cross-correlation function

Figure 4.3.4: Cross-correlation structure within the original and synthetic data.

4.3.5 Conditional Joint Distribution Analysis On The Original Scale

We now examine the joint distribution of the metocean variables with respect to the preceding values of the variables. Figures 4.3.5 - 4.3.7 depict each lagged joint distribution up to the first three time steps for the windspeed, significant wave heights and wave periods respectively. We include horizontal and vertical lines marking our chosen operating restrictions of 15m/s for windspeed and 1.5m for significant wave height in order to emphasise their presence in the body of the distribution. Observations from the original dataset are marked in black and have density contours coloured blue, whilst observations from the synthetic data are red with green density contours. The contours are overlaid on top of the observations for visual clarity.

The lagged joint distributions produce a good visual matching between the original and synthetic data across all of the variables. Figure 4.3.5a is clearly a symmetric distribution with the windspeed threshold of 15m/s occurring in the centre of the distribution. The density contours also closely match. Lagged joint distributions involving significant wave height are shown in Figure 4.3.6 and again resemble the original data well. In contrast to the windspeed threshold 1.5m corresponds to values larger than those in the area of highest density for the significant wave heights. We capture the spread of the data well as synthetic observations are created both in the bulk of the distribution and the upper tails of the marginals and joint distribution respectively. The distributions involving wave period in Figure 4.3.7 are the most irregular as our data contains a mixture of swell and wind-sea waves. This is

most obvious in Figure 4.3.7a where there is clear evidence of two distinct underlying processes: a wave period in excess of 6 seconds can be generated from either a low windspeed below 10m/s or a high windspeed above 20m/s. We note that our kernel density based method allows us to generate new observations but on average they do not greatly exceed the maximum values found in the original data.

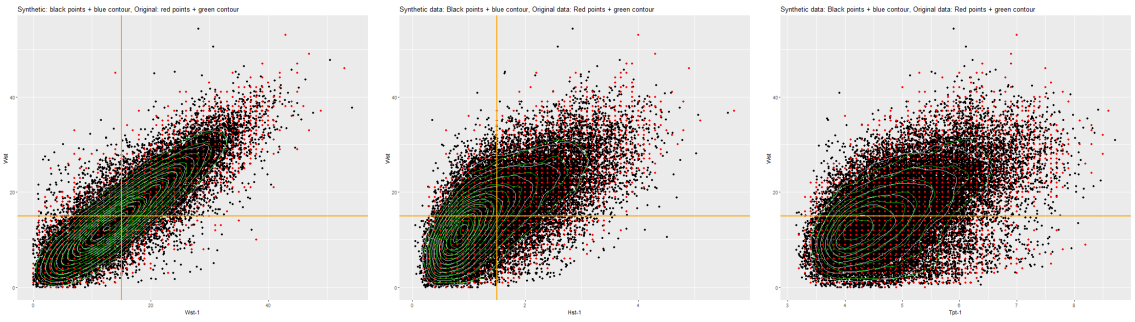

(a) Ws_{t-1} vs Ws_t (b) Hs_{t-1} vs Ws_t (c) Tp_{t-1} vs Ws_t

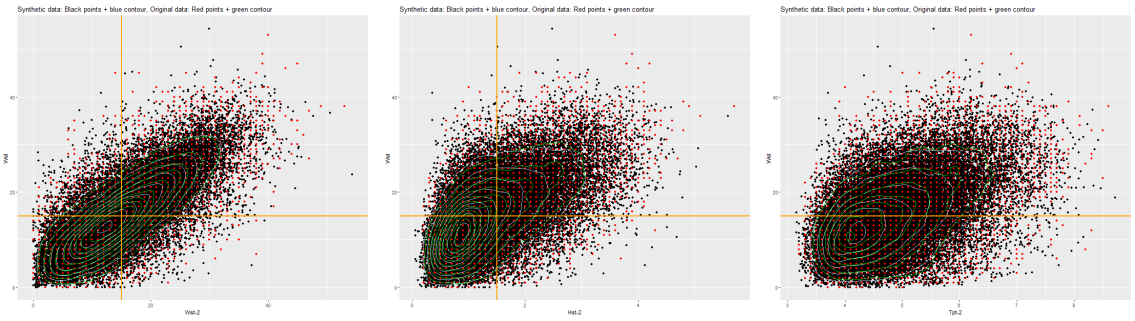

(d) Ws_{t-2} vs Ws_t (e) Hs_{t-2} vs Ws_t (f) Tp_{t-2} vs Ws_t

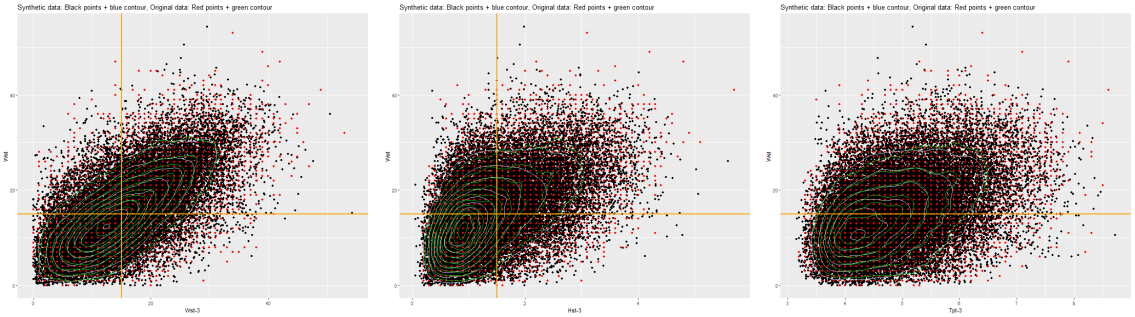

(g) Ws_{t-3} vs Ws_t (h) Hs_{t-3} vs Ws_t (i) Tp_{t-3} vs Ws_t

Figure 4.3.5: Conditional joint distributions involving windspeed for the original and synthetic data.

