

Evaluation of low-frequency wind variations on wind turbine fatigue damage estimations

MSc Thesis

Marcos García Bravo



EWEM



Evaluation of low-frequency wind variations on wind turbine fatigue damage estimations

Marcos García Bravo

to obtain the degrees of:

Master of Science
in Aerospace Engineering
at Delft University of Technology

Master of Science
in Wind Energy
at Technical University of Denmark

Supervisors:

Oscar Ramírez Requesón	Siemens Gamesa
Taeseong Kim	DTU
Wei Yu	TU Delft
Nirav Dangi	TU Delft



Acknowledgments

I would like to thank Siemens Gamesa for giving me the opportunity to carry out this thesis in such an inspiring professional environment. I am especially grateful to Oscar Ramírez for his guidance, support and trust throughout the project, and to everyone at the company who contributed their time, knowledge or advice at any stage.

Many thanks also go to my supervisors, Wei Yu and Nirav Dangi at TU Delft and Taeseong Kim at DTU, for their guidance, availability and the many constructive discussions that shaped this work. I would also like to thank the EWEM programme for the opportunity to study in two countries, experience different academic environments and grow through such a rewarding journey.

Finally, I am grateful to my family and friends for their unconditional support and encouragement throughout this process.

*Marcos García Bravo
Copenhagen, July 2026*

Abstract

Wind turbine fatigue assessments commonly represent load histories as independent 10-minute records, although slow atmospheric and operational changes can generate cycles that cross these boundaries. The resulting fatigue error and its dependence on structural response and material sensitivity remain insufficiently established.

This thesis determined how independent 10-minute processing affects fatigue estimates compared with longer continuous histories and whether the missing cycles can be recovered efficiently. Six months of load and operational measurements from a 7 MW offshore turbine were analysed for blade root flapwise and edgewise bending, and for two tower-bottom bending directions. Continuous rainflow counting, conventional independent-block processing and residue-aware low-frequency fatigue dynamics (LFFD) processing were compared across observation periods and Wöhler exponents. BHawC aeroelastic simulations compared 600 and 3600 s histories, continuous and segmented operating-state sequences, and concatenated short simulations.

LFFD reproduced the continuous reference while enabling analysis of the full six-month record. Independent 10-minute processing underestimated the flapwise damage equivalent load (DEL) by approximately 16% to 20%, and the tower bottom DELs by 3% to 20%, depending on the Wöhler exponent, while the edgewise difference remained below 0.2%. Daily residue aggregation recovered approximately 95%–96% of the full-record flapwise DEL and weekly aggregation about 99%. Cycles connecting different operating regimes contributed approximately 61% of the residue-aware flapwise fatigue contribution.

Extending stationary simulations from 600 to 3600 s increased the combined DEL by 3.62% for flapwise and 1.76% for fore-aft tower bending, whereas continuously processing an idling–startup–production–shutdown sequence increased them by 19.19% and 9.74%, respectively. Concatenated short simulations did not consistently reproduce directly simulated one-hour histories.

The results show that 10-minute storage remains practical, but processing those blocks independently can produce non-conservative estimates for aerodynamically driven responses. Residue-aware processing and targeted continuous transition simulations therefore provide more reliable fatigue estimates.

Contents

Acknowledgements	i
Abstract	ii
1 Introduction	1
1.1 Motivation	1
1.2 Problem statement and research objectives	4
1.2.1 Research questions	5
1.3 Scope and contribution	5
1.4 Thesis structure	6
2 Literature review	7
2.1 Introduction	7
2.2 Fatigue assessment workflow in wind energy	7
2.3 Fatigue evaluation using aeroelastic simulations	14
2.4 Rainflow residue processing and low-frequency fatigue	18
2.5 Operational transitions and fatigue damage	20
2.6 Reliability of long-term fatigue estimates	21
2.7 Synthesis and identified research gap	21
3 Measurement-based fatigue methodology	23
3.1 Measurement dataset	23
3.2 Fatigue-processing strategies	25
3.3 Implementation of the residue-aware workflow	26
3.4 Rainflow implementation and validation	28
3.5 Fatigue measures and comparison metrics	29
3.6 Analysis across time scales, Wöhler exponents and cycle ranges	30
3.7 Operational-state attribution	31
4 BHawC aeroelastic model and simulation methodology	32
4.1 BHawC aeroelastic model	32
4.2 Simulation inputs and environmental conditions	34
4.3 Simulation campaign	40
4.4 Concatenation of 600 s simulations	43
4.5 Methodological limitations	45
5 Measurement-based fatigue results	47
5.1 Measured load behaviour across different time scales	47
5.2 Validation of the rainflow and residue-recovery workflow	49
5.3 Blade root flapwise bending moment	51
5.3.1 Development of the segmentation effect	51
5.3.2 Sensitivity to the Wöhler exponent	52
5.3.3 Contribution of different aggregation scales	53
5.3.4 Cycle-range distribution	55
5.4 Comparison between structural response channels	57

5.4.1	Blade root edgewise bending moment	57
5.4.2	Tower-bottom bending moments	59
5.5	Operational origin of the recovered cycles	63
5.5.1	Fatigue attribution by operating regime	63
5.5.2	Damage matrices and transition cycles	64
6	Simulation-based results	67
6.1	Baseline results	67
6.1.1	Normal-production comparison	67
6.1.2	Combined baseline DLCs	73
6.2	Continuous sequence of design load cases	77
6.3	Concatenated simulation results	82
7	Conclusions and recommendations	84
7.1	Conclusions	84
7.2	Practical implications	87
7.3	Limitations and future work	88
7.4	Final statement	89
	References	90
	Use of Artificial Intelligence	92
A	Measurement data processing	93
A.1	Measurement data filtering	93
A.2	Visualising the LFFD process	95
A.3	Distribution of operating regimes from measurement data	98
B	Additional load-channel simulation results	100
B.1	Baseline simulations	100
B.2	DLC combination simulations	111
B.3	Concatenation simulations	117

List of Figures

1.1	Example of power production design load cases defined within the IEC 61400 framework [4].	2
1.2	Spectral density of wind speed according to the classical Van der Hoven spectrum [6, 8].	2
1.3	Wind speed and structural load variations over a one hour period divided into consecutive 10-minute intervals.	3
2.1	Fatigue damage assessment workflow.	8
2.2	Typical power-law Wöhler curve used in fatigue assessment.	9
2.3	Schematic relationship between a variable-amplitude load history and the corresponding closed material hysteresis loops [20].	10
2.4	Illustration of turning point extraction, candidate ranges and the four-point rainflow closure rule.	12
2.5	Damage equivalent load concept, where a constant-amplitude load produces the same damage as the original variable-amplitude history.	13
2.6	Example discretisation of a Weibull distribution into wind speed bins.	15
2.7	Residue based recovery of fatigue cycles extending across the original processing windows, based on Sadeghi et al. [12].	19
2.8	Damage matrix showing cycles classified by the turbine condition at their start and end.	20
3.1	Continuous-reference, conventional blockwise and residue-aware fatigue-processing strategies.	26
3.2	Computational workflow used to process closed cycles and propagate the ordered rainflow residue between consecutive 10-minute blocks.	28
4.1	Main categories of inputs and outputs used in the BHawC simulation workflow.	34
4.2	Schematic representation of the Mann turbulence box and its coordinate system.	35
4.3	Representative Mann turbulence-box diagnostic showing the centre-point wind-speed history, a vertical slice through the box, two spatial cross-sections and the corresponding centre-point power spectrum.	38
4.4	Extension of the original wave height time series. The initial wave record is shown in orange and the mirrored extension in blue.	39
4.5	Detailed view of the transition between the original and mirrored height wave records.	39
4.6	Comparison of the wave height spectra before and after extending the time series. The initial wave record is shown in orange and the mirrored extension in blue.	40
4.7	Concatenation of individual 600s simulations concept.	43
5.1	Thirty-day record of the blade root flapwise bending moment.	47
5.2	Thirty-day flapwise bending-moment record with the slow variation highlighted using a low-pass-filtered curve.	48

5.3	First 10-minute block of the flapwise record, showing the shorter cycles associated with blade rotation.	48
5.4	Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the blade root flapwise channel at $m = 10$. Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.	51
5.5	Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the blade root flapwise channel at $m = 10, 11$ and 12 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.	53
5.6	Schematic of splitting of record depending on aggregation scales.	54
5.7	Fraction of the six-month flapwise DEL recovered using different residue-aggregation periods for $m = 10, 11$ and 12	55
5.8	Cycle-count distribution for the blade root flapwise channel before and after residue recovery.	56
5.9	Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the blade root edgewise channel at $m = 10, 11$ and 12 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.	57
5.10	Cycle-count distribution for the blade root edgewise channel before and after residue recovery.	58
5.11	Tower coordinate convention and site wind climate.	59
5.12	Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the tower-bottom X direction at $m = 3, 4$ and 5 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.	60
5.13	Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the tower-bottom Y direction at $m = 3, 4$ and 5 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.	60
5.14	Cycle-count distribution for the tower-bottom X direction before and after residue recovery.	61
5.15	Cycle-count distribution for the tower-bottom Y direction before and after residue recovery.	62
5.16	Fatigue share associated with each operating regime for the blade root flapwise channel at $m = 10$, before and after residue recovery.	64
5.17	Folded operating-regime damage matrices for the blade root flapwise channel at $m = 10$. Left: conventional blockwise processing. Right: residue-aware processing.	65
6.1	Blade root flapwise bending-moment histories from paired 600 s and 3600 s DLC 1.2 simulations at 12 m/s using the same turbulence seed.	68
6.2	Mean blade root flapwise DEL for the 600 s and 3600 s DLC 1.2 simulations. The shaded regions indicate the minimum-to-maximum range across turbulence seeds.	68
6.3	Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the blade root flapwise moment.	69
6.4	Weibull-weighted equivalent-load contribution for the blade root flapwise moment.	70
6.5	Ratio of the 3600 s and 600 s blade root flapwise DELs for each DLC 1.2 wind-speed bin. The shaded envelope spans the minimum and maximum seed-wise ratios.	71
6.6	Seed-matched comparison of the 600 s and 3600 s blade root flapwise DELs. Colour indicates the nominal wind speed, while the dashed line represents equal DEL for both durations.	72

6.7	Weighted equivalent-load contribution of each baseline DLC for the blade root flapwise moment. Error bars indicate turbulent seed variability.	74
6.8	Blade root flapwise bending-moment histories from paired 600 s and 3600 s DLC 7.2 simulations at 12 m/s using the same turbulence seed.	74
6.9	Contribution of each baseline DLC to the total annual blade root flapwise fatigue damage. Error bars indicate turbulent seed variability.	75
6.10	Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for the blade root flapwise moment.	77
6.11	Mean blade root flapwise DEL from continuous and DLC-segmented processing. The shaded regions indicate the minimum-to-maximum range across turbulence seeds.	78
6.12	Ratio of blade root flapwise DEL from continuous and DLC-segmented processing. The shaded regions indicate the minimum-to-maximum range across turbulence seeds.	79
6.13	Mean contribution of each operational state to the total blade root flapwise fatigue damage when the simulation is divided into separate DLC segments.	80
6.14	Ratio of blade root flapwise DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases. The shaded band indicates seed variance of the one-hour long simulations.	82
A.1	Simplified workflow used for measurement data pre-processing.	93
A.2	One-hour blade root flapwise bending moment divided into six consecutive 10-minute blocks.	95
A.3	Rainflow residues from the individual 10-minute blocks overlaid on the original load signal.	96
A.4	Residues from the six 10-minute blocks joined in chronological order.	96
A.5	Cycle ranges obtained by rainflow counting the combined residue.	97
A.6	Measured blade pitch, rotor speed, grid power and blade root flapwise bending moment as functions of 10-minute mean wind speed, together with the classified operating regimes.	98
A.7	Time share of the simplified operating regimes over the measurement period.	99
B.1	DEL as a function of wind speed for DLC 1.2 for the root blade edge moment.	100
B.2	DEL as a function of wind speed for DLC 1.2 for the tower-bottom fore-aft moment.	101
B.3	DEL as a function of wind speed for DLC 1.2 for the tower-bottom side-side moment.	101
B.4	Wind-speed-dependent damage-equivalent loads and Weibull weighting for the root blade edge moment.	102
B.5	Wind-speed-dependent damage-equivalent loads and Weibull weighting for the tower-bottom fore-aft moment.	102
B.6	Wind-speed-dependent damage-equivalent loads and Weibull weighting for the tower-bottom side-side moment.	103
B.7	Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the root blade edge moment.	103
B.8	Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the tower-bottom fore-aft moment.	104
B.9	Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the tower-bottom side-side moment.	104

B.10 Ratio of the 3600 s and 600 s blade root edge weighted DELs for each DLC 1.2 wind-speed bin.	105
B.11 Ratio of the 3600 s and 600 s tower-bottom fore-aft weighted DELs for each DLC 1.2 wind-speed bin.	105
B.12 Ratio of the 3600 s and 600 s tower-bottom side-side weighted DELs for each DLC 1.2 wind-speed bin.	106
B.13 Comparison of 600s vs. 3600s DEL per simulation seed for root blade edge moment.	106
B.14 Comparison of 600s vs. 3600s DEL per simulation seed for tower-bottom fore-aft moment.	107
B.15 Comparison of 600s vs. 3600s DEL per simulation seed for tower-bottom side-side moment.	107
B.16 Equivalent-load contribution of each design load case for the root blade edge moment.	108
B.17 Equivalent-load contribution of each design load case for the tower-bottom fore-aft moment.	108
B.18 Equivalent-load contribution of each design load case for the tower-bottom side-side moment.	109
B.19 Contribution of each design load case to the total fatigue damage of the root blade edge moment.	109
B.20 Contribution of each design load case to the total fatigue damage of the tower-bottom fore-aft moment.	110
B.21 Contribution of each design load case to the total fatigue damage of the tower-bottom side-side moment.	110
B.22 Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for blade-root edgewise moment.	111
B.23 Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for tower-bottom fore-aft moment.	111
B.24 Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for tower-bottom side-side moment.	112
B.25 Mean blade-root edgewise DEL from continuous and DLC-segmented processing.	112
B.26 Mean tower-bottom fore-aft DEL from continuous and DLC-segmented processing.	113
B.27 Mean tower-bottom side-side DEL from continuous and DLC-segmented processing.	113
B.28 Ratio of blade-root edgewise DEL from continuous and DLC-segmented processing.	114
B.29 Ratio of tower-bottom fore-aft DEL from continuous and DLC-segmented processing.	114
B.30 Ratio of tower-bottom side-side DEL from continuous and DLC-segmented processing.	115
B.31 Mean contribution of each operational state to the total blade-root edgewise fatigue damage when the simulation is divided into separate DLC segments.	115
B.32 Mean contribution of each operational state to the total tower-bottom fore-aft fatigue damage when the simulation is divided into separate DLC segments.	116

B.33 Mean contribution of each operational state to the total tower-bottom side-side fatigue damage when the simulation is divided into separate DLC segments. . 116

B.34 Ratio of blade root edgewise DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases. 117

B.35 Ratio of tower bottom fore-aft DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases. 117

B.36 Ratio of tower bottom side-side DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases. 118

List of Tables

2.1	Typical workflow for aeroelastic fatigue assessment.	14
2.2	Main differences between the Kaimal and Mann turbulence models, based on [4, 22, 23].	16
4.1	Main characteristics of the BHawC turbine model.	33
4.2	BHawC load responses and fatigue parameters used in the comparisons. . . .	34
4.3	Mann turbulence-box parameters used in the BHawC simulation campaign. . .	37
4.4	Baseline BHawC simulation campaign.	40
4.5	Operating-state sequence used in the combined one-hour simulations.	42
4.6	Comparison of DLC weights used in the conventional fatigue limit state calculation, the operating-state shares from the measurements, and the time shares prescribed in the combined one-hour simulations.	42
4.7	Simulation cases used in the concatenation study.	44
5.1	Validation ratios for the blade root flapwise channel at $m = 10$	49
5.2	Indicative computation times for the blade root flapwise channel at $m = 10$. . .	50
5.3	Approximate six-month DEL ratios for the measured load channels.	63
6.1	Percentage change in DEL from the 600 s to the 3600 s simulations for each load channel and DLC.	76
6.2	Percentage increase in the Weibull-weighted DEL obtained from continuous processing relative to DLC-segmented processing.	80
A.1	Measurement channels used in the fatigue analysis.	94

1

Introduction

1.1. Motivation

Fatigue damage is one of the main factors limiting the operational lifetime of modern wind turbines. Repeated loading affects the rotor blades, tower, drivetrain and other mechanical components throughout the turbine lifetime. For offshore wind turbines, the support structure is also subjected to the combined effects of wind and wave loading, resulting in complex and continuously varying structural responses [1, 2].

