U.PORTO

FEUP FACULDADE DE ENGENHARIA UNIVERSIDADE DO PORTO

OPTIMISATION OF LARGE SCALE CO-LOCATED WAVE-WIND ENERGY ARRAYS USING GENETIC ALGORITHMS AND CONSIDERING MULTIPLE SPATIAL CONSTRAINTS

ANDRÉ RODRIGUES MARQUES

Dissertation submitted in partial fulfilment of the requirements for the degree of

MASTER IN CIVIL ENGINEERING — SPECIALISATION IN HYDRAULICS, WATER RESOURCES AND ENVIRONMENT

Supervisor: Professor Doctor Paulo Jorge Rosa Santos

Co-Supervisor: Doctor Felipe Teixeira Duarte

Co-Supervisor: Doctor Ajab Gul Majidi

MARCH 2026

MESTRADO EM ENGENHARIA CIVIL 2025/2026

DEPARTAMENTO DE ENGENHARIA CIVIL E GEORRECURSOS

Tel. +351-22-508 1901

✉ m.ec@fe.up.pt

Edited by

FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Rua Dr. Roberto Frias

4200-465 PORTO

Portugal

Tel. +351-22-508 1400

✉ feup@fe.up.pt

http://www.fe.up.pt

Partial reproductions of this document are authorized provided that the Author is acknowledge and reference is made to the *Master's Degree in Civil Engineering - 2025/2026 - Department of Civil Engineering and Georesources, Faculty of Engineering of the University of Porto, Porto, Portugal, 2026*.

The opinions and information contained in this document are solely those of the Author. The Publisher assumes no legal or other responsibility for any error or omissions that may exist.

This document was produced from the electronic version provided by the Author.

AGRADECIMENTOS

Em primeiro lugar, gostaria de agradecer ao meu orientador, Professor Doutor Paulo Rosa Santos, pela orientação, exigência científica, disponibilidade e acompanhamento ao longo de todo o desenvolvimento deste trabalho. A sua experiência, o seu sentido crítico e as suas sugestões foram fundamentais para a definição e consolidação desta dissertação. Agradeço também aos meus coorientadores, Doutor Felipe Teixeira Duarte e Doutor Ajab Gul Majidi, por toda a ajuda, pela disponibilidade, pela partilha de conhecimento e pelo contributo prestado em diferentes fases deste percurso.

Gostaria igualmente de agradecer aos professores e docentes do DEC, que ao longo do curso contribuíram, de diferentes formas, para a minha formação académica e pessoal. Em particular, ao Professor Tiago Ferradosa, ao Professor Francisco Piqueiro e à Esmeralda, pela disponibilidade constante, pelo apoio e pela forma como sempre estiveram presentes quando foi necessário. Ao Professor Francisco Piqueiro deixo um agradecimento pela atenção que sempre demonstrou ao longo deste percurso.

Quero também agradecer à a3S, e em particular ao Engenheiro João Burmester, ao Engenheiro Álvaro Pereira e ao Miguel Adão, por toda a ajuda, disponibilidade e, sobretudo, por todos os ensinamentos ao longo destes últimos anos. O apoio que me deram, tanto a nível profissional como pessoal, teve um significado muito importante para mim e merece, sem dúvida, ficar aqui reconhecido.

Naturalmente, esta caminhada também não teria sido a mesma sem os amigos que me acompanharam ao longo dos anos. Ao Manel, Mike e Maty pela amizade de sempre e por estar presente ao longo de todo este percurso, dentro e fora da faculdade. Muito deste percurso deve-se a vocês. Ao Jacinto pela disponibilidade e pela ajuda sempre que foi necessário. Deixo também um agradecimento ao grupo "churrascada" e a todos os amigos e colegas que fizeram parte desta fase da minha vida.

Um agradecimento muito especial à Margarida, pela presença constante, pela paciência e pelo apoio ao longo de todo este percurso. Obrigado por estares sempre ao meu lado e por tornares tudo mais leve mesmo nos momentos de maior pressão. Obrigado também por ouvires falar, provavelmente mais vezes do que qualquer pessoa deveria ser obrigada, sobre ondas, algoritmos, otimizações e tudo o que esta tese acabou por trazer pelo meio. No meio de prazos, dúvidas e dias mais exigentes, foste sempre uma presença segura, tranquila e muito importante para mim. Esse apoio, e tudo o que fazes por mim, em todas as suas formas, significa mais do que consigo pôr em palavras.

Aos meus avós, Avó Preciosa, Avô Leonel, Papi e Mamie, deixo um agradecimento muito especial por tudo aquilo que sempre fizeram por mim. Pelo carinho, pelo apoio, pela dedicação e por tudo o que me proporcionaram ao longo da vida. Muito do que sou hoje, e do caminho que consegui fazer até aqui, também vos devo a vocês.

Ao meu irmão, deixo um agradecimento muito especial, simplesmente por ser quem é. Pela presença constante, pela forma genuína como sempre esteve ao meu lado e por nunca ter faltado quando realmente era importante. Há ligações que não precisam de grandes explicações, esta é uma delas.

Por fim, e acima de tudo, aos meus pais. A vocês devo muito mais do que aquilo que alguma página conseguirá dizer. Obrigado por todo o esforço, por todas as oportunidades e por nunca deixarem faltar o que realmente importava. Obrigado por acreditarem em mim e por serem os principais responsáveis por eu ter chegado até aqui. Tudo aquilo que sou, a forma como penso, a maneira como enfrento as dificuldades e o valor que dou ao trabalho, ao compromisso e à persistência devo-o, em grande parte, a vocês. Não há forma verdadeiramente justa de agradecer tudo como merece, mas há pelo menos a certeza de que nada disto teria sido possível sem vocês.

A todos, o meu sincero obrigado.

RESUMO

A crescente necessidade de descarbonização do setor energético tem reforçado o interesse na exploração sustentável de recursos renováveis em ambiente *offshore*. Neste contexto, os sistemas combinados de energia eólica *offshore* e energia das ondas surgem como uma solução promissora para melhorar a utilização do espaço marítimo, aumentar a eficiência das infraestruturas e contribuir para uma maior estabilidade da produção energética. A presente dissertação investiga a otimização da disposição espacial de conversores de energia das ondas em co-colocação com parques eólicos *offshore* flutuantes, considerando condições realistas de implantação ao largo da costa portuguesa.

Para esse efeito, foi utilizada uma metodologia que combina modelação hidrodinâmica, representação agregada do desempenho de conjuntos de conversores de energia das ondas e otimização evolutiva baseada em algoritmos genéticos. A propagação da agitação marítima e a interação entre ondas e dispositivos foram simuladas utilizando o modelo espectral SWAN, na sua extensão SNL-SWAN, permitindo representar os conversores de energia das ondas como obstáculos extratores de energia no campo de ondas. Devido à escala do problema e às limitações computacionais associadas, os conjuntos CorPack foram representados através de uma formulação agregada, recorrendo a duas representações hidrodinâmicas alternativas: uma matriz de potência e uma formulação baseada na largura de captura relativa.

A metodologia foi aplicada a um caso de estudo localizado numa das áreas da costa atlântica portuguesa a concessionar, para a qual foi definida uma configuração fixa de parque eólico *offshore* flutuante e um domínio admissível para a implantação dos conversores de energia das ondas. A otimização da disposição espacial destes dispositivos foi realizada através de um algoritmo genético, considerando explicitamente restrições espaciais relevantes, nomeadamente o polígono de implantação e as zonas de exclusão associadas às turbinas eólicas. Foram analisados diferentes cenários de otimização, permitindo avaliar a influência da representação hidrodinâmica adotada, do número de dispositivos e dos principais parâmetros de controlo do algoritmo genético.

Os resultados demonstraram que a metodologia desenvolvida é capaz de identificar disposições viáveis e hidrodinamicamente eficientes para os conversores de energia das ondas em parques *offshore* co-colocados. Verificou-se que as duas representações hidrodinâmicas conduzem a níveis de desempenho global semelhantes, embora originem diferenças na organização espacial das soluções ótimas. Observou-se igualmente que a densidade de dispositivos influencia o desempenho hidrodinâmico do sistema, podendo configurações com menor número de dispositivos conduzir a soluções mais eficientes. A análise de sensibilidade evidenciou ainda que os parâmetros do algoritmo genético influenciam a qualidade das soluções e o comportamento de convergência, embora a metodologia tenha demonstrado uma robustez global consistente.

Conclui-se, assim, que a integração de modelação hidrodinâmica, restrições espaciais e otimização evolutiva constitui uma abordagem eficaz para apoiar o planeamento e o dimensionamento preliminar de sistemas *offshore* co-colocados de energia eólica *offshore* e energia das ondas.

PALAVRAS-CHAVE: energia das ondas, energia eólica *offshore* flutuante, sistemas co-colocados, algoritmos genéticos, restrições espaciais.

ABSTRACT

The growing need to decarbonise the energy sector has intensified interest in the sustainable exploitation of renewable resources in offshore environments. In this context, co-located offshore wind and wave energy systems have emerged as a promising solution to improve the use of the marine space, enhance infrastructure efficiency, and contribute to greater stability in offshore energy production. This dissertation investigates the optimisation of the spatial arrangement of wave energy converters co-located with floating offshore wind farms, considering realistic deployment conditions off the Portuguese coast.

To address this objective, an integrated methodology was used combining hydrodynamic modelling, aggregated performance representation of wave energy converter clusters, and evolutionary optimisation based on genetic algorithms. Wave propagation and wave-device interaction were simulated using the SWAN spectral wave model, through its SNL-SWAN extension, which enables wave energy converters to be represented as energy-extracting obstacles within the wave field. Due to the scale of the problem and associated computational constraints, CorPack clusters were represented through an aggregated formulation using two alternative hydrodynamic representations: a power matrix and a formulation based on Relative Capture Width.

The methodology was applied to a case study located within a future concession area along the Portuguese Atlantic coast, where a fixed floating offshore wind farm configuration and an admissible deployment domain for the wave energy converters were defined. The optimisation of the spatial arrangement of these devices was carried out using a genetic algorithm, explicitly accounting for relevant spatial constraints, namely the deployment polygon and the exclusion zones associated with the wind turbines. Several optimisation scenarios were analysed to evaluate the influence of the hydrodynamic representation, the number of devices, and the main control parameters of the genetic algorithm.

The results demonstrated that the proposed methodology could identify feasible and hydrodynamically efficient layouts for wave energy converters within co-located offshore parks. The two hydrodynamic representations were found to produce similar levels of overall performance, although noticeable differences were observed in the spatial organisation of the optimal solutions. It was also observed that device density influences the hydrodynamic performance of the system, with configurations using fewer devices sometimes leading to more efficient layouts. The sensitivity analysis further showed that genetic algorithm parameters influence both the quality of the solutions and the convergence behaviour, although the optimisation framework demonstrated consistent overall robustness.

It is therefore concluded that the integration of hydrodynamic modelling, spatial constraints, and evolutionary optimisation provides an effective approach to support the planning and preliminary design of co-located offshore wind and wave energy systems.

KEYWORDS: wave energy, floating offshore wind energy, co-located systems, genetic algorithms, spatial constraints.

TABLE OF CONTENTS

AGRADECIMENTOS I

RESUMO III

ABSTRACT V

1. INTRODUCTION 1

1.1. CONTEXT OF OFFSHORE ENERGY 1

1.2. WIND-WAVE COMPLEMENTARITY 2

1.3. MOTIVATION AND OBJECTIVES 3

1.4. STRUCTURE 4

2. STATE OF THE ART 5

2.1. OFFSHORE WIND ENERGY: FROM FIXED TO FLOATING SYSTEMS 5

2.1.1. FIXED-BOTTOM FOUNDATIONS 5

2.1.2. FLOATING OFFSHORE WIND SYSTEMS 7

2.1.2.1. Floating Platform Concepts and Technological Characteristics 7

2.1.2.2. Turbine Scaling and Emerging High-Capacity Offshore Classes 9

2.1.3. WAKE AND AERODYNAMIC CONSTRAINTS 11

2.1.4. ECONOMIC AND INFRASTRUCTURE CONSTRAINTS 14

2.1.5. FLOATING-SPECIFIC SPATIAL CONSTRAINTS 16

2.1.6. REFERENCE OFFSHORE WIND FARMS LAYOUTS 18

2.2. WAVE ENERGY CONVERSION TECHNOLOGY 22

2.2.1. CLASSIFICATION AND OPERATION PRINCIPLES 22

2.2.2. PHYSICS OF POINT ABSORBERS 24

2.2.3. TARGET TECHNOLOGY: CORPOWER C4 AND CORPACK 26

2.3. CO-LOCATED WAVE-WIND ENERGY SYSTEMS 30

2.3.1. CONCEPT DEFINITION: CO-LOCATED VS HYBRID 30

2.3.2. ECONOMIC SYNERGIES AND LEVELIZED COST OF ENERGY 31

2.3.3. HYDRODYNAMIC INTERACTIONS AND OPERATIONAL SYNERGIES IN CO-LOCATED FARMS 32

2.4. OPTIMISATION OF ARRAY LAYOUTS USING GENETIC ALGORITHMS 33

2.4.1. THE ARRAY OPTIMISATION PROBLEM 33

2.4.2. GENETIC ALGORITHMS METHODOLOGY 34

2.4.3. MULTI-OBJECTIVE FITNESS FUNCTION: THE WLA INDEX 35

2.5. SUMMARY AND IDENTIFIED GAPS ... 36

3. METHODOLOGY ... 39

3.1. OVERVIEW OF THE OPTIMISATION FRAMEWORK ... 39

3.2. DEFINITION OF THE STUDY DOMAIN, GRID AND SPATIAL LAYERS ... 41

3.3. HYDRODYNAMIC REPRESENTATION OF THE CORPACK CLUSTER ... 43

3.3.1. PERFORMANCE REPRESENTATION OF THE CORPACK CLUSTER ... 43

3.3.2. DIRECTIONAL DEPENDENCE OF CLUSTER PERFORMANCE ... 44

3.3.3. NUMERICAL MODELLING FRAMEWORK FOR CLUSTER RESPONSE SIMULATIONS ... 45

3.3.4. CORPACK POWER MATRIX REPRESENTATION ... 47

3.3.5. CORPACK RCW REPRESENTATION ... 49

3.3.6. VALIDATION OF THE CORPACK PERFORMANCE ... 51

3.4. REPRESENTATIVE SEA-STATES GENERATION FOR OPTIMISATION ... 54

3.5. HYDRODYNAMIC MODELLING SETUP ... 60

3.6. OBJECTIVE METRIC DEFINITION ... 61

3.7. GENETIC ALGORITHM STRUCTURE AND OPTIMISATION PROCEDURE ... 62

3.8. CONSTRAINTS HANDLING STRATEGY ... 65

3.9. OPTIMISATION CONVERGENCE AND EVALUATION PARAMETERS ... 66

3.10. METHODOLOGICAL WORKFLOW SUMMARY ... 67

4. CASE STUDY ... 69

4.1. THE PORTUGUESE CONTEXT: SPATIAL CONSTRAINTS AND RESOURCE ASSESSMENT ... 69

4.1.1. PORTUGUESE CONTINENTAL SHELF AND STRATEGIC CONTEXT ... 69

4.1.2. REGULATORY FRAMEWORK AND SPATIAL CONSTRAINTS FOR OPTIMISATION ... 70

4.1.3. OFFSHORE WIND RESOURCE ASSESSMENT ... 71

4.1.4. OFFSHORE WAVE RESOURCE ASSESSMENT ... 74

4.1.5. WIND-WAVE COMPLEMENTARITY ... 77

4.1.6. OPERATIONAL VIABILITY AND MAINTENANCE WEATHER WINDOWS ... 82

4.2. ZONE P3 SPECIFIC CHARACTERIZATION ... 83

4.2.1. GEOGRAPHIC DEFINITION AND PHYSICAL CONSTRAINTS ... 83

4.2.2. OFFSHORE WIND CHARACTERIZATION ... 88

4.2.3. OFFSHORE WAVE CHARACTERIZATION ... 90

4.2.4. WIND-WAVE COMPLEMENTARITY ... 91

4.2.5. LOGISTICAL ACCESSIBILITY AND OPERATIONAL FEASIBILITY ... 92

4.3. FROM CONCESSION-SCALE P3 TO PROJECT-SCALE STUDY AREA 95
4.3.1. RATIONALE FOR PROJECT-SCALE SUBDIVISION 95
4.3.2. DEFINITION OF THE SELECTED STUDY AREA 96
4.4. REFERENCE FLOATING WIND FARM DEFINITION 98
4.4.1. CAPACITY TARGET AND TURBINE CLASS 98
4.4.2. WIND FARM GEOMETRY AND ORIENTATION 99
4.5. TURBINE EXCLUSION ZONES AND NET DEPLOYMENT AREA 101
4.5.1. EXCLUSION-RADIUS DEFINITION 101
4.5.2. AVAILABLE AREA FOR WEC DEPLOYMENT 102
4.6. DEFINITION OF THE WAVE-ENERGY LAYER 103
4.6.1. CORPACK MODULE DEFINITION 103
4.6.2. SELECTION OF THE CORPACK DENSITY 105

5. RESULTS AND DISCUSSION 109
5.1. OVERVIEW OF THE OPTIMISATION RESULTS 109
5.2. HYDRODYNAMIC Q-FACTOR EVALUATION 111
5.3. INFLUENCE OF THE HYDRODYNAMIC REPRESENTATION: CASE 1 VS. CASE 2 113
5.4. INFLUENCE OF THE NUMBER OF CORPACKS: CASE 1 VS. CASE 3 117
5.5. INFLUENCE OF THE POPULATION SIZE: CASE 3 VS. CASE 4 121
5.6. INFLUENCE OF THE CROSSOVER RATE: CASE 3 VS. CASE 5 125
5.7. INFLUENCE OF THE MUTATION RATE: CASE 3 VS. CASE 6 129
5.8. WAVE-HEIGHT REDUCTION DUE TO THE PRESENCE OF CORPACKS 133

6. CONCLUSIONS 135
6.1. MAIN FINDINGS 135
6.2. SCIENTIFIC AND METHODOLOGICAL CONTRIBUTIONS 136
6.3. ENGINEERING AND PLANNING IMPLICATIONS 137
6.4. CONTRIBUTION TO THE UNITED NATIONS SUSTAINABLE DEVELOPMENT GOALS 139
6.5. LIMITATIONS AND RECOMMENDATIONS FOR FUTURE WORKS 140

APPENDIX A - CORPOWER POWER MATRIX 150
APPENDIX B - FINAL LAYOUT FOR EACH OPTIMISATION CASE 151
APPENDIX C - H_s FIELDS FOR EACH OPTIMISATION CASE 152

LIST OF FIGURES

Figure 1 – Types of fixed-bottom foundations. 5

Figure 2 - Platform degrees of freedom. 7

Figure 3 - Type of FOWT foundations. 8

Figure 4 - Indicative evolution of offshore wind turbine capacity and rotor dimensions. 9

Figure 5 - Installation of the CTG-Goldwind 20MW offshore wind turbine. 10

Figure 6 - Turbine wake effect. 12

Figure 7 - Princess Amalia Wind Park turbine locations with indicative directional spacing expressed in D. 13

Figure 8 - Cable layout for co-located array. 14

Figure 9 - Safety zone around offshore substation. 15

Figure 10 - Different floating offshore mooring systems. 16

Figure 11 - Typical spread mooring design with catenary lines. 17

Figure 12 - Exclusive zone (red rectangle) for the WindFloat Atlantic farm. 18

Figure 13 . Alpha Ventus' offshore wind farm layout. 20

Figure 14 - Development of wind farm patterns over time. 20

Figure 15 – The various wave energy technologies. 23

Figure 16 – Examples of wave energy converters: (a) Oscillating Water Column; (b) Overtopping; (c) Oscillating Bodies. 23

Figure 17 - Example of a point absorber. 25

Figure 18 - CorPower C4 overview. 27

Figure 19 - Power matrix for the CorPower C4 device. 28

Figure 20 - CorPack array overview. 29

Figure 21 - Classification of combined wave-wind technologies. 30

Figure 22 - Cost model comparison of array types. 31

Figure 23 - The main evolutionary cycle. 35

Figure 24 - Optimisation workflow. 40

Figure 25 - Conceptual representation of the spatial-domain construction used in the optimisation framework. 42

Figure 26 - Representation of the CorPack layout implemented in the SWAN computational domain. 46

Figure 27 - Workflow adopted for the construction of the aggregated CorPack power matrix. 47

Figure 28 - Discretised sea-states bins used to define the CorPack cluster simulations. 48

Figure 29 - CorPower CorPack power matrix. 48

Figure 30 - Absorbed power values obtained, organised as a function of T_p and incident wave direction. 50

Figure 31 - CorPower CorPack RCW curve. 50

Figure 32 - Sea states used in the SafeWAVE modelling study for the period 15 Oct 2022 to 30 Oct 2022. 51

Figure 33 - Sea states used in the present SWAN simulations for the period 15 Oct 2022 to 30 Oct 2022. 52

Figure 34 - Simplified representation of the CorPack layout, as a single obstacle, implemented in the SWAN computational domain. 53

Figure 35 - Directional distribution of H_s for the study domain. 55

Figure 36 - Statistical distributions of the main wave parameters (H_s, T_p and θ). 56

Figure 37 - Three-dimensional wave clustered wave dataset. 57

Figure 38 - Elbow method for the optimal number of clusters. 58

Figure 39 - Power representativeness of the reduced wave dataset as function of the number of clusters. 59

Figure 40 - Genetic Algorithm workflow [53]. 62

Figure 41 - Schematic representation of crossover and mutation operators applied to layout chromosomes. 63

Figure 42 - Summary of the methodological workflow adopted. 68

Figure 43 – Portugal's Exclusive Economic Zone and Extended Continental Shelf. 69

Figure 44 - PAER zones P1 to P6. 70

Figure 45 - Spatial distribution of wind speed (left) and wind power density (right) along mainland Portugal. 71

Figure 46 - Spatial distribution of MVI for wind power density along mainland Portugal. 72

Figure 47 - Wind roses for PAER zones. 73

Figure 48 - Spatial distribution of mean wave power density along mainland Portugal. 74

Figure 49 - Spatial distribution of MVI for wave power density along mainland Portugal. 75

Figure 50 - Monthly distribution of wave power density for PAER zones. 76

Figure 51 - Wave roses for PAER zones. 77

Figure 52 - Spatial distribution of CCT(0) along mainland Portugal. 78

Figure 53 - c_T distributions for PAER zones. 79

Figure 54 - CLF distribution for PAER zones. 80

Figure 55 - Boxplot illustrating the monthly variations of wind and wave power densities for the PAER zones. 81

Figure 56 - The spatial distributions of the operational time of OWTs, the operational time of WECs, and the O&M weather window, from left to right, respectively. 83

Figure 57 - PAER Zone P3 polygon. 84

Figure 58 - Site slope map. 85

Figure 59 - Site seabed composition. 86

Figure 60 - Regulatory and environmental constraints. 87

Figure 61 - Spatial distribution of offshore wind speed for zone P3. 88

Figure 62 - Spatial distribution of offshore wave intensity for zone P3. 90

Figure 63 - Spatial distribution of CSI along mainland Portugal. 93

Figure 64 - Conceptual internal subdivision of PAER Zone P3 into five representative project-scale blocks, with the selected study block highlighted. 95

Figure 65 - Map of the selected study polygon. 96

Figure 66 - Final wind farm layout. 99

Figure 67 - Exclusion zones around each turbine (red circles). 101

Figure 68 - Schematic representation of the residual deployable area populated conceptually by CorPack modules (blue rectangles). 105

Figure 69 - Optimal WECS layouts: Case 1 (top-left); Case 2 (top-mid); Case 3 (top-right); Case 4 (bottom-left); Case 5 (bottom-mid); Case 6 (bottom-right). Wind turbines (triangles); Exclusion zones (circles); CorPacks (red lines); PAER limits (black polygon). 110

Figure 70 - q-factor values obtained for the six optimisation cases under the representative sea states. 111

Figure 71 - H_s values for Case 1 (left) and Case 2 (right). 114

Figure 72 - CF evolution for Case 1 (top) and Case 2 (bottom). 115

Figure 73 - Average absorbed power evolution for Case 1 (top) and Case 2 (bottom). 115

Figure 74 - CF range and coefficient of variation evolution for Case 1 (top) and Case 2 (bottom). 116

Figure 75 - P2M evolution for Case 1 (top) and Case 2 (bottom). 116

Figure 76 - H_s values for Case 1 (left) and Case 3 (right). 118

Figure 77 - CF evolution for Case 1 (top) and Case 3 (bottom). 118

Figure 78 - Average absorbed power evolution for Case 1 (top) and Case 3 (bottom). 119

Figure 79 - CF range and coefficient of variation evolution for Case 1 (top) and Case 3 (bottom). 120

Figure 80 - P2M evolution for Case 1 (top) and Case 3 (bottom). 120

Figure 81 - H_s values for Case 3 (left) and Case 4 (right). 122

Figure 82 - CF evolution for Case 3 (top) and Case 4 (bottom). 122

Figure 83 - Average absorbed power evolution for Case 3 (top) and Case 4 (bottom). 123

Figure 84 - CF range and coefficient of variation evolution for Case 3 (top) and Case 4 (bottom). 124

Figure 85 - P2M evolution for Case 3 (top) and Case 4 (bottom). 124

Figure 86 - H_s values for Case 3 (left) and Case 5 (right). 126

Figure 87 - CF evolution for Case 3 (top) and Case 5 (bottom). 126

Figure 88 - Average absorbed power evolution for Case 3 (top) and Case 5 (bottom). 127

Figure 89 - CF range and coefficient of variation evolution for Case 3 (top) and Case 5 (bottom). 127

Figure 90 - P2M evolution for Case 3 (top) and Case 5 (bottom). 128

Figure 91 - H_s values for Case 3 (left) and Case 6 (right). 129

Figure 92 - CF evolution for Case 3 (top) and Case 6 (bottom). 130

Figure 93 - Average absorbed power evolution for Case 3 (top) and Case 6 (bottom). 131

Figure 94 - CF range and coefficient of variation evolution for Case 3 (top) and Case 6 (bottom). 131

Figure 95 - P2M evolution for Case 3 (top) and Case 6 (bottom). 132

Figure 96 - H_s field in the study domain without WECs. 133

LIST OF TABLES

Table 1 - Summary of spatial constraints governing offshore wind farm layout design. 19

Table 2 - Layout characteristics of representative large-scale offshore wind farms. 21

Table 3 - Validation of the aggregated CorPack representations against SafeWAVE study. 53

Table 4 - Representative sea states for the optimisation study. 59

Table 5 - Parameters of the six cases presented in this study. 64

Table 6 - Vertices coordinates. 84

Table 7 - Key descriptors of zone P3. 94

Table 8 - Coordinates of the selected study area. 97

Table 9 - Design basis of the reference floating wind farm. 98

Table 10 - First-order area balance after turbine exclusion. 103

Table 11 - CorPack representation adopted for the wave energy layer. 104

Table 12 - Qualitative interpretation of CorPack density in the residual deployable domain. 106

Table 13 - Summary of optimisation results. 110

Table 14 - Minimum, maximum and mean q-factor obtained for the final layout of each optimisation case. 113

Table 15 - Quantitative comparison between Case 1 and Case 2. 117

Table 16 - Quantitative comparison between Case 1 and Case 3. 121

Table 17 - Quantitative comparison between Case 3 and Case 4. 125

Table 18 - Quantitative comparison between Case 3 and Case 5. 128

Table 19 - Quantitative comparison between Case 3 and Case 6. 132

Table 20 - Contribution of the dissertation to the United Nations Sustainable Development Goals, including associated targets and performance indicators. 139

UNITS:

GW - Gigawatts
h - Hours
km - Kilometres
km^2 - Square Kilometres
kW - Kilowatts
m - Metres
m^2 - Square Metres
MW - Megawatts
nm - Nautical Miles
s - Seconds
W - Watts

SYMBOLS:

b	[m]	- CorPack Width
C_k	[-]	- Set Points in Cluster k
CCT(0)	[-]	- Cross-Correlation at Zero Time-Lag
c_τ	[-]	- Cross-Correlation at a Specific Time
D	[m]	- Rotor Diameter
g	[m/s^2]	- Gravitational Acceleration
H_s	[m]	- Significant Wave Height
$H_{s,i}$	[m]	- Significant Wave Height at node i (without WEC)
$H_{s,\,WEC,\,i}$	[m]	- Significant Wave Height at node i (with WEC)
J	[kW/m]	- Wave Power per Unit Crest Lenght
K	[-]	- Number of Clusters
m_0	[m^2]	- Spectral moment of order 0
m_{-1}	[m^2.s]	- Spectral moment of order -1
N(HRA)	[-]	- Normalized Value of HRA
N(q)	[-]	- Normalized Value of q
p	[-]	- Weighting Factor Used to Define the Importance Assigned to Power Absorption
P_{abs}	[kW]	- Absorbed Power
P_{ARRAY}	[kW]	- Total Power Absorbed by the Array
P_{RATED}	[kW]	- Rated Power
P_{WEC}	[kW]	- Power Absorbed by a Single WEC
$P(H_s, T_p)$	[kW]	- Power Absorbed at a Specific H_s-T_p Combination
P_w	[kW/m]	- Wave Power Density
q	[-]	- q-Factor (Interaction Factor)
s	[s]	- Weighting Factor Used to Define the Importance Assigned to Shielding
S	[-]	- Within Cluster Variance
T_e	[s]	- Wave Energy Period
T_p	[s]	- Wave Peak Period
W_p	[W/m^2]	- Wind Power Density
x_i	[-]	- Vector Parameter Associated with Observation i
θ	[º]	- Wave Direction
ρ	[kg/m^3]	- Water Density
μ_k	[-]	- Centroid of Cluster k
σ	[-]	- Standard Deviation

ACRONYMS:

CAPEX	- Capital Expenditure
CF	- Capacity Factor
CLF	- Co-Location Feasibility Level
CoV	- Coefficient of Variation
CSI	- Composite Suitable Index
EEZ	- Exclusive Economic Zone
EU	- European Union
FOW	- Floating Offshore Wind
GA	- Genetic Algorithm
GWEC	- Global Wind Energy Council
HRA	- Height Reduction Area
IEA	- International Energy Agency
LCoE	- Levelized Cost of Energy
MVI	- Monthly Variability Index
NEPS	- Number of Operating Hours Equivalent to the Nominal Power
OWC	- Oscillating Water Column
O&M	- Operations and Maintenance
PAER	- Allocation Plan for Offshore Renewable Energies
PTO	- Power Take-Off
P2M	- Peak-to-Mean Power Ratio
RCW	- Relative Capture Width
SNL-SWAN	- Sandia National Laboratories - Simulating Waves Nearshore
SWAN	- Simulating Waves Nearshore
TLP	- Tension Leg Platform
WEC	- Wave Energy Converter
WGS84	- World Geodetic System 1984
WLA	- Wave Energy Park Layout Assessment

1. INTRODUCTION

1.1. CONTEXT OF OFFSHORE ENERGY

The decarbonisation of the global energy system has become an urgent and irreversible priority in the face of escalating climate risks and growing energy insecurity [1]. The transition away from fossil fuels is no longer viewed solely as an environmental necessity but also as a foundation for economic stability, technological innovation, and long-term resilience [2]. Within this transformation, offshore renewable energy, particularly wind and wave power, has emerged as a central pillar due to the abundance, persistence, and predictability of oceanic resources [3], [4], [5].

Offshore technologies offer distinct advantages compared to their land-based counterparts. They can access stronger and more consistent resources, occupy vast marine areas with minimal spatial conflict, and avoid the land-use and visual constraints associated with onshore developments [6], [7]. This capacity to harness the untapped potential of ocean environments makes offshore renewables a cornerstone of global and regional decarbonisation strategies.

According to the International Energy Agency (IEA) [1], global installed wind capacity surpassed 1,170 gigawatts (GW) in 2023, of which less than 10% corresponded to offshore wind. Although this represents a small percentage, offshore wind is expanding more rapidly than onshore due to falling costs and ambitious policy frameworks.

Over the past decade, the rapid expansion of offshore wind has profoundly reshaped the global renewable energy landscape, with Europe assuming a leading role in this transformation. According to the Global Wind Energy Council (GWEC) [4], more than 12GW of new offshore wind capacity were installed worldwide in 2024, bringing total global capacity to approximately 83GW, of which over 21 GW are located within the European Union.

Nowhere has this progress been more evident than in Europe, which has firmly established itself as the global leader in offshore renewable energy. The European Union (EU) has maintained its global leadership through a combination of ambitious policy frameworks, industrial innovation, and sustained investment. Core legislative initiatives, such as the European Green Deal, the Offshore Renewable Energy Strategy, and the Renewable Energy Directive (RED III), set binding targets of at least 60 GW of offshore wind and 1 GW of ocean energy by 2030, rising to 300 GW and 40 GW, respectively, by 2050 [3], [8]. These goals align with the European Green Deal, which aims to achieve climate neutrality by mid-century while ensuring energy security, economic growth, and industrial competitiveness.

These objectives illustrate not only a heightened level of climate ambition but also a growing recognition of the broader strategic, economic, and social importance of offshore renewables. The expanding blue economy, which encompasses marine energy, maritime transport, coastal infrastructure, and emerging

ocean technologies, has become a key driver of sustainable growth. In 2023, this sector accounted for nearly five million jobs and generated over €260 billion in gross value added across the EU [7].

While Europe's initial offshore wind developments were concentrated in the North Sea and Baltic Sea, and member states such as the United Kingdom, Germany, the Netherlands and Denmark already host large-scale offshore wind farms, a new phase of expansion is now emerging, with floating concepts in deep-water, particularly in France, Spain, and Portugal [9].

The Atlantic Ocean is increasingly regarded as both a promising frontier and a natural laboratory for offshore renewable energy innovation. Portugal, in particular, provides a compelling example. With its extensive and highly exposed western coastline, the country benefits from some of the most energetic and persistent wind and wave conditions in Europe [9].

Over the past two decades, the Portuguese coast has hosted several landmark offshore renewable energy demonstration projects, particularly in the Aguçadoura test site, which has served as a real-scale experimental platform for multiple pioneering technologies. These include early wave energy converters such as Pelamis [10], Archimedes Wave Swing [10], and more recently CorPower Ocean [11], as well as floating offshore wind demonstrators, most notably the WindFloat concept [9].

Collectively, these projects have played a decisive role in validating innovative offshore energy technologies under harsh Atlantic conditions, while simultaneously highlighting the technical, operational, and economic challenges associated with large-scale marine energy deployment [10].

While offshore wind has reached a commercial stage, wave energy remains in the early stages of technological development. Reports from the International Renewable Energy Agency [12] identify wave energy as highly promising but still constrained by low Technology Readiness Levels (TRLs). Pilot projects are underway in several European countries, as well as in the United States and the Asia-Pacific region, but large-scale deployment has not yet been achieved. Nonetheless, wave energy is increasingly valued for its complementarity with offshore wind, as it can contribute to a more stable and predictable aggregated generation profile [13].

Collectively, these developments underscore the pivotal role of offshore renewables within Europe's long-term vision for climate neutrality, energy security, and industrial competitiveness. As policy frameworks evolve and practical expertise grows, the successful and sustainable integration of offshore wind and wave energy will become a cornerstone in Europe's transition towards a resilient, low-carbon energy future.

1.2. WIND-WAVE COMPLEMENTARITY

Within the broader offshore context, the idea of resource complementarity, the notion that wind and wave energies are not only abundant but may reinforce each other, has emerged as a central theme in both science and strategy [6]. While both resources originate from the same atmospheric processes, their temporal and spatial behaviours differ substantially. Wind events typically precede high waves, and wave fields, by inheriting energy from distant meteorological systems, often persist when wind events have diminished [14], [15]. This physical distinction creates periods of asynchronous availability, during which one resource compensates for the other's variability.

This temporal complementarity in the development of integrated marine energy systems yields more regular and reliable power than single-source technologies. The potential synergy between these resources suggests that combined wind-wave exploitation can significantly smooth overall power generation and reduce variability in offshore energy output [16], [17], [18]. Beyond temporal balancing,

spatial complementarity also plays an important role. Optimal zones for wind and wave energy do not always coincide, offering opportunities for coordinated spatial planning and more efficient use of maritime areas [19]. Such complementarity supports broader objectives of system resilience, grid integration, and infrastructure optimisation, which are essential for the sustainable expansion of offshore renewables in the coming decades.

Although further technical and economic exploration is essential to fully quantify and operationalise these synergies, the existing body of research consistently identifies hybrid wind-wave complementarity as a promising pathway towards a more stable, predictable, and resilient offshore energy mix [20]. This principle provides the scientific and strategic rationale for exploring co-located wind and wave energy systems, which aim to harness the advantages of both resources within a shared marine environment, a topic examined in greater depth in Chapter 2.

1.3. MOTIVATION AND OBJECTIVES

The global momentum towards decarbonisation has placed unprecedented emphasis on the diversification and integration of renewable energy sources. At this stage, offshore environments stand out as a frontier for achieving large-scale, low-carbon energy generation. Yet, despite the remarkable growth of offshore wind in recent years, the potential for joint exploitation of marine resources remains only partially realised. Most current developments are technologically and spatially segregated, overlooking the synergies that could arise from coordinated use of wind and wave energy within the same maritime space. The motivation for this research stems from the need to explore these synergies more comprehensively.

Floating offshore wind turbines have matured into commercially viable technologies capable of harnessing deep-water wind resources, whereas wave energy converters remain at earlier stages of development but possess considerable potential to complement wind generation. Their combination, through co-located deployment, offers a promising pathway to enhance energy yield stability, improve the efficiency of infrastructure utilisation, and reduce variability in power output. Furthermore, such systems align with the wider European agenda for marine spatial planning and the sustainable expansion of the blue economy, in which optimising the use of ocean space is increasingly essential.

Against this background, there is a clear need for integrated, site-specific assessments that go beyond generic resource evaluations and explicitly address the combined behaviour of wind and wave systems.

Despite these promising prospects, significant uncertainties remain regarding the technical, economic, and environmental feasibility of co-located offshore wind and wave energy systems. In particular, a robust understanding of how wind and wave resources can be combined under realistic offshore conditions, including their temporal variability, complementarity, and mutual influence, is essential for informing effective deployment strategies. This need is especially pronounced in highly energetic and dynamic environments such as the Portuguese Atlantic coast.

In this context, this thesis aims to investigate the optimisation of large-scale co-located offshore wind and wave energy arrays under realistic Portuguese offshore conditions. Rather than performing a broad feasibility assessment of co-located systems, the work focuses on the spatial design of the wave energy layer within a predefined floating offshore wind farm configuration, using a quantitative optimisation framework that explicitly accounts for key deployment constraints.

More specifically, the thesis refines and applies an existing methodology to a new Portuguese offshore case study, that combines spectral wave modelling, aggregated WEC performance representation, and genetic algorithm optimisation to identify feasible and hydrodynamically efficient WEC layouts within

a selected offshore case-study area. The work evaluates how the optimisation outcomes are influenced by the adopted hydrodynamic representation of the WEC cluster, the number of WEC units deployed, and the main control parameters of the genetic algorithm, while respecting the spatial restrictions imposed by the host wind farm and the admissible deployment domain.

By integrating hydrodynamic modelling, spatial constraint handling, and evolutionary optimisation, this thesis seeks to contribute to the methodological validation of optimisation tools for co-located marine renewable energy systems and to support more informed layout design strategies for future offshore wave-wind projects.

1.4. STRUCTURE

This thesis is organised to guide the reader from the broad context of offshore renewable energy through the specific challenges and solutions associated with co-located wind and wave energy systems.

Following this introduction, which introduces the global and European context of offshore renewable energy, emphasising the growing relevance of wind-wave complementarity and outlining the motivation and objectives of this study, Chapter 2 presents a comprehensive review of the state of the art, detailing the technological, economic, and environmental aspects of floating offshore wind turbines and wave energy converters, as well as the current approaches to co-located system design and assessment.

Chapter 3 describes the methodological framework adopted for this research, including resource characterisation, modelling procedures, and performance evaluation criteria. Chapter 4 applies this methodology to a case study along the Portuguese Atlantic coast, analysing the energy potential and interaction between offshore wind and wave resources. Chapter 5 presents the main findings and assesses their implications for system performance, spatial planning, and policy development.

Finally, Chapter 6 summarises the study's key conclusions and proposes recommendations for future research and implementation pathways.

.

2. STATE OF THE ART

2.1. OFFSHORE WIND ENERGY: FROM FIXED TO FLOATING SYSTEMS

2.1.1. FIXED-BOTTOM FOUNDATIONS

The genesis of the offshore wind industry is rooted in the adaptation of marine engineering principles originally developed for the offshore oil and gas sectors, specifically regarding the design and installation of bottom-fixed structures [21]. For most of the deployed offshore wind capacity to date, the industry has relied on foundations that are rigidly connected to the seabed, providing a stable platform that mimics the dynamic behaviour of onshore turbines [22]. These fixed-bottom systems, presented in Figure 1, have underpinned the explosive growth of the sector in the North Sea and the Baltic Sea, where large areas of shallow water, typically less than 40 m, with favourable soil conditions are readily available [23]. Understanding the mechanics and limitations of these structures is essential to appreciating the technological leap required for the transition to the deep-water environments characteristic of the Portuguese continental shelf [9].

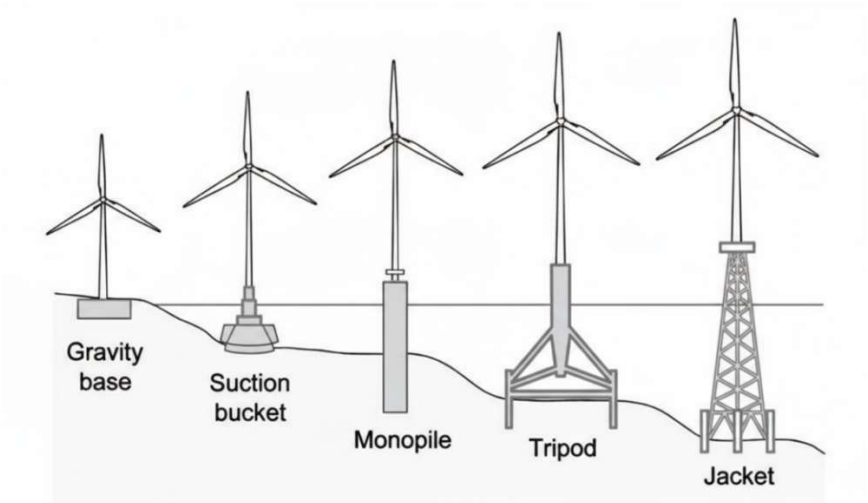


Figure 1 – Types of fixed-bottom foundations [24].

The most predominant foundation type in the industry is the monopile, a simple yet effective design consisting of a large-diameter hollow steel cylinder that is driven deep into the seabed [22]. Monopiles have historically accounted for the largest share of installed offshore wind foundations due to their relative ease of fabrication and the straightforward nature of their installation process, which typically involves hydraulic hammering or drilling, depending on the substrate [21]. Structurally, a monopile

supports the tower and the wind turbine generator as a cantilever beam, with its stability derived from soil-structure interaction and the bending stiffness of the steel pile itself [22].

However, the applicability of monopiles is strictly limited by water depth and turbine size. A monopile behaves as a slender cantilever subjected to combined lateral loading and a significant concentrated mass at its free end, corresponding to the tower, nacelle, and rotor assembly. As water depth and hub height increase, the unsupported length of the pile and the associated bending moments at the seabed grow substantially [22]. In parallel, the structure's dynamic response becomes increasingly critical. To control lateral displacements and ensure adequate separation between the natural frequencies of the system and dominant excitation frequencies associated with wind, waves, and rotor rotation, a significant increase in structural stiffness is required, typically achieved through larger pile diameters and increased wall thickness [22].

These constraints become particularly limiting in water depths exceeding approximately 30 to 40 m [25]. At greater depths, such as those typical of the Portuguese Atlantic coast [26], the amount of steel required to satisfy both static and dynamic performance criteria increases rapidly, rendering monopile foundations economically inefficient when compared to alternative support solutions [25].

Gravity-Based Structures represent another alternative, relying on the sheer mass of a concrete caisson filled with ballast to resist overturning moments caused by wind and wave loading [22]. While this solution avoids the need for pile driving, thereby mitigating underwater noise pollution, they require substantial seabed preparation to ensure a flat foundation interface and is generally limited to shallower waters due to the colossal size and weight required to maintain stability in deeper environments [21].

To extend the operational envelope of fixed-bottom wind into transitional water depths, typically 30 to 60 m, the industry adopted space-frame structures such as jackets and tripods [22]. These designs utilize a lattice of smaller-diameter steel tubes to distribute loads across a wider footprint, offering a higher stiffness-to-weight ratio than monopiles [22].

Jacket foundations, consisting of four legs connected by cross-bracing, function similarly to conventional offshore oil platforms and are anchored to the seabed using piles at each foot [22]. While these structures allow for deployment in deeper waters, they introduce significant complexity in terms of manufacturing, requiring numerous welded joints, and installation, often necessitating larger heavy-lift vessels and prolonged offshore operations [26].

Despite the maturity of these fixed-bottom systems, they confront a hard technical and economic limit as water depths approach 60 m [21]. Beyond this depth, the capital expenditure (CAPEX) associated with the fabrication and installation of rigid structures rises non-linearly, destroying the project's financial viability [26]. This depth limitation is particularly problematic for nations with narrow continental shelves, such as Portugal [9]. Unlike the North Sea, the Portuguese coast is characterized by a steep bathymetric profile where water depths drop rapidly to several hundred metres only a short distance from the shore [9].

Recent resource assessments of the Portuguese coast indicate that the prime areas for wind energy extraction are predominantly located in waters exceeding the feasible limit for fixed-bottom technology, exhibiting an average water depth ranging from 100 m to over 150 m [9]. In these deep-water environments, fixed foundations are technically impossible or economically prohibitive [27]. Consequently, the continued expansion of offshore wind in Portugal and similar Atlantic jurisdictions necessitates a fundamental decoupling of the energy capture system from the seabed, driving the transition toward floating offshore wind (FOW) systems [27].

2.1.2. FLOATING OFFSHORE WIND SYSTEMS

2.1.2.1. Floating Platform Concepts and Technological Characteristics

To overcome the depth limitations inherent to fixed-bottom foundations and to access the superior wind resources located in deep-water environments, the industry has transitioned towards the development of FOW sytems [26]. FOW systems fundamentally differ from their fixed counterparts by decoupling the support structure from the seabed, relying instead on a balance of buoyancy and mooring forces to maintain stability [22], as shown in Figure 2. This decoupling allows for the deployment of wind turbines in water depths ranging from 60 m to over 1,000 m, effectively unlocking most of the technical offshore wind potential in the Atlantic Ocean [9].

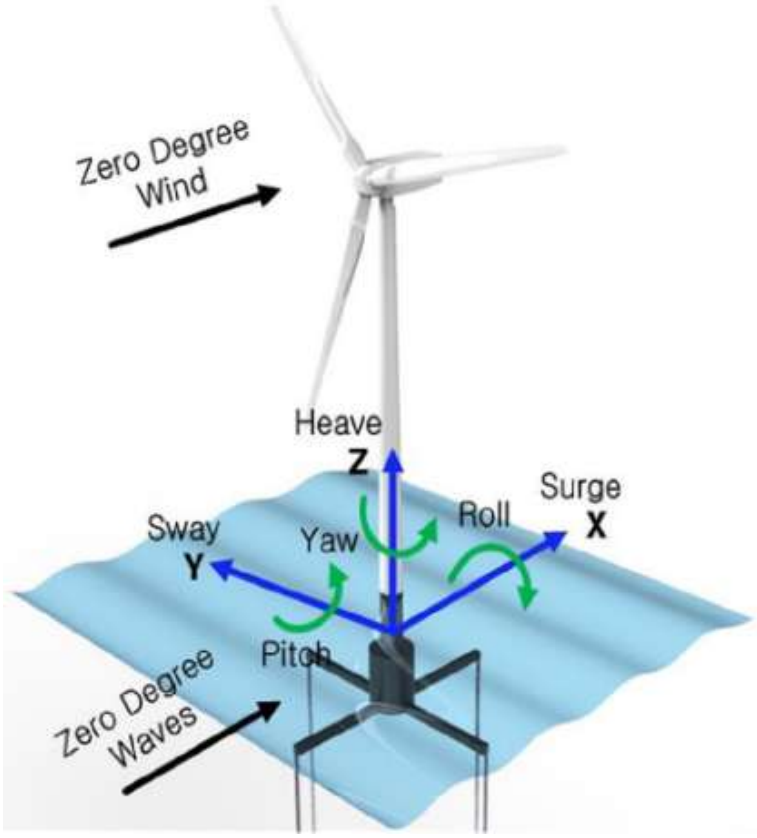


Figure 2 - Platform degrees of freedom [28].

However, this transition introduces a significant increase in dynamic complexity, whereas a fixed tower behaves as a stiff cantilever with negligible base motion, a floating platform is a compliant structure subject to rigid-body motion in six degrees of freedom: surge, sway, heave, roll, pitch, and yaw [22].

The design of a FOW platform is governed by the necessity to limit these motions, particularly pitch and roll, to ensure the aerodynamic efficiency of the turbine and to prevent structural fatigue [29]. To achieve this, floating concepts are generally categorized into three primary stability classes based on the physical mechanism used to provide the restoring moment against overturning forces: ballast-stabilized, waterplane-area-stabilized, and mooring-stabilized [22].

The application of these stability concepts has led to four principal floating offshore wind foundation typologies: the barge, the semi-submersible, the spar, and the tension leg platform (TLP). These concepts are presented schematically in Figure 3 and are discussed individually in the following sections.

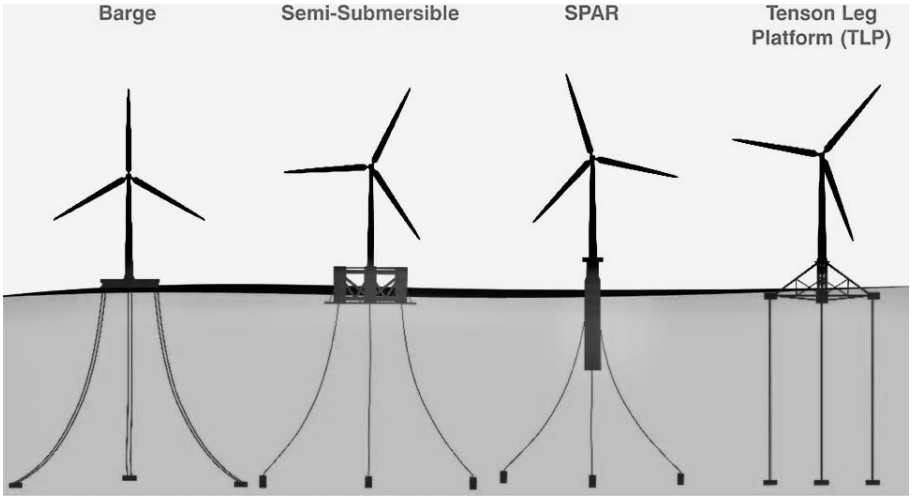


Figure 3 - Type of FOWT foundations [30].

The Spar Buoy, which is ballast-stabilized, is typically a deep-draft, vertical cylindrical structure, exemplified by the Hywind concept deployed by Statoil (now Equinor) [29]. The stability of a spar is derived from a low centre of gravity achieved by placing heavy ballast material, such as water and rocks, at the bottom of the cylinder, ensuring that the centre of gravity remains well below the centre of buoyancy [29]. This configuration creates a large righting moment that naturally resists pitch and roll motions [22]. The simple geometry of the spar offers advantages for manufacturing and hydrodynamic modelling. However, the requirement for a very deep draft, often exceeding 100 m, poses significant logistical challenges [29], [31]. The mooring system for a spar usually consists of a catenary configuration with drag-embedment anchors, which, while reliable, results in a substantial seabed footprint [22].

Semi-submersible platforms, which are waterplane-area-stabilized, rely on the principle of distributed buoyancy to achieve stability [32]. These structures consist of multiple vertical columns connected by horizontal pontoons or bracing members [22]. By spacing the columns far apart, the platform maximizes its second moment of waterplane area, which provides a high hydrostatic stiffness to resist overturning moments [22]. A key advantage of semi-submersibles is their shallow draft, often less than 20 m when de-ballasted, which allows for the full assembly of the turbine onto the platform at the quayside of a standard port [31], [32]. The integrated unit can then be towed to the offshore site using conventional tugboats, significantly reducing the dependency on scarce and expensive heavy-lift crane vessels [26], [31]. To mitigate the heave motion induced by waves, semi-submersibles often incorporate heave plates at the base of the columns, which increase the hydrodynamic added mass and viscous damping of the structure [32]. This technology has been validated in the Portuguese context through the WindFloat Atlantic project, demonstrating its suitability for the Atlantic wave climate [27].

The Tension Leg Platform, classified as a mooring-stabilized type, differs fundamentally from the other two concepts as it relies on the tension in its mooring system for stability rather than hydrostatic properties alone [33]. A TLP is a highly buoyant structure that is pulled down into the water by vertical tendons anchored to the seabed, creating a large excess buoyancy that keeps the lines taut [33]. This tension prevents the platform from heaving or pitching significantly, resulting in a very stiff system that behaves similarly to a fixed foundation in the vertical degrees of freedom [33]. While TLPs offer the advantage of a minimal seabed footprint and low vertical motion, they are operationally complex [31]. The mooring system must be installed with high precision, otherwise, the platform becomes unstable if the tendons fail [33]. Furthermore, the vertical load on the anchors is extreme, requiring expensive pile or suction anchors rather than simple drag anchors [22], [31].

Barge platforms, an emerging design alternative, represent one of the simplest floating foundation concepts for offshore wind turbines, characterized by a large, flat-bottomed hull with a substantial waterplane area [34]. Stability is primarily achieved through hydrostatic restoring forces generated by the wide waterplane, rather than by deep ballast concentration or taut mooring systems. As a result, barge-type platforms typically exhibit large platform dimensions and relatively shallow drafts, which can offer advantages in terms of simplified fabrication, transportation, and installation, particularly in shallow to moderate water depths [34].

However, the large waterplane area also makes barge platforms inherently more sensitive to wave-induced pitch motions, requiring careful hydrodynamic and structural design to ensure acceptable dynamic performance. Consequently, despite their conceptual simplicity and potential cost benefits, barge-type foundations have seen limited commercial deployment to date, with most applications remaining at the prototype or pilot scale [34].

For the specific context of the Portuguese continental shelf and the co-located arrays proposed in this thesis, semi-submersible platforms are often favoured due to their flexibility in installation and compatibility with local port infrastructure, as evidenced by the strategic choice for the Viana do Castelo developments [27]. However, regardless of the specific FOW concept, the deployment of these large floating structures in arrays introduces the critical challenge of mooring line management, as the spread of anchors can occupy a seabed area significantly larger than the turbine spacing itself, constraining the placement of co-located devices [27].

2.1.2.2. Turbine Scaling and Emerging High-Capacity Offshore Classes

The progressive increase in offshore wind turbine size has been one of the most significant design trends in the evolution of the sector. Early offshore reference studies were centred on turbine classes that are now comparatively modest by current standards. A well-known example is the NREL offshore baseline defined by *Jonkman* [33], which adopted a 5 MW machine with a 126 m rotor diameter as a representative large offshore turbine for coupled aero-hydro-servo-elastic analysis, even though the broader rating range considered during the initial sizing discussion extended from 5 MW up to 20 MW [33]. This illustrates that, even at an early stage, offshore wind development was already understood as a field in which continued scaling would play a central role.

To clarify the scale progression observed in offshore wind turbine development, Figure 4 summarises representative turbine classes and associated rotor diameters reported in the literature.

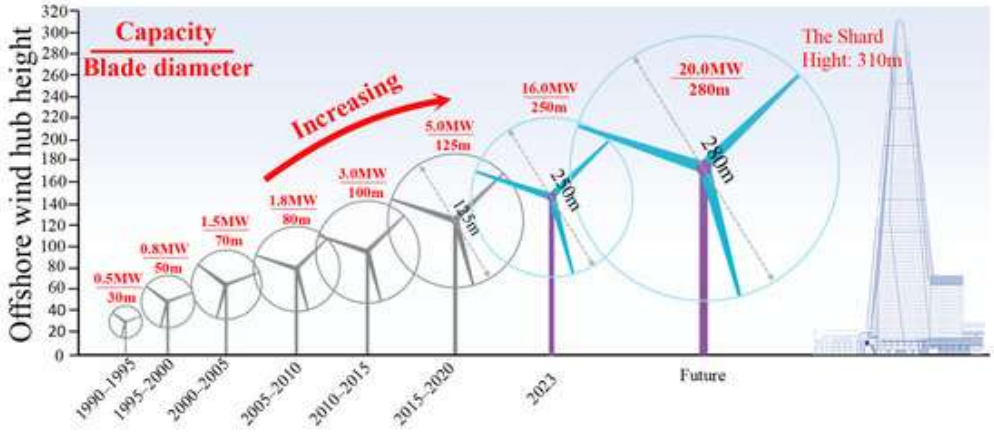


Figure 4 - Indicative evolution of offshore wind turbine capacity and rotor dimensions [35].

After Figure 4, the historical trend becomes clear. Offshore wind turbine development has progressed not only toward higher rated power, but also toward substantially larger rotor diameters and overall envelope dimensions. In the *NYSERDA* [36] layout study, generic offshore wind turbine classes of 8, 10, 12, and 15 MW were defined to support early-stage layout and spacing analysis, with proposed rotor diameters of 164, 177, 194, and 217 m, respectively. The same study notes that these values were selected as representative design cases rather than as exact commercial products, reflecting the way offshore wind planning commonly uses turbine class assumptions to define indicative spacing, tip height, and area requirements [36].

This scaling trend has direct engineering consequences for offshore wind farm design. As turbine rating increases, fewer machines are required to achieve a given installed capacity, which may reduce the total turbine count within a project. At the same time, larger rotor diameters increase the characteristic spacing required between turbines when separation distances are expressed in multiples of rotor diameter, and they also affect tip height, air-gap requirements, support-structure mass, installation demands, and electrical system design [36]. In floating wind applications, this trend is especially relevant because the increase in turbine size also propagates into platform sizing, mooring loads, dynamic response, and the effective spatial footprint of each unit.

The transition toward higher-capacity machines is also reflected in current floating-wind modelling practice. Recent floating offshore wind optimisation and cable-layout studies already use the IEA15MW offshore reference FOWT as the turbine model for representative large-scale analyses, indicating that the 15 MW class has become a credible benchmark for forward-looking floating wind engineering studies [37]. This is important from a state of the art perspective, because it shows that large-capacity turbines are no longer only conceptual extrapolations, but are already embedded in the modelling frameworks used to study realistic floating wind farm behaviour.

The upper end of this scaling trajectory is further illustrated by the recent emergence of the 20 MW class in industry reporting. Although this power level should not be interpreted as the current dominant commercial standard, it provides a useful indication of the direction in which offshore wind turbine development is moving, particularly in the context of next-generation large-scale offshore projects.

A recent example of this latest scaling step is the installation of a 20 MW turbine, with a 300m rotor, by the Chinese company China Three Gorges, as shown in Figure 5.



Figure 5 - Installation of the CTG-Goldwind 20MW offshore wind turbine [38].

As illustrated by Figure 5, the 20 MW class is more appropriately interpreted as the current frontier of offshore turbine upscaling than as an already established market baseline [38]. Its significance within the state of the art lies not in suggesting that this rating has become standard practice, but in demonstrating that the long-term trajectory of offshore wind development continues to move toward larger and more powerful individual machines.

This ongoing scaling trend reflects a broader sectoral shift toward reducing the number of turbines required to achieve a given installed capacity while simultaneously increasing the spatial and structural influence of each unit within the overall wind farm configuration [22], [36].

As turbine ratings and rotor dimensions continue to increase, the consequences extend well beyond energy capture alone, directly influencing turbine spacing requirements, wake interaction patterns, support-structure dimensions, installation and maintenance logistics, and the overall layout logic of future offshore projects [36]. In the context of floating offshore wind, these implications become even more pronounced, since larger turbines also propagate into increased platform dimensions, higher structural demands, more complex mooring and anchoring systems, and wider operational exclusion areas [37].

For these reasons, the emergence of the 20 MW class is particularly relevant, as it provides a clear indication of the design direction shaping next-generation offshore wind farms and of the spatial and engineering challenges that will increasingly govern their configuration.

2.1.3. Wake and Aerodynamic Constraints

Wake interactions represent the dominant aerodynamic constraint governing the spatial layout of offshore wind farms [22], [39]. When a wind turbine extracts kinetic energy from the incoming flow, it generates a wake characterised by a reduction in the mean wind velocity and an increase in turbulence intensity downstream of the rotor [22]. These wake-induced flow disturbances persist over multiple rotor diameters and directly influence both the energy production and the structural loading of downstream turbines, thereby imposing minimum spacing requirements that fundamentally shape feasible wind farm layouts [40], [41].

From an energy perspective, wake effects are particularly significant because wind turbine power output scales with the cube of wind speed. Even moderate velocity deficits within a wake can therefore lead to substantial reductions in downstream power production [22], [41]. In parallel, the high turbulence levels within wakes amplify unsteady aerodynamic loads acting on the rotor, drivetrain, and support structure, accelerating fatigue damage accumulation and reducing component lifetimes [22], [39]. Consequently, wake interactions affect not only array efficiency but also the long-term reliability of offshore wind turbines. In addition, wake modelling literature has shown that wake-added turbulence and wake meandering can have a major impact on wind turbine loading, reinforcing turbine separation as a structural-reliability driver and not only an energy-yield driver [42].

A key factor controlling wake development and recovery is the ambient (free stream) turbulence intensity. Higher turbulence levels promote enhanced mixing between the wake and the surrounding flow, resulting in faster wake recovery. Conversely, low-turbulence environments produce more coherent and persistent wakes that extend further downstream. Experimental investigations by *Gambuzza and Ganapathisubramani* [40] demonstrate that reduced free-stream turbulence intensity leads to slower wake recovery and increased wake meandering, with wake effects persisting over multiple downstream distances expressed in rotor diameters.

These findings are particularly relevant for offshore wind farms, where atmospheric conditions are typically characterised by relatively low turbulence intensity and high directional persistence [40], [43]. Under such conditions, wakes tend to remain coherent over longer distances than in many onshore environments, increasing the likelihood of wake overlap when turbines are placed too closely together, as presented in Figure 6 [40]. As a result, offshore wind farm layouts generally require larger inter-turbine spacing than onshore layouts to limit wake-induced energy losses and excessive dynamic loading [39], [40], [41].

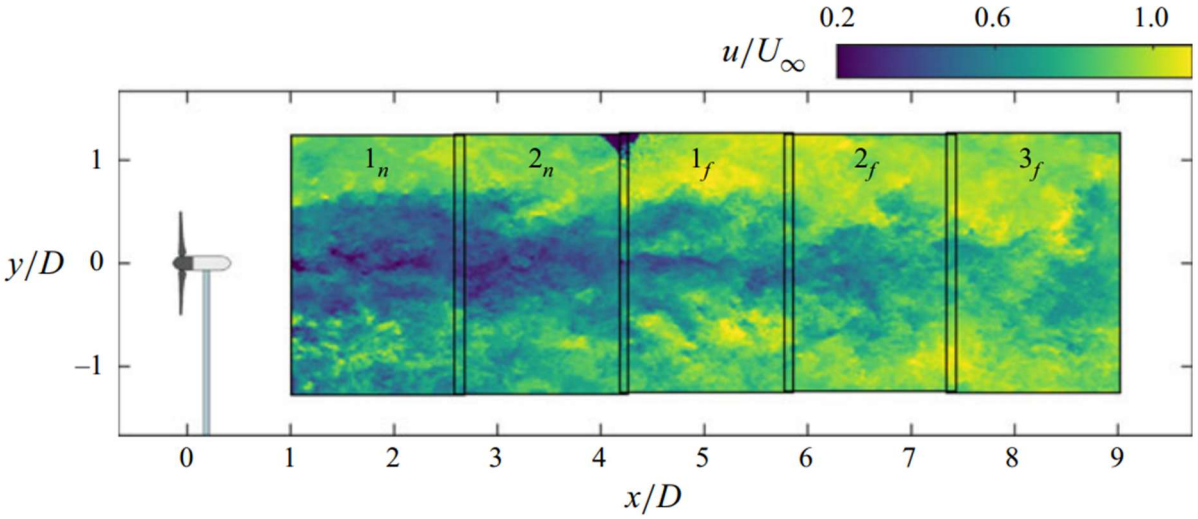


Figure 6 - Turbine wake effect [40].