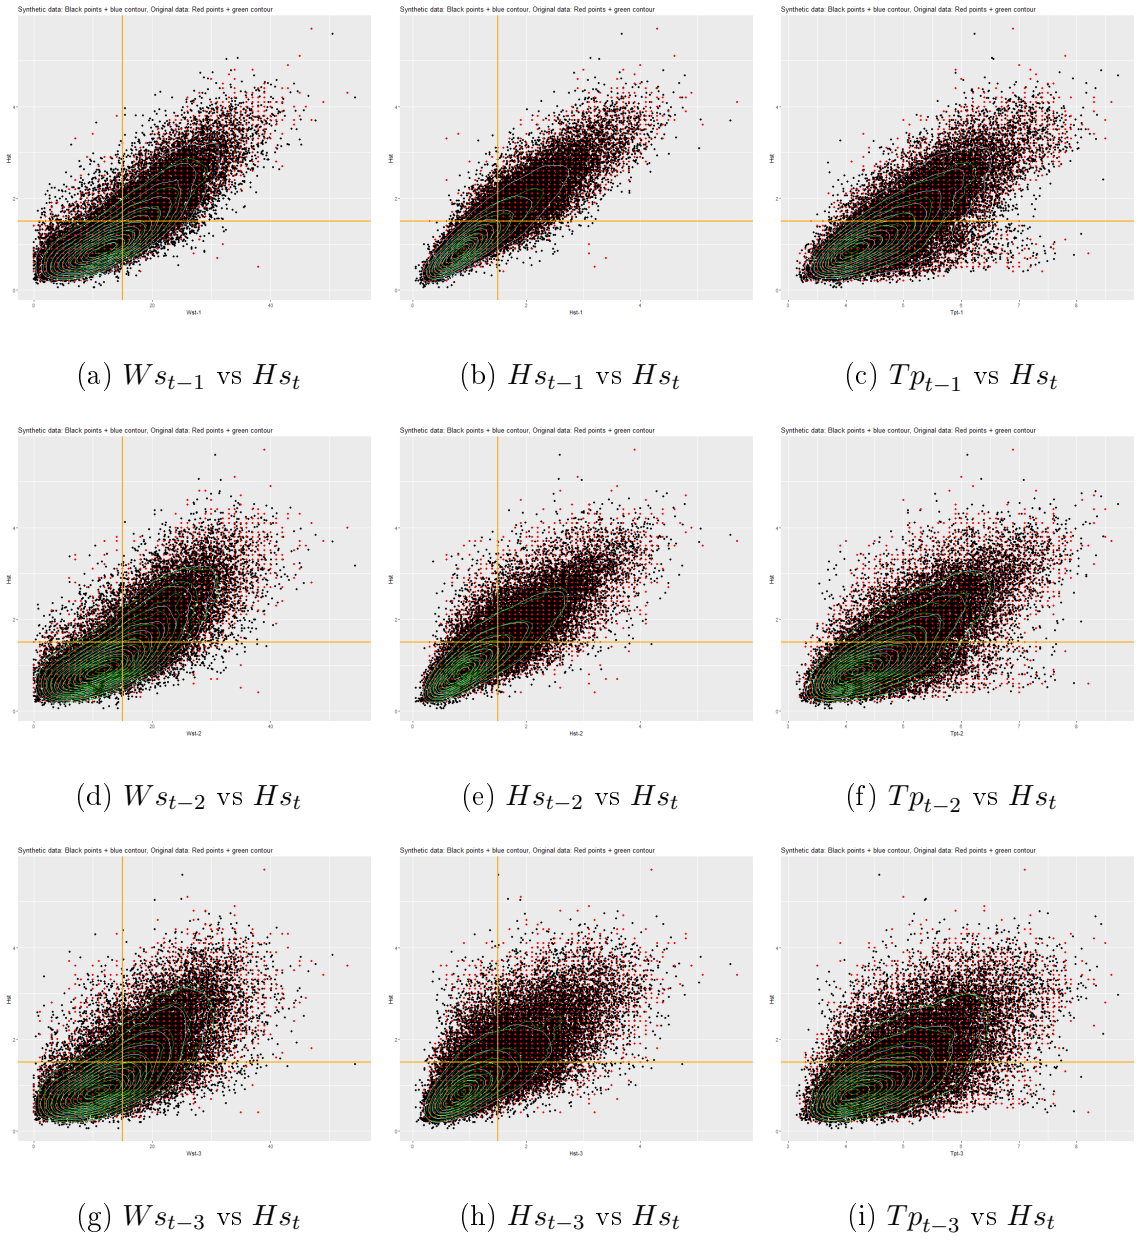

(a) Ws_{t-1} vs Hs_t (b) Hs_{t-1} vs Hs_t (c) Tp_{t-1} vs Hs_t

(d) Ws_{t-2} vs Hs_t (e) Hs_{t-2} vs Hs_t (f) Tp_{t-2} vs Hs_t

(g) Ws_{t-3} vs Hs_t (h) Hs_{t-3} vs Hs_t (i) Tp_{t-3} vs Hs_t

Figure 4.3.6: Conditional joint distributions involving significant wave height for the original and synthetic data.

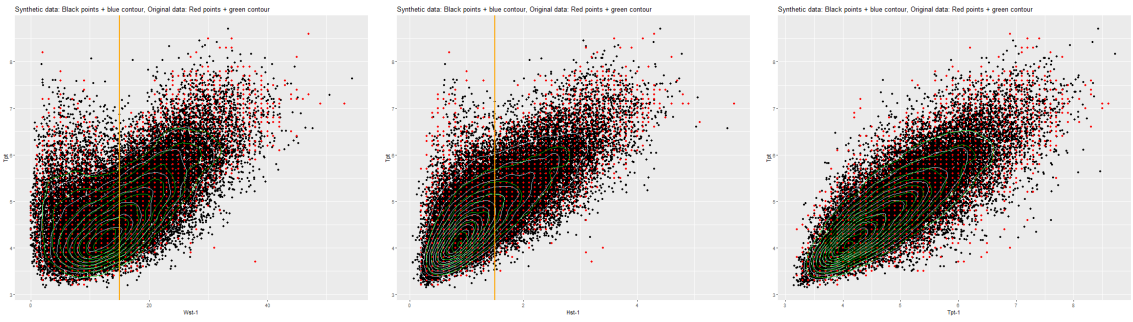

(a) Ws_{t-1} vs Tp_t (b) Hs_{t-1} vs Tp_t (c) Tp_{t-1} vs Tp_t

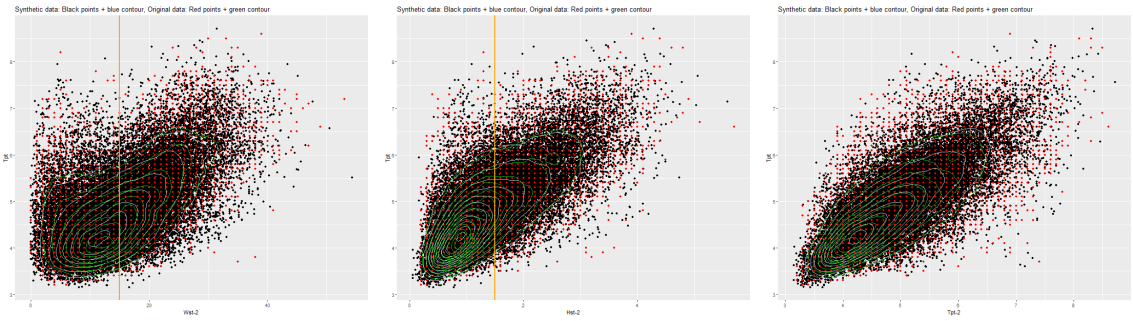

(d) Ws_{t-2} vs Tp_t (e) Hs_{t-2} vs Tp_t (f) Tp_{t-2} vs Tp_t

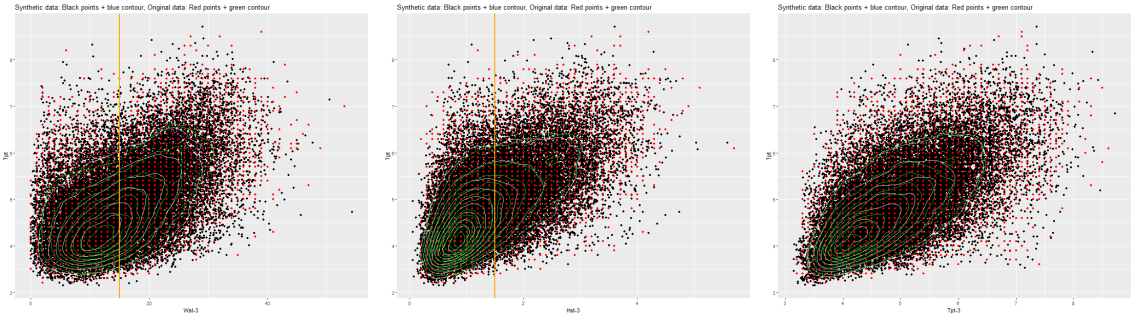

(g) Ws_{t-3} vs Tp_t (h) Hs_{t-3} vs Tp_t (i) Tp_{t-3} vs Tp_t

Figure 4.3.7: Conditional joint distributions involving wave period for the original and synthetic data.

4.3.6 Weather Window Statistics

Following the approach of Leontaris et al. (2016) we seek to further verify characteristics of our synthetic data that pertain to its practical domain of use. Given the aim of the process is to generate metocean data for access to offshore wind farms, we evaluate the frequency and persistence of weather windows based on typical operational limits. A common threshold for crew transfer vessel (CTV) access is 1.5m for significant wave height and a threshold of 15m/s for wind speed, however other thresholds could be easily examined since our method creates continuous data rather than pre-defined weather states. We compare the persistence and waiting time distribution between the original and synthetic datasets visually through the histograms presented in Figures 4.3.8a and 4.3.9a.

We observe that the distribution of weather window lengths in the synthetic dataset matches well with the distribution in the original dataset. The system spent a total of 43938 hours below the operating restrictions in the original data whilst 42180 hours occurred in the synthetic data, an approximate decrease of 4%. This suggests that the cumulative total hours of weather windows are well calibrated. It is possible to observe some minor differences between the original and synthetic data. It appears that we may overestimate the quantity of very short weather windows which can clearly be seen through the large spike at the three hour weather window value. Furthermore it is conceivable that we underestimate the frequency of very long weather windows (>100hrs), however as these are very rare events in the original data we do not consider it particularly important. This information is more clearly seen

through the q-q plot in Figure 4.3.8b which indicates a linear relation between the two quantiles of the two samples. The line initially follows a 45 degree gradient but diverges before continuing parallel to the red line indicating identical distributions. It should be noted that our approach is compared with the empirical weather windows and may still perform better than a different method for forecasting weather windows. Our model as such provides a conservative estimate on the long term distributions of weather window lengths.