Fatigue damage describes the progressive weakening of a material or structure under repeated loading, even when the individual load levels remain below the ultimate strength of the material. Although each individual cycle may cause only a very small amount of damage, the accumulation of a large number of cycles can eventually lead to crack initiation, crack growth and structural failure [1].

In engineering practice, fatigue assessment generally relies on three main elements:

1. A representative description of the environmental and operational loading
2. A method for reducing a load time series into a set of fatigue cycles
3. A material model and cumulative damage rule used to relate those cycles to fatigue damage

This general workflow is used both during turbine design and certification, as well as during the assessment of turbines already in operation. In the latter case, measured load or strain signals can be used to estimate accumulated fatigue damage and support decisions related to inspection, maintenance, lifetime extension and continued operation [3].

In current design and certification practice, wind turbines are assessed using a series of design load cases (DLCs) defined in the International Electrotechnical Commission (IEC) 61400 standards. IEC 61400-1 describes the main design requirements, while IEC 61400-13 addresses the measurement of mechanical loads [4, 5]. The DLCs represent different combinations of wind conditions, turbine operating states and fault situations. They include, among others, normal power production, startup, shutdown, idling, parked conditions and fault events.

An example of the design load cases associated with power production is shown in Figure 1.1. The cases cover several inflow conditions that may occur while the turbine is operating between cut-in and cut-out wind speed. These include normal turbulence, extreme turbulence, extreme coherent gusts combined with wind-direction changes, and extreme wind shear. Each

condition is evaluated using either fatigue or ultimate-load analysis, as appropriate. Together, these cases form part of a wider certification framework in which the turbine must be shown to withstand both repeated operational loading and less frequent extreme events throughout its intended design lifetime.

Design situation	DL C	Wind condition	Other conditions	Type of analysis	Partial safety factors
1) Power production	1.1	NTM $V_{in} < V_{hub} < V_{out}$	For extrapolation of extreme events	U	N
	1.2	NTM $V_{in} < V_{hub} < V_{out}$		F	*
	1.3	ETM $V_{in} < V_{hub} < V_{out}$		U	N
	1.4	ECD $V_{hub} = V_r - 2 \text{ m/s}, V_r, V_r + 2 \text{ m/s}$		U	N
	1.5	EWS $V_{in} < V_{hub} < V_{out}$		U	N

Figure 1.1: Example of power production design load cases defined within the IEC 61400 framework [4].

However, a practical limitation of the conventional fatigue assessment procedure is its reliance on 10-minute time windows. Wind statistics such as mean wind speed and turbulence intensity are generally calculated over 10-minute periods, while aeroelastic fatigue simulations are also commonly carried out over the same duration [4, 5].

The use of 10-minute intervals represents a practical compromise. The interval is sufficiently long to estimate basic wind statistics under approximately stationary conditions, while remaining short enough to limit computational cost and data storage. It is also broadly consistent with the low-energy spectral gap illustrated by the classical Van der Hoven spectrum in Figure 1.2, which separates short-term atmospheric turbulence from slower mesoscale variations. In the original spectrum, this gap corresponds approximately to periods between 6 minutes and 1 hour, or frequencies from 2.8×10^{-4} to 2.8×10^{-3} Hz [6]. A 10-minute interval therefore falls within this range, making it a reasonable choice of time length for load analyses. It is worth noting that the exact range and prominence of the gap are not universal, however, and can vary with atmospheric conditions and measurement location [7].




Figure 1.2: Spectral density of wind speed according to the classical Van der Hoven spectrum [6, 8].

Nevertheless, the use of a 10-minute interval does not mean that all fatigue-relevant loading occurs within this period. Wind speed, turbulence intensity, wind direction and turbine operating state can vary over periods ranging from several minutes to several days. These variations

can lead to changes in the mean structural load and may generate load cycles with periods longer than a single 10-minute window.

The relevance of these slower variations is expected to increase as wind turbines become larger. Modern rotors interact with an increasing portion of the atmospheric boundary layer, while the natural frequencies of blades, towers and support structures generally decrease as the structures become larger and more flexible [2, 9]. Large eddy simulation (LES) studies indicate that coherent atmospheric structures can influence blade and tower loading for large turbines, with some of the biggest differences between conventional spectral turbulence models and higher-fidelity inflow models occurring at low frequencies [9].

As shown in Figure 1.2, the wind spectrum contains contributions from turbulence occurring over periods of seconds and minutes, together with slower atmospheric variations extending over hours, days and longer periods [6]. Conventional 10-minute simulations are intended primarily to represent the turbulent part of the spectrum. Longer-term variability is not generally simulated as one continuous time history. Instead, it is represented indirectly through separate simulations at different mean wind speeds, whose fatigue contributions are combined using long-term occurrence probabilities [4, 10].

This separation is convenient for design calculations, but it may remove cycles that arise from changes between the independently simulated conditions. If the mean load level changes between two consecutive periods, a load cycle can extend across the boundary separating those periods. Such a cycle cannot be identified if each interval is counted independently [11, 12].

The practical effect is illustrated in Figure 1.3. The figure shows an example wind speed signal and its corresponding structural load response over one hour, divided into six consecutive 10-minute intervals. Although each individual interval contains high-frequency load fluctuations, the mean wind speed and mean load also vary across the complete hour.




Figure 1.3: Wind speed and structural load variations over a one hour period divided into consecutive 10-minute intervals.

When the full signal is considered continuously, the slow variation in the load level may form large fatigue cycles. When the six intervals are processed separately, these cycles can be divided into smaller ranges or left as unresolved residues at the segment boundaries. Since

fatigue damage increases nonlinearly with load range, a small number of missed large cycles can have a substantial effect on the estimated damage [11, 12].

The same issue is relevant for measured turbine data. Structural health monitoring systems commonly store or process signals in 10-minute blocks. This makes the data manageable and aligns it with the format of standard wind statistics. However, independently processing these blocks may prevent cycles extending across block boundaries from being represented correctly [2, 12].

1.2. Problem statement and research objectives

Conventional fatigue assessment assumes that long-term fatigue damage can be reconstructed by combining the results from independently processed short-duration records. This assumption is valid only when the relevant cycles are fully contained within each record, or when unresolved cycles at the record boundaries make a negligible contribution to the total damage. Neither condition has been systematically verified for wind turbine loading.

Existing literature shows that this assumption is not always valid. Slow changes in wind speed, wind direction and turbine operating state can create load cycles that last longer than the standard 10-minute interval [2, 11, 12]. What is still missing is a clear picture of the *size and consistency* of the resulting error. Prior work has shown that these low-frequency cycles exist and that they can be recovered with residue-aware counting methods. However, it has not shown, across a common set of structural response channels and material assumptions, how large the resulting damage difference usually is. Moreover, it has also not shown which conditions make the effect stronger or weaker, or whether an effect measured in operational data can be reproduced under the controlled conditions of an aeroelastic simulation. This is the gap this thesis addresses. The question is not whether low-frequency cycles exist. The question is whether they matter enough, and under which conditions, to justify moving away from the conventional independent-block procedure.

This distinction matters in practice. Structural responses that are strongly affected by changes in aerodynamic mean loading are expected to be more sensitive to this effect. Blade root flapwise bending and tower bottom fore-aft bending are good examples. Responses that are dominated by higher-frequency dynamics, such as edgewise or side-side bending, are expected to be less sensitive. Furthermore, fatigue damage accumulation is also a nonlinear function of cycle range, since it follows a power law. As a result, the assumed Wöhler exponent is expected to influence how strongly missed low-frequency cycles affect the final damage estimate. Neither of these two dependencies has been quantified consistently on a shared dataset so far.

Measurement data can be used to quantify these effects under real operating conditions. However, measurements do not give full control over the wind field, the turbine state or the occurrence of specific events. Aeroelastic simulations offer a complementary perspective, since they allow selected environmental and operational conditions to be reproduced systematically. Combining the two approaches makes it possible to connect fatigue differences observed in operational data to the conditions and processing assumptions that produce them. It also makes it possible to test whether trends found in the field are simply artefacts of one dataset, or whether they reflect a more general feature of the underlying physics.

The aim of this thesis is therefore twofold. First, it aims to quantify how 10-minute segmentation and the independent representation of environmental and operational conditions affect wind-turbine fatigue estimates. Second, it aims to determine whether the resulting low-frequency cycles can be reproduced, isolated and interpreted using aeroelastic simulations. Throughout

the thesis, particular attention is given to the influence of structural response channel, Wöhler exponent, record duration and transitions between turbine operating states.

1.2.1. Research questions

Based on the problem identified above, the main research question of this thesis is:

RQ: How does representing wind-turbine load histories as independent 10-minute records affect fatigue estimates compared to longer, continuous histories?

This question is further addressed through the following subquestions:

- RQ1:** How does the difference between independent 10-minute processing and continuous processing of measured load histories develop as the observation period increases?
- RQ2:** How does the effect of 10-minute segmentation vary with structural response channel and material fatigue sensitivity?
- RQ3:** To what extent can the fatigue estimates of a continuous measured load history be reproduced without processing the full record at once?
- RQ4:** How do fatigue estimates from directly simulated one-hour histories compare with those obtained from independent and concatenated 10-minute simulations?
- RQ5:** How does preserving continuity across turbine operating-state transitions affect fatigue estimates from one-hour simulations?

1.3. Scope and contribution

To answer the research questions, this thesis combines operational measurements with controlled aeroelastic simulations. The analysis considers four structural response channels:

1. Blade root flapwise bending moment
2. Blade root edgewise bending moment
3. Tower bottom fore–aft bending moment
4. Tower bottom side–side bending moment

The measurement-based analysis investigates how the segmentation of a longer load history into independent 10-minute records affects fatigue estimates. The measured signals are therefore processed using three approaches. The first treats each selected observation period as one continuous load history and provides the reference result. The second represents the conventional approach, in which the corresponding 10-minute blocks are processed independently. The third retains information from the boundaries of successive blocks so that cycles extending across them can be accounted for without processing the complete record at once.

These approaches are applied over observation periods ranging from one hour to several months. Their results are compared across the four response channels and for different levels of material fatigue sensitivity. This makes it possible to determine whether the effect of segmentation develops with observation duration and whether it depends on the type of structural response or the assumed relationship between load range and fatigue damage.

The simulation-based analysis examines whether similar effects arise when short-duration aeroelastic simulations are used to represent longer loading histories. Directly simulated 3600 s histories are compared with results obtained from conventional 600 s simulations. The

600 s simulations are considered both independently and after concatenation into one-hour histories. Additional continuous one-hour simulations are used to investigate changes between turbine operating states, including startup and shutdown, and are compared with equivalent representations in which the individual operating segments are simulated and processed separately.

The study is limited to the selected measurement period, turbine model, structural response channels and simulated environmental and operational conditions. It does not provide a complete certification assessment or a general correction factor applicable to all turbines and sites. The focus is instead on the relative differences caused by segmentation, simulation duration and the representation of transitions between operating conditions.

The main contribution of this thesis lies in combining measured and simulated load histories to assess the consequences of representing fatigue loading as independent 10-minute records. Measurements are used to quantify the effect under real operating conditions and to evaluate an efficient method for retaining cycles that cross block boundaries. Controlled simulations are then used to distinguish the influence of record duration, concatenation and transitions between operating states. Together, the two approaches show how sensitivity to longer-duration loading varies across response channels and material fatigue assumptions, and indicate when conventional 10-minute processing may require additional treatment.

1.4. Thesis structure

Following this introduction, chapter 2 presents the theoretical background and published research relevant to the study. It introduces fatigue damage calculation, damage equivalent loads, aeroelastic fatigue simulation, rainflow cycle counting, the treatment of unresolved cycles and the influence of low-frequency environmental and operational variations.

Chapter 3 presents the measurement dataset and the associated fatigue-processing methodology. It describes the selected response channels, data preparation, operational classification and the three processing approaches used to represent continuous and segmented load histories. This methodology provides the basis for answering RQ1–RQ3.

Next, chapter 4 introduces the simulation methodology and the BHawC (Bonus Horizontal Axis Wind Turbine Code) model, Siemens Gamesa's in-house aero-servo-elastic simulation tool for wind-turbine load calculations. This chapter describes the turbine model, environmental inputs, selected design load cases, fatigue-processing assumptions and the procedures used to compare 600 and 3600 s simulations. It also presents the concatenation approach and the construction of continuous simulations containing transitions between turbine operating states. This methodology provides the basis for answering RQ4 and RQ5.

The measurement results are then presented in chapter 5 and are used to quantify how the effect of 10-minute segmentation develops with observation period, structural response channel and material fatigue sensitivity. The performance of the boundary-preserving processing approach is also evaluated against the continuously processed measurement records.

Subsequently, chapter 5 presents the aeroelastic simulation findings. It compares directly simulated one-hour histories with independent and concatenated 10-minute simulations and examines the influence of preserving continuity across operating-state transitions.

Finally, chapter 7 integrates the findings from the measurement and simulation analyses. It answers the main research question and its subquestions, discusses the extent to which the trends observed in the measurements are supported by the controlled simulations, and presents the main conclusions, limitations and recommendations for future work.

2

Literature review

2.1. Introduction

Chapter 1 established the practical context of representing wind-turbine loading through independent 10-minute records. This chapter develops the theoretical and empirical background needed to evaluate the consequences of that representation.