In practical wind farm design and optimisation studies, wake effects are commonly incorporated through explicit minimum inter-turbine spacing constraints, expressed as multiples of the turbine rotor diameter, D. These constraints are directly motivated by wake persistence and by the need to maintain acceptable levels of power loss and structural loading. In the context of floating offshore wind, *Tao et al.* [44] impose a minimum spacing constraint of 5D between floating wind turbines to ensure safe and feasible operation, reflecting conservative offshore engineering practice.

This choice is further supported by the experimental results of *Gambuzza and Ganapathisubramani* [40], whose measurements indicate that, beyond approximately 5D downstream, the mean velocity deficit begins to recover significantly toward free-stream conditions, marking a transition toward reduced wake influence.

A few other optimisation studies explore tighter spacing limits for comparative or methodological purposes. For example, *Clark et al.* [39] adopt a minimum spacing of 3D in certain reliability-based optimisation scenarios. While such values are useful for sensitivity analyses, they are generally associated with stronger wake interactions and are less representative of layouts intended for commercial-scale offshore deployment.

Beyond optimisation-only constraints, operational offshore wind farms also illustrate that spacing is often implemented as a direction-dependent design variable, with larger separations aligned with dominant wind directions rather than a single isotropic distance. A representative example is the Princess Amalia offshore wind park, , for which the published turbine-location map explicitly indicates different characteristic separations in terms of rotor diameters within the array [45].

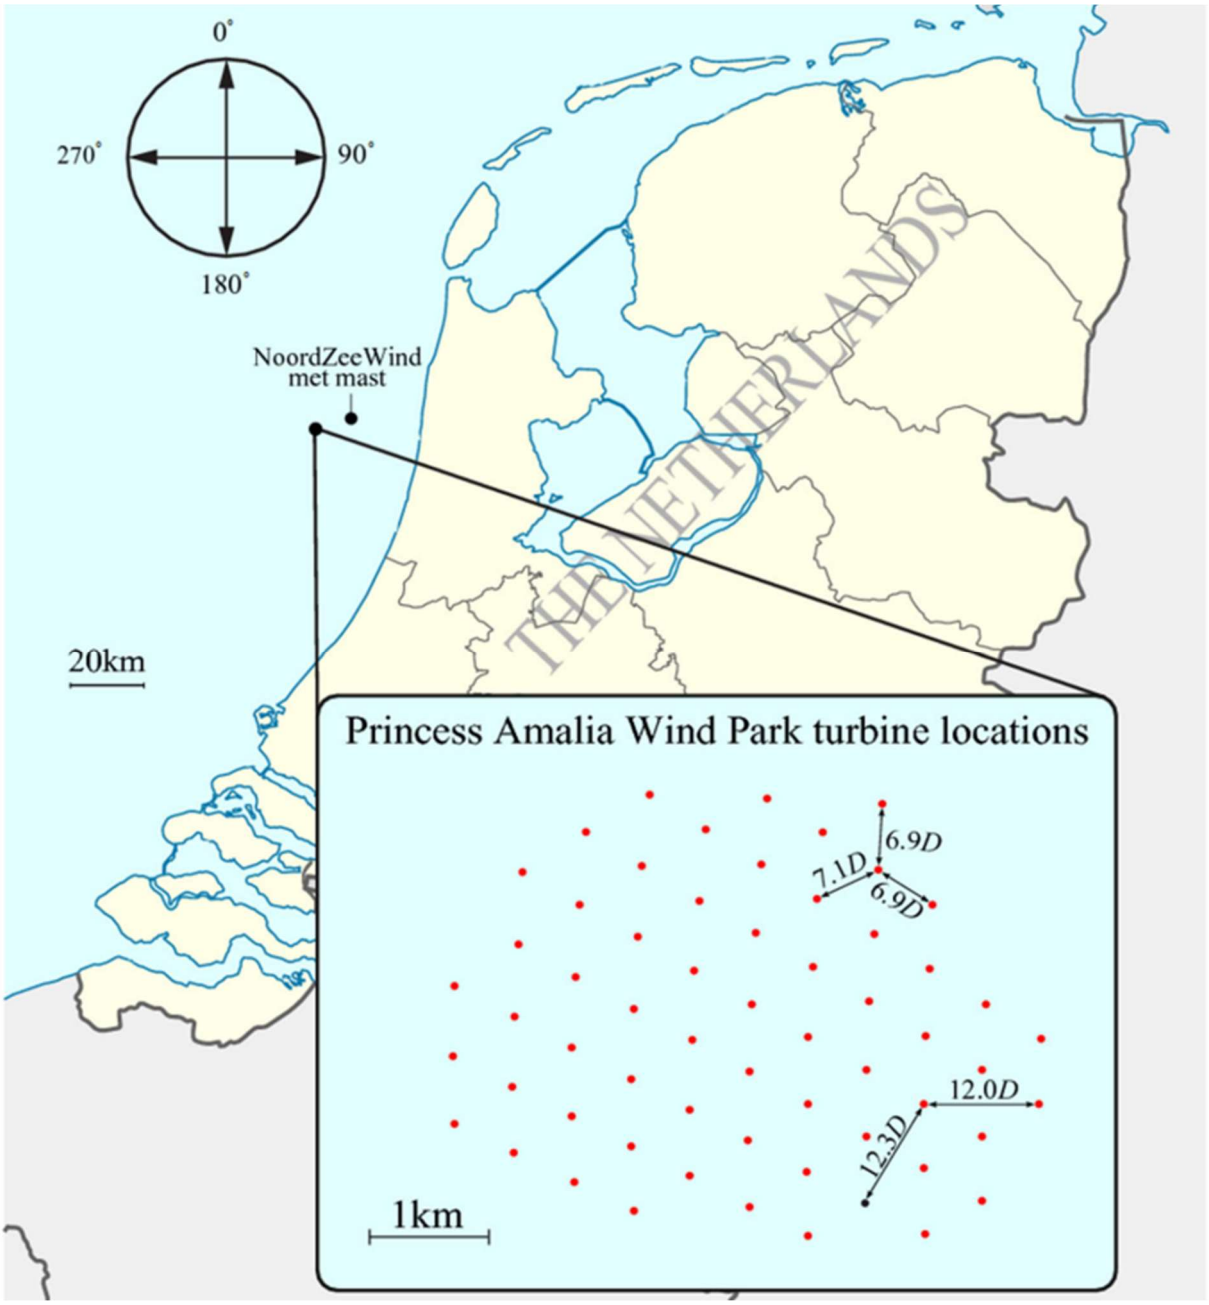


Figure 7 - Princess Amalia Wind Park turbine locations with indicative directional spacing expressed in D [45].

For floating offshore wind farms, rotor-diameter-based spacing remains a practical aerodynamic descriptor, but its interpretation should also account for floating-specific effects that may alter the effective turbine-to-turbine separation during operation. In particular, mean platform offsets within the mooring watch circle can shift turbine positions relative to their nominal design coordinates, potentially modifying wake overlap and wake-loading patterns, therefore, conservative spacing assumptions and/or offset-aware spacing checks are recommended during early-stage floating-array definition [46].

Briefly, the literature indicates that wake effects constitute a fundamental aerodynamic constraint that must be accounted for when defining offshore wind farm layouts. Although the optimal spacing depends on site-specific wind conditions, turbine characteristics, and design objectives, the use of minimum spacing thresholds expressed in rotor diameters provides a practical and widely adopted mean of

embedding wake considerations into layout design. Detailed spacing ranges and additional examples from operational offshore wind farms are consolidated later in Section 2.1.6 and Table 2.

2.1.4. Economic and Infrastructure Constraints

In addition to aerodynamic considerations, offshore wind farm layouts are strongly constrained by economic and infrastructure-related factors, among which the electrical system plays a central role. The spatial arrangement of turbines directly affects the length, routing, and complexity of inter-array cabling, the placement of offshore substations, and the definition of cable corridors, all of which have a measurable impact on capital expenditure and spatial feasibility [26], [41].

Figure 8 presents a conventional offshore wind farm in which turbines are not connected individually to the shore. Instead, electrical power is collected via inter-array cables that connect multiple turbines in series, each operating at medium voltage. These strings are then connected to one or more offshore substations, where voltage is stepped up before transmission to shore via export cables. This radial or tree-type collection architecture is widely adopted because it limits the number of long-distance cables, reduces material usage, and simplifies protection and maintenance strategies [47], [48].

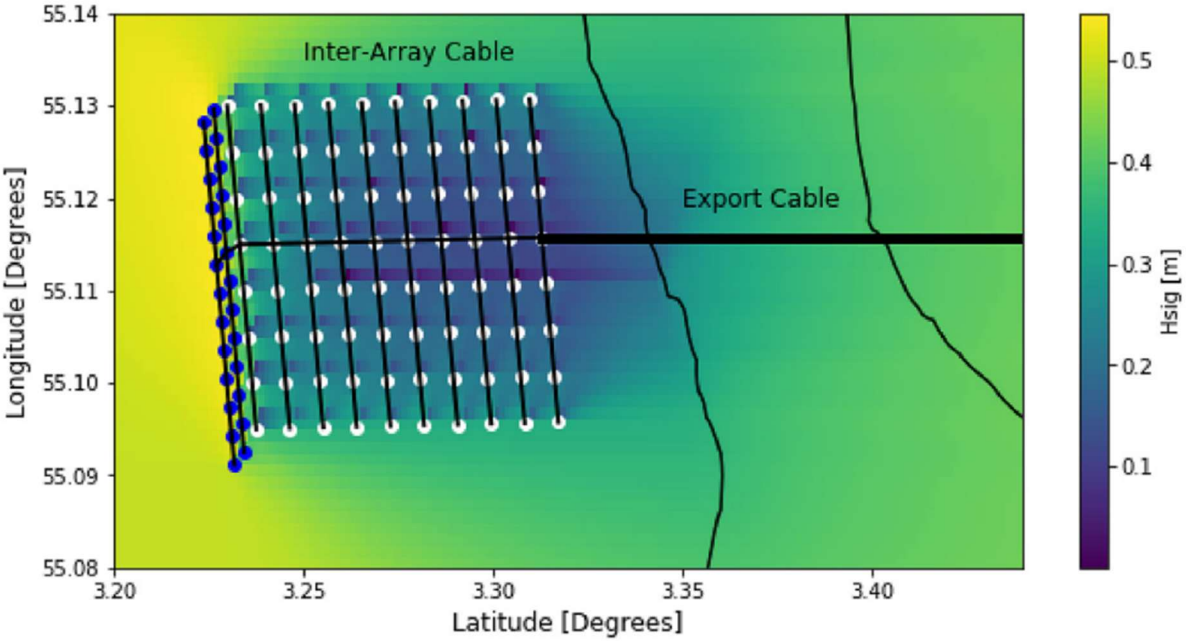


Figure 8 - Cable layout for co-located array [41].

From an economic standpoint, the electrical infrastructure represents a significant share of the total project cost. Analytical cost models for co-located floating wind-wave energy parks show that the electrical system, including inter-array cables, export cables, and offshore substations, can account for approximately 20% of the total CAPEX in floating offshore wind projects [37], [41]. This contribution is typically higher than in fixed-bottom projects due to greater water depths, longer cable lengths, and the use of dynamic cables to accommodate platform motions.

The dependence between layout geometry and electrical cost is therefore direct. Increasing inter-turbine spacing to reduce wake losses inevitably increases the total length of inter-array cables and the distance

between turbines and offshore substations. *Pérez-Rúa et al.* [37] demonstrate that the optimisation of inter-array cable networks in floating offshore wind farms is a combinatorial problem, in which turbine positions, cable routing, and substation placement are tightly coupled. Their results highlight that layouts which are favourable from an aerodynamic perspective may become economically inefficient once dynamic cable lengths and installation constraints are accounted for [37].

The number and location of offshore substations further structure the wind farm layout. Large commercial offshore wind projects typically employ one to three offshore substations, depending on total installed capacity, collection voltage, and export system design [48]. Substations are generally located within or near the centre of the wind farm, or distributed across the site in larger developments, in order to minimise inter-array cable lengths and avoid excessive cable crossings [37], [47]. As fixed electrical hubs, offshore substations act as geometric anchors that define preferred cable corridors and influence the orientation and extent of turbine strings.

Beyond cost considerations, electrical infrastructure introduces hard spatial constraints. Inter-array and export cables must be routed along seabed corridors that remain accessible for installation, protection, inspection, and repair throughout the project lifetime. In addition, offshore substations require safety buffers for navigation, installation, and operational reasons. Regulatory and industry guidance commonly adopts conservative safety zones of up to 500m around offshore transformer stations, as shown in Figure 9, effectively creating local exclusion areas where additional structures cannot be installed [49].

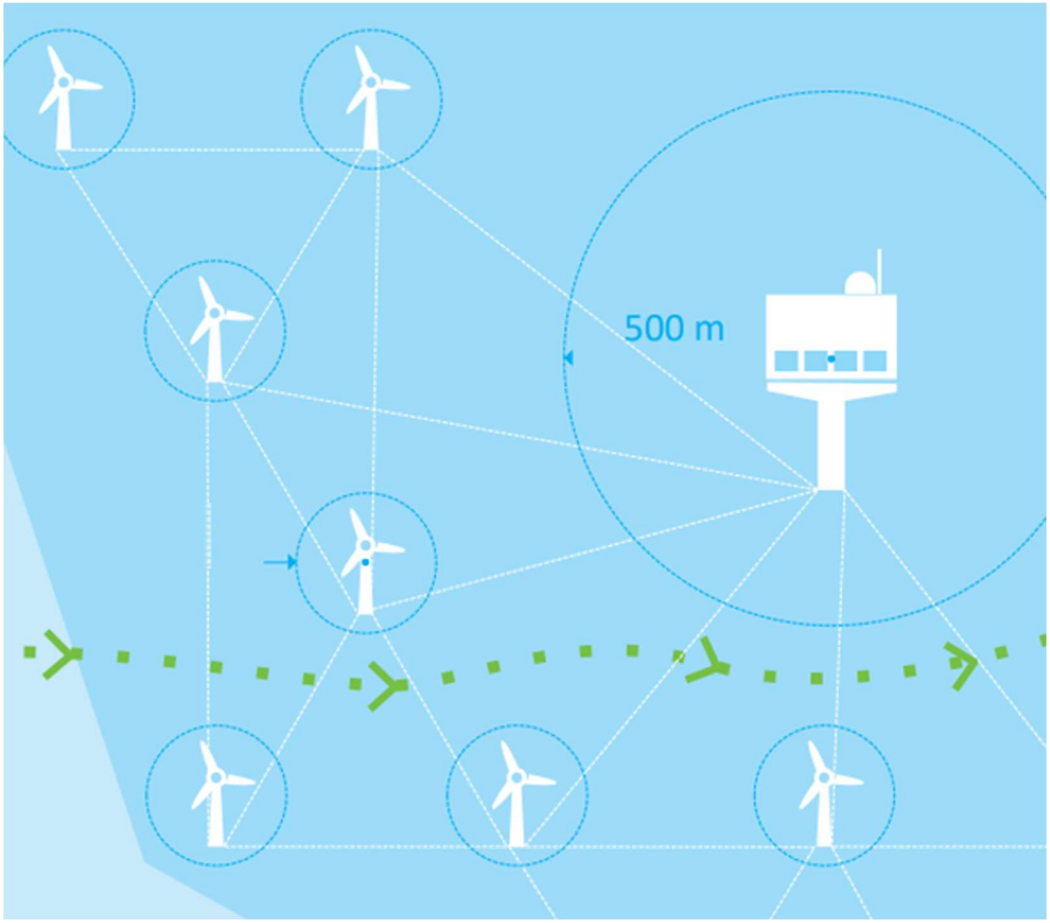


Figure 9 - Safety zone around offshore substation [49].

Taken together, these factors indicate that economic and infrastructure constraints do not merely penalise dispersed layouts financially but also reduce the effectively usable seabed area within the wind farm. Inter-array cabling costs limit excessive turbine dispersion, offshore substations introduce fixed spatial anchors, and cable corridors and safety buffers generate exclusion regions that must be respected. Consequently, realistic offshore wind farm layouts emerge from a compromise between wake-driven spacing requirements and the need to control electrical infrastructure cost and spatial complexity [26], [37], [41].

In co-located developments, the same corridors and exclusion buffers also constrain the internal allocation of space for wave-energy arrays, as WEC modules must remain separated not only from turbine positions but also from cable routes and substation safety zones, which can fragment the wind farm interior into a smaller admissible deployment domain.

2.1.5. Floating-Specific Spatial Constraints

Floating offshore wind farms are subject to spatial constraints that differ fundamentally from those governing fixed-bottom installations, as the station-keeping systems required for floating platforms significantly increase the seabed area influenced by each turbine [22], [26]. While fixed-bottom turbines occupy a relatively small and well-defined footprint at the seabed, floating wind turbines rely on mooring systems and dynamic cables whose geometry and motion introduce additional spatial constraints that must be explicitly accounted for in layout design [9], [26]. Floating wind turbines are typically stabilised using mooring systems, most commonly catenary or semi-taut configurations, as illustrated in Figure 10, connecting the floating platform to anchors on the seabed [22].

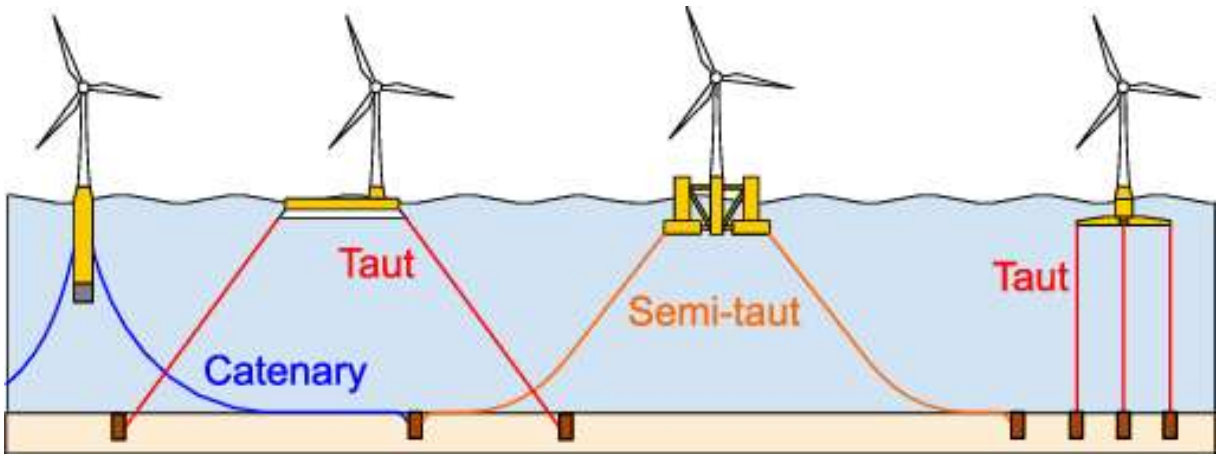


Figure 10 - Different floating offshore mooring systems [50].

In these systems, mooring lines extend radially from the platform and may rest partially on the seabed, generating a substantial horizontal footprint whose size depends primarily on water depth, mooring line length, and environmental loading. Recent offshore engineering studies ([31], [51]) provide explicit order-of-magnitude guidance for this footprint, indicating that the horizontal extent of mooring systems is typically on the order of 4 to 8 times the local water depth for preliminary layout and spacing assessments, as presented in Figure 11, from the work made by *Ma K. et al. (2023)* [31].

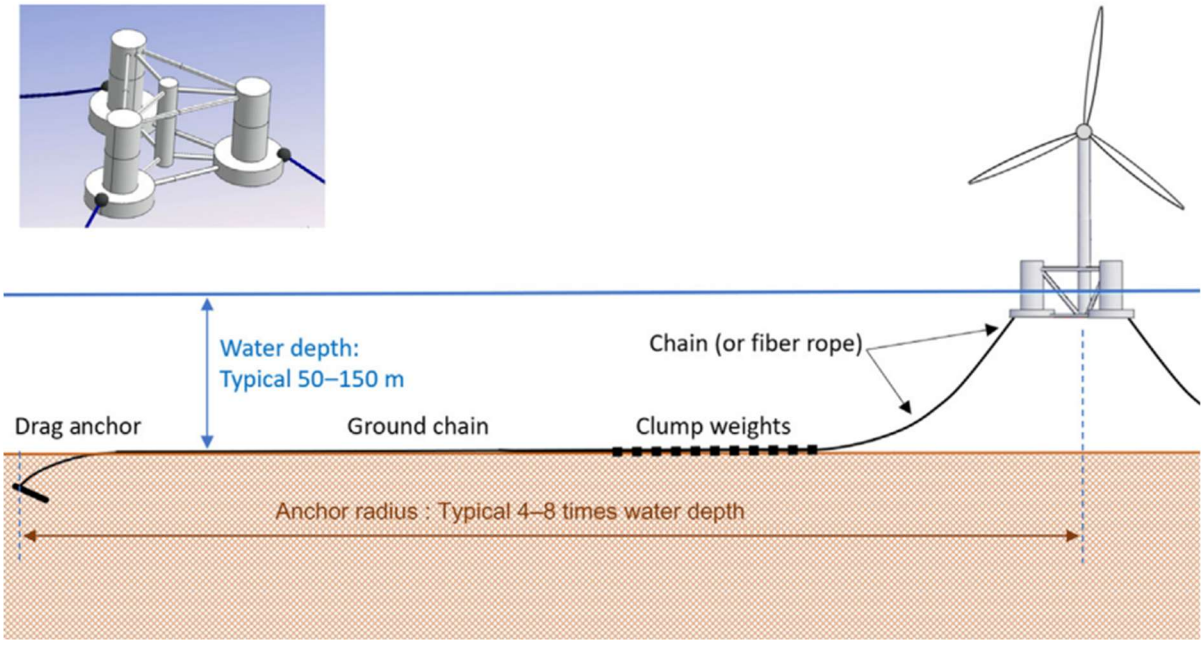


Figure 11 - Typical spread mooring design with catenary lines [31].

For representative water depths of 80 to approximately 150 m, which are characteristic of many prospective floating wind sites along the Portuguese Atlantic margin, this rule-of-thumb implies a mooring footprint radius in the range of approximately 320 to 1200 m from the turbine location. These distances correspond to the region potentially occupied by mooring lines and anchors under static or quasi-static conditions and already exceed typical aerodynamic spacing requirements derived from wake considerations alone [9], [52].

Beyond the static footprint of the mooring system, floating offshore wind farm layouts must account for dynamic exclusion zones associated with platform motions and environmental extremes. Floating platforms experience surge, sway, and yaw motions under combined wind and wave loading, as presented in Figure 2, which induce horizontal excursions of mooring attachment points and dynamic inter-array cables [21]. To prevent line crossing, entanglement, or excessive bending during extreme events, additional spatial clearance beyond the nominal mooring footprint is required [37].

Recent work by *Pérez-Rúa et al. (2024)* [37] provides quantitative evidence of this requirement by defining safety zones around floating wind turbines using coupled aero-hydro-servo-elastic simulations that account for mooring systems and dynamic electrical cables. Their results demonstrate that, when realistic platform motions and cable dynamics are considered, additional exclusion distances of several tens to hundreds of metres beyond the nominal mooring footprint may be required to ensure safe operation and prevent mechanical interference [37].

Integrated spatial assessments of floating wind and co-located wind-wave energy systems further confirm that these floating-specific constraints can become the dominant factor controlling feasible layouts, often exceeding wake-driven spacing constraints in terms of occupied seabed area [9], [26]. Even when turbine spacing satisfies aerodynamic and electrical infrastructure requirements, large portions of the seabed between turbines may remain unavailable due to overlapping mooring footprints, dynamic cable excursions, and associated safety buffers [9].

From a spatial planning perspective, these findings imply that the apparent free space within a floating offshore wind farm is significantly reduced once mooring systems and exclusion zones are properly

accounted for. This reduction has been explicitly highlighted in recent layout optimisation and site-selection studies, which show that neglecting mooring and safety constraints can lead to layouts that are infeasible from an installation or operational standpoint [37], [52].

Overall, floating-specific spatial constraints impose hard geometric limitations that fundamentally shape the feasible layouts of offshore wind farms. The large seabed footprint associated with mooring systems, combined with additional exclusion margins required to accommodate platform motions and dynamic cable behaviour, must be explicitly incorporated alongside wake and infrastructure-related constraints when defining realistic layouts for floating offshore wind developments [9], [26], [37].

Figure 12 presents an example of the exclusive zone used by *Teixeira-Duarte et. al (2025)* [53] in the study of the WindFloat Atlantic farm.

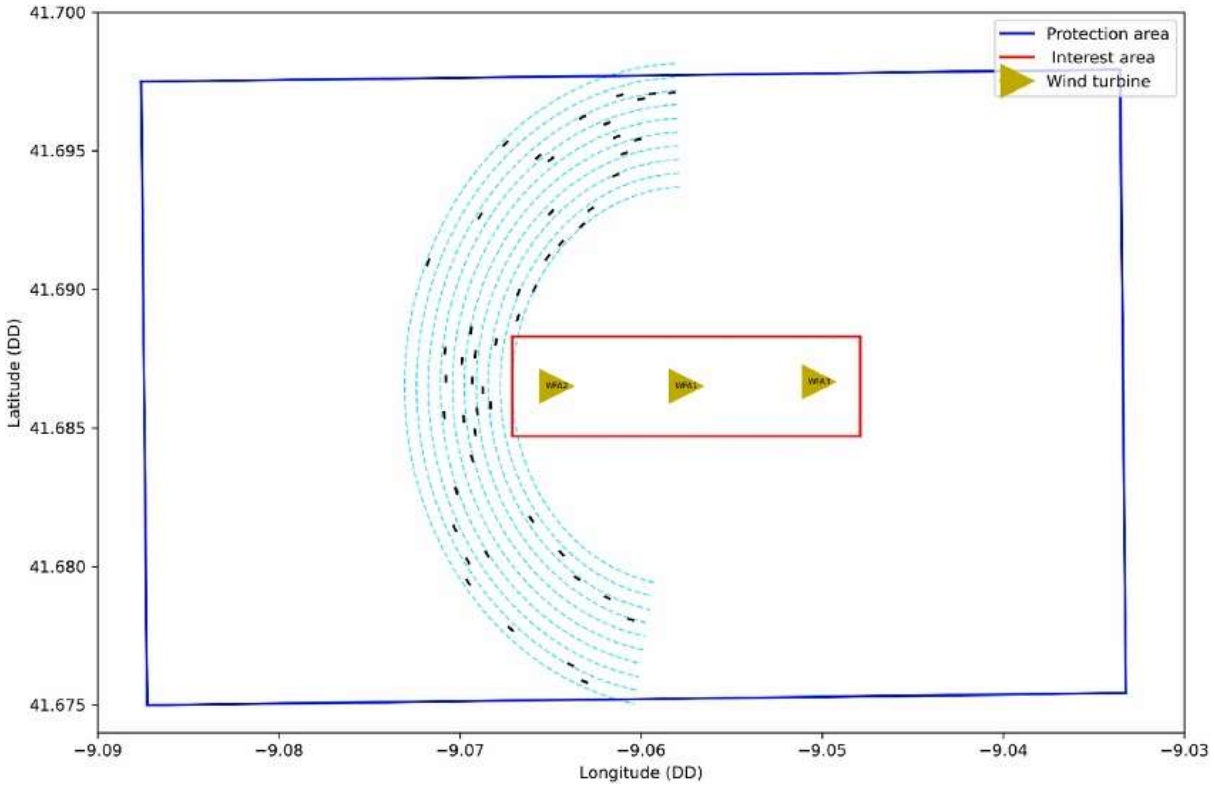


Figure 12 - Exclusive zone (red rectangle) for the WindFloat Atlantic farm [53].

2.1.6. Reference Offshore Wind Farms Layouts

The definition of a reference offshore wind farm layout must remain consistent with the set of aerodynamic, electrical, and floating-specific constraints identified in the previous subsections. Feasible offshore layouts are not determined by turbine spacing alone, but by the combined influence of wake interactions, inter-array and export cable routing, offshore substation safety clearances, navigation requirements, and, in floating applications, the additional seabed occupation associated with mooring systems and dynamic cable excursions [37], [54]. A realistic layout therefore emerges from the simultaneous consideration of these constraints rather than from a single-spacing rule.

Before examining how offshore wind farms are arranged in practice, the principal layout-defining constraints may be consolidated in Table 1. This synthesis provides the boundary conditions within which practical offshore layouts are developed and clarifies that spacing values must be interpreted together with infrastructure and safety requirements, rather than in isolation.

Table 1 - Summary of spatial constraints governing offshore wind farm layout design.

Layout constraint	Spatial meaning for layout	Typical/Reference value	Source
Streamwise turbine spacing	Minimum spacing along the dominant wind direction	5D - 10D	[40], [42], [44]
Crosswind turbine spacing	Minimum spacing between adjacent rows	5D - 7D	[44], [45]
Diagonal effective spacing	Governing distance for oblique wind direction	9D - 11D	[45]
Substation safety buffer	No-installation zone around the offshore substation	500 m	[49]
Cable corridors	Reserved seabed space for inter-array /export cables	Tens of metres	[37]
Mooring footprint	Area occupied by anchors and mooring lines	4 to 8 times the water depth	[31]
Additional floating safety buffer	Clearance for motions and dynamic interactions	Tens to hundreds of metres	[37]

When co-location is considered, these constraints should also be interpreted as a set of internal ‘no-placement’ structures that fragment the wind farm footprint into a reduced admissible space for wave energy modules.

With these constraints established, it becomes possible to examine how offshore wind farms are typically organised in commercial developments and planning studies. In offshore environments, where the wind resource is usually more spatially uniform than onshore, layout design is less concerned with locating isolated high-resource points and more concerned with arranging turbines so that wake losses remain acceptable without creating excessive electrical, installation, or operational complexity [55].

Commercial offshore wind farms have historically been developed predominantly as regular or lightly skewed grid-based arrays, because repeated geometric patterns simplify installation sequencing, vessel access, cable routing, navigation management, and long-term operation and maintenance [55]. In this type of arrangement, turbines are commonly distributed in repeated rows and columns, and the resulting array may be orthogonal, slightly rotated, or staggered while still preserving a clear constructional logic [55]. This structural regularity is also relevant for internal co-location, as it tends to create predictable corridors between turbine rows that may serve as candidate regions for wave-energy placement once turbine, mooring, and cable buffers are accounted for.

An early operational example of this type of row-based offshore organisation is illustrated by the Alpha Ventus layout, shown in Figure 13, which provides a clear visual reference of a compact offshore array structured through repeated turbine positioning and ordered spacing.

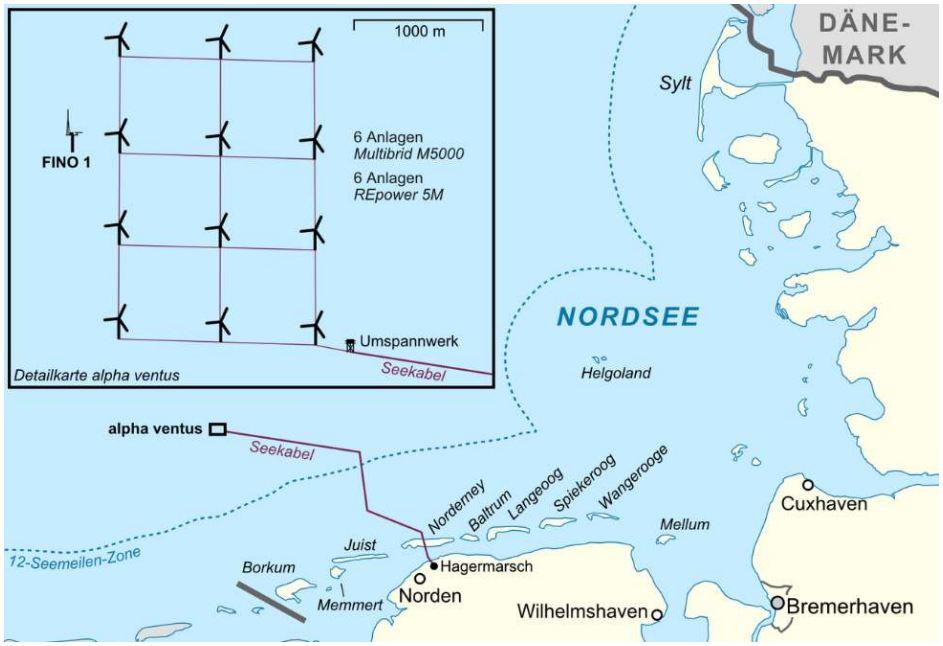


Figure 13 . Alpha Ventus' offshore wind farm layout [56].

As shown in Figure 13, early offshore wind farms often adopted simple and highly legible geometric arrangements, making them useful reference examples of regular offshore layout practice before larger and more recent projects introduced greater geometric variation.

A useful geometric interpretation of this practice is provided by the regular-layout formulation described by *Yong and Shengli (2016)* [55], in which offshore arrays are represented through a repeated parallelogram defined by two spacing directions, the angle between them, and the overall array orientation. This is relevant because it shows that even when an offshore layout is regular, it does not need to be a simple square grid; regularity may also be expressed through staggered equally spaced arrangements, provided that the spacing logic and row organisation remain consistent [55].

The historical evolution of offshore wind farm geometric patterns is illustrated in Figure 14.

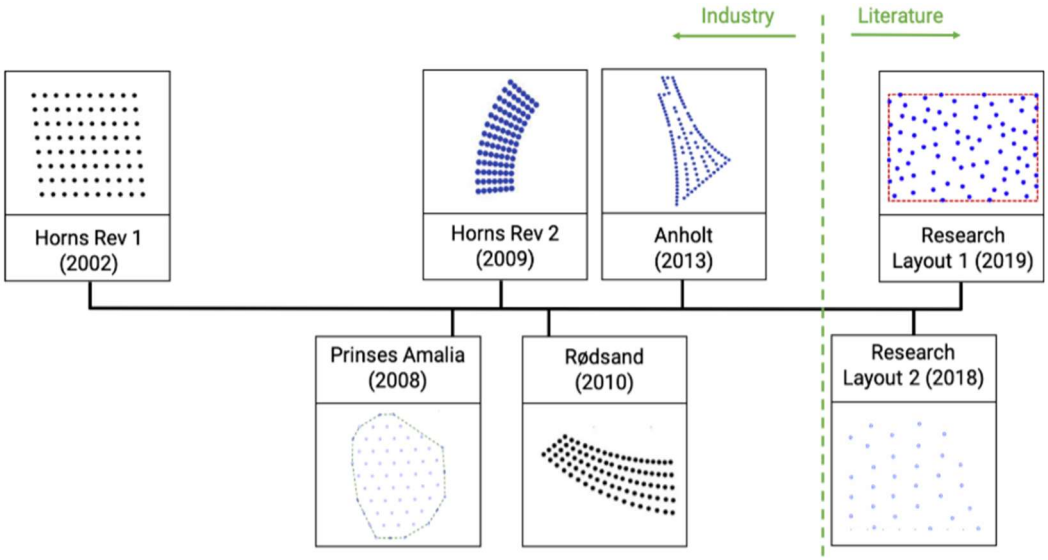


Figure 14 - Development of wind farm patterns over time [57].

Figure 14 is useful because it shows that earlier operational offshore wind farms were characterised by strongly regular patterns, whereas newer and larger projects exhibit greater geometric variation, including partial irregularity, while optimisation studies frequently generate even more irregular arrangements [57]. This distinction is important: the existence of more irregular layouts in optimisation literature does not imply that highly irregular geometries should be treated as the default industrial reference. In the comparison performed by *Sickler et al. (2023)* [57], irregular layouts improved overall energy production and reduced directional sensitivity of power output, but they also increased local effective turbulence and fatigue at the most affected turbine locations unless minimum spacing was carefully controlled. As a result, the literature indicates that regularity remains relevant not because it always maximises yield, but because it offers a robust compromise across multiple performance criteria [57].

Operational and planning examples reinforce this interpretation. Horns Rev 1 is one of the most widely cited benchmark offshore wind farms and is repeatedly used in wake studies as an example of a highly regular array, with approximately 7D spacing in the principal layout directions [54], [58]. Hornsea Project One demonstrates that the same underlying layout logic remains compatible with much larger developments, combining 1218 MW of installed capacity with 174 turbines in a single utility-scale project [58].

Dogger Bank planning documents similarly show that very large offshore projects continue to be governed by explicit separation rules. In the Hornsea Project One basis cited within the Dogger Bank documentation, the densest layout scenario considered a minimum turbine spacing of 924m, while broader Dogger Bank-related planning studies also use explicit minimum-separation criteria such as 750m or multiples of rotor diameter [59].

Planning-scale studies provide the same signal in a more general form. The *NYSERDA (2018)* [36] turbine layout study structures development using repeated 800 MW site areas and explicitly derives base-case layouts from a set of design drivers that include power density, inter-site distance, and turbine size assumptions. In that study, the base-case layouts are intentionally conservative, with spacing of approximately 13D in the prevailing wind direction and 10D in the cross-prevailing direction, while the sample of 27 European offshore wind farms reviewed in the same report shows average spacing values of approximately 7.5D downwind and 5.9D crosswind [36].

Representative project and planning benchmarks are summarised in Table 2.

Table 2 - Layout characteristics of representative large-scale offshore wind farms.

Wind farm	Capacity	Main Layout Signal	Source
Horns Rev 1	160 MW	80 turbines; 7D spacing in both directions	[54]
Hornsea One	1218 MW	174 turbines; spacing of 7D - 9D	[58]
Dogger Bank	Up to 1200 MW per project	minimum separation of approximately 6D	[59]
NYSERDA sites	800 MW per site	Conservative spacing of 10D - 13D	[36]

Overall, these examples indicate that offshore wind layout design is governed less by the search for a single optimal universal geometry and more by the repeated use of a consistent engineering logic.

Turbine separations are commonly expressed in rotor diameters, arrays are typically structured through repeated rows, and the final arrangement reflects a compromise between wake mitigation, electrical routing, constructability, safe navigation, and operational accessibility [36], [54], [55]. Even where project-specific adaptations are introduced, the underlying layout language remains one of ordered spacing and controlled geometric repetition rather than arbitrary turbine placement [36], [54], [55].

A second recurring pattern in the literature is that large offshore development zones are frequently translated into project-scale blocks rather than being implemented as single immediate full-zone buildouts.

The Hornsea documentation presents Hornsea Project One as a defined project within the broader Hornsea Zone, while the NYSERDA planning study explicitly creates repeated 800MW site areas as the basis for subsequent layout development [36], [58].

The same NYSERDA review also compiles operational projects in the several-hundred-megawatt range, including Meerwind (288 MW), Nordsee Ost (295.2 MW), Borkum Riffgrund 1 (308 MW), Nordsee 1 (332.1 MW), Galloper (336 MW), Merkur Offshore (396 MW), BARD Offshore 1 (400 MW), and Horns Rev 3 (406.7 MW), while also listing substantially larger projects such as Dogger Bank Teesside B (1200 MW) and Hornsea 3 (2400 MW) in the planning pipeline [36]. This indicates that offshore wind development spans a broad range of project scales, but that several-hundred-megawatt projects are recurrent, particularly in operational benchmark datasets [36].

Accordingly, the literature supports describing a reference offshore wind farm layout as a project-scale, regular or lightly staggered, row-based array whose spacing is defined in rotor-diameter terms and whose final geometry reflects a balance between aerodynamic performance and non-aerodynamic engineering constraints.

For floating offshore wind, this bottom-fixed-derived geometric baseline remains especially useful because it provides a well-documented starting point onto which floating-specific constraints, such as mooring footprints, exclusion zones, and dynamic cable clearances, can subsequently be superimposed [37], [55], [57].

2.2. Wave Energy Conversion Technology

2.2.1. Classification and Operation Principles

In stark contrast to the offshore wind industry, which has rapidly converged on a dominant design configuration, the wave energy sector remains characterized by a high degree of technological diversity and a lack of a single standard architecture [10]. This divergence stems fundamentally from the complexity and variability of the wave resource itself, which differs significantly in terms of water depth, wave direction, and frequency characteristics across different deployment sites [60].

Unlike wind, which is a unidirectional flow of air, ocean waves involve complex orbital motions of water particles that decay with depth, necessitating a wide variety of engineering approaches to efficiently capture this energy [10]. Consequently, wave energy converters (WECs) are generally classified into three main categories based on their working principle, presented in Figure 15, and in Figure 16 with more detail, Oscillating Water Columns, Overtopping Devices, and Oscillating Bodies [10].

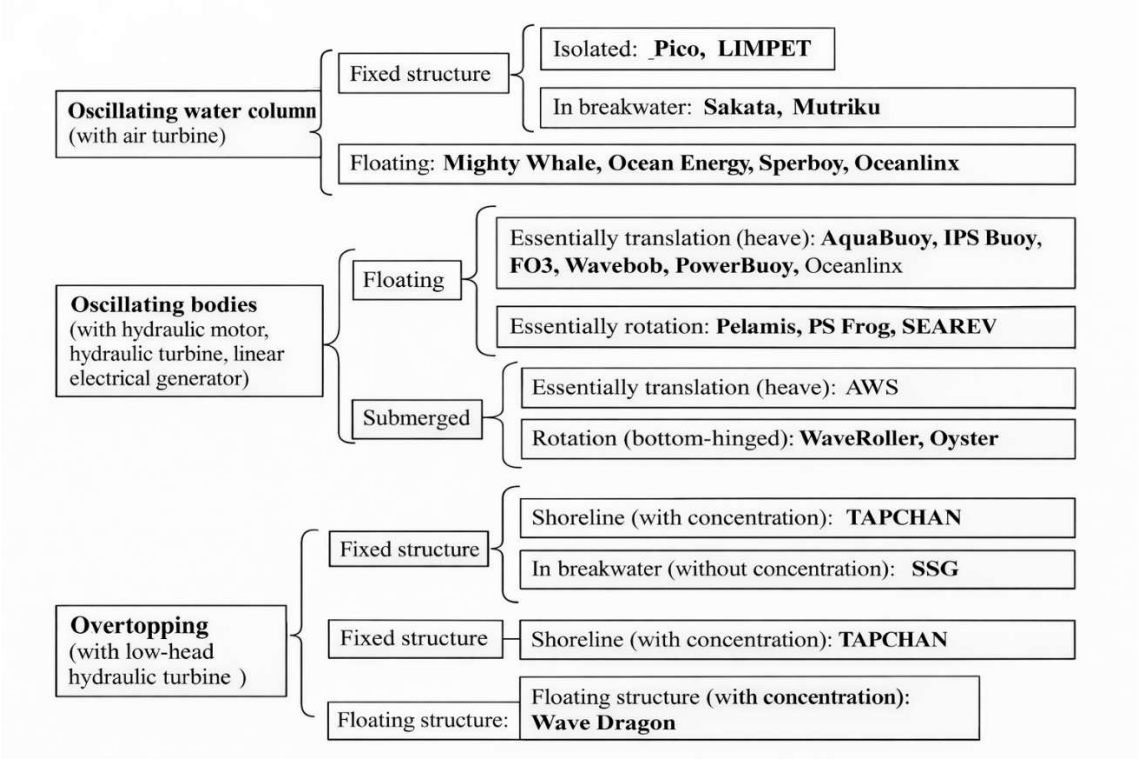


Figure 15 – The various wave energy technologies [10].

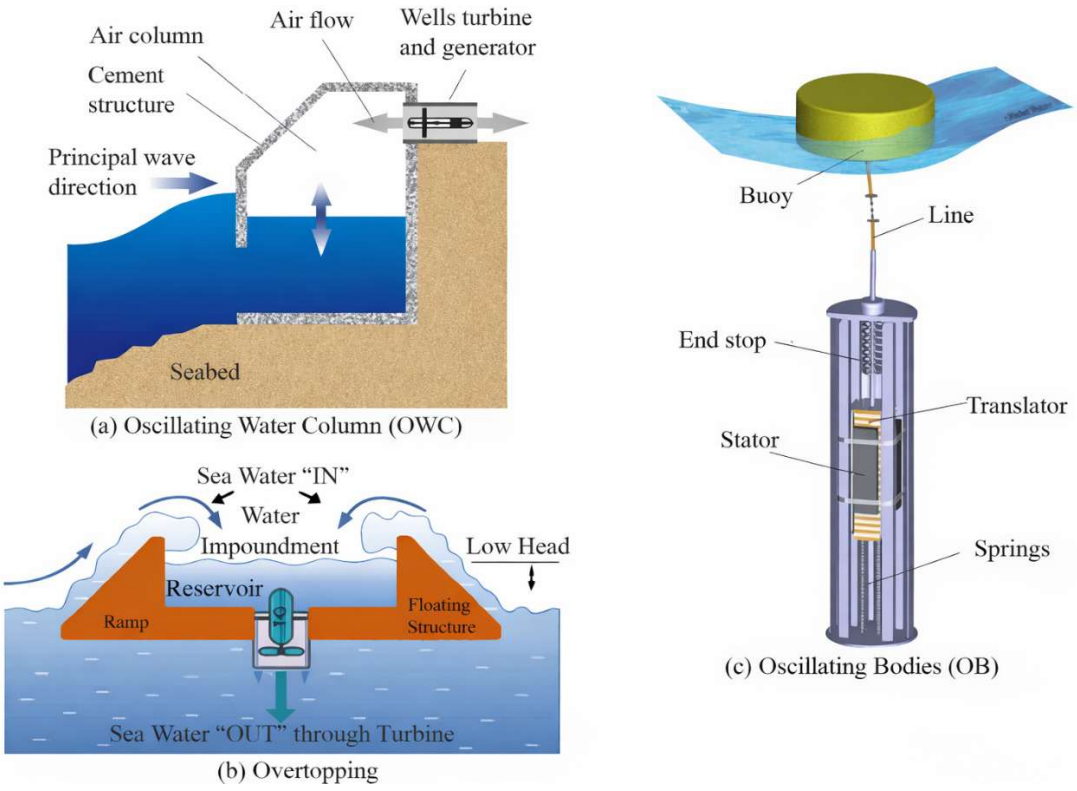


Figure 16 – Examples of wave energy converters: (a) Oscillating Water Column; (b) Overtopping; (c) Oscillating Bodies [61].

The oscillating water column (OWC) is the most extensively studied concept and typically consists of a partially submerged concrete or steel chamber with an opening to the sea below the waterline [10]. As incident waves enter and exit the chamber, the internal water column oscillates, acting like a piston that compresses and decompresses the air trapped above it [10]. This moving air is forced through a turbine, typically a self-rectifying well turbine or an impulse turbine, which drives a generator to produce electricity [10]. The primary advantage of the OWC concept lies in its mechanical simplicity, with the rotor of the turbine being the only moving part, which is located above the water surface, thereby removing complex mechanical systems from the corrosive marine environment [10]. However, their overall efficiency is often constrained by the thermodynamic losses inherent in the air compression process and the aerodynamic limitations of the self-rectifying turbine, which must accept airflow in both directions [10].

Overtopping devices function on a principle similar to a hydroelectric dam, by capturing water from incident waves into a reservoir located above the surrounding sea level [10]. These devices utilize a ramp or a tapered channel to elevate the waves, allowing them to spill over into the storage basin. The potential energy stored in the reservoir is then converted into electricity using conventional low-head hydro turbines as the water is released back into the sea [10]. While these systems are technologically robust and utilize mature turbine components, they generally require large physical structures to capture a sufficient volume of water, making them capital-intensive and less suitable for deep-water offshore arrays where structural mass is a primary cost driver [10].

The third and most diverse category, the oscillating bodies, encompasses devices that extract energy from the wave-induced motion of a floating or submerged structure relative to a reference point. This category includes attenuators, which are long, multi-segment floating structures aligned parallel to the wave direction, exemplified by the Pelamis device, where energy is captured from the relative motion between the hinged segments [10]. More significantly for the context of this thesis, this category includes point absorbers.

Point absorbers are defined as devices with horizontal dimensions that are small relative to the incident wavelength, enabling them to capture energy from a wavefront that is wider than their own physical dimensions [10]. For the co-located arrays proposed in this research, point absorbers are considered the most viable technological option due to their compact size, high energy density, and ability to absorb energy from any wave direction, which simplifies the layout optimisation process in complex offshore environments [62].

2.2.2. Physics of Point Absorbers

A point absorber is defined in the wave energy literature as a wave energy converter whose characteristic horizontal dimensions are small compared to the wavelength of the incident sea state, so that the device behaves as a compact oscillator interacting with the wave field over a limited physical footprint [10]. In contrast to attenuators or terminators, which are elongated to intercept a wave front over a larger span, point absorbers are typically designed to operate as oscillating bodies, most commonly exploiting heave motion, although wave energy converters may also extract power from surge and pitch, or a combination of modes, depending on the concept and its power take-off (PTO) architecture [63]. A key practical advantage of the point absorber approach is that energy can be harvested without requiring complex directional alignment systems, which makes these devices particularly attractive for offshore environments where wave directionality can be variable [10]. Figure 17 presents a detailed point absorber WEC.

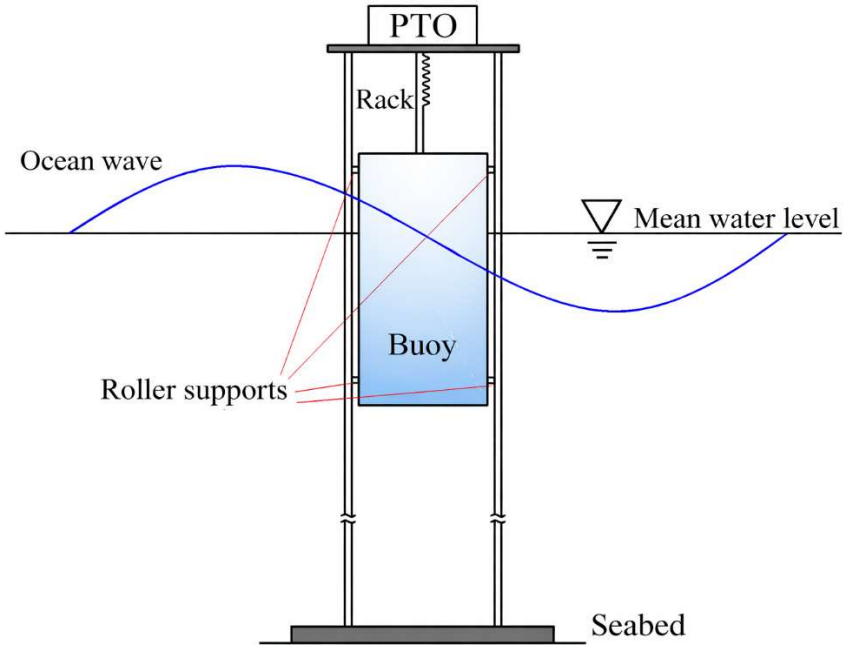


Figure 17 - Example of a point absorber [64].

From a modelling perspective, and as a first-order and simplified approximation, the heave dynamics of a point absorber can be described using a linear equation of motion analogous to a single-degree-of-freedom mass spring damper system, in which the device response results from the balance between incident-wave excitation, hydrostatic restoring, radiation forces associated with wave generation by the moving body, and the PTO damping that controls energy extraction [63]. This linearised representation is widely used in early-stage design and optimisation because it captures the dominant physics near resonance and allows parametric assessment of performance under regular and irregular wave conditions, even though additional effects may become important in realistic operation [63], [65].

A central concept in point absorber theory is that the absorbed power depends strongly on dynamic tuning and PTO control. In simplified linear theory, high capture performance is associated with operating near resonance and selecting an appropriate PTO damping level that balances motion amplitude and absorbed power [10], [65]. However, in real sea states, the incident wave spectrum spans a range of frequencies, therefore, practical devices must balance tuning for dominant conditions with maintaining acceptable performance across variable wave climates, while also respecting survivability constraints [10], [63]. This is particularly relevant offshore, where performance in irregular waves and the sensitivity of output to sea-state variability are key considerations in realistic power assessment [63].

Energy absorption by a point absorber is intrinsically linked to wave radiation. As the buoy oscillates, it generates outgoing wave components that propagate away from the device. These waves are inherently omnidirectional, in contrast to the predominantly downstream character of wind turbine wakes [66]. The radiated wave field interacts with the incident waves and with the wave fields generated by neighbouring converters, producing patterns of constructive and destructive interference. These interaction phenomena, commonly termed array effects, mean that the performance of a wave farm cannot be inferred by simply multiplying the single-device power by the number of converters [66].

This hydrodynamic coupling provides a direct physical motivation for array layout optimisation, and depending on spacing and geometric arrangement, interactions may increase the absorbed power of some devices through constructive effects or reduce overall farm output through destructive interference and sheltering effects [66].

Consequently, the spatial configuration of a WEC array becomes a primary design variable. *Child and Venugopal* [67] explicitly demonstrate that the array power output is strongly dependent on device placement, and that optimisation of configuration can improve overall performance compared to non-optimised arrangements, reinforcing that point absorber farms should be treated as a coupled hydrodynamic system rather than a collection of independent units [67]. This is consistent with modern design approaches, in which layout optimisation and device parameter selection are often combined to maximise energy yield while controlling interaction effects and satisfying deployment constraints [65], [66].

In practical wave energy assessment and numerical modelling, the performance of wave energy converters is commonly represented through a power matrix, which expresses the electrical power output of the device as a function of sea-state parameters, typically the significant wave height (H_s) and a characteristic wave period such as the peak period (T_p) or the energy period (T_e). This representation provides a convenient way to translate the hydrodynamic response of a device into energy production estimates by linking device performance to the statistical distribution of sea states at a given site. In many resource assessments, the energy period is derived from the wave spectrum through spectral moments. In particular, the energy period is defined as:

$$T_e = \frac{m_{-1}}{m_0}; \quad Te \approx 0.9\, T_p \tag{1}$$

where m_{-1} denotes the spectral moment of order -1 and m_0 the spectral moment of order 0. When detailed spectral information is not available, practical applications may approximate the energy period from the peak period using empirical relationships based on typical wave spectra. Power matrices are therefore widely used in wave resource assessments, techno-economic analyses, and array-level simulations to compute expected power production under realistic wave climates [68], [69].

In the context of wave farm modelling, power matrices also provide a practical interface between device-level hydrodynamic behaviour and numerical wave models, allowing WEC performance to be incorporated into large-scale simulations without explicitly resolving the internal dynamics of each converter. For example, numerical frameworks based on SWAN can represent WEC arrays by linking device power absorption to a power-matrix description of the converter performance, thereby coupling wave propagation with energy extraction processes [68]. An example of a representative power matrix for a point absorber device is presented in the following section, where the specific WEC technology considered in this work is introduced.

2.2.3. Target Technology: CorPower C4 and CorPack

For the specific purpose of the optimisation modelling conducted in this thesis, the CorPower Ocean C4 has been selected as the reference WEC. This device represents the current state of the art in point absorber technology, having been rigorously tested and validated in full-scale offshore deployments, including campaigns in the Portuguese waters offshore Aguçadoura [11]. The C4 is a heaving buoy point absorber with a diameter of 9 m and a dry mass of approximately 60 tonnes, constructed from a lightweight glass fibre reinforced composite hull to minimize inertial mass and material costs [11]. The device is nominally rated at 300 kW, consistent with manufacturer specifications [11], whereas experimental and modelling studies report maximum electrical outputs of approximately 400 kW under favourable sea states [69]. Furthermore, short-duration instantaneous power peaks exceeding 600 kW

have been recorded during offshore trials on the Portuguese coast, highlighting the strong dependence of power production on wave conditions and operational tuning [70].

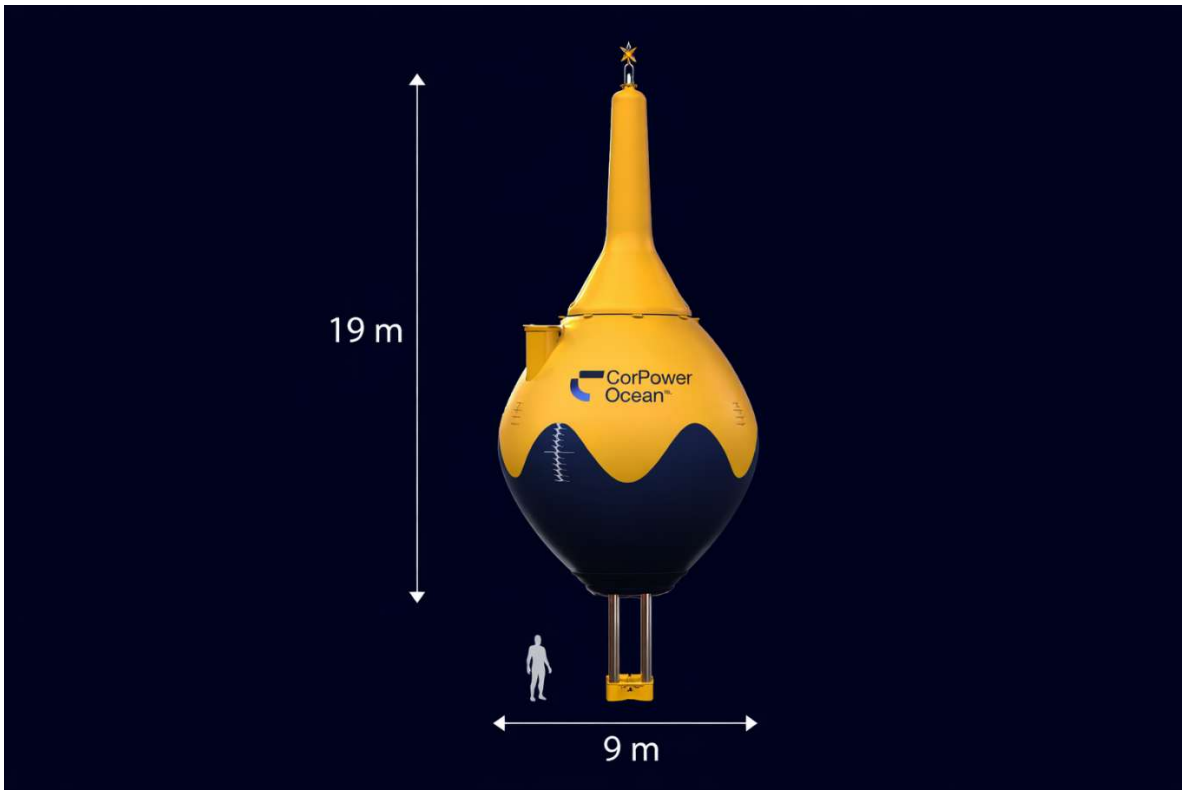


Figure 18 - CorPower C4 overview [11].

The defining technological advancement of the C4 system is its novel wave spring phase control technology, which addresses the fundamental bandwidth limitations of traditional passive point absorbers [11]. The wave spring module functions as a pneumatic negative spring, providing a restoring force that counters the natural hydrostatic stiffness of the buoy [65]. By actively controlling this negative stiffness, the system effectively reduces the resonant frequency of the device, allowing it to be tuned to match the variable period of incident ocean waves [11]. This resonance tuning amplifies the heave motion of the buoy compared to a conventional passive floater and increases the effective capture performance across broader range of operating conditions [11].

Equally critical for the feasibility of large-scale offshore deployment is the device's survivability strategy. Under extreme storm conditions, the wave spring system is deactivated or shifted into a survival mode, which detunes the device from the wave frequency [11]. In this detuned state, the buoy becomes effectively "transparent" to long-period storm waves, moving in a compliant phase with the water surface rather than resisting against it [11]. This operating mode drastically reduces the structural loads on the hull, and the peak tension forces on the mooring system, allowing the device to withstand wave heights of up to 18.5 m without requiring the heavy, over-engineered steel structures typical of earlier WEC generations [11].

As introduced in the previous section, the sea-state-dependent power absorption of a WEC is commonly represented through a power matrix, where each bin of H_s and characteristic wave period (typically T_p or T_e) is associated with a corresponding electrical power output. In practice, this matrix can be interpreted as a lookup surface. For a given sea state (H_s, T_p), the device power is read from the relevant

bin and subsequently aggregated across a sea-state dataset to estimate mean power, energy yield, or capacity factor. An example of a CorPower C4 power matrix is shown in Figure 19, where the output is expressed in kW.

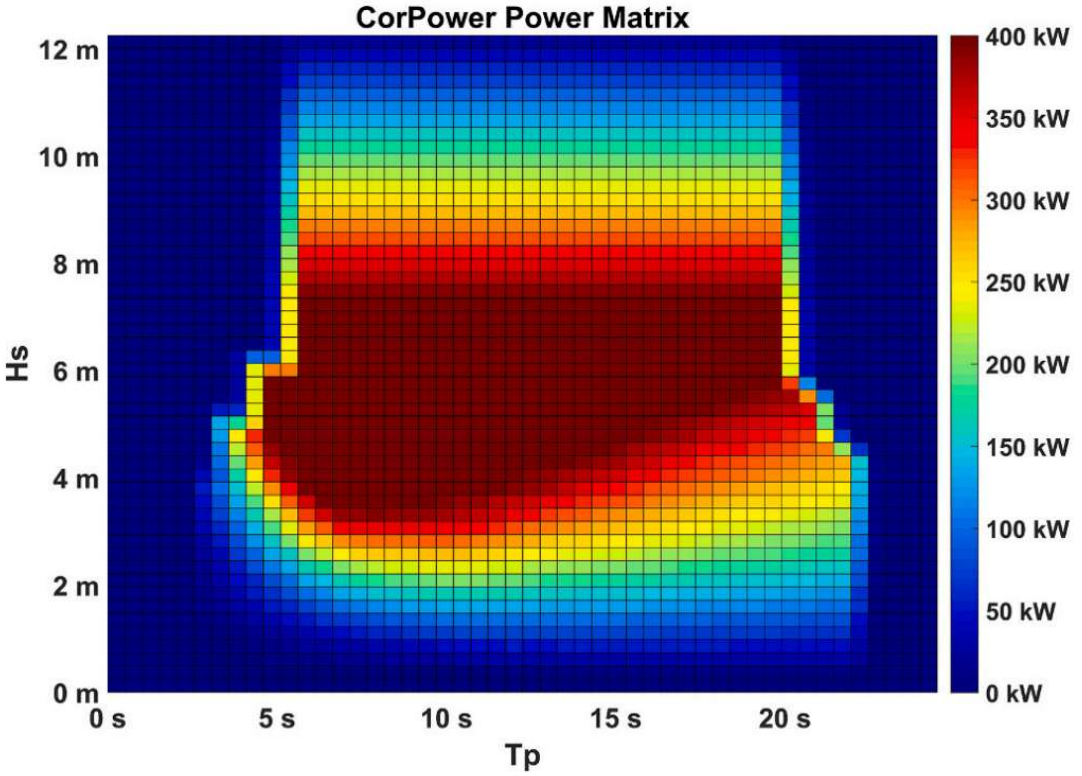


Figure 19 - Power matrix for the CorPower C4 device [69].

Although the exact numerical values of a power matrix may vary slightly between studies due to modelling assumptions, control strategies, or efficiency factors, matrices corresponding to the same device technology generally present similar performance patterns and operating regions. For the CorPower device considered here, the matrix indicates that the converter reaches its highest performance for intermediate to long wave periods, approximately between 6 s and 19 s, and for moderate to high significant wave heights, typically between 3 m and 7.5 m. Within this region, the device approaches its rated output of approximately 400 kW, forming a broad high-performance region in the central portion of the matrix.

Outside this operational envelope, power production decreases significantly. At very low wave heights the available wave energy is insufficient to drive large buoy motions, while extremely long periods or very high sea states correspond to operating regimes where the device control system limits energy absorption in order to maintain structural safety and survivability.

To improve the practicality of farm-scale deployment, CorPower has developed a modular array concept known as the CorPack, which is illustrated in Figure 20 [71]. Rather than treating each WEC as a fully independent unit, this approach groups multiple devices into a predefined cluster designed to operate as a single repeatable deployment entity [71]. The concept is intended to simplify offshore installation, reduce the complexity of electrical collection and export, and provide a more integrated basis for large-scale wave farm development [71]. In this way, the CorPack translates individual device technology into a configuration that is more compatible with commercial offshore implementation [71].