The distribution of time the system waits before the occurrence of a weather window again shows a close fit to the original data. A total of 43632 and 45327 hours are spent waiting for a weather window in the original data and the synthetic data respectively. The short term synthetic durations appear to have a slight but consistent overestimation in their frequency. This is expected as the weather windows and waiting time until weather windows are complements of each other. It also appears that the data matches well even into the tail of the distribution. The q-q plot of the waiting times before a weather window reinforces this point with the observations fitting accurately to the background line indicating they follow the same distribution.

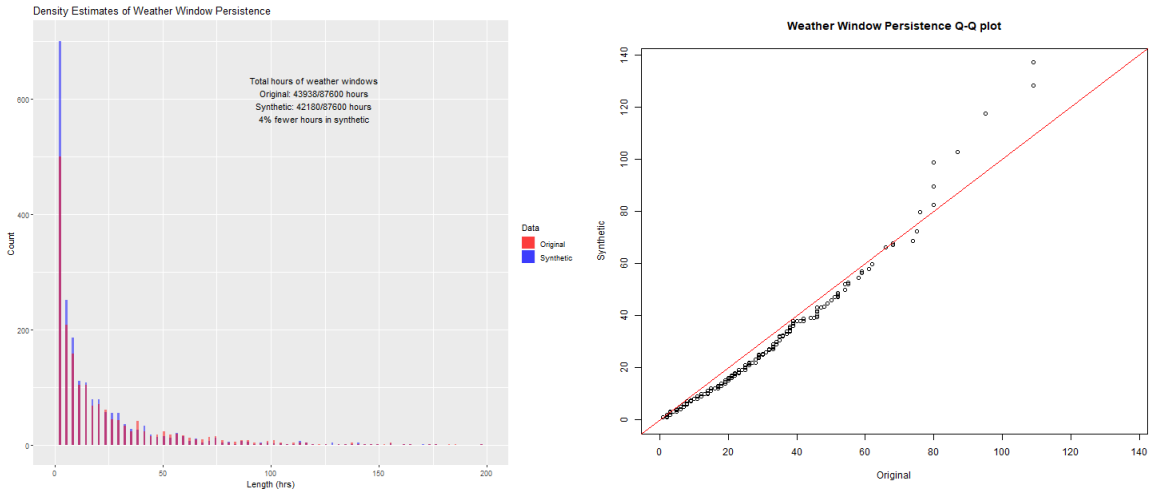

(a) Histogram (b) Q-Q plot

Figure 4.3.8: Weather window persistence comparison

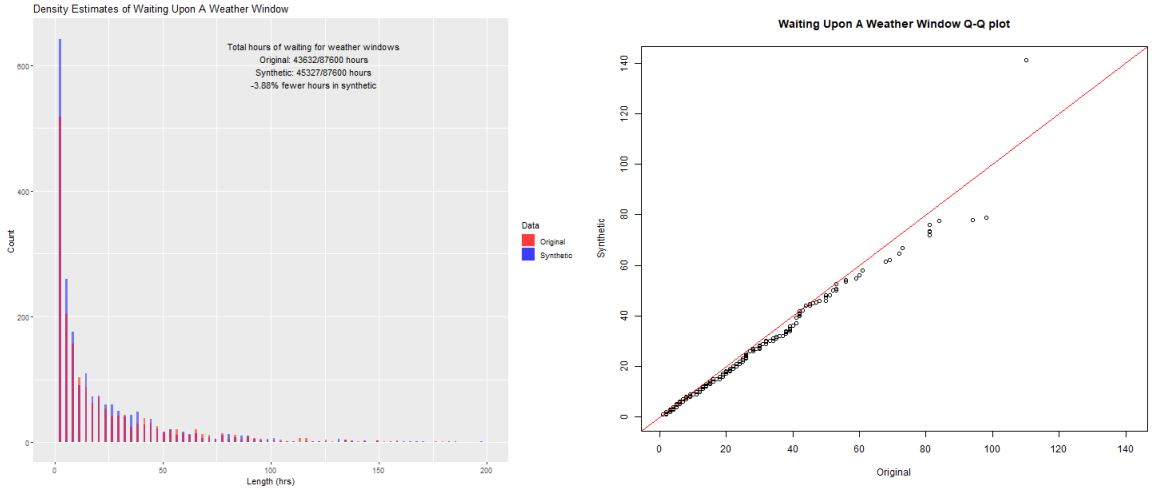

(a) Histogram (b) Q-Q plot

Figure 4.3.9: Waiting upon a weather window comparison

In order to gain deeper insights into the performance of our method we now examine the weather windows based on only one of the two thresholds previously used. Figures 4.3.10 and 4.3.11 are based on the results from only considering a windspeed

threshold of 15m/s, whilst Figures 4.3.12 and 4.3.13 are derived from a single significant wave height threshold of 1.5m. Our results indicate that the modelling inaccuracies in the weather window distributions are caused from both the joint distribution and marginal distribution of weather windows.

Weather windows purely based on windspeed

The q-q plot shows that the fit for pure windspeed weather windows is very good and matches the equivalent q-q plot for the full weather windows described earlier. This may be because it is a more restrictive threshold than using just significant wave height as seen by the total hours of weather windows being approximately 10,000 fewer. Further evidence for modelling the windspeed windows correctly is the fact that the total hours of weather windows only differs for the synthetic dataset by around 0.3%. The waiting time before a weather window is a noticeably worse fit, as there is a noticeable deviation around the central quantiles which is not present in the corresponding joint threshold q-q plot.

Weather windows purely based on significant wave height

The persistence of significant wave height weather windows is shown in Figure 4.3.12. The q-q plot presented is similar to that of the joint weather window for the first 15 percentiles, before a systematic miscalibration begins to occur. This is true despite the relative difference in axis scales. We note that the significant wave height weather windows have a few much longer weather windows in excess of 150 hours. The waiting time before a weather window is much closer to that

found for the joint weather threshold, although there is still a minor deviation similar to that found in Figure 4.3.11b. This could indicate a minor issue with our model as it occurs for both marginal variables.