The review first presents the fatigue-calculation framework used throughout this investigation, including cumulative damage, Wöhler curves, rainflow counting and damage equivalent loads. It then examines how aeroelastic simulations represent the long-term wind climate and turbine operating conditions, with particular attention to how independently simulated cases are combined.

The second part of the chapter reviews methods for treating unresolved rainflow cycles and the evidence for low-frequency fatigue contributions in measured turbine responses. Previous work on operating-state transitions and the reliability of fatigue extrapolations is then considered. The chapter concludes by synthesising the limitations of the available literature and relating them to the research questions introduced in subsection 1.2.1.

2.2. Fatigue assessment workflow in wind energy

Wind-turbine fatigue assessment converts a measured or simulated load history into a set of cycles, relates those cycles to a material fatigue model and accumulates their individual damage contributions [1]. The principal steps are summarised in Figure 2.1.

However, even though the individual steps are well established, the resulting fatigue estimate remains sensitive to the selected material model, cycle-counting procedure and treatment of incomplete cycles.




Figure 2.1: Fatigue damage assessment workflow.

From load cycles to fatigue damage

Fatigue resistance is commonly described using the number of cycles that a material can withstand at a specified stress or load range. For a constant-amplitude stress range S_i , the number of cycles to failure is denoted by N_i .

Under the Palmgren–Miner linear damage accumulation rule, one cycle at this amplitude range contributes a damage amount

$$d_i = \frac{1}{N_i}. \quad (2.1)$$

If the same range occurs n_i times, the corresponding damage is

$$D_i = \frac{n_i}{N_i}. \quad (2.2)$$

The total accumulated damage is obtained by summing the contributions from all ranges:

$$D = \sum_i \frac{n_i}{N_i}. \quad (2.3)$$

Failure is commonly associated with a cumulative damage value of one, although the exact interpretation depends on the material model, safety factors and design procedure. The Palmgren–Miner rule assumes that the damage contributions can be added linearly and that their order does not affect the final result. Despite these simplifications, the method remains widely used because it provides a practical way to combine large numbers of fatigue cycles obtained from measurements or simulations [1, 13].

Wöhler curves

The relationship between load range and the number of cycles to failure N_i is commonly represented using a Wöhler curve. A Wöhler curve describes how the fatigue life of a material decreases as the applied cyclic load range increases. A power-law form can be written as

$$N_i = \left(\frac{S_0}{S_i} \right)^m, \quad (2.4)$$

where S_0 is a reference value and m is the Wöhler exponent.

A typical Wöhler curve is shown in Figure 2.2 [14]. In logarithmic coordinates, the slope of the curve is related to the exponent m .




Figure 2.2: Typical power-law Wöhler curve used in fatigue assessment.

Substituting Equation 5.2 into the Palmgren–Miner expression gives

$$D = \sum_i \frac{n_i}{(S_0/S_i)^m} = \frac{1}{S_0^m} \sum_i n_i S_i^m. \quad (2.5)$$

This expression shows that the contribution of a cycle increases with the load range raised to the power m . When m is large, relatively small changes in the largest cycle ranges can produce very large changes in damage.

This is particularly relevant for wind turbine blades, where composite material fatigue calculations commonly use larger Wöhler exponents than those used for steel tower components [15]. A small number of large cycles can therefore dominate blade fatigue, even when they represent only a very small part of the complete cycle population.

The value of m therefore controls how strongly a change in the upper part of the cycle-range distribution influences the calculated fatigue damage.

Cycle counting

The fatigue damage calculation requires the load ranges S_i and the number of times each range occurs, n_i . These quantities must therefore be extracted from the measured or simulated load time series.

A basic approach is to calculate the range between each adjacent maximum and minimum. Although this is simple, it does not reliably separate the different cycles contained within an irregular signal. In particular, smaller cycles may be nested inside a larger load variation. Treating every adjacent maximum and minimum as an independent cycle can therefore give a cycle distribution that does not represent the material response correctly [16, 17].

Consequently, rainflow counting was developed to identify these cycles more realistically. It groups the reversals in the load signal into cycles that correspond to the closed hysteresis loops produced by repeated material loading. The method was originally proposed by Matsuishi and Endo and is commonly explained using the analogy of rain flowing down a pagoda roof [18, 19]. The relationship between the load history and the corresponding material response is illustrated in Figure 2.3.




Figure 2.3: Schematic relationship between a variable-amplitude load history and the corresponding closed material hysteresis loops [20].

The upper part of Figure 2.3 shows the load history, while the lower part shows the corresponding stress and strain response. When the loading returns through an earlier level, the material response can form a closed hysteresis loop. Rainflow counting identifies these closed loops and represents them as fatigue cycles.

The first step is to reduce the signal to its turning points. These are the local maxima and minima at which the direction of the signal changes. Points located along a continuously increasing or decreasing part of the signal are removed because they do not introduce an additional reversal.

The original rainflow method can be understood by rotating the turning point history so that time runs vertically downwards, as shown in Figure 2.3. Rain is imagined to begin at each turning point and flow down the signal. A rain path stops when one of the following conditions is reached:

1. the path reaches the end of the load history;
2. the path meets rain flowing from an earlier turning point; or
3. the path reaches a later reversal that is equal to or larger than the reversal from which it started.

The paths stopped by surrounding reversals define closed cycles. Paths that reach the end of the available record without closing remain unresolved and may subsequently be counted as half-cycles, or retained as part of the rainflow residue. The graphical rainflow rules explain the basic idea, while numerical implementations apply equivalent conditions directly to the turning point sequence.

For a sequence of turning points $x_1, x_2, \dots, x_n$, the movement from one turning point to the next is called a load excursion. Its range is the absolute difference between the two points:

$$R_i = |x_{i+1} - x_i|. \quad (2.6)$$

Here, R_i denotes a possible cycle range and corresponds to the range S_i used in the fatigue damage calculation. At this stage, however, the range has not yet been confirmed as a complete cycle.

A common numerical implementation is the four-point method. Four consecutive turning points x_1, x_2, x_3, x_4 define three ranges:

$$R_{12} = |x_2 - x_1|, \quad R_{23} = |x_3 - x_2|, \quad R_{34} = |x_4 - x_3|. \quad (2.7)$$

For a four-point sequence, the inner range R_{23} forms a closed cycle when it is no larger than either surrounding range:

$$R_{23} \leq R_{12}, \quad R_{23} \leq R_{34}. \quad (2.8)$$

When this condition is satisfied, the load movement between x_2 and x_3 is enclosed by the surrounding reversals. It is therefore counted as one closed cycle, and the points x_2 and x_3 are removed from the sequence. The process is then repeated using the remaining turning points. If the condition is not satisfied, the points are retained until more of the signal becomes available [16, 17]. This process can be shown visually in Figure 2.4.




Figure 2.4: Illustration of turning point extraction, candidate ranges and the four-point rainflow closure rule.

This procedure allows smaller cycles contained within a larger load variation to be counted without losing the surrounding reversal. Rainflow counting is therefore well suited to variable-amplitude loading. Each counted cycle can then be described by its range, mean value and count [14].

A cycle that closes fully within the analysed record is assigned a count of one. Turning points that do not form a complete cycle remain unresolved. Depending on the selected treatment, they may be counted as half cycles or retained for later processing. The unresolved turning points remaining after the closed cycles have been removed form the rainflow residue. The residue contains load changes that have started but have not yet closed within the available record. These points should not automatically be discarded because they may form complete cycles when more data are included [11].

The residue is therefore not merely a numerical remainder. It stores the ordered turning points that have not yet formed complete cycles and may be required when subsequent parts of the load history are considered [2, 12]. Alternative methods for treating this information are reviewed in section 2.4.

Damage equivalent loads

Fatigue damage is useful for comparing different sets of counted cycles, but it is not always the clearest way to present the structural loads themselves. Wind turbine fatigue results are therefore often expressed using a damage equivalent load (DEL) [1, 21].

The damage equivalent load represents the constant-amplitude load range that would produce the same fatigue damage as the original variable-amplitude signal when repeated for a specified number of reference cycles.

An illustration of the concept is shown in Figure 2.5 [14].




Figure 2.5: Damage equivalent load concept, where a constant-amplitude load produces the same damage as the original variable-amplitude history.

Let R_i represent the rainflow counted load ranges and let n_i represent the associated cycle counts. The equivalent load R_{eq} is defined by

$$\sum_i n_i R_i^m = n_{eq} R_{eq}^m, \quad (2.9)$$

where n_{eq} is the selected reference number of cycles.

Solving for the equivalent load gives

$$R_{eq} = \left(\frac{1}{n_{eq}} \sum_i n_i R_i^m \right)^{1/m}. \quad (2.10)$$

The formulation above is based on load range alone and does not explicitly account for the mean value of each counted cycle. In stress-based fatigue assessment, cycles with the same range can produce different damage when their mean stresses differ. Tensile mean stress generally reduces fatigue resistance, while compressive mean stress may increase it. A mean stress correction, commonly represented using a Goodman or Haigh relationship, can therefore be applied to convert each cycle into an equivalent range before the DEL is calculated. Thomsen showed that mean load correction can have a major influence on lifetime equivalent loads for wind turbines [21].

In the present thesis, DELs are calculated directly from the counted load ranges without applying a component-specific mean correction, because the objective is a consistent relative comparison between processing methods. The same treatment is applied to every method and simulation case. The resulting DELs should therefore be interpreted as comparative load

measures rather than complete component-specific fatigue life predictions in cases where mean stress effects must be included.

The damage equivalent load has the same units as the original load range. For a bending moment channel, the equivalent load is therefore also expressed as a bending moment. The numerical value itself depends on the selected Wöhler exponent and the reference number of cycles. DEL values should therefore only be compared directly when the same values of m and n_{eq} have been used.

When several segments are combined, their DEL values cannot be averaged directly. If segment j has an equivalent load $R_{eq,j}$ and a reference count $n_{eq,j}$, the equivalent load for the combined duration is

$$R_{eq,tot} = \left(\frac{\sum_j n_{eq,j} R_{eq,j}^m}{\sum_j n_{eq,j}} \right)^{1/m}. \quad (2.11)$$

This expression is important when combining simulations from different wind-speeds, turbulence seeds or design load cases. The individual contributions must be combined at the damage level before the final equivalent load is calculated.

2.3. Fatigue evaluation using aeroelastic simulations

Aeroelastic simulations generate the load histories used in design-stage fatigue assessment by representing the coupled response of the inflow, rotor aerodynamics, structural dynamics and turbine controller. A finite set of simulations is used to represent the environmental and operational conditions expected over the turbine lifetime [4]. A simplified assessment procedure is shown in Table 2.1.

Table 2.1: Typical workflow for aeroelastic fatigue assessment.

Step	Main activity
1	Define the long-term wind climate
2	Divide the wind speed range into bins
3	Generate turbulent wind fields and independent seeds
4	Run aeroelastic simulations for each design load case
5	Extract structural load time histories
6	Apply rainflow counting and calculate fatigue damage
7	Weight the damage according to long-term occurrence
8	Combine all contributions into lifetime damage and DEL

Wind climate description

A long-term fatigue calculation requires a statistical description of the wind climate at a given site or for a specific wind class. The probability of different mean wind speeds is commonly represented using a Weibull distribution [4].

The Weibull probability density function can be written as

$$f(v; k, \lambda) = \frac{k}{\lambda} \left(\frac{v}{\lambda} \right)^{k-1} e^{-\left(\frac{v}{\lambda}\right)^k}, \quad (2.12)$$

where v is wind speed, k is the shape parameter and λ is the scale parameter. The shape parameter controls the form and spread of the distribution, while the scale parameter controls the characteristic wind-speed.

The operational wind speed range is divided into discrete bins. A set of aeroelastic simulations is then carried out for each bin, and the resulting fatigue damage is weighted by the probability of that wind-speed occurring. The discretisation of a continuous wind distribution into wind speed bins is illustrated in Figure 2.6.




Figure 2.6: Example discretisation of a Weibull distribution into wind speed bins.

The probability associated with a wind speed bin determines how often the corresponding simulated loading is expected to occur over the turbine lifetime. Wind speeds that occur frequently can therefore contribute strongly to lifetime damage even if the load ranges are not the largest. Conversely, a less frequent wind-speed can still be important when it produces sufficiently large load cycles. The final contribution depends on the combination of the occurrence probability and the damage rate calculated from the simulations.

Turbulent wind fields and stochastic seeds

For each wind speed bin, turbulent wind fields are generated to represent the random fluctuations present in atmospheric inflow. Unlike a uniform wind input, a turbulent wind field describes how the wind-speed varies with time and across the rotor area. This is important because different parts of the rotor experience different velocities, which creates varying aerodynamic loads on the blades, rotor and tower.

Spectral turbulence models are used to generate these wind fields. They define how the turbulent energy is distributed across different frequencies and spatial scales, while also reproducing the target wind speed variance, length scales and spatial coherence. The main differences between two of the most commonly used turbulence models are summarised in Table 2.2 [4, 22, 23].

Table 2.2: Main differences between the Kaimal and Mann turbulence models, based on [4, 22, 23].

Aspect	Kaimal model	Mann model
Main approach	Describes the turbulence components using separate wind-speed spectra.	Describes the turbulence components together using a three-dimensional spectral model.
Spatial variation	Spatial coherence across the rotor is added using a separate coherence model.	Spatial relationships are included directly in the model formulation.
Velocity components	The longitudinal, lateral and vertical components are mainly defined separately.	The relationships between the three velocity components are included in the model.
Main inputs	Uses parameters such as turbulence intensity, length scales and coherence.	Uses parameters describing turbulence strength, length scale and shear distortion.
General characteristic	Provides a relatively simple and widely used representation of turbulent inflow.	Provides a more detailed three-dimensional representation of turbulent inflow.

A wind field generated at a given mean wind speed represents only one possible realisation of the turbulent inflow. A different random seed produces a different sequence and spatial distribution of turbulent structures, even when the same mean wind speed, turbulence intensity and model parameters are used. These differences can change the timing and magnitude of the loads experienced by the turbine. Using only one seed could therefore make the calculated fatigue load strongly dependent on one particular turbulent realisation. Multiple seeds are used to reduce this dependence and provide a more representative estimate. The fatigue results can then be averaged across the seeds, while the spread between them gives an indication of the variability caused by turbulence.

Increasing the number of seeds generally improves the statistical reliability of the result, but it also increases the required simulation time. Studies focused on uncertainty may use hundreds of seeds to describe the full distribution of possible responses [24]. Industrial simulation campaigns normally use fewer seeds and aim to achieve a reasonable balance between statistical convergence and computational cost [24, 25].

Design load cases

The IEC design load cases specify the operating and environmental conditions that must be represented during turbine assessment for certification. Normal power production under turbulent inflow is one of the main fatigue cases, but other relevant conditions include startup, shutdown, idling, parked operation and electrical faults [4].