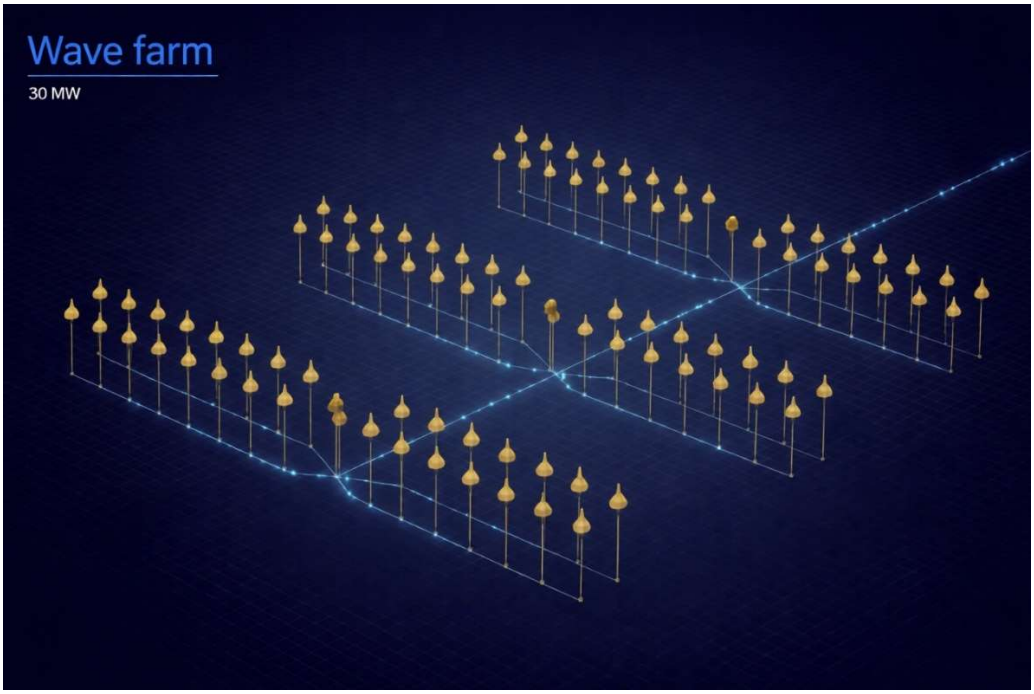


Figure 20 - CorPack array overview [71].

A CorPack constitutes the fundamental building block of a wave farm, consisting of a predefined cluster of WECs with a combined capacity typically ranging from 10 MW to 20 MW [71]. This clustering approach integrates multiple devices within a shared anchoring and mooring configuration, with an inter-device spacing of approximately 150 m, while also incorporating a consolidated electrical collection hub that exports power through a single cable [71]. In practical terms, this arrangement is intended to balance hydrodynamic interaction effects, mooring feasibility, and electrical integration within a standardised offshore deployment unit [71].

Although an individual CorPower point absorber is designed to capture wave energy without requiring directional alignment, the farm-scale module is not orientation-neutral. In fact, the CorPack concept is an elongated cluster whose main axis is positioned perpendicular to the predominant wave direction, in order to maximise exposure to the incoming waves and to manage hydrodynamic interactions within the cluster [69]. In addition, the internal layout is conceived as a staggered configuration, in which the devices in the rear row are laterally offset relative to those in the front row, so that each downstream device is positioned between two upstream units [69]. When this elongated module is oriented perpendicular to the incident wave direction, the sheltered region immediately behind the front-row devices does not fall directly onto the devices in the rear row. This reduces the direct propagation of local shadowing effects through the cluster and helps limit excessive hydrodynamic interference while preserving a compact deployment footprint [67], [69].

This arrangement is consistent with established array hydrodynamics, whereby the incident-wave direction strongly influences constructive and destructive interference patterns between neighbouring devices, making array orientation a relevant design variable even for omnidirectional point absorbers [66], [67]. The CorPack therefore represents a practical engineering compromise, with devices placed sufficiently close to enable shared anchoring and consolidated electrical export, while remaining separated enough to avoid severe hydrodynamic shadowing and excessive mutual interference [71].

From a modelling perspective, this concept allows the wave farm to be represented as a set of discrete operational units rather than as a collection of isolated devices, which is advantageous for layout

optimisation while preserving the essential physics of array interaction and simplifying installation, operation, and maintenance logistics [71]. Because the analysis is performed at the farm scale rather than at the level of internal device arrangement, the CorPack is treated as the deployable optimisation unit, while the number of C4 devices is used only to express the equivalent installed capacity represented by each cluster.

2.3. CO-LOCATED WAVE-WIND ENERGY SYSTEMS

2.3.1. CONCEPT DEFINITION: CO-LOCATED VS HYBRID

The integration of offshore wind and wave energy technologies represents a strategic frontier in marine renewable energy, aiming to maximise the use of the marine space and improve the economic viability of offshore projects [72]. In the scientific literature, combined systems are generally categorized into two distinct technological families: hybrid systems and co-located systems [72], as presented in Figure 21. Hybrid systems involve a high degree of physical and structural integration, where the WEC is mounted directly onto the substructure of the wind turbine, effectively sharing the same foundation and stability mechanism [72]. While this concept offers the theoretical advantage of a single support structure, it introduces significant engineering complexity regarding structural coupling, increased hydrodynamic loading on the wind turbine platform, and the logistical challenge of aligning the maintenance schedules of two disparate technologies with different maturity levels [72].

In contrast, co-located systems, which constitute the focus of this thesis, are defined as an independent wind and wave arrays that are deployed within the same leased maritime area [72]. In this configuration, FOWTs and the WECs are structurally separate entities, each possessing their own dedicated mooring system [73]. The integration in a co-located system occurs at the array level rather than the device level, primarily through the sharing of the electrical infrastructure and the joint utilization of the sea space [73]. This decoupled approach allows for the use of mature, off-the-shelf technologies for both wind and wave components without the need to develop custom, unproven hybrid platforms [32]. Furthermore, co-location offers significantly greater flexibility in layout design, allowing the WEC array to be positioned optimally to exploit hydrodynamic interactions without being constrained by the rigid geometry of the wind turbine array [74].

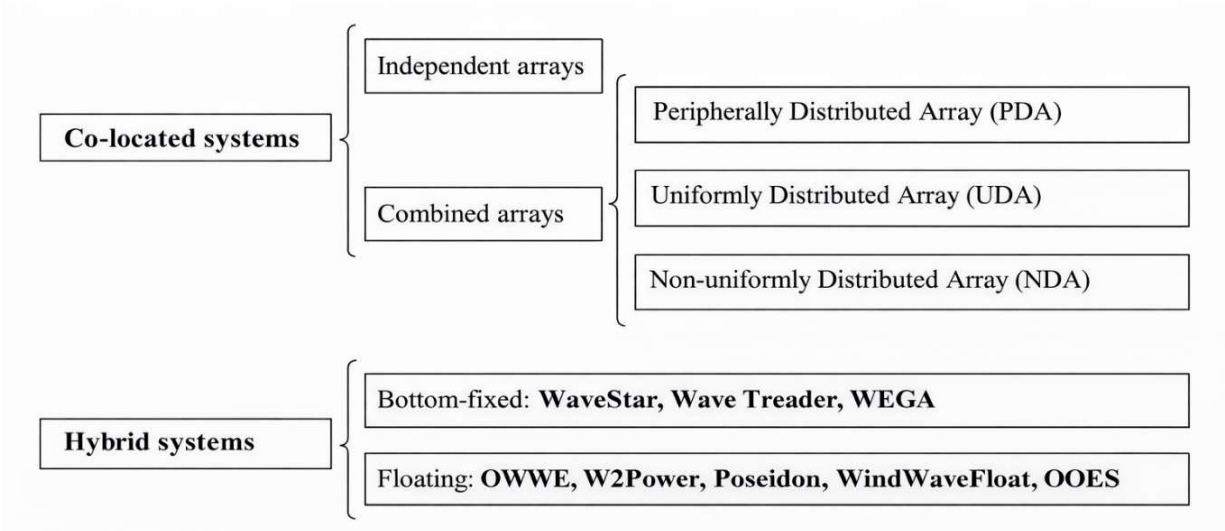


Figure 21 - Classification of combined wave-wind technologies [72].

2.3.2. Economic Synergies and Levelized Cost of Energy

The primary economic driver for the development of co-located farms is the potential for significant cost reduction, measured by the levelized cost of energy (LCoE) [75]. Offshore energy projects are inherently capital-intensive, with grid connection and infrastructure accounting for a substantial portion of the total CAPEX [41].

In a co-located farm, the wind and wave arrays share the same offshore substation, export cable, and onshore grid connection point [41]. Since the peak production of wind and wave resources rarely occurs simultaneously, due to the time lag between the wind generation event and the subsequent development of the wave field, the combined power output is smoother and more consistent than that of a standalone wind farm [76]. This enables a higher capacity factor (CF) for the shared transmission infrastructure, thereby reducing the per-unit transmission cost [41]. The CF provides an indication of how effectively the installed capacity is utilized under real sea conditions, and is computed as:

$$CF = \frac{P_{abs}}{P_{rated}} \tag{2}$$

where P_{abs} is the absorbed power of the device (or array) and P_{rated} its rated power.

Analytical cost models developed for co-located floating arrays demonstrate that these shared infrastructure costs can be reduced by up to 30% compared to separate developments [41]. Additionally, co-location dilutes the soft costs of development, such as site surveys, environmental impact assessments, and permitting fees, over a larger total installed capacity [73]. From a resource perspective, co-location significantly increases the energy yield per square kilometre (km^2) of the leased ocean area [73]. This is particularly relevant in the context of marine spatial planning, where high-value sea space is scarce [7]. By filling the interstitial spaces between widely spaced wind turbines with compact WEC arrays, the overall energy density of the concession zone is maximised without expanding the project's physical footprint [73].

These cost-reduction mechanisms are quantitatively illustrated in the analytical case study presented by *Clark et al. (2019)* [41] in Figure 22.

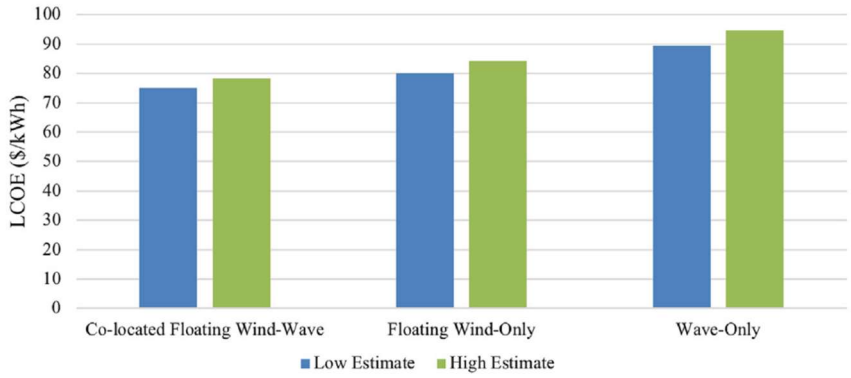


Figure 22 - Cost model comparison of array types [41].

This study evaluates a co-located wind-wave farm in the North Sea comprising 80 Vestas V-80 (2.0MW) wind turbines and 26 WaveCat WECs. The study assumes idealised unidirectional wind and wave conditions, along with 90% turbine efficiency and 95% turbine availability, for both turbines and WECs.

Within this framework, wind and wave devices are connected to a common offshore substation and export system, allowing both technologies to exploit the same electrical and grid-connection assets [41].

Despite the simplified environmental assumptions, the results suggest that co-located floating wind-wave arrays may achieve lower LCoE than equivalent floating wind-only or wave-only configurations, primarily due to reductions in shared transmission infrastructure and balance-of-plant costs [41]. This comparative analysis reinforces the economic rationale for co-location, highlighting the potential of integrated offshore energy systems to improve asset utilisation and enhance overall project viability.

2.3.3. Hydrodynamic Interactions and Operational Synergies in Co-Located Farms

Beyond the economic benefits associated with shared infrastructure, co-located wave-wind systems exhibit a range of physical and operational interactions arising from the modification of the local wave field by the WECs. Rather than acting solely as external barriers, point absorber arrays deployed within offshore wind farms fundamentally alter the incident wave climate through diffraction and radiation processes, giving rise to complex patterns of constructive and destructive interference [66], [77]. These hydrodynamic interactions influence both the performance of neighbouring WECs, and the environmental conditions experienced by nearby wind turbines, making array layout a critical design variable.

In contrast to co-location concepts based on peripheral wave farms acting as external breakwaters, the present work focuses explicitly on internal co-location, whereby WECs are deployed within the wind farm footprint, occupying the inter-turbine spaces between floating wind units. This configuration is motivated by the objective of maximising energy extraction per unit area while preserving the original wind farm boundaries, and it places primary emphasis on optimising WEC placement for wave energy capture rather than on deliberate wave shielding of the wind turbines.

As individual WECs extract energy from the incoming wave field, part of the wave energy is converted into electricity. At the same time, the remainder is redistributed through radiated waves generated by device motion. When multiple converters are deployed in proximity, their radiated and diffracted wave fields overlap, producing interaction effects that can either enhance or degrade the overall power capture of the array, depending on device spacing and geometric configuration [66], [78]. Consequently, the total energy yield of a wave farm is not simply the sum of isolated device outputs, but a coupled response governed by layout-dependent hydrodynamics. This behaviour provides a strong physical motivation for layout optimisation, as demonstrated by recent multi-objective studies showing that appropriate spatial arrangements can significantly improve farm-level performance while respecting practical deployment constraints [53], [74].

Within co-located wind-wave farms, these same interaction mechanisms also partially attenuate wave energy within the array. While this effect is less pronounced than in perimeter wave farms designed explicitly for coastal protection, it nevertheless results in a measurable reduction of significant wave height in regions occupied by wind turbines [77]. Importantly, in the context of internal co-location, where WECs are placed between turbines rather than forming an external breakwater, this wave attenuation is a by-product of optimised energy extraction rather than a primary design objective.

One operational consequence of this local wave attenuation is improved accessibility for offshore operations and maintenance (O&M). Access to floating wind turbines is strongly constrained by sea-state limits, with conventional crew transfer vessels typically restricted to significant wave heights below approximately 1.5 m [79]. In energetic offshore environments, such limits frequently result in weather-induced downtime, during which maintenance interventions are delayed, and turbines may remain

offline. By reducing local wave heights within the farm, co-located WEC arrays increase the frequency of weather windows suitable for vessel access. Accessibility studies indicate that this effect can lead to increases in operational access of more than 15% in high-energy sea states [79]. This improvement results in shorter repair delays, reduced turbine downtime, and, consequently, higher realised availability of wind turbines, thereby increasing annual energy production.

In addition to operational benefits, wave attenuation also reduces hydrodynamic loading on floating wind substructures and mooring systems. Lower wave-induced motions can decrease fatigue accumulation in critical structural components, potentially extending service life or enabling lighter platform designs [80]. However, within the scope of this work, these effects are considered secondary outcomes of array hydrodynamics rather than primary optimisation targets.

2.4. Optimisation of Array Layouts Using Genetic Algorithms

2.4.1. The Array Optimisation Problem

The design of a co-located array layout constitutes a complex engineering challenge that fundamentally seeks to determine the optimal spatial coordinates for each WEC and FOWT, within a defined maritime area [78]. Unlike conventional wind farm design, where wake effects are primarily directional and downstream, hydrodynamic interactions in a wave farm are omnidirectional and coupled, the motion of each WEC generates radiated waves that influence the power absorption of every other device in the array [67]. Consequently, the total power output of the array is not simply the sum of individual isolated outputs [81].

This collective behaviour is commonly quantified by a interaction factor, the q-factor (q). This parameter evaluates whether the presence of neighbouring devices enhances or reduces the total energy absorption of the array when compared with the performance of isolated devices. Therefore, it provides a concise measure of how device spacing and geometric configuration govern array-level energy extraction through radiation and diffraction effects [81].

It is defined as the ratio of the total power absorbed by the WEC array to the sum of the power each device would produce when operating independently under the same sea state:

$$q = \frac{P_{array}}{\sum P_{WEC}} \tag{3}$$

where P_{array} represents the total power absorbed by the array and $\sum P_{WEC}$ corresponds to the sum of the power output of each individual device considered in isolation.

The values of q greater than unity indicate constructive hydrodynamic interference between radiated and incident waves, whereby array interactions enhance overall energy capture, whereas values below unity reflect destructive interference and reduced collective performance. A value equal to the unity represents the neutral case in which the array behaves as a simple aggregation of independent devices [62].

The magnitude of the q-factor depends strongly on array geometry, device spacing, incident wave direction, and the hydrodynamic characteristics of the WEC concept. As a result, it is frequently used as a key performance indicator in layout optimisation studies, allowing different configurations to be compared in terms of their ability to enhance or mitigate interaction effects within the farm.

From a mathematical standpoint, the layout optimisation problem is formulated as a nonlinear, non-convex optimisation problem with a continuous search space [81]. The associated objective function is typically multi-modal, exhibiting a rugged landscape with numerous local optima. Under such conditions, conventional gradient-based optimisation methods tend to become trapped in suboptimal solutions and are unable to reliably identify global optima [67].

In addition, the problem is subject to multiple physical and operational constraints [53]. Candidate layouts must satisfy minimum inter-device spacing requirements to avoid collisions and ensure navigational safety, while also respecting exclusion zones imposed by the mooring footprints of floating wind turbines. These constraints significantly restrict the feasible design space and further increase optimisation complexity [53].

As the number of devices in the array increases, the computational complexity of evaluating the hydrodynamic interactions grows exponentially, or higher, depending on the solver, making a brute force search of all possible coordinate combinations computationally impossible [78]. To navigate this vast and complex search space effectively, heuristic or meta-heuristic approaches are required. Among these, genetic algorithms have emerged as the standard optimisation tool in the wave energy sector due to their robustness in handling continuous variables and their ability to escape local optima through stochastic processes [81].

2.4.2. Genetic Algorithms Methodology

Genetic algorithms (GAs) are a class of stochastic search heuristics within the broader field of evolutionary computing, inspired by the biological principles of natural selection and Darwinian evolution. The fundamental premise of a GA is to evolve a population of candidate solutions towards an optimal, or near-optimal, solution by mimicking the survival of the fittest mechanism observed in nature [82].

In the specific context of array layout optimisation, each candidate solution is represented by a set of vectors containing the spatial coordinates of every WEC deployed within the farm [81]. The algorithm is initialized with a randomly generated population of individuals, while ensuring that all layouts satisfy the predefined spatial constraints and remain within the designated maritime area.

The evolutionary cycle proceeds through an iterative sequence of evaluation, selection, and reproduction [82]. First, the performance of each individual is assessed through a fitness function, implemented here as a numerical model that computes both the energy production and the wave-shielding performance associated with a given layout configuration [83]. Based on the resulting fitness scores, a selection operator identifies the best-performing individuals to serve as parents for the next generation [82].

Figure 23 schematically summarizes the evolutionary cycle adopted in this study, highlighting the flow from population initialization to parent selection, genetic recombination and mutation, and subsequent survivor selection.

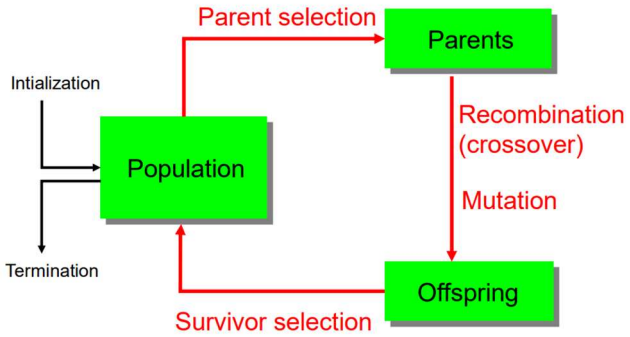


Figure 23 - The main evolutionary cycle [82].

Once the parents are selected, the reproduction phase employs two primary genetic operators, the crossover and mutation [82]. Crossover involves combining the genetic information of two parents to produce offspring, effectively swapping coordinate sets between layouts to explore new areas of the solution space [81]. Mutation, conversely, introduces small, random perturbations to the coordinates of the offspring [82]. This stochastic element is critical for maintaining genetic diversity within the population, preventing the algorithm from converging prematurely on a local optimum [82]. The cycle repeats for a pre-defined number of generations or until a convergence criterion is met, with the expectation that the average fitness of the population will improve over time, eventually yielding a highly optimised array configuration [83].

2.4.3. Multi-Objective Fitness Function: The WLA Index

In the design of co-located offshore energy parks, the optimisation problem is inherently multi-objective, requiring a balance between two conflicting goals: maximizing the electrical energy production of the WEC array and maximizing the wave shielding protection provided to the wind farm [74]. To enable the genetic algorithm to navigate this trade-off, a composite fitness function known as the wave energy park layout assessment (WLA) index has been developed by *Teixeira-Duarte et al.* [53]. The WLA index integrates the hydrodynamic efficiency of the array and its wave attenuation performance into a single scalar value, allowing for the ranking and selection of candidate layouts during the evolutionary search [53].

The WLA is defined as the weighted sum of two normalised parameters, the q-factor and the height reduction area index (HRA) [53]. The q-factor is a dimensionless measure of the array's hydrodynamic interference, explained in Section 2.4.1, while the HRA index quantifies the effectiveness of the array in reducing the significant wave height in the lee of the farm, within a defined area of interest located in the lee of the WEC array and overlapping the wind-farm region of interest [74]. Specifically, it measures the surface area of the shadow zone within the wind farm where the wave height is reduced below a specific threshold required for safe maintenance access [74].

The HRA index is computed by comparing the baseline H_s at each calculation node within the area of interest, without WECs, to the H_s obtained when the WEC array is present, and averaging the relative reduction over all nodes:

$$HRA = \frac{100}{n}\sum_{i=1}^{n}\frac{H_{s,i} - H_{s,WEC,i}}{H_{s,i}} \tag{4}$$

where i denotes a calculation node, n the total number of nodes, $H_{s,i}$ is the significant have weight at node i for the baseline scenario (without WEC), and $H_{s,WEC,i}$ is the significant wave height ate the same

node when the WEC array is present. To combine both competing objectives into a single fitness value, the WLA index is computed as:

$$WLA = pN(q) + sN(HRA) \tag{5}$$

where p and s are weighting factors used to define the relative importance assigned to power absorption (through q) versus shielding performance (through HRA), respectively, and the sum of them is equal to 1. As a result, different optimisations priorities can be represented through different weight choices. For example, assigning equal importance to power absorption and shielding corresponds to a value of 0.5 for each component, while increasing p (and therefore decreasing s) shifts the optimisation emphasis towards maximising array hydrodynamic efficiency at the expense of shielding performance, and vice versa. In the context of the present work, the optimisation procedure is intentionally focused on maximising the WEC array energy output. Consequently, the shielding term is deactivated, as discussed later.

The terms N(q) and N(HRA) denote the min-max normalized values of each metric, ensuring both indicators on a consistent scale between 0 and 1.

The min-max normalization applied to a generic metric x is given by:

$$N(x) = \frac{x - \min(x)}{\max(x) - \min(x)} \tag{6}$$

where min(x) and max(x) are the minimum and maximum values of that metric within the set of candidates' layouts being compared [53].

The evaluation of these complex hydrodynamic parameters for every candidate layout requires a robust numerical modelling tool [84]. For this purpose, the Sandia National Laboratories - Simulating Waves Nearshore (SNL-SWAN) model is widely utilised in the research community [85]. SNL-SWAN is a modified version of the spectral wave model Simulating Waves Nearshore (SWAN), developed by Sandia National Laboratories to explicitly account for the presence of WECs [85]. Unlike standard SWAN, which models wave propagation and dissipation, SNL-SWAN represents WECs as obstacles with a frequency-dependent transmission coefficient, allowing the model to simulate the energy absorption of the devices and the resulting modification of the wave field in the lee of the array [85]. This coupling of a genetic optimiser with the SNL-SWAN solver provides the computational engine necessary to identify layouts that deliver both high energy yields (high q) and effective wave shielding within the wind farm [85].

2.5. SUMMARY AND IDENTIFIED GAPS

This chapter has presented a comprehensive review of the state of the art in floating offshore wind, wave energy converters, and co-located offshore energy systems. The literature confirms that while floating wind technology is rapidly progressing toward commercial maturity, wave energy remains at a pre-commercial stage characterised by a wide diversity of device concepts and deployment strategies. Among these, point absorber technologies, particularly advanced systems such as CorPower, have emerged as promising candidates for co-location due to their compact footprint, modular deployment potential, and high energy density. The review has further demonstrated that co-located wind-wave

farms offer substantial economic and operational synergies. Multiple studies show that shared electrical infrastructure, export cabling, offshore substations, and development processes can significantly reduce overall project costs and levelized cost of energy. In parallel, hydrodynamic interactions within WEC arrays can lead to local wave attenuation, with potential secondary benefits for offshore accessibility and structural loading.

Beyond these general advantages, recent research has shown that a complex set of spatial and infrastructural factors constrains the design of offshore farms. Wake-induced spacing requirements primarily dictate wind turbine layouts, while electrical infrastructure introduces additional geometric constraints through inter-array cabling corridors, substation safety buffers, and export cable routes. In floating offshore wind, these challenges are compounded by the large seabed footprint of mooring systems, which create exclusion zones extending several hundred metres from each turbine and impose strict minimum separation distances between energy devices. Typical offshore wind farms, therefore, adopt regular grid-based layouts aligned with dominant wind directions, balancing aerodynamic performance against cabling cost and navigational safety.

Although these constraints and layout principles are increasingly well documented individually, a critical gap remains in their integrated application to co-located floating wind-wave systems. Most existing co-location studies rely on simplified geometrical arrangements, neglect realistic mooring footprints, or assume idealised spacing conditions. Likewise, many optimisation frameworks focus either on wave energy array performance or on wind farm design in isolation, without simultaneously accounting for wake effects, cable routing, exclusion zones, and internal WEC placement within an operational floating wind farm.

Furthermore, while high-fidelity numerical tools such as SNL-SWAN enable detailed modelling of wave-device interactions, they are rarely coupled with optimisation algorithms in a manner that reflects the practical spatial constraints of commercial-scale floating wind farms. In particular, there is limited research on internal co-location, in which WECs are deployed within the wind farm footprint between turbines rather than as peripheral wave barriers.

Consequently, a methodological gap exists between theoretical co-location concepts and deployable engineering solutions. There is a need for optimisation frameworks that integrate realistic wind farm layouts, floating-specific spatial constraints, wave energy device clustering, and site-specific environmental conditions into a unified design approach.

This thesis seeks to address this gap by developing a layout optimisation framework for internally co-located floating wind-wave farms that explicitly incorporates turbine spacing rules, mooring exclusion zones, infrastructure buffers, and environmental forcing within the regulatory context of the Portuguese “Plano de Afetação para as Energias Renováveis Portuguesas” (PAER) zones. By applying a genetic algorithm to optimise the placement of CorPack wave energy modules inside a representative floating wind farm, this work aims to provide practical insight into how co-located systems can be designed to maximise energy extraction while respecting real-world deployment constraints.

3. METHODOLOGY

3.1. OVERVIEW OF THE OPTIMISATION FRAMEWORK

The design of co-located offshore renewable energy farms requires the integration of multiple modelling components, including environmental characterisation, hydrodynamic simulation, and spatial optimisation techniques. In the present work, these components are combined within a unified computational framework aimed at determining efficient layouts of WECs within a predefined offshore wind farm configuration, for the Portuguese offshore conditions, integrating site-specific constraint layers (PAER boundaries, turbine exclusions) and validated CorPack representations. This integrated approach allows the simultaneous consideration of wave climate conditions, device interactions, and spatial deployment constraints within a single optimisation workflow, thereby enabling a realistic assessment of the performance of large-scale co-located energy systems. Unlike previous applications , this study incorporates Portuguese regulatory constraints and high-capacity 20 MW wind turbines, providing a more realistic co-location scenario.

The optimisation methodology adopted in this thesis couples spectral wave modelling with an evolutionary optimisation algorithm. Wave propagation and wave-structure interactions are simulated using the SWAN model, a third-generation spectral wave model widely used for coastal and offshore wave prediction [86]. To represent the interaction between ocean waves and WECs, the simulations are performed using the SNL-SWAN extension, which enables the representation of WEC arrays as energy-extracting obstacles within the spectral wave model [84]. This modelling approach allows the numerical model to account for the modification of the wave field induced by the presence of WECs while simultaneously estimating the power absorbed by the devices.

The spatial optimisation component of the framework is based on a GA specifically developed by *Teixeira-Duarte et al.* [53] for the optimisation of WEC arrays and co-located wave-wind offshore energy systems. The algorithmic structure follows the optimisation framework in which evolutionary operators are combined with hydrodynamic simulations to evaluate the performance of candidate array configurations within a realistic wave propagation environment [53], [74]. Rather than relying on simplified interaction models, this approach evaluates each candidate layout through direct hydrodynamic simulations, allowing the optimisation process to account for wave transformation processes and device-array interactions. Consequently, the optimisation procedure can capture the complex hydrodynamic behaviour associated with large arrays of energy converters.

Within this modelling framework, the optimisation procedure consists of repeatedly generating candidate WEC array layouts, simulating their hydrodynamic response under representative sea-state conditions, and evaluating their performance according to a predefined objective metric. The evolutionary search process progressively improves the spatial configuration of the array through

successive generations of candidate solutions. Through this iterative process, the algorithm explores the spatial design space of possible array configurations while identifying layouts that maximise the performance of the wave energy system under the considered environmental conditions.

The overall structure of the modelling framework implemented in this work is illustrated in Figure 24.

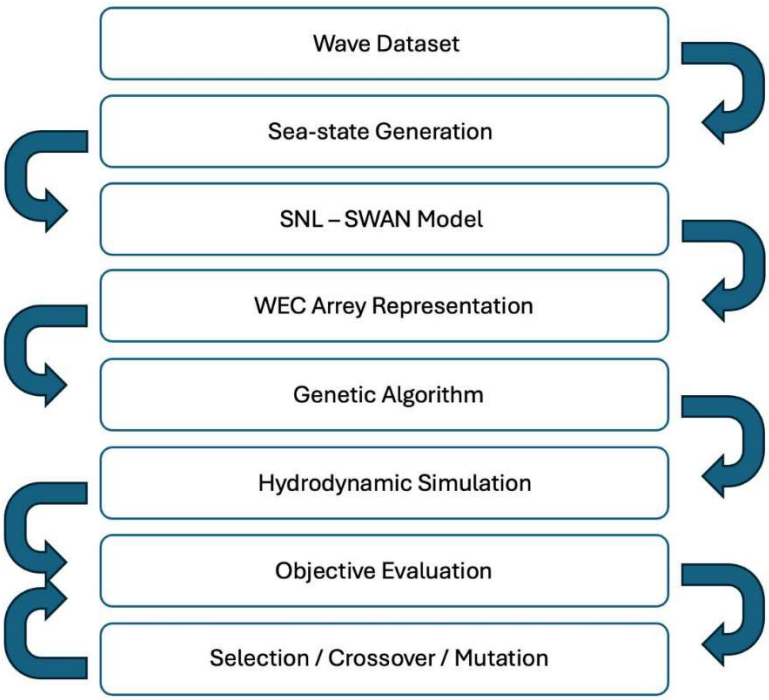


Figure 24 - Optimisation workflow.

As illustrated in Figure 24, the optimisation framework can be divided into four main stages. The first stage consists of the characterisation of the environmental forcing conditions through the derivation of representative sea states from long-term wave climate data. This step reduces the dimensionality of the environmental dataset while preserving the main statistical characteristics of the wave climate relevant for energy production.

The second stage involves the hydrodynamic modelling of wave propagation within the specific study domain. In this stage, the SNL-SWAN model is used to simulate the wave field while accounting for wave energy extraction by WEC arrays. To enable computationally efficient simulations of large arrays, an aggregated representation of the CorPower CorPack system is adopted. The methodology used to derive this aggregated representation is presented in Section 3.3.

The third stage corresponds to the spatial optimisation process. In this step, the GA generates candidate layouts of WEC arrays within the available deployment area while enforcing spatial constraints associated with device separation distances and exclusion zones around offshore wind turbines. Each candidate layout is evaluated through hydrodynamic simulations to quantify its performance under the selected sea-state conditions.

Finally, the fourth stage consists of the evolutionary optimisation loop. Based on the performance evaluation of each candidate configuration, the genetic algorithm applies selection, crossover, and mutation operators to generate new populations of layouts. Through successive generations, the

algorithm progressively explores the spatial design space and converges towards increasingly efficient WEC array configurations.

The following sections describe the individual components of this modelling framework in detail. Section 3.2 introduces the spatial definition of the study domain and the implementation of deployment constraints. Section 3.3 presents the aggregated representation of the CorPack WEC system adopted in the hydrodynamic model. Section 3.4 describes the generation of representative sea states used in the optimisation simulations. Section 3.5 details the hydrodynamic modelling setup, while Sections 3.6 to 3.9 present the optimisation methodology, including the objective function, GA implementation, and constraint handling strategy.

3.2. Definition of the Study Domain, Grid and Spatial Layers

The optimisation framework described in the previous section operates within a spatially defined modelling domain representing the marine area available for the deployment of WECs within the co-located wind-wave energy system considered in this study. To evaluate candidate layouts and simulate their hydrodynamic response, this domain must be represented in a form that is compatible with both the spectral wave model and the spatial optimisation algorithm. Consequently, the study area is described through a set of spatial layers that collectively define the feasible deployment region explored during the optimisation process.

The geographical context of the case study and the broader characteristics of the offshore wind concession are discussed in Chapter 4. The objective of the present section is therefore limited to the definition of the modelling domain as required by the optimisation framework, focusing on how the physical environment and operational constraints are represented within the computational workflow.

Hydrodynamic simulations are performed using the SWAN spectral wave model, which requires the study area to be discretised through a computational grid covering the region of interest. This numerical grid provides the spatial framework within which wave propagation processes and wave-structure interactions are evaluated. The construction of the grid, together with the associated bathymetric processing and interpolation procedures, is described in detail in Section 3.5. At this stage, it is sufficient to note that the grid allows the continuous offshore domain to be represented in a numerical form suitable for repeated simulations during the optimisation procedure.

Bathymetric information is incorporated into the modelling framework to account for depth-induced wave transformation processes. The seabed elevation field used in the simulations is obtained from the General Bathymetric Chart of the Oceans (GEBCO) dataset and subsequently interpolated onto the computational grid. Through this bathymetric representation, the wave model can capture the influence of seabed depth on wave propagation characteristics within the study domain. As with the grid definition, the specific processing steps used to construct the bathymetric input are presented later in Section 3.5.

In addition to the physical description of the marine environment, the optimisation framework requires the explicit definition of spatial constraints that restrict the locations where WEC units may be deployed. Offshore energy developments are typically subject to multiple classes of constraints, including regulatory boundaries, environmental restrictions, geotechnical limitations, and infrastructure protection zones. Within the context of this work, these restrictions are represented conceptually as a set of spatial layers that progressively reduce the area available for device deployment.

The outer boundary of the modelling domain corresponds to the region considered for energy development within the case study. This boundary defines the maximum spatial extent within which

candidate layouts may be generated. Rather than treating the domain as an unrestricted plane, the optimisation space is therefore explicitly bounded and subject to the operational limitations associated with the offshore energy park.

Within this boundary, engineering-related constraints are applied to account for the presence of floating offshore wind turbines. These turbines introduce safety buffers associated with mooring systems, dynamic cables, and operational clearance requirements. In the optimisation framework, such constraints are represented through fixed radial exclusion zones, rather than explicit dynamic cable modelling. Areas falling within these buffers are removed from the deployable region prior to layout generation, ensuring that candidate WEC placements do not interfere with wind-farm infrastructure.

Because of these spatial restrictions, the feasible deployment region becomes irregular and potentially fragmented. Such geometrical complexity is typical of real offshore planning environments and poses challenges for deterministic optimisation approaches. For this reason, heuristic search techniques, such as the GA employed in this work, are particularly well suited to exploring candidate solutions within constrained spatial domains.

To enable numerical optimisation, the admissible region defined by the spatial layers is used as the domain within which candidate WEC layouts are generated. During the optimisation process, each candidate configuration must satisfy the geometric constraints imposed by the deployment boundary and the infrastructure exclusion zones. Additional geometric conditions are enforced to ensure physically consistent spacing between WEC units, preventing unrealistic overlap between devices.

The layered spatial representation described above forms the geometric basis of the optimisation framework. Candidate arrays generated by the GA are therefore guaranteed to satisfy the spatial restrictions associated with the co-located offshore energy park before their hydrodynamic performance is evaluated. A schematic illustration of the spatial-domain construction adopted in this study is presented in Figure 25.

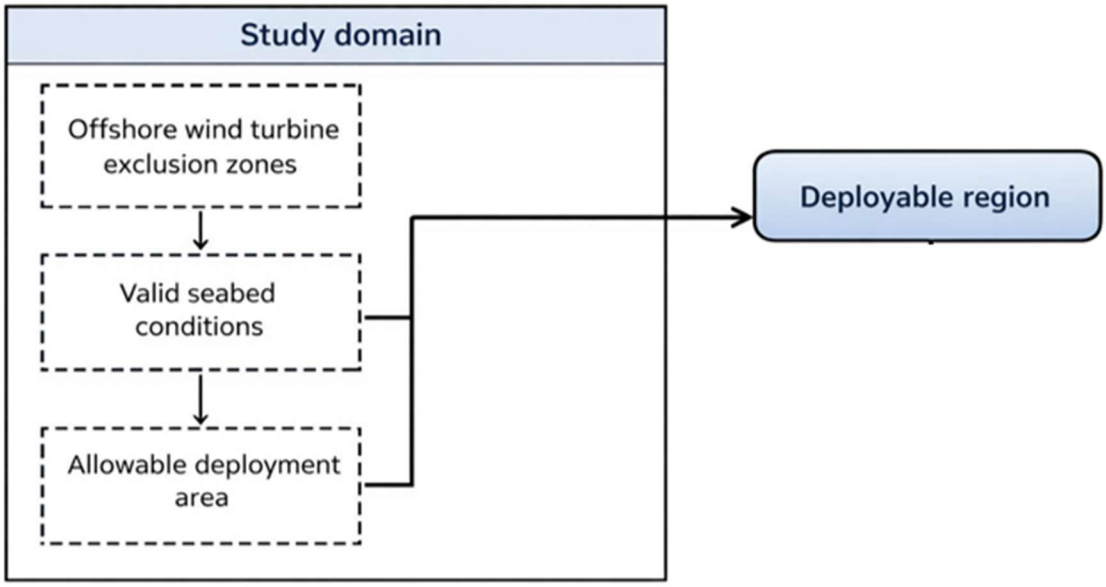


Figure 25 - Conceptual representation of the spatial-domain construction used in the optimisation framework.

The implementation of the WEC representation within the hydrodynamic model, including the performance parameterisations adopted in the optimisation, is described in the following section.

3.3. HYDRODYNAMIC REPRESENTATION OF THE CORPACK CLUSTER

3.3.1. PERFORMANCE REPRESENTATION OF THE CORPACK CLUSTER

The accurate representation of wave energy converter arrays within large-scale spectral wave models remains a significant challenge for integrated offshore energy system studies. While detailed hydrodynamic simulations can resolve the behaviour of individual devices within a wave energy farm, such approaches become computationally impractical when combined with large spatial domains and optimisation frameworks requiring thousands of model evaluations. Consequently, simplified yet physically consistent representations of wave energy converter clusters are required to enable the integration of wave energy systems into large-scale spatial optimisation studies.

Spectral wave models represent the interaction between wave energy converters and the incident wave field through parameterised formulations that remove energy from the wave spectrum according to the device performance characteristics. In the SWAN spectral wave model and its WEC-oriented extension SNL-SWAN, this interaction is typically defined through device performance inputs describing absorbed power as a function of the sea state [84], [86], [87]. It should be noted that SWAN and its SNL-SWAN extension are linear spectral wave models. As such, they do not explicitly resolve radiation effects or strong non-linear hydrodynamic interactions, such as those associated with wave breaking, overtopping, or viscous dissipation processes. Consequently, the model framework provides a simplified representation of wave-device interactions, and the extrapolation of the results to very high H_s should be interpreted with caution, particularly under extreme sea-state conditions.

For individual wave energy converters, performance is commonly expressed through a power matrix defined as a function of H_s and T_p. Such matrices are widely used in wave-energy resource assessments and numerical modelling studies to characterise the power absorption behaviour of a device under different sea states [10], [88], [89]. Within spectral wave models, these matrices allow absorbed power to be determined by linking the local sea-state parameters resolved by the model to experimentally or numerically derived device performance curves [84].

However, when multiple converters operate within a cluster or array, hydrodynamic interactions between devices modify the local wave field experienced by each unit. Wave shadowing effects generated by upstream absorbers and radiation interactions between neighbouring devices alter the distribution of wave energy within the array, meaning that the aggregate response of a cluster cannot be obtained through a simple linear scaling of the performance of a single device [90], [91].

In the case of the CorPower technology considered in this work, the deployment concept is based on the CorPack configuration, in which multiple point absorber devices are arranged in a regular cluster layout. Although the hydrodynamic behaviour and performance characteristics of individual CorPower C4 converters have been investigated through numerical modelling and experimental testing campaigns [69], publicly available literature does not provide a complete power matrix describing the aggregate performance of the full cluster configuration under varying sea states.

One of the few studies addressing cluster-scale CorPack dynamics is the SafeWAVE marine dynamics modelling study [68], which simulated the hydrodynamic behaviour of a 28-device cluster under realistic wave forcing conditions. In that work, the full cluster was explicitly represented by modelling each individual converter within a high-resolution numerical domain. While the study provides valuable insight into cluster-level behaviour and reports performance indicators such as capacity factor, it does not provide a complete sea-state power matrix that could be directly implemented in spectral wave models.

For the present study, the absence of an aggregated CorPack performance matrix represents a methodological limitation. The modelling framework adopted in this thesis requires repeated wave simulations over a large offshore domain in which the WEC layer must be represented in a computationally efficient manner. Explicitly resolving all individual devices within each CorPack directly in the full-domain simulations would introduce a prohibitive computational burden, particularly when simulations must be repeated across multiple scenarios.

Consequently, an equivalent aggregated representation of the CorPack cluster was required. The objective of the methodological procedure presented in this chapter is therefore to derive a cluster-level performance representation that can be used within SNL-SWAN to model the CorPack as a single equivalent obstacle while preserving the hydrodynamic behaviour of the explicit multi-device configuration as closely as possible.

To achieve this, a dedicated numerical modelling framework was developed to simulate the CorPack response under controlled sea states. The results of these simulations were used to construct two alternative aggregated performance representations:

i. a conventional power matrix;
ii. a RCW representation.

Both representations were subsequently validated against the SafeWAVE reference simulations to assess their ability to reproduce the performance of the explicit 28-device configuration. The methodological procedure used to derive these representations, and the validation results are presented in the following sections.

3.3.2. Directional Dependence of Cluster Performance

In spectral wave modelling, the absorbed power of a wave energy converter is typically represented as a function of the incident sea state through a two-dimensional power matrix defined in terms of significant wave height H_s and T_p, as illustrated in Figure 19. This formulation is appropriate for many single-device converters because their hydrodynamic response is largely independent of the direction of wave propagation when considered in isolation, particularly in the case of axisymmetric point absorbers [88], [89]. Consequently, device performance can often be described adequately through a direction-independent $P(H_s, T_p)$ relationship.

However, this assumption becomes less appropriate when multiple converters operate within a cluster configuration. In such cases, the relative orientation between the incoming wave field and the spatial arrangement of the devices determines the order in which converters interact with the wave field. Devices located upstream with respect to the incident wave direction partially absorb and redistribute wave energy, thereby modifying the wave conditions experienced by downstream devices. This phenomenon, commonly referred to as wave shadowing, results in spatial variability of the wave field inside the array and therefore affects the distribution of absorbed power among the devices [90], [91].

For a cluster such as the CorPack, this implies that the aggregate power absorption of the cluster is not solely a function of H_s and T_p but also depends on the incident wave direction. Different propagation directions produce different internal interaction patterns within the cluster, altering the degree of shadowing between devices and thus modifying the total absorbed power.

From a modelling perspective, this leads to a conceptual difficulty. The most physically consistent way to represent the cluster response would be through a three-dimensional performance function $P(H_s, T_p, \theta)$, where θ denotes the incident wave direction. Such a representation would allow the cluster

performance to vary with the relative orientation between the wave field and the device layout. However, the performance input formulations implemented in SNL-SWAN do not support a three-dimensional performance matrix. Instead, the model requires either a conventional two-dimensional (H_s, T_p) power matrix or a one-dimensional relative capture width curve as input [84], [87].

As a result, an approximation is required to represent the aggregated behaviour of the cluster within the modelling framework. In the present work, two complementary strategies were adopted to address this limitation.

The first strategy retains the conventional (H_s, T_p) power matrix format but incorporates directional information indirectly. For each (H_s, T_p) bin derived from the wave dataset, a representative incident direction is assigned based on the most frequently occurring direction within that bin. Cluster response simulations are then performed using this representative direction, allowing the resulting absorbed power to reflect typical directional conditions associated with each sea state.

The second strategy explicitly accounts for directional effects during the pre-processing stage by analysing cluster response as a function of wave direction and period. The absorbed power obtained from these directional simulations is then converted into a RCW representation, which relates the absorbed power to the incident wave energy flux and characteristic device width. Directional variability is subsequently incorporated through a weighted averaging procedure based on the directional occurrence frequencies in the wave climate.

These two approaches provide alternative aggregated representations of the CorPack response that can be implemented within SNL-SWAN. The first preserves the conventional performance matrix used in many wave-energy studies, while the second provides a directionally informed RCW parameterisation.

The methodology used to derive each representation is described in the following sections.

3.3.3. Numerical Modelling Framework for Cluster Response Simulations

To derive an aggregated performance representation of the CorPack cluster, a dedicated numerical modelling framework was developed in which the full 28-device configuration was explicitly simulated under controlled sea-state conditions. The objective of this modelling stage was not to represent the full offshore domain considered later in the case study, but rather to obtain a robust description of the hydrodynamic response of the cluster that could subsequently be translated into an equivalent aggregated representation suitable for implementation in SNL-SWAN.

The simulations were performed using the SWAN spectral wave model and its wave-energy extension SNL-SWAN. The model solves the spectral action balance equation, which describes the evolution of wave energy density in geographic and spectral space under the influence of physical processes such as wave propagation, refraction, depth-induced breaking, bottom friction and nonlinear wave-wave interactions [87]. The SNL-SWAN extension introduces additional formulations that allow wave energy converters to be represented within the spectral wave framework through obstacle-based energy extraction mechanisms [84].

Since the objective of this stage was to capture the interaction between devices within the CorPack cluster rather than large-scale wave transformation processes, a simplified computational domain was constructed. The domain was generated using Delft3D tools and consisted of a rectangular grid with dimensions of 3000 m in the y direction and 600 m in the x direction. This configuration provides sufficient fetch for the development of the wave field upstream of the cluster while allowing downstream wave shadowing effects to be resolved within the computational domain.

The computational grid was discretised with a resolution of 1.5 m in y direction and 2.0 m in the x direction, resulting in a high-resolution grid capable of resolving the spatial distribution of wave energy in the vicinity of the devices. Such resolution is necessary to capture the local attenuation and redistribution of wave energy generated by the interaction of multiple absorbers within the cluster.

Bathymetry within the modelling domain was defined as a uniform water depth of 150 m. This value corresponds approximately to the mean water depth of the offshore study region considered in this work and is representative of deep-water conditions for the wave periods of interest. Since the objective of the simulations was to evaluate cluster-scale hydrodynamic interactions rather than bathymetric wave transformations, the adoption of a constant depth allows the modelling framework to isolate the effects associated with device interactions while avoiding additional complexity associated with depth variability.

Within this computational domain, the CorPack cluster was represented explicitly by modelling each of the 28 individual C4 converters. Each device was implemented as a linear obstacle segment with a characteristic length of 9 m (on y direction), consistent with the device representation adopted in SNL-SWAN obstacle formulations. The devices were arranged according to the configuration described in Section 2.2.3, with a spacing of 150 m between neighbouring converters, as illustrated in Figure 26.

This explicit representation allows the simulations to capture intra-cluster hydrodynamic interactions, including wave shadowing and redistribution of wave energy among the devices.

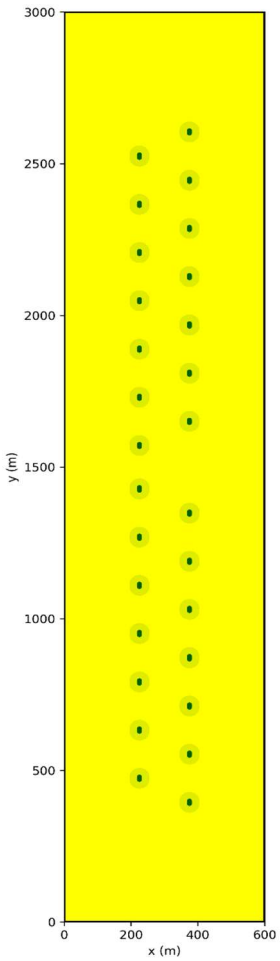


Figure 26 - Representation of the CorPack layout implemented in the SWAN computational domain.

By resolving the individual devices within the simplified domain, the resulting simulations provide a physically consistent estimate of the total absorbed power of the CorPack cluster under specified sea-state conditions. These results form the basis for constructing the aggregated performance representations described in the following sections.

3.3.4. CorPack Power Matrix Representation

Due to the spatial scale of the optimisation domain considered in this work, representing each individual device of the CorPack cluster explicitly within the spectral wave model would lead to an excessive computational burden. A reduced representation of the cluster performance is therefore required. To address this limitation, an aggregated power matrix was constructed to describe the hydrodynamic behaviour of the CorPack cluster as a function of the incident sea state. The objective of this procedure was to construct a cluster-level power matrix expressing absorbed power as a function of H_s and T_p, analogous to the power matrix used to describe individual C4 devices.

The methodology adopted to construct this aggregated representation is summarised in Figure 27, which illustrates the sequence of steps used to derive the cluster power matrix from the wave dataset and the numerical simulations.

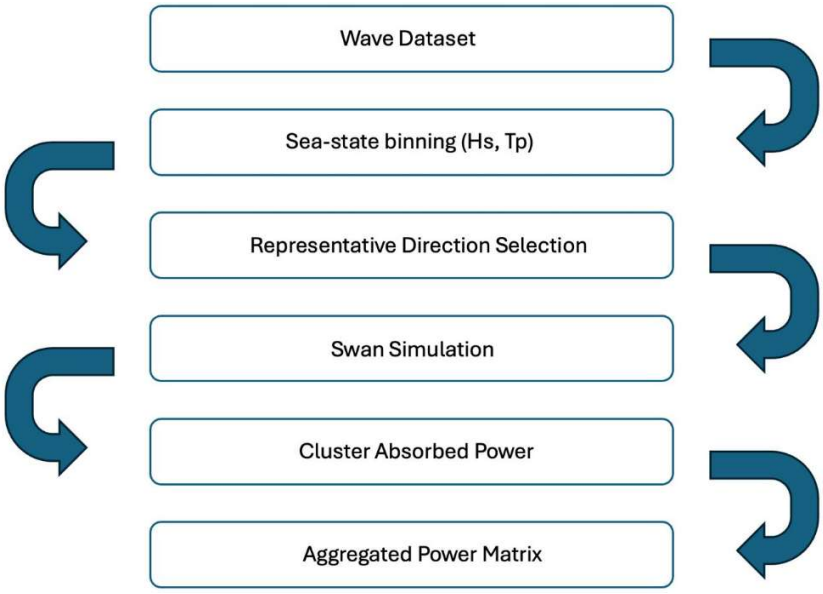


Figure 27 - Workflow adopted for the construction of the aggregated CorPack power matrix.

The starting point for this procedure was the long-term wave dataset used in the resource characterisation of the study area. This dataset was discretised into bins of H_s and T_p to define a representative set of sea states covering the range of conditions relevant for the offshore site. H_s was discretized from 0 to 11 m, using intervals of 0.5 m, while T_p was discretized from 2.5 to 21.0 s using 0.5 s intervals, as presented in Figure 28. This discretization provides a sufficiently detailed description of the sea-state space relevant to the offshore site while maintaining a manageable number of numerical simulations.

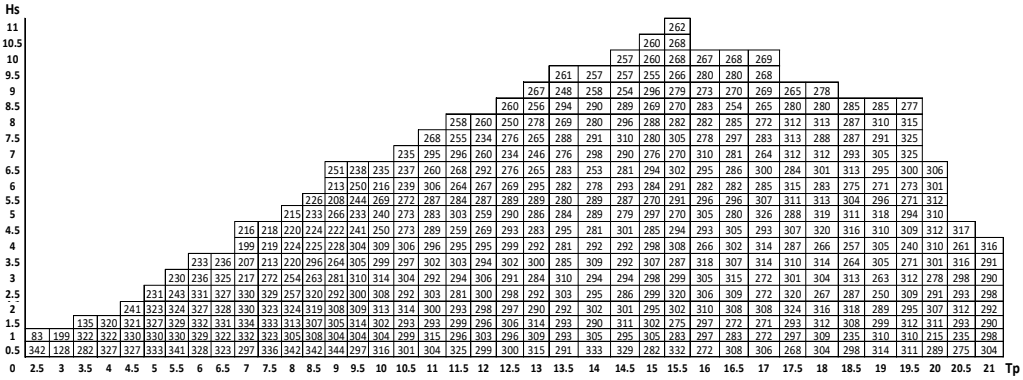


Figure 28 - Discretised sea-states bins used to define the CorPack cluster simulations.

As discussed in Section 3.3.2, the hydrodynamic response of a multi-device cluster depends not only on the sea-state parameters previously referred, but also on the incident wave direction, which determines the order in which devices interact with the incoming wave field and therefore the internal shadowing patterns within the cluster. To incorporate directional effects while preserving the conventional two-dimensional power matrix structure, a representative incident direction was assigned to each discretized sea-state bin.

For each (H_s, T_p) bin, the incident wave direction used in the simulations was selected as the most frequently occurring direction observed in the underlying wave dataset for that specific combination of H_s and T_p. This procedure ensures that the directional conditions associated with each sea state are represented in a statistically consistent manner while avoiding the introduction of an additional dimension in the performance representation. Using the numerical modelling framework described in Section 3.3.3, stationary SWAN simulations were then performed for each (H_s, T_p, θ) combination defined by the discretized bins and their representative directions. In all simulations, identical boundary conditions were imposed along the four boundaries of the computational domain to generate a spatially uniform incident wave field within the domain. The incoming wave spectra was defined using a JONSWAP spectral formulation, with the sea-states parameters prescribed according to the values presented in Figure 28. The simulations' results dataset was subsequently organized into a two-dimensional matrix, representing the aggregated absorbed power of the CorPack cluster under the range of sea states considered. These results are illustrated in Figure 29, where absorbed power is shown as a function of H_s and T_p. For clarity and improved visual inspection, the same layouts are also provided in larger format in Appendix A.

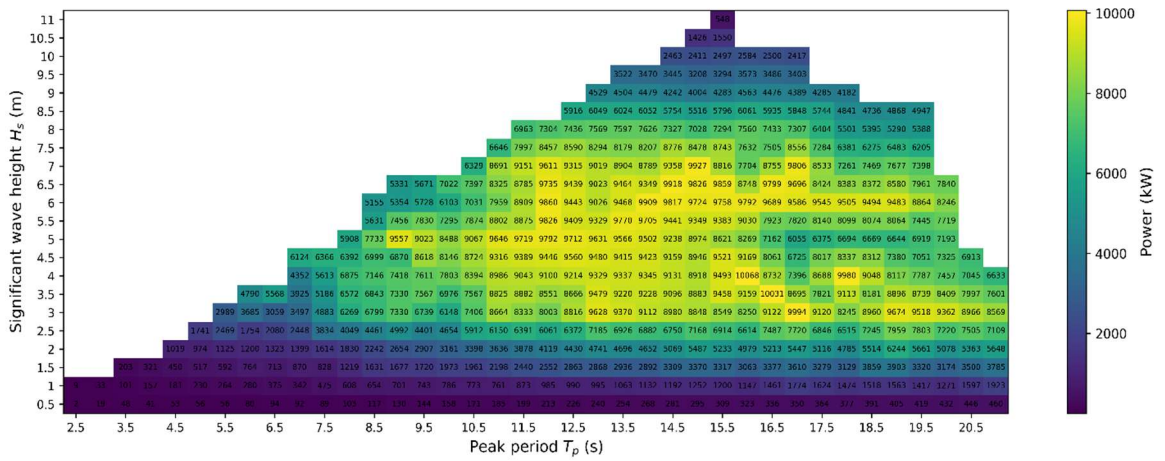


Figure 29 - CorPower CorPack power matrix.

The aggregated matrix shown in Figure 29 provides a compact representation of the hydrodynamic behaviour of the CorPack cluster across the discretised sea-state space. Increasing absorbed power is observed for larger wave heights and peak periods, while a saturation region corresponding to the rated operating range of the cluster can be identified at higher energy levels. This matrix therefore constitutes the performance representation used to characterise the cluster response in subsequent modelling steps.

3.3.5. CorPack RCW Representation

While the aggregated power matrix described in the previous section provides a complete representation of the CorPack cluster performance as a function of significant wave height and peak wave period, certain numerical modelling frameworks require a simplified representation of WEC based on relative RCW. In this approach, the interaction between the wave field and the device is described through a capture coefficient relating the absorbed power to the incident wave energy flux.

The RCW is defined as the ratio between the absorbed power and the available wave power across the characteristic width of the device [88]. For irregular waves, the relationship can be expressed as:

$$RCW = \frac{P_{abs}}{J * b} \tag{7}$$

where P_{abs} is the absorbed power of the cluster, b represents the characteristic width of the cluster interacting with the wave field, in this case 2210 m, and J is the incident wave energy flux per unit crest length, and it depends on the spectral wave conditions and can be expressed for irregular waves as:

$$J = \frac{\rho g^2}{64\pi} {H_s}^2 T_e \tag{8}$$

where ρ is the water density, g the gravitational acceleration, H_s the significant wave height and T_e the energy period of the spectrum.

Unlike the power matrix representation developed in the previous section, the RCW representation was constructed using a binning strategy that explicitly prioritises the directional characteristics of the wave climate. Both approaches were derived from the same sea-state dataset, but they differ in the variables used to define the simulation bins. In the power matrix approach the primary binning variables were H_s and T_p, whereas the RCW representation was organised primarily as a function of wave direction and T_p.

Since the CorPack cluster is geometrically symmetric with respect to both horizontal axes, the hydrodynamic response of the array is invariant for equivalent wave directions mirrored across these axes. This property allows the directional space to be reduced to a single quadrant without loss of generality. Consequently, the original directional dataset was first transformed such that all incident wave directions were mapped to the range between 270° and 360°. After this transformation, the sea-state data were organised into directional and period bins. Peak wave periods were grouped using the same interval definition adopted for the power matrix representation, while wave directions were grouped into 10° directional sectors covering the quadrant between 270° and 360°.

For each combination of directional bin and T_p bin, the most frequently occurring value of H_s within the original dataset was selected as the representative value for that bin. This procedure ensures that the resulting simulations reflect the most typical sea-state conditions associated with each directional-period combination. The resulting binning strategy therefore differs fundamentally from that used for the power matrix representation. In the power matrix approach, a single representative wave direction was associated with each (H_s, T_p) bin, whereas in the RCW formulation the representative variable is the H_s associated with each (θ, T_p) bin. In this sense, the power matrix representation primarily emphasises the joint distribution of wave height and period, while the RCW representation prioritises the directional structure of the wave climate.

Once the representative sea states were defined, a series of SWAN simulations were performed following the same modelling procedure described in Section 3.3.4. Each simulation corresponds to a specific combination of T_p and incident wave direction, with the representative H_s value assigned to that bin. The absorbed power obtained from these simulations was organised into a two-dimensional matrix, presented in Figure 30.

Dir																				
355		37.9	72.9	59.8		142.4	130.1	146.7	210.9	282.8	170.4	57.1		940.3						
345		25.2	162.8	255.5	41.4	51.8	478.2	162.7	191.6	387.4	271.8	230.9	176.4	237.2	713.2					
335		38.8	19.5	544.9	886.4	70.4	1025.4	636.5	925.7	811.1	932.9	5547.6	322.4	82.9	90.6		400.4			
325	0.6	74.6	209.8	34.2	1395.5	339.0	534.7	840.8	512.9	689.3	122.6	1237.8	2228.5	0.0	444.4	1604.2	541.9	842.6		
315	0.0	42.2	60.1	545.7	115.4	1185.3	2422.3	2903.9	1590.8	2002.4	451.4	4555.8	2788.3	705.5	572.1	3758.5	695.7	1685.7		
305	0.8	23.5	129.0	142.7	240.7	540.1	2604.3	2384.8	2171.5	1196.9	8130.1	4164.7	3027.6	0.0	730.8	6060.3	958.9	1023.1		4687.5
295	1.7	75.7	88.1	119.2	216.9	373.7	1099.2	2101.3	2947.6	9004.4	9079.4	9108.3	8590.7	9145.8	7393.5	2953.1	6813.0	1276.0		3540.6
285	3.2	30.6	71.0	122.2	628.4	521.2	1326.7	1409.1	4226.2	7784.4	9425.0	9721.1	4097.6	10001.2	9674.6	9672.8	7622.1	8392.9		3406.5
275	0.0	76.7	44.6	103.1	323.5	666.3	1057.6	869.0	3952.3	6524.3	9920.9	10024.2	10047.0	10051.0	6391.1	4111.0	2235.9	4045.2	9241.2	6683.0
	2.5	3.5	4.5	5.5	6.5	7.5	8.5	9.5	10.5	11.5	12.5	13.5	14.5	15.5	16.5	17.5	18.5	19.5	20.5	21.5 Tp

Figure 30 - Absorbed power values obtained, organised as a function of T_p and incident wave direction.

The absorbed power values obtained from these simulations were subsequently converted into RCW values using the relation defined in equations (7) and (8). Although the directional RCW values provide a detailed description of the cluster response, the simplified modelling framework for this work requires a single RCW value for each T_p. To obtain such a representation, the directional RCW values corresponding to each peak-period bin were aggregated through a weighted averaging procedure.

The weighting factors used in this averaging correspond to the relative occurrence frequency of each directional bin within the long-term wave climate dataset. As a result, directions that occur more frequently in the observed wave climate exert a proportionally greater influence on the resulting RCW value. The final RCW curve obtained through this process is shown in Figure 31.

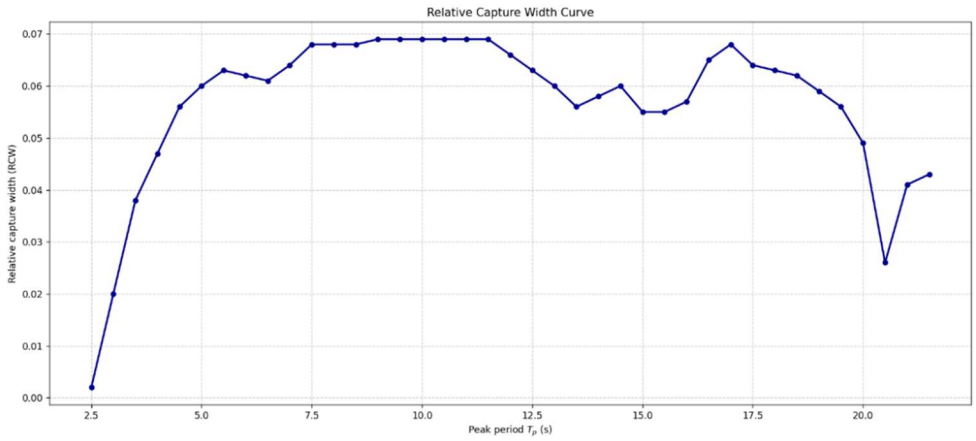


Figure 31 - CorPower CorPack RCW curve.

This curve provides a compact representation of the hydrodynamic efficiency of the CorPack cluster as a function of T_p while preserving the statistical influence of the directional wave climate. The resulting RCW representation is subsequently used as a simplified performance description of the aggregated wave energy converter cluster in the numerical modelling framework adopted in this work.