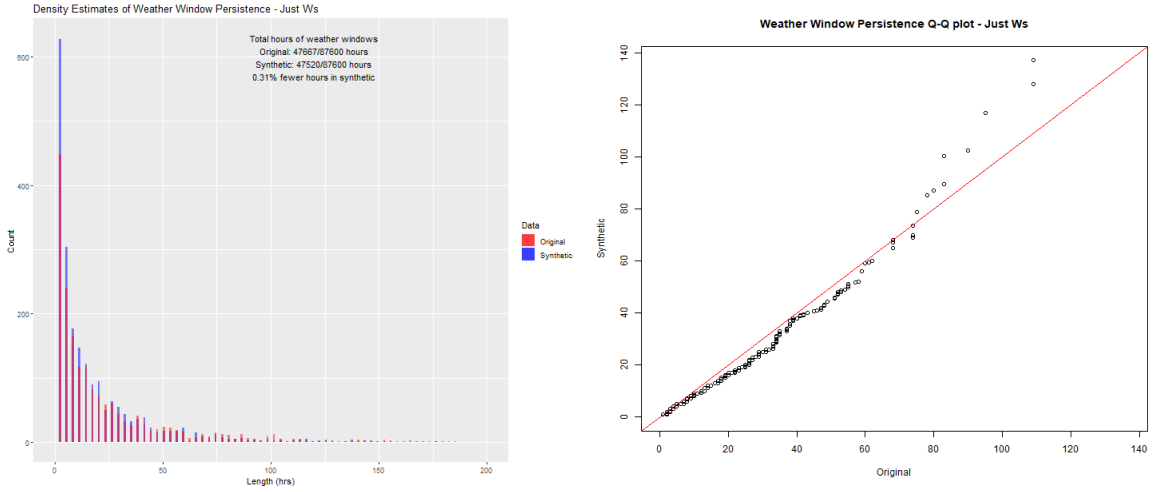

(a) Histogram (b) Q-Q plot

Figure 4.3.10: Windspeed weather window persistence comparison

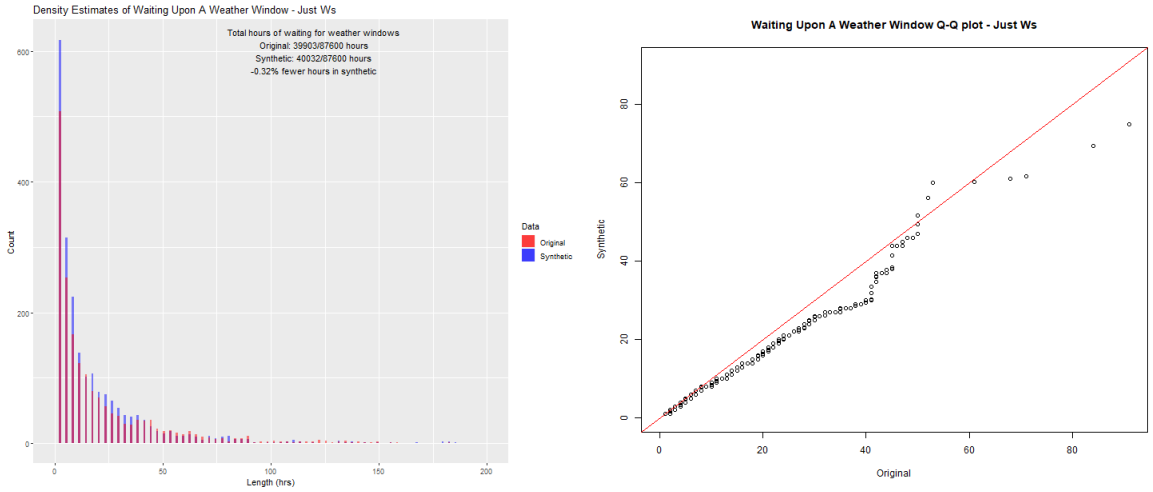

(a) Histogram (b) Q-Q plot

Figure 4.3.11: Waiting upon a windspeed weather window comparison

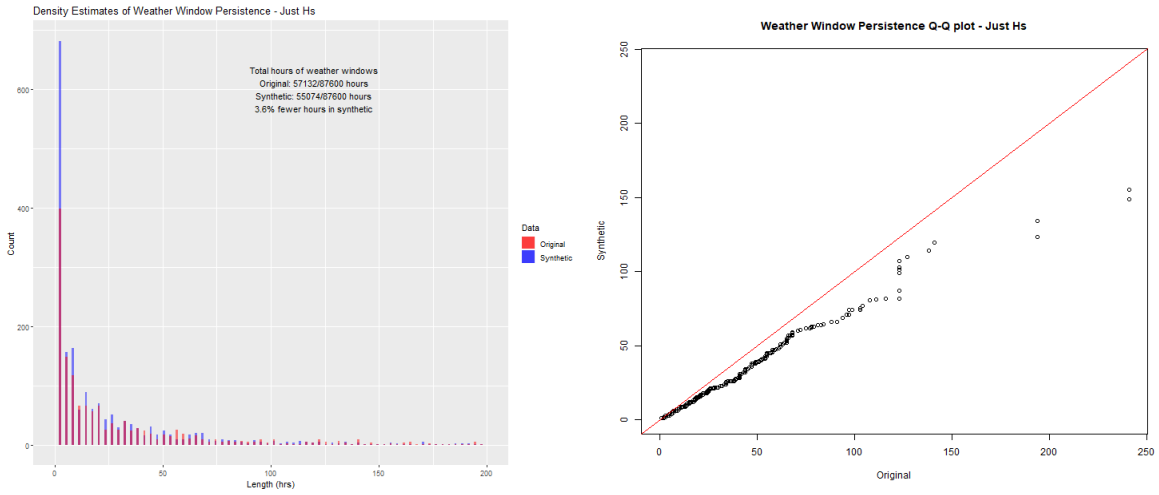

(a) Histogram (b) Q-Q plot

Figure 4.3.12: Significant wave height weather window persistence comparison

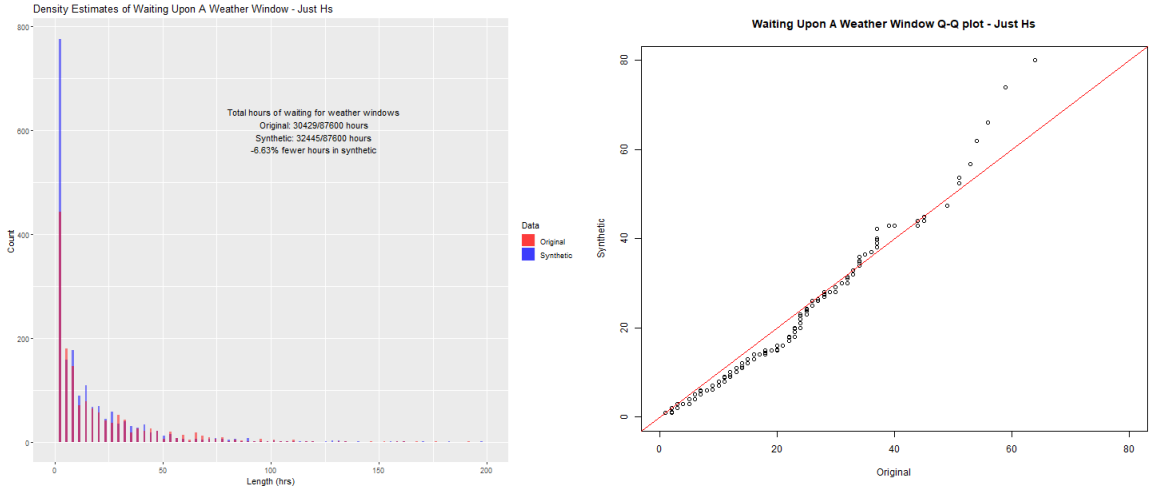

(a) Histogram (b) Q-Q plot

Figure 4.3.13: Waiting upon a significant wave height weather window comparison

Finally we now consider the effect of increasing the wind speed and significant wave height thresholds by 10% on our weather window comparisons. The new thresholds now become 16.5 m/s and 1.65 m for the windspeed and significant wave heights

respectively. This causes a significant change in the weather windows for both the original historical dataset and our synthetic dataset. A total of 48684 hours are spent inside weather windows in the historical data based on the increased weather thresholds. The slight underestimation highlighted on the original weather window graphs is more pronounced on Figure 4.3.14 for the weather window persistence. It therefore appears that we have a worse fit for the less restrictive weather windows. The q-q plot for the waiting time before a weather window for the increased thresholds, Figure 4.3.15b, also exhibits a similar fit to the previous graph. The lower quantiles are well calibrated with this extending well into the upper tail of the distribution.