For each design load case, simulations are carried out over the required wind speed range according to the standards. These simulations produce time histories of the structural response, including quantities such as blade root moments, tower bottom moments, shaft torque and bearing loads. Fatigue simulations are commonly run for 600 seconds, which matches the 10-minute interval used for wind statistics and allows a large number of independent turbulence seeds to be generated within a practical computational effort.

The individual simulations are then treated as representative samples of the turbine response under each condition. Results are averaged across the available seeds to reduce the influ-

ence of any single turbulent realisation, after which the damage contributions are weighted according to how often each operating condition is expected to occur.

Lifetime fatigue estimation

The previous subsections described how the long-term wind climate, turbulent seeds and design load cases are represented through a finite set of aeroelastic simulations. The final step is to combine the resulting damage values into an estimate of the total fatigue accumulated over the turbine lifetime.

For normal power production, let $\dot{D}(V_i)$ denote the damage rate for wind speed bin i , averaged across the available turbulence seeds, and let $P(V_i)$ denote the long-term probability of that wind-speed. The lifetime damage from power production is then

$$D_{\text{production}} = T_{\text{life}} \sum_i P(V_i) \dot{D}(V_i), \quad (2.13)$$

where T_{life} is the assumed turbine lifetime.

The remaining design load cases must also be included according to how they occur. Continuous states, such as idling or parked operation, are weighted by the expected time spent in those conditions. Discrete events, such as startup, shutdown or electrical faults, are weighted by their expected number of occurrences. The total lifetime damage can therefore be expressed as

$$D_{\text{lifetime}} = D_{\text{production}} + \sum_j D_{\text{state},j} + \sum_k N_k D_{\text{event},k}, \quad (2.14)$$

where $D_{\text{state},j}$ is the lifetime damage associated with operating state j , $D_{\text{event},k}$ is the damage caused by one occurrence of event k , and N_k is the expected number of occurrences during the turbine lifetime. Once all contributions have been combined at the damage level, the total can be converted into a lifetime damage equivalent load using the selected Wöhler exponent and reference number of cycles.

Uncertainty in simulation-based fatigue estimates

Lastly, aeroelastic fatigue estimates are affected by uncertainty at each stage of the assessment. The main sources include the long-term wind climate, turbulence model, number of stochastic seeds, turbine controller, structural properties, material model and lifetime extrapolation procedure.

Some of these uncertainties arise from the way the environmental conditions are represented. The selected wind speed bin width determines how closely the discrete simulations reproduce the continuous wind climate, while the number of seeds determines how well the variation between turbulent realisations is captured. The selected design load cases also control which operating states and transient events contribute to the final lifetime estimate.

Other uncertainties are introduced by the turbine model and fatigue calculation. Differences in controller behaviour or structural properties can change the simulated load response, while the selected Wöhler exponent and damage accumulation assumptions influence how the counted cycles are converted into fatigue damage. As a result, an accurate wind input does not by itself guarantee an accurate lifetime fatigue estimate.

Recent studies have explored several ways to improve the reliability and efficiency of simulation-based fatigue estimates. Probabilistic approaches account for the full range of possible turbulence conditions and structural responses, rather than relying only on a small number of representative values [25]. Numerical quadrature methods, which approximate a weighted integral using a limited set of carefully selected evaluation points, can reduce the number of simulations required while still providing accurate equivalent fatigue loads [26]. Other studies combine aeroelastic simulations with measurements or statistical models to represent operational loading more realistically [27].

These studies mainly address uncertainty within the individual simulations or in the weighting of their results. The present thesis considers a different but related issue. Even when each simulation accurately represents its assigned condition, processing short simulations independently can exclude fatigue cycles associated with slow changes between conditions. The resulting error is therefore not caused by the accuracy of a single simulation, but by the way the separate load histories are divided, counted and combined.

This distinction concludes the conventional simulation-based fatigue assessment described in this section. The following sections examine how finite processing windows and rainflow residue treatment affect the recovery of longer fatigue cycles and, consequently, the estimated lifetime damage.

2.4. Rainflow residue processing and low-frequency fatigue

As introduced in Equation 2.2, rainflow counting generally leaves an ordered sequence of unresolved turning points. The treatment of this residue determines whether it remains available for combination with later parts of the load history.

Rainflow counting is formalised in standards such as ASTM E1049 [16]. However, differences can still arise from the reduction of turning points, the comparison of ranges and the treatment of the final residue. Amzallag et al. proposed a standardised procedure to improve the consistency and reproducibility of rainflow results [17]. Their work also emphasised that the residue contains useful information and should not automatically be discarded.

Marsh et al. reviewed several methods for processing rainflow residues and showed that the selected treatment can have a considerable effect on fatigue damage, especially for non-stationary signals containing slow changes or large transitions [11]. The main difficulty is that a conventional cycle histogram stores the ranges and counts of the identified cycles but does not preserve the order of the unresolved turning points. Once this order has been lost, cycles crossing the segment boundaries cannot be reconstructed.

A residue-aware procedure avoids this problem by storing two outputs from each segment. The cycles that close within the segment are accumulated directly, while the unresolved turning points are retained in their original order. Residues from consecutive segments can then be joined and rainflow counted again over a longer period. This allows cycles that began in one segment and closed in another to be recovered.

In wind turbine research, the effect of these longer cycles is often described as low-frequency fatigue dynamics. In this context, low-frequency does not refer specifically to a structural eigenfrequency. It refers more generally to fatigue cycles whose duration is longer than the processing window, which is commonly 10 minutes. These cycles can result from gradual changes in wind-speed, wind direction, turbulence level or turbine operating state. They can also be associated with wakes, weather systems, curtailment, startup, shutdown and transitions between idling and power production [2, 28]. While rotor-related cycles may close repeatedly within a

10-minute interval, cycles linked to these slower changes can extend across several intervals. Sadeghi et al. investigated this effect using offshore wind turbine strain measurements stored in 10-minute blocks [12]. Their method first rainflow counted each block and stored its unresolved residue. The residues were then joined in chronological order and counted again over a longer period. The additional cycles identified from the combined residue represented the fatigue contribution that was not recovered by independent blockwise processing. A schematic of this procedure is shown in Figure 2.7.




Figure 2.7: Residue based recovery of fatigue cycles extending across the original processing windows, based on Sadeghi et al. [12].

The effect of residue recovery can be quantified using the DEL ratio between the standard approach, and that which includes low-frequency fatigue dynamics (LFFD) recovery:

$$r_{\text{DEL}} = \frac{\text{DEL}_{\text{STA}}}{\text{DEL}_{\text{LFFD}}}. \quad (2.15)$$

A value of one indicates no difference between the methods, while a value below one indicates that independent blockwise processing gives a lower DEL, and thus a lower fatigue damage estimation.

Sadeghi et al. later extended the analysis to several offshore turbines and different structural response channels [2]. The results showed that the importance of low-frequency cycles can vary considerably between turbines and between channels on the same turbine. This variation can partly be explained by the physical source of each load. A tower response measured in a fixed direction may be sensitive to gradual changes in wind direction. Blade flapwise loading may respond strongly to variations in aerodynamic thrust. Blade edgewise loading, however, may be dominated by gravity related rotor cycles that close many times within each 10-minute block. Furthermore, the effect also increases with the Wöhler exponent because larger values of m give greater importance to the largest recovered cycles, indicating that material choice also has a significant impact.

The available literature therefore does not support the use of a single universal low-frequency

correction factor. Its magnitude depends on the turbine, response channel, operating conditions, processing duration and material assumptions. DEL ratios can quantify the overall effect, but they do not explain which cycles are recovered or why they are damaging. This requires a closer examination of the cycle range distributions, recovery time scales and corresponding turbine operating states.

2.5. Operational transitions and fatigue damage

The literature on low-frequency fatigue indicates that part of the additional cycle population is associated with changes between turbine operating states. Startup, shutdown, curtailment and transitions between idling and production alter rotor speed, blade pitch, aerodynamic thrust and the mean structural response over time.

Faria et al. investigated the relation between low-frequency fatigue damage and operating conditions using measured tower strain data [28] for both onshore and offshore turbines. The load cycles were linked to operational information derived from SCADA signals such as power, rotor speed and pitch.

Their analysis used damage matrices, in which each fatigue cycle was classified according to the wind speed or operating condition at its start and end. Cycles that begin and end in the same condition appear on the diagonal of the matrix. Cycles that cross between different conditions appear outside the diagonal. Consequently, the off-diagonal terms provide a direct representation of transition cycles. An illustration of this is shown in Figure 2.8. These transition cycles cannot be identified correctly when each 10-minute block is treated as an isolated record and assigned to only one operating state.




Figure 2.8: Damage matrix showing cycles classified by the turbine condition at their start and end.

The results showed that transitions between normal operation, idling, curtailment and shut-down can contribute substantially to the recovered damage [28]. The effect was particularly

relevant for offshore turbines, where wave-induced loading introduces slower variations that can persist beyond the standard processing window.

This finding is relevant to both measurement processing and aeroelastic simulation. In a conventional design load case campaign, each operating condition is represented by a separate simulation, after which the resulting fatigue contributions are combined statistically. Although this approach captures the damage accumulated within each operating state, it does not automatically account for cycles formed as the turbine transitions between them. Dedicated startup and shutdown simulations represent part of this behaviour, but they are still commonly processed independently from the conditions that precede and follow them.

Continuous simulations containing several operating states can therefore provide additional information. By preserving the structural and controller response throughout the full sequence, they allow rainflow cycles to develop across the transitions and capture load changes that would otherwise be separated between individual simulations.

2.6. Reliability of long-term fatigue estimates

Fatigue estimates derived from measurements or simulations are generally extrapolated over periods much longer than the available records. Their reliability therefore depends on whether the analysed period represents the environmental and operational variability relevant to long-term turbine operation.

Hübner et al. investigated the reliability of fatigue damage extrapolation using offshore strain measurements and showed that both the lifetime estimate and its uncertainty depend on the selected binning and extrapolation procedure [3]. Their results also indicated that the available measurement period must be sufficiently long before the estimated fatigue behaviour becomes representative. For the dataset considered, approximately nine to ten months of measurements were required for the extrapolated damage to stabilise [3].

This is directly relevant to low-frequency fatigue assessment. The difference between conventional and residue-aware processing should be evaluated over a sufficiently long period, since a ratio derived from only a few days of data may be strongly influenced by a single storm, operational transition or other isolated event. However, convergence of this ratio does not necessarily mean that the overall lifetime damage estimate is representative. The available record may still underrepresent particular seasons, wind directions, sea states or turbine operating conditions.

Residue recovery therefore addresses only one part of the problem. It improves the completeness of the rainflow count by recovering cycles that extend across the selected processing windows, but it cannot compensate for environmental or operational conditions that are absent from the measurement period. A reliable lifetime assessment consequently requires both an appropriate treatment of rainflow residues and a sufficiently representative observation period.

2.7. Synthesis and identified research gap

The reviewed literature provides a well-established framework for converting wind-turbine load histories into fatigue estimates. Rainflow counting identifies cycles from ordered load reversals, Wöhler curves determine their relative fatigue importance, and Palmgren–Miner accumulation allows contributions from different conditions to be combined. Aeroelastic fatigue assessments apply these methods to a finite set of environmental and operational cases intended to represent the turbine lifetime.

Studies of rainflow residue processing show that unresolved turning points contain information

that is not retained by cycle histograms alone [11, 17]. Preserving their chronological order allows later parts of the load history to complete cycles that remain unresolved within an individual segment. Applications to turbine measurements demonstrate that the additional fatigue contribution can vary between turbines, response channels, processing durations and Wöhler exponents [2, 12].

Three limitations remain in the measurement-based evidence. First, the development of the difference between independent and continuous processing has not been evaluated consistently over a common sequence of increasing observation periods. Second, the dependence of that difference on structural response channel and material fatigue sensitivity has not been established using the same underlying records and processing assumptions. Third, although residue-preserving methods have been applied successfully, their agreement with a fully continuous reference requires explicit verification for the channels and observation periods considered here. These limitations correspond to RQ1, RQ2 and RQ3.

A related limitation concerns aeroelastic simulations. Conventional fatigue campaigns combine damage from independently simulated wind-speeds, stochastic seeds and operating conditions, but there is limited direct evidence on whether a one-hour simulation produces the same fatigue estimate as an equivalent collection of 10-minute simulations. Concatenating short histories preserves a longer sequence during post-processing, but it does not create continuous inflow, structural motion or controller states at the original file boundaries. The comparison required by RQ4 must therefore distinguish directly simulated, independently processed and concatenated histories.

Measurement-based damage matrices further indicate that cycles connecting different operating states may contribute to the longer-duration fatigue response [28]. Existing studies can associate these cycles with operational transitions, but they provide limited control over the transition sequence and surrounding conditions. The effect of preserving structural and controller continuity throughout a complete operating sequence consequently remains uncertain. This is the gap addressed by RQ5.

Collectively, the literature establishes the relevance of retaining information beyond individual processing windows but does not determine consistently when doing so materially changes a wind-turbine fatigue estimate. The research questions therefore examine the development of the effect with observation length, its dependence on structural and material characteristics, the performance of residue-preserving processing, its representation in aeroelastic simulations and the influence of operating-state continuity. The following chapters present the methods used to investigate these questions.

Measurement-based fatigue methodology

The first part of the investigation uses measured load histories from an operational wind turbine to determine how representing the data as independent 10-minute records affects the resulting fatigue estimates under real operating conditions.

The measured histories are processed using three approaches: continuous rainflow counting, conventional independent 10-minute counting and residue-aware LFFD processing. Their results are compared to quantify the effect of segmentation and assess whether the cycles identified in the continuous record can be retained without processing the complete history at once. The comparison is repeated across several structural response channels, Wöhler exponents and observation periods.

In addition to the resulting damage and DEL values, the cycles responsible for the differences are examined according to their range, aggregation scale and associated turbine operating state. This provides insight into both the magnitude of the segmentation effect and the loading conditions from which it originates.

The purpose of the analysis is to isolate the influence of the processing method rather than to produce a certified lifetime estimate for the turbine. All approaches are therefore applied to the same measured intervals using consistent Wöhler exponents, reference cycle counts and DEL normalisation. Differences between the results can consequently be attributed to the treatment of the load history rather than to differences in data duration or calculation assumptions.

The measurement results subsequently guide the BHawC investigation by identifying the response channels, timescales and operating transitions that are most relevant for examination under controlled simulation conditions.

3.1. Measurement dataset

The dataset contains measured load and SCADA data from a 7 MW offshore direct-drive wind turbine with a rotor diameter of 154 m. The selected measurement period extends from 1 May 2022 to 1 November 2022 and forms an approximately six-month continuous record without significant data gaps.

The blade root bending moments were obtained from LOPC (Load-Optimized Pitch Control) strain sensors, while the tower bottom bending moments were derived from strain gauges

installed near the tower base and calibrated in two orthogonal directions. As the different sensors were originally recorded at different sampling frequencies, all load signals were re-sampled onto a common 20 Hz time grid before the fatigue analysis. This ensured consistent timestamps and temporal alignment between the channels. The resulting record was then trimmed to an integer number of complete 10-minute intervals so that each processing method analysed exactly the same data. A more detailed explanation on the measurement filtering used in this investigation can be found in section A.1.