3.3.6. VALIDATION OF THE CORPACK PERFORMANCE

To evaluate the reliability of the CorPack performance representations developed in the previous sections, a validation exercise was conducted by comparing the energy absorption predicted by the aggregated models with the hydrodynamic performance reported in the detailed CorPack modelling study carried out within the SafeWAVE project [68].

The SafeWAVE study provides a high-fidelity numerical characterisation of the hydrodynamic behaviour of CorPower C4 devices and their interaction within a CorPack cluster configuration. Since that study resolves the interaction between individual devices and the surrounding wave field, its results provide a suitable benchmark against which simplified aggregated representations can be evaluated.

The validation was performed by reproducing the same storm period analysed in the SafeWAVE study, covering a 15-day interval between 15 October 2022 and 30 October 2022. During this period, the wave conditions were prescribed at hourly resolution through boundary forcing in the SWAN modelling framework.

The evolution of the sea-state conditions used in the SafeWAVE study is shown in Figure 3.6-1. The figure presents the temporal evolution of H_s, T_p and mean wave direction measured at the "Instituto Hidrográfico Leixões" wave buoy during the selected storm sequence.



Figure 32 - Sea states used in the SafeWAVE modelling study for the period 15 Oct 2022 to 30 Oct 2022.

The storm sequence presented in Figure 32 is characterised by several energetic events occurring between 17 October and 24 October, during which H_s increase progressively and reach values close to

6 - 7 m during the most intense peak around 23 October. Outside the storm peak, calmer conditions are observed, with H_s generally between 1.5 m and 3 m, followed by a moderate increase again towards the end of the analysed period.

For the purposes of the present study, the same temporal window was reconstructed using the available data and processed to generate the sea-state used in the SWAN simulations. The resulting time series of H_s, T_p and mean direction used in the present modelling framework is illustrated in Figure 33.

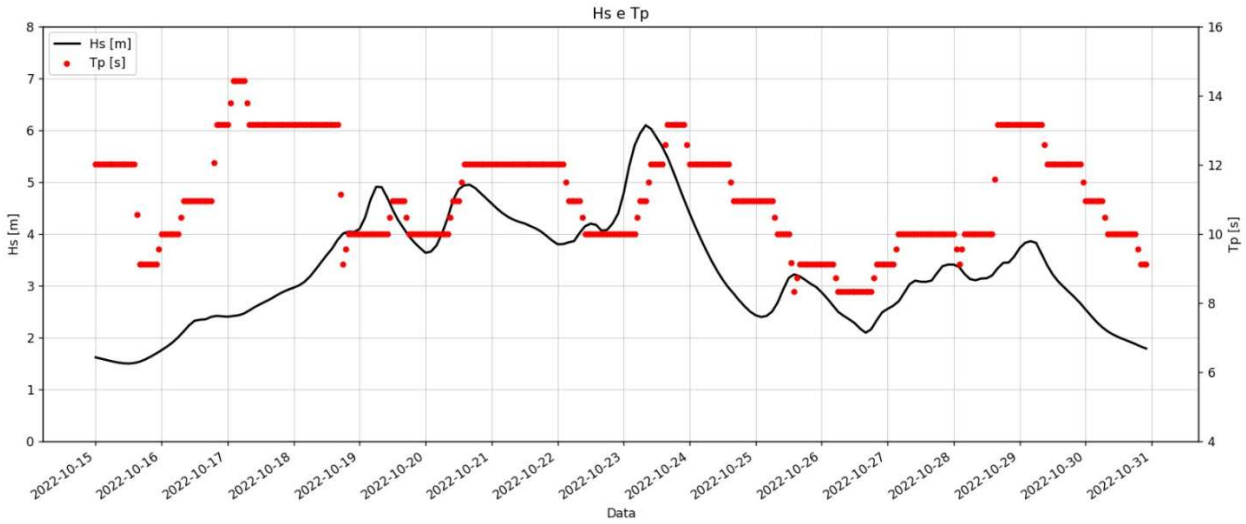


Figure 33 - Sea states used in the present SWAN simulations for the period 15 Oct 2022 to 30 Oct 2022.

A comparison between Figure 32 and Figure 33 shows that the reconstructed series reproduces the overall temporal evolution of the storm sequence. The timing of the main energetic events and the magnitude of the dominant wave peaks remain consistent between the two representations, particularly the main storm event occurring around 23 October, where H_s reach values close to 6 m in both datasets.

However, some differences between the two representations can be observed. The SafeWAVE series exhibits a higher variability in wave direction and T_p, reflecting the direct use of buoy measurements.

In contrast, the sea-state sequence used in the present simulations displays a smoother evolution of H_s and a more discretised representation of T_p values. This behaviour results from the preprocessing required to convert the observational dataset into a format compatible with the SWAN boundary forcing, where representative spectral parameters must be assigned at each hourly time step.

It is important to note that reproducing the exact sea-state sequence used in the SafeWAVE simulations is inherently difficult. The original modelling study employed a specific preprocessing methodology and boundary forcing configuration that are not fully documented in the publicly available deliverable.

Consequently, the storm sequence implemented in the present work represents the closest practical reconstruction of the same energetic period using the available sea-states data and the same temporal window. Although small differences in the exact values of H_s and T_p are therefore unavoidable, the reconstructed dataset preserves the timing, magnitude and spectral characteristics of the dominant storm events, ensuring that the validation simulations are conducted under physically comparable conditions.

To evaluate the performance of the CorPack representations, the validation simulations were performed using a simplified configuration in which the entire CorPack cluster was represented as a single equivalent obstacle, as illustrated in Figure 34, within the numerical domain. This configuration allows

the performance description itself to be isolated from other modelling factors and enables a direct comparison between the different approaches.

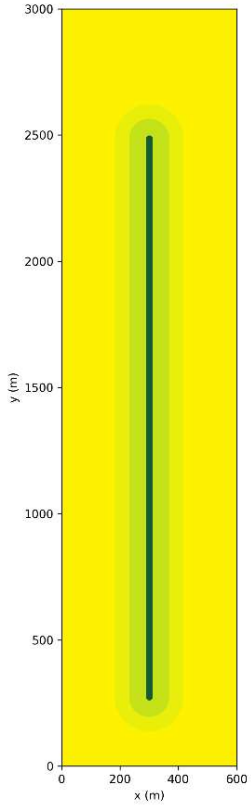


Figure 34 - Simplified representation of the CorPack layout, as a single obstacle, implemented in the SWAN computational domain.

Two alternative performance representations were evaluated:

i. the power matrix representation, described in Section 3.3.4;
ii. the RCW representation, described in Section 3.3.5..

For each case, the absorbed energy produced by the cluster during the 15-day simulation period was calculated and converted into the corresponding capacity factor using the nominal cluster rating. In the present study, the CorPack cluster is assumed to have a nominal power of 11.2 MW, which corresponds to a nominal production of 4032 MWh over the analysed 15-day period.

The resulting comparison between the SafeWAVE reference results and the modelling approaches developed in this work is presented in Table 3.

Table 3 - Validation of the aggregated CorPack representations against SafeWAVE study.

Case	Energy over 15 days (MWh)	Nominal 15-day energy (MWh)	CF	Relative error (%)
SafeWAVE	2620.12	4032	0.65	-
Power Matrix	2378.88	4032	0.59	-9.2
CorPack RCW	2775.51	4032	0.69	+6.0

The comparison with the SafeWAVE reference demonstrates that both aggregated representations reproduce the overall absorption behaviour of the CorPack cluster with satisfactory accuracy for the purposes of large-scale system modelling. Over the analysed 15-day period, the power matrix representation slightly underestimates the absorbed energy (2378.88 MWh, corresponding to a CF of 0.59) when compared with the SafeWAVE reference (2620.12 MWh, CF = 0.65), whereas the RCW-based representation leads to a moderate overestimation (2775.51 MWh, CF = 0.69). The deviations remain within approximately ±10%, which is considered acceptable given the strong model-order reduction involved in representing a cluster of twenty-eight individual wave energy converters as a single equivalent obstacle in a spectral wave model.

These differences are consistent with the distinct physical assumptions embedded in each representation. The power-matrix formulation preserves the dependence of the absorbed power on the local wave parameters Hs and Tp, while the RCW-based approach expresses the device behaviour through an equivalent capture width that is subsequently translated by the SNL-SWAN obstacle formulation into a transmission coefficient Kt^2. As a result, the two approaches distribute absorption differently across the sea-state spectrum, which explains the small deviations observed in the integrated energy estimates.

From the perspective of the optimisation framework adopted in this work, both representations provide a physically consistent and computationally efficient description of the hydrodynamic interaction between waves and the CorPack cluster. However, power matrix preserves H_s-T_p dependence, while RCW is T_p only (directionally weighted), affecting directional sensitivity.

For this reason, the two performance representations are retained for the subsequent modelling steps of this work. In the following chapter, the validated CorPack representation is introduced as a modular wave-energy unit within the project-scale case-study configuration, enabling its integration with the offshore wind farm layout and the spatial optimisation framework.

3.4. Representative Sea-States Generation for Optimisation

The optimisation of WEC array layouts requires the evaluation of candidate configurations under representative environmental forcing conditions. In offshore environments, wave conditions vary continuously in time and are typically described through long-term time series of wave parameters, including H_s, T_p, and mean wave direction. Directly simulating the full time series within the optimisation loop would be computationally prohibitive, since each candidate layout generated by the optimisation algorithm must be evaluated through hydrodynamic simulations. Consequently, the continuous wave climate must be reduced to a limited set of representative sea states capable of capturing the dominant characteristics of the long-term wave resource while maintaining a tractable number of simulations.

To address this challenge, a data-reduction approach is adopted in which the original wave dataset is condensed into a smaller number of representative sea states using statistical clustering techniques. Such approaches have been widely applied in wave energy resource assessment and offshore energy system modelling as a means of preserving the statistical and energetic characteristics of the wave climate while significantly reducing computational requirements [92]. Similar methodologies have also been applied in the optimisation of co-located wave-wind offshore energy systems, where representative sea states are required to evaluate candidate layouts within hydrodynamic optimisation frameworks [53], [74].

The initial dataset consists of long-term wave records describing the offshore wave climate at the study location. Each record is defined by the triplet (H_s, T_p, θ). These parameters provide the environmental

forcing conditions used to drive the spectral wave simulations performed during the optimisation process.

Prior to the clustering procedure, a directional filtering step is applied to isolate the dominant wave-propagation sector influencing the study area. This filtering is performed by analysing the statistical distribution of wave directions and removing directional outliers outside the predominant sector. The procedure is based on statistical thresholds derived from the mean and standard deviation of the directional distribution, ensuring that the retained dataset reflects the primary wave energy pathways reaching the site. The directional characteristics of the filtered dataset are illustrated in Figure 35, which presents the wave-rose representation of the local wave climate. The figure highlights the predominant wave directions affecting the study area as well as the relative frequency of different wave-height classes.

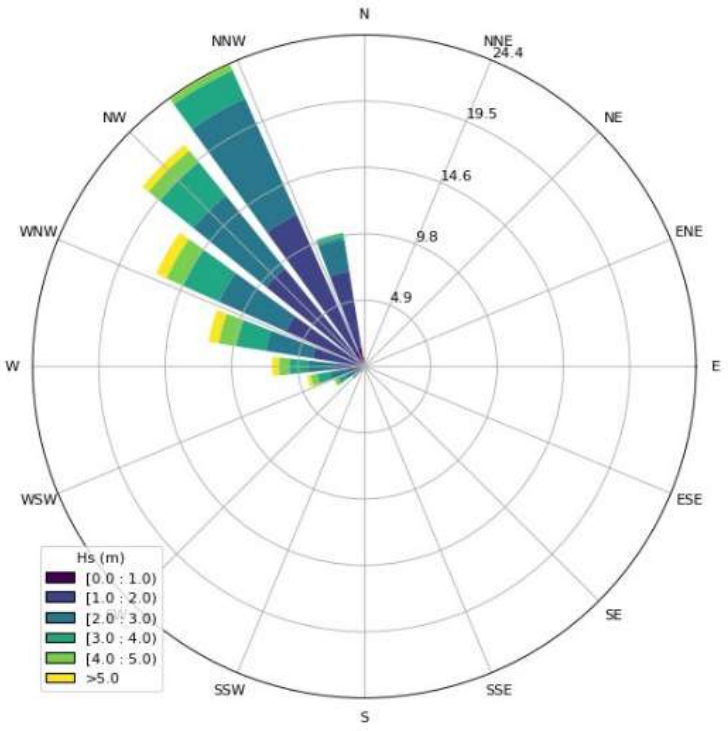


Figure 35 - Directional distribution of H_s for the study domain.

The figure shows that the most energetic wave conditions are associated with the prevailing north-western wave climate, which justifies the directional filtering applied before the clustering stage.

Following the directional filtering, the statistical distributions of the wave parameters are analysed to characterise the variability of the wave climate. Figure 36 presents the histograms of wave direction, Hs, and T_p obtained from the filtered dataset. These distributions provide a preliminary overview of the environmental forcing conditions used to define representative sea states and help identify the most common ranges of each variable prior to clustering

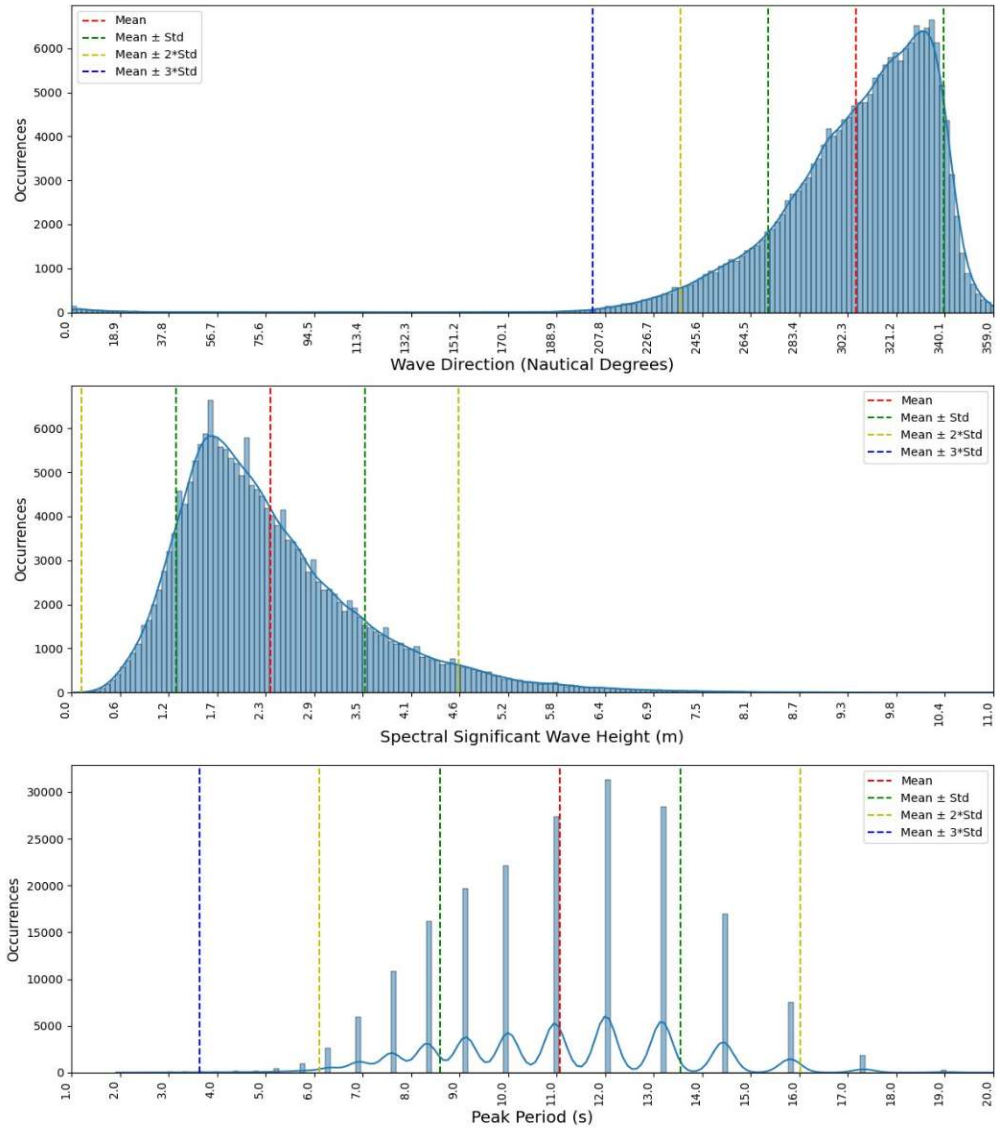


Figure 36 - Statistical distributions of the main wave parameters (H_s, T_p and θ).

The histograms summarise the central tendency and spread of the filtered wave dataset and provide the statistical basis for the subsequent clustering procedure.

To reduce the dataset to a manageable number of representative conditions, a clustering procedure based on the k-means is applied to the three-dimensional parameter space defined by (H_s, T_p, θ). K-means clustering is a widely used unsupervised classification method that partitions a dataset into groups of similar observations by minimising the intra-cluster variance according to the Euclidean distance between data points and cluster centroids [93], [94], and is defined as:

$$S = \sum_{k=1}^{K} \sum_{i \in C_k} \|x_i - \mu_k\|^2 \tag{9}$$

where K is the number of clusters, C_k represents the set of observations belonging to cluster k, x_i denotes the vector parameter associated with observation i, and μ_k is the centroid of cluster k. The algorithm

iteratively updates the cluster centroids and reassigns observations until convergence is achieved and the value of the objective function is minimised. Prior to applying the clustering algorithm, the wave parameters are normalised to ensure that each variable contributes equally to the distance calculation. This normalisation step prevents variables with larger numerical ranges from dominating the clustering process and ensures a balanced representation of wave height, wave period, and wave direction within the classification procedure [53], [92].

The clustering structure of the wave dataset is illustrated in Figure 37, which shows the distribution of sea states in the three-dimensional parameter space defined by H_s, T_p, and θ. Each cluster groups wave conditions with similar hydrodynamic characteristics, while the centroid of each cluster represents one of the representative sea states adopted in the optimisation simulations. The complementary 2D projections included in the figure also help visualise the separation between cluster in the H_s-T_p, H_s-θ and T_p-θ planes.

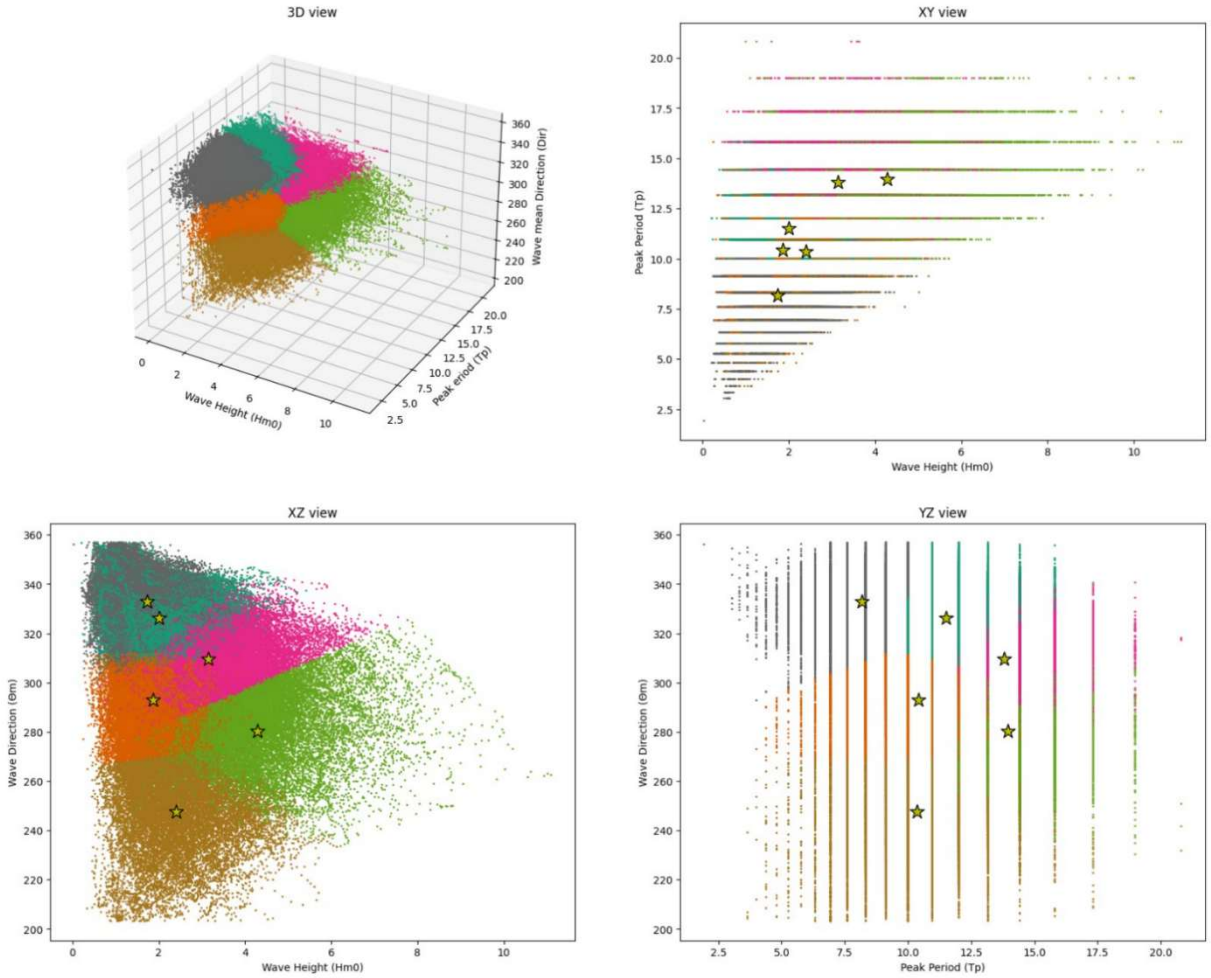


Figure 37 - Three-dimensional wave clustered wave dataset.

The coloured point clouds represent the clusters identified by the k-means, while the yellow markers indicate the centroids corresponding to the representative sea states used in the optimisation. The four panels show the full 3D distribution and its projections onto the three principal planes.

An important aspect of the clustering procedure is the selection of the number of clusters used to represent the wave climate. Increasing the number of clusters generally improves the representativeness

of the reduced dataset but also increases the computational cost of the optimisation process. This trade-off between accuracy and computational efficiency is commonly addressed by evaluating clustering performance for different cluster numbers [53], [92]. The variation of clustering cost as a function of the number of clusters is illustrated in Figure 38 using the elbow method. This approach identifies the point beyond which increasing the number of clusters produces only marginal improvements in dataset representation.

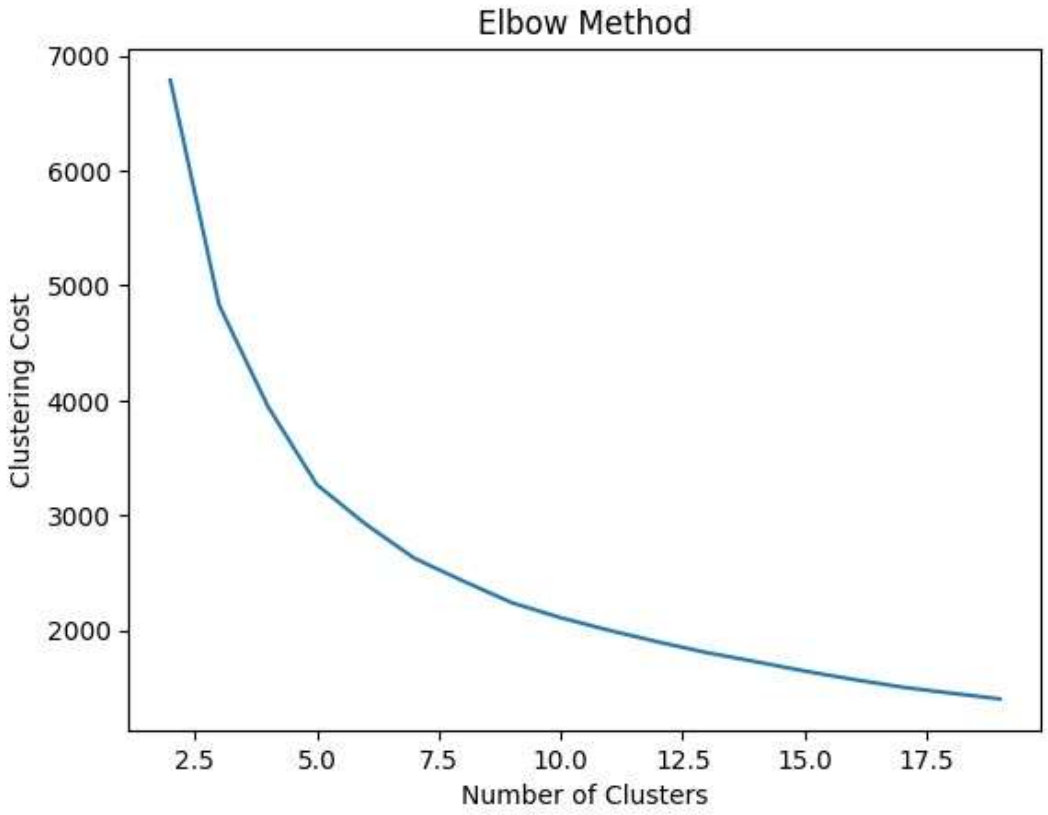


Figure 38 - Elbow method for the optimal number of clusters.

The figure illustrates how the clustering cost decreases as the number of clusters increases, allowing the identification of the point beyond which the gain in representation becomes progressively less significant.

In addition to clustering cost, the representativeness of the reduced dataset is evaluated by comparing the wave-energy representation of the clustered sea states with that of the original dataset. This approach follows common wave-climate reduction practices in which the final number of representative sea states is selected not only on statistical grounds, but also according to the amount of wave energy preserved by the reduced dataset [53], [92]. Figure 39 presents the relationship between the number of clusters and the resulting representation of the total wave energy available at the site.

This procedure allows the selection of a cluster configuration that provides an adequate representation of the wave resource while maintaining a manageable number of simulations within the optimisation framework.

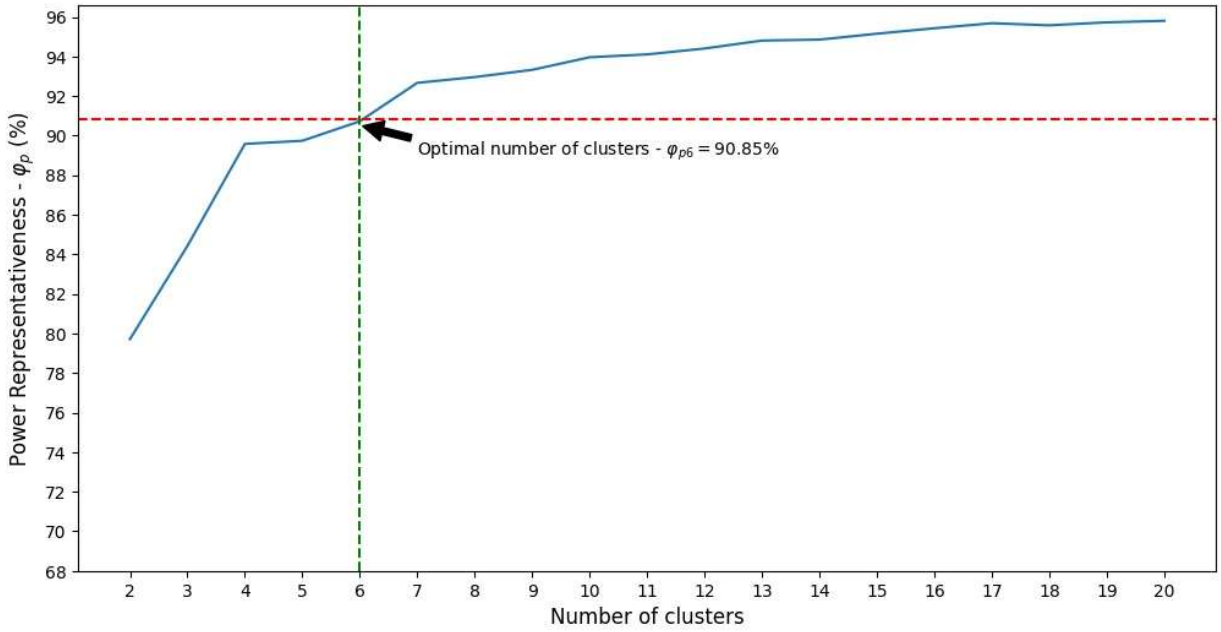


Figure 39 - Power representativeness of the reduced wave dataset as function of the number of clusters.

The dashed horizontal line indicates the selected representativeness threshold, while the vertical dashed line identifies the number of clusters adopted in this study. The figure shows that six representative sea states provide an adequate balance between computational efficiency and wave-energy representation.

The final representative sea states correspond to the centroids of the selected clusters in (H_s, T_p, θ) parameter space. These centroids define the environmental forcing conditions used in the SNL-SWAN simulations during the optimisation process. By reducing the continuous wave dataset to a limited number of representative sea states, the optimisation framework maintains a realistic description of the wave climate while keeping the computational cost of the optimisation within practical limits [53], [74].

At the end, the six representative sea states considered for this work are presented in .

Table 4 - Representative sea states for the optimisation study.

Cluster	H_s (m)	T_p (s)	θ (º)
1	2.00	11.51	326.17
2	1.86	10.41	292.91
3	3.14	13.79	309.70
4	4.29	13.94	280.40
5	2.40	10.35	247.47
6	1.73	8.17	332.84

3.5. HYDRODYNAMIC MODELLING SETUP

The hydrodynamic evaluation of candidate WEC layouts is performed using the SWAN spectral wave model and its SNL-SWAN extension. In the present work, this modelling chain constitutes the numerical core of the optimisation procedure, since each candidate array generated by the GA is evaluated through repeated SNL-SWAN simulations under the representative sea-state conditions defined in Section 3.4.

The numerical simulations are carried out on a curvilinear computational grid covering the study area introduced in Section 3.2. The use of a curvilinear grid allows the model domain to conform to the offshore geometry while preserving a sufficiently regular numerical structure for the repeated simulations required by the optimisation framework [86]. In the present case, the effective spatial resolution of the grid is approximately 30 m in the longitudinal east-west direction and 50 m in the transverse north-south direction. This resolution was selected as a compromise between spatial detail and computational efficiency, ensuring that wave transformation processes and array-induced perturbations could be resolved while maintaining simulation times compatible with the large number of model evaluations required by the optimisation procedure. Similar grid resolutions have been adopted in SWAN-based wave energy studies, where computational grids on the order of 50 m are used to resolve wave conditions near energy devices, while finer grids may be applied for detailed near-field analysis. [95].

The geographical extent of the SWAN domain is defined by the four corner coordinates adopted for the study area, namely (41.2242583, -9.28355), (41.2248917, -9.1530056), (41.1189583, -9.15095) and (41.1189583, -9.33055). These coordinates define a modelling region of approximately 15 km in the east-west direction and 12 km in the north-south direction, which is sufficient to encompass both the deployable WEC region, and the surrounding wave field required to resolve the hydrodynamic response of the array.

Bathymetric input is derived from the GEBCO global bathymetric dataset, which provides seabed elevation data for offshore regions and is widely used in regional wave modelling studies. The GEBCO bathymetry was interpolated onto the SWAN computational grid to generate the depth matrix used by the model. This interpolated bathymetric field constitutes the seabed input file adopted in the simulations and enables the model to account for depth-induced wave transformations, including shoaling and refraction, within the study domain.

The environmental forcing applied in the hydrodynamic simulations is based on the representative sea states derived from the long-term wave dataset, as described in Section 3.4. For each representative condition, a dedicated SWAN input file is automatically generated from a common template file. In this process, the H_s, T_p, and θ corresponding to each selected sea state are inserted into the boundary-condition definition of the model input. This automated procedure produces one SWAN input file per representative sea state and enables a consistent generation of the background and optimisation simulations.

Before the optimisation runs are carried out, a set of background simulations are performed to compute the reference wave conditions in the absence of WEC arrays. These simulations are generated using the python script, which reads the representative sea states from the sea-state table, constructs the corresponding SWAN input files from the base template, executes the simulations, and extracts the mean H_s over the reference integration area from the output file. The resulting values are stored in the background column of the sea-state table and are later used to compute indicators for each candidate array layout.

The representation of the wave energy converters within SNL-SWAN follows the aggregated CorPack methodology presented in Section 3.3. In practical terms, candidate WEC arrays are written to the model input as linear obstacles of fixed length, allowing SNL-SWAN to associate each obstacle with the corresponding hydrodynamic transmission behaviour and energy extraction characteristics.

The distinction between the two hydrodynamic representations considered in this work is implemented directly in the SNL-SWAN model configuration. In the power matrix approach, the array obstacle is linked to an absorbed-power formulation based on the aggregated $P(H_s, T_p)$ description of the CorPack.

In the RCW-based approach, the same obstacle geometry is associated with a RCW input file, which SNL-SWAN uses to determine the corresponding transmission behaviour. This distinction allows the effect of the adopted hydrodynamic representation on the final optimised layouts to be investigated explicitly, which constitutes one of the methodological contributions of the present study.

For each candidate layout and each representative sea state, the model outputs are post-processed automatically using two python scripts. One script extracts the hydrodynamic response of each simulation, including mean H_s and absorbed power, while the other aggregates the results across all representative sea states using the sea-state occurrence weights. This produces the final performance indicators associated with each candidate array and provides the quantitative basis for the optimisation loop described later in this chapter.

To improve computational efficiency, the multiple SNL-SWAN simulations required for each optimisation generation are executed through an automated shell workflow, which manages the parallel execution of model runs and organises the associated outputs. This automated execution procedure enables the optimisation framework to evaluate large numbers of candidate layouts while maintaining a manageable overall computational time.

3.6. OBJECTIVE METRIC DEFINITION

The optimisation of WEC array layouts requires the definition of a quantitative performance metric capable of evaluating the effectiveness of each candidate configuration generated by the GA. Within the present optimisation framework, the performance of each layout is assessed through hydrodynamic simulations carried out for the representative sea states defined in Section 3.4, and the resulting model outputs are aggregated to compute a single performance indicator used as the optimisation objective.

Several performance metrics have been proposed in the literature for evaluating WEC array behaviour, including wave-height reduction indicators, absorbed power, and wave-attenuation indices [53], [74].

In this work, the CF is used as the primary fitness function guiding the GA search. For each candidate layout generated during the optimisation process, the absorbed power obtained is then aggregated using the statistical weights associated with the occurrence probability of each sea state. This weighted aggregation produces an average performance indicator representing the expected long-term behaviour of the array under realistic wave-climate conditions. The resulting value constitutes the fitness score assigned to the candidate layout within the optimisation procedure. Layouts producing higher CF values are therefore favoured by the evolutionary search process, which progressively improves the spatial configuration of the WEC array over successive generations.

The objective function defined above therefore provides the quantitative basis for the evolutionary optimisation procedure described in the following section.

3.7. GENETIC ALGORITHM STRUCTURE AND OPTIMISATION PROCEDURE

The spatial optimisation problem addressed in this work is characterised by a high-dimensional, non-linear search space subject to hard geometric constraints and repeated hydrodynamic evaluations. Under these conditions, deterministic or gradient-based optimisation approaches are not well suited, as they are generally unable to navigate discontinuous feasible domains efficiently and may converge prematurely toward local optima. For this reason, as mentioned on the beginning of this present chapter, the optimisation engine adopted in this thesis is a GA, whose general structure follows the continuous-domain hydrodynamic optimisation framework proposed by *Teixeira-Duarte et al.* [53] for WEC arrays and co-located wave-wind energy systems.

Within this framework, each individual of the population represents a complete candidate WEC layout, encoded through the spatial position of the deployable wave-energy units inside the admissible domain defined in Section 3.2. The overall optimisation loop follows the structure illustrated in Figure 40.

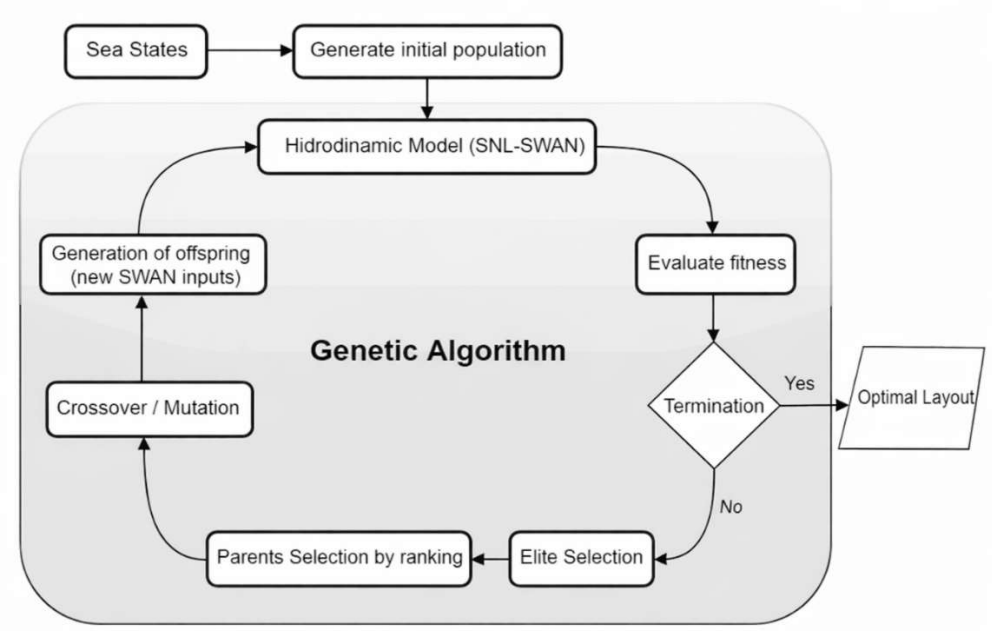


Figure 40 - Genetic Algorithm workflow [53].

As shown in Figure 40, the optimisation procedure begins with the generation of representative sea states and the construction of an initial population of candidate layouts.

After this initial stage, the optimisation proceeds iteratively through successive generations. For each generation, the hydrodynamic model is called to evaluate all candidate layouts under the representative sea-state conditions. The model outputs are processed into layout-level performance indicators, which are then used by the GA to rank the individuals and determine which solutions should contribute genetic information to the next generation. In the framework proposed by *Teixeira-Duarte et al.* [53], this ranking is based on the fitness values obtained after hydrodynamic evaluation, and parent selection is performed probabilistically according to ranking rather than absolute fitness magnitude, which promotes stable convergence behaviour while preserving diversity within the population.

The implementation adopted in this work preserves the same overall logic but is specialised to the present case study and optimisation objective. In particular, the algorithm uses an elite preservation strategy, ensuring that the best-performing layouts of a given generation are carried directly into the

next one without modification. This mechanism prevents the loss of high-quality solutions through stochastic reproduction and guarantees that the best-known layout is not degraded during the evolutionary process. In addition to elite preservation, parent selection is performed by ranking, and new offspring layouts are generated through crossover and mutation operations.

The logical structure of these genetic operators is shown in Figure 41. This figure is particularly relevant in the present work because the optimisation is carried out in a continuous spatial domain, where candidate layouts must satisfy geometric constraints such as minimum separation distances and admissibility within the deployment polygon. Under such conditions, reproduction is not a simple exchange of coordinate vectors, but a constrained geometric operation that must preserve layout feasibility throughout the creation of new offspring.

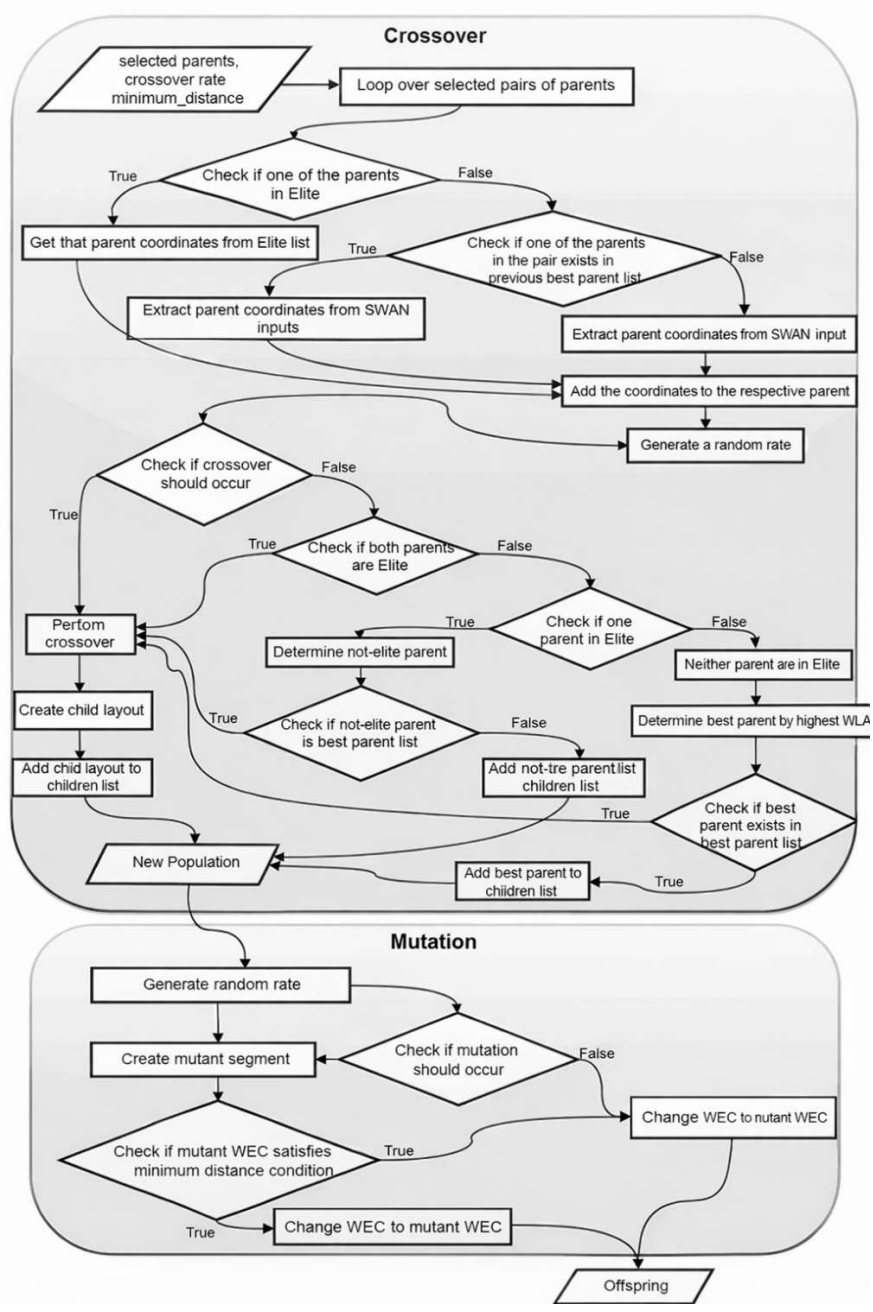


Figure 41 - Schematic representation of crossover and mutation operators applied to layout chromosomes [53].

In the original *Teixeira-Duarte et al.* [53] implementation, the child layout is created by combining WEC positions from the selected parents while enforcing the minimum-distance constraint between units, and a check is introduced to avoid repeated individuals in the offspring population. This logic is also present in the implementation used in the present work. A newly selected CorPack position is accepted only if it remains admissible with respect to the deployment constraints and the minimum separation rule relative to previously placed units in the offspring layout. If this condition is not satisfied, the algorithm retries placement until a valid child is obtained or the reproduction process is restarted. This procedure is essential for continuous-domain optimisation, where the geometric feasibility of the child cannot be guaranteed a priori.

Mutation is then applied to the offspring population to introduce additional diversity into the search process. In the framework proposed by *Teixeira-Duarte et al.* [53], mutation acts by introducing random changes in the WEC segments while respecting the pre-defined layout constraints, thereby injecting novelty into the population and promoting exploration of the solution space. In the present work, the same principle is retained: mutation alters the position of selected CorPacks while preserving the hard geometric restrictions imposed by the deployment domain and minimum-separation condition. Together, crossover and mutation provide the balance between exploitation of high-performing layouts and exploration of new regions of the search space.

Although the general optimisation structure follows that proposed by *Teixeira-Duarte et al.* [53], the present implementation is tailored to the specific objective, spatial constraints, and case-study scale of this work. In particular, the optimisation study was designed to compare both physical modelling assumptions and algorithmic sensitivity parameters. A total of six optimisation cases were considered, and the parameters of each one is presented in Table 5.

Table 5 - Parameters of the six cases presented in this study.

Case	Hydrodynamic representation	Number of CorPacks	Population size	Crossover rate	Mutation rate	Main purpose
1	Power matrix	40	30	0.70	0.01	Evaluation of power matrix
2	RCW	40	30	0.70	0.01	Evaluation of RCW
3	Power matrix	30	30	0.70	0.01	Sensitivity to number of CorPacks
4	Power matrix	30	50	0.70	0.01	Sensitivity to number of arrays
5	Power matrix	30	30	0.95	0.01	Sensitivity to crossover rate
6	Power matrix	30	30	0.70	0.025	Sensitivity to mutation rate

In the context of GA, the crossover rate defines the probability that two selected parent individuals exchange genetic information to produce new offspring solutions. During the crossover operation, portions of the parent layouts are combined to generate a new candidate configuration that inherits

characteristics from both parents. This operator promotes the exploitation of favourable traits already present in the population and is therefore one of the main mechanisms driving convergence toward improved solutions [96], [97]. The mutation rate, in contrast, controls the probability that random modifications are introduced into an offspring solution after the crossover stage.

Mutation typically alters one or more genes of the candidate individual, which in the present case corresponds to modifying the spatial position of one or more units while maintaining the geometric feasibility constraints imposed by the deployment domain. The primary role of mutation is to maintain genetic diversity within the population and to prevent premature convergence of the optimisation process toward local optima [96], [98]. In practical terms, the crossover rate defines the probability that selected parent pairs undergo recombination to produce offspring layouts, while the mutation rate defines the probability of random alterations being introduced into the offspring chromosomes. In the optimisation considered in this work, crossover values of 0.70 and 0.95 and mutation values of 0.01 and 0.025 were examined to assess the sensitivity of the search process to the balance between exploitation and exploration [96], [97].

This optimisation study serves therefore two complementary purposes. First, it enables a physical comparison between the power matrix and RCW representations under the same array size and GA settings (case 1 vs. case 2). Second, it evaluates the sensitivity of the optimisation results to the number of CorPacks and to the main GA control parameters, namely the number of candidate arrays, the crossover rate, and the mutation rate (cases 3 to 6).

Through this structure, the optimisation framework proposed by *Teixeira-Duarte et al.* [53] is not only applied to a new large-scale co-located case study but also extended to assess the robustness of the resulting layouts under different modelling and algorithmic assumptions.

3.8. CONSTRAINTS HANDLING STRATEGY

The optimisation of WEC array layouts within a co-located offshore energy farm requires the enforcement of several geometric and operational constraints to ensure that all candidate solutions generated by the GA remain physically feasible and compatible with the infrastructure of the wind farm.

In the optimisation framework adopted in this work, these constraints are implemented directly within the layout-generation procedures of the GA, ensuring that only admissible configurations are evaluated by the hydrodynamic model.

The constraint-handling strategy follows the continuous-domain optimisation approach proposed by *Teixeira-Duarte et al.* [53], in which feasibility conditions are incorporated into the generation of candidate layouts rather than applied as external penalty functions [53], [74]. Under this approach, the algorithm generates new candidate configurations only within the admissible spatial domain defined in Section 3.2, and each newly proposed device position is validated against the relevant geometric constraints before being accepted as part of the layout.

The first constraint corresponds to the deployment polygon, which defines the spatial region where WEC devices may be placed. This polygon represents the effective deployable area obtained after applying the spatial masks and exclusion buffers described previously. During the generation of candidate layouts, the algorithm samples potential device positions within the bounding region of the deployment domain and verifies whether each proposed location lies inside the admissible polygon. Positions falling outside the domain are rejected and replaced by new random samples until a valid location is obtained.

A second constraint arises from the presence of offshore wind turbines within the co-located energy farm. To avoid interference between WEC units and wind turbine mooring systems, circular exclusion zones are defined around each turbine location. These exclusion buffers represent safety distances that prevent the placement of devices in the immediate vicinity of turbine structures. In the present implementation, these regions are represented geometrically as buffered circular areas centred on the turbine coordinates, and they are subtracted from the deployable region before the optimisation begins.

In addition to the spatial restrictions imposed by the wind farm infrastructure, the optimisation algorithm enforces a minimum separation distance between CorPacks units. This constraint ensures that individual devices within the array do not overlap and prevents unrealistic hydrodynamic interactions caused by excessively close placement of neighbouring units. During layout generation, the algorithm verifies that each newly proposed segment satisfies the minimum-distance condition relative to all previously placed units in the array. If this condition is not satisfied, the proposed position is rejected, and a new candidate location is generated.

Because WEC units are represented as linear obstacles corresponding to CorPack segments, an additional geometric feasibility condition is required to ensure that the entire segment remains within the admissible deployment region. This validation step verifies that the full extent of each segment lies inside the deployable polygon and does not intersect any turbine exclusion zones or previously placed devices.

Through this constraint-handling procedure, the optimisation algorithm guarantees that all candidate layouts evaluated during the evolutionary search satisfy the geometric feasibility conditions imposed by the case study. By embedding these checks directly within the layout-generation routines, the optimisation avoids the need for penalty functions and prevents computational resources from being spent on infeasible solutions. This strategy has been shown to improve the stability and efficiency of evolutionary optimisation in spatial planning problems involving complex deployment domains [53], [74].

The constraint set described above therefore defines the admissible solution space explored by the genetic algorithm during the optimisation analysis.

3.9. Optimisation Convergence and Evaluation Parameters

The optimisation process described in the previous sections is executed through successive generations of candidate layouts, each evaluated through hydrodynamic simulations and ranked according to the objective metric defined in Section 3.6. To ensure a stable and computationally tractable optimisation process, several parameters controlling the population size, evaluation procedure, and convergence behaviour of the genetic algorithm must be specified.

The optimisation procedure begins with the generation of an initial population of candidate layouts, which serves as the starting point for the evolutionary search. In the present implementation, an initial set of 500 candidate layouts is generated randomly within the admissible deployment domain defined in Section 3.2. These layouts are then evaluated hydrodynamically for all representative sea states using the modelling framework described in Section 3.5.

Following this evaluation stage, the best-performing individuals are retained to form the effective population of the first generation of the optimisation. Specifically, the initial population is reduced to the 200 layouts with the highest fitness values, which are stored, and constitute generation gen000 of the optimisation process. This initial selection step ensures that the subsequent evolutionary search begins from a population composed of feasible and relatively high-performing configurations.

After this initial generation, the optimisation proceeds iteratively through successive generations. For each new generation, the GA produces a new set of candidate layouts through the crossover and mutation operators described in Section 3.7. In the optimisation study considered in this work, the number of candidate layouts generated and retained at each generation depends on the analysed case: 30 arrays were used in cases 1, 2, 3, 5, and 6, while 50 arrays were used in case 4. Each layout is simulated under the full set of representative sea states, and the resulting performance indicators are aggregated to compute the fitness value used for ranking the individuals.

At the end of each generation, the evaluated layouts are ranked according to their fitness values, and only the best-performing individuals up to the predefined population size of the corresponding case are retained to form the population of the next generation.

The optimisation process continues iteratively until the termination criterion defined in the algorithm configuration is satisfied, that is, until no significant improvement is observed over a prescribed number of consecutive generations. In the present implementation, this convergence check is based on the stability of the best performance indicators obtained from one generation to the next. A maximum limit of 200 generations is nevertheless imposed as a stopping safeguard, ensuring that the algorithm does not continue indefinitely in cases where full convergence is not reached within a reasonable computational time. The convergence behaviour of the optimisation is analysed later through the evolution of the fitness values across successive generations.

3.10. METHODOLOGICAL WORKFLOW SUMMARY

The methodology adopt in this work combines spectral hydrodynamic modelling and evolutionary optimisation into a unified framework for the design of co-located wave-wind energy layouts. Rather than treating these components separately, the proposed workflow integrates them into a sequential modelling chain in which the output of each stage provides the input required for the next stage of the analysis.

The process begins with the representation of the wave energy system itself through an hydrodynamic description of the CorPack cluster, as presented in Section 3.3. This step provides the performance parameterisations required by the spectral wave model and ensures that the WEC array can be represented in a computationally efficient manner while preserving the main hydrodynamic characteristics relevant to the optimisation process.

The second stage consists of reducing the long-term wave dataset to a limited number of representative sea states, as described in Section 3.4. Through directional filtering, variable normalisation, and k-means clustering, the continuous offshore wave climate is condensed into a reduced set of forcing conditions that retain the dominant energetic and directional characteristics of the site. This reduction is essential to make the repeated hydrodynamic simulations required by the optimisation computationally feasible.

Once the representative sea states are defined, the SNL-SWAN modelling framework is used to simulate wave propagation and WEC-array interaction within the computational domain. As explained in section 3.5, this hydrodynamic model provides the numerical basis for evaluating the performance of each candidate layout generated during the optimisation process. The resulting model outputs are then post-processed and aggregated into layout-level performance indicators. The optimisation itself is performed through a GA specifically developed by *Teixeira-Duarte et al.* [53] for continuous-domain WEC layout optimisation in offshore environments. Within this framework, candidate layouts are generated inside the admissible deployment region, evaluated through hydrodynamic simulations, and ranked according to the objective metric defined in Section 3.6. Successive generations are then produced through elite

preservation, selection, crossover, and mutation, while respecting the geometric feasibility constraints imposed by the co-located wind-wave farm.

The complete sequence of operations implemented in this thesis is synthesised in Figure 42.

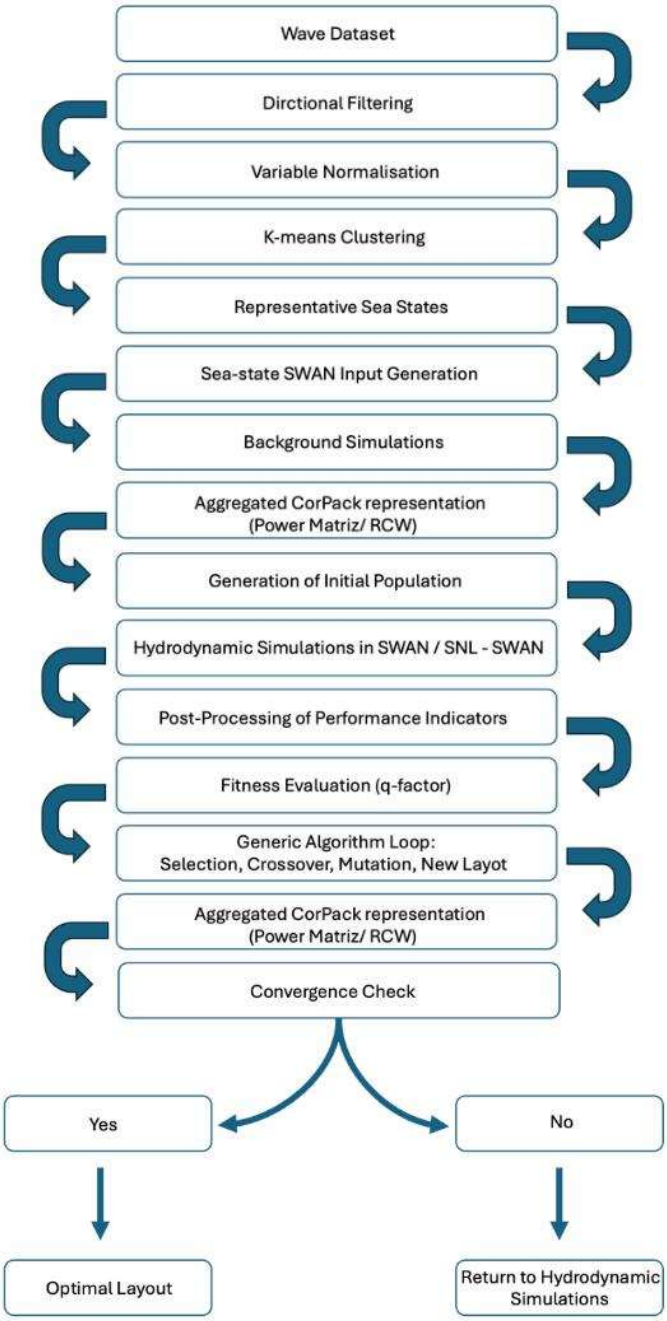


Figure 42 - Summary of the methodological workflow adopted.

Overall, the methodology described in this chapter provides a consistent framework for evaluating and optimising large-scale co-located wave-wind energy arrays under realistic environmental forcing and spatial constraints. The following chapter applies this framework to the selected case-study area and defines the project-scale assumptions used in the optimisation.

4. CASE STUDY

4.1. THE PORTUGUESE CONTEXT: SPATIAL CONSTRAINTS AND RESOURCE ASSESSMENT

4.1.1. PORTUGUESE CONTINENTAL SHELF AND STRATEGIC CONTEXT

Portugal presents a highly favourable offshore setting for the development of large-scale marine renewable energy systems, driven by its extensive Exclusive Economic Zone (EEZ), as illustrated in Figure 43, and its direct exposure to the energetic North Atlantic wind and wave climate [9], [60]. The west-facing continental margin lies within the main Atlantic storm corridor, promoting persistent wind regimes and long-period swell systems that propagate toward the Iberian Peninsula with limited energy dissipation [9], [18].



Figure 43 – Portugal's Exclusive Economic Zone and Extended Continental Shelf [99].

Long-term reanalysis and hindcast datasets confirm that offshore wind and wave resources exhibit strong spatial coherence across the Portuguese continental shelf, with particularly energetic conditions observed along the northern and central sectors [9], [60]. These characteristics have positioned Portugal among the most promising regions in Europe for floating offshore wind deployment and for integrated wind-wave energy concepts [9], [72].

To support the structured development of offshore renewables, Portugal has established the PAER, which defines six concession zones (P1 to P6), as presented in Figure 44, distributed along the western coastline, based on metocean resource availability, bathymetric suitability for floating systems, proximity to ports and grid connection points, and compatibility with existing maritime activities [27], [100]. Together, these zones provide the regulatory and spatial framework for future commercial offshore renewable energy projects [27].

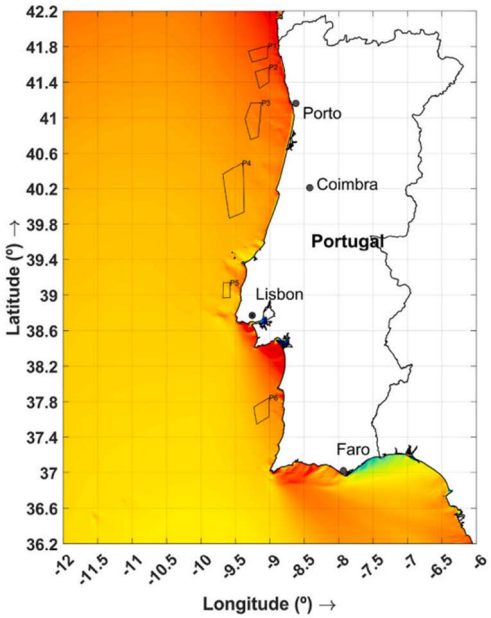


Figure 44 - PAER zones P1 to P6 [9].

Integrated assessments of offshore wind and wave resources across these concession zones reveal clear geographical trends, with northern and central areas generally exhibiting higher mean wind power density and wave energy flux than southern regions [9]. In addition to average resource magnitude, temporal variability and directional persistence also differ between zones, influencing both expected energy production and operational accessibility [9]. Such spatial heterogeneity highlights the importance of site-specific characterisation when designing co-located offshore energy systems [27].

Beyond resource magnitude alone, the Portuguese offshore environment also exhibits favourable wind-wave complementarity. Atmospheric wind events typically precede wave energy arrival due to the propagation time of swell systems, introducing a temporal offset between wind and wave power peaks and contributing to smoother combined energy output in co-located systems [9], [18].

This section provides a high-level characterisation of the Portuguese offshore context and planning framework, establishing the physical and regulatory background for the present case study. The following sections progressively narrow this analysis toward the selected PAER concession zone, where detailed environmental conditions and deployment constraints are examined to define the specific optimisation domain.

4.1.2. Regulatory Framework and Spatial Constraints For Optimisation

The transition from resource assessment to deployable offshore energy systems is governed by a robust regulatory and spatial planning framework. In Portugal, the PAER, approved through the Council of Ministers Resolution no. 19/2025, designates six concession zones, P1 to P6, along the Atlantic coast for commercial offshore wind and wave energy development [100].

The PAER was developed within the broader Maritime Spatial Planning framework, following a rigorous screening process that excluded environmentally sensitive areas, major shipping routes, and zones reserved for military activities [27]. Water depths across the designated areas range from approximately 85 m to over 200 m, effectively precluding fixed-bottom foundations and reinforcing the necessity for floating systems [9].

From an optimisation perspective, the PAER boundaries constitute hard spatial constraints. Layout algorithms must strictly respect concession zone polygons, local bathymetry, and navigational exclusion areas. At the micro-scale, additional constraints govern inter-device spacing, minimum safety distances between wave energy converters, and exclusion zones surrounding floating wind turbines, defined by mooring system geometry and platform motion envelopes [53].

In this thesis, these spatial and operational constraints are explicitly encoded into the genetic algorithm framework. This ensures that the resulting optimised layouts are not only energetically efficient, but also physically deployable, operationally safe, and compliant with the regulatory and spatial realities of offshore renewable energy development on the Portuguese continental shelf.

4.1.3. Offshore Wind Resource Assessment

The offshore wind climate along mainland Portugal is governed primarily by large-scale atmospheric circulation over the North Atlantic, combined with regional coastal dynamics that promote persistent northerly flow during summer and more variable westerly regimes during winter [9]. The reduced surface roughness of the marine environment relative to land allows wind speeds to remain higher and more spatially coherent offshore, creating favourable conditions for large-scale floating wind deployment [100].

To characterise the national offshore wind resource, *Majidi et al.* [9] performed a long-term ERA5-based reanalysis covering the period 1979-2022 to derive spatial fields of hub-height wind speed and wind power density across the Portuguese continental shelf and the PAER concession zones, as illustrated in Figure 45 [9].

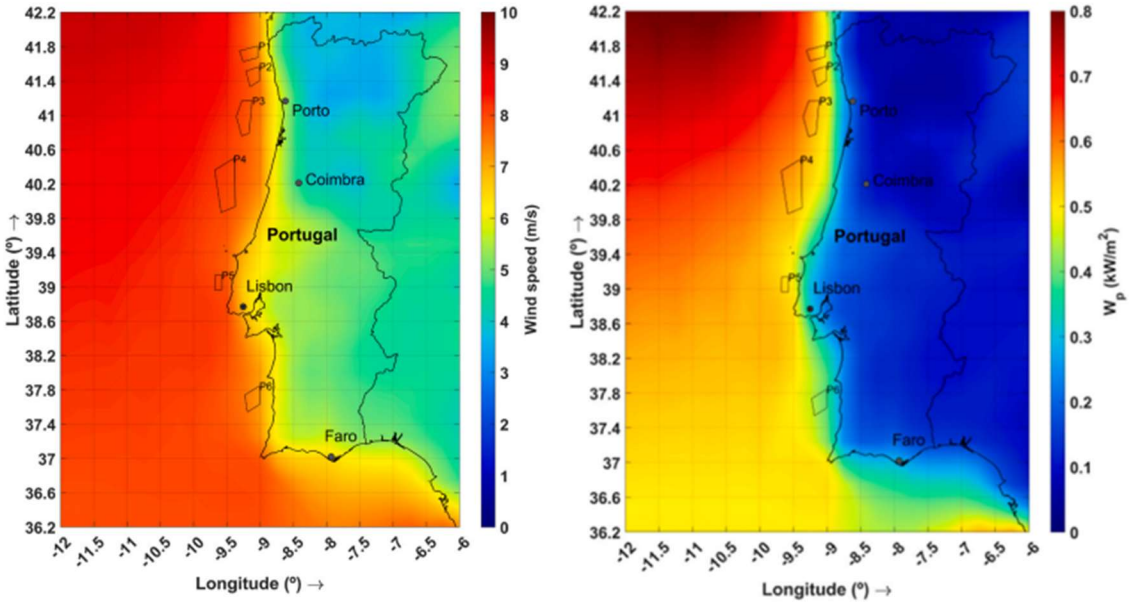


Figure 45 - Spatial distribution of wind speed (left) and wind power density (right) along mainland Portugal [9].