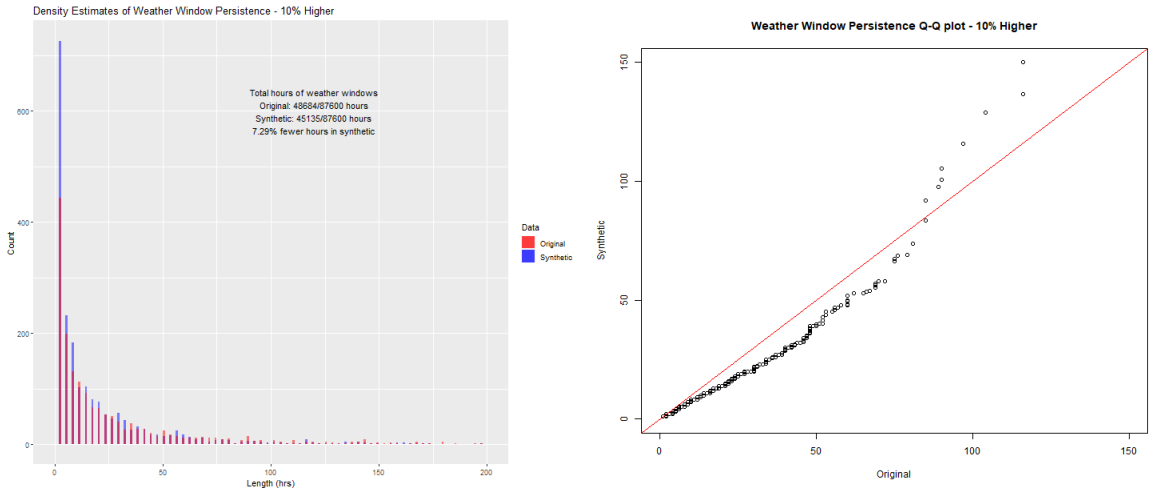

(a) Histogram

(b) Q-Q plot

Figure 4.3.14: Weather window persistence comparison with 10% higher thresholds

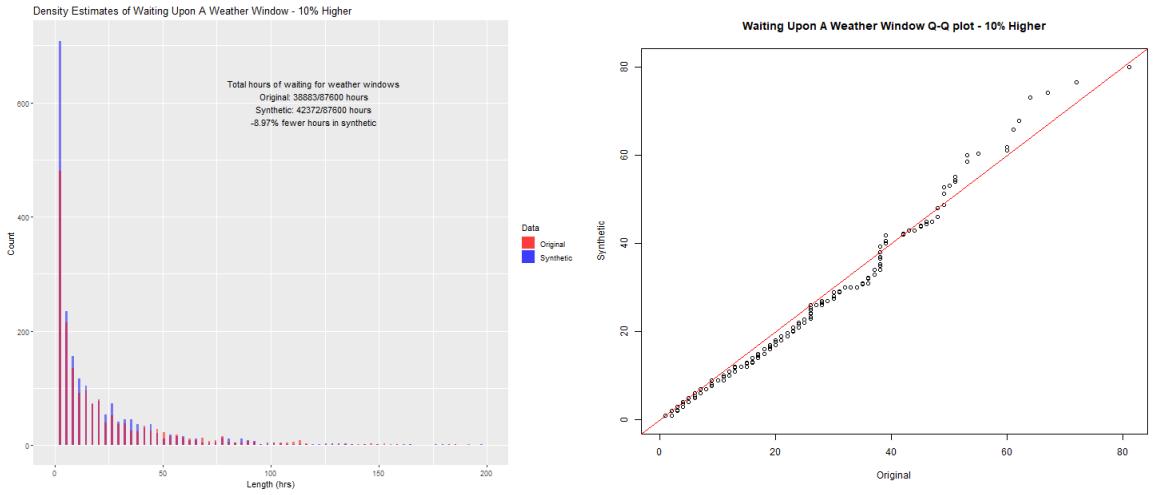

(a) Histogram

(b) Q-Q plot

Figure 4.3.15: Waiting upon a weather window comparison with 10% higher thresholds

4.3.7 Conclusions and Further Work

In this chapter we have proposed a method for generating a realistic joint distribution of metocean variables, such as wind speed and significant wave height, for which substantial values restrict offshore vessels access to offshore locations. The use of a joint distribution allows our model to fully capture the key features of metocean conditions, since the environmental variables involved are clearly dependent on each other. Our statistical model can be applied to any point in the year (or future year) and takes into account the autocorrelation and cross-correlation between variables whilst accounting for seasonal trends. An important method for validating the resulting time series generated by our method is through the frequency and persistence of weather windows. This has been performed on our target dataset with our results indicating a good fit when a graphical model is fitted to a tenth order Markov chain with the

removal of several higher order terms.

For future work it would be useful to compare the persistence of and waiting time before weather windows computed from our synthetic time series with other methods. We currently compare with the empirical weather window distributions of the dataset provided to us, so it would be of interest to see how our method compares with alternative methods, such as the Markov chain approach of Bruijn et al. (2019), that purely forecast weather windows. We could also investigate how our validation measures differ as we include more environmental variables such as wind and wave direction.

Chapter 5

Conclusions and Future Work

In this thesis we have presented a new exact mathematical improvisation model of the offshore maintenance routing problem and a custom heuristic method for solving larger problem instances. We have also provided a statistical model for generating alternative metocean scenarios based on historical data from offshore wind farms. These models have had input from JBA Consulting with the goal of extending the scope of their ForeCoast ® Marine software. We now summarise our main conclusions and outline some ideas for further work.

5.1 Optimisation Models

In Chapter 2 we presented our mixed integer programming formulation of the routing and scheduling problem present at offshore wind farms. We included important characteristics such as distinction between corrective and preventive maintenance tasks and the possibility of task preemption. The relative performance of the model was

assessed in a variety of resource restricted situations that highlights the additional complexity from our inclusion of task prioritisation.

One avenue of future research could involve the relaxation of the assumption that only one task can appear on a turbine in the planning horizon. An initial approach could be to bundle multiple concurrent tasks on a given turbine into a single bulk task specified by the minimum number of technicians and work duration required to complete all said tasks. This would imply that the tasks are completed in parallel, however it could be possible to extend the model to allow for technicians moving around the system to help perform specific tasks within a task bundle. A better global solution could be obtained if technicians are not required to work on an entire activity bundle and can perform a subset of tasks before being transported to another turbine. Further complications could arise if predecessor-successor relations are included between tasks within activity bundles.

Another aspect of the model that could be enhanced is the question of how to quantify the relative benefit of performing preventive maintenance tasks versus immediate corrective maintenance tasks. We have proceeded on the basis that corrective maintenance is significantly more valuable than preventive maintenance, so that our solutions choose to heavily prioritise completing corrective maintenance tasks as soon as possible. This could be explored with either a strict lexicographic approach or a proper evaluation of the long term benefit of preventive maintenance. Given that our focus was on day to day operations we have not thoroughly investigated this, but our modelling framework is flexible enough to consider its integration as another future direction of research. This could form the basis of a coveted condition based

maintenance strategy with the inclusion of a turbine state of health index mapped to expected downtime losses if faults do occur.

Our novel technician-vessel maximum safety range constraint could be evolved to better capture offshore behaviour. In practice there will be some uncertainty in vessel paths so that they are not always the direct straight lines we assume them to be. It is likely that routes will contain some natural curvature which we could build upon in the safety constraint. If deviations to straight vessel paths are allowed forming curved arcs there may be an added benefit of increasing the time vessels spend within the safety radius. We could expand the mathematical model or an equivalent sub-problem to determine whether it is beneficial to deviate these routes to prolong periods of vessel coverage to offshore technicians. This optimisation procedure would also have to consider if several turbines are covered by the resulting vessel relocation.

Our investigation of the merits of a simple stochastic solution to deal with the uncertainty in weather conditions could also be expanded. It would be expected that a full evaluation on a longer term scenario would yield greater benefits in the value of a stochastic solution. The most interesting examples would likely be the borderline cases where there is a realistic chance of either exceeding or falling short of an operating threshold. Other instances with conditions expected to persist below thresholds could see value from a focus on the uncertainty in the completion time of corrective maintenance tasks, which would likely have a greater impact on the problem.

In Chapter 3 we provided an alternative formulation of the offshore wind farm maintenance routing problem which we decompose into two sub-problems. The first

ordering sub-problem is associated with generating the sequence of visits made by each vessel, whilst the timing sub-problem determines the exact arrival and departure times of vessels in such a way that the most valuable balance of corrective and preventive maintenance is performed. An adaptive large neighbourhood search heuristic is employed to solve the problem and contains 7 removal operators and 3 repair operators. The performance of different acceptance criteria are statistically evaluated with our results confirming the benefit of the adaptive layer to operator selection.

A simple extension to our model lies in the incorporation of technician skill types. Most tasks requiring technicians actually require more specialised technicians with different skill types such as electrical or mechanical technicians. Our decomposition structure means that restrictions on the amounts of skilled technicians can be efficiently evaluated as additional feasibility checks in the timing sub-problem. Therefore not only could the total amount of skilled technicians be restricted but also the equipment loads of vessels if desired.

Our basic timing sub-problem heuristic was implemented as we experienced difficulties in the practical implementation of solving the optimisation model repeatedly within the adaptive large neighbourhood search. It is anticipated that future work could overcome these difficulties to achieve the benefits we previously outlined, since it amounts effectively to solving a small linear programming model. For example, solving the timing sub-problem exactly would allow for a simple incorporation of technicians, of different abilities, to be picked up and dropped off by different vessels. If binary variables associated with task completion are considered within the timing sub-problem then significantly more routes could be pruned from the sequencing prob-

lem as the timing sub-problem would show which tasks are more valuable to perform and which are optimal to omit. Alternatively time could be dedicated to developing an improved heuristic for the timing problem.

Another interesting extension would be to investigate whether a similar mathematical model and heuristic approach could be employed to consider the routing problem of offshore SOVs and motherships. These vessels are capable of staying offshore overnight and do not require our assumption that vessels must return to port at the end of each shift. These craft are both capable of traversing the wind farm to transport technicians to maintenance tasks, and in the case of motherships launch smaller daughter vessels to complete tasks in different areas. The mathematical challenge of routing daughter craft in conjunction with larger motherships could provide another research direction.

5.2 Statistical Modelling

Chapter 4 detailed a statistical model capable of generating a realistic joint distribution of the metocean variables for any point in the year (or future year). This provides the basis for simulating alternative representations of short-term meteorological and oceanographic scenarios offshore.

Typically planned offshore activities, such as wind turbine installation, are evaluated on the basis of weather forecasts or a limited range of historical weather data. These are valid only as far as one month ahead at most, whereas a statistical model allows for inferences to be made much further ahead in time than using pure weather

forecasts.

The statistical model we have developed could be used to generate significantly more metocean scenarios from the historical data. Multivariate kernel density estimation is used in conjunction with graphical modelling to reproduce the temporal dependence structure of the joint distributions of key metocean variables such as wind speed and significant wave height. We also formally introduced the definition of a weather window and waiting time before a weather window as key metrics for offshore wind farm accessibility.