Four structural response channels are considered:

1. Blade root flapwise bending moment
2. Blade root edgewise bending moment
3. Tower bottom bending moment in the X direction
4. Tower bottom bending moment in the Y direction

These channels provide a comparison between different physical loading mechanisms. Blade flapwise and tower bending responses are strongly influenced by aerodynamic thrust, turbulence, wind direction and changes in turbine operation. Blade edgewise loading is more strongly governed by gravity-induced rotor cycles, which repeat several times within a conventional 10-minute interval. The edgewise channel therefore provides a useful comparison against the responses expected to be more sensitive to slow load variations.

The available tower measurements were provided as bending moments in two fixed orthogonal directions, X and Y. The Y axis is aligned with geographical North, while the X axis points East. This coordinate system is appropriate for tower-base design assessment because the tower itself does not rotate with the nacelle. As a result, tower design loads are commonly examined in fixed X and Y directions rather than only in fore-aft and side-side directions, which are defined relative to the instantaneous nacelle and rotor orientation. In principle, the fixed-direction loads could be transformed into fore-aft and side-side components using the corresponding nacelle yaw angle, however, this information was not available for the turbine considered in this study.

Using fixed geographical axes also preserves the influence of wind direction in the measured response. As the nacelle yaws, the aerodynamic thrust vector changes its projection onto the X and Y axes, causing the relative loading in the two tower channels to vary. The measured X and Y moments therefore provide both the orthogonal bending components required for the fatigue comparison and information about longer-term changes in wind direction and turbine orientation.

The accompanying SCADA data include generator speed, blade pitch reference, grid power and 10-minute mean wind speed. Following a simplified version of the approach used in [28], each 10-minute interval is classified into one of three operating regimes:

1. *Fault*
2. *Idling*
3. *Operational*

The classification is based primarily on generator speed, pitch reference and grid power, since wind speed alone does not uniquely describe the turbine state. This is described in further detail in section A.3. The selected thresholds are intended to distinguish the main operating conditions associated with different structural loading levels rather than reproduce every con-

troller mode in detail. Curtailed periods may therefore remain within the Operational category when the rotor continues to turn and the turbine remains connected to the grid.

3.2. Fatigue-processing strategies

The core comparison applies three fatigue-processing strategies to the same measured signal. These represent a continuous reference, the conventional independent 10-minute approach and a residue-aware LFFD approach.

Continuous-reference processing

For the continuous-reference method, the selected measurement interval is treated as one uninterrupted signal. Rainflow counting is applied directly to the full load history, allowing cycles to form across all internal 10-minute boundaries.

The continuous result is the most complete cycle count available for the selected finite record and is therefore used as the reference. It is not an absolute representation of lifetime fatigue because cycles may still remain unresolved at the beginning and end of the available interval. Nevertheless, it provides an appropriate benchmark for evaluating the effect of dividing that same interval into shorter blocks.

The verified MATLAB `rainflow` function, included in the Signal Processing Toolbox, is used to obtain the reference results for the validation windows, hereafter referred to as `FLS.rainflow`. Continuous processing becomes less convenient as the data volume increases, which is one reason why long-term monitoring data are commonly stored and analysed in shorter intervals.

Conventional 10-minute processing

The conventional method divides the measured history into consecutive 10-minute blocks. Each block is rainflow counted independently and its fatigue contribution is accumulated with the contributions from the remaining blocks.

The unresolved residue from one block is not propagated into the next. Consequently, a load excursion that begins in one block and closes in a later block cannot be reconstructed as one continuous cycle. Depending on the residue treatment used within the individual blocks, it may instead appear as several smaller excursions or remain unresolved.

This method represents the standard segmented workflow and provides the baseline against which residue-aware processing is evaluated.

Residue-aware LFFD processing

The residue-aware method retains two outputs from every 10-minute block: the cycles that close within the block, and the ordered sequence of unresolved turning points. The closed-cycle contributions are accumulated directly, while the residues are joined in chronological order and rainflow counted again.

This second counting stage allows turning points from different blocks to form complete cycles. This procedure is further explained in section A.2 using real measurement data from the turbine being evaluated. The method therefore preserves the practical 10-minute data structure without requiring the original high-frequency blocks to be concatenated and recounted in full, following the approach in [11, 12, 17]. This greatly reduces the required computational storage, and thus cost.

Figure 3.1 illustrates the three approaches. The continuous reference counts the full interval as one signal. The conventional approach treats each 10-minute block independently. The

residue-aware approach accumulates the cycles closed within each block and separately combines the ordered residues to recover cycles crossing the original boundaries.




Figure 3.1: Continuous-reference, conventional blockwise and residue-aware fatigue-processing strategies.

3.3. Implementation of the residue-aware workflow

The residue-aware workflow is implemented in two stages. First, each 10-minute block is processed independently and its relevant fatigue information is stored. The stored results are then combined over the required analysis period to recover cycles that cross the original block boundaries. This avoids repeatedly loading and rainflow counting the complete high-frequency signal.

Block-level processing and storage

The measurement record is divided into consecutive 600 second blocks. For each block, repeated values are removed and the load history is reduced to an alternating sequence of local maxima and minima. The resulting turning-point sequence is then processed using the four-point rainflow algorithm.

The algorithm separates the block into closed cycles and an unresolved residue. For each Wöhler exponent m , the fatigue contribution of the closed cycles is stored as

$$D_{\text{closed},k}(m) = \sum_{i \in \mathcal{C}_k} n_i R_i^m, \quad (3.1)$$

where $\mathcal{C}_k$ is the set of cycles closing within block k , R_i is the cycle range and n_i is its count. The ordered residue sequence is stored separately. The original high-frequency samples are no longer required once these two outputs have been obtained.

This preprocessing is carried out once for each load channel. For every 10-minute block, it

saves the fatigue contribution from cycles that close within the block together with the unresolved turning points needed to recover longer cycles later.

Recovery over an extended period

For each 10-minute block k , the rainflow procedure returns the damage contribution from cycles that close within the block and an ordered residue vector

$$\mathbf{r}_k = [r_{k,1}, r_{k,2}, \dots, r_{k,N_{r,k}}], \quad (3.2)$$

where $r_{k,i}$ is the i -th unresolved turning point in block k , and $N_{r,k}$ is the number of turning points remaining in its residue.

For an analysis period extending from block a to block b , the closed-cycle contributions are first added directly:

$$D_{\text{closed}}(m) = \sum_{k=a}^b D_{\text{closed},k}(m). \quad (3.3)$$

The corresponding residue vectors are then concatenated in chronological order:

$$\mathbf{r}_{a:b} = [\mathbf{r}_a, \mathbf{r}_{a+1}, \dots, \mathbf{r}_b]. \quad (3.4)$$

After concatenation, duplicate and non-alternating points at the edges are removed so that the sequence remains a valid turning-point history. The combined residue is then rainflow counted once using the same four-point algorithm. Cycles identified at this stage form the recovered low-frequency contribution,

$$D_{\text{LF}}(m) = \sum_{i \in \mathcal{C}_{\text{LF}}} n_i R_i^m. \quad (3.5)$$

These cycles contain turning points originating from different 10-minute blocks and would therefore be absent from conventional independent processing.

A final residue may remain even after the combined sequence has been counted. In the present implementation, consecutive excursions within this final residue are treated as half cycles:

$$D_{\text{tail}}(m) = \frac{1}{2} \sum_{j=1}^{N_r-1} |r_{j+1} - r_j|^m, \quad (3.6)$$

where N_r is the number of turning points remaining in the final residue.

The complete residue-aware fatigue contribution is therefore

$$D_{\text{LFFD}}(m) = D_{\text{closed}}(m) + D_{\text{LF}}(m) + D_{\text{tail}}(m). \quad (3.7)$$

Once each 10-minute block has been processed, the saved results can be reused for any longer analysis period. To study different time scales, the blocks are grouped into intervals

such as one day, one week or one month. Residue recovery is then applied within each interval, and the fatigue contributions from all intervals are added together. A visual representation of this process is shown in Figure 3.2.




Figure 3.2: Computational workflow used to process closed cycles and propagate the ordered rainflow residue between consecutive 10-minute blocks.

This implementation follows the residue-concatenation approach described in [11, 12, 17]. Since the recovery procedure relies on the custom four-point algorithm, its agreement with the established rainflow implementation is verified in the following section.

3.4. Rainflow implementation and validation

The residue-aware procedure requires access to the ordered turning points remaining after the closed cycles have been extracted. These intermediate outputs are not available from the existing `FLS.rainflow` workflow. A custom four-point rainflow implementation, referred to as RF4PT, was therefore developed.

RF4PT first reduces the signal to its turning points and then applies the four-point closure criterion introduced in Equation 2.2. In addition to the closed cycles, it returns the final ordered residue and the time indices associated with the cycle endpoints. The latter are required to connect recovered cycles to the corresponding turbine operating states. Due to confidentiality restrictions, the source code and detailed intermediate outputs of the implementation cannot be reproduced in this thesis. The underlying algorithm and validation procedure are nevertheless described in sufficient detail in the following sections.

The custom implementation was validated before it was used for the main analysis. Continuous signals were counted using both `RF4PT` and `FLS.rainflow`, and the resulting cycle sets and fatigue sums were compared. The validation used 20 randomly selected windows and several fixed durations ranging from 10 minutes to one month. This tested both numerical agreement and sensitivity to the position of the analysis interval within the six-month record. The corresponding results are reported in section 5.2.

The complete residue-aware workflow was also compared with the continuous reference. Agreement between these two methods confirmed that the blockwise storage of closed cycles and ordered residues preserved the fatigue information contained in the original signal. Once this agreement had been established, differences between the conventional and LFFD results could be attributed to residue propagation rather than to the underlying rainflow implementation.

3.5. Fatigue measures and comparison metrics

For a counted set of cycle ranges R_i , cycle counts n_i and Wöhler exponent m , the quantity governing the relative fatigue contribution is

$$D = \sum_i n_i R_i^m. \quad (3.8)$$

For a fixed S–N curve, this quantity is proportional to Palmgren–Miner damage. The material intercept is not required for the present comparisons because it is common to all methods and cancels in the ratios. The absolute values of D are therefore not interpreted as physical consumed-life fractions.

The corresponding damage equivalent load is calculated as

$$\text{DEL} = \left(\frac{D}{N_{\text{ref}}} \right)^{1/m}. \quad (3.9)$$

The reference cycle count N_{ref} is defined for the measurements as

$$N_{\text{ref}} = f_{\text{eq}} T, \quad (3.10)$$

where $f_{\text{eq}} = 1 \text{ Hz}$ and T is the valid duration of the analysed interval. Consequently, a 10-minute interval is normalised using $N_{\text{ref}} = 600$ cycles. The same duration and equivalent frequency are used for all methods within each comparison, so the reference count cancels in the resulting DEL ratios.

During validation, the continuous `FLS.rainflow` reference is used as the denominator:

$$r_{\text{RF4PT}} = \frac{\text{DEL}_{\text{RF4PT}}}{\text{DEL}_{\text{REF}}}, \quad r_{\text{STA}} = \frac{\text{DEL}_{\text{STA}}}{\text{DEL}_{\text{REF}}}, \quad r_{\text{LFFD}} = \frac{\text{DEL}_{\text{LFFD}}}{\text{DEL}_{\text{REF}}}. \quad (3.11)$$

Once the LFFD method has been shown to reproduce the continuous reference, the main segmentation effect is expressed more directly through

$$r_{\text{DEL}} = \frac{\text{DEL}_{\text{STA}}}{\text{DEL}_{\text{LFFD}}} \quad (3.12)$$

A value of one indicates that the conventional and residue-aware methods produce the same result. A value below one indicates that independent 10-minute processing gives a lower fatigue estimate. The DEL ratio is therefore used as the main engineering comparison moving forward.

3.6. Analysis across time scales, Wöhler exponents and cycle ranges

The segmentation effect is examined across different observation lengths, residue aggregation periods and Wöhler exponents. The cycle-range distributions are also compared to identify which parts of the cycle population are responsible for the resulting fatigue differences.

First, the measurement period is gradually extended and the cumulative ratios in Equation 3.12 are recalculated. This shows how the difference between conventional and residue-aware processing develops as more data and operating events are included. If the ratio approaches a stable value, the result has become less sensitive to individual events within the available record. However, this does not necessarily mean that the six-month period is fully representative of the turbine lifetime [3].

The analysis also investigates how long the residues must be combined before most of the missing fatigue contribution is recovered. Residue recovery is repeated using grouping periods of 10 minutes, 1 hour, 1 day, 1 week, 2 weeks, 1 month, 3 months and the complete six-month record. For each grouping period T , the recovered fatigue share is calculated as

$$f_{\text{DEL}}(T) = \frac{\text{DEL}_{\text{LFFD}}(T)}{\text{DEL}_{\text{LFFD}}(T_{\text{full}})}. \quad (3.13)$$

where $\text{DEL}_{\text{LFFD}}(T)$ is obtained by applying residue recovery within intervals of length T , and $\text{DEL}_{\text{LFFD}}(T_{\text{full}})$ is obtained using the complete six-month record. A value of $f_{\text{DEL}}(T) = 1$ means that the selected grouping period recovers the same fatigue contribution as the full record. Lower values indicate that cycles requiring a longer aggregation period remain unresolved.

The complete analysis is repeated for several Wöhler exponents. The blade channels are evaluated using

$$m = 10, 11, 12, \quad (3.14)$$

while the tower channels are evaluated using

$$m = 3, 4, 5. \quad (3.15)$$

These values represent the higher exponents commonly associated with composite blade materials and the lower range relevant to steel tower components respectively. The purpose is not to define one exact material model, but to determine how strongly the segmentation effect depends on the weighting assigned to the largest cycle ranges.

Finally, cycle-count distributions are compared before and after residue recovery using the same range bins. The histograms show whether the difference between the methods is spread across the complete cycle population or concentrated in a small number of cycles. They also reveal whether residue recovery introduces previously missing ranges or redistributes shorter excursions into larger cycles.

The cycle distributions are interpreted together with the selected Wöhler exponent because fatigue scales with R^m . Consequently, even a small number of recovered cycles at the upper end of the range distribution may dominate the fatigue contribution, particularly for the blade channels where larger values of m are used.

3.7. Operational-state attribution

As discussed in section 2.5, changes between turbine operating states may generate large load cycles that extend across several 10-minute blocks. To investigate this effect, the time indices retained by RF4PT are used to identify the SCADA state at the beginning and end of every counted cycle, following a simplified version of the damage-matrix approach used in [28].

Each cycle is classified by the state pair

$$a \rightarrow b, \quad (3.16)$$

where a is the operating state at the beginning of the cycle and b is the state at its end. For instance, a cycle beginning and ending during normal operation is therefore classified as Operational–Operational, while a cycle beginning during idling and ending during operation is classified as Idling–Operational.