Figure 45 shows that the northern and central offshore sectors exhibit consistently higher wind speeds and wind power density compared to southern regions. Wind power density (W_p) represents the mean kinetic energy flux of the wind per unit area and provides a direct measure of theoretical wind energy availability at a site. Because wind power scales approximately with the cube of wind speed, even moderate increases in velocity translate into substantial gains in available energy. Offshore wind

climates with mean power densities exceeding approximately 500 W/m² are typically considered highly suitable for commercial-scale wind farm deployment.

Across the PAER zones, incident W_p ranges from approximately 350 W/m² in southern areas to values approaching 650 W/m² in the northern and central offshore regions [100]. These values place much of the western Portuguese shelf within the upper tier of European offshore wind resources, confirming its strategic relevance for floating wind development.

Complementing power density, hub-height wind speed at 100 m above the mean sea level provides a practical indicator of turbine productivity. National planning documentation reports offshore wind speeds between approximately 6.5 m/s and 8.0 m/s across the PAER zones [100], which is consistent with the results obtained by *Majidi et al.* [9], Figure 45. Wind speeds above 7 m/s correspond to regimes where modern offshore turbines operate frequently within their efficient aerodynamic range, supporting high annual productivity, making these ranges particularly attractive from an energy-yield perspective.

To quantify expected annual production more directly, the Portuguese Government introduced the metric "Número de horas de funcionamento equivalentes à potência nominal" (NEPS), representing the equivalent number of hours per year that a turbine would need to operate at rated power to match its actual annual energy output under variable wind conditions [100]. NEPS values between approximately 3000 and 4250 h/year are reported across the offshore zones, with higher values concentrated in the northern and central regions. In practical terms, NEPS serves as a proxy for capacity factor and allows early-stage comparison of offshore productivity without detailed turbine modelling.

Beyond mean magnitude, temporal variability strongly influences system performance and grid integration. Monthly variability of wind power density is quantified using the monthly variability index (MVI), which measures the degree of seasonal fluctuation in energy availability [27]. Lower MVI values indicate more uniform month-to-month production, while higher values imply strong seasonal concentration. For offshore wind, reduced variability is advantageous because it improves predictability, supports steadier power delivery, and enhances utilisation of shared electrical infrastructure. Figure 46 presents the spatial distribution of MVI for the PAER zones [9].

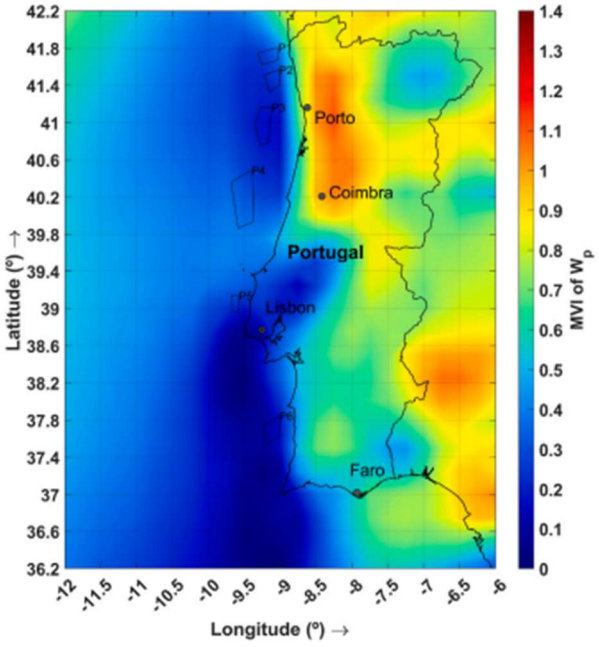


Figure 46 - Spatial distribution of MVI for wind power density along mainland Portugal [9].

Across the Portuguese concession zones, wind MVI values generally remain low, reflecting the relatively stable nature of offshore wind forcing [9]. This characteristic distinguishes wind as the primary contributor to baseload in co-located systems.

Directional structure is another critical parameter in offshore wind farm design, as wake development and optimal turbine-row orientation depend directly on prevailing wind directions. Wind power roses derived for each PAER zone, by *Majidi et al.* [9], presented in Figure 47, reveal a clear predominance of winds originating from the northern sector across most locations.

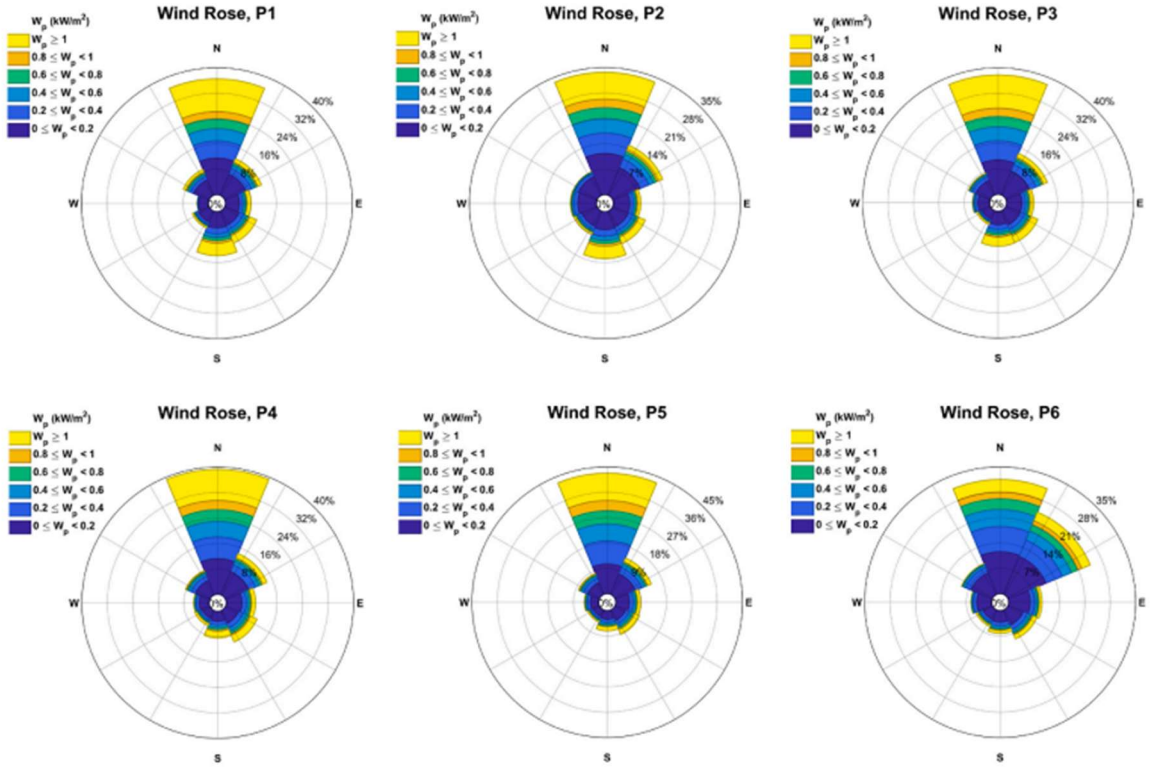


Figure 47 - Wind roses for PAER zones [9].

Figure 47 illustrates that offshore wind forcing along the Portuguese west coast retains a persistent directional signature. A prevailing trend emerges for most sites, excluding P6, where the dominant wind direction, in terms of total duration, originates from the North, accounting for approximately 35% to 40% of the total recorded period [9]. Consistently, the highest W_p intensity typically emanates from the North, succeeded by contributions from the South and Southeast directions [9]. Concession zone P6 presents a localized exception, with the dominant wind direction tending slightly towards the North-Northeast in terms of duration and exhibiting high W_p originating from both the North and the Northeast sectors [9]. This directional coherence is essential for realistic array layout design, as it provides a dominant axis along which turbine spacing strategies can be optimised to mitigate wake losses [9].

Taken together, offshore wind conditions along mainland Portugal are characterised by:

i. high hub-height wind speeds (approximately 6.5 - 8.0 m/s);
ii. high incident wind power densities (up to 650 W/m²);
iii. substantial equivalent full-load hours (up to 4250 h/year);
iv. relatively low monthly variability;
v. a persistent northerly-dominated directional regime.

These characteristics establish the Portuguese west coast as a mature floating wind resource region and provide the national-scale wind framework for subsequent site-specific assessments.

4.1.4. Offshore Wave Resource Assessment

The offshore wave climate along mainland Portugal is dominated by long-period swell generated in the North Atlantic, supplemented by locally generated wind seas during energetic weather systems, producing one of the most powerful wave environments in continental Europe [9], [100]. The west-facing continental margin lies directly exposed to Atlantic storm tracks, allowing swell systems to propagate toward the Iberian Peninsula with limited dissipation and resulting in sustained high wave energy levels across large portions of the continental shelf [9].

Long-term wave conditions were characterised, by *Majidi et al.* [9] using a 44-year hindcast dataset (1979-2022), from which spatial distributions of different parameters were derived for the Portuguese offshore region and PAER concession zones. The hindcast analyses indicate that mean H_s in northern offshore regions typically range between 2.0 m and 2.5 m, with extreme values exceeding 10 m to 12 m during severe storm events [27]. In contrast, southern regions experience lower mean wave heights and reduced extremes, leading to diminished overall wave energy potential. Another important parameter is the wave power density, (P_w). P_w, expressed in kW per metre of wave crest (kW/m), represents the rate at which wave energy propagates through a vertical plane perpendicular to wave direction and serves as the standard indicator of theoretical wave energy availability. Offshore environments with mean wave power exceeding approximately 25 to 30 kW/m are generally classified as high-energy wave climates suitable for utility-scale wave energy deployment [27]. Figure 48, shows the spatial variation of the P_w across the PAER zones.

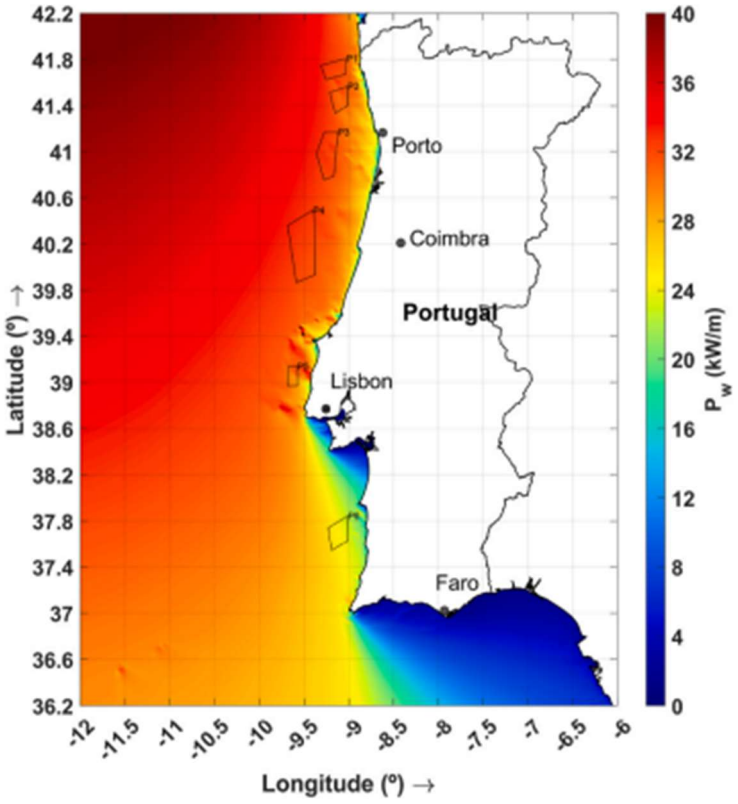


Figure 48 - Spatial distribution of mean wave power density along mainland Portugal [9].

Figure 48 shows that across the Portuguese continental coast, the mean offshore wave power exceeds 25 kW/m in northern and central areas, decreasing progressively toward the south [9], [100]. The PAER documentation reports representative values of approximately 30 to 35 kW/m for the northern and central zones, placing these regions among the most energetic wave environments in Europe [100]. This spatial gradient reflects the attenuation of Atlantic swell energy with latitude and increasing sheltering toward southern Portugal.

In addition to mean magnitude, wave energy exhibits pronounced seasonal variability driven by the North Atlantic storm activity. Winter months are characterised by high significant wave heights and wave power fluxes, while summer periods show substantially reduced energy levels and narrower directional spreading [9]. This behaviour is quantified through the MVI applied to wave power density, which captures the degree of seasonal concentration of energy, as illustrated in Figure 49 [9].

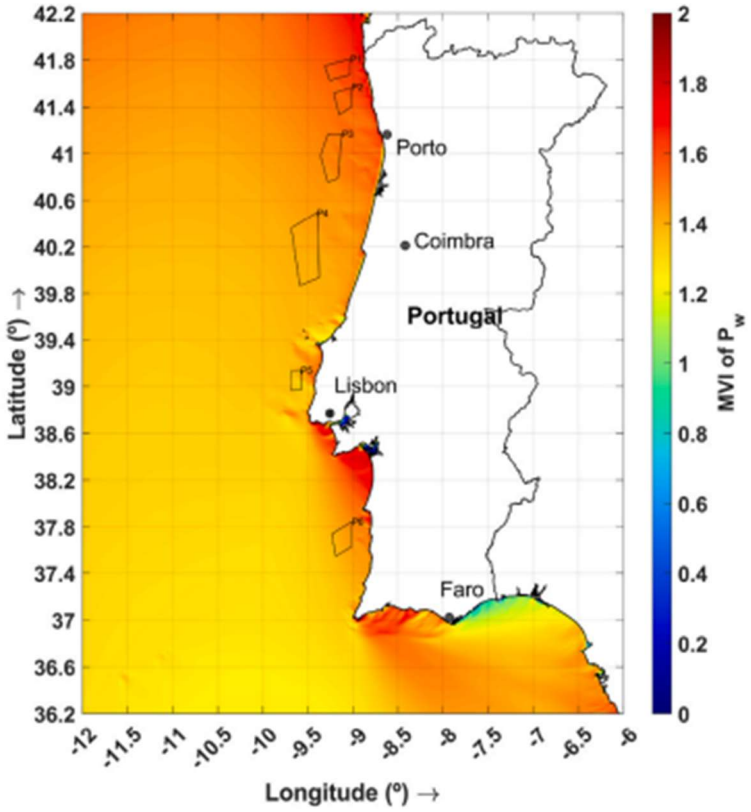


Figure 49 - Spatial distribution of MVI for wave power density along mainland Portugal [9].

Unlike offshore wind, wave MVI values along the Portuguese coast are relatively high, reflecting a pronounced intra-annual variability of the wave resource [27]. This characteristic reflects the resource's swell-driven nature and has essential implications for wave energy systems, as it introduces greater intermittency and concentrates energy extraction into shorter annual windows.

To clarify this temporal behaviour, Figure 50 presents monthly descriptors of wave power for the PAER zones, including the coefficient of variation, skewness, and kurtosis. The coefficient of variation measures relative variability (standard deviation normalised by the mean), the skewness indicates asymmetry of the distribution, and the kurtosis quantifies tail heaviness and the prevalence of extremes.

Together, these indicators provide a compact description of both the seasonal cycle and the intermittency/extremeness of wave power throughout the year.

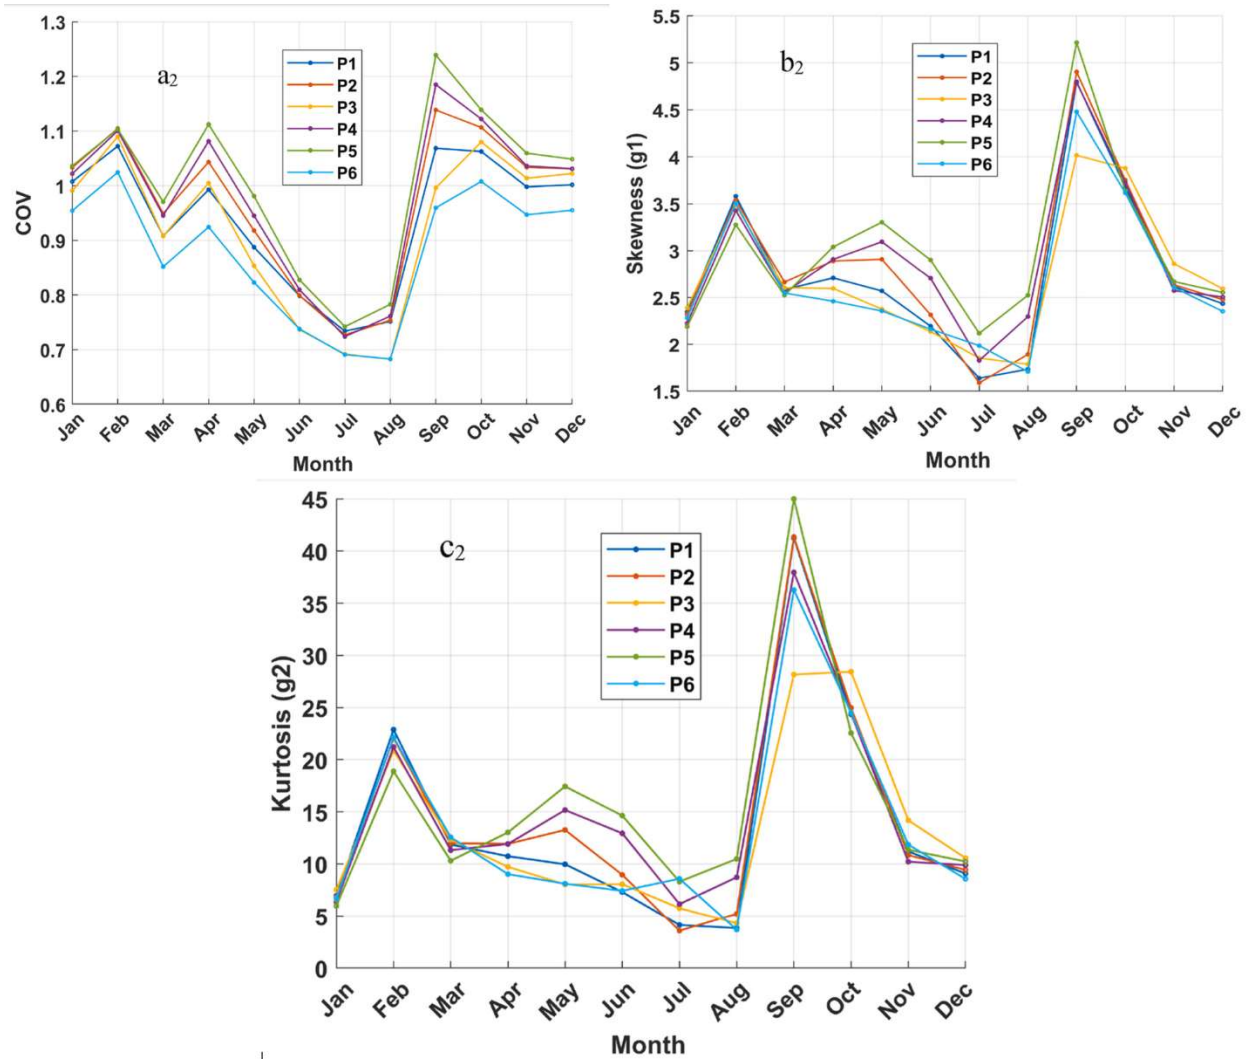


Figure 50 - Monthly distribution of wave power density for PAER zones [9].

Figure 50 shows that the offshore zones exhibit a pronounced winter intensification, with larger variability and stronger non-Gaussian behaviour during the high-energy season. In particular, the sharp increases in skewness and kurtosis in late summer and the elevated CoV values in the energetic season reflect the increased contribution of extreme events to monthly wave power. Conversely, summer months are characterised by substantially lower wave power levels and reduced variability. This seasonal asymmetry explains the elevated MVI values observed in Figure 49. The annual wave-energy potential is temporally concentrated into a limited set of high-energy months, leading to strong intra-annual variation compared to the more stable offshore wind resource. This contrast underpins the potential for wind-wave complementarity explored later in this chapter.

Extreme sea states further characterise the Atlantic exposure of the Portuguese coast. Maximum H_s exceeding 10 m are reported across several offshore zones, indicating severe storm conditions that must be accommodated in structural design, mooring systems, and operational strategies for both floating wind turbines and wave energy converters [27]. These extremes provide first-order survivability constraints and directly influence platform sizing and safety margins.

Directional characteristics also play a critical role in the exploitation of wave energy. Wave power roses derived for the offshore zones (Figure 51) show a dominant northwest to west sector, corresponding to

Atlantic swell propagation toward the Iberian coast, with secondary contributions from westerly systems during winter storm events [9].

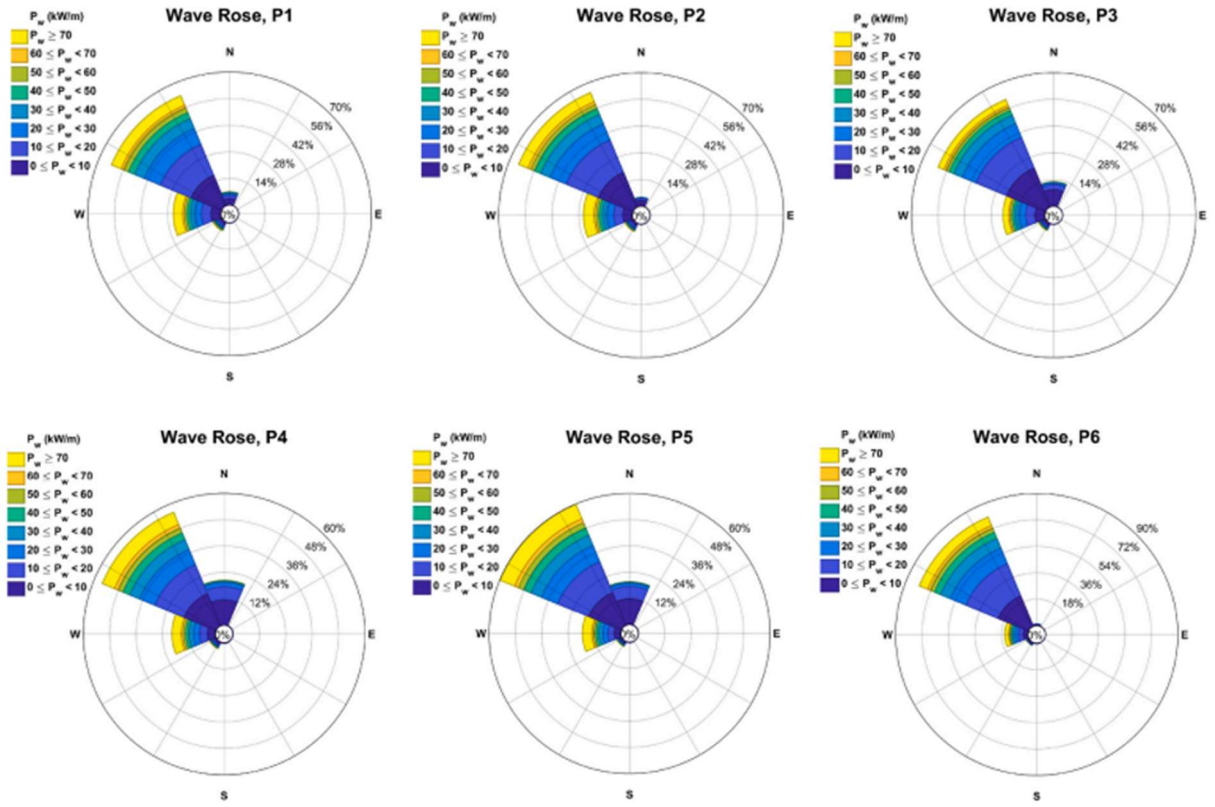


Figure 51 - Wave roses for PAER zones [9].

Figure 51 illustrates that incident wave energy arrives along a relatively narrow angular band for most of the year. This directional persistence is advantageous for wave energy array design, as it allows device layouts to be optimised relative to a prevailing wave direction rather than assuming fully isotropic forcing. At the same time, this dominant wave direction differs from the primary wind sector, introducing directional decoupling that becomes important for co-located system design [9].

Overall, the Portuguese offshore wave climate is characterised by:

i. high mean wave power density (25 - 35 kW/m) in northern and central regions;
ii. strong seasonal variability with winter-dominated energy production;
iii. exposure to severe extreme sea states (H_s >10 m);
iv. a persistent northwest-west directional regime associated with Atlantic swell.

These attributes confirm that mainland Portugal offers substantial wave energy potential but with markedly higher temporal variability and more severe extremes than offshore wind. This combination establishes wave energy as a high yield yet strongly intermittent resource, whose integration with floating offshore wind motivates the co-located system concept examined in subsequent sections.

4.1.5. Wind-Wave Complementarity

Beyond the individual magnitudes of offshore wind and wave resources, their temporal and directional relationships strongly influence the feasibility and performance of co-located energy systems. Along the

Portuguese continental coast, wind and wave energy are not perfectly synchronous, since atmospheric wind events typically precede the arrival of swell-generated wave energy due to the finite propagation time of ocean waves from their generation regions in the North Atlantic [9], [100]. This physical delay introduces a systematic offset between wind and wave power peaks, commonly referred to as wind-wave time lag.

A first indicator of instantaneous coupling between wind and wave power is the zero-lag cross-correlation coefficient (CCT(0)), which quantifies how strongly both resources vary simultaneously in time [27]. Values close to 1 indicate strong synchrony, while lower values indicate increasing temporal decoupling. For co-located systems, lower CCT(0) values are favourable because they imply that wind and wave production peaks rarely coincide, reducing instantaneous power congestion and improving utilisation of shared electrical infrastructure.

Spatial distributions of CCT(0) across the PAER zones reveal generally weak instantaneous correlation along most of the Portuguese offshore, as presented in Figure 52.

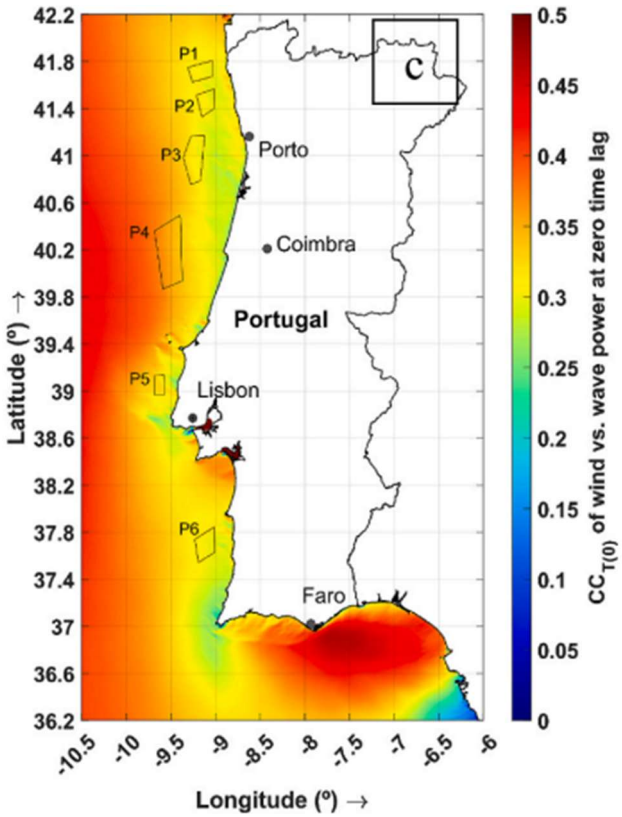


Figure 52 - Spatial distribution of CCT(0) along mainland Portugal [27].

Figure 52 shows that CCT(0) values remain moderate to low across the northern and central offshore regions, indicating limited simultaneous peaking of wind and wave power. This behaviour reflects the swell-dominated nature of the wave climate, in which wave energy impacting the coast is often generated remotely and arrives with a delay relative to local wind forcing.

Temporal complementarity is further characterised through lagged cross-correlation analysis, which evaluates how wind and wave power align when time shifts are introduced. The cross-correlation at a specific time, denoted c_τ, provides a quantitative measure of the typical propagation delay between wind generation and wave arrival, and is illustrated in Figure 53 for the PAER zones [9].

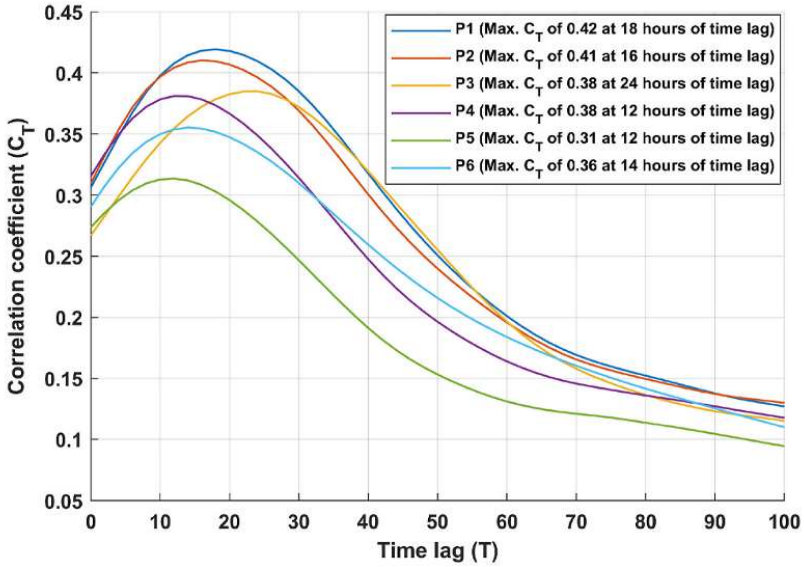


Figure 53 - cτ distributions for PAER zones [9].

Figure 53 indicates that the maximum correlation generally occurs at positive lags of several hours, confirming that wave energy systematically follows wind energy. This delayed response contributes to the smoothing of aggregated offshore generation, as periods of declining wind output may be partially compensated by incoming swell, while storm-driven wind peaks are followed by high wave energy, extending high-production windows over longer durations [9].

Directional complementarity further conditions array behaviour. As shown in Sections 4.2 and 4.3, dominant wind and wave energy sectors along the Portuguese coast are not collinear, reflecting the combined influence of regional atmospheric circulation and long-period Atlantic swell [9]. This angular separation implies that aerodynamic wake development and hydrodynamic energy flux act along different principal axes. From an engineering perspective, this decoupling requires multidirectional layout strategies, since turbine spacing optimised for wake mitigation cannot be assumed to align with optimal WEC positioning.

To synthesise wind magnitude, wave magnitude, and their temporal relationships into a single planning metric, a co-location feasibility level (CLF) was introduced, integrating resource strength, variability, and correlation characteristics. The CLF value for each zone is presented in Figure 54.

To synthesise wind magnitude, wave magnitude, and their temporal relationships into a single planning metric, a co-location feasibility level (CLF) index was introduced, integrating resource strength, variability, and correlation characteristics [9]. In the formulation adopted by *Majidi et al.* [9], the CLF is computed as a weighted combination of normalised indicators for mean wind power density, mean wave power, instantaneous cross-correlation at zero-time lag, and the standard deviations of wind and wave power [9]. The CLF index is expressed as:

$$CLF_i = \alpha W_p \frac{W_{pi} - W_{pmin}}{W_{pmax} - W_{pmin}} + \alpha P_w \frac{P_{wi} - P_{wmin}}{P_{wmax} - P_{wmin}} + \alpha_{c\tau\,(0)} \left(1 - \frac{c\tau\,(0)_{max} - c\tau\,(0)_i}{c\tau\,(0)_{max} - c\tau\,(0)_{min}}\right) + \\ + \alpha_{\sigma Wp} \left(1 - \frac{\sigma Wp_{max} - \sigma Wp_{min}}{\sigma Wp_{max} - \sigma Wp_i}\right) + \alpha_{\sigma Pw} \left(1 - \frac{\sigma Pw_{max} - \sigma Pw_{min}}{\sigma Pw_{max} - \sigma Pw_i}\right) \quad (10)$$

where the subscript i denotes the value at the analysed location, and the subscripts min and max correspond to the minimum and maximum values across the set of candidate locations considered in the assessment [9]. The coefficients α are weighting factors that reflect the relative importance assigned to each term. In Majidi et al. [9], the weights are set to 0.35 for α_{Wp} and α_{Pw}, 0.20 for $\alpha_{c(0)}$ and 0.05 for $\alpha_{\sigma Wp}$ and $\alpha_{\sigma Pw}$, following the weighting scheme proposed by *Astariz and Iglesias* [73].

This construction increases CLF for locations combining high mean wind and wave power, low instantaneous correlation and lower variability (smaller σ) in both resources, thereby favouring sites expected to deliver stronger and more stable combined generation potential [9].

The CLF value for each zone is presented in Figure 54.

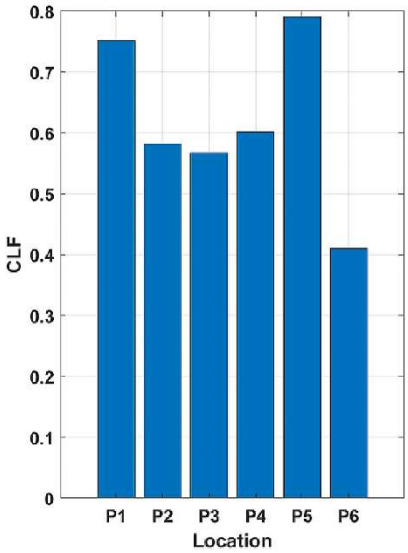


Figure 54 - CLF distribution for PAER zones [9].

Figure 54 shows the CLF values obtained for the six PAER concession zones, providing an integrated indicator of co-location suitability that combines resource magnitude, variability and the degree of temporal coupling between wind and wave conditions [9]. The results indicate that the highest CLF values are found in the northern and central offshore sectors, reflecting locations where strong individual resources coincide with favourable temporal decoupling, and therefore emerge as particularly promising candidates for co-located floating wind-wave developments from a metocean perspective. Nevertheless, the ranking differences between the central zones are moderate, and zones such as P3 still present competitive feasibility levels when considered alongside practical project drivers, like grid connection, port proximity and concession constraints, that are addressed later in the case-study definition.

To further illustrate the seasonal behaviour and relative variability of wind and wave resources, monthly boxplots of wind power density and wave power density are presented in Figure 55. This representation provides a compact statistical summary of central tendency (median), dispersion (interquartile range), and the occurrence of extreme events throughout the year, enabling a direct visual comparison of the temporal structures governing both energy streams [9]. Figure 55 highlights that the wind resource exhibits comparatively smaller seasonal modulation, with monthly medians and spreads remaining within a narrower band, whereas wave power displays a markedly stronger seasonal cycle, characterised by higher winter-time values and broader dispersion associated with energetic storm events.

This contrast supports the complementarity rationale discussed in this chapter where periods of enhanced wave power availability tend to coincide with months when wind conditions do not necessarily

peak to the same extent, suggesting that co-located exploitation can contribute to a more balanced utilisation of offshore transmission and shared infrastructure across the annual cycle.

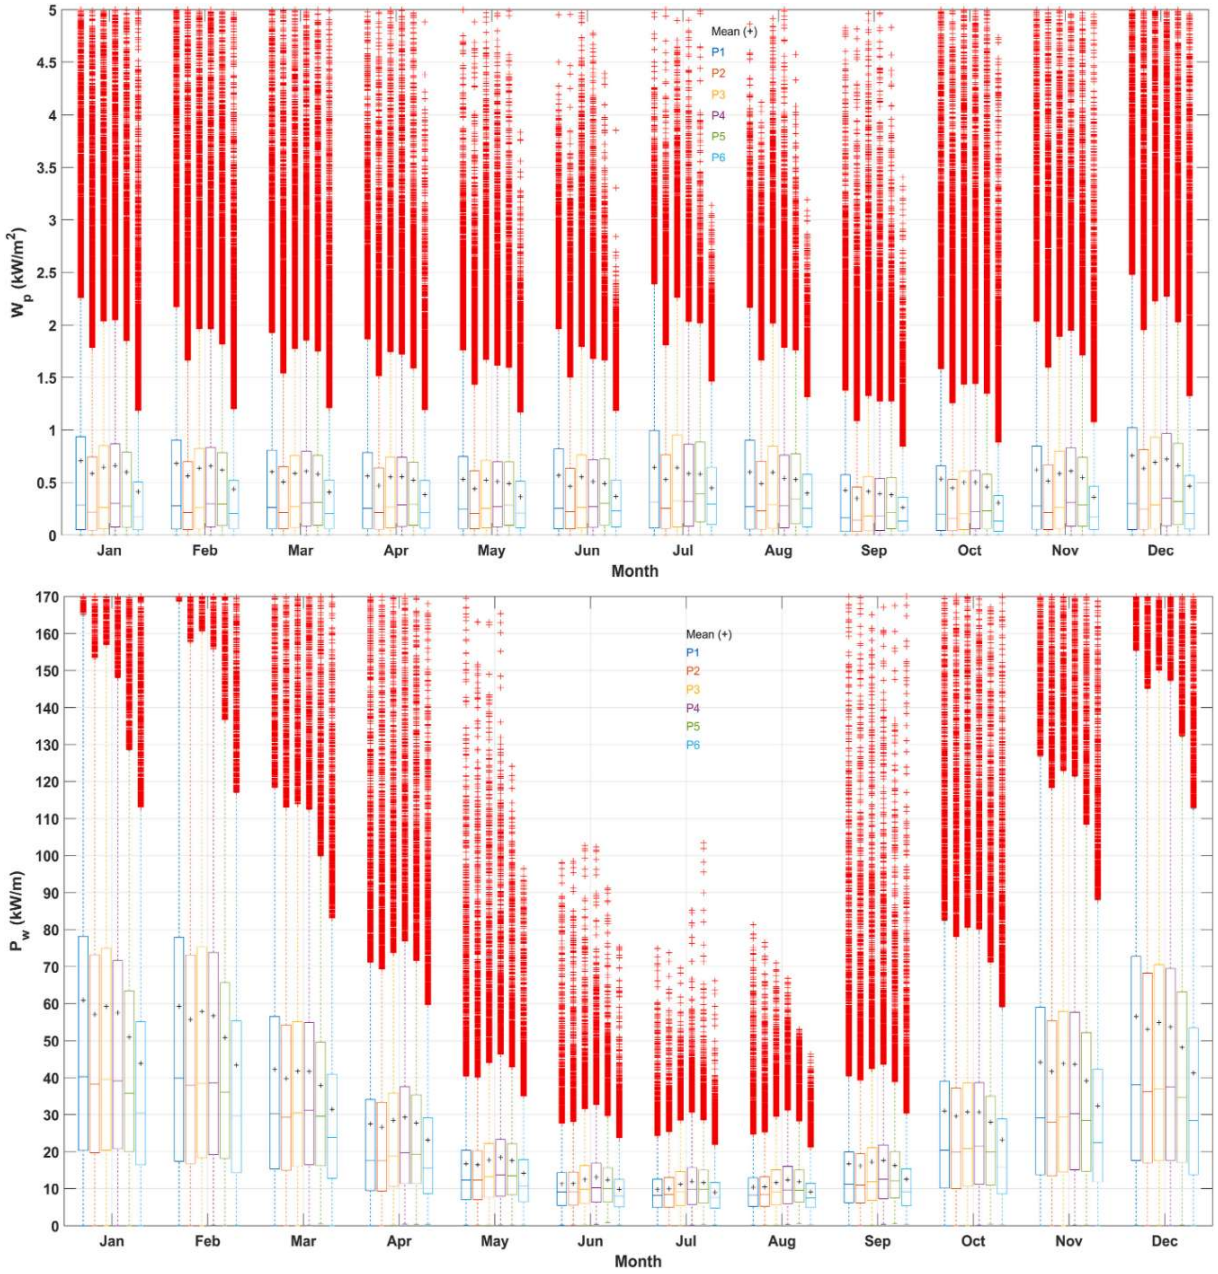


Figure 55 - Boxplot illustrating the monthly variations of wind and wave power densities for the PAER zones [9].

The wind power density boxplots in the upper panel show a moderate seasonal modulation, with higher central tendencies during winter months and lower values during summer. However, the interquartile ranges remain relatively compact throughout the year, indicating a comparatively stable offshore wind regime in statistical terms [9]. In contrast, the wave power density boxplots in the lower panel reveal a much stronger seasonal contrast, characterised by substantially larger dispersion and more frequent extreme values during winter months, reflecting the storm-driven character of the North Atlantic swell climate [9].

Taken together, the two panels highlight clear differences in the variability structure of the two resources. While offshore wind provides a relatively consistent baseline contribution, wave energy availability is more episodic and concentrated in energetic winter periods. This difference in temporal structure further reinforces the potential benefits of co-location, as combining both resources can contribute to a more balanced utilisation of shared offshore infrastructure and improve the stability of aggregated energy production over the annual cycle [9].

Overall, the Portuguese offshore environment exhibits weak instantaneous wind-wave correlation, a systematic positive time lag between wind and wave peaks, and distinct dominant forcing directions [9]. These characteristics provide a strong physical basis for co-located energy systems capable of increasing energy density per unit area while smoothing aggregate production profiles, motivating the subsequent narrowing of analysis toward individual PAER concession zones.

4.1.6. Operational Viability and Maintenance Weather Windows

High energy availability alone does not guarantee the feasibility of offshore renewable energy systems. From an engineering and economic standpoint, operational and maintenance constraints play a decisive role in determining lifecycle performance and cost-effectiveness. For floating offshore wind turbines, operational windows are defined by wind speeds between typical cut-in and cut-out thresholds, approximately 3 m/s and 25 m/s, respectively [27].

Across the Portuguese offshore concession zones, wind operational availability typically ranges from 88% to 96% of annual hours, reflecting the persistence of moderate wind conditions and the relatively low frequency of extreme wind events [27]. This indicates that wind conditions rarely impose operational limitations on floating wind turbines in Portuguese waters.

Wave energy converters exhibit similarly high operational availability, as offshore wave climates remain within typical device-specific operational limits for more than 95% of the year [27]. However, the distinction between operational availability and maintenance accessibility is critical. O&M weather windows, defined by sea states suitable for safe personnel transfer and installation activities, are typically limited to periods when significant wave height remains below approximately 1.5 m [73]. In highly energetic offshore regions, such conditions are available for only 10% to 20% of annual hours, whereas more sheltered southern nearshore zones may experience availability exceeding 30% [27].

This introduces a fundamental engineering trade-off. Sites that maximise energy yield simultaneously impose more restrictive maintenance conditions, increasing logistical complexity and operation and maintenance costs. These constraints must therefore be explicitly incorporated into optimisation frameworks targeting realistic offshore deployment scenarios.

To further assess the practical operability of offshore renewable energy systems along the Portuguese coast, Figure 56 presents spatial distributions of three operational indicators derived from the metocean dataset [27]:

a) the percentage of time during which wind speeds fall within the typical operational range of offshore wind turbines (3 - 25 m/s);
b) the percentage of time during which wave conditions remain within operational limits for wave energy converters (H_s = 1 - 8 m) and (T_p = 4 - 17 s);
c) the availability of O&M weather windows defined by H_s below 1.5 m.

Together, these indicators provide a spatial overview of the environmental conditions governing both energy production and accessibility for offshore maintenance activities across the PAER concession zones.

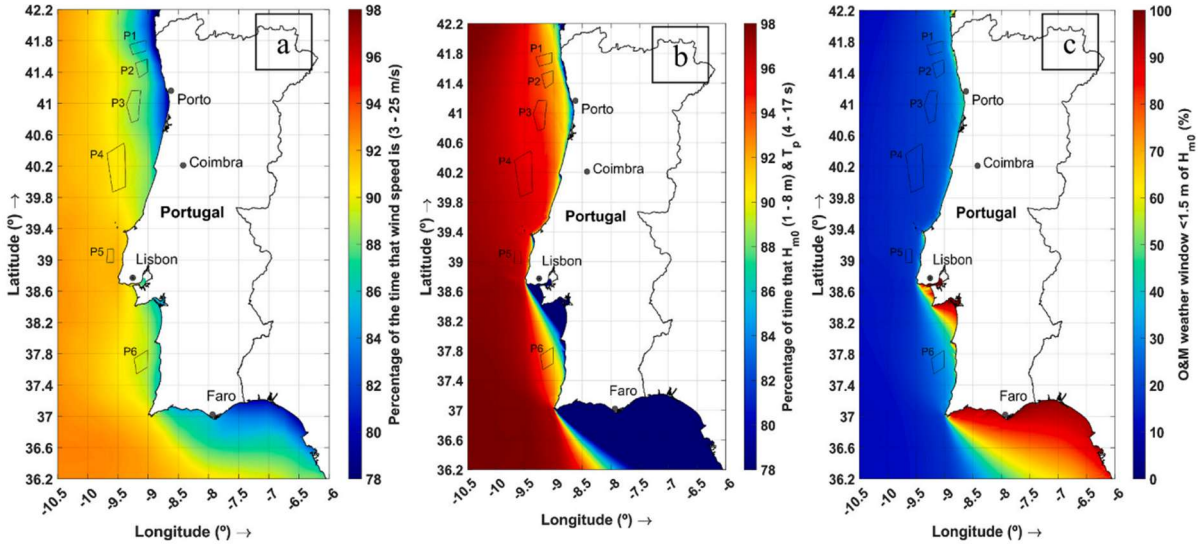


Figure 56 - The spatial distributions of the operational time of OWTs, the operational time of WECs, and the O&M weather window, from left to right, respectively [9].

These spatial patterns provide the practical motivation for including operational viability indicators in the site selection process. While the northern and central offshore sectors generally offer stronger wind and wave resources, they also tend to present reduced accessibility windows for offshore operations, whereas more sheltered southern areas provide improved O&M access but at the expense of lower resource intensity [27]. The maps in Figure 56 therefore support the need for a balanced selection that combines high energy potential with a workable maintenance regime, consistent with realistic project development constraints.

In the context of this work, the P3 concession zone is selected as the case-study area because it represents a technically meaningful compromise between these competing considerations. P3 lies within the northern-central offshore sector characterised by strong wind and wave resource levels (Sections 4.1.3 to 4.1.5), while still maintaining non-negligible operational availability and O&M-accessible weather windows compared with the most energetic northern extremes [27]. In addition, P3 benefits from favourable proximity to established coastal infrastructure in the Porto/Leixões region, as discussed later, supporting the logistical requirements of installation and maintenance. For these reasons, P3 is adopted as the reference zone for the detailed project-scale definition and optimisation analysis developed in the following sections.

4.2. Zone P3 Specific Characterization

4.2.1. Geographic Definition and Physical Constraints

PAER Zone P3, Leixões, is located offshore of the Porto metropolitan region and is formally defined in the PAER as a polygon referenced in the GCS WGS84 coordinate system [100]. The concession zone covers an offshore surface area of 722 km² and extends between a minimum distance to the coastline of approximately 31 km and a maximum distance of approximately 61 km, placing it fully beyond the

nearshore zone while preserving practical access to the Port of Leixões and surrounding industrial infrastructure, enabling the development of up to 2.5 GW of installed capacity [100].

The polygon geometry provides a continuous and spatially coherent administrative development area, preserving cable routing, mooring-field coherence, and internal navigation corridors at concession scale. However, for project-scale engineering definition, internal conceptual subdivision may still be introduced in subsequent sections to define realistic development blocks within the broader polygon. Its elongated north-south orientation is also aligned with the dominant atmospheric circulation along the Portuguese west coast, which favours layouts oriented perpendicular to prevailing northerly winds and supports efficient turbine row alignment relative to wake development [9].

Table 6 presents the PAER vertices coordinates, and Figure 57 the official polygon.

Table 6 - Vertices coordinates.

Vertex	Latitude	Longitude
1	41º 13’ 27.33’’N	9º 17’ 00.78’’W
2	41º 13’ 29.61’’N	9º 09’ 10.82’’W
3	40º 51’ 15.07’’N	9º 08’ 46.75’’W
4	40º 47’ 06.13’’N	9º 08’ 46.31’’ W
5	40º 47’ 06.18’’N	9º 16’ 43.24’’W
6	41º 01’ 31.95’’N	9º 22’ 18.51’’W

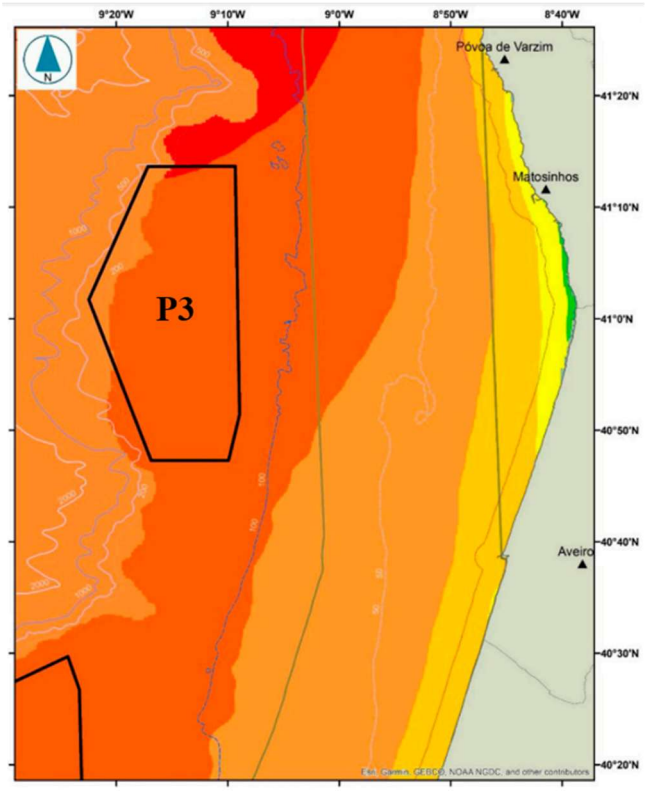


Figure 57 - PAER Zone P3 polygon [100].

From a bathymetric perspective, P3 lies entirely within depth ranges incompatible with bottom-fixed foundations, with most of the zone falling at depths around 150 m, situating the entire concession within the operational envelope of floating offshore wind systems, particularly favourable for semi-submersible and spar-type platforms combined with conventional catenary mooring systems [9], [100].

To characterise seabed gradients and identify potential installation constraints, the PAER provides a slope map for Leixões, presented in Figure 58.

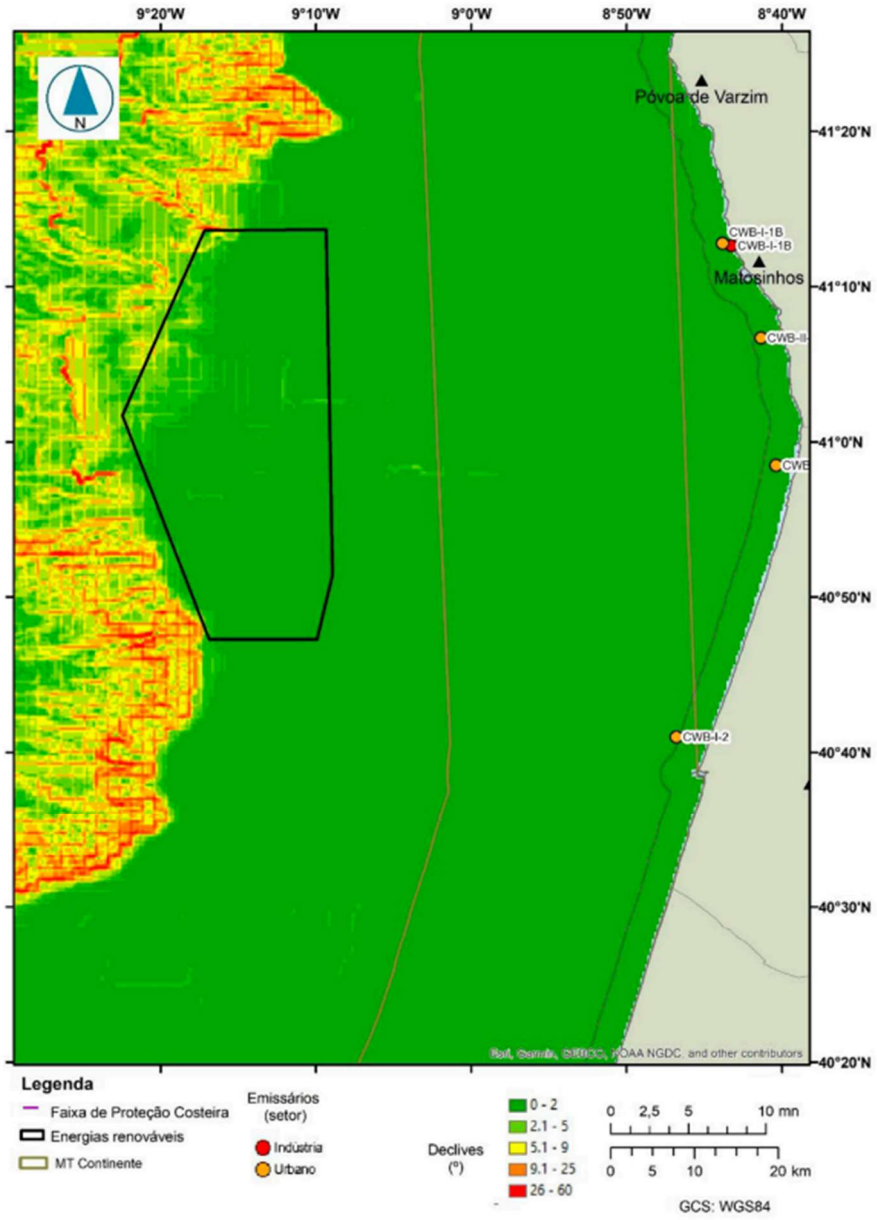


Figure 58 - Site slope map [100].

As illustrated in Figure 58, the majority of P3 is characterised by relatively gentle seabed slopes, although localised regions of increased gradient are present. While slope alone does not determine technical feasibility, steeper gradients can affect anchor embedment performance, mooring-line touchdown behaviour, and installation complexity. Consequently, slope distributions represent a first-order spatial constraint that must be considered when defining feasible layout domains for floating

systems [27]. Seabed composition further conditions floating deployment strategies. According to the PAER seabed classification, approximately 74% of P3 consists of muddy sand, 20% of sand, and 6% of rock or gravel, indicating predominantly soft sediments with smaller areas of more rigid substrate, as shown in Figure 59 [100].

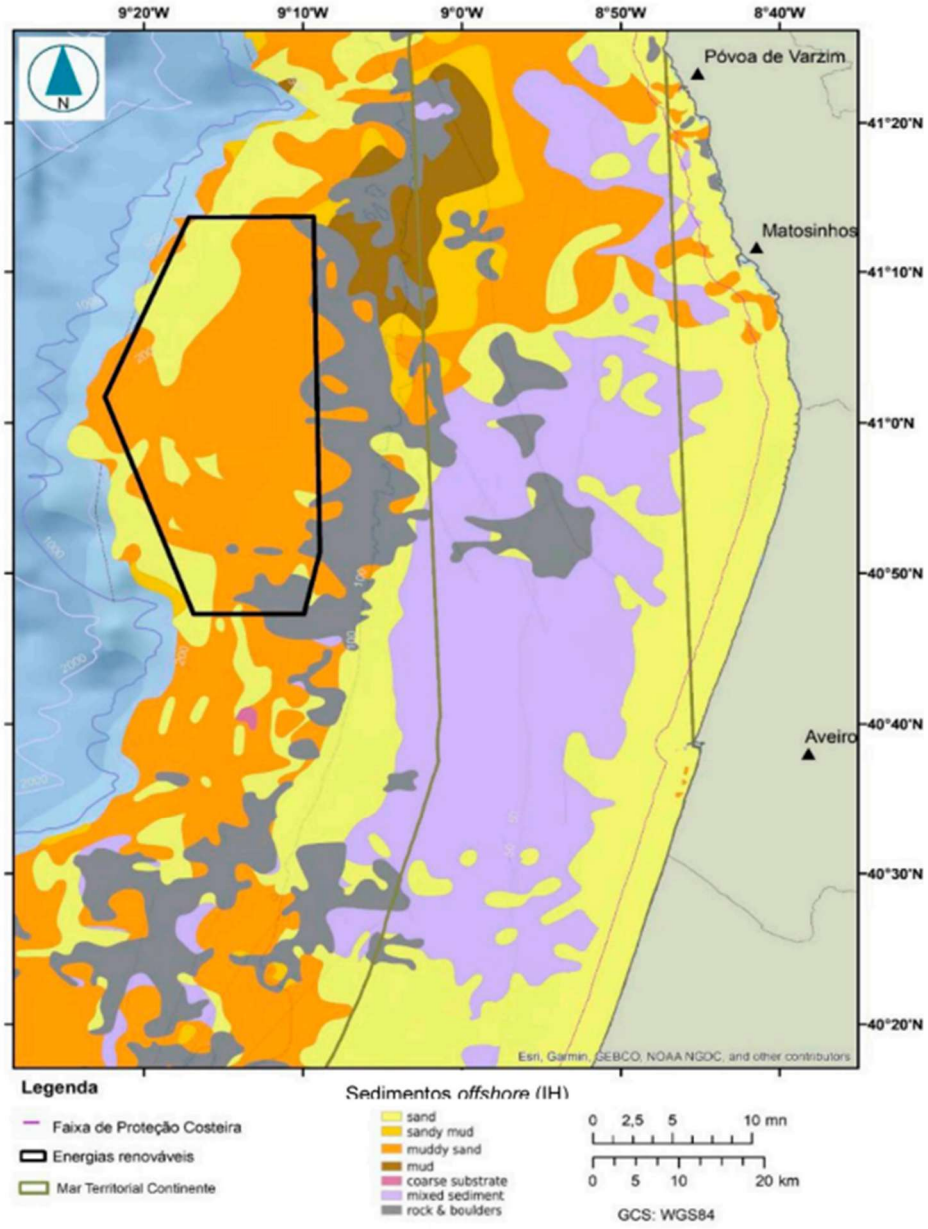


Figure 59 - Site seabed composition [100].

This heterogeneous distribution of sediment has direct implications for anchor typology and spatial planning. Soft sediment regions are generally compatible with drag anchors and suction piles, whereas rocky patches may require alternative anchoring solutions or exclusion from deployment, reinforcing the need for spatial filtering before layout optimisation [27]. Beyond purely geophysical characteristics, the PAER also identifies regulatory and environmental constraints within P3, including navigation corridors and protected areas, presented in Figure 60.

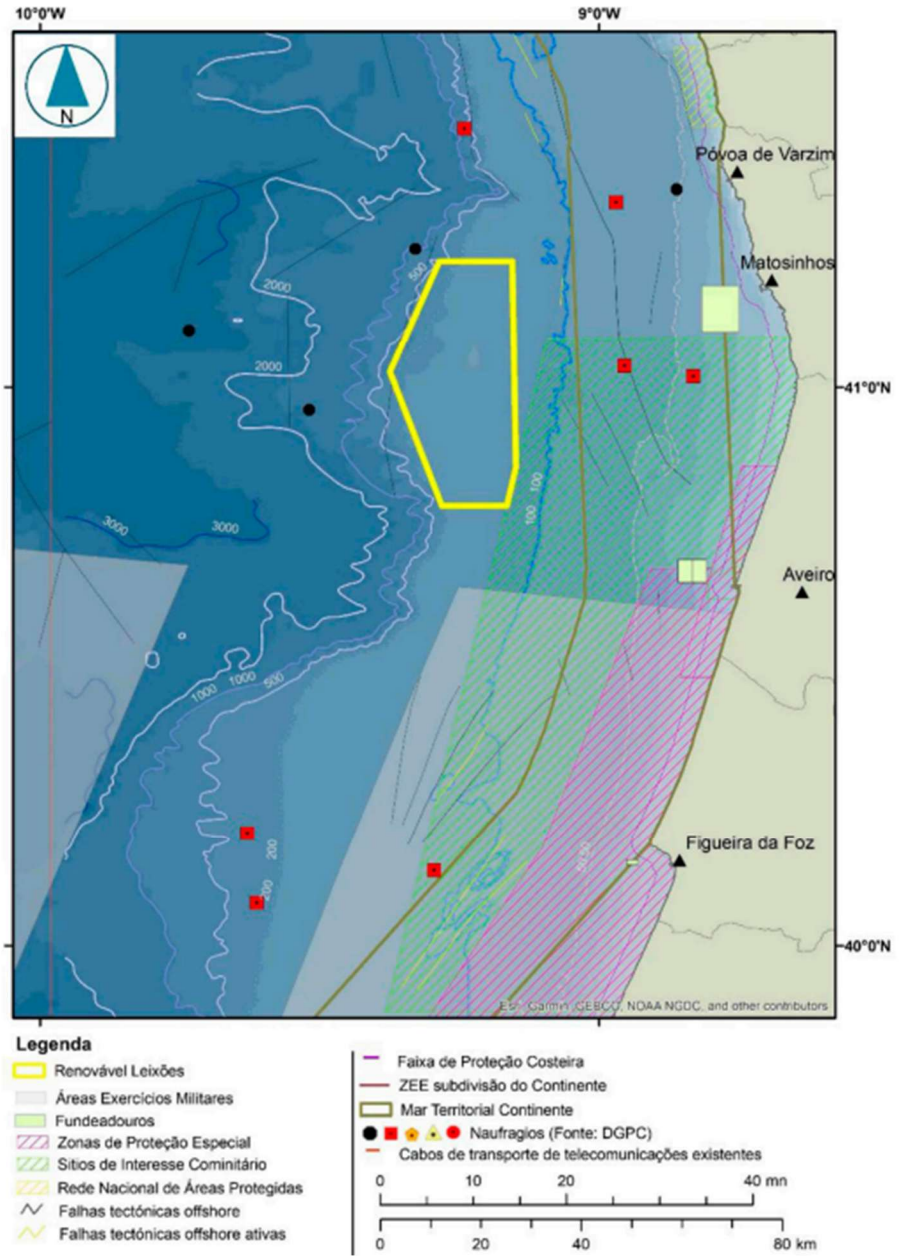


Figure 60 - Regulatory and environmental constraints [100].

Figure 60 highlights zones subject to spatial restrictions that must be excluded from any engineering layout. These planning-level constraints define the initial feasible envelope for offshore development and serve as a foundational layer upon which technical constraints, such as turbine spacing, mooring footprints, and cable safety corridors, are subsequently applied [100].

Taken together, the geometry of P3, its offshore distances, depth profile, seabed characteristics, and regulatory exclusions establish a spatially complex but technically viable environment for floating offshore energy systems. The combination of large contiguous area, favourable depth ranges, predominantly soft sediments, and proximity to port infrastructure provides a realistic physical foundation for co-located floating wind-wave deployment, while simultaneously motivating the need for optimisation-based layout design to navigate the multiple interacting spatial constraints inherent to the site [9], [27], [100].

4.2.2. OFFSHORE WIND CHARACTERIZATION

The offshore wind resource in PAER Zone P3 is consistently classified as a high-quality development area in both the national planning documentation and concession-scale quantitative assessments [9], [27], [100]. In the PAER characterisation for Leixões, the mean horizontal wind speed at 100 m hub height is reported to range between 7.51 m/s and 8.0 m/s [100]. This interval is favourable for offshore wind exploitation because it lies well above typical turbine cut-in speeds and in a regime where modern offshore turbines operate frequently in their efficient aerodynamic range, supporting high-capacity factors and stable annual energy production [9], [100].

To support this site-scale spatial characterisation, the PAER includes a wind map for the Leixões concession, as illustrated in Figure 61.