We note that our wave modelling is somewhat naive in the statistical model as our method makes no distinction between swell and wind-sea waves. Whilst the model presented in Chapter 4 somewhat captures these different wave generation processes through kernel density estimation, Figures 4.3.7a and 4.3.7b show that they could be better represented as a mixture of two distributions. These are the locally generated swell waves and the longer-ranged wind-sea waves. Swell waves tend to have longer wave periods given a particular significant wave height than wind-sea waves. We can therefore infer from Figure 4.3.7b that the observations with wave period greater than 6 seconds and significant wave height less than 1.5m are examples of swell waves. Our existing approach makes no distinction between the two and so conditional distributions are not informed by which wave state they are currently in. We chose to ignore this as the dataset provided to us did not split significant wave heights into swell and wind-sea components. A possible extension to the model could be to implement methods for partitioning significant wave heights into their components, whilst also minimising the amount of switching between states to provide

a more realistic depiction of the sea conditions.

A drawback of our current approach is that it is heavily reliant on the existing observations found in our dataset. The kernel density estimation approach will accurately interpolate values between existing datapoints, but is unlikely to extrapolate the information to generate values outside of the dataset. In practice there exists a small chance of abnormally large or extreme events occurring offshore which might be expected to occur every few decades. These extreme values will not have a significant effect on the durations of calm or stormy weather, but will help to make our synthetic data more realistic and thus could be of value for other planning scenarios. One way of modelling these extreme values is by fitting a parametric model to the tail of the kde.

The generalised Pareto distribution is a parametric tail model that is theoretically justified by extreme value theory, (Coles et al., 2001). It is used to model the occurrences of infrequent and unusually large observations. Informally it models the exceedance of observations above a given threshold level. The conditional excess distribution function can be used to replace the upper tail of a probability distribution with that of a generalised Pareto distribution. This adjusts the kde to include the possibility of generating values outside the range of the data using extreme value theory instead of the arbitary choice of a kernel function. The replacement tail value $F(y)$ above a threshold u is given by,

$$F(y) = 1 - \lambda \left(1 + \xi \left(\frac{y-u}{\sigma}\right)\right)^{-\frac{1}{\xi}}, \quad y > u \tag{5.2.1}$$

with $\lambda = 1 - \hat{F}(u)$ where $\hat{F}$ is the kernel estimated distribution function from Eq. 4.2.8

for a given time t, so λ is approximately the fraction of observations within the sample that exceed the threshold. The model parameters are $\sigma > 0$ the scale parameter and $\xi \in \mathcal{R}$ the shape parameter. Both the shape and scale parameters are estimated based on maximum likelihood estimates for the sample's fit by a generalised Pareto distribution in the upper tail above the threshold. This concept has been illustrated with the upper tail as it is of more interest to us. The same procedure could also be applied to the lower tail, albeit with the caveat that there is a physical lower bound of zero on the windspeed, significant wave heights etc.

The choice of threshold value u is informed by the desire to maximise the information within the tail subject to it remaining a good fit to the sample. A basic method to determine u is through visual inspection of fits and quantile-quantile (Q-Q) plots. A more sophisticated approach involves investigating the scale, σ and shape, ξ parameter stabilities for different potential threshold choices via the method presented in Coles et al. (2001). The aim is to select the smallest threshold u, such that $\xi(u)$ remains within the confidence interval for $\xi(u^*), \forall u^* \geq u$.

Our statistical model relies on having access to information derived from the offshore wind farm site that it will be applied to. Gauged wind farms regularly record the required metocean data by employing one or two wind and wave buoys in the wind farm. In the rare case that an offshore site is ungauged, the historical data needed for our statistical model will be lacking. This means that we will be unable to perform kernel density estimation based on past information from the desired site and need to adapt our approach. One option would be to borrow information from nearby gauged sites and apply spatial interpolation methods to estimate the variables at the desired

location. Alternatively we could use spatial statistical methods to account for the spatial variation in the offshore conditions and to give estimates at ungauged sites.

Several of our key assumptions about the data will also start to break down if climate change has a significant impact on the offshore environment. We have constructed our statistical model under the belief that the deseasonalised data can be considered identically distributed over time. The kernel density estimation process currently pools local data and their equivalent data points in other years of the dataset. The latter concept cannot be implemented if we assume some underlying upwards trend due to climate change. Recent research has indicated that average wind speeds have risen by 1.5 m/s and wave heights by 30cm over the last 30 years in some locations, (Young and Ribal, 2019). Extending our model to account for the effect of climate changes within the historical time series available would be of great value to the offshore wind industry, particularly given that the lifetime of offshore wind farms is expected to be 20-25 years. This is clearly exceedingly difficult to model and would likely require external information produced by another climate model creating projections into the future. This could provide an estimate of both of the upwards trend and rate of change of the trend which our model could use as an input. Thus our hybrid model could be informed of the global patterns and the localised effects for a given offshore site.

A long-term continuation of this project could involve the integration of our model with short term weather forecasts available for future time periods. These short-term weather forecasts are known to be highly accurate for a few days into the future meaning that incorporating this information into the model would likely produce a

more accurate prediction of metocean conditions in the short term. Our statistical model gives information at longer scales which weather forecasts cannot, so combining both gives the option of combining both short and long term information in one model. A joint statistical and numerical model should be able to quantify the uncertainty in short term numerical predictions and balance it with longer term climate patterns such as monthly seasonality. Furthermore, if an ensemble forecasting approach cannot produce forecasts for all the desired variables or if historical data are missing due to sensor failures, we can utilise the statistical approach to make up the shortfall. The synthesis of numerical and statistical models could be achieved with a Bayesian framework using forecasts as prior beliefs to update an existing statistical model.

In Section 4 we develop a statistical model of the joint distribution of the metocean variables that restrict access to and work on offshore wind turbines. This was done in recognition of the need to capture the time dependent nature of the metocean conditions incorporated into the stochastic extension of our model in Section 2. This can be considered as part of a rolling horizon approach that we outlined earlier to consider the impact of uncertain metocean conditions on offshore wind farm operations. Rather than considering a global optimisation taking into account every possible combination of weather scenarios across multiple shifts, we instead attempt to create a good plan for the first shift taking into account some of the future shifts. This amounts to solving the optimisation model on a daily basis with uncertain weather conditions, but only using the solution for the current shift in practice. After this the schedule rolls over to the next day based on the work completed and is updated with any additional new tasks or predictions of offshore conditions entering the system.

In a deterministic model the metocean conditions and associated parameters were assumed to be deterministic. In practice these parameters are probabilistic and as such the uncertainty surrounding them should be propagated through a proper model. A simple stochastic model based on uncertain metocean conditions for a single shift of activities could be modelled as thus. Let us denote the metocean conditions as $\boldsymbol{\theta}$ which are themselves uncertain. The simulation of offshore conditions currently finds the best estimate of these parameters $\hat{\boldsymbol{\theta}}$ and uses this to inform a model of offshore maintenance operations. Decisions about which actions x to make are compared through the use of a test statistic T (likely to be the objective function of an optimisation model) given the metocean conditions, $T(x|\hat{\boldsymbol{\theta}})$. As we wish to incorporate the uncertainty in the metocean conditions into the evaluation of the costs of choosing actions x, the metocean conditions become the prior distribution in the Bayesian sense $\pi(\boldsymbol{\theta})$. Then the cost of choosing actions x are given below as,

$$T(x) = \int_{\boldsymbol{\theta}} T(x|\boldsymbol{\theta})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}. \tag{5.2.2}$$

This framework can be adjusted to include information related to the weather forecasts y. The impact of short or long term weather forecasts could be conceptualised with the change of $\pi(\boldsymbol{\theta}) \rightarrow \pi(\boldsymbol{\theta}|y)$ in Equation 5.2.2.

In practice we would evaluate the rolling horizon model with the knowledge of forecasts and expected work patterns for several days ahead. This would allow us to consider the impact of scheduling in a realistic setting involving multiple shifts. Furthermore as we have shown in Section 4 there is a large degree of seasonality present in the metocean conditions which has an impact on offshore operations. This

can be accounted for by changing our prior beliefs on $\boldsymbol{\theta}$ to become time dependent $\boldsymbol{\theta}_t$. We can also include some notion of time dependency to the actions performed given that tasks can be left partially incomplete between shifts in our multiple shift models. It is also likely that new tasks can be generated within the planning horizon so that a longer plan involves the actions in a subsequent periods. This leads to $T(x|\boldsymbol{\theta})$ becoming $T(x, t|\hat{\boldsymbol{\theta}}_t)$ once we consider later shifts at time t.