The fatigue contribution associated with each state pair is expressed as a fraction of the total damage:

$$M_{ab} = \frac{\sum_{c \in \mathcal{C}_{ab}} n_c R_c^m}{\sum_{c \in \mathcal{C}} n_c R_c^m}, \quad (3.17)$$

where $\mathcal{C}_{ab}$ contains the cycles classified as $a \rightarrow b$, while $\mathcal{C}$ contains the complete cycle population. The values M_{ab} are arranged in a damage matrix, with the state at the beginning of the cycle shown along one axis and the state at its end shown along the other. Entries on the diagonal represent cycles that begin and end in the same operating state. Off-diagonal entries represent cycles crossing between different states and therefore provide a direct measure of the fatigue associated with operational transitions, similarly to Figure 2.8.

For conventional processing, each 10-minute block is assigned a single operating-state label and counted independently for simplicity. All cycles identified within that block are therefore classified as beginning and ending in the same state, meaning that the resulting matrix contains only diagonal terms. Off-diagonal transition cycles can thus only be identified after the ordered residues have been combined across consecutive blocks.

Moving forwards, the measurement results are subsequently used to identify the response channels, aggregation timescales and operating transitions examined in the BHawC study. The two analyses are compared at the level of channel sensitivity and physical loading mechanism rather than through event-by-event reproduction, because the measured three-dimensional inflow, nacelle heading and complete controller state required for a defensible one-to-one comparison are not available.

4

BHawC aeroelastic model and simulation methodology

The measurement-based investigation established how dividing a continuous load history into independent 10-minute blocks can alter the counted cycle population. It also identified the structural channels, cycle ranges and operating transitions most relevant to the recovered low-frequency fatigue contribution. However, the measured data does not allow the wind field, turbine state and controller response to be varied independently. BHawC simulations are therefore used to examine the same mechanisms under controlled conditions.

The simulation study begins by examining the effect of simulation duration. Standard 600 s simulations are compared with directly simulated 3600 s load histories under the same design load cases. The analysis then moves to continuous 3600 s simulations to investigate the additional fatigue cycles caused by transitions between operating states. Finally, concatenating existing 600 s simulations is assessed as a practical alternative to directly simulating longer load histories.

The purpose is not to perform a complete certification assessment or independently validate every part of the aeroelastic model. Instead, a consistent turbine model and fatigue-processing framework are used to determine whether simulation duration, signal segmentation and the representation of changing conditions affect the resulting fatigue estimate.

4.1. BHawC aeroelastic model

General model description

BHawC is a time-domain aero-servo-elastic simulation tool developed for the calculation of wind turbine dynamics and design loads. The turbine is represented through interconnected structural submodels describing the blades, hub, drivetrain, nacelle, tower and foundation. Structural motion is coupled to the aerodynamic loading, gravitational forces, controller actions and environmental inputs throughout the simulation [29, 30].

The structural response can be expressed in general form as

$$\mathbf{f}_{\text{inertia}}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}) + \mathbf{f}_{\text{damping}}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{f}_{\text{internal}}(\mathbf{q}) = \mathbf{f}_{\text{external}}, \quad (4.1)$$

where $\mathbf{q}$ contains the model degrees of freedom. The external force vector includes aerody-

dynamic loading, gravity, controller-related forces and, if enabled, hydrodynamic and other environmental loads. BHawC mainly represents the flexible components using beam elements within a co-rotational formulation, allowing the large rigid-body rotations of the rotor to be combined with elastic blade and tower deformation [29, 30].

Aerodynamic forces are calculated along the blades following blade element momentum (BEM) theory. At each aerodynamic calculation point, the local relative velocity and angle of attack are determined from the incoming wind, rotor motion and structural deformation. The resulting airfoil forces are then transferred to the structural model. Corrections for effects such as tip and root losses, high induction and unsteady aerodynamic behaviour are included within the existing turbine model [30].

The turbine controller is coupled directly to the aeroelastic calculation. Generator torque and blade pitch commands therefore respond to the simulated rotor speed and operating state as time develops. This coupling is particularly important for the present study because changes in pitch, rotor speed and generator loading influence the mean structural response during startup, production and shutdown.

Turbine configuration

The simulations represent the same 7 MW offshore direct-drive turbine platform considered in the measurement study. The model information is summarised in Table 4.1. Only the details required to interpret the fatigue comparisons are reported, while proprietary structural and controller parameters are not shown.

Table 4.1: Main characteristics of the BHawC turbine model.

Parameter	Model description
Rated power	7 MW
Drivetrain configuration	Direct drive
Rotor diameter	154 m
Number of blades	3
Hub height	117.9 m
Installation type	Offshore
Support structure	Monopile
BHawC version	7.2

The same structural model, aerodynamic data and controller configuration are retained throughout the simulation campaign. Differences between the simulation sets therefore arise from the duration, environmental input or prescribed operating sequence rather than from changes to the underlying turbine model template.

Selected response channels

The simulation analysis uses the same four general structural responses as the measurement study. These comprise blade root flapwise and edgewise bending, and tower-bottom bending in two orthogonal directions, which are now aligned with the varying wind direction unlike before. Their BHawC channels and fatigue parameters are listed in Table 4.2. The differences in m reflect the material associated with each load channel, while N_{ref} follows the Siemens Gamesa standard convention used for lifetime fatigue assessment of these channels.

Table 4.2: BHawC load responses and fatigue parameters used in the comparisons.

Simulated response	m	N_{ref}
Blade root flapwise moment	10.25	2×10^6
Blade root edgewise moment	10.25	2×10^6
Tower-bottom fore-aft moment	4	1×10^7
Tower-bottom side-side moment	4	1×10^7

The blade root flapwise moment is mainly driven by aerodynamic loading and is therefore expected to respond strongly to changes in wind speed, turbulence and blade pitch. The edgewise moment contains a larger gravity-driven component, which repeats with rotor rotation and provides a useful comparison with the more aerodynamically sensitive responses.


The simulated tower channels represent fore-aft and side-side bending. This differs from the measurements, where the tower-bottom moments are recorded in fixed geographical X and Y directions. As the nacelle rotates, the aerodynamic loading is distributed differently between these two fixed axes. The measured channels therefore depend on both the structural response and the nacelle orientation.

The simulated and measured tower channels are consequently not directly equivalent in direction. They are compared mainly in terms of their broader fatigue behaviour, including sensitivity to segmentation, cycle ranges and low-frequency loading. A direct directional comparison would require the measured moments to be transformed into fore-aft and side-side coordinates using the nacelle heading, which is not available in the measurement dataset.

4.2. Simulation inputs and environmental conditions

Input-file structure

A BHawC simulation combines the turbine definition with the environmental and operational inputs required for a particular design load case. The main simulation file references the structural properties, aerodynamic data, controller settings, wind field, wave conditions, event definitions and requested output channels. The general input and output structure is illustrated in Figure 4.1.

**Figure 4.1:** Main categories of inputs and outputs used in the BHawC simulation workflow.

The simulations are executed in batches covering the required wind speeds, stochastic realisations and yaw conditions. The output files follow a consistent naming convention from which the design load case, wind speed, seed and yaw angle can be identified. The required channels are subsequently extracted from the BHawC result files using an automated post-processing workflow.

Turbulent wind-field generation

The turbulent inflow used in the fatigue simulations is represented using three-dimensional Mann turbulence fields. The Mann model describes the relationship between the longitudinal, lateral and vertical turbulent velocity components through a three-dimensional spectral tensor. It therefore produces a spatially correlated wind field that can be applied across the complete rotor area during the aeroelastic simulation [23]. The target turbulence level is defined using the Normal Turbulence Model for the selected turbine and site conditions. For each mean hub-height wind speed, the model specifies the longitudinal turbulence standard deviation based on site-specific turbulence parameters [4].

Several statistically independent wind fields are generated for each mean wind speed. Although these fields have the same target mean wind speed and turbulence level, different random seeds produce different spatial and temporal arrangements of the turbulent fluctuations. Using several seeds reduces the dependence of the fatigue results on one particular turbulence realisation.

Each turbulence box is generated on a three-dimensional grid, as illustrated in Figure 4.2. The longitudinal x -axis is aligned with the mean wind direction. The lateral y -axis and vertical z -axis span the rotor plane. The corresponding box dimensions are denoted by L_x , L_y and L_z . Under the frozen-turbulence assumption, the wind field is transported through the rotor at the prescribed mean wind speed. For a constant mean advection speed $\bar{U}$, the longitudinal distance required for a box duration T_{box} is approximately

$$L_x = \bar{U}T_{\text{box}}. \quad (4.2)$$




Figure 4.2: Schematic representation of the Mann turbulence box and its coordinate system.

A 3600 s simulation therefore requires a considerably longer turbulence box than a 600 s simulation at the same mean wind speed. The lateral and vertical dimensions are selected to cover the complete rotor area with sufficient margin, as well as having sufficiently high resolution based on the number of lateral and vertical grid points.

The spatial resolution of the rotor-plane grid is checked against the requirements of IEC 61400-1 [4]. For a uniform grid with lateral and vertical spacing Δy and Δz , the diagonal distance across one grid cell is

$$d_{\text{grid}} = \sqrt{(\Delta y)^2 + (\Delta z)^2}. \quad (4.3)$$

The standard requires this distance to satisfy

$$d_{\text{grid}} < 0.25\Lambda_1, \quad d_{\text{grid}} \leq 0.15D, \quad (4.4)$$

where Λ_1 is the longitudinal turbulence scale parameter and D is the rotor diameter. These limits provide a relatively broad resolution requirement and do not define one unique grid configuration.

To avoid generating a separate 3600 s turbulence box for every wind-speed bin, three reference boxes are used following the standard approach. Their reference speeds represent typical rated, cut-out and storm wind conditions.

$$U_{\text{ref}} = \begin{cases} 11 \text{ m/s}, & \bar{U} < 11 \text{ m/s}, \\ 25 \text{ m/s}, & 11 \leq \bar{U} < 25 \text{ m/s}, \\ 52 \text{ m/s}, & 25 \leq \bar{U} < 52 \text{ m/s}. \end{cases} \quad (4.5)$$

Each reference box is generated with a longitudinal length of

$$L_{x,\text{ref}} = U_{\text{ref}}T_{\text{box}}, \quad T_{\text{box}} = 3600 \text{ s}. \quad (4.6)$$

For a simulation at mean wind speed $\bar{U}$, the next highest reference speed is selected. For example, a simulation at 18 m/s uses the box generated at 25 m/s. Since the selected reference speed is always equal to or greater than the simulated wind speed, the box contains at least 3600 s of turbulent inflow when transported at $\bar{U}$.

The same number of grid points is kept for all three reference boxes. Since their physical lengths differ, a width-ratio scaling factor is applied when defining the box dimensions. This keeps the numerical grid size unchanged while adjusting the physical spacing between grid points for each reference speed. This approach greatly reduces the number of turbulence boxes that must be generated and stored. The main drawback is a small reduction in effective temporal resolution when a box is reused at a lower wind speed than its reference value.

The turbulence amplitude is adjusted to the required standard deviation using

$$S_{\text{turb}} = \frac{\sigma_{\text{target}}(\bar{U})}{\sigma_{\text{ref}}}. \quad (4.7)$$

The same factor is applied to the longitudinal, lateral and vertical turbulent velocity components. The target mean wind speed is then added separately to the zero-mean fluctuations. In this way, one reference box can be reused for several wind-speed bins while preserving the required turbulence level and simulation duration.

The main parameters used to generate the Mann fields are summarised in Table 4.3. The turbulence parameters and cross-sectional grid are kept consistent between comparable simulations, while the longitudinal box length changes with the required duration and mean wind speed.

Table 4.3: Mann turbulence-box parameters used in the BHawC simulation campaign.

Parameter	Value
Turbulence-energy parameter, $\alpha\varepsilon^{2/3}$	1.0
Length-scale parameter, L	35.4
Anisotropy parameter, Γ	3.9
Longitudinal grid points, N_x	8192
Lateral grid points, N_y	32
Vertical grid points, N_z	32
Longitudinal box length, L_x	Velocity dependent
Lateral box width, L_y	155 m
Vertical box height, L_z	155 m
Number of independent seeds	90

A representative turbulence-box diagnostic is shown in Figure 4.3. The upper plot shows the streamwise turbulent velocity at the centre of the box as the wind field passes the turbine. The mean wind speed is not included. The vertical slice shows how the wind field varies with both time and height, while the two lower-left plots show cross-sections at different longitudinal positions. Together, these plots illustrate the temporal and spatial variation contained within the generated field.

The power spectrum in the lower-right panel shows how the turbulent energy is distributed across frequency. Most of the energy is contained at low frequencies, corresponding to the largest turbulent structures. Under the frozen-turbulence assumption, the intermediate-frequency range is expected to approximately follow the Kolmogorov $f^{-5/3}$ slope associated with the inertial subrange. The smallest dissipative turbulence scales are not resolved by the grid, and the highest represented frequency is instead limited by the spatial and temporal resolution of the turbulence box. This plot can be inspected for unexpected narrow peaks, loss of frequency content and signs of periodic repetition prior to running the aeroelastic simulations.




Figure 4.3: Representative Mann turbulence-box diagnostic showing the centre-point wind-speed history, a vertical slice through the box, two spatial cross-sections and the corresponding centre-point power spectrum.

Wave conditions

The turbine model represents an offshore installation and includes hydrodynamic loading through the wave and support-structure inputs. When wave loading is active, the wave files define the sea state, wave direction and stochastic realisation used in each simulation.

A preliminary review showed that wave loading was unlikely to have a noticeable influence on the comparisons made in this work. Therefore, existing wave realisations from the original 600 s simulations could be reused. Unfortunately, these files were too short for the 3600 s cases, while generating a new set of wave realisations would have added computational cost without being expected to affect the main results.

The existing wave time series were therefore extended by repeatedly mirroring them until the required duration was reached. This approach was taken since otherwise, direct repetition could create a jump between the end of one segment and the beginning of the next. Mirroring avoids this discontinuity because each new segment starts from the final value of the previous one. The full extended signal is shown in Figure 4.4, and Figure 4.5 shows the continuity at one of the joining points.




Figure 4.4: Extension of the original wave height time series. The initial wave record is shown in orange and the mirrored extension in blue.




Figure 4.5: Detailed view of the transition between the original and mirrored height wave records.

The effect of this mirroring treatment was checked by comparing the wave spectra before and after extension. As shown in Figure 4.6, the extended signal retains the same overall spectral shape and dominant frequency range as the original record. However, the extended spectrum contains stronger local oscillations, which result from the imposed repetition and phase coherence of the mirrored segments, together with the finer frequency resolution of the longer time series. Regardless, no meaningful shift in the overall energy distribution was observed.

This approach does not produce a fully independent 3600 s stochastic wave realisation and introduces repeated structures into the wave history. It was nevertheless considered sufficient for the present study because the preliminary assessment indicated that wave loading had only a minor influence on the investigated response channels. The resulting assumptions should be reconsidered in studies focused specifically on hydrodynamic loading, support-structure response or wave-driven fatigue.