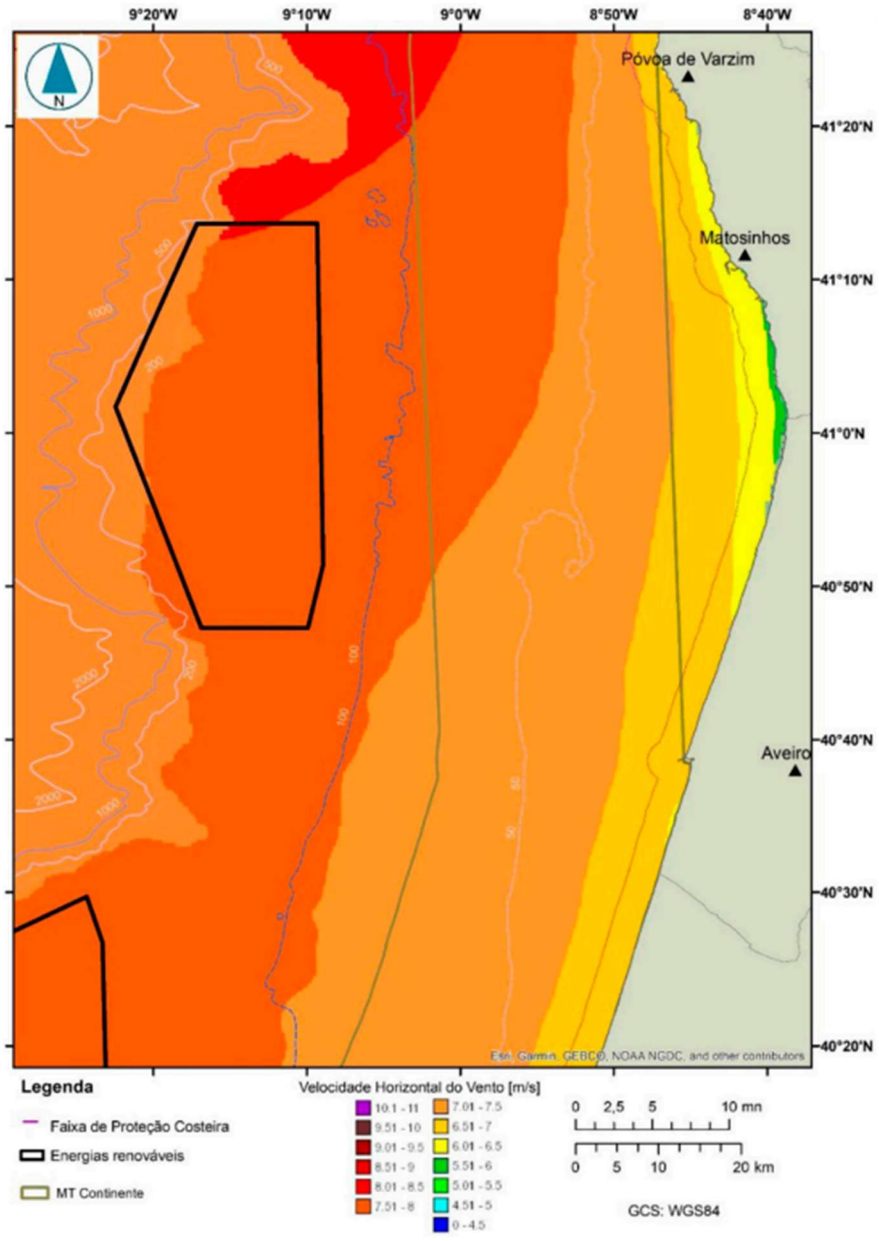


Figure 61 - Spatial distribution of offshore wind speed for zone P3 [100].

Figure 61 shows that wind speed levels are spatially coherent across most of the P3 polygon, which is important at the array scale. A wind farm layout can be developed across a broad portion of the concession without being constrained to a narrow "high-wind corridor" [100].

In addition to wind speed, the PAER reports the incident W_p for P3 as 500 to 600 W/m² [100]. This range aligns closely with the concession-scale mean W_p of 0.56 kW/m² reported for P3 by *Majidi et al.* [27], placing Leixões among the highest-ranked Portuguese offshore zones in terms of average wind energy availability.

Expected annual productivity is reflected in the reported NEPS values of 3600 h/year to 3900 h/year, indicating that turbines installed in P3 would generate the equivalent of rated power output for nearly four thousand hours annually under the local wind regime. This level of equivalent full-load operation supports commercially viable offshore wind deployment without requiring optimistic turbine assumptions [100].

A key reason P3 is particularly attractive for wind-wave system integration is not only the strength of its wind resource, but also its temporal stability, quantified by the MVI. As discussed in Section 4.1.3., lower MVI values indicate reduced seasonal concentration and more uniform month-to-month energy availability. *Majidi et al.* [27] report an MVI of wind power of 0.18 for P3, Figure 46, one of the lowest among all PAER zones. This means that the wind resource in Leixões is comparatively stable over the year, a desirable trait for grid export utilisation and for reducing uncertainty in annual production profiles [27].

Majidi et al. [9] reinforce this stability interpretation by analysing the monthly distributions of wind power density across the six zones, as illustrated in Figure 55. On that figure, P3 exhibits a consistent distribution shape across months, with winter intensification but without extreme dispersion compared to other zones [9]. This confirms that Leixões combines strong mean values with a comparatively regular seasonal structure, which is particularly relevant for this work because the optimisation relies on a realistic, representative case where production is not dominated by rare extremes but by persistent usable conditions.

For completeness, extreme wind exposure at the site is characterised by a maximum wind speed of nearly 32.0 m/s. This parameter is relevant from a survivability and operational envelope standpoint, since it provides a concession-scale indication of extreme wind exposure that offshore wind turbines and floating support structures must tolerate [27].

Finally, beyond magnitude and seasonality, directional structure is essential for realistic offshore wind farm design because wake development and optimal turbine row orientation depend strongly on prevailing wind direction [9]. P3's directional characteristics follow the dominant northerly wind sector identified in Section 4.1.3., Figure 47, associated with the Iberian coastal jet, with secondary westerly contributions during winter storm periods [9]. This persistent directional structure supports conventional offshore array orientations aimed at mitigating wake losses and provides a physically consistent basis for defining reference layouts in subsequent optimisation stages.

Overall, the offshore wind resource in PAER Zone P3 combines:

i. strong hub-height wind speeds (7.51 - 8.0 m/s);
ii. high incident wind power flux (500 - 600 W/m²);
iii. high equivalent full-load hours (3600 - 3900 h/year);
iv. high mean wind power density (0.56 kW/m²);
v. relatively low month-to-month variability (MVI = 0.18);
vi. a persistent directional regime.

These attributes confirm Leixões as a robust offshore wind site characterised by substantial resource magnitude, favourable temporal stability, and clear directional forcing, providing a technically sound foundation for the co-located floating wind-wave case study.

4.2.3. OFFSHORE WAVE CHARACTERIZATION

The offshore wave climate within PAER Zone P3 reflects its direct exposure to North Atlantic swell systems and is consistently classified among the most energetic along the Portuguese continental shelf [9], [27], [100]. At concession scale, the mean P_w for P3 reaches 32.16 kW/m [27]. As discussed in Section 4.1.4, P_w represents the rate at which wave energy propagates through a unite crest length and is commonly used to characterize the theoretical energy production of a wave climate. Values exceeding 30 kW/m are characteristic of energetic Atlantic environments and place Leixões firmly within the upper tier of Portuguese offshore wave resource zones.

To illustrate how this resource is distributed across the concession area, the PAER provides a spatial mapping of Pw for P3, shown in Figure 62 [100].

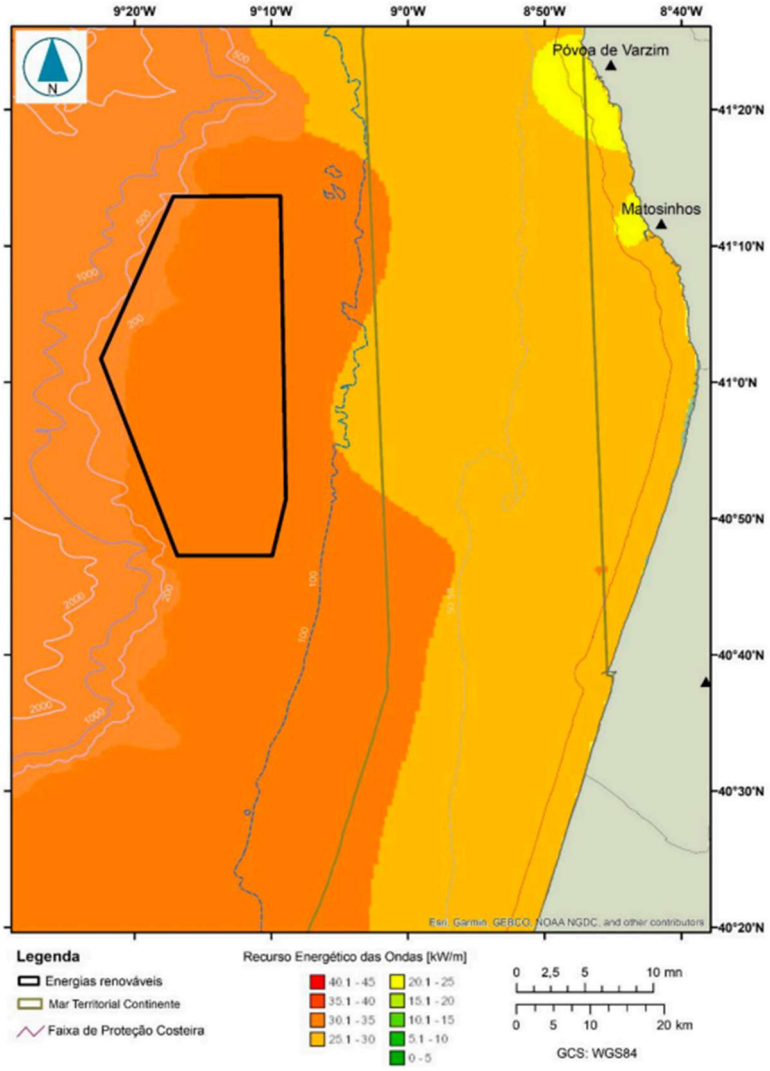


Figure 62 - Spatial distribution of offshore wave intensity for zone P3 [100].

Spatially, the PAER wave energy mapping illustrated in Figure 62 shows that high wave power density levels extend across most of the P3 concession area rather than being confined to narrow high-energy corridors [100]. Values within the polygon fall predominantly within the 30 - 35 kW/m range, with slightly higher levels toward the western boundary where the site is more directly exposed to Atlantic swell propagation. This spatial coherence indicates that the wave resource is relatively uniform at the concession scale, meaning that WEC placement is unlikely to be constrained by strong intra-zone gradients in resource magnitude and can instead be guided primarily by engineering, bathymetric and spatial planning considerations.

Extreme wave exposure at P3 is characterised by a maximum H_s of 11.17 m, observed in the hindcast dataset analysed by *Majidi et al.* [27]. This value represents the highest sea state recorded within the analysed period and provides a concession-scale indicator of survivability under severe Atlantic storm conditions, implying that both WECs and floating wind platforms must be designed to withstand such events. Such extremes influence platform sizing, mooring system design, and operational strategies, but remain within the envelope addressed by contemporary floating offshore systems.

Directional characteristics of the wave climate follow the dominant northwest-to-west sector identified in Section 4.1.4., Figure 51. This persistent swell direction reflects long-period Atlantic wave propagation toward the Iberian margin and results in a relatively narrow angular band of incoming wave energy. From an array-design perspective, this directional consistency is advantageous because WEC layouts can be optimised relative to a prevailing wave direction rather than assuming fully isotropic forcing [9].

In summary, the wave resource in PAER Zone P3 is defined by:

i. high mean wave power (32.16 kW/m);
ii. substantial seasonal variability (MVI = 1.49);
iii. severe but realistic extreme conditions ($H_{s,max}$ = 11.17 m);
iv. a dominant NW-W directional regime.

These characteristics confirm that Leixões offers substantial wave energy potential under a highly energetic Atlantic climate, while simultaneously presenting realistic variability and extreme-event conditions that must be accounted for in the design of co-located floating wind–wave systems.

4.2.4. Wind-Wave Complementarity

The temporal and directional interaction between wind and wave resources in PAER Zone P3 provides a key physical basis for co-located offshore energy deployment.

Instantaneous coupling between wind and wave power, quantified through the CCT(0) reaches 0.31 for P3 [27]. This relatively low value indicates weak synchrony between wind and wave production peaks, meaning that high wind power and high wave power rarely occur simultaneously. Such behaviour is favourable for co-located systems, as it reduces instantaneous power congestion and improves the effective utilisation of shared export infrastructure.

Temporal complementarity is further characterised by lagged correlation analysis. For P3, the maximum correlation coefficient reaches approximately $c\tau = 0.38$ at a lag of about 24 hours, shown in Figure 53, [9]. This reflects the swell-dominated nature of the Portuguese west coast, where energetic wave systems affecting Leixões are typically generated by wind events occurring offshore in the North Atlantic and arrive at the site after a one-day propagation delay. This delayed response directly contributes to smoothing aggregated offshore generation.

Directional complementarity further conditions array design in P3. As established in Sections 4.1.3. and 4.1.4., dominant wind forcing is associated primarily with northerly sectors, whereas wave energy arrives predominantly from the northwest-west. This angular separation implies that aerodynamic wake development and hydrodynamic energy flux act along different principal axes. Realistic array design in P3 must therefore accommodate oblique forcing conditions, reinforcing the need for optimisation frameworks capable of handling coupled but directionally distinct wind-wave interactions [9].

To synthesise multiple complementarity indicators into a single planning metric, as mentioned in Section 4.1.5, the CLF, is used. For P3, this index reaches approximately 2.0, placing Leixões among the highest-ranked Portuguese offshore zones for combined wind-wave development [9]. This classification reflects the simultaneous presence of strong individual resources and favourable temporal decoupling.

Taken together, PAER Zone P3 is characterised by:

i. weak instantaneous wind-wave coupling (CCT(0) = 0.31);
ii. a clear 24-hour lag between wind and wave peaks ($c\tau = 0.38$);
iii. distinct dominant directional sectors for wind and wave forcing;
iv. elevated co-location feasibility (CLF = 2.0).

These attributes indicate a site where wind and wave resources are both individually strong and temporally complementary, creating favourable conditions for co-located offshore energy systems that seek to increase energy density per unit area while smoothing aggregate power output.

4.2.5. Logistical Accessibility and Operational Feasibility

Beyond resource availability and wind-wave complementarity, the practical viability of offshore developments in PAER Zone P3 is governed by logistical distances, environmental accessibility, and the capacity to sustain reliable operation and maintenance under energetic Atlantic conditions [27].

For P3, the mean distance to major ports is 46.38 km, while the mean distance to grid connection points is 75.92 km [27]. These distances define baseline expectations for offshore transit times, export cable lengths, and electrical connection complexity. Although not minimal, both values fall within ranges routinely addressed by existing offshore wind developments and are compatible with regular marine operations supported by the Port of Leixões and surrounding industrial infrastructure [27], [100].

Environmental availability provides a first-order estimate of how frequently offshore systems can operate under prevailing metocean conditions. As presented in Figure 56(a), for P3, the operational time of offshore wind turbines is approximately 89%, indicating that turbines are expected to remain operational for nearly nine-tenths of the time when considering environmental constraints alone. Wave energy converters exhibit an even higher operational time of approximately 95%, Figure 56(b), reflecting their broader operational envelope under elevated sea states [27]. This asymmetry highlights the inherent robustness of WECs under energetic wave climates and supports their integration within floating wind farms.

Access for maintenance activities is quantified through the O&M weather window index. For P3, as illustrated in Figure 56(c), this index reaches approximately 20% [27]. This indicates that Leixões combines vigorous resource intensity with acceptable accessibility levels when evaluated at concession scale.

To integrate resource strength, temporal variability, extreme environmental conditions, logistical distances, and operational accessibility into a single planning indicator, *Majidi et al.* [27] developed a

composite suitability index (CSI) based on a multi-criteria decision-making framework. In total, seventeen parameters were considered, including wind and wave resource intensity, variability metrics, extreme conditions, distance to ports and grid infrastructure, and operational accessibility. Each parameter is first normalized and then weighted according to its relative importance for offshore energy development. The resulting CSI is calculated as the weighted sum of the normalized criteria:

$$CSI = \sum_{i=1}^{n} w_i N_i \tag{11}$$

where N_i is the normalized value of criterion i, w_i the weight assigned to criterion i and n the total number of criteria, with the weights' sum being equal to 1. The resulting CSI values range between 0 and 1, with higher values indicating locations that simultaneously exhibit strong renewable resource availability and favourable deployment conditions. Its spatial distribution along the Portuguese coast is presented in Figure 63.

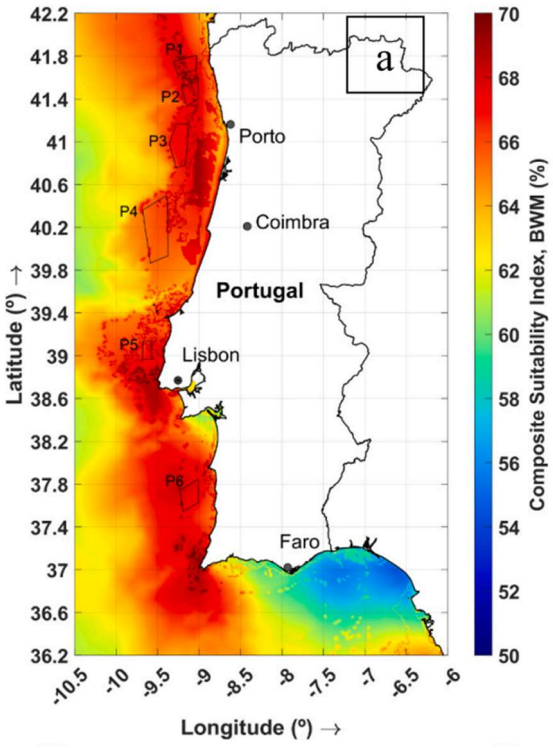


Figure 63 - Spatial distribution of CSI along mainland Portugal [27].

Figure 63 shows that the northern and central Portuguese offshore regions form statistically significant clusters of high suitability, with P3 located within one of these high-ranking corridors. This result confirms that Leixões is not only resource-rich but also embedded within a broader zone characterised by favourable technical and logistical conditions for co-located floating energy systems.

Taken together, P3 exhibits:

i. moderate distances to port (46.38 km) and grid connection (75.92 km);
ii. high operational availability for both wind turbines (88.63%) and WECs (94.52%);
iii. realistic O&M access windows under Atlantic sea states (19.94%);
iv. inclusion within a CSI-defined high-suitability offshore cluster.

These combined logistical and operational characteristics demonstrate that PAER Zone P3 represents not merely a high-energy offshore site but a realistically serviceable development area in which floating wind-wave systems can be installed, accessed, and maintained under representative North Atlantic conditions. Finally, to eliminate ambiguity and provide a single quantitative site sheet for subsequent chapters, Table 7 consolidates the complete set of key P3 descriptors.

Table 7 - Key descriptors of zone P3.

Descriptor	P3 Value	Source
Area	722 km^2	[100]
Possible installed capacity	2.5 GW	[100]
Distance to shore	31 - 61 km	[100]
Wind speed at 100 m height	7.51 - 8.0 m/s	[100]
Wind power flux	500 - 600 W/m^2	[100]
NEPS	3600 - 3900 h/year	[100]
Mean wind power density	0.56 kW/m^2	[27]
MVI of wind power	0.18	[27]
Maximum wind speed	31.16 m/s	[27]
Mean wave power	32.16 kW/m	[27]
MVI of wave power	1.49	[27]
Maximum H_s	11.17 m	[27]
Wind-wave synchrony CCT(0)	0.31	[27]
Time-lag complementarity	0.38 at 24h	[9]
Mean water depth	155.28 m	[27]
Mean seabed slope	0.30%	[27]
Distance to main ports	46.38 km	[27]
Distance to grid connection points	75.92 km	[27]
Operational time (OWTs)	88.63%	[27]
Operational time (WECs)	94.52%	[27]
O&M weather window	19.94%	[27]

Overall, Leixões combines:

i. high wind energy levels with comparatively low monthly variability;
ii. remarkably high wave energy potential with stronger seasonality;
iii. persistent directional regimes;
iv. measurable wind-wave complementarity with a 24h lag;
v. favourable bathymetric and logistical conditions compatible with floating developments.

These characteristics establish P3 as a technically robust and representative Atlantic offshore case study site for evaluating co-located floating wind-wave layouts under realistic environmental, operational, and spatial constraints.

4.3. FROM CONCESSION-SCALE P3 TO PROJECT-SCALE STUDY AREA

4.3.1. RATIONALE FOR PROJECT-SCALE SUBDIVISION

As established in the previous sections, PAER Zone P3 is defined as a single continuous concession polygon and constitutes the official administrative development area considered in this work. However, although the concession is spatially continuous at planning scale, the present case study does not assume the simultaneous occupation of the full P3 envelope. The complete polygon, with a total area of 722 km² and an indicative maximum capacity of 2.5 GW, represents the upper planning boundary of the zone rather than the direct footprint of a single project-scale optimisation case.

For the purposes of engineering definition, it is therefore necessary to distinguish between the administrative continuity of the concession and the project-scale domain required for a specific layout study. In this sense, the fact that P3 is treated as a continuous polygon does not preclude the adoption of an internal conceptual subdivision when moving from concession-scale characterization to project-scale design. The objective is not to redefine the official PAER limits, but to establish a realistic and operationally manageable development block within the broader concession area.

To clarify this transition from concession-scale characterization to project-scale definition, Figure 64 presents a conceptual internal subdivision of P3 into five representative development blocks, with the selected block highlighted for the subsequent case-study definition.

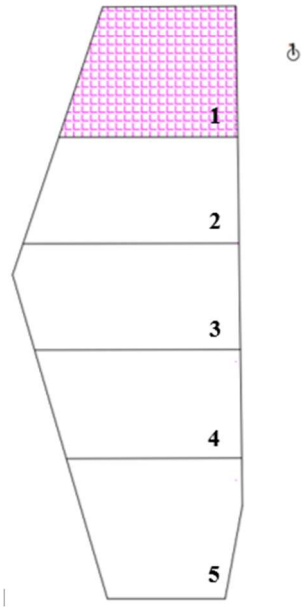


Figure 64 - Conceptual internal subdivision of PAER Zone P3 into five representative project-scale blocks, with the selected study block highlighted.

As shown in Figure 64, the proposed internal subdivision does not alter the official external boundary of P3. Instead, it provides a schematic partition of the concession into a set of equivalent internal blocks that make it possible to translate the full planning-scale polygon into a more practical project-scale interpretation. The highlighted upper block should therefore be understood as a representative

development unit adopted as the basis for the present case-study, rather than as an official regulatory subdivision of the concession.

Under this first-order conceptual interpretation, dividing the total concession area into five approximately equivalent internal blocks yields a reference area of about 145 km² per block. The same logic applied to the indicative maximum capacity of 2.5 GW results in a corresponding project-scale order of magnitude of approximately 500 MW per block. This provides a coherent engineering basis for defining a single representative development unit within P3 while preserving consistency with the broader planning characteristics established previously.

The value of this subdivision lies in creating a project-scale reference that remains compatible with the physical, spatial, and logistical characteristics already identified for P3, while also defining a manageable footprint for explicit layout development. Rather than attempting to represent the entire concession in a single optimisation problem, the case study is therefore structured around one representative internal block. This makes it possible to define a fixed floating wind farm, apply turbine-related exclusion zones, and preserve sufficient residual area for wave energy deployment within a spatial scale that remains both technically credible and computationally practical.

Accordingly, the internal subdivision adopted here should be understood as a modelling abstraction introduced to support project-scale engineering specification. The selected study area presented in the following subsection is therefore derived as one representative block within the broader concession, subsequently refined in geometric terms to accommodate the fixed wind farm layout and the internal optimisation space required for the wave energy layer.

4.3.2. Definition of the Selected Study Area

Following the conceptual subdivision established in the previous subsection, the present case study is developed within the highlighted upper block of P3, which is adopted as the basis for the project-scale engineering definition. Within that representative block, a specific study polygon was then explicitly delimited to accommodate the fixed floating wind farm layout and, simultaneously, preserve sufficient internal space for the subsequent optimisation of the wave energy layer. The selected area is therefore not an arbitrary extraction from the concession, but a deliberately defined project-scale domain whose geometry was chosen to support both the fixed wind layer and the variable WEC layer under realistic spacing constraints.

To identify this domain spatially, Figure 65 presents the selected study area and its position within the representative project-scale block adopted from the broader P3 concession.

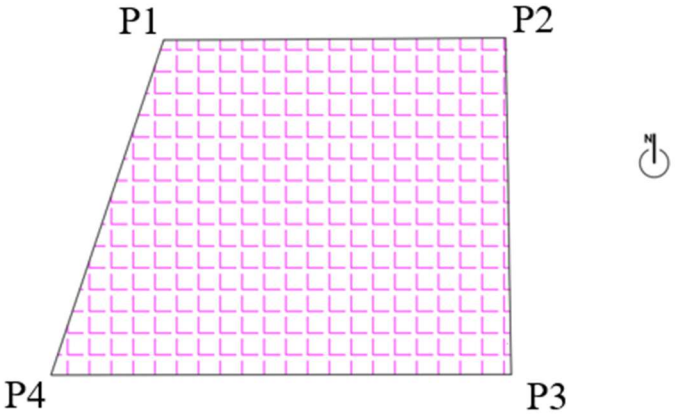


Figure 65 - Map of the selected study polygon.

As shown in Figure 65, the selected study area is defined as a four-vertex polygon with a predominantly north-south extent and a slightly widening footprint toward the southern boundary. This geometry was selected because it provides a compact but sufficiently regular project-scale envelope, allowing the reference wind farm to be positioned with internal alignment and peripheral margin while still preserving a substantial residual area for wave energy deployment. The upper and lower boundaries remain approximately aligned in the east-west direction, while the lateral edges provide the additional width required to avoid excessive confinement of the fixed wind layout in the southern portion of the domain.

The boundary of the selected study area is defined by four vertices, ordered clockwise, and should be stated explicitly to ensure full geometric reproducibility of the case study. Table 8 presents the adopted coordinates of the study polygon in WGS84 geographic coordinates, corresponding to the boundary used for the project-scale layout definition.

Table 8 - Coordinates of the selected study area.

Vertex	Latitude (DMS)	Longitude (DMS)	Latitude (Decimal Degrees)	Latitude (Decimal Degrees)
P1	41º 13’ 27.33’’N	9º 17’ 00.78’’W	41.224258	-9.283550
P2	41º 13’ 29.61’’N	9º 09’ 10.82’’W	41.224892	-9.153006
P3	41º 07’ 08.25’’N	9º 09’ 03.42’’W	41.118958	-9.150950
P4	41º 07’ 08.25’’N	9º 19’ 49.98’’ W	41.118958	-9.330550

As indicated in Table 8, the selected polygon is defined directly from its geographic coordinates, which provide the reproducible boundary reference for the case study. For geometric operations involving area and distance, however, the study area should be handled in a projected coordinate system consistent with the location of P3. In the present work, the boundary is therefore interpreted computationally in a metric projection after its definition in WGS84, so that spacing constraints, exclusion distances, and area calculations can be applied consistently in metres.

When the polygon defined by the four vertices is converted into projected coordinates and its enclosed area is calculated, the resulting surface is approximately 152.45 km². For the purposes of case-study definition, this value is adopted in rounded form as 152 km². This final area remains close to the first-order target derived from the conceptual subdivision described in Section 4.3.1, while providing a slightly enlarged and more practical envelope for explicit layout development.

This refinement is important for two reasons. First, it preserves consistency with the project-scale logic established in the previous subsection, since the final study area remains close to the 145 km² representative block obtained from the five-part conceptual subdivision of P3. Second, the slight increase in area provides additional geometric flexibility, which is necessary to accommodate the fixed wind farm arrangement, the turbine-related exclusion zones introduced later, and the residual admissible space required for meaningful WEC optimisation. The selected study polygon should therefore be understood as the final project-scale domain used throughout the remainder of the case study.

4.4. REFERENCE FLOATING WIND FARM DEFINITION

4.4.1. CAPACITY TARGET AND TURBINE CLASS

The purpose of this chapter is to translate the methodological framework developed previously into a concrete project-scale case study. To achieve this, the study first defines the host offshore wind farm and its associated spatial constraints and subsequently introduces the wave energy layer within the residual deployable domain, ultimately producing an optimisation-ready configuration for the co-located wind-wave system.

Once the project-scale study area has been defined, the next step is to establish the reference floating wind farm that will act as the fixed spatial layer of the co-located case study. In the present work, the purpose of the wind farm is not to constitute the optimisation target itself, but to provide the host infrastructure within which the wave energy layer will later be positioned. For that reason, the wind farm is introduced here as a predefined engineering configuration whose geometry and spatial footprint are fixed in advance, while the optimisation procedure is reserved for the placement of the wave energy modules.

As discussed previously in Chapter 2 and in the project-scale subdivision presented in Section 4.3.1, the intended order of magnitude for the case study was a project in the several-hundred-megawatt range, with a target close to 500 MW. This value was used as a design objective rather than as a rigid requirement, since the final wind farm definition also had to remain compatible with the available study area, the spacing logic adopted for a high-capacity turbine class, and the need to preserve sufficient residual space for the subsequent wave energy deployment. The selected wind farm should therefore be understood as the result of a controlled geometric balance between capacity, footprint, and co-location feasibility. The main design basis parameters adopted for the reference wind farm are summarised in Table 9.

Table 9 - Design basis of the reference floating wind farm.

Parameter	Adopted value	Role in case-study definition
Target project scale	~ 500 MW	Preference project-scale objective derived from P3 internal division
Selected turbine design class	20 MW class	Forward-looking high-capacity design assumption adopted for the wind layer
Reference rotor diameter	300 m	Basis for converting spacing in D into metric spacing values
Methodological role	Fixed spatial layer	Wind farm geometry is prescribed. Only the wave layer is optimised

As indicated in Table 9, the selected turbine class for the reference wind farm is the 20 MW class, adopted as a forward-looking design assumption consistent with the turbine-scaling discussion developed previously in Chapter 2. The use of this turbine class does not imply that it is treated as the current market standard, but rather that it provides a realistic high-capacity basis for defining a near-term floating offshore configuration. This is particularly relevant at project scale, since the use of larger individual turbines makes it possible to reach a substantial installed capacity with a reduced number of

units, which is advantageous when the same study area must also accommodate the internal deployment of wave energy modules.

The adoption of a 300 m rotor diameter is directly associated with this turbine class and is especially important for the geometric definition that follows. Because the spacing of the wind farm is established in multiples of rotor diameter, the selected rotor dimension becomes the reference parameter from which all metric spacing values are derived. In this sense, the turbine class is not only a capacity assumption, but also the dimensional basis that governs the spatial footprint of the host wind farm.

At this stage, the wind farm is therefore defined in terms of its intended project scale, its turbine design class, and its dimensional reference basis. The detailed geometric arrangement required to translate these assumptions into a specific number of turbines and a fixed internal layout is presented in the following subsection. That subsequent step is where the target project scale is converted into an explicit row-based wind farm configuration inside the 152 km² study polygon established in Section 4.3.2.

Accordingly, the methodological role of the wind farm is already fixed from this point onward. It constitutes the prescribed infrastructure layer of the case study. Its detailed geometry will be defined in the next subsection, but its function within the analysis remains unchanged. The wind farm is introduced as the host spatial framework that constrains the optimisation domain, whereas the decision variables of the optimisation are restricted to the wave energy modules deployed within the residual admissible area.

4.4.2. Wind Farm Geometry and Orientation

Once the design basis of the reference wind farm has been established, the next step is to translate that basis into a fixed internal layout within the study polygon. At a nominal turbine rating of 20 MW, a strict project target of 500 MW would correspond to approximately 25 turbines. In practice, however, the final turbine count could not be defined from capacity alone. It also had to remain compatible with the geometry of the selected study area, the spacing logic adopted for the wind farm, and the need to preserve sufficient residual space for the subsequent deployment of the wave energy layer. For that reason, the final number of turbines was determined through geometric layout definition rather than by forcing an exact round-number capacity target.

To clarify how these design constraints were translated into the final fixed wind layer, Figure 66 presents the adopted wind farm layout within the selected study polygon, including the in-row and inter-row spacing values used in the case-study definition.

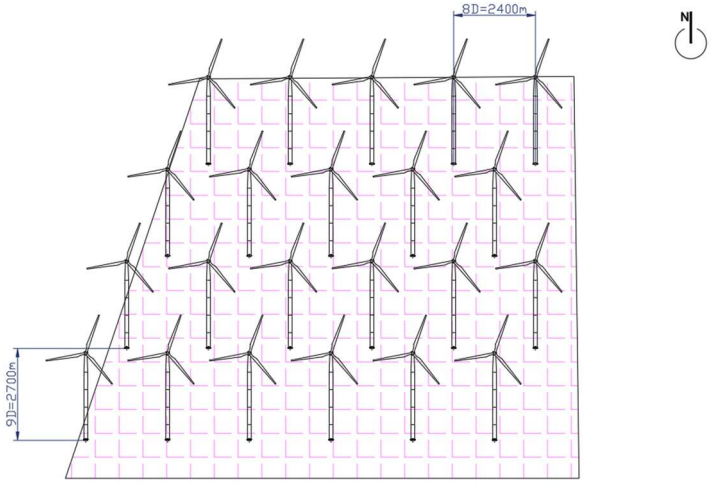


Figure 66 - Final wind farm layout.

As shown in Figure 66, the final wind farm is arranged in four rows and contains a total of 22 turbines, and the row population is defined as 5 / 5 / 6 / 6, from north to south. This arrangement results directly from the shape of the selected study area. The northern portion of the polygon is more restrictive in terms of usable horizontal extent, which limits the practical row length that can be accommodated while preserving adequate edge margins. By contrast, the southern portion of the polygon provides a wider and more permissive footprint, allowing the two downstream rows to accommodate one additional turbine each without compromising the intended spacing logic. In this way, the final layout reflects not only the target project scale, but also the spatial adaptation of the wind farm to the geometry of the selected domain.

Within each row, the turbines are spaced at 8D, corresponding to 2400 m for the adopted rotor diameter of 300 m. This provides a regular internal row structure and preserves a clear and reproducible geometric basis for the fixed wind layer. The spacing between consecutive rows is set to 9D, corresponding to 2700 m. Together, these two spacing values define the primary dimensional framework of the wind farm and establish the characteristic distances that govern its footprint inside the study polygon.

The final arrangement is not a perfectly aligned rectilinear grid. Instead, it adopts a staggered configuration, in which each row is laterally offset relative to the adjacent one by approximately half of the in-row spacing. In practical terms, this means that a turbine in a downstream row is positioned approximately between two turbines in the upstream row, rather than directly behind one of them. This offset preserves the regularity of the layout while reducing direct wake superposition along the principal downstream direction. It also allows the row-to-row spacing to remain compact without introducing a fully aligned wake corridor through the farm.

This staggered geometry is particularly important because the selected site is characterised by a dominant northerly wind regime, as established previously in Chapter 4. For that reason, the wind farm is oriented so that the rows extend approximately in the east-west direction, while the principal downstream progression of the layout occurs from north to south. Under this arrangement, the dominant wind approaches the farm approximately normal to the row direction, and wake development is distributed across the staggered downstream rows rather than being transferred directly from one turbine to another in a fully aligned sequence. The adopted orientation therefore preserves coherence between the concession-scale site characterization and the project-scale engineering layout.

The final four-row arrangement should therefore be understood as the geometric solution that best reconciles the different requirements of the case study. A higher turbine count could have moved the installed capacity closer to the nominal 500 MW target, but would have required either a more crowded distribution, a less balanced row structure, or a greater occupation of the study polygon, all of which would reduce the internal flexibility later required for wave energy deployment. The adopted 22-turbine configuration, corresponding to 440 MW, was selected because it provides a sufficiently high-capacity fixed wind layer while maintaining a regular spacing framework and preserving a meaningful residual area for the co-located optimisation problem.

At this stage, the reference floating wind farm is fully defined in geometric terms. Its row structure, spacing values, staggered configuration, and directional orientation are now fixed inputs to the case study. The next step is therefore to convert this geometric layout into an effective spatial constraint by introducing the turbine-related exclusion zones and quantifying the net area that remains available for wave energy placement.

4.5. Turbine Exclusion Zones and Net Deployment Area

4.5.1. Exclusion-Radius Definition

With the geometry of the fixed wind farm now established, the case study can move from simple layout definition to effective spatial constraint definition. This is done by assigning an exclusion zone around each turbine, within which wave energy modules cannot be placed. These protected envelopes convert the wind farm from a purely geometric arrangement into an operational footprint that must be respected by the subsequent wave energy optimisation.

As discussed previously in Chapter 2, floating offshore wind turbines require a larger effective protected area than bottom-fixed units because the occupied footprint is not limited to the tower location alone. In addition to the turbine position itself, the surrounding area must also accommodate the operational influence of the floating support structure, the mooring system, and the dynamic behaviour of associated cables. For that reason, the turbine-related exclusion zone is represented here as a simplified radial buffer, used to define the minimum surrounding area that must remain reserved around each unit before the wave layer is introduced.

The exclusion-radius criterion adopted in this work follows the depth-dependent rule established earlier in chapter 2, where indicative floating-wind protection distances were defined in the order of 4 to 8 times the local water depth (Figure 11). Since the P3 site was previously characterised in chapter 4 as having a mean water depth of approximately 150 m, this criterion yields an admissible first-order interval between approximately 600 m and 1200 m. This range defines the practical envelope within which the project-scale turbine buffer is selected.

The application of this criterion to the fixed wind farm is shown in Figure 67, where the adopted turbine exclusion envelopes are superimposed on the layout defined in the previous subsection.

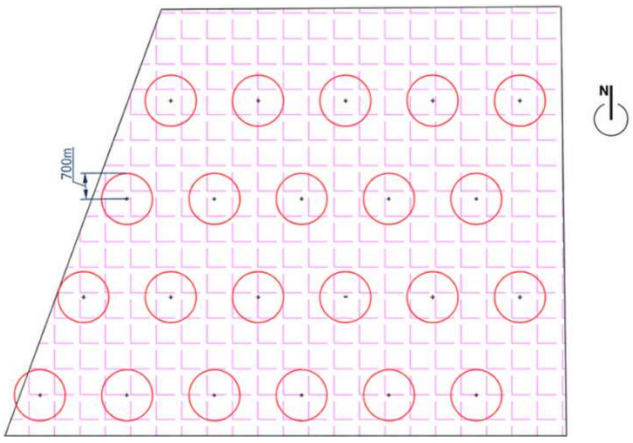


Figure 67 - Exclusion zones around each turbine (red circles).

Figure 67 shows the exclusion zones as circular buffers centred on each turbine position. This simplified representation is adopted to provide a clear and uniform geometric rule for the case study, allowing the turbine-related protected areas to be incorporated directly into the spatial definition of the optimisation domain. At this stage, the objective is not to reproduce the detailed geometry of mooring lines or dynamic cable paths, but rather to establish a practical and consistent no-placement envelope around each turbine that can be applied uniformly across the full wind farm.

Within the admissible interval for the turbines' exclusion zone, a value of 700 m is adopted for the present case study. This value is deliberately positioned in the lower-middle part of the acceptable range.

It is therefore neither the minimum possible buffer, which would reduce conservatism, nor the most restrictive choice, which would substantially reduce the area available for wave energy deployment. Instead, it represents an intermediate engineering compromise that preserves a realistic protection margin around the floating wind units while maintaining a sufficiently large residual domain for the subsequent optimisation of the wave layer.

This choice is particularly important in the context of the present work because the objective is not simply to define a conservative wind farm buffer, but to define one that remains compatible with a meaningful co-location problem. A radius too close to the lower bound could underrepresent the effective protected footprint of the floating turbines. By contrast, a value too close to the upper bound would excessively constrain the area available for wave energy deployment, reducing the optimisation problem to a narrow residual geometry rather than a realistic co-location design space. Therefore, the adopted 700 m radius therefore preserves both engineering plausibility and optimisation relevance.

For the purposes of the case study, this exclusion distance is applied uniformly to all 22 turbines of the fixed wind farm. Each turbine is therefore associated with the same circular no-placement envelope, ensuring that the constraint is implemented in a simple, transparent, and reproducible manner across the full layout. The cumulative effect of these exclusion zones on the study area, together with the resulting admissible domain for wave energy placement, is quantified in the following subsection.

4.5.2. Available Area for WEC Deployment

With the turbine exclusion envelopes now defined, the next requirement is to determine how much of the 152 km² study polygon remains available for wave energy deployment. This step is essential because it converts the exclusion-radius criterion introduced in the previous subsection into a quantifiable deployable domain, making it possible to understand the actual spatial margin within which the CorPack layer can later be positioned. At this stage, the study area is no longer treated as a single continuous polygon, but as the combination of two distinct spatial components: the protected area occupied by the wind layer and the residual area that remains admissible for wave energy placement.

The application of the turbine buffers removes a substantial portion of the central and upper area of the polygon, while still preserving a continuous residual domain around and between the protected envelopes. This remaining space constitutes the first-order admissible area available for WEC placement. Its importance is not only quantitative, but also geometric. The residual domain must remain sufficiently large and sufficiently connected to allow the wave energy layer to be positioned with meaningful spatial flexibility rather than being restricted to a small number of isolated pockets.

The area removed by each turbine buffer is obtained from the area of a circle with radius 700 m which corresponds to approximately 2.01 km^2. With the 22 turbines adopted, the first-order total exclusion area associated with the wind layer is approximately 44.23 km^2.

A key point in this calculation is that the exclusion envelopes can be treated, at first order, as non-overlapping. This simplification is consistent with the spacing adopted for the fixed wind layout. The in-row distance of 2400 m and the inter-row distance of 2700 m are both substantially greater than the 1600 m diameter of each exclusion circle. In addition, the staggered arrangement further increases the effective centre-to-centre distance between turbines in adjacent rows. Under this geometry, the protected envelopes remain separated from one another, allowing the total excluded area to be approximated as the sum of the individual turbine buffers without requiring overlap corrections.

The main area-balance values derived from this first-order calculation are summarised in Table 10.

Table 10 - First-order area balance after turbine exclusion.

Parameter	Value
Total study area	152 km^2
Exclusion radius per turbine	700 m
Excluded area per turbine	2.01 km^2
Number of turbines	22
Total first-order turbine exclusion area	44.23 km^2
First-order residual deployment area	107.77 km^2

As indicated in Table 10, subtracting the total turbine exclusion area from the full study polygon yields a first-order residual deployable area of approximately 107.77 km². This value represents the main admissible spatial envelope available for the subsequent placement of wave energy modules. It is this residual area, rather than the original 152 km² polygon, that defines the effective scale of the co-location problem addressed in the remainder of the case study.

This result is particularly important for the interpretation of the following sections, since it is large enough to preserve a meaningful optimisation domain, but it is also sufficiently reduced relative to the original study polygon to make clear that the wind layer imposes a substantial spatial constraint on the wave layer. In other words, the co-location problem is not defined in an unrestricted project area, but within a domain that has already been significantly structured by the presence of the fixed floating wind farm and its associated protected envelopes.

The significance of this subsection is therefore twofold. On one hand, it quantifies the direct spatial impact of the fixed wind farm on the study area. On the other, it provides the first explicit basis for judging whether the wave layer adopted later in the chapter is spatially credible. Without this intermediate step, the subsequent choice of CorPack density would appear disconnected from the actual available area. By explicitly calculating the residual deployable domain here, the wave-energy definition that follows can be framed as a direct consequence of the spatial constraints already imposed by the fixed wind layer.

With the residual deployable domain established, the next section introduces the wave energy module adopted for the case study and specifies its optimisation-ready definition.

4.6. Definition of the Wave-Energy Layer

4.6.1. CorPack Module Definition

With the first-order deployable area now established, the case study can move from the definition of the fixed wind layer to the explicit representation of the wave energy layer. At this stage, the objective is not to define the internal arrangement of individual WEC devices, but to introduce the deployable module through which the wave farm will be represented in the optimisation framework. The wave layer is therefore formulated in a way that remains consistent with the farm-scale abstraction adopted previously in Chapter 2, where the CorPack was identified as the practical deployment unit rather than the individual CorPower C4 device.

The wave energy layer adopted in the present case study is based on the CorPack as the fundamental deployable unit. Each CorPack is treated here as one discrete wave-energy module within the optimisation domain, meaning that the placement problem is defined at the level of modular farm-scale units rather than at the level of individual buoys. This representation is consistent with the engineering and modelling logic established earlier, since it preserves the physical meaning of the CorPack as a repeatable deployment cluster while allowing the co-located wave farm to be described through a manageable number of spatial entities.

The main parameters used to define the wave layer are summarised in Table 11.

Table 11 - CorPack representation adopted for the wave energy layer.

Parameter	Value	Role in case-study definition
Deployable wave module	1 CorPack	Basic spatial unit used in the optimisation
Number of CorPacks	40	Total number of deployable wave modules considered
Equivalent capacity per CorPack	10 MW	Nominal wave capacity assigned to each module
Total equivalent installed wave capacity	400 MW	Aggregate nominal wave capacity of the wave layer
Equivalent number of C4 devices per CorPack	28	Device count used to express module composition
Total equivalent number of C4 devices	1120	Aggregate number of individual devices represented by the full wave layer

As indicated in Table 11, a total of 40 CorPacks are adopted for the project-scale study area. Using the modular representation established earlier in the work and taking 10 MW as the nominal equivalent capacity of each CorPack, this corresponds to a total equivalent wave capacity of 400 MW. The device count associated with each module, and the corresponding aggregate equivalent count for the full wave layer, are retained as descriptive indicators of the physical scale represented by the modular abstraction (as previously defined in Chapter 3), rather than as direct optimisation variables.

This representation is particularly important because it preserves consistency between the technology description presented in Chapter 2 and the farm-scale optimisation framework adopted in the present chapter. The optimisation problem is therefore not formulated in terms of the exact position of each individual point absorber, but in terms of the placement of modular CorPack units inside the residual deployable area quantified in Section 4.6.1. The individual C4 count remains relevant as an expression of equivalent installed capacity and of the physical scale represented by each module, but it does not constitute the direct optimisation variable.

Figure 68 provides a schematic representation of how the wave layer is interpreted at case-study level, with each CorPack treated as one deployable unit inside the residual admissible domain defined by the fixed wind farm and its exclusion envelopes.

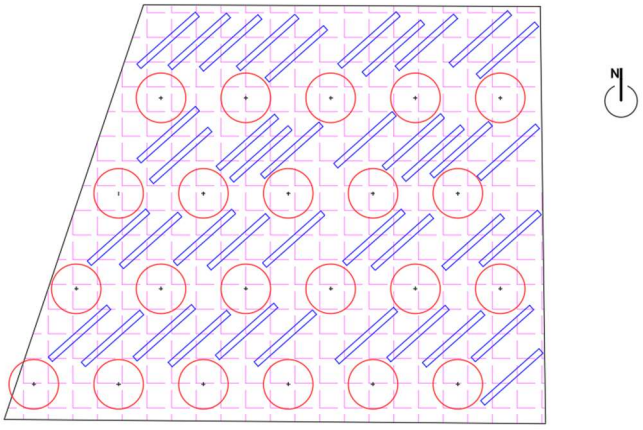


Figure 68 - Schematic representation of the residual deployable area populated conceptually by CorPack modules (blue rectangles).

Figure 68 is not intended to represent the final optimised solution, but rather the modelling abstraction adopted for the wave layer. Its purpose is to show that the deployable entities considered in the optimisation are modular CorPack units distributed within the residual domain left available by the fixed wind farm. In this sense, the figure serves to bridge the geometric area-balance established in the previous section and the optimisation-ready representation developed in the present one.

The adoption of the CorPack as the deployable unit also contributes to maintaining the problem at a realistic engineering scale. If the wave layer were represented directly through the placement of individual devices, the spatial problem would become excessively detailed and would no longer correspond to the modular deployment logic established for the technology itself. By working at the level of CorPack modules, the optimisation remains aligned with the practical implementation scale of the wave farm while preserving a clear and computationally manageable set of decision variables.

At this point, the wave layer is fully defined in terms of its deployable unit, total number of modules, and equivalent installed capacity. What remains to be explained is why the specific value of 40 CorPacks was adopted. That justification is developed in the following subsection, where the selected module density is interpreted in relation to the residual deployable area and to the balance required between realistic area occupation and optimisation flexibility within the residual domain.

4.6.2. Selection of the CorPack Density

With the wave layer now defined in terms of its deployable unit, the remaining step is to justify the number of modules introduced into the case study. This point is particularly important because the total number of units determines how strongly the wave layer occupies the residual deployable area identified in Section 4.6.1 and, therefore, how demanding the co-location problem becomes in spatial terms. The selected value cannot be interpreted only as a capacity assumption, it is also a geometric and methodological choice that directly influences the realism and difficulty of the optimisation problem.

The adopted density of 40 CorPacks was established from preliminary geometric placement tests carried out within the residual deployable domain, considering both the available area and the modular scale of the CorPack defined in Section 2.2.3. The objective at this stage was not to determine an optimised final arrangement, but to identify a module count that would provide a credible level of occupation without excessively saturating the admissible space. In this sense, the selected number had to satisfy two simultaneous requirements: it had to be large enough for the wave layer to constitute a meaningful

project-scale presence inside the study area, and it had to remain low enough to preserve sufficient freedom for the optimisation algorithm to reposition modules in a non-trivial way.

The practical implications of different CorPack density levels are summarised in Table 12.

Table 12 - Qualitative interpretation of CorPack density in the residual deployable domain.

Density level	Spatial effect	Role in case-study definition
Lower than the adopted value	The deployable area remains underused and large portions of the residual domain remain empty	Spatial competition is weak, feasible solutions are abundant, and the optimisation becomes less realistic
Adopted value	A substantial portion of the residual domain is occupied while preserving internal freedom for module rearrangement	The optimisation remains constrained enough to be meaningful, but still retains sufficient flexibility to explore alternative layouts
Higher than the adopted value	The deployable area becomes increasingly saturated and free corridors between modules are progressively reduced	Feasible solutions become scarcer, geometric crowding dominates the search, and the optimisation becomes constrained primarily by placement infeasibility rather than by layout quality

As indicated in Table 12, the adopted value corresponds to the intermediate case in which the wave layer produces a substantial occupation of the available area without exhausting the geometric flexibility required by the optimisation framework. If the number of CorPacks were significantly lower, the study area would be only partially used, and the resulting layouts would not adequately reflect the spatial competition expected in a realistic co-located farm. In that situation, the optimisation would tend to operate in an excessively permissive domain, where many feasible configurations would exist simply because the wave layer would not exert enough pressure on the residual area.

The opposite extreme would be equally problematic. A substantially higher number of modules would move the case study toward an over-constrained configuration in which the residual deployable area would become excessively crowded. Under those conditions, the optimisation process would be dominated by geometric infeasibility, local crowding, and restricted module mobility, rather than by a meaningful search for improved spatial arrangements. In practical terms, the optimisation problem would become less informative, because the resulting layouts would reflect a near-saturated domain rather than a realistic balance between occupation and layout freedom.

The selected number therefore represents an intermediate and deliberate engineering choice. It ensures that the wave layer occupies a meaningful share of the residual deployable domain while preserving enough internal space for modules to be repositioned during optimisation. This is precisely the balance required for the present case study. The domain must be sufficiently populated to make co-location spatially competitive, but it must not be so saturated that the optimisation loses the ability to distinguish between poor and improved layouts.

A further consequence of this choice is that the resulting wave layer remains at a comparable project scale to the wind layer already defined in Section 4.4. With 10 MW assigned to each module, it yields an equivalent wave capacity of 400 MW, which places the wave layer in the same several-hundred-megawatt order of magnitude as the host wind farm. Even though the primary basis for selecting the module count is spatial rather than electrical, this outcome reinforces the internal coherence of the case study by ensuring that both layers remain comparable in overall project scale.

For these reasons, the adopted value of 40 CorPacks should be understood as the module density that best reconciles area occupation and optimisation flexibility within the residual deployable domain. It is neither a minimal representation of the wave layer nor a near-saturation upper bound. Rather, it is the value that provides the most credible basis for a farm-scale co-location problem in which the wave layer is spatially significant, yet still optimisable. With this choice established, the full project-scale definition of the case study is complete.

5. RESULTS AND DISCUSSION

5.1. OVERVIEW OF THE OPTIMISATION RESULTS

This chapter presents the results obtained from the optimisation described in the previous sections. The aim of this analysis is to assess how different modelling assumptions and optimisation parameters influence both the spatial arrangement of CorPacks and the overall performance of the co-located wave-wind energy farm. Six optimisation cases were defined, each designed to investigate a specific aspect of the modelling framework. In particular, the analysed scenarios explore the influence of the hydrodynamic representation adopted for the aggregated WEC cluster, the effect of varying the number of deployed devices within the farm, and the sensitivity of the optimisation process to key parameters of the genetic algorithm.

The chapter begins with an overview of the optimal layouts produced by each optimisation run. These layouts correspond to the final configurations identified by the GA after the optimisation process reached convergence.

Following this overview, the subsequent sections analyse the optimisation results in greater detail through a series of pairwise comparisons between cases. The behaviour of the optimisation process is examined using several performance indicators. These include the evolution of the CF during the optimisation, the convergence of the average absorbed power per device, the behaviour of P2M, and metrics describing the diversity of the candidate solutions explored by the GA. Together, these indicators provide insight into both the quality of the final solutions and the internal dynamics of the optimisation process. In addition to these optimisation metrics, the hydrodynamic effects associated with the optimised layouts are also evaluated. This is achieved by analysing the H_s fields generated for each configuration, which illustrate how the presence and spatial distribution of the arrays modify wave propagation within the farm area.

Figure 69 presents the optimal layouts obtained for the six optimisation cases considered in this work. Each subplot illustrates the final configuration produced by the GA for a given case within the study domain. The black triangles represent the fixed locations of the wind turbines and the surrounding exclusion zones that prevent WEC deployment in their vicinity. The red line segments correspond to the aggregated WEC elements that form the optimised wave-energy array.

.

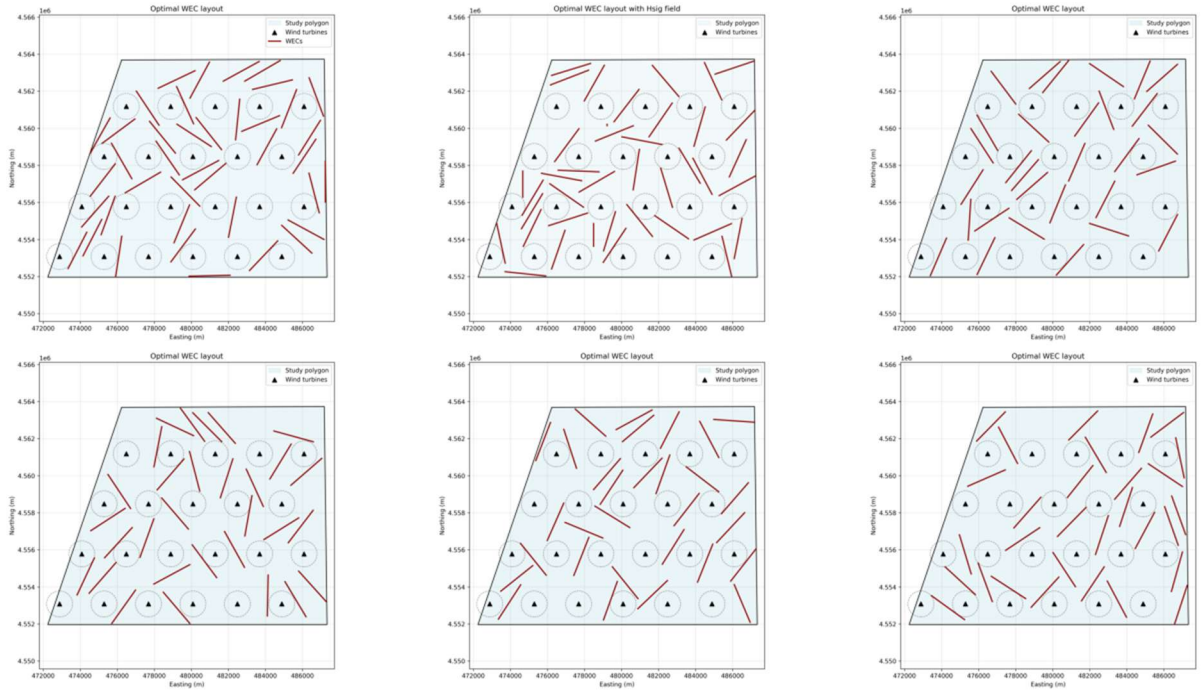


Figure 69 - Optimal WECS layouts: Case 1 (top-left); Case 2 (top-mid); Case 3 (top-right); Case 4 (bottom-left); Case 5 (bottom-mid); Case 6 (bottom-right). Wind turbines (triangles); Exclusion zones (circles); CorPacks (red lines); PAER limits (black polygon).

All configurations satisfy the spatial constraints imposed by the wind-turbine exclusion buffers and remain fully contained within the deployable region of the study polygon. Although the overall deployment area is identical for all cases, the layouts exhibit noticeable differences in the orientation and spatial distribution of the CorPack segments. These variations reflect the different modelling assumptions and optimisation parameters explored in each run. A more detailed interpretation of these layouts and their hydrodynamic implications will be provided in the subsequent comparative sections of this chapter.

The main quantitative results obtained from each optimisation are summarised in Table 13. The table reports the best CF achieved, the corresponding maximum absorbed power of the array, the generation at which the best solution was identified, the total number of generations performed, and the associated computational times. These values provide an overall picture of the optimisation performance across the six analysed cases.

Table 13 - Summary of optimisation results.

Case	Best CF	Generation of best solution	Total generations	Mean time per generation	Total simulation time
1	0.291	10	70	17 min 38 s	20 h 34 min 24 s
2	0.299	10	88	17 min 38 s	25 h 51 min 50 s
3	0.307	45	57	17 min 38 s	16 h 45 min 10 s
4	0.301	12	71	30 min 42 s	36 h 19 min 29 s
5	0.304	20	129	17 min 38 s	37 h 54 min 50 s
6	0.304	30	108	17 min 38 s	31 h 44 min 31 s

The layouts shown in Figure 69 provide an overview of the spatial solutions obtained by the optimisation process. For clarity and improved visual inspection, the same layouts are also provided in larger format in Appendix B.

In the following sections, the different optimisation cases will be compared in pairs in order to isolate the influence of each modelling assumption or algorithm parameter. These comparisons will include the analysis of the corresponding H_s fields, which illustrate how the optimised arrays modify wave propagation across the farm. Enlarged versions of these Hs plots are provided in Appendix C for improved readability.

5.2. HYDRODYNAMIC Q-FACTOR EVALUATION

The hydrodynamic efficiency of a wave energy array cannot be evaluated solely from the absorbed power obtained by the optimisation procedure, since the interaction between neighbouring devices may significantly modify the collective response of the array. For this reason, the q-factor provides an important complementary indicator for assessing array behaviour. In the present work, the q-factor was calculated according to equation (2), previously introduced in Chapter 2.

In this section, the objective is to evaluate the hydrodynamic interaction behaviour of the final optimised layouts obtained in the six optimisation cases presented earlier. Instead of analysing a single sea state, the q-factor was evaluated for each layout under the six representative sea states used during the optimisation process. These sea states correspond to the reduced wave climate defined in Table 4, which represents the dominant wave conditions at the study location.

For each optimisation case, the final layout obtained from the GA was simulated under each of the six sea states. The absorbed power of the full array was then compared with the sum of the absorbed power of the corresponding isolated devices, allowing the q-factor to be calculated for every combination of layout and sea state. As a result, each optimisation case produced six q-factor values, one for each representative sea state. From these values, the maximum, minimum, and mean q-factor were determined to characterise the overall hydrodynamic interaction behaviour of each optimised array.

The q-factor values obtained for all cases and sea states are presented in Figure 70, while the summary indicators extracted from these results are reported in Table 14.

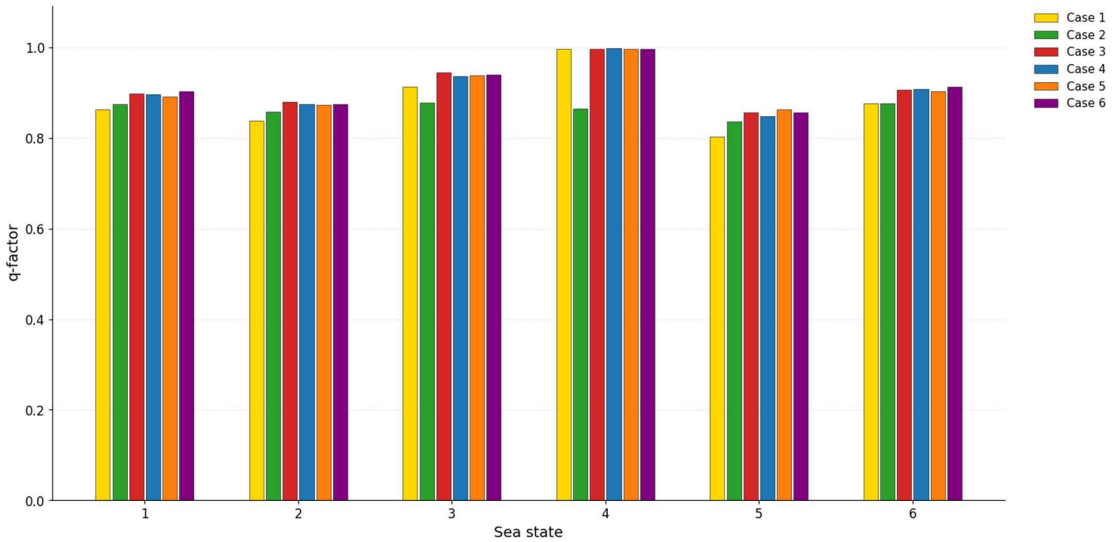


Figure 70 - q-factor values obtained for the six optimisation cases under the representative sea states.

The results show that the q-factor varies systematically with the incident sea state, confirming that the interaction behaviour of the arrays depends strongly on the incoming wave conditions. Across all optimisation cases, the highest q-factor values occur for Sea State 4, for which the layouts reach values very close to unity. In contrast, the lowest q-factor values are generally observed for Sea State 5, and to a lesser extent for Sea State 2

The variation of the q-factor across the six sea states reflects the combined influence of the wave characteristics and the spatial configuration of the CorPack arrays. In particular, the interaction behaviour of the layouts is governed not only by the incident H_s and T_p, but also by the relative orientation between the incoming wave direction and the geometry of the WECs arrangement within the farm. This directional relationship determines how the wave field propagates through the array and therefore controls the intensity of constructive or destructive interference between neighbouring devices.

The consistently high q-factor values obtained for Sea State 4 indicate that this combination of wave conditions and incident direction interacts favourably with the geometry of the optimised layouts. Under these conditions, the incoming wave field propagates through the farm in a manner that limits strong shadowing effects between adjacent devices and promotes a relatively uniform distribution of wave energy across the array. In addition, the spatial placement and orientation of the segments within the layouts align more effectively with this dominant wave direction, reducing local interference between neighbouring elements. As a result, the arrays operate under interaction conditions that remain very close to the ideal isolated-device response, which explains the q-factor values approaching unity observed for this sea state.

A second important observation concerns the relative behaviour of the different optimisation cases. The layouts corresponding to Cases 3 to 6, which contain 30 WEC units, generally produce higher q-factor values than the layouts obtained for Cases 1 and 2, which include 40 CorPack units. This tendency is clearly visible for Sea States 1, 3, 4 and 6. The difference is consistent with the spatial characteristics of the final layouts shown in Section 5.1. The arrays with 30 units exhibit larger average spacing between devices and a less crowded occupation of the deployable region, which reduces the intensity of hydrodynamic interaction between neighbouring converters. Consequently, the collective response of these layouts remains closer to the sum of isolated-device behaviour, producing q-factor values nearer to unity.

The comparison between Case 1 and Case 2 also provides insight into the influence of the hydrodynamic representation adopted in the optimisation. While both cases contain the same number of devices, the first case uses the power matrix representation, whereas the second case adopts the RCW-based formulation. The q-factor differences between the two cases vary depending on the sea state, indicating that the interaction behaviour captured by the two hydrodynamic representations remains sensitive to the specific wave conditions. Rather than producing a uniform shift in interaction efficiency, the two formulations lead to slightly different responses depending on the incident wave characteristics.