Bibliography

A2Sea. A2sea. sea installer. http://www.a2sea.com/fleet/sea-installer/, 2016. Accessed: 10-08-2016.

Yossiri Adulyasak, Jean-François Cordeau, and Raf Jans. Optimization-based adaptive large neighborhood search for the production routing problem. *Transportation Science*, 48(1):20–45, 2012.

Deniz Aksen, Onur Kaya, F Sibel Salman, and Özge Tüncel. An adaptive large neighborhood search algorithm for a selective and periodic inventory routing problem. *European Journal of Operational Research*, 239(2):413–426, 2014.

Alicat Workboats. Offshore wind. dalby offshore returns to alicat workboats for another wfsv. http://www.offshorewind.biz/2016/01/14/dalby-offshore-returns-to-alicat-workboats-for-another-wfsv/, 2016. Accessed: 10-08-2016.

K Anastasiou and C Tsekos. Persistence statistics of marine environmental parameters from markov theory, part 1: Analysis in discrete time. *Applied ocean research*, 18 (4):187–199, 1996.

Francois Besnard, Michael Patriksson, Ann-Brith Strömberg, Adam Wojciechowski, and Lina Bertling. An optimization framework for opportunistic maintenance of offshore wind power system. In *Proc. of IEEE PowerTech*, 2009.

François Besnard, Michael Patriksson, A Stromberg, Adam Wojciechowski, Katharina Fischer, and Lina Bertling. A stochastic model for opportunistic maintenance planning of offshore wind farms. In *PowerTech, 2011 IEEE Trondheim*, pages 1–8. IEEE, 2011.

Burak Bilgin, Ender Özcan, and Emin Erkan Korkmaz. An experimental study on hyper-heuristics and exam timetabling. In *International Conference on the Practice and Theory of Automated Timetabling*, pages 394–412. Springer, 2006.

Willem EL Bruijn, Jolien Rip, Antoon JH Hendriks, Pieter HAJM van Gelder, and Sebastiaan N Jonkman. Probabilistic downtime estimation for sequential marine operations. *Applied Ocean Research*, 86:257–267, 2019.

Jose Caceres-Cruz, Pol Arias, Daniel Guimarans, Daniel Riera, and Angel A Juan. Rich vehicle routing problem: Survey. *ACM Computing Surveys (CSUR)*, 47(2): 1–28, 2014.

Fatih Camci. Maintenance scheduling of geographically distributed assets with prognostics information. *European Journal of Operational Research*, 245(2):506 – 516, 2015. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2015.03.023. URL `http://www.sciencedirect.com/science/article/pii/S0377221715002301`.

Ann Melissa Campbell, Dieter Vandenbussche, and William Hermann. Routing for relief efforts. *Transportation Science*, 42(2):127–145, 2008.

Carbon Brief. Analysis: Uk renewables generate more electricity than fossil fuels for first time. `https://www.carbonbrief.org/analysis-uk-renewables-generate-more-electricity-than-fossil-fuels-for-first-ti` 2019. Accessed: 02-04-2020.

James Carroll, Alasdair McDonald, and David McMillan. Failure rate, repair time and unscheduled o&m cost analysis of offshore wind turbines. *Wind Energy*, 2015.

CATAPULT. Cost of offshore wind energy falls sharply. `https://ore.catapult.org.uk/press-release/cost-of-offshore-wind-energy-falls-sharply/`, 2015. Accessed: 01-09-2016.

José E Chacón, Tarn Duong, and MP Wand. Asymptotics for general multivariate kernel density derivative estimators. *Statistica Sinica*, pages 807–840, 2011.

I-Ming Chao, Bruce L Golden, and Edward A Wasil. The team orienteering problem. *European journal of operational research*, 88(3):464–474, 1996.

Marielle Christiansen, Kjetil Fagerholt, Bjørn Nygreen, and David Ronen. Ship routing and scheduling in the new millennium. *European Journal of Operational Research*, 228(3):467–483, 2013.

Stuart Coles, Joanna Bawa, Lesley Trenner, and Pat Dorazio. *An introduction to statistical modeling of extreme values*, volume 208. Springer, 2001.

Global Wind Energy Council. Global wind reports (gwec), 2019.

Lijuan Dai, Magnus StÃělhane, and Ingrid B. Utne. Routing and scheduling of maintenance fleet for offshore wind farms. *Wind Engineering*, 39(1):15–30, 2015. doi: 10.1260/0309-524X.39.1.15. URL `https://doi.org/10.1260/0309-524X.39.1.15`.

Yalcin Dalgic, Iraklis Lazakis, and Osman Turan. Vessel charter rate estimation for offshore wind o&m activities. 2013.

Yalcin Dalgic, Iraklis Lazakis, Osman Turan, and Sol Judah. Investigation of optimum jack-up vessel chartering strategy for offshore wind farm o&m activities. *Ocean Engineering*, 95:106–115, 2015.

R. Dawid, D. McMillan, and M. Revie. Heuristic algorithm for the problem of vessel routing optimisation for offshore wind farms. *The Journal of Engineering*, 2017 (13):1159–1163, 2017. doi: 10.1049/joe.2017.0511.

Rafael Dawid, David McMillan, and Matthew Revie. Decision support tool for offshore wind farm vessel routing under uncertainty. *Energies*, 11(9), 2018. ISSN 1996-1073. doi: 10.3390/en11092190. URL `http://www.mdpi.com/1996-1073/11/9/2190`.

Giulia De Masi, Roberto Bruschi, and Michele Drago. Synthetic metocean time series generation for offshore operability and design based on multivariate markov model. In *OCEANS 2015-Genova*, pages 1–6. IEEE, 2015.

Martin Desrochers and Gilbert Laporte. Improvements and extensions to the miller-tucker-zemlin subtour elimination constraints. *Operations Research Letters*, 10(1): 27–36, 1991.

José Antonio Domínguez-Navarro, Iain Dinwoodie, and David McMillan. Statistical forecasting for offshore wind helicopter operations. In *Probabilistic Methods Applied to Power Systems (PMAPS), 2014 International Conference on*, pages 1–6. IEEE, 2014.

Gunter Dueck. New optimization heuristics: The great deluge algorithm and the record-to-record travel. *Journal of Computational physics*, 104(1):86–92, 1993.

Energy and House of Commons Climate Change Committee. 2020 renewable heat and transport targets. 2016.

Morten W Fagerland and Leiv Sandvik. The wilcoxon–mann–whitney test under scrutiny. *Statistics in medicine*, 28(10):1487–1497, 2009.

Fiberline. Fiberline components. helihoist platforms. `https://fiberline.com/helihoist-platforms`, 2016. Accessed: 10-08-2016.

Martina Fischetti and David Pisinger. Mathematical optimization and algorithms for offshore wind farm design: An overview. *Business & Information Systems Engineering*, Apr 2018. ISSN 1867-0202. doi: 10.1007/s12599-018-0538-0. URL `https://doi.org/10.1007/s12599-018-0538-0`.

Renewable Energy Focus. Renewable energy focus. offshore wind farm maintenance vessel designed. `http://www.renewableenergyfocus.com/view/15310/offshore-wind-farm-maintenance-vessel-designed`, 2011. Accessed: 09-08-2016.

AurÃľlien Froger, Michel Gendreau, Jorge E. Mendoza, Eric Pinson, and Louis-Martin Rousseau. A branch-and-check approach for a wind turbine maintenance scheduling problem. *Computers & Operations Research*, 88:117 – 136, 2017. ISSN 0305-0548. doi: https://doi.org/10.1016/j.cor.2017.07.001. URL `http://www.sciencedirect.com/science/article/pii/S030505481730165X`.

Philippe Grangier, Michel Gendreau, Fabien Lehuédé, and Louis-Martin Rousseau. An adaptive large neighborhood search for the two-echelon multiple-trip vehicle routing problem with satellite synchronization. *European Journal of Operational Research*, 254(1):80–91, 2016.

Andres Felipe Gutierrez, Laurence Dieulle, Nacima Labadie, and Nubia Velasco. Wind farm maintenance scheduling model and solution approach. 2017.

GL Garrad Hassan. A guide to uk offshore wind operations and maintenance. *Scottish Enterprise and the Crown Estate*, 2013.

Chandra Ade Irawan, Djamila Ouelhadj, Dylan Jones, Magnus StÃĕlhane, and Iver Bakken Sperstad. Optimisation of maintenance routing and scheduling for offshore wind farms. *European Journal of Operational Research*, 256(1):76 – 89, 2017. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2016.05.059. URL `http://www.sciencedirect.com/science/article/pii/S0377221716303964`.