Figure 4.6: Comparison of the wave height spectra before and after extending the time series. The initial wave record is shown in orange and the mirrored extension in blue.

4.3. Simulation campaign

With the turbine model, environmental inputs and fatigue-processing approach established, the BHawC simulation campaign is organised in a series of increasingly complex comparisons. The analysis begins with conventional stationary design load cases, which provide a controlled reference for examining the effect of simulation duration. Several operating states are then combined within a single continuous simulation to investigate cycles that develop across startup and shutdown transitions. Existing 600 s simulations are finally concatenated as a practical diagnostic approximation and compared with directly simulated one-hour cases.

Baseline design load cases

The baseline campaign contains the stationary production, idling and event simulations used throughout the subsequent comparisons. The selected cases are summarised in Table 4.4.

Table 4.4: Baseline BHawC simulation campaign.

DLC	Operating condition	Wind input	Wind speeds [m/s]	Duration	Seeds
1.2	Normal production	NTM	3–31, in 1 m/s steps	600 and 3600 s	6
3.1	Startup	NWP	3, 12 and 25	300 s	1
4.1m	Manual shutdown	NWP	3, 12, 25 and 31	150 s	1
6.4l	Low-wind idling	NTM	1 and 2	600 and 3600 s	6
6.4h	High-wind idling	NTM	32	600 and 3600 s	6
7.2	Idling following loss of grid	NTM	2–32, in 2 m/s steps	600 and 3600 s	6

The analysis begins with DLC 1.2, which represents normal power production and covers the complete operating wind-speed range from 3 to 31 m/s, which includes a High Wind Ride Through (HWRT) region starting at 25 m/s. At each wind speed, the conventional 600 s simulations are compared with directly simulated 3600 s histories. This provides the clearest initial test of whether a longer simulation changes the estimated fatigue rate when the operating condition and nominal wind speed remain unchanged. The comparison is first performed separately for each wind-speed bin. The resulting damage rates are then combined using the

adopted Weibull wind-speed distribution. This provides an overall production fatigue estimate that accounts for both the response at each wind speed and the expected frequency with which that wind speed occurs.

Once the normal-production comparison has been established, the remaining design load cases are included. DLC 6.4 and DLC 7.2 represent different idling conditions, while DLC 3.1 and DLC 4.1m represent startup and manual shutdown respectively. These cases extend the analysis beyond stationary power production and introduce operating conditions with different pitch settings, rotor speeds, controller actions and mean structural loads.

For the combined lifetime comparison, each DLC is first weighted over the relevant wind-speed distribution, as was done previously for normal operation, and then scaled according to its overall occurrence during the turbine lifetime. Continuous conditions such as production and idling are weighted by the fraction of time spent in each state, while startup and shutdown cases are weighted by their expected number of events. Where possible, these state fractions and event frequencies are derived from the measurement campaign rather than taken only from generic design assumptions, making the lifetime estimate more representative of the turbine considered in this investigation.

For normal operation and idling, six turbulence seeds are available in each wind-speed bin. The DEL curves show the lowest and highest values obtained across these seeds, illustrating the variation between individual simulations. The aggregated bar charts instead show the mean value with an uncertainty bar of ± 1 standard error, calculated as the sample standard deviation divided by the square root of the number of seeds. The two measures describe different aspects of the results. The minimum-to-maximum range shows how much the individual simulations vary, while the standard error shows how precisely the mean has been estimated. Uncertainty bars are omitted when no comparable set of seed-level results is available, as is the case for startup or shutdown.

The baseline campaign forms the foundation of the simulation study. It provides a controlled assessment of how simulation duration affects fatigue under stationary conditions and establishes the contribution of each DLC to the overall lifetime estimate. These results then serve as the basis for the following sections, where the analysis is extended to continuous one-hour simulations with changing operating conditions.

Continuous operating-state sequence

The baseline simulation campaign treats each DLC as an independent load history. During real operation, however, the turbine moves continuously between operating states, and the structural response at the end of one state forms the initial condition of the next. Processing each state independently may therefore omit fatigue cycles that begin before a transition and close only after the turbine has entered a different operating condition.

To investigate this effect, a one-hour simulation was constructed by combining idling, startup, normal production and manual shutdown within a single continuous aeroelastic load history. These conditions represent the main progression into and out of power production and involve substantial changes in rotor speed, blade pitch, controller action and structural loading. The sequence was designed to include one complete startup and shutdown while retaining normal production as the dominant operating condition within the one-hour simulation.

The prescribed sequence is summarised in Table 4.5. The simulations were performed for mean wind speeds from 3 to 25 m/s in 1 m/s increments, using six turbulence seeds at each wind speed.

Table 4.5: Operating-state sequence used in the combined one-hour simulations.

Time [s]	DLC	Operating condition	Duration [s]
0–200	7.2	Idling following loss of grid	200
200–320	3.1	Startup	120
320–3500	1.2	Normal power production	3180
3500–3600	4.1m	Manual shutdown	100

For context, Table 4.6 compares the time shares prescribed in the combined sequence from Table 4.5, with the normalised DLC weights used in the conventional fatigue limit state calculation, and with the operating-state distribution obtained from the measurements in Figure A.7. These distributions are not directly equivalent because they were developed for different purposes and use different classifications. Nevertheless, the combined sequence was designed to approximate the conventional and measurement-based shares as closely as possible within a one-hour simulation, while accounting for the fixed durations required for startup and shutdown.

Table 4.6: Comparison of DLC weights used in the conventional fatigue limit state calculation, the operating-state shares from the measurements, and the time shares prescribed in the combined one-hour simulations.

DLC	Operating condition	Conventional DLC weight [%]	Measurement share [%]	Combined-sequence share [%]
1.2	Normal power production	93.33	82.19	88.33
3.1	Startup	1.96		3.33
4.1m	Manual shutdown	2.00	16.89	2.78
6.4h	High-wind idling	0.01		–
6.4l	Low-wind idling	2.37		–
7.2	Idling following loss of grid	0.33	0.92	5.56
Total		100.00	100.00	100.00

The measurement-based 16.89% idling share includes startup, shutdown and other non-producing periods because of the simplified classification in section A.3, which only distinguishes between normal production, idling and fault conditions. It therefore cannot be divided directly among DLC 3.1, DLC 4.1m and DLC 6.4.

In addition, for the combined sequence the simulated wind-speed range of 3–25 m/s does not include the below-cut-in and above-cut-out conditions represented by DLC 6.4l and DLC 6.4h. DLC 7.2 was therefore used as a practical representation of the initial idling segment. Its prescribed share should not be interpreted as solely the expected lifetime occurrence of loss-of-grid events, but should also represent normal idling conditions.

The resulting load history for the combined sequence simulation was evaluated in two ways. First, the complete 3600 s signal was rainflow counted as one continuous record. Then, the four operating portions were separated, counted independently, and their fatigue contributions were combined. Since both calculations use the same underlying aeroelastic response, any difference between them is caused by cycles that extend across the operating-state boundaries. This comparison therefore isolates the fatigue contribution that would be omitted when the individual DLCs are processed independently.

4.4. Concatenation of 600 s simulations

Purpose and interpretation

Many existing simulation databases contain only conventional 600 s results. Repeating the full aeroelastic campaign with 3600 s simulations would require substantial computational effort and storage, particularly when several wind speeds, seeds and operating conditions are considered. Concatenating existing files therefore offers a practical way to investigate whether longer load histories can be approximated using simulations that are already available. This can be visually interpreted as follows:




Figure 4.7: Concatenation of individual 600s simulations concept.

The main purpose is to assess whether this approach can recover some of the fatigue effects associated with longer records, especially cycles that may form across the boundaries between individual 600 s segments. The concatenated histories are compared with directly simulated one-hour seeds to determine how closely they reproduce the resulting cycle distributions and DELs.

However, it is worth noting that the resulting signal is not equivalent to a physically continuous one-hour aeroelastic simulation. Each source file begins from an independently initialised turbine state and turbulent wind field. Structural displacement, controller memory, rotor azimuth and inflow phase are therefore not continuous across the joins. Concatenation is consequently treated as a post-processing sensitivity study rather than as a complete representation of continuous turbine operation.

To test the validity of this approach, six independent 600 s DLC 1.2 simulations are selected at each wind speed from 3–31 m/s to form an approximately 3600 s history. All source files represent normal production at the same nominal wind speed, but use different turbulent realisations. Duplicate files are excluded, and each segment is used only once within a sequence to avoid artificial repetition of the same load history.

The segments are then joined using the boundary-matching procedure described below. For

each pair of simulations, short windows near the end of the preceding segment and the beginning of the following segment are compared. The joining locations are selected from within these overlapping windows so that the change in the load signal and relevant operational channels is minimised.

Allowing a small part of each 600 s record to overlap provides flexibility in the exact position of the join and helps avoid sharp discontinuities. The objective is not to reproduce full physical continuity, since the source simulations remain independently initialised, but to construct the smoothest possible combined load history and reduce the risk of artificial cycles being introduced at the segment boundaries. The cases included in the concatenation study are summarised in Table 4.7.

Table 4.7: Simulation cases used in the concatenation study.

Case	DLC	Operating condition	Wind speeds [m/s]	Source duration	Final duration	Realisations
Concatenated signal	1.2	Normal production	3–31, in 1 m/s steps	600 s	$\approx$ 3600 s	6 segments
One-hour reference	1.2	Normal production	3–31, in 1 m/s steps	3600 s	3600 s	6

To ensure a consistent comparison, the independently processed segments, the concatenated histories and the directly simulated 3600 s cases are evaluated using the same rainflow implementation, Wöhler exponent, reference cycle count and DEL normalisation. Differences between the results can therefore be attributed to the treatment and duration of the load history rather than to changes in the fatigue-processing settings.

Boundary matching

The boundary-matching procedure determines where each 600 s segment should end and where the next segment should begin. Rather than joining the files exactly at their original endpoints, a short search window is defined at the end of the preceding simulation and another at the beginning of the following simulation. Different pairs of cut points within these two windows are then tested.

For each candidate pair, the values immediately before and after the proposed join are compared. The comparison includes the bending moment being analysed, together with wind speed, rotor speed and blade pitch. The bending moment indicates whether the structural response remains smooth, while the other channels indicate whether the turbine is in a similar operating condition on both sides of the boundary. These differences are combined into a single mismatch score, with weighting factors controlling the influence of the operational channels. The pair of cut points with the lowest score is selected. The preceding segment is retained up to its selected cut point, and the following segment is appended from its selected cut point onwards.

Two aspects of the procedure must be selected: the length of the search windows and the weighting assigned to the operational channels. These settings are calibrated through a bounded parameter study, described below.

The calibration operates at two distinct levels, which should not be confused with each other. At each individual boundary, the algorithm always selects the cut points that produce the smallest local mismatch, as described above, regardless of which search-window length and channel weights are currently being tested. Separately, across the full parameter study, different combinations of search-window length and channel weights are compared against each other, to find the single combination that performs best overall. This second, higher-level comparison requires its own criterion, which is defined below. Once the best combination is identified, the

same final settings are used for every remaining sequence, and are not adjusted separately for each wind speed or realisation.

To evaluate a given combination of settings, all six 600 s segments are joined for each calibration wind speed using that combination. The resulting concatenated history is rainflow counted, and its DEL is compared with the DEL from a directly simulated one-hour case at the same wind speed. This gives one measure of agreement per wind speed. The overall agreement for a given combination of settings is quantified using

$$\Phi = \frac{1}{N_V} \sum_{i=1}^{N_V} \left| \frac{\text{DEL}_{\text{concat}}(V_i)}{\text{DEL}_{1\text{hour,ref}}(V_i)} - 1 \right|, \quad (4.8)$$

where N_V is the number of wind speeds used in the calibration.

To see what this means in practice, consider one calibration wind speed V_i . $\text{DEL}_{\text{concat}}(V_i)$ is the DEL obtained from the six concatenated 600 s segments at that wind speed, and $\text{DEL}_{1\text{hour,ref}}(V_i)$ is the DEL from the directly simulated one-hour reference case at the same wind speed. Dividing one by the other and subtracting 1 gives the relative error between the two, expressed as a fraction. For example, a value of 0.05 means the concatenated result is 5% higher or lower than the reference. The absolute value is taken so that overestimates and underestimates count equally, rather than cancelling each other out when averaged.

This relative error is calculated separately for every calibration wind speed, and Φ is the average of these errors across all of them. A low Φ therefore means that, on average, the concatenated one-hour DELs are close to the directly simulated one-hour DELs across the full range of calibration conditions, not just at a single wind speed. Φ is exactly the criterion used to compare candidate combinations of search-window length and channel weights: each combination is scored using Φ , and the combination with the lowest Φ is selected as the final setting.

If a candidate combination fails to produce a valid concatenated history for one or more of the required wind speeds, for example because no acceptable cut point is found within the search window, that combination is assigned a fixed penalty value instead of a Φ score. This ensures that incomplete or failed combinations are never mistaken for good ones simply because they were never fully evaluated.

4.5. Methodological limitations

The BHawC model is a proprietary industrial turbine model. The structural distributions, aerodynamic coefficients and detailed controller logic cannot therefore be reported in full. The model description is limited to the information required to understand the simulation campaign and fatigue comparisons.

The study also uses a finite number of turbulent realisations. Seed averaging reduces the influence of individual fields but does not remove stochastic uncertainty completely. Differences between the 600 and 3600 s cases must therefore be interpreted together with the observed seed-to-seed variation. Even longer simulations would provide a stronger basis for assessing low-frequency fatigue behaviour and convergence with record duration, but the associated increase in computational cost and data volume was not feasible within the scope of this study. The 3600 s cases therefore represent a practical compromise between improved temporal coverage and available computational resources.

Furthermore, another limitation is that the conventional Mann fields reproduce the prescribed turbulence statistics but do not represent every form of long-term atmospheric variability. In particular, the baseline cases remain stationary around a prescribed mean wind speed and do not reproduce mesoscale changes in inflow. The results should therefore not be interpreted as a universal correction for all wind conditions.

Moreover, the combined-DLC simulations preserve the turbine response across a prescribed sequence of operating states, allowing fatigue cycles that develop during startup, production and shutdown to be retained. However, the selected event timing, sequence and occurrence frequencies remain modelling assumptions. Their contribution to a lifetime estimate therefore depends on how often similar transitions occur in practice and how representative the prescribed sequence is of actual turbine operation.

The concatenated histories have an additional limitation because they join independently initialised simulations. Boundary matching reduces visible discontinuities but cannot recreate continuous turbulence, rotor phase, structural motion or controller memory. These histories are therefore used as diagnostic approximations rather than substitutes for continuous aeroelastic simulations.

Finally, the lifetime results depend on assumptions regarding the selected Weibull wind-speed distribution, the fraction of time spent in each operating state and the annual number of startup and shutdown events. These inputs affect the absolute magnitude of the estimated lifetime fatigue loads. The main conclusions are therefore based primarily on relative comparisons performed using identical assumptions. This allows the effects of duration, segmentation and operating-state continuity to be separated from uncertainty in the absolute lifetime model.