Among the 30-device layouts, the variation in q-factor values is comparatively small. This indicates that once the device density is reduced, the influence of the GA parameters on the hydrodynamic interaction behaviour becomes secondary compared with the influence of array density itself. Although minor differences remain visible, the layouts produced by the sensitivity analyses preserve interaction characteristics like those of the reference 30-CorPack configuration.

An important feature of the results is that all q-factor values remain close to unity, typically ranging between approximately 0.80 and 1.00 depending on the sea state. This indicates that none of the optimised arrays exhibits severe destructive interaction effects. At the same time, the fact that the q-factor does not significantly exceed unity confirms that the layouts do not rely on unrealistically strong

constructive interference to achieve high performance. Instead, the optimisation process identifies configurations that balance spatial feasibility, hydrodynamic interaction control, and overall energy extraction.

For clarity, the main summary indicators derived from the six sea-state simulations are presented in Table 14.

Table 14 - Minimum, maximum and mean q-factor obtained for the final layout of each optimisation case.

Case	q-factor min	q-factor máx	q-factor mean
1	0.80	0.99	0.88
2	0.83	0.87	0.86
3	0.85	0.99	0.91
4	0.84	0.99	0.91
5	0.87	0.99	0.91
6	0.85	0.99	0.91

Although Table 14 summarises the overall behaviour of each case, the graphical representation remains essential for interpreting how the q-factor evolves across different sea states. Figure 70 shows that some layouts maintain relatively stable interaction behaviour across the six representative conditions, whereas others exhibit stronger sensitivity to the incident wave direction and spectral characteristics.

Overall, the q-factor analysis confirms that the optimised layouts maintain hydrodynamically consistent interaction behaviour across the representative wave climate, while also highlighting the importance of device density and wave-direction alignment in determining the collective performance of the array.

These results complement the CF and absorbed-power analyses presented in the following sections, providing an additional perspective on the physical behaviour of the optimised wave-energy farms.

5.3. Influence of the Hydrodynamic Representation: Case 1 vs. Case 2

The comparison between Case 1 and Case 2 was conducted to investigate how the hydrodynamic representation adopted for the aggregated WEC obstacle influences the optimised configuration of the co-located wind-wave farm. In both cases the same number of CorPacks is deployed and identical GA parameters are used. Consequently, the only intentional difference between the two optimisation runs lies in the physical description of wave-device interaction implemented in the spectral model. Case 1 employs the power matrix representation, while Case 2 adopts the RCW formulation. This comparison therefore allows the impact of the hydrodynamic modelling strategy itself to be isolated and evaluated, both in terms of the spatial arrangement of the devices and the performance indicators produced by the optimisation.

The optimal layouts previously presented in Figure 69 of Section 5.1 already indicate that the two modelling approaches lead to somewhat different deployment patterns. In both cases the segments occupy the same admissible region and fully respect the exclusion buffers surrounding the wind turbines. Nevertheless, differences appear in the orientation and clustering of the devices. The layout obtained

with the power matrix appears slightly more dispersed throughout the available corridors, whereas the RCW-based configuration displays several areas where the segments tend to align differently relative to neighbouring devices. Although these variations may appear subtle, they illustrate how the hydrodynamic formulation can influence the optimisation process by modifying how energy extraction is interpreted within the model.

To examine how the resulting layouts interact with the incoming wave field, Figure 71 shows the spatial distribution of significant wave height for the optimised configurations of Case 1 and Case 2 under the representative sea state considered, sea state 4 from Table 4. For the followings sections of this present chapter, the same representative sea state is considered.

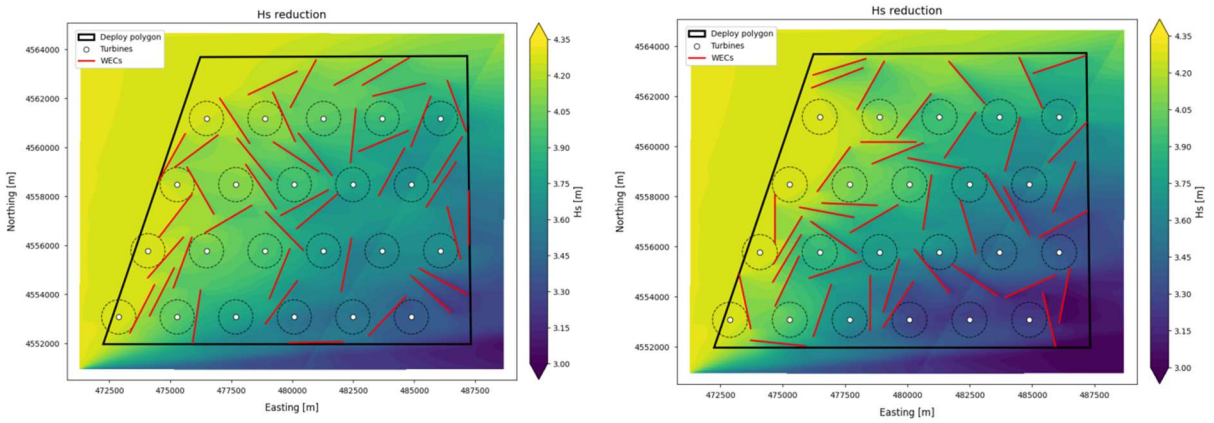


Figure 71 - H_s values for Case 1 (left) and Case 2 (right).

The two maps clearly exhibit the same large-scale propagation behaviour, confirming that both simulations are subjected to identical offshore wave conditions. In both cases, higher H_s values appear along the north-western portion of the domain, gradually decreasing towards the south-eastern side as waves propagate across the farm. Overlaying this general trend, the influence of the CorPack arrays becomes visible through the formation of local attenuation zones.

Although the overall attenuation footprint is similar in both cases, some differences can be observed in the structure of these shadow regions. In the power matrix configuration, Figure 71 (left), the disturbances appear relatively distributed across the domain, with several localized zones of reduced wave height that remain relatively distinct from each other. In contrast, the RCW-based layout, Figure 71 (right), produces slightly stronger attenuation in parts of the central and eastern sectors of the domain, leading to more connected regions of wave-height reduction. This suggests that the RCW formulation tends to emphasize the energy extraction associated with some device alignments, producing a more pronounced modification of the local wave field.

Additional insight into the optimisation behaviour can be obtained by analysing the convergence of the CF throughout the GA iterations. Figure 72 illustrates the evolution of this indicator for both cases, including the best solution found at each generation, the population mean, the elite mean, and the corresponding min-max range.

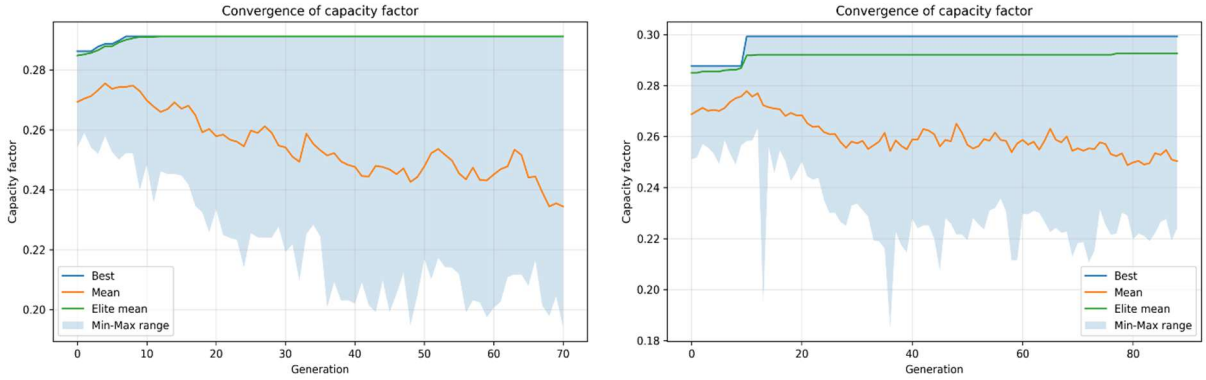


Figure 72 - CF evolution for Case 1 (top) and Case 2 (bottom).

The convergence histories reveal that both optimisation runs reach their best solution relatively early in the search process. In Case 1, the CF increases slightly during the first iterations and stabilizes after roughly 10 generations, reaching a final value close to 0.291. A similar behaviour can be observed in Case 2, where the optimum is also identified during the early part of the optimisation, but the final value reaches approximately 0.299. This corresponds to an increase of about 0.008 in absolute CF, equivalent to roughly 3% relative to the power matrix case.

As indicated in Table 13, the optimisation of Case 1 was completed after 70 generations, whereas Case 2 required 88 generations. Although the best solution appears early in both runs, the RCW-based case therefore required a longer overall search process before the stopping criteria were satisfied. Despite this longer run, the final solution obtained with the RCW representation exhibits a slightly higher energy-conversion efficiency. The same behaviour is observed when examining the evolution of the average absorbed power per CorPack, shown in Figure 73.

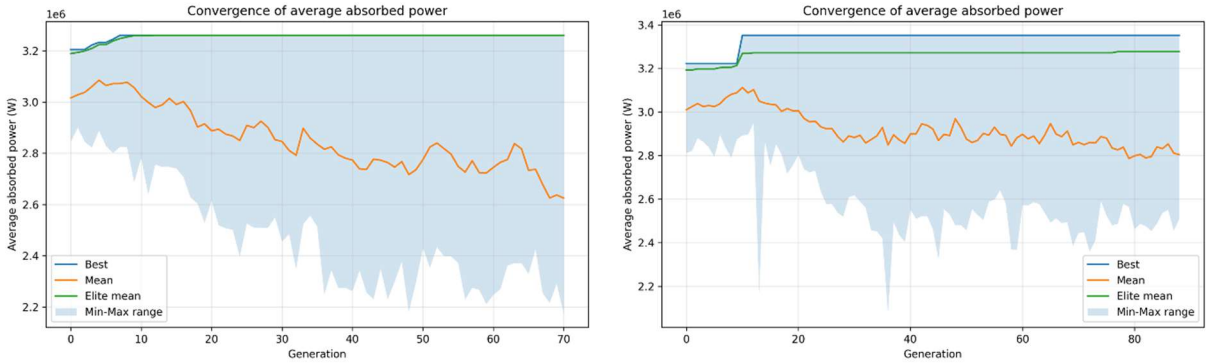


Figure 73 - Average absorbed power evolution for Case 1 (top) and Case 2 (bottom).

The power trends mirror the results obtained for the CF. In Case 1, the optimal configuration reaches an average absorbed power of approximately 3.26 MW per CorPack, whereas in Case 2 the optimised layout reaches about 3.35 MW. The difference of roughly 0.09 MW per device corresponds once again to an increase of around 3%. Expressing the results in terms of absorbed power highlights the practical meaning of this improvement: under the RCW formulation each device, on average, captures slightly more wave energy than in the power matrix case. This indicates that the alternative hydrodynamic representation modifies the way energy extraction is evaluated by the model, ultimately favouring layouts associated with somewhat higher power capture.

The internal dynamics of the optimisation process can also be interpreted through the diversity indicators presented in Figure 74, which show the evolution of the CF range and the coefficient of variation within the population.

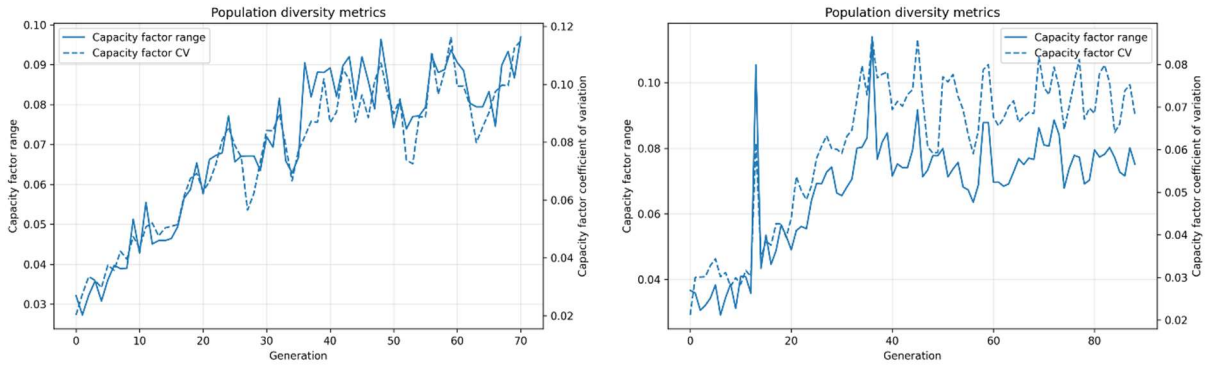


Figure 74 - CF range and coefficient of variation evolution for Case 1 (top) and Case 2 (bottom).

In both cases the diversity indicators gradually increase during the optimisation, reflecting the fact that the GA continues to generate candidate layouts with varying performance levels as the search progresses. However, the nature of this diversity differs slightly between the two runs. In Case 1 the increase in diversity is accompanied by a relatively wide spread of fitness values across the population. In Case 2, although diversity also increases, the relative dispersion of performance remains somewhat lower. This indicates that the RCW formulation produces a population in which candidate layouts remain distinct but tend to cluster more closely around higher-performance solutions. In practical terms, the hydrodynamic representation therefore influences not only the best layout identified by the algorithm but also the overall structure of the optimisation landscape.

Further information on the internal behaviour of the array can be obtained from the P2M, shown in Figure 75.

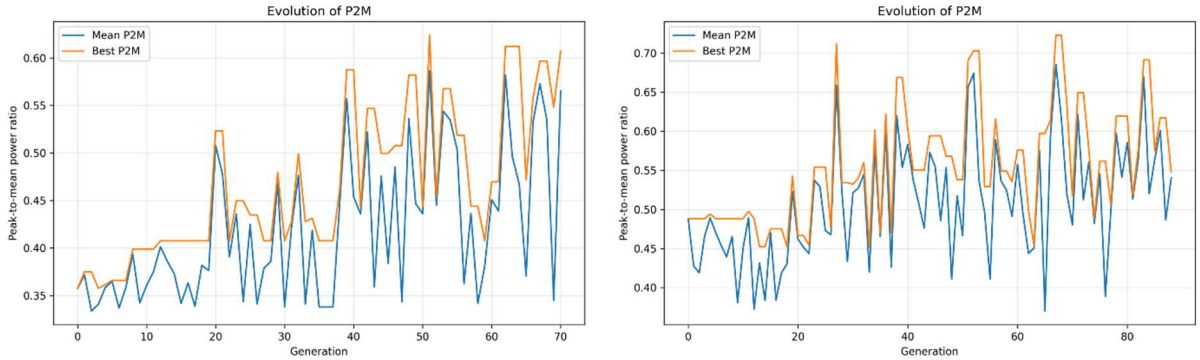


Figure 75 - P2M evolution for Case 1 (top) and Case 2 (bottom).

This metric provides an indication of how evenly the absorbed power is distributed among the individual devices. The difference between the two cases is evident. In Case 1, the highest P2M values reach approximately 0.62, while in Case 2 the maximum values increase to around 0.72. This represents an increase of roughly 16% when the RCW formulation is used. Such an increase suggests that the RCW-based configuration produces a less uniform distribution of absorbed power, with certain devices capturing noticeably more energy than others. In contrast, the power matrix case leads to a slightly more

balanced distribution of power among the WECs. This distinction is consistent with the stronger attenuation patterns observed in the wave-height maps and indicates that the hydrodynamic formulation can influence not only the total energy captured but also how that energy is distributed across the farm.

For clarity, the principal quantitative differences between the two cases are summarised in Table 15.

Table 15 - Quantitative comparison between Case 1 and Case 2.

Metric	Case 1	Case 2	Relative difference
Best CF	0.291	0.299	+3.0%
Best average absorbed power per unit	3.26 MW	3.35 MW	+3.0%
Maximum P2M	0.62	0.72	+16%
Generations to reach optimum	10	10	Very Similar Convergence

Taken together, the results of this comparison indicate that the RCW-based representation leads to a modest but consistent improvement in the predicted performance of the optimised layout. The H_s fields show somewhat stronger attenuation zones, while the CF and absorbed-power results indicate an increase of approximately 3% in device-level efficiency. At the same time, the P2M analysis reveals a less uniform distribution of power across the array, suggesting stronger local differences between devices. In addition, the RCW case requires a longer optimisation run, with 88 generations compared to 70 generations in the power matrix case.

5.4. Influence of the Number of CorPacks: Case 1 vs. Case 3

The comparison between Case 1 and Case 3 was carried out to evaluate the influence of the number of deployed CorPacks on the optimised co-located farm configuration. In both cases, the same hydrodynamic representation was adopted through the power matrix approach, and the same GA settings were maintained, so that the only intentional difference between the two runs was the number of units considered in the optimisation. Case 1 corresponds to the configuration with 40 devices, whereas Case 3 considers 30. This comparison is particularly relevant because it allows the effect of device density to be assessed directly, both in terms of the final spatial arrangement and in terms of the associated hydrodynamic and optimisation performance.

As previously shown in Figure 69, the optimal layouts obtained for Case 1 and Case 3 already suggest different spatial organization patterns. Although both layouts fully respect the wind turbine exclusion zones and occupy the same deployable region, the 40-CorPack configuration necessarily produces a denser occupation of the residual area, with a larger number of devices distributed between turbine corridors and near the polygon boundaries. In contrast, the 30-device configuration exhibits a slightly more open arrangement, with larger average spacing between neighbouring devices and a less congested overall deployment pattern. This distinction is important because the density of deployment directly affects the strength of wave-device interactions inside the farm.

To illustrate how these two optimised layouts modify the incident wave field, Figure 76 presents the corresponding H_s distributions for Case 1 and Case 3 under the analysed representative sea state.

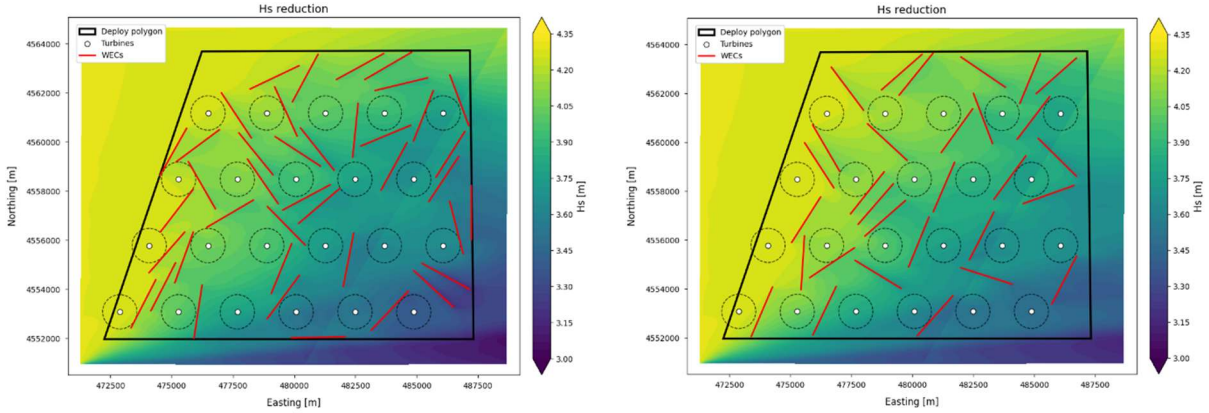


Figure 76 - H_s values for Case 1 (left) and Case 3 (right).

The H_s fields confirm that the two cases are subjected to the same general incoming wave conditions, since both maps display the same large-scale propagation trend across the domain, with higher H_s values on the north-western side and lower values towards the south-eastern region. At the same time, the local disturbances introduced by the CorPack arrays reveal meaningful differences between the two optimised solutions. In the first case, the farm produces a more fragmented attenuation pattern, with a larger number of local shadow zones and a greater tendency for these zones to overlap, especially in the central parts of the domain where the device concentration is higher. In the second case, the attenuation pattern remains clearly visible, but it appears slightly smoother and more spatially distributed, with fewer overlaps between neighbouring interaction zones. In practical terms, this indicates that the lower-density array preserves more continuous wave propagation corridors through the farm, which is consistent with a reduction in local hydrodynamic interference. Although the overall wave attenuation footprint remains similar in both cases, the 30-device layout appears to generate a less perturbed wave field, suggesting that the reduction in device density alleviates at least part of the interaction penalty associated with the more crowded arrangement.

This tendency becomes even clearer when the evolution of the capacity factor during the optimisation process is examined. Figure 77 shows the convergence of the capacity factor for both cases, including the best solution found in each generation, the population mean, the elite mean, and the min-max range of the population.

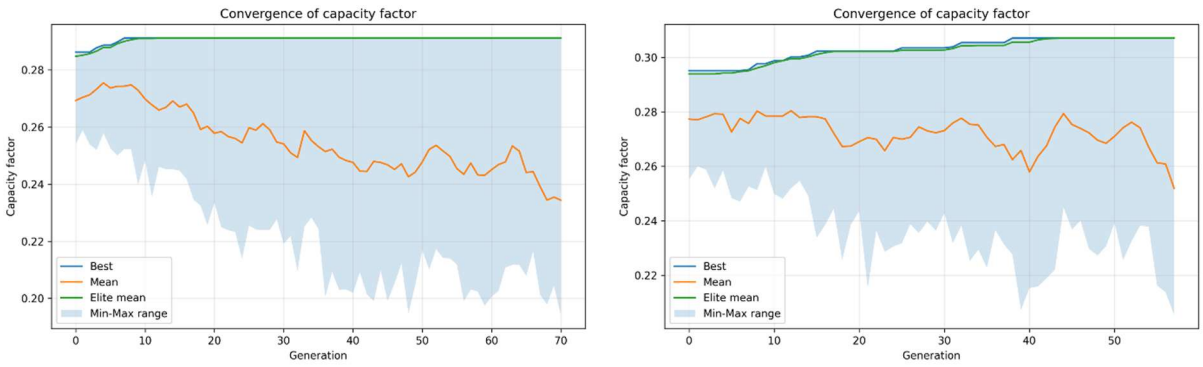


Figure 77 - CF evolution for Case 1 (top) and Case 3 (bottom).

The convergence histories reveal an interesting trade-off between optimisation speed and final performance. In Case 1, the best CF increases only slightly from its initial value and stabilizes rapidly, reaching its optimum after approximately 10 generations. Although the optimal solution is identified

relatively early, the genetic algorithm continued evolving the population until the predefined stopping criterion was reached, resulting in a total of 70 generations for this optimisation run, as reported in Table 13. The final best CF obtained in this case is approximately 0.291.

In Case 3, the convergence occurs more gradually, with the best solution continuing to improve until approximately generation 38, after which the optimal layout stabilizes near a CF of 0.307. In this case the optimisation terminated after 57 generations, again according to the stopping criteria defined for the algorithm. The improvement relative to Case 1 corresponds to an increase of about 0.016 in absolute capacity factor, or approximately 5.2%.

Therefore, although the second case converges more slowly, it ultimately achieves a more efficient final layout on a per-device basis. This behaviour suggests that when fewer devices are deployed, the optimisation retains greater spatial freedom to explore alternative arrangements, and this broader effective search space allows the algorithm to identify layouts with improved hydrodynamic performance. In the more densely packed case, the higher occupancy of the deployable area seems to constrain the search more strongly, leading to quicker stabilization but also to a lower final efficiency.

The same conclusion is supported by the evolution of the average absorbed power per CorPack, presented in Figure 78.

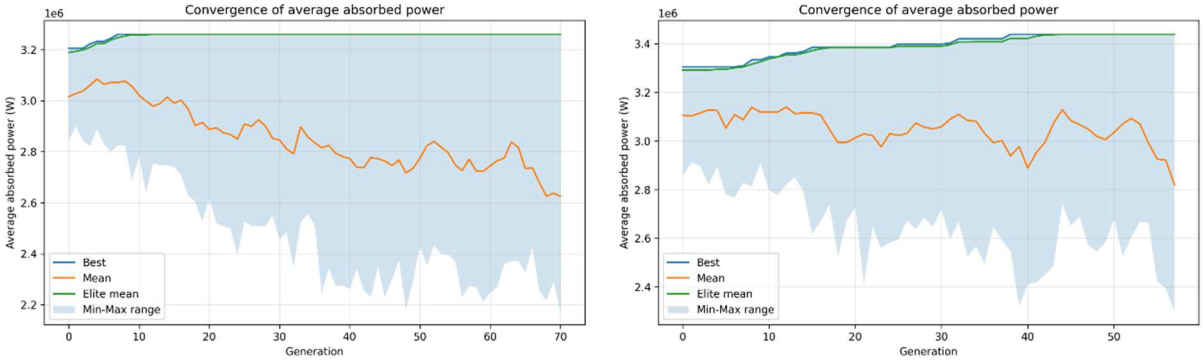


Figure 78 - Average absorbed power evolution for Case 1 (top) and Case 3 (bottom).

The trends shown in this figure are fully consistent with the CF results. In Case 1, the optimal layout reaches an average absorbed power of approximately 3.26 MW per device, whereas in Case 3 the final optimal layout reaches about 3.44 MW. This represents an increase of roughly 0.18 MW per device, again corresponding to approximately 5.2%. Since the definition of CF used here is directly linked to the average absorbed power normalized by the rated power of a single device, the agreement between the two figures is expected. However, expressing the result in power terms makes the practical significance more evident: reducing the number of CorPacks from 40 to 30 enables each device, on average, to absorb more energy. This indicates that the lower-density configuration reduces the hydrodynamic competition among neighbouring devices and allows the available wave resource to be utilized more effectively by each converter. It is also worth noting that the mean performance of the population remains consistently higher in Case 3 than in Case 1, which suggests that, beyond the best solution itself, a broader set of candidate layouts also performs better when the array is less crowded.

The behaviour of the population during the optimisation can be further interpreted through the fitness-diversity indicators shown in Figure 79, where the evolution of the CF range and coefficient of variation are plotted for both cases.

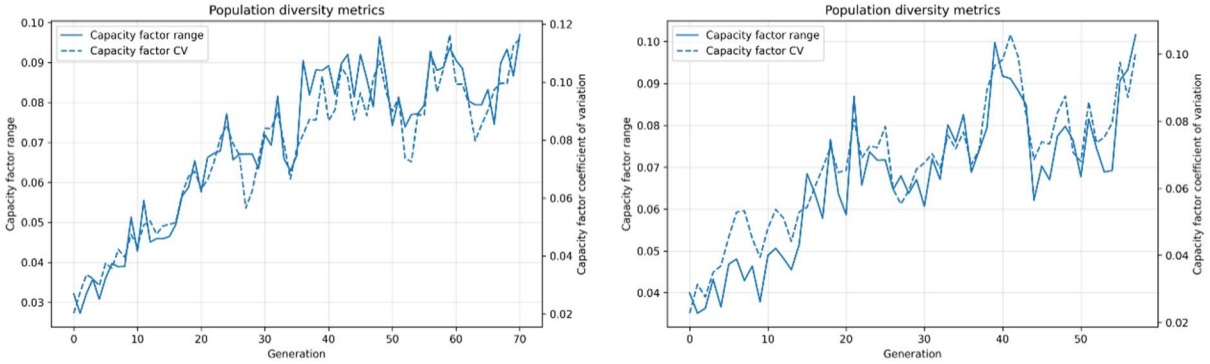


Figure 79 - CF range and coefficient of variation evolution for Case 1 (top) and Case 3 (bottom).

In both cases, the diversity metrics increase over the course of the optimisation, which indicates that the algorithm continues to generate candidate layouts with a broad spread of fitness values as the search evolves. However, the meaning of this diversity is slightly different in each case. In Case 1, the increase in diversity is accompanied by a gradual decrease in the population mean, suggesting that many candidate solutions generated during the search are penalized by stronger hydrodynamic interference and therefore perform significantly worse than the best layouts. In Case 3, the diversity also remains significant, but it is associated with a generally higher level of performance across the population. In other words, the 30-device configuration still allows the algorithm to explore a wide range of solutions, but many of these candidate layouts remain reasonably efficient because the reduced density mitigates the risk of severe interaction losses. This distinction is important because it shows that the lower-density array does not simply produce a single better solution, rather, it creates a more favourable optimisation landscape in which good-performing layouts are easier to obtain.

A complementary perspective is provided by the evolution of the P2M, shown in Figure 80.

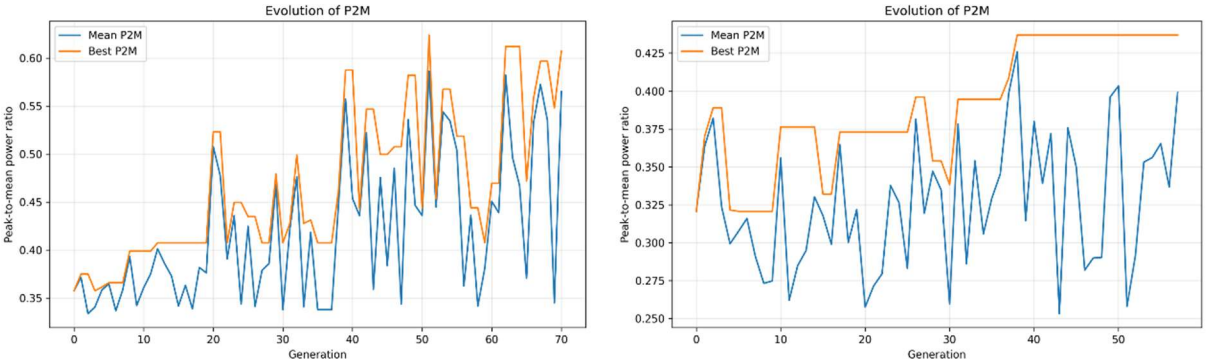


Figure 80 - P2M evolution for Case 1 (top) and Case 3 (bottom).

This metric highlights how uniformly the absorbed power is distributed among the devices in each layout. The difference between the two cases is particularly clear. In Case 1, the best P2M values reach approximately 0.62, whereas in Case 3 the maximum values remain around 0.43. This corresponds to a reduction of nearly 30% when the number of units is reduced from 40 to 30. Such a decrease indicates that the 30-device layout achieves a much more balanced distribution of power absorption across the farm, whereas the denser 40-CorPack case tends to produce stronger disparities between the best-performing and worst-performing devices. From a physical point of view, this is again consistent with the higher degree of interaction expected in the denser array, where some units may benefit from local constructive effects while others are more strongly affected by shadowing and reduced energy availability. In the second case, the larger spacing between devices contributes to a more homogeneous

wave field at device scale, which in turn leads to a more even distribution of absorbed power. This result is important not only from an energy-efficiency perspective but also from an engineering point of view, since lower P2M values are associated with more uniform operating conditions across the array.

For clarity, the main quantitative differences between the two cases are summarised in Table 16.

Table 16 - Quantitative comparison between Case 1 and Case 3.

Metric	Case 1	Case 3	Relative difference
Best CF	0.291	0.307	+5.2%
Best average absorbed power per unit	3.26 MW	3.44 MW	+5.2%
Maximum P2M	0.62	0.43	-30%
Generations to reach optimum	10	38	Slower convergence in Case 3

Overall, the results of this comparison show that reducing the number of CorPacks leads to a consistent improvement in the hydrodynamic performance of the optimised layout at device level. The H_s maps indicate slightly less pronounced local interaction effects, the CF and average-power plots show an increase of about 5% in final performance, and the P2M analysis indicates a substantially more balanced power distribution, with a reduction of approximately 30% in the maximum value of this metric. The main drawback is that the optimisation takes longer to converge, with the 30-device case requiring about 45 generations to reach the optimum compared to around 10 generations in the 40-device case. This suggests that the less crowded deployment offers more spatial flexibility and therefore more opportunities for improvement, but it also demands a longer search before the best solution is identified.

5.5. Influence of the Population Size: Case 3 vs. Case 4

The comparison between Case 3 and Case 4 was carried out to evaluate the influence of the number of candidate arrays explored per generation by the GA. In both cases, the same physical problem was considered: the same study domain, the same turbine layout and exclusion zones, the same number of deployed CorPacks, and the same hydrodynamic representation based on the power matrix approach. The only parameter modified between the two optimisations was the number of candidate layouts generated and evaluated at each generation. In Case 3, the optimisation was performed with 30 arrays per generation, whereas in Case 4 this value was increased to 50 arrays per generation. This comparison is therefore intended to isolate the effect of increasing the search effort of the GA, without changing the physical characteristics of the co-located farm itself.

As shown previously in Figure 69, the optimal layouts obtained for Case 3 and Case 4 remain globally similar. The increase in the number of candidate arrays per generation produces some visible differences in the final arrangement of the CorPacks. In Case 4, the layout appears slightly more irregular, with some devices shifting towards the central corridors between turbine rows and others adopting orientations that differ from those observed in Case 3. These changes are consistent with the larger exploration capability of the GA when more candidate solutions are evaluated at each generation. In practice, increasing the number of arrays does not alter the physical problem being solved, but it expands the portion of the design space sampled by the optimizer during each generation.

To examine whether these layout differences lead to noticeable hydrodynamic changes, Figure 81 presents the H_s fields corresponding to the optimal solutions of Case 3 and Case 4 under the analysed representative sea state.

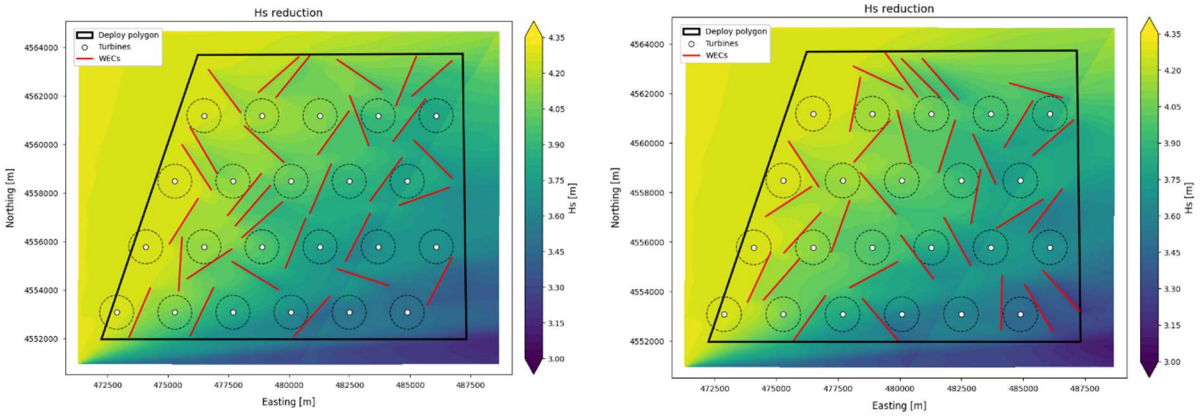


Figure 81 - H_s values for Case 3 (left) and Case 4 (right).

The two H_s fields confirm that the large-scale wave propagation pattern remains essentially unchanged between the two cases. In both maps, higher H_s values are observed in the north-western sector of the domain and progressively lower values towards the south-east, reflecting the same incoming wave conditions. The attenuation zones associated with the device array are also present in both configurations and occupy broadly similar regions of the farm. Nevertheless, some subtle differences can be identified.

In Case 4, the local disturbance pattern appears slightly more concentrated in some of the central deployment corridors, which is consistent with the somewhat more reorganized placement of several devices. By contrast, the attenuation zones in Case 3 appear slightly more evenly distributed across the domain. These differences are relatively modest, which is an important result. Increasing the number of candidate arrays per generation does not fundamentally change the hydrodynamic footprint of the final farm layout. Instead, its main effect is expected to appear in the optimisation behaviour rather than in the physical wave field generated by the resulting array.

The effect of this difference becomes clearer when examining the evolution of the capacity factor during the optimisation process. Figure 82 shows the convergence histories of the CF for Case 3 and Case 4, including the best individual in each generation, the population mean, the elite mean, and the min-max range of the population.

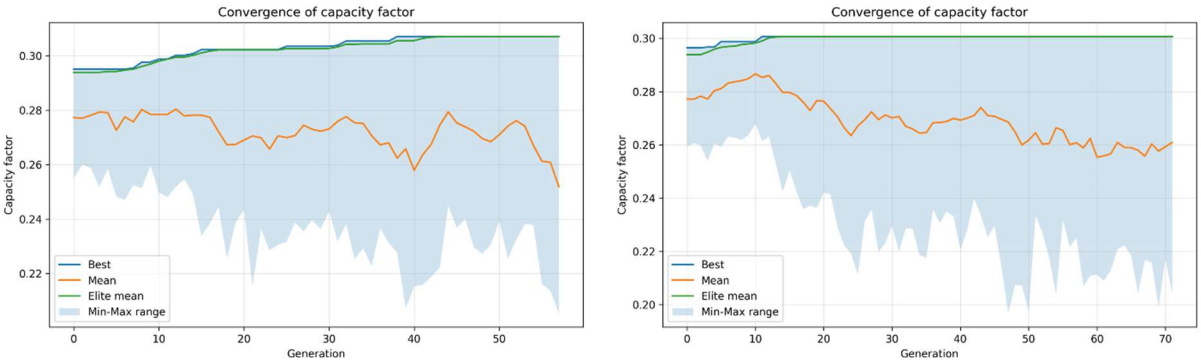


Figure 82 - CF evolution for Case 3 (top) and Case 4 (bottom).

The convergence histories show that the increase in the number of candidate arrays per generation has a clear effect on the search dynamics of the GA. In Case 4, the optimal region of the search space is identified much earlier, with the best solution stabilizing after approximately 12 generations, 33 less generations than Case 3. This indicates that the higher number of candidate layouts evaluated per generation substantially accelerates the discovery of high-performing solutions. However, this faster convergence does not translate into a better final optimum. The best CF obtained in Case 4 is approximately 0.301, 2% lower than the obtained in Case 3.

This is a relatively small variation, but it shows that greater search intensity does not necessarily improve the final solution quality. Instead, it mainly improves the efficiency with which the algorithm reaches a good-performing region of the design space.

The total number of generations executed in each optimisation run is reported in Table 13 of Section 5.1, where the main metrics for all cases are summarised. As indicated there, the optimisation proceeded for a total of 71 generations in Case 4. Therefore, although the best solution in Case 4 was reached much earlier, the full optimisation process continued for a larger total number of generations. This distinction between the generation at which the optimum is found and the total number of generations performed is important for interpreting the computational effort associated with each case.

A complementary perspective is provided by the evolution of the average absorbed power per unit, shown in Figure 83.

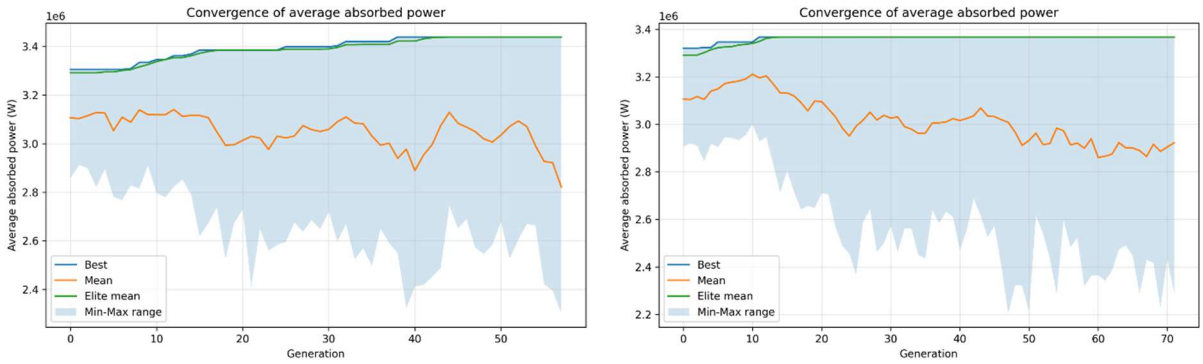


Figure 83 - Average absorbed power evolution for Case 3 (top) and Case 4 (bottom).

The trends observed in this figure are consistent with the CF behaviour. In Case 4, the optimal layout reaches an average absorbed power of approximately 3.35 MW per device, which corresponds to a reduction of roughly 2.5% relative to Case 3. Although this decrease is relatively small, it indicates that the layouts identified when a larger number of arrays are explored per generation do not necessarily translate into improved hydrodynamic performance at device level. Instead, the broader search process appears to maintain a larger set of intermediate solutions across generations without producing a substantially more efficient final configuration.

Further insight into the optimisation dynamics can be obtained from the population diversity indicators presented in Figure 84, where the evolution of the CF range and coefficient of variation are shown for both cases.

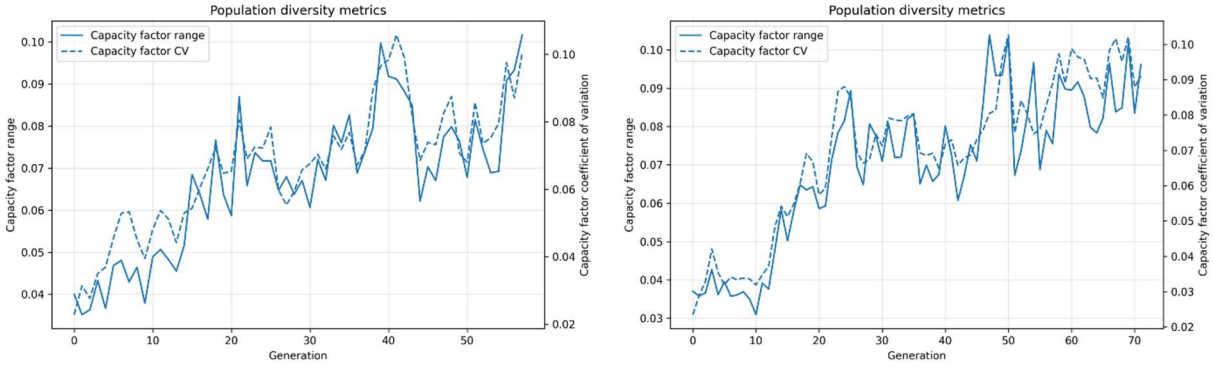


Figure 84 - CF range and coefficient of variation evolution for Case 3 (top) and Case 4 (bottom).

The diversity metrics reveal that Case 4 maintains a slightly larger spread of candidate solutions during most of the optimisation process. Both the coefficient of variation and the CF range tend to reach marginally higher values than those observed in Case 3, particularly in the later generations. This behaviour confirms that increasing the number of candidate arrays per generation enhances the exploratory capability of the GA by preserving a broader diversity of layouts within the population. The larger population reduces the likelihood of premature convergence and allows the algorithm to continue exploring alternative spatial configurations even after high-performing layouts have already been identified.

A final indicator of the optimisation behaviour is provided by the P2M, shown in Figure 85.

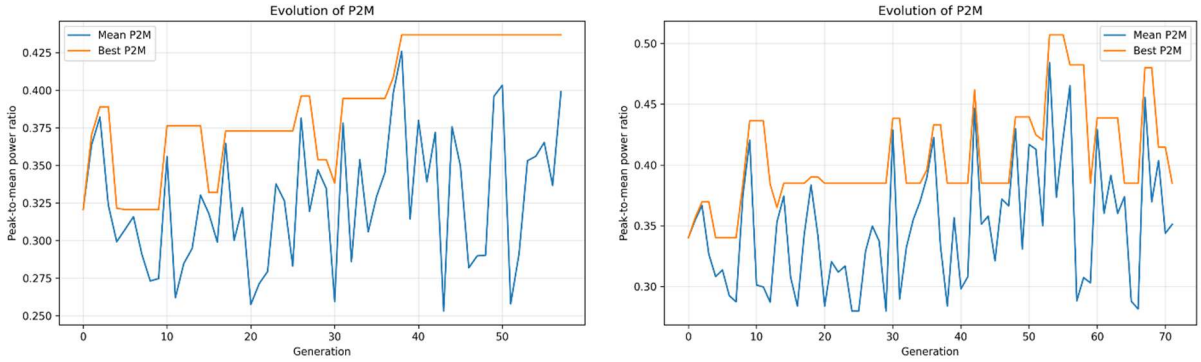


Figure 85 - P2M evolution for Case 3 (top) and Case 4 (bottom).

The results indicate that the layouts generated in Case 4 tend to produce higher P2M values than those obtained in the reference case. The maximum P2M reached in Case 4 is approximately 0.51, which represents an increase of roughly 19% relative to Case 3.

This increase indicates that the distribution of absorbed power among devices becomes slightly less uniform when the optimisation is performed with a larger candidate population. From a physical perspective, this behaviour suggests that the layouts identified in Case 4 may contain more localized clusters of highly productive devices, while other converters operate under slightly less favourable wave conditions. Although this does not necessarily reduce the overall performance of the farm significantly, it does indicate a somewhat larger disparity in device-level energy extraction.

For clarity, the main quantitative differences between the two cases are summarised in Table 17.

Table 17 - Quantitative comparison between Case 3 and Case 4.

Metric	Case 3	Case 4	Relative difference
Best CF	0.307	0.301	-2.0%
Best average absorbed power per unit	3.44 MW	3.35 MW	-2.5%
Maximum P2M	0.43	0.51	+19%
Generations to reach optimum	38	12	Slower convergence in Case 3

In the end, the results of this comparison show that increasing the number of candidate arrays evaluated per generation primarily affects the behaviour of the optimisation process rather than the final hydrodynamic efficiency of the farm. The larger population explored in Case 4 increases the diversity of candidate solutions and prolongs the search process, but it ultimately produces a final layout with slightly lower device-level efficiency and a somewhat less uniform distribution of absorbed power. These results therefore suggest that, within the configuration analysed here, increasing the number of candidate arrays per generation improves exploration but does not necessarily translate into a superior optimal solution.

5.6. Influence of the Crossover Rate: Case 3 vs. Case 5

The comparison between Case 3 and Case 5 was performed to evaluate the influence of the crossover rate on the optimised co-located farm configuration. In both cases, the same hydrodynamic representation was adopted through the power matrix approach, the same number of CorPacks was considered, and the same number of candidate arrays per generation was maintained, so that the only intentional difference between the two runs was the crossover rate used in the GA. Case 3 corresponds to the reference configuration with a crossover rate of 0.70, whereas Case 5 increases this value to 0.95. This comparison is particularly relevant because the crossover operator plays a central role in the exchange of information between parent layouts, and therefore directly affects how efficiently the optimisation process combines favourable spatial features from different candidate solutions.

As previously shown in Figure 69, the optimal layouts obtained for Case 3 and Case 5 already suggest slight differences in spatial organization. Although both configurations fully respect the wind turbine exclusion zones and occupy the same deployable region, the layout obtained with the higer crossover rate appears somewhat less structured in its overall arrangement, with a slightly more irregular distribution of segment orientations across the farm. In contrast, Case 3 presents a more coherent organization of the devices within the available corridors between turbines. This distinction is important because the crossover rate controls how aggressively the algorithm recombines existing candidate layouts, and a lower value of this parameter is therefore expected to limit the rate at which favourable spatial features are propagated through the population.

To illustrate how the two optimised layouts modify the incident wave field, Figure 86 presents the corresponding H_s distributions for Case 3 and Case 5 under the analysed representative sea state.

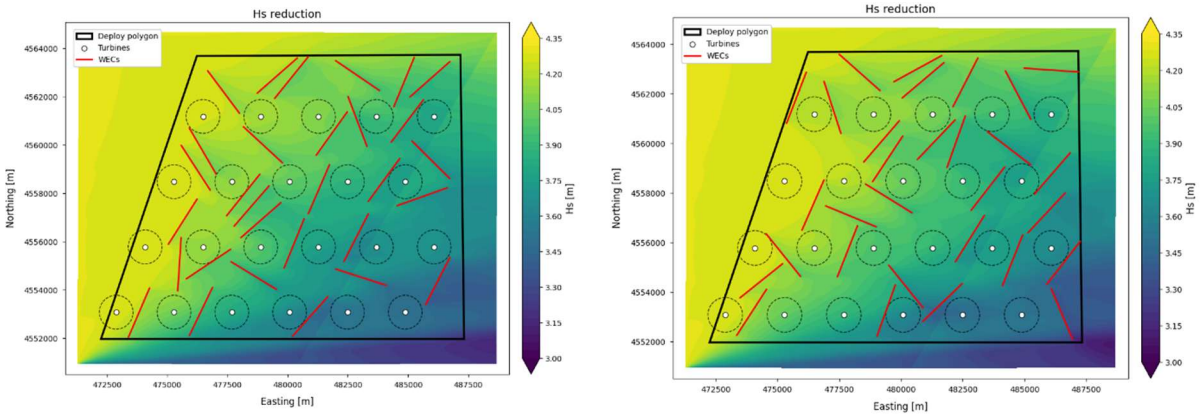


Figure 86 - H_s values for Case 3 (left) and Case 5 (right).

The local disturbances introduced by the arrays reveal some differences between the two optimised solutions. In the reference case, Case 3, the attenuation pattern is relatively coherent across the central part of the domain, with several local shadow regions linking more clearly along the dominant wave propagation direction. In the higher-crossover case, the attenuation zones remain well defined but appear slightly more fragmented and less continuously connected across the farm. In practical terms, this indicates that increasing the crossover rate leads to a final arrangement in which the local hydrodynamic interactions are still present but somewhat less effectively coordinated at array scale. Although the overall wave attenuation footprint remains broadly similar in both cases, the Case 5 layout appears to generate a slightly less organized interaction pattern, suggesting that the higher rate of recombination may reduce the algorithm's ability to assemble highly coherent spatial solutions from previously successful layouts.

This tendency becomes clearer when the evolution of the CF during the optimisation process is examined. Figure 87 shows the evolution of it.

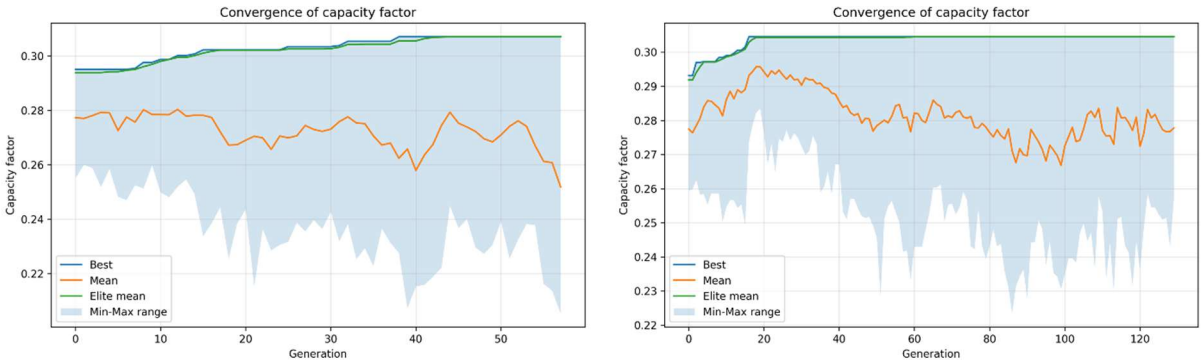


Figure 87 - CF evolution for Case 3 (top) and Case 5 (bottom).

The convergence histories reveal an interesting contrast between the speed at which a competitive solution is first identified and the overall duration of the optimisation. In Case 5, the best CF reaches approximately 0.304, corresponding to a value about 1% lower than that obtained in Case 3. The best individual improves rapidly in the first stages of the search and stabilizes relatively early, which indicates that the higher crossover rate does not prevent the algorithm from finding a good candidate solution. However, as reported in Table 13, the optimisation proceeds for a total of 129 generations, substantially longer than in the reference case. This means that although a near-optimal solution is found

early, the population requires a much longer search process to satisfy the stopping criteria. Such behaviour suggests that increasing the crossover rate weakens the efficiency with which favourable solution features are recombined and propagated through the population. As a result, the algorithm is still capable of identifying good layouts, but the overall refinement of the population becomes slower and less effective. The same conclusion is supported by the evolution of the average absorbed power per device, presented in Figure 88.

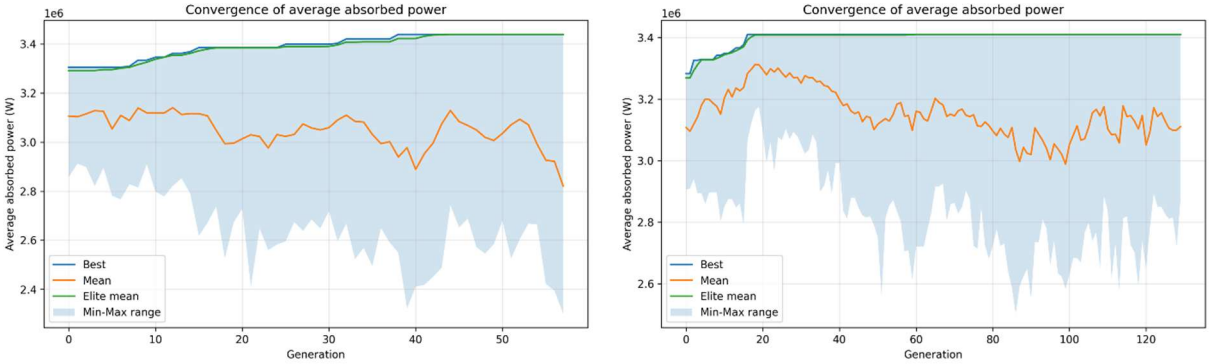


Figure 88 - Average absorbed power evolution for Case 3 (top) and Case 5 (bottom).

The trends shown in this figure are fully consistent with the CF results. In Case 5, the optimal layout reaches an average absorbed power of approximately 3.41 MW per device, which corresponds to a decrease of roughly 1% relative to Case 3. Increasing the crossover rate does not produce a substantial loss in device-level performance, but it does lead to a slightly less efficient final layout. It is also noteworthy that the mean performance of the population remains comparatively high over a large number of generations, even though the best solution shows only limited improvement after the initial stage of the optimisation. This behaviour reinforces the interpretation that the increased crossover rate limits the capacity of the algorithm to combine beneficial spatial traits from different layouts into progressively better offspring.

The behaviour of the population during the optimisation can be further interpreted through the fitness-diversity indicators shown in Figure 89.

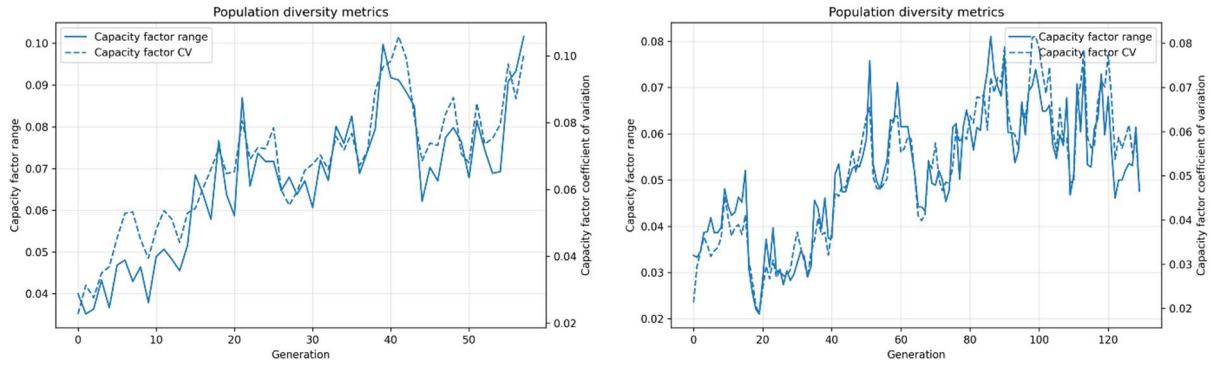


Figure 89 - CF range and coefficient of variation evolution for Case 3 (top) and Case 5 (bottom).

In both cases, the diversity metrics increase over the course of the optimisation, indicating that the GA continues to generate candidate layouts with a spread of fitness values as the search evolves. However, the diversity maintained in Case 5 is generally lower than that observed in the reference case. Both the

CF range and the coefficient of variation remain more limited, and their peaks are visibly reduced relative to Case 3. This suggests that increasing the crossover rate reduces the breadth of effective exploration of the design space. In other words, the algorithm still generates variability within the population, but the increased recombination between parent layouts restricts its ability to create strongly differentiated offspring. This distinction is important because it shows that the higher-crossover configuration does not simply delay convergence; rather, it modifies the search mechanism itself, leading to a more conservative exploration of alternative layouts.

A complementary perspective is provided by the evolution of the P2M, shown in Figure 90.

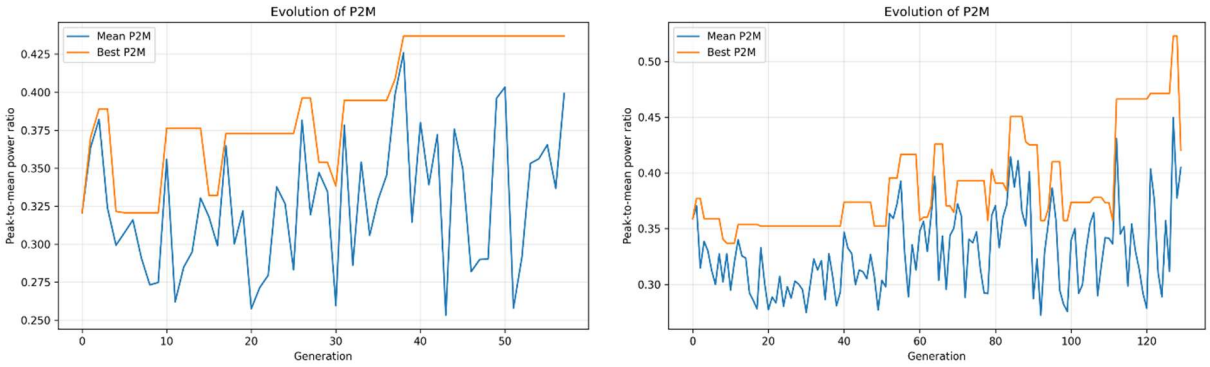


Figure 90 - P2M evolution for Case 3 (top) and Case 5 (bottom).

This metric highlights how uniformly the absorbed power is distributed among the devices in each layout. The difference between the two cases is again informative. In Case 5, the maximum P2M reaches approximately 0.53, corresponding to an increase of about 23% relative to Case 3. Such an increase indicates that the higher-crossover layout produces a less balanced distribution of absorbed power across the farm, with stronger disparities between the most productive and least productive devices. From a physical point of view, this is consistent with the slightly less coordinated attenuation pattern observed in the H_s maps, where the arrangement of the units appears to exploit local wave conditions somewhat unevenly. In the reference case, the more effective recombination of promising layout features appears to favour solutions with a more homogeneous spatial exploitation of the resource, which in turn leads to a more balanced power distribution. The higher P2M observed in Case 5 is therefore important not only as an energy-performance indicator but also as a sign that the higher crossover rate may lead to less uniform operating conditions across the array.

For clarity, the main quantitative differences between the two cases are summarised in Table 18.

Table 18 - Quantitative comparison between Case 3 and Case 5.

Metric	Case 3	Case 5	Relative difference
Best CF	0.307	0.304	-1.0%
Best average absorbed power per unit	3.44 MW	3.41 MW	-1.0%
Maximum P2M	0.43	0.53	+23%
Generations to reach optimum	38	20	Slower convergence in Case 3

Taken together, the results of this comparison show that increasing the crossover rate leads to only a slight decrease in final hydrodynamic performance at device level, but it substantially alters the behaviour of the optimisation process. The H_s maps indicate a somewhat less coordinated attenuation pattern, the CF and average-power plots show a reduction of approximately 1% in final performance, and the P2M analysis indicates a noticeably less balanced power distribution, with an increase of about 23% in the maximum value of this metric. At the same time, the total number of generations required by the optimisation increases markedly, showing that the higher-crossover configuration slows down the overall refinement of the population even when a competitive best solution is found early. This suggests that a lower crossover rate is advantageous not only because it supports efficient exploration of the design space, but also because it helps the algorithm combine favourable layout characteristics into more coherent and hydrodynamically balanced solutions.

5.7. Influence of the Mutation Rate: Case 3 vs. Case 6

The comparison between Case 3 and Case 6 was undertaken to evaluate the influence of the mutation rate on the optimised co-located farm configuration. Case 3 corresponds to the reference configuration with a mutation rate of 0.01, whereas Case 6 increases this value to 0.025. This comparison is particularly relevant because the mutation operator is responsible for introducing random variations into the offspring layouts and therefore plays a central role in preserving exploration capability during the optimisation. A lower value of this parameter is expected to decrease the rate at which new spatial traits are introduced into the population, potentially affecting both the diversity of the search and the robustness of the final solution.

As previously shown in Figure 69, the optimal layouts obtained for Case 3 and Case 6 already suggest some differences in spatial organisation. The layout obtained in Case 6 appears slightly less coherent in the way the CorPack segments are arranged across the farm, their orientations appear more irregular and the array seems less evenly organized across the domain, while in the reference case, the devices tend to occupy the available corridors with a somewhat more structured distribution.

This distinction is important because the mutation rate affects the generation of new candidate configurations throughout the optimisation process, and a reduction in this parameter may therefore limit the ability of the algorithm to escape from locally favourable but globally suboptimal patterns.

To illustrate how these two optimised layouts modify the incident wave field, Figure 91 presents the corresponding H_s distributions for Case 3 and Case 6 under the analysed representative sea state.

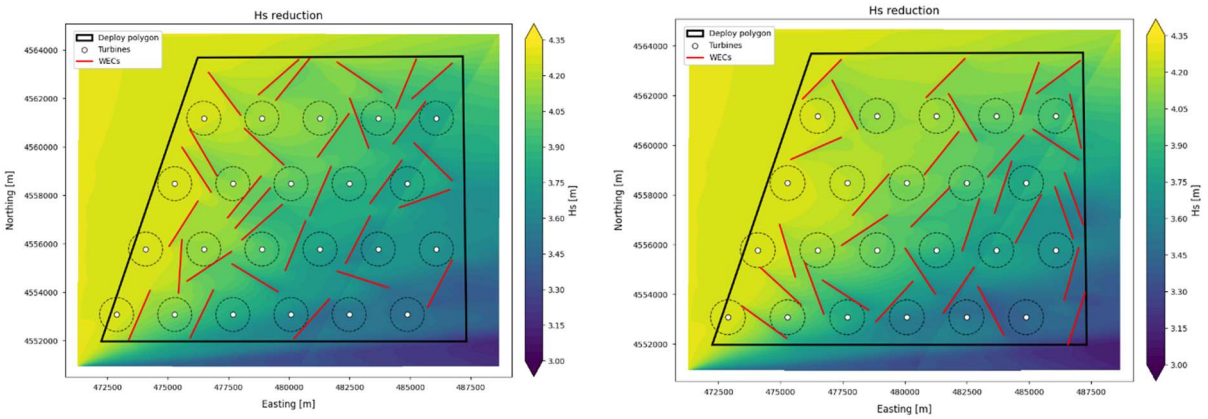


Figure 91 - H_s values for Case 3 (left) and Case 6 (right).

As shown in Figure 91, the local disturbances introduced by the arrays reveal meaningful differences between the two optimised solutions. In the reference case, the attenuation pattern across the farm remains relatively coherent, with several local shadow regions linking more smoothly along the dominant wave propagation direction.

In Case 6, the attenuation zones remain clearly visible but appear more fragmented, particularly in the central and southern portions of the domain, where the reduction pattern is less continuous and more irregular. In practical terms, this indicates that the higher mutation rate leads to a final arrangement in which the local wave-device interactions are still present but less effectively coordinated at array scale.

Although the overall wave attenuation footprint remains similar in both cases, the Case 6 layout appears to produce a somewhat less organized interaction field, suggesting that the increased introduction of new spatial variations may constrain the algorithm's ability to refine the array towards a more hydrodynamically balanced configuration.

This tendency becomes clearer when the evolution of the CF during the optimisation process is examined. Figure 92 shows the convergence of the CF for both cases.