Chandra Ade Irawan, Majid Eskandarpour, Djamila Ouelhadj, and Dylan Jones. Simulation-based optimisation for stochastic maintenance routing in an offshore wind farm. *European Journal of Operational Research*, 2019.

Kevin Kennedy, Paul Walsh, Thomas W Mastaglio, and Ted Scully. Genetic optimisation for a stochastic model for opportunistic maintenance planning of offshore wind farms. In *Environment Friendly Energies and Applications (EFEA), 2016 4th International Symposium on*, pages 1–6. IEEE, 2016.

AndrÃąs Kovács, GÃąbor ErdÅŚs, Zsolt JÃąnos Viharos, and LÃąszlÃş Monostori. A system for the detailed scheduling of wind farm maintenance. *CIRP Annals*, 60(1):497 – 501, 2011. ISSN 0007-8506. doi: https://doi.org/10.1016/j.cirp.2011.03.049. URL `http://www.sciencedirect.com/science/article/pii/S0007850611000503`.

William H Kruskal. Historical notes on the wilcoxon unpaired two-sample test. *Journal of the American Statistical Association*, 52(279):356–360, 1957.

William H Kruskal and W Allen Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.

Georgios Leontaris, Oswaldo Morales-Nápoles, and ARM Rogier Wolfert. Probabilistic scheduling of offshore operations using copula based environmental time series–an application for cable installation management for offshore wind farms. *Ocean Engineering*, 125:328–341, 2016.

Yuan Li, Haoxun Chen, and Christian Prins. Adaptive large neighborhood search for the pickup and delivery problem with time windows, profits, and reserved requests. *European Journal of Operational Research*, 252(1):27–38, 2016.

Alexis Mérigaud and John V Ringwood. Condition-based maintenance methods for

marine renewable energy. *Renewable and Sustainable Energy Reviews*, 66:53–78, 2016.

Clair E Miller, Albert W Tucker, and Richard A Zemlin. Integer programming formulation of traveling salesman problems. *Journal of the ACM (JACM)*, 7(4):326–329, 1960.

Ellen Karoline Norlund and Irina Gribkovskaia. Environmental performance of speed optimization strategies in offshore supply vessel planning under weather uncertainty. *Transportation Research Part D: Transport and Environment*, 57:10–22, 2017.

Ellen Karoline Norlund, Irina Gribkovskaia, and Gilbert Laporte. Supply vessel planning under cost, environment and robustness considerations. *Omega*, 57:271–281, 2015.

North Sea Wind Power Hub. North sea wind power hub project, 2018. URL `https://northseawindpowerhub.eu/project/#anker2`.

Samuel Nucamendi-Guillén, Iris Martínez-Salazar, Francisco Angel-Bello, and J Marcos Moreno-Vega. A mixed integer formulation and an efficient metaheuristic procedure for the k-travelling repairmen problem. *Journal of the Operational Research Society*, 67(8):1121–1134, 2016.

offshoreWindBiz. Breaking: Uk offshore wind strike prices slide down to gbp 39.65/mwh. `https://www.offshorewind.biz/2019/09/20/uk-offshore-wind-strike-prices-slide-down-to-gbp-39-65-mwh/`, 2019. Accessed: 20-12-2019.

Orsted. Hornsea one - powering one million homes with green electricity. https://hornseaprojectone.co.uk/, 2018. Accessed: 15-3-2020.

Hans-Peter Piepho. An algorithm for a letter-based representation of all-pairwise comparisons. *Journal of Computational and Graphical Statistics*, 13(2):456–466, 2004.

Nora Tangen Raknes, Katrine Ødeskaug, Magnus Stålhane, and Lars Magnus Hvattum. Scheduling of maintenance tasks and routing of a joint vessel fleet for multiple offshore wind farms. *Journal of Marine Science and Engineering*, 5(1):11, 2017.

RenewableUK. Offshore wind prices tumble in record-breaking auction, 2017. URL https://www.renewableuk.com/news/362971/Offshore-wind-prices-tumble-in-record-breaking-auction-results--cheaper-than-nu htm.

RenewableUK. Recard breaking amount of new uk offshore wind capacity, 2018. URL https://www.renewableuk.com/news/430793/Record-breaking-amount-of-new-UK-offshore-wind-capacity-installed-in-2018-. htm.

Glaydston Mattos Ribeiro and Gilbert Laporte. An adaptive large neighborhood search heuristic for the cumulative capacitated vehicle routing problem. *Computers & operations research*, 39(3):728–735, 2012.

Juan Carlos Rivera, H Murat Afsar, and Christian Prins. A multistart iterated local

search for the multitrip cumulative capacitated vehicle routing problem. *Computational Optimization and Applications*, 61(1):159–187, 2015.

Stefan Ropke and David Pisinger. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation science*, 40 (4):455–472, 2006a.

Stefan Ropke and David Pisinger. A unified heuristic for a large class of vehicle routing problems with backhauls. *European Journal of Operational Research*, 171 (3):750–775, 2006b.

Albert H. Schrotenboer, Michiel A.J. uit het Broek, Bolor Jargalsaikhan, and Kees Jan Roodbergen. Coordinating technician allocation and maintenance routing for offshore wind farms. *Computers & Operations Research*, 98:185 – 197, 2018. ISSN 0305-0548. doi: https://doi.org/10.1016/j.cor.2018.05.019. URL `http://www.sciencedirect.com/science/article/pii/S0305054818301424`.

Mahmood Shafiee. Maintenance logistics organization for offshore wind energy: Current progress and future perspectives. *Renewable Energy*, 77:182–193, 2015.

Paul Shaw. A new local search algorithm providing high quality solutions to vehicle routing problems. *APES Group, Dept of Computer Science, University of Strathclyde, Glasgow, Scotland, UK*, 1997.

Bernard W Silverman. *Density estimation for statistics and data analysis*. Routledge, 2018.

Iver Bakken Sperstad, Elin E Halvorsen-Weare, Matthias Hofmann, Lars Magne Nonås, Magnus Stålhane, and MingKang Wu. A comparison of single-and multi-parameter wave criteria for accessing wind turbines in strategic maintenance and logistics models for offshore wind farms. *Energy Procedia*, 53:221–230, 2014.

Iver Bakken Sperstad, Fiona Devoy McAuliffe, Magne Kolstad, and Severin Sjømark. Investigating key decision problems to optimize the operation and maintenance strategy of offshore wind farms. 2016.

Magnus Stålhane, Lars Magnus Hvattum, and Vidar Skaar. Optimization of routing and scheduling of vessels to perform maintenance at offshore wind farms. *Energy Procedia*, 80:92 – 99, 2015. ISSN 1876-6102. doi: https://doi.org/10.1016/j.egypro.2015.11.411. URL `http://www.sciencedirect.com/science/article/pii/S1876610215021438`. 12th Deep Sea Offshore Wind R&D Conference, EERA DeepWind'2015.

Clym Stock-Williams and Siddharth Krishna Swamy. Automated daily maintenance planning for offshore wind farms. *Renewable Energy*, 2018. ISSN 0960-1481. doi: https://doi.org/10.1016/j.renene.2018.08.112. URL `http://www.sciencedirect.com/science/article/pii/S0960148118310620`.

UK Energy Statistics. Statistical press release: 2014 uk energy statistics 2015, & q4 2015. `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/513244/Press_Notice_March_2016.pdf`, 2016. Accessed: 20-08-2016.

UK Government. 2014 uk greenhouse emissions. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/496946/2014_Final_Emissions_Statistical_Summary_Infographic.pdf, 2015. Accessed: 01-09-2016.

USA Department of Energy. United states - land based and offshore annual average wind speed at 100 meters. https://www.energy.gov/eere/wind/downloads/united-states-land-based-and-offshore-annual-average-wind-speed-100-meters, 2017. Accessed: 27-10-2019.

GJW Van Bussel, AR Henderson, CA Morgan, B Smith, R Barthelmie, K Argyriadis, A Arena, G Niklasson, and E Peltola. State of the art and technology trends for offshore wind energy: operation and maintenance issues. In *Offshore Wind Energy EWEA special topic conference*, 2001.

Joe Whittaker. Graphical models in applied multivariate analysis. *Chichester New York et al: John Wiley & Sons*, 1990.

Ian R Young and Agustinus Ribal. Multiplatform evaluation of global trends in wind speed and wave height. *Science*, 364(6440):548–552, 2019.