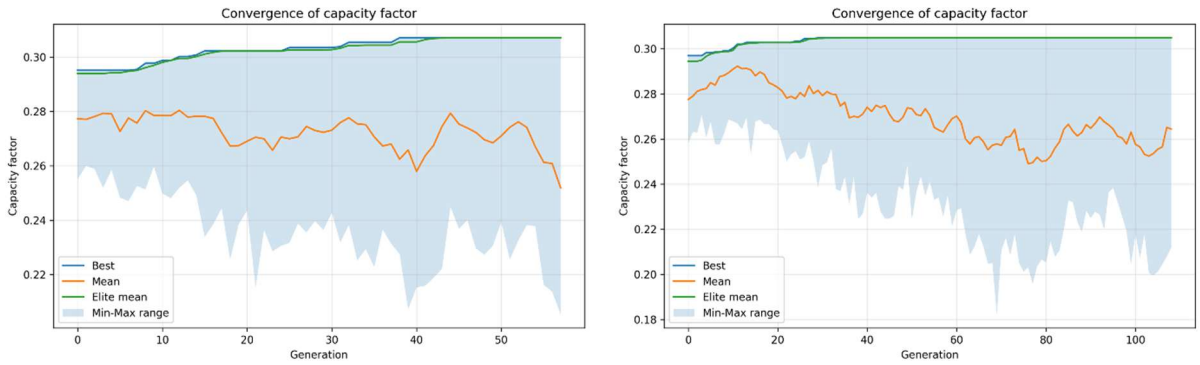


Figure 92 - CF evolution for Case 3 (top) and Case 6 (bottom).

The convergence histories indicate that increasing the mutation rate changes both the speed and the nature of the search process. In Case 6, the best CF reaches approximately 0.304, corresponding to a value about 1% lower than that obtained in Case 3. The best individual improves rapidly in the first stages of the search and stabilizes after roughly 30 generations, indicating that the optimisation can identify a competitive solution relatively early. However, as reported in Table 13, the run continues for a total of 108 generations, considerably longer than the reference case.

This behaviour suggests that the higher mutation rate allows the GA to locate a near-optimal solution but reduces its ability to keep generating sufficiently distinct alternatives to improve the population more efficiently afterwards. In other words, the optimisation still converges towards a good-performing layout, but the broader refinement of the search becomes slower and less effective when the amount of random variation introduced at each generation is reduced.

The same conclusion is supported by the evolution of the average absorbed power per CorPack, presented in Figure 93.

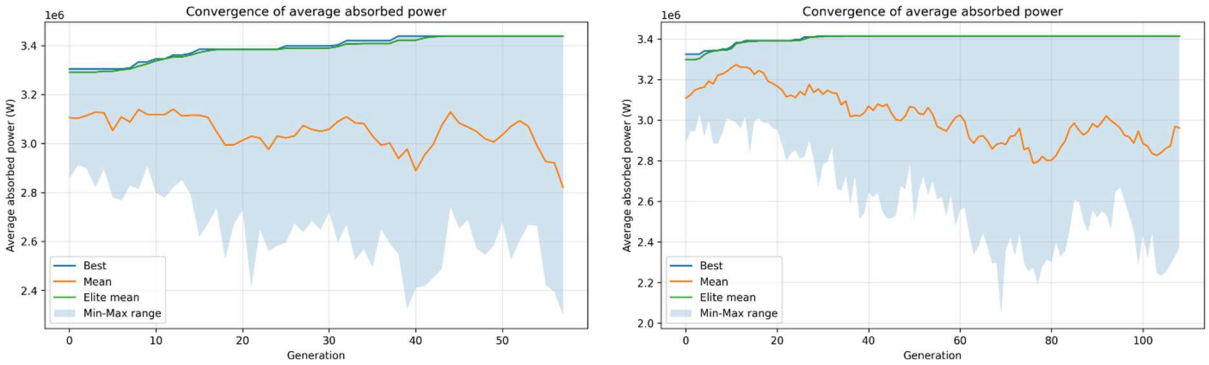


Figure 93 - Average absorbed power evolution for Case 3 (top) and Case 6 (bottom).

The trends shown in this figure are fully consistent with the CF results. In Case 6, the optimal layout reaches an average absorbed power of approximately 3.41 MW per device, corresponding again to a decrease of about 1% relative to Case 3. Increasing the mutation rate does not drastically reduce the final energy captured by each converter, but it does lead to a slightly less efficient solution. It is also worth noting that the mean performance of the population remains reasonably high across many generations, even though the best solution exhibits only limited gains after the early part of the optimisation. This reinforces the interpretation that the higher mutation rate weakens the capacity of the GA to keep introducing beneficial variations once the first competitive layouts have already emerged.

The behaviour of the population during the optimisation can be further interpreted through the fitness-diversity indicators shown in Figure 94, where the evolution of the CF range and the coefficient of variation are plotted for both cases.

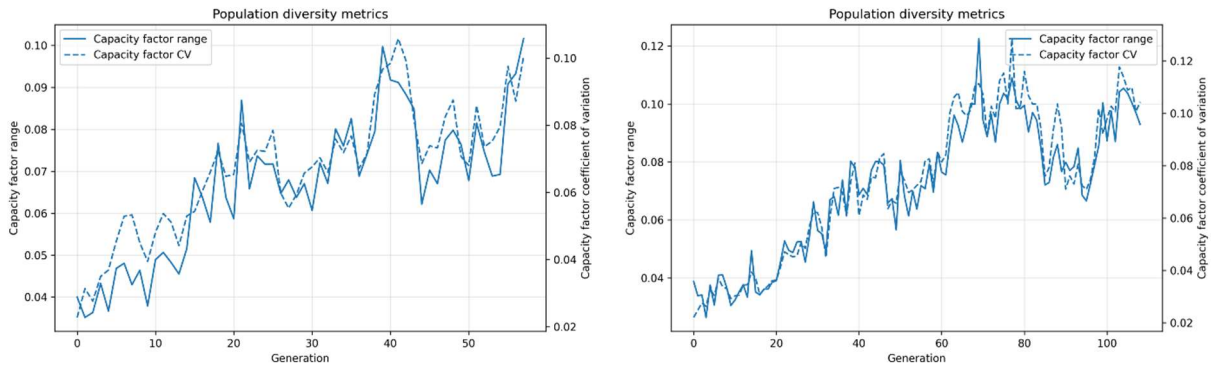


Figure 94 - CF range and coefficient of variation evolution for Case 3 (top) and Case 6 (bottom).

In both cases, the diversity metrics increase over the course of the optimisation, which indicates that the algorithm continues to generate candidate layouts with a broad spread of fitness values as the search evolves. However, the diversity maintained in Case 6 becomes generally higher than that observed in the reference case, particularly in the intermediate and later stages of the run. Both the CF range and the coefficient of variation rise to values above those reached in Case 3, showing that the population retains a wider dispersion of performance levels throughout the search. This behaviour is significant because it indicates that increasing the mutation rate does not necessarily reduce diversity in a simple monotonic sense. Instead, it appears to delay the effective refinement of the population, allowing a broader spread of partially improved but unevenly performing layouts to persist for longer. This distinction is important because it shows that the higher-mutation configuration does not simply slow the search, it also alters

the way in which the population evolves, with the algorithm taking longer to consolidate the most favourable solution traits into a more uniformly high-performing set of candidate layouts.

A complementary perspective is provided by the evolution of the P2M, shown in Figure 95.

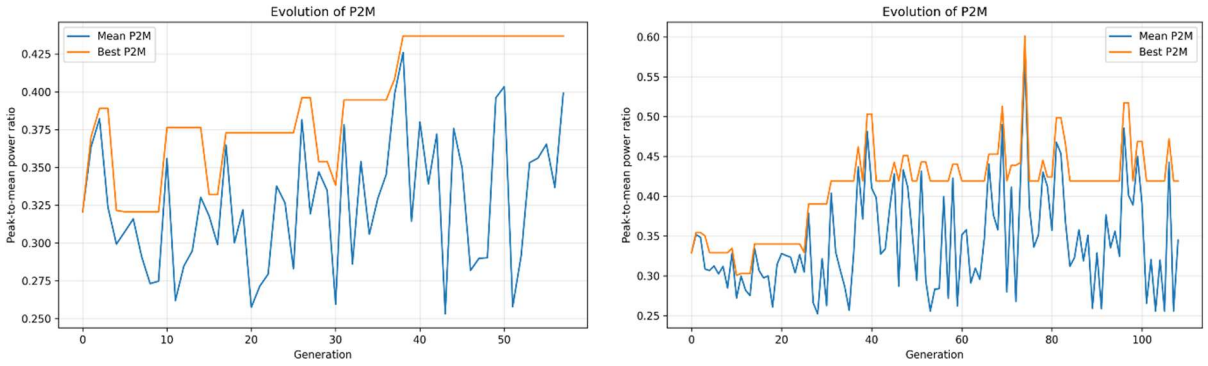


Figure 95 - P2M evolution for Case 3 (top) and Case 6 (bottom).

This metric highlights how uniformly the absorbed power is distributed among the devices in each layout. The difference between the two cases is particularly clear. In Case 6, the maximum P2M reaches approximately 0.60, which corresponds to an increase of about 40% relative to Case 3. Such an increase indicates that the higher-mutation layout produces a much less balanced distribution of absorbed power across the farm, with stronger disparities between the most productive and least productive devices. From a physical point of view, this is consistent with the more fragmented attenuation pattern observed in the H_s maps, where the arrangement of the devices appears to exploit the incident wave field less uniformly. In the reference case, the lower mutation rate seems to support a more balanced spatial exploration, favouring solutions in which the resource is distributed more homogeneously across the array. The much higher P2M observed in Case 6 is therefore relevant not only as an energy-performance indicator but also as evidence that increasing the mutation rate can compromise the uniformity of operating conditions within the farm.

For clarity, the main quantitative differences between the two cases are summarised in Table 19.

Table 19 - Quantitative comparison between Case 3 and Case 6.

Metric	Case 3	Case 6	Relative difference
Best CF	0.307	0.304	-1.0%
Best average absorbed power per WEC	3.44 MW	3.41 MW	-1.0%
Maximum P2M	0.43	0.60	+40%
Generations to reach optimum	38	30	Slower convergence in Case 3

Overall, the results of this comparison show that increasing the mutation rate leads to only a slight decrease in final hydrodynamic performance at device level, but it substantially modifies the behaviour of the optimisation process. The H_s maps indicate a somewhat less coordinated attenuation pattern, the CF and average-power plots show a reduction of approximately 1% in final performance, and the P2M analysis indicates a markedly less balanced power distribution, with an increase of about 40% in the

maximum value of this metric. At the same time, the total number of generations required by the optimisation increases substantially, showing that the higher-mutation configuration slows the overall evolution of the search even when a competitive solution is found relatively early. This suggests that maintaining a sufficiently lower mutation rate is beneficial not only because it supports continued exploration of the design space, but also because it helps the algorithm refine layouts towards solutions that are both hydrodynamically effective and more evenly balanced at device level.

5.8. Wave-Height Reduction due to the Presence of CorPacks

Although the optimisation framework in this work was not explicitly formulated to minimise H_s within the wind farm domain, the reduction of it induced by the presence of the device arrays remains an important aspect of the final co-located configurations and therefore deserves specific discussion. In addition to their energy-conversion role, WECs modify the incident wave field through energy extraction, producing localised attenuation zones downstream of the deployed arrays. From an engineering perspective, this effect is relevant because it may contribute to milder wave conditions within parts of the farm, which in turn may be beneficial for offshore operations.

In particular, lower local H_s may reduce the hydrodynamic loading experienced by FOW turbines and associated infrastructure, while also contributing to more favourable access conditions for maintenance and inspection activities. Although this aspect was not used as an optimisation objective in the present work, it represents a potentially valuable secondary effect of wave-wind co-location.

To visualise this effect more clearly, an additional hydrodynamic simulation was performed for the same incident conditions adopted in the wave-height plots discussed in the previous comparative sections, but without the presence of WECs. This reference case allows the wave field inside the study domain to be compared directly with the corresponding cases in which CorPack arrays are present. The H_s distributions associated with the optimised layouts were previously presented in Figure 71, Figure 76, Figure 81, Figure 86 and Figure 91, and are also compiled in Appendix C, whereas the corresponding reference case without WECs is shown in Figure 96.

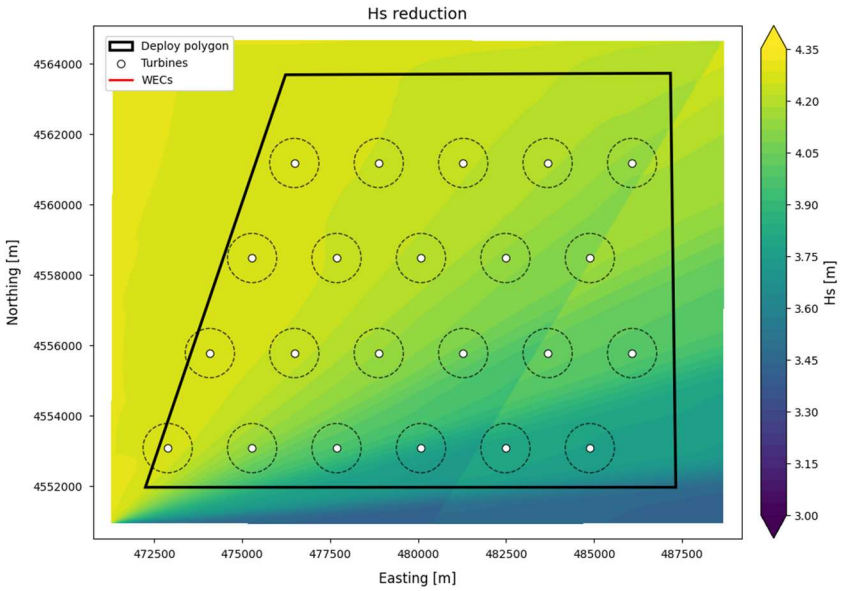


Figure 96 - H_s field in the study domain without WECs.

The reference wave field without WECs displays the expected large-scale propagation pattern across the domain. In this configuration, the wave field remains comparatively smooth, with no additional internal attenuation zones beyond those associated with the natural transformation of the incoming sea state across the domain.

When this reference case is compared with the H_s fields obtained in the presence of WECs, the influence of the wave-energy layer becomes immediately visible. In all optimised layouts, the presence of the CorPack arrays introduces additional zones of wave attenuation within the deployable region, producing a more structured internal H_s field than in the no-WEC case. These reductions are particularly noticeable along the wave-propagation path downstream of the segments, where localised shadow regions develop and partially overlap depending on the layout geometry. As a result, the wave field within the farm is no longer governed solely by offshore propagation and bathymetric transformation, but also by the cumulative extraction of wave energy by the WEC layer.

From a global perspective, the most important point is therefore not the specific difference between individual optimisation cases, but the consistent contrast between the presence and absence of WECs.

In the no-WEC configuration, the farm interior remains exposed to the incident wave field with comparatively limited internal attenuation. By contrast, in all six optimised layouts, the addition of the WEC layer produces a measurable reduction in H_s across parts of the domain. The exact spatial pattern varies according to the distribution and orientation of the segments, but the overall effect remains the same: the co-located wave-energy layer acts as an internal attenuating element that modifies the wave climate experienced within the wind-farm area.

This point is relevant because it shows that the integration of WECs within a FOW farm may offer benefits beyond electricity production alone. By extracting wave energy before it propagates further into the farm, the CorPack arrays can contribute to locally reduced sea-state severity in some areas of the domain. Such a reduction may be advantageous from an operational perspective, especially in relation to turbine accessibility and offshore intervention windows. In practical terms, if local wave attenuation contributes to a greater proportion of time during which the threshold condition presented near the end of Section 2.3.3. ($H_s < 1.5$ m) is satisfied, then the accessible time available for maintenance and service operations may also increase. Likewise, lower incident wave conditions within part of the farm may contribute to reduced wave loading on floating structures and associated components.

At the same time, this effect should be interpreted carefully. The present optimisation framework was to maximise array performance under realistic spatial constraints, not to optimise coastal protection or operational accessibility. Consequently, the H_s reduction observed in the results should be regarded as a secondary but relevant outcome of the co-location strategy, rather than as the primary target of the design process. Even so, its consistent presence across all cases indicates that the wave-energy layer does not merely coexist spatially with the wind farm but also modifies the internal hydrodynamic environment in a way that may have practical engineering implications.

6. CONCLUSIONS

6.1. MAIN FINDINGS

This dissertation investigated the optimisation of large-scale co-located offshore wind-wave energy arrays under realistic Portuguese offshore conditions, with particular focus on the spatial design of the wave energy layer within a predefined floating offshore wind farm. Rather than treating co-location as a generic conceptual possibility, the work addressed it as a concrete project-scale design problem in which the host wind farm defines the primary infrastructure and the wave energy system must subsequently be integrated within the residual admissible space. In this sense, the study was intentionally structured to reflect a realistic offshore development sequence, where the wind layer is established first and the wave layer is then positioned within the remaining spatial envelope. The results show that this approach is not only methodologically coherent, but also capable of producing feasible and hydrodynamically efficient co-located layouts at a meaningful engineering scale.

A central outcome of the work is the demonstration that WECs can be incorporated within the constrained internal space of a floating offshore wind farm in a way that remains compatible with realistic geometric and operational restrictions. In the adopted case study, the fixed wind farm layer was defined through a 22-turbine configuration corresponding to approximately 440 MW, while the wave layer was represented through 40, or 30, CorPack modules, corresponding to an additional equivalent wave capacity of approximately 400 MW and 1120 equivalent C4 devices. After applying the turbine exclusion zones, the original 152 km² study area was reduced to a first-order residual deployable area of approximately 107.77 km², and it is within this remaining domain that the wave layer was optimised. This result is particularly significant because it shows that the apparent “dead spaces” left by the wind farm are not merely unusable residual zones but can instead constitute a meaningful opportunity for additional energy production within the same offshore concession area.

Another important conclusion concerns the hydrodynamic representation of the CorPack system. Since no aggregated power matrix for the CorPack cluster was available in the literature, a new representation was constructed in this work to enable its implementation within the optimisation procedure. An RCW-based representation was also derived for the same cluster, and both formulations were validated against the SafeWAVE reference results. The comparison showed that both approaches reproduce the overall absorption behaviour of the CorPack with acceptable agreement for project-scale modelling, with the power matrix representation slightly underestimating and the RCW formulation slightly overestimating the reference result, while remaining within an approximately ±10% deviation interval.

This validation step is important because it provides confidence that the reduced-order representations adopted in the optimisation preserve the essential hydrodynamic behaviour of the CorPack system while remaining computationally tractable.

The optimisation campaign further showed that the adopted workflow can identify stable and consistent layout solutions across different modelling assumptions and algorithm settings. The analysed cases converged toward CF values in a relatively narrow interval, approximately between 0.29 and 0.31, indicating that the optimal arrays obtained under the different scenarios remain broadly comparable in hydrodynamic terms. The best results were obtained for the case with 30 units using the power matrix representation and for the case with increased mutation rate, both reaching CF values of approximately 0.306. At the same time, the comparison between the power matrix and RCW formulations showed that similar global performance levels may still lead to layouts with different spatial characteristics, highlighting that modelling choices influence not only the predicted energy absorption but also the geometric structure of the final solution.

In summary, the results of this dissertation support a clear final conclusion: co-located offshore wind-wave systems should not be understood simply as two technologies sharing the same maritime zone, but as a spatially coupled design problem in which the geometry of the floating wind farm, the remaining deployable area, the hydrodynamic representation of the wave system, and the optimisation strategy all interact to determine the final layout. Within this context, the present work shows that it is possible to use the residual space of a floating offshore wind farm to incorporate a substantial wave energy layer and thereby increase the total energy-production potential of the same offshore development area.

More broadly, the thesis contributes to the view that the future use of offshore concession zones may benefit from integrated planning approaches in which wave energy is considered not as an external add-on, but as a technology that can be meaningfully embedded within the spatial structure of next-generation floating wind projects.

6.2. SCIENTIFIC AND METHODOLOGICAL CONTRIBUTIONS

The main scientific and methodological contribution of this dissertation lies in the application of an established optimisation workflow to a new large-scale co-located offshore wind-wave case study, supported by the construction and validation of problem-specific hydrodynamic representations for the wave energy system. In this sense, the originality of the work does not reside in the development of a new optimisation framework, but rather in the way an existing methodology was adapted and implemented to address a realistic project-scale scenario in which CorPacks are integrated within the constrained internal space of a FOW farm.

One of the most relevant contributions of the thesis is the development of an aggregated power matrix representation for the CorPack cluster. Since no such representation was available in the literature, its construction was a necessary step for enabling the inclusion of this technology within the spectral wave modelling and optimisation procedure adopted in the present work. This contribution is particularly significant because it transforms a technology concept that was not directly available in a form suitable for large-scale optimisation into a hydrodynamic representation that can be consistently implemented within the SNL-SWAN environment. In parallel, an RCW-based representation of the same aggregated CorPack system was also derived, providing a second simplified formulation for the same device cluster and enabling a structured comparison between two alternative modelling approaches.

A second important contribution lies in the validation of these two aggregated CorPack representations against the detailed hydrodynamic results available in the literature. Given the limited number of studies available for this specific system, this validation step was essential to establish confidence in the simplified formulations adopted for project-scale modelling. The results showed that both the power matrix and the RCW representations reproduce the global absorption behaviour of the CorPack cluster

with acceptable agreement for the intended application, thereby supporting their use as reduced-order representations in large-scale optimisation studies.

This aspect is particularly relevant because it provides a methodological basis for future work involving CorPack-type systems in spectral wave models, while also clarifying the degree of agreement that can be expected when simplified representations are used in place of more detailed hydrodynamic descriptions.

The dissertation also contributes through the implementation of a large-scale co-located optimisation application under realistic offshore conditions. The analysed case combines a floating wind farm of substantial scale with a large wave energy layer represented through multiple CorPack clusters, resulting in a project configuration whose total energy capacity approaches the scale of future commercial offshore developments. By working at this scale, the thesis moves beyond small conceptual array studies and shows how co-located wave-wind optimisation can be explored within an engineering context that is much closer to practical offshore planning.

Another relevant contribution of the work is the direct comparison between the power matrix and RCW representations within the same optimisation study. The results showed that both hydrodynamic formulations lead to similar levels of global performance in terms of q-factor, while at the same time producing layouts with different spatial characteristics. This is an important outcome because it demonstrates that the hydrodynamic representation selected for the wave energy system may influence not only the predicted absorption behaviour, but also the geometry of the final optimal layouts. The thesis therefore contributes to a more careful interpretation of optimisation results, showing that similar scalar performance indicators do not necessarily imply equivalent spatial solutions.

A further contribution of the dissertation lies in the analysis of how the wave energy layer behaves when deployed within the residual space left by a predefined floating offshore wind farm. By treating the wind farm as the primary infrastructure layer and positioning the wave energy system within the remaining admissible area, the work reflects a realistic offshore development sequence and provides a methodological basis for studying how co-location can be implemented under practical spatial constraints. In this respect, the thesis contributes not only to optimisation analysis, but also to the broader understanding of how so-called unused or residual areas inside offshore wind farms may be reinterpreted as zones with additional energy-production potential.

Finally, the optimisation study also contributes through the structured assessment of sensitivity to key case-study and algorithm parameters. By comparing different numbers of deployed WEC clusters and different GA settings, the work provides additional insight into the robustness of the obtained solutions and into the extent to which layout results depend on modelling and optimisation choices. This is relevant because it complements the physical comparison between hydrodynamic representations with a methodological evaluation of the optimisation itself, helping to identify which assumptions have the greatest influence on the final layout configurations obtained for large-scale co-located offshore renewable systems.

6.3. Engineering and Planning Implications

Beyond the methodology presented in this work, the results obtained in this dissertation also carry important implications for the engineering design and planning of future offshore renewable energy systems. In particular, the findings highlight the potential of co-located wave-wind energy farms as a strategy to improve the utilisation of offshore space while increasing the overall energy output of marine energy projects.

One of the most relevant outcomes of this study is the demonstration that WECs can be deployed within the residual spatial envelope of FOW farms without interfering with the primary wind infrastructure. In the case study considered, a FOW composed of 22 turbines with an installed capacity of approximately 440 MW defines the main energy-generation layer. After accounting for turbine exclusion zones and operational buffers, a substantial portion of the study domain remains available for secondary energy deployment. The optimisation results show that this residual space can accommodate arrays of aggregated WEC units representing an additional wave-energy capacity of approximately 400 MW under the adopted modelling assumptions. In practical terms, this indicates that a co-located offshore development of roughly 840 MW of combined wind and wave capacity could be achieved within a single concession area while maintaining spatial compatibility between the two technologies.

This result is particularly relevant from a spatial-planning perspective. Offshore wind farms are typically designed with large internal spacing between turbines to reduce aerodynamic wake interactions and ensure safe operation of floating platforms and mooring systems. Therefore, significant portions of the concession area remain unoccupied by physical infrastructure. The present study shows that part of this internal space may be used to host WECs, effectively transforming the wind farm into a multi-layer offshore energy system. From an engineering standpoint, this approach allows the available maritime area to be utilised more efficiently while maintaining the operational integrity of the wind farm.

Another important implication concerns the role of optimisation-based design in the early stages of offshore project development. The results demonstrate that the spatial organisation of WEC arrays within wind farms cannot be determined solely through geometric considerations or simple grid-based layouts. Instead, the performance of the system depends strongly on the interaction between wave propagation, energy extraction by individual devices, and the spatial constraints imposed by the wind infrastructure. By coupling hydrodynamic modelling with evolutionary optimisation, the methodology used in this work provides a practical tool for identifying deployment configurations that balance these competing factors. Such approaches may therefore assist project developers and planners during the conceptual design phase of multi-technology offshore farms.

The results also suggest that device density within the wave-energy layer plays a key role in determining the efficiency of the system. While higher numbers of devices increase the total installed capacity, they also intensify hydrodynamic interactions between neighbouring converters. The optimisation cases analysed in this study indicate that reducing device density can lead to improved energy absorption per unit and a more balanced distribution of power across the array. From a design perspective, this highlights the importance of carefully balancing installed capacity with interaction effects when determining the optimal number of devices within a co-located system.

At a broader scale, the methodology and results presented in this dissertation contribute to the discussion on the future development of multi-technology offshore renewable energy systems. As offshore wind deployment continues to expand globally, the integration of additional marine energy technologies within existing wind-farm infrastructure is increasingly being considered as a strategy to enhance the overall productivity of offshore concession areas. The co-location concept explored in this work aligns with this perspective by demonstrating that wave-energy systems can be integrated within wind farms in a structured and technically informed manner.

In the specific context of the Portuguese Atlantic coast, where wave resources are among the most energetic in Europe, such integrated developments could play a particularly relevant role. The results obtained in this study illustrate how optimisation-based methodologies can support the evaluation of potential co-located configurations within designated offshore development zones, such as the areas considered under the Portuguese offshore renewable-energy planning framework. In this sense, the

present work provides a methodological basis that may assist future feasibility assessments of hybrid offshore developments combining floating wind and wave energy technologies.

Overall, the findings reinforce the idea that offshore renewable energy systems should not be approached as isolated technological solutions occupying separate areas of the ocean. Instead, they can be conceptualised as integrated energy landscapes in which multiple technologies coexist and interact within the same maritime space. By exploring this perspective, the present dissertation contributes to the ongoing transition toward more efficient and spatially coordinated offshore energy systems.

6.4. CONTRIBUTION TO THE UNITED NATIONS SUSTAINABLE DEVELOPMENT GOALS

This section presents an analysis of the alignment of the present dissertation with the United Nations Sustainable Development Goals (SDGs), highlighting its contribution to sustainable development within the context of offshore renewable energy systems.

The main contributions of this dissertation are mapped against the most relevant SDGs, their corresponding targets, and associated performance indicators, as summarised in Table 20.

Table 20 - Contribution of the dissertation to the United Nations Sustainable Development Goals, including associated targets and performance indicators.

SDG	Goal	Contribution	Performance indicators and metrics
7 - Affordable and Clean Energy	7.2 - Increase substantially the share of renewable energy	The dissertation uses an optimisation framework for co-located offshore wind and wave energy systems, enabling increased energy production within the same maritime area.	Total installed capacity; Increase in energy production; Capacity factor; Power-to-mean ratio.
	7.3 - Improve energy efficiency	The optimisation of CorPack layouts improves the hydrodynamic efficiency of wave energy extraction.	q-factor; Average absorbed power; Optimisation convergence performance.
9 - Industry, Innovation and Infrastructure	9.5 - Enhance scientific research and innovation	The dissertation introduces an innovative methodology combining hydrodynamic modelling and genetic algorithms for offshore energy systems.	Number of optimisation cases analysed; Robustness of solutions; Computational performance.
13 - Climate Action	13.2 - Integrate climate measures into policies and planning	The work supports the integration of renewable offshore systems into energy planning, contributing to decarbonisation strategies.	Renewable energy production; CO_2 emissions avoided.
14 - Life Below Water	14.2 - Sustainably manage marine ecosystems	The methodology incorporates spatial constraints (e.g., exclusion zones), promoting environmentally responsible deployment of offshore infrastructure.	Area of constrained vs usable space; Compliance with spatial constraints; Reduction of unnecessary marine footprint

The increasing need for decarbonisation of the energy sector and the efficient use of marine space reinforces the importance of engineering solutions that not only improve energy production but also promote environmental sustainability and responsible resource management. In this context, the optimisation framework used in this work, based on the co-location of offshore wind and wave energy systems, directly addresses key global challenges related to clean energy generation, climate change mitigation, and sustainable use of ocean resources.

6.5. LIMITATIONS AND RECOMMENDATIONS FOR FUTURE WORKS

While the results presented in this dissertation provide useful insights into the optimisation of large-scale co-located offshore wind-wave energy systems, several limitations associated with the modelling approach and the scope of the case study should be acknowledged.

A first limitation concerns the hydrodynamic representation of the wave energy system. To enable project-scale optimisation within a computationally feasible framework, the CorPack clusters were represented through aggregated performance formulations rather than through explicit modelling of each individual WEC unit. Although both the power matrix and the RCW representations were validated against available reference results and shown to reproduce the overall absorption behaviour of the cluster with acceptable accuracy, these simplified formulations do not capture the detailed intra-cluster hydrodynamic interactions that may occur between individual devices. Future studies could therefore investigate higher-resolution modelling approaches in which aggregated representations are complemented by detailed simulations for selected configurations.

A second limitation relates to the representation of the offshore wave climate used in the optimisation process. The modelling approach relies on a reduced set of representative sea states derived from the long-term dataset to maintain a tractable number of SWAN simulations during the optimisation procedure. While this strategy is widely adopted in array optimisation studies, it inevitably introduces a simplified description of the full variability of the offshore wave climate. Future work could explore alternative clustering techniques or larger representative sea-state sets to evaluate the potential influence of this reduction on the optimisation results.

The spatial configuration of the FOW farm was also kept fixed throughout the optimisation. In the present work, the objective was to analyse how a wave energy layer could be incorporated within a predefined wind farm geometry, reflecting a realistic development sequence in which the wind infrastructure is established first and the wave system is integrated afterwards. However, future offshore projects may benefit from more integrated design strategies in which both the wind and wave layers are considered simultaneously. Extending the methodology toward coupled wind-wave layout optimisation could therefore represent a valuable direction for further research.

Additional simplifications are associated with the spatial constraints considered in the optimisation procedure. The present study primarily accounts for turbine exclusion zones and for the overall deployable polygon of the study area. In practical offshore developments, however, additional constraints may also influence the feasible layout of co-located systems, including electrical infrastructure routing, mooring footprints, navigation corridors, environmental restrictions, and installation logistics. Incorporating such constraints into the optimisation framework would allow future studies to further approximate real project planning conditions.

The scope of the present work is also limited to a single case study location within the Portuguese P3 offshore concession zone. Although the methodological approach adopted here is general and could be applied to other offshore environments, the quantitative results obtained in this dissertation remain site-

specific and reflect the particular characteristics of the selected study area, including its bathymetry, wave climate, and spatial constraints. Extending the analysis to additional offshore sites would therefore help evaluate the generality of the trends identified in this work.

Finally, the optimisation objective used in this study focuses primarily on hydrodynamic performance through the CF metric. Other aspects that are highly relevant for real offshore developments, such as economic performance, electrical infrastructure requirements, installation costs, and maintenance logistics, were not explicitly included in the optimisation formulation. Future work could therefore incorporate techno-economic indicators into the optimisation objective to evaluate co-located offshore renewable systems from a broader project-development perspective.

Overall, these considerations suggest several promising directions for future research. Further work could focus on improving the hydrodynamic representation of WEC clusters, expanding the treatment of environmental conditions, exploring coupled wind-wave layout optimisation strategies, and incorporating additional engineering and economic constraints. Such developments would contribute to bridging the gap between conceptual optimisation studies and the practical design of large-scale co-located offshore renewable energy projects.

REFERÊNCIAS BIBLIOGRÁFICAS

[1] IEA, *Renewables 2024*, 2024. Available: www.iea.org

[2] REN21, *Global status report 2025*, 2025.

[3] European Commission, *European Comission (2025) - Offshore renewable energy*, 2025.

[4] GWEC, *Global Wind Report 2025*, Clean, 2025. Available: www.gwec.net

[5] D. Clemente, V. Ramos, F. Teixeira-Duarte, F. V. C. Taveira-Pinto, P. Rosa-Santos, and F. Taveira-Pinto, *Assessment of electricity production and coastal protection of a nearshore 500 MW wave farm in the north-western Portuguese coast*, Appl. Energy, vol. 379, p. 124950, 2025, doi: 10.1016/J.APENERGY.2024.124950.

[6] N. Gonzalez *et al.*, *Offshore wind and wave energy can reduce total installed capacity required in zero-emissions grids*, Nat. Commun., vol. 15, no. 1, 2024, doi: 10.1038/s41467-024-50040-6.

[7] EU blue economy report, *The EU blue economy report 2025*, 2025.

[8] European Parliament, *EU strategy for offshore renewable energy sources*, 2024.

[9] A. G. Majidi, V. Ramos, T. Calheiros-Cabral, P. R. Santos, L. das Neves, and F. Taveira-Pinto, *Integrated assessment of offshore wind and wave power resources in mainland Portugal*, Energy, vol. 308, p. 132944, Nov. 2024, doi: 10.1016/J.ENERGY.2024.132944.

[10] A. F. de O. Falcão, *Wave energy utilization: A review of the technologies*, 2010. doi: 10.1016/j.rser.2009.11.003.

[11] Primre, *CorPower C4 WEC*. Available: https://openei.org/wiki/PRIMRE/Databases/Projects_Database/Devices/CorPower_Ocean_C4_WEC

[12] IRENA, *Global renewables outlook: energy transformation 2050*. International Renewable Energy Agency, 2020.

[13] R. A. Gideon and E. Bou-Zeid, *Collocating offshore wind and wave generators to reduce power output variability: A Multi-site analysis*, Renew. Energy, vol. 163, pp. 1548–1559, Jan. 2021, doi: 10.1016/J.RENENE.2020.09.047.

[14] A. S. H. Md Yasir, M. A. Omar, N. A. Mohd Azhar, A. Mukhtar, and M. Z. Yusoff, *A comprehensive review of recent advancements and challenges in breakwater oscillating water column technology for wave energy conversion*, Energy Reports, vol. 13, pp. 4095–4113, 2025, doi: 10.1016/J.EGYR.2025.03.045.

[15] N. Y. Sergiienko, L. Xue, L. S. P. da Silva, B. Ding, and B. S. Cazzolato, *Statistical analysis of floating hybrid wind–wave energy systems*, Appl. Energy, vol. 401, p. 126669, Dec. 2025, doi: 10.1016/J.APENERGY.2025.126669.

[16] K. Kardakaris, I. Boufidi, and T. Soukissian, *Offshore wind and wave energy complementarity in the greek seas based on ERA5 data*', Atmosphere (Basel)., vol. 12, no. 10, 2021, doi: 10.3390/atmos12101360.

[17] S. Astariz, A. Vazquez, M. Sánchez, R. Carballo, and G. Iglesias, *Co-located wave-wind farms for improved O&M efficiency*, Ocean Coast. Manag., vol. 163, pp. 66–71, 2018, doi: 10.1016/J.OCECOAMAN.2018.04.010.

[18] A. Ribeiro *et al.*, *Assessment of Hybrid Wind-Wave Energy Resource for the NW Coast of Iberian Peninsula in a Climate Change Context*, Applied Sciences, vol. 10, no. 21, 2020, doi: 10.3390/app10217395.

[19] A. Azzellino, C. Lanfredi, L. Riefolo, P. Contestabile, and D. Vicinanza, *Combined exploitation of offshore wind and wave energy in the Italian seas: A spatial planning approach*, Front. Energy Res., vol. 7, no. APR, 2019, doi: 10.3389/fenrg.2019.00042.

[20] N. Jamaleddin, V. A. D. de Faria, A. R. de Queiroz, and M. A. Gabr, *Optimizing co-located offshore wave and wind energy systems: Feasibility, design, and portfolios assessment*, J. Clean. Prod., vol. 523, p. 146461, 2025, doi: 10.1016/J.JCLEPRO.2025.146461.

[21] B. Skaare, F. G. Nielsen, T. D. Hanson, R. Yttervik, O. Havmøller, and A. Rekdal, *Analysis of measurements and simulations from the Hywind Demo floating wind turbine*, Wind Energy, vol. 18, no. 6, pp. 1105–1122, 2015, doi: https://doi.org/10.1002/we.1750.

[22] J. M. Jonkman, *Dynamics Modeling and Loads Analysis of an Offshore Floating Wind Turbine*, 2007. doi: https://doi.org/10.2172/921803.

[23] WindEurope, *Wind Energy in Europe: 2024 Statistics and the Outlook for 2025-2030*, 2025.

[24] D. Moulas, M. Shafiee, and A. Mehmanparast, *Damage analysis of ship collisions with offshore wind turbine foundations*, Ocean Engineering, vol. 143, pp. 149–162, 2017, doi: 10.1016/J.OCEANENG.2017.04.050.

[25] W. Musial, S. Butterfield, and B. Ram, *Energy from Offshore Wind*, 2006. doi: https://doi.org/10.4043/18355-MS.

[26] L. Castro-Santos, E. Martins, and C. Guedes Soares, *Cost assessment methodology for combined wind and wave floating offshore renewable energy systems*, Renew. Energy, vol. 97, pp. 866–880, 2016, doi: 10.1016/j.renene.2016.06.016.

[27] A. G. Majidi, V. Ramos, P. Rosa-Santos, A. Akpınar, L. das Neves, and F. Taveira-Pinto, *Development of a multi-criteria decision-making tool for combined offshore wind and wave energy site selection*, Appl. Energy, vol. 384, 2025, doi: 10.1016/j.apenergy.2025.125422.

[28] T. T. Tran and D. H. Kim, *A CFD study into the influence of unsteady aerodynamic interference on wind turbine surge motion*, Renew. Energy, vol. 90, pp. 204–228, 2016, doi: 10.1016/J.RENENE.2015.12.013.

[29] B. Skaare, F. G. Nielsen, T. D. Hanson, R. Yttervik, O. Havmøller, and A. Rekdal, *Analysis of measurements and simulations from the Hywind Demo floating wind turbine*, Wind Energy, vol. 18, no. 6, pp. 1105–1122, 2015, doi: https://doi.org/10.1002/we.1750.

[30] K. Ha, H. V. A. Truong, T. D. Dang, and K. K. Ahn, *Recent Control Technologies for Floating Offshore Wind Energy System: A Review*, International Journal of Precision Engineering and Manufacturing-Green Technology, vol. 8, no. 1, pp. 281–301, 2021, doi: 10.1007/s40684-020-00269-5.

[31] K.-T. Ma, Y. Luo, T. Kwan, and Y. Wu, *Mooring for floating wind turbines*, Mooring System Engineering for Offshore Structures, pp. 299–315, 2019, doi: 10.1016/B978-0-12-818551-3.00015-6.

[32] V. Ramos, G. Giannini, T. Calheiros-Cabral, M. López, P. Rosa-Santos, and F. Taveira-Pinto, *Assessing the Effectiveness of a Novel WEC Concept as a Co-Located Solution for Offshore Wind Farms*, J. Mar. Sci. Eng., vol. 10, no. 2, 2022, doi: 10.3390/jmse10020267.

[33] J. M. Jonkman, *Dynamics Modeling and Loads Analysis of an Offshore Floating Wind Turbine*, 2007. doi: https://doi.org/10.2172/921803.

[34] Y. Zhou *et al.*, *Initial Design of a Novel Barge-Type Floating Offshore Wind Turbine in Shallow Water*, J. Mar. Sci. Eng., vol. 11, no. 3, 2023, doi: 10.3390/jmse11030464.

[35] X. Su, X. Wang, W. Xu, L. Yuan, C. Xiong, and J. Chen, *Offshore Wind Power: Progress of the Edge Tool, Which Can Promote Sustainable Energy Development*, Sustainability, vol. 16, no. 17, 2024, doi: 10.3390/su16177810.

[36] Nyserda, *Analysis of Turbine Layouts and Spacing Between Wind Farms for Potential New York State Offshore Wind Development*, 2018.

[37] J. A. Pérez-Rúa, R. S. Lund, D. R. Verelst, A. B. Abrahamsen, and K. Dykes, *Exact optimization of inter-array dynamic cable networks for Floating Offshore Wind Farms*, Renew. Energy, vol. 237, 2024, doi: 10.1016/j.renene.2024.121647.

[38] offshoreWIND.biz, *World's First 20MW Wind Turbine Installed Offshore*, https://www.offshorewind.biz/2026/01/15/worlds-first-20-mw-wind-turbine-installed-offshore/.

[39] C. E. Clark, G. Barter, K. Shaler, and B. DuPont, *Reliability-based layout optimization in offshore wind energy systems*, *Wind Energy*, vol. 25, no. 1, pp. 125–148, 2022, doi: https://doi.org/10.1002/we.2664.

[40] S. Gambuzza and B. Ganapathisubramani, *The influence of free stream turbulence on the development of a wind turbine wake*, J. Fluid Mech., 2023, doi: 10.1017/jfm.2023.302.

[41] C. E. Clark, A. Miller, and B. DuPont, *An analytical cost model for co-located floating wind-wave energy arrays*, Renew. Energy, vol. 132, pp. 885–897, 2019, doi: 10.1016/J.RENENE.2018.08.043.

[42] J. N. Sørensen, R. F. Mikkelsen, D. S. Henningson, S. Ivanell, S. Sarmast, and S. J. Andersen, *Simulation of wind turbine wakes using the actuator line technique*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, vol. 373, no. 2035, p. 20140071, 2015, doi: 10.1098/rsta.2014.0071.

[43] H. Baki, S. Basu, and G. Lavidas, *Estimating the offshore wind power potential of Portugal by utilizing gray-zone atmospheric modeling*, Journal of Renewable and Sustainable Energy, vol. 16, no. 6, 2024, doi: 10.1063/5.0222974.

[44] S. Tao, J. Yang, G. Zheng, R. He, and A. E. Feijóo-Lorenzo, *A two-stage optimization model for floating offshore wind farm layout and control*, Renew. Energy, vol. 253, 2025, doi: 10.1016/j.renene.2025.123614.

[45] P. A. Fleming, A. Ning, P. M. O. Gebraad, and K. Dykes, *Wind plant system engineering through optimization of layout and yaw control*, Wind Energy, vol. 19, no. 2, pp. 329–344, 2016, doi: https://doi.org/10.1002/we.1836.

[46] IEA Wind TCP Task 49, *The IEA Wind Task 49 Reference Floating Wind Array Design Basis*, Jun. 2024. Available: www.nrel.gov/publications.

[47] P. Lakshmanan, R. Sun, and J. Liang, *Electrical collection systems for offshore wind farms: A review*, 2021, Institute of Electrical and Electronics Engineers Inc. doi: 10.17775/CSEEJPES.2020.05050.

[48] WINDExchange, *Offshore Wind Guide-wind-energy-guide Introduction to the Guide*, 2024. Available: https://windexchange.energy.gov/offshore

[49] Rijkswaterstaat, *Rules and safety tips: Code of conduct for safe passage through offshore wind farms*, 2018.

[50] B. Cerfontaine, S. M. Gourvenec, and D. J. White, *Specificities of floating offshore wind turbines for risk and safety evaluation of anchoring systems*, 2024, pp. 154–166. doi: 10.1201/9781003431749-10.

[51] P. Xu, X. Chang, and Y. E. Liu, *Novel Exact Solution of Elastic Catenary and Applications on Floating Wind Turbine Mooring Systems*, Journal of Offshore Mechanics and Arctic Engineering, vol. 146, no. 4, 2024, doi: 10.1115/1.4063852.

[52] M. Hall, M. Biglu, S. Housner, K. Coughlan, M. Y. Mahfouz, and E. Lozon, *Floating Wind Farm Layout Optimization Considering Moorings and Seabed Variations*, in Journal of Physics: Conference Series, Institute of Physics, 2024. doi: 10.1088/1742-6596/2767/6/062038.

[53] F. Teixeira-Duarte, P. Rosa-Santos, and F. Taveira-Pinto, *Multi-objective optimization of co-located wave-wind farm layouts supported by genetic algorithms and numerical models*, Renew. Energy, vol. 241, p. 122362, 2025, doi: 10.1016/J.RENENE.2025.122362.

[54] R. J. Barthelmie *et al.*, *Quantifying the Impact of Wind Turbine Wakes on Power Output at Offshore Wind Farms*, J. Atmos. Ocean. Technol., vol. 27, no. 8, pp. 1302–1317, 2010, doi: 10.1175/2010JTECHA1398.1.

[55] Z. Yong and Q. Shengli, *A Practical Optimization for Offshore Wind Farm Layout*, 2016.

[56] Wikimedia Commons, *Alpha Ventus Offshore Wind Farm*, https://upload.wikimedia.org/wikipedia/commons/c/ca/Windpark_alpha_ventus_Lagekarte.png

[57] M. Sickler, B. Ummels, M. Zaaijer, R. Schmehl, and K. Dykes, *Offshore wind farm optimisation: a comparison of performance between regular and irregular wind turbine layouts*, Wind Energy Science, vol. 8, no. 7, pp. 1225–1233, 2023, doi: 10.5194/wes-8-1225-2023.

[58] Orsted, *Hornsea Project One Project Overview Presentation for Ofgem OFTO TR6 Launch Event*, 2018.

[59] Forewind, *Environmental Statement Chapter 9 Appendix A Assessment of Effects Technical Report*, 2014.

[60] D. Silva, P. Martinho, and C. Guedes Soares, *Wave energy distribution along the Portuguese continental coast based on a thirty-three years hindcast*, Renew. Energy, vol. 127, pp. 1064–1075, 2018, doi: 10.1016/j.renene.2018.05.037.

[61] Syarif Arief Irfan, Aria Pria Utama I Ketut, Hantoro Ridho, Prananda Juniarko, and Muhammad Megawan Alfa, *Computational Fluid Dynamics (CFD) Simulation for Designing Mooring Bitts Position at the Barge for Wave Energy Conversion (WEC)*, E3S Web Conf., vol. 190, p. 17, 2020, doi: 10.1051/e3sconf/202019000017.

[62] F. Teixeira-Duarte, D. Clemente, G. Giannini, P. Rosa-Santos, and F. Taveira-Pinto, *Review on layout optimization strategies of offshore parks for wave energy converters*, 2022, Elsevier Ltd. doi: 10.1016/j.rser.2022.112513.

[63] M. T. Rahmati and G. A. Aggidis, *Numerical and experimental analysis of the power output of a point absorber wave energy converter in irregular waves*, Ocean Engineering, vol. 111, pp. 483–492, 2016, doi: 10.1016/j.oceaneng.2015.11.011.

[64] R. Berriel, M. Shadman, Z. Wu, R. Dias, R. Stephan, and S. Estefen, *Hardware-in-the-loop development of a heaving point absorber wave energy converter using inertia emulation*, Electrical Engineering, vol. 103, 2021, doi: 10.1007/s00202-021-01244-2.

[65] M. Blanco, M. Lafoz, D. Ramirez, G. Navarro, J. Torres, and L. Garcia-Tabares, *Dimensioning of point absorbers for wave energy conversion by means of differential evolutionary algorithms*, IEEE Trans. Sustain. Energy, vol. 10, no. 3, pp. 1076–1085, 2019, doi: 10.1109/TSTE.2018.2860462.

[66] M. Folley *et al.*, *A review of numerical modelling of wave energy converter arrays*, 2012.

[67] B. F. M. Child and V. Venugopal, *Optimal configurations of wave energy device arrays*, Ocean Engineering, vol. 37, no. 16, pp. 1402–1417, 2010, doi: 10.1016/j.oceaneng.2010.06.010.

[68] I. de Santiago, P. Liria, R. Garnier, J. Bald, J. Leitão, and J. Ribeiro, *Deliverable 3.3 Marine Dynamic Modelling*, 2023. doi: 10.13140/RG.2.2.36716.74886.

[69] R. Satymov, D. Bogdanov, M. Dadashi, G. Lavidas, and C. Breyer, *Techno-economic assessment of global and regional wave energy resource potentials and profiles in hourly resolution*, Appl. Energy, vol. 364, p. 123119, Jun. 2024, doi: 10.1016/J.APENERGY.2024.123119.

[70] M. J. Dickson, S. Pennock, H. Todalshaug, and K. Rebenius, *Results and Lessons Learned from CorPower Ocean's Full Scale C4 Wave Energy Converter First Ocean Deployment Pe*riod, 2025.

[71] CorPower, '*CorPack*'. Accessed: Dec. 14, 2025. Available: https://corpowerocean.com/corpack/

[72] C. Pérez-Collazo, D. Greaves, and G. Iglesias, *A review of combined wave and offshore wind energy*, 2015, Elsevier Ltd. doi: 10.1016/j.rser.2014.09.032.

[73] S. Astariz and G. Iglesias, *Selecting optimum locations for co-located wave and wind energy farms. Part I: The Co-Location Feasibility index*, Energy Convers. Manag., vol. 122, pp. 589–598, 2016, doi: 10.1016/j.enconman.2016.05.079.

[74] F. Teixeira-Duarte, V. Ramos, P. Rosa-Santos, and F. Taveira-Pinto, *Multi-objective decision tool for the assessment of co-located wave-wind offshore,* 2024.

[75] L. Castro-Santos, E. Martins, and C. Guedes Soares, *Cost assessment methodology for combined wind and wave floating offshore renewable energy systems*, Renew. Energy, vol. 97, pp. 866–880, 2016, doi: 10.1016/J.RENENE.2016.06.016.

[76] L. Rusu, *An evaluation of the synergy between the wave and wind energy along the west Iberian nearshore*, Energy Conversion and Management: *X*, vol. 20, 2023, doi: 10.1016/j.ecmx.2023.100453.

[77] G. Chang, J. Magalen, C. Jones, and J. Roberts, *Wave Energy Converter Effects on Wave Fields: Evaluation of SNL-SWAN and Sensitivity Studies in Monterey Bay CA*, 2014.

[78] M. M. Moarefdoost, L. V. Snyder, and B. Alnajjab, *Layouts for ocean wave energy farms: Models, properties, and optimization*, Omega (United Kingdom), vol. 66, pp. 185–194, 2017, doi: 10.1016/j.omega.2016.06.004.

[79] S. Astariz, J. Abanades, C. Perez-Collazo, and G. Iglesias, *Improving wind farm accessibility for operation & maintenance through a co-located wave farm: Influence of layout and wave climate*, Energy Convers. Manag., vol. 95, pp. 229–241, 2015, doi: 10.1016/j.enconman.2015.02.040.

[80] J. Abanades, D. Greaves, and G. Iglesias, *Coastal defence through wave farms*, Coastal Engineering, vol. 91, pp. 299–307, 2014, doi: 10.1016/j.coastaleng.2014.06.009.

[81] M. Giassi and M. Göteman, *Layout design of wave energy parks by a genetic algorithm*, Ocean Engineering, vol. 154, pp. 252–261, 2018, doi: 10.1016/j.oceaneng.2018.01.096.

[82] A. E. Eiben, *Introduction to Evolutionary Computing II with thanks to the EvoNet Training Committee and its "Flying Circus"*, 2015. Available: http://www.cs.vu.nl/~gusz/

[83] M. Neshat, B. Alexander, and M. Wagner, *A hybrid cooperative co-evolution algorithm framework for optimising power take off and placements of wave energy converters*, Inf. Sci. (N Y)., vol. 534, pp. 218–244, 2020, doi: 10.1016/j.ins.2020.03.112.

[84] K. Ruehl, A. Porter, C. Chartrand, H. Smith, G. Chang, and J. Roberts, *Development of SNL-SWAN, a Validated Wave Energy Converter Array Modeling Tool*, 2013.

[85] K. Ruehl, A. Porter, C. Chartrand, H. Smith, G. Chang, and J. Roberts, *Development, Verification and Application of the SNL-SWAN Open Source Wave Farm Code*, 2015.

[86] N. Booij, R. Ris, and L. Holthuijsen, *A third-generation wave model for coastal regions, Part I, Model description and validation*, J. Geophys. Res., vol. 104, pp. 7649–7656, Mar. 1999.

[87] Delft University of Technology, *SWAN scientific and technical documentation*.

[88] J. Falnes, *Ocean Waves and Oscillating Systems*, pp. 286. ISBN 0521782112. Cambridge, UK: Cambridge University Press, April 2002., vol. 1, 2002, doi: 10.1017/CBO9780511754630.

[89] A. Babarit, *A database of capture width ratio of wave energy converters*, Renew. Energy, vol. 80, pp. 610–628, 2015, doi: https://doi.org/10.1016/j.renene.2015.02.049.

[90] A. Babarit, *Impact of long separating distances on the energy production of two interacting wave energy converters*, Ocean Engineering, vol. 37, no. 8, pp. 718–729, 2010, doi: https://doi.org/10.1016/j.oceaneng.2010.02.002.

[91] M. Folley, *Numerical modelling of wave energy converters*, 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:114314758

[92] P. Camus, F. Mendez, and R. Medina, *A hybrid efficient method to downscale wave climate to coastal areas*, Coastal Engineering, vol. 58, pp. 851–862, 2011, doi: 10.1016/j.coastaleng.2011.05.007.

[93] A. K. Jain, *Data clustering: 50 years beyond K-means*, Pattern Recognit. Lett., vol. 31, no. 8, pp. 651–666, 2010, doi: https://doi.org/10.1016/j.patrec.2009.09.011.

[94] J. Macqueen, *Some Methods for Classification and Analysis of Multivariate Observations*.

[95] J. R. Joubert, *Design and development of a novel wave energy converter*, 2013.

[96] M. Mitchell, *An introduction to genetic algorithms*, 1996. Available: https://api.semanticscholar.org/CorpusID:25869059

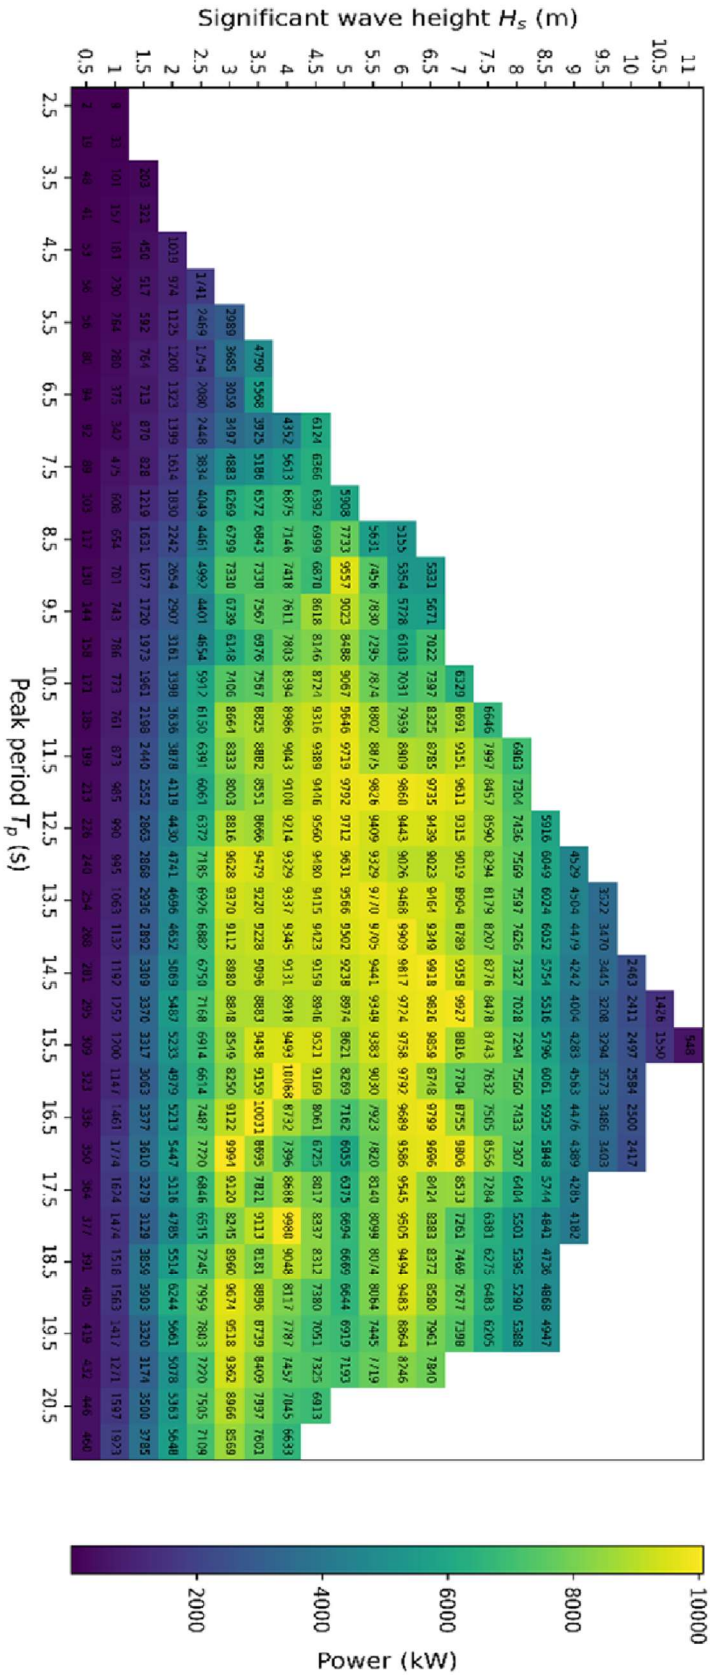
[97] D. E. Goldberg, *Genetic algorithms in search, optimization & machine learning*.

[98] J. H. Holland, *Adaptation in Natural and Artificial Systems*.

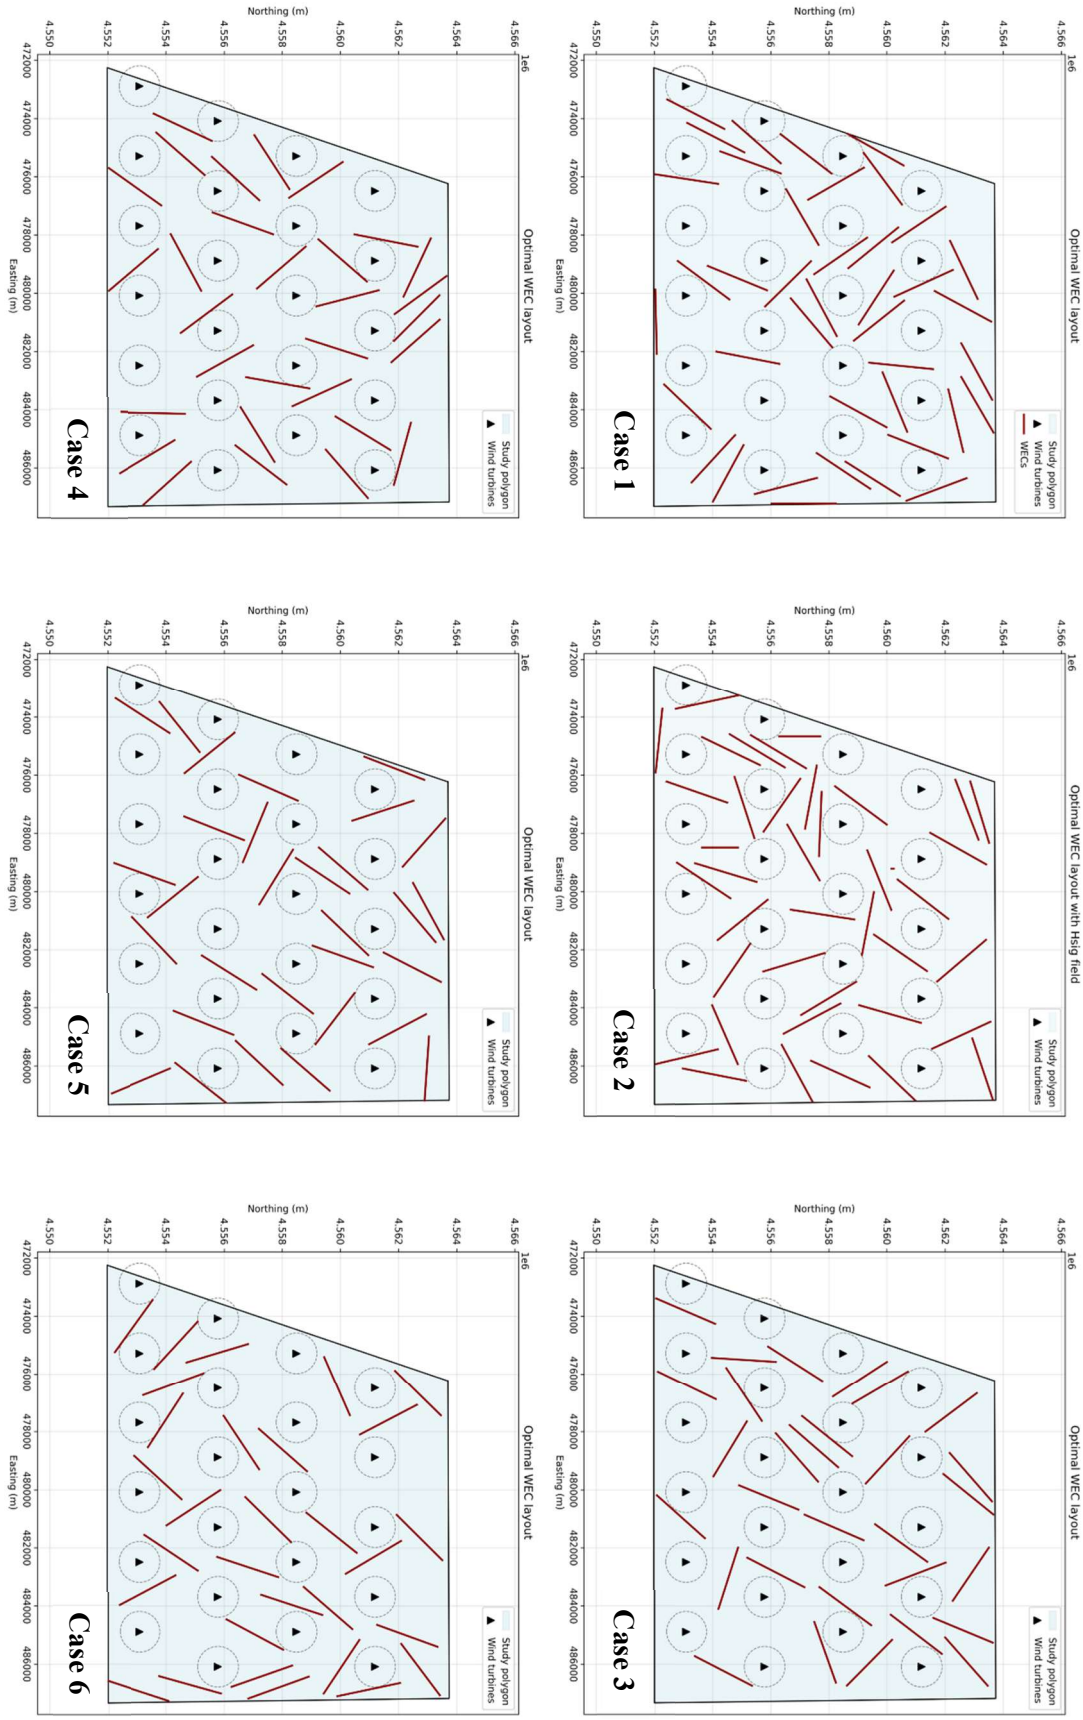
[99] The Resident, *Portugal is 95% water*, https://www.portugalresident.com/portugal-is-95-water/.

[100] Diário da República, *Presidência do conselho de ministros - Resolução do Conselho de Ministros n.º 19/2025*, 2025.

APPENDIX A - CORPOWER POWER MATRIX



APPENDIX B - FINAL LAYOUT FOR EACH OPTIMISATION CASE



APPENDIX C - H_S FIELDS FOR EACH OPTIMISATION CASE