Measurement-based fatigue results

This chapter presents the results obtained from the measured turbine data using the methodology introduced in chapter 3. The analysis begins by illustrating how the load response changes when it is viewed over different time scales. Next, the custom four-point rainflow implementation and residue-aware workflow are validated against continuous rainflow counting. With this, the effect of independent 10-minute processing is examined in detail for the blade root flapwise bending moment before the comparison is extended to the blade edgewise and tower-bottom channels. Finally, the recovered cycles are linked to the turbine operating state to identify the physical origin of the missing fatigue contribution. Together, these results establish which cases should be prioritised in the subsequent BHawC investigation.

5.1. Measured load behaviour across different time scales

The segmentation problem is first illustrated using the blade root flapwise bending moment. Figure 5.1 shows a 30-day section of the measured signal. The response varies considerably over the selected period, both in its oscillation amplitude and in its mean load level. These changes occur over time scales ranging from several hours to several days.




Figure 5.1: Thirty-day record of the blade root flapwise bending moment.

The slow evolution of the response is highlighted in Figure 5.2 using a low-pass-filtered curve. This makes the gradual changes in the mean load level more visible and shows that the flap-wise response cannot be assumed to be stationary over the selected period. Note that the filtered curve is included only as a visual aid. As described in section 3.2, no trend or low-frequency content is removed from the signal before rainflow counting.




Figure 5.2: Thirty-day flapwise bending-moment record with the slow variation highlighted using a low-pass-filtered curve.

In contrast, Figure 5.3 shows only the first 10 minutes of the same record. At this scale, the response is dominated by repeated cycles associated with blade rotation, with a typical peak-to-peak range of approximately 5000 kNm. The signal appears to oscillate around a comparatively stable local mean, and the longer variation visible over the 30-day period is no longer evident, where variations of more than 20000 kNm could be seen.




Figure 5.3: First 10-minute block of the flapwise record, showing the shorter cycles associated with blade rotation.

This comparison demonstrates why a 10-minute block cannot describe the complete load history. Each block contains many short cycles, but only a small part of the slower change in structural loading. A fatigue cycle may begin during one block and close only after several later blocks. If the blocks are counted independently, the turning points on either side of the boundaries cannot form one continuous rainflow cycle.

5.2. Validation of the rainflow and residue-recovery workflow

Before examining the effect of 10-minute segmentation, the rainflow and residue-recovery procedures were validated. For each selected time window, the same blade flapwise load history was processed in three ways:

1. the complete window was counted using the established `FLS.rainflow` implementation, providing the continuous reference, DEL_{REF} ,
2. the same complete window was counted using the custom RF4PT implementation, DEL_{RF4PT} ,
3. the window was divided into 10-minute blocks and processed using:
 - (a) conventional independent counting, DEL_{STA} ,
 - (b) residue-aware LFFD processing, DEL_{LFFD} .

This comparison answers two separate questions. The RF4PT result shows whether the custom rainflow algorithm reproduces the established reference count. The LFFD result shows whether processing the signal in 10-minute blocks, while keeping track of the ordered residues at each boundary, can still recover the same result as counting the whole window in one pass.

Table 5.1 presents a representative validation sequence for the blade root flapwise channel using $m = 10$, with each row showing the average ratio across 20 randomly selected windows of that length. Each value is divided by the continuous `FLS.rainflow` result for the same window, so a ratio of one indicates exact agreement with the reference.

Table 5.1: Validation ratios for the blade root flapwise channel at $m = 10$.

Window length	Continuous RF4PT DEL_{RF4PT}/DEL_{REF}	Independent 10 min DEL_{STA}/DEL_{REF}	Residue-aware LFFD DEL_{LFFD}/DEL_{REF}
10 min	1.000	1.000	1.000
1 hour	1.000	0.941	1.000
1 day	1.000	0.900	1.000
1 week	1.000	0.844	1.000
2 weeks	1.000	0.810	1.000
1 month	1.000	0.825	1.000

Each row corresponds to a different, independently selected group of windows, so the table should not be read as a single convergence curve. The ratio is not expected to decrease smoothly and monotonically as the window gets longer, since every window contains a different sequence of load excursions. Three trends stand out in the table, one per column.

First, the continuous RF4PT ratio is 1.000 for every window length, and the same agreement was found for all the randomly selected validation windows tested. This confirms that the custom rainflow implementation reproduces the established `FLS.rainflow` reference exactly, so any differences seen elsewhere in this investigation are not caused by the counting algorithm itself.

Second, the residue-aware LFFD ratio is also 1.000 throughout. Even though the signal is split into 10-minute blocks and processed separately, keeping the ordered residues from each block and recombining them recovers exactly the same result as counting the full window in one go, to the specified degree of accuracy. This confirms that the LFFD approach is valid, and specifically that the block-level storage procedure keeps all the information needed to identify cycles that cross the original 10-minute boundaries.

Third, the independent 10-minute ratio matches the reference only for the 10-minute window itself, where the window and the block are identical. For every longer window, this ratio falls clearly below one. For example, at one month the ratio is 0.825, meaning the conventional independent-block DEL is 17.5% below the continuous reference for that interval. The exact ratio varies between window lengths, again because each one contains a different sequence of large load excursions, but the underestimation is consistent across all of them. This is the central result of the validation: independent 10-minute processing systematically misses fatigue contributions that are present in the continuous signal, producing a non-conservative estimate. The corresponding damage-sum ratios are far more extreme than the DEL ratios shown here, since DEL is obtained by taking the m th root of the normalised fatigue sum. For $m = 10$, a moderate difference in DEL therefore corresponds to a much larger difference in damage sum.

Beyond accuracy, the choice of processing method also has a practical cost in computation time. Table 5.2 reports the time required to process the same validation windows using each of the four methods, for the blade root flapwise channel.

Table 5.2: Indicative computation times for the blade root flapwise channel at $m = 10$.

Window length	REF [s]	STA [s]	RF4PT [s]	LFFD [s]
10 min	0.004	0.002	0.007	0.008
1 hour	0.005	0.007	0.029	0.038
1 day	0.049	0.077	0.125	0.259
1 week	0.346	0.511	1.459	2.792
2 weeks	0.801	0.913	4.229	5.779
1 month	1.435	1.969	5.236	10.769

The computation times in Table 5.2 exclude data loading, resampling, figure generation and file writing, and should be read as indicative implementation timings rather than hardware-independent benchmarks.

REF is the fastest method at virtually every window length, even compared with STA. This is despite STA using the same underlying algorithm as REF but with the additional step of splitting the signal into 10-minute blocks before counting, which shows that this splitting overhead is small relative to the counting itself. RF4PT is consistently slower than REF for the same task. This is expected, since REF is a heavily optimised, built-in function from the MATLAB Signal Processing Toolbox rather than a custom implementation.

LFFD is noticeably slower than the other three methods, but this comparison should be made against RF4PT specifically, since the two share the same underlying counting algorithm. Against RF4PT, LFFD is approximately twice as slow at the longest duration tested. This is consistent with LFFD requiring a second rainflow counting pass on the combined ordered residues, in addition to the block-level counting already performed by RF4PT.

As expected, LFFD is slower than simply summing independently. However, the block-level preprocessing only needs to be done once, and the stored residues can then be reused for many different analysis windows. This makes the method practical for long monitoring records without needing to reload the complete high-frequency history each time. This reuse also matters for memory, not just speed. Direct continuous counting requires the entire high-frequency signal to be held in memory at once, which in the present implementation limited continuous processing to periods of roughly one month, as longer intervals exceeded the available memory. LFFD avoids this limitation, since it only stores the closed-cycle contributions and ordered residues from each 10-minute block rather than the full signal.

Since LFFD reproduces the continuous reference within the validated range and is the only practical way to process the full six-month measurement record, it is used for all subsequent measurement results. On the contrary, for the BHawC results presented in chapter 6, the reference single-pass rainflow counting algorithm is used instead, since the one-hour simulations can still be processed directly as continuous histories.

5.3. Blade root flapwise bending moment

The blade root flapwise response is examined first because it contains substantial slow variation and is strongly influenced by aerodynamic loading. Composite blade materials are also commonly represented using relatively high Wöhler exponents, making the fatigue result particularly sensitive to the largest recovered cycle ranges.

5.3.1. Development of the segmentation effect

Figure 5.4 shows the flapwise cumulative DEL ratio over the six-month measurement period for $m = 10$. At each point, the conventional and residue-aware DELs are recalculated using all data available from the beginning of the record up to that time.




Figure 5.4: Cumulative DEL ratio $\text{DEL}_{\text{STA}}/\text{DEL}_{\text{LFFD}}$ for the blade root flapwise channel at $m = 10$. Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.

The first point is equal to one, as expected. For a single 10-minute window, both methods process the same isolated load record and LFFD has no neighbouring residues to combine. It therefore gives the same result as standard independent processing.

For each analysis duration, several windows of equal length are taken from different positions within the measurement record. These windows are allowed to overlap, which means that multiple samples remain available even for durations approaching the full record length. The shaded band thus shows plus or minus one standard error,

$$SE = \frac{s_R}{\sqrt{n}}, \quad (5.1)$$

where s_R is the sample standard deviation of the window ratios and n is the number of windows.

The band is wide at short durations because the result depends strongly on which operating events and large cycles are included in each window. As the duration increases, the windows contain a larger and more similar portion of the record, causing their ratios to converge and the band to narrow.

By the end of the six-month period, the DEL ratio is approximately 0.84. Conventional, independent 10-minute processing therefore produces an equivalent load approximately 16% below the residue-aware result. This represents a substantial and non-conservative difference in fatigue estimation for the flapwise response. The analysis extends the trend shown in the middle column of Table 5.1 by evaluating the ratio at daily increments over the complete six-month record, rather than only at selected window lengths.

A similar cumulative behaviour has been reported for offshore wind turbine support-structure measurements. Sadeghi et al. and Marsh et al. also observed considerable variation over short observation periods, followed by stabilisation as progressively more data were included [11, 12]. Their analyses considered damage ratios, steel Wöhler exponents and different structural responses, so the numerical values and convergence periods are not directly comparable with the present blade result. However, the similar development supports the interpretation that short observation periods can give unstable estimates of the segmentation effect, and that the magnitude of these destabilizing effects reduces over time.

Importantly, the ratio does not gradually return towards one. Instead, it remains below one and approaches a comparatively stable plateau. Within the available dataset, the difference is therefore a systematic consequence of independent 10-minute processing rather than an effect caused only by one unusual event near the beginning of the record.

This apparent convergence should not be interpreted as evidence that six months are sufficient for a complete lifetime assessment. As discussed in section 2.6, stabilisation of the processing ratio does not guarantee that the record contains every relevant season, wind direction or operating condition. It does, however, show that the difference between the two methods persists as the observation period increases. This supports the conclusion that the flapwise response contains fatigue-relevant load variations that develop over periods longer than 10 minutes and are therefore not fully captured when each block is processed independently.

5.3.2. Sensitivity to the Wöhler exponent

Figure 5.5 extends the cumulative comparison to Wöhler exponents of $m = 11$ and $m = 12$, allowing the influence of the assumed fatigue sensitivity to be examined.




Figure 5.5: Cumulative DEL ratio $\text{DEL}_{\text{STA}}/\text{DEL}_{\text{LFFD}}$ for the blade root flapwise channel at $m = 10, 11$ and 12 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.

For all three exponents, the ratio remains below one and approaches a comparatively stable level as the observation period increases. The final DEL ratio decreases from approximately 0.84 at $m = 10$ to around 0.80 at $m = 12$. Conventional 10-minute processing therefore underestimates the residue-aware DEL by approximately 16% to 20%, depending on the selected Wöhler exponent.

The increasing difference with m indicates that the recovered cycles are located mainly at the upper end of the cycle-range distribution. A larger Wöhler exponent gives these high-range cycles greater importance relative to the more numerous smaller cycles that already close within each 10-minute block. It is likely then, that the segmentation effect is controlled by a limited number of large, slow-moving load excursions rather than by a substantial change in the overall cycle count.

The change in DEL ratio between the three exponents appears relatively modest because the DEL calculation includes the m -th root of the accumulated fatigue sum. The underlying difference in pure fatigue damage contribution is considerably larger. This shows that the assumed material sensitivity is important when evaluating the consequences of segmentation, particularly for composite blade structures where high Wöhler exponents are commonly applied.

5.3.3. Contribution of different aggregation scales

The previous results show that important flapwise cycles extend beyond a single 10-minute block. The next logical step is to determine how long the residues must be retained before most of this missing fatigue is recovered.

For this comparison, the six-month record is divided into groups of increasing duration. At the 10-minute level, every block is processed independently and no residue is passed across a boundary. At the one-day level, the residues from all 10-minute blocks within each day are

combined, but cycles are not allowed to cross from one day into the next. The same procedure is repeated using weekly and monthly groups. The result from processing the complete six-month residue as one continuous sequence is used as the reference. Visually, this can be represented as follows:




Figure 5.6: Schematic of splitting of record depending on aggregation scales.

Figure 5.7 shows, for each aggregation period ranging from 10 minutes to 6 months, the fatigue sum divided by the full six-month reference. A value of 1 therefore means that the selected aggregation period recovers the same fatigue contribution as allowing residues to interact across the entire measurement period. A lower value means that some cycles remain separated by the boundaries between the aggregation groups.




Figure 5.7: Fraction of the six-month flapwise DEL recovered using different residue-aggregation periods for $m = 10, 11$ and 12 .

The 10-minute values agree with the final cumulative ratios in Figure 5.5, as expected, because both represent 10 minute independent processing of the complete record. Depending on the Wöhler exponent, this approach recovers approximately 80–84% of the full-record LFFD DEL, as was shown previously.

Allowing residues to interact within each day increases the recovered share to approximately 95–96%, while weekly aggregation recovers about 99%. This shows that most of the fatigue contribution missed by independent 10-minute processing comes from cycles that cross one or more 10-minute boundaries but close within a day. Only a small additional contribution is associated with cycles extending over several days, and an even lower proportion spans several weeks.

Monthly aggregation reproduces almost perfectly the complete six-month result, and longer periods provide little further improvement. The figure therefore shows that retaining residues for about one week is sufficient to recover nearly all of the missing flapwise DEL in this dataset, while monthly aggregation gives an almost exact match.

5.3.4. Cycle-range distribution

The aggregation analysis shows that the missing flapwise contribution is formed by cycles that require more than 10 minutes to close. To understand why these relatively slow cycles have such a large influence on the fatigue result, their ranges must also be examined. Figure 5.8 compares the cycle-count distributions obtained from conventional independent processing and residue-aware LFFD processing.




Figure 5.8: Cycle-count distribution for the blade root flapwise channel before and after residue recovery.

The two distributions are nearly identical across the low and moderate range bins. This shows that conventional 10-minute processing captures the large population of shorter cycles that close within each block, including the repeated loading associated with blade rotation. Residue recovery therefore does not substantially change the main body of the cycle distribution.

The difference appears mainly at the upper end of the range spectrum. LFFD processing recovers numerous cycles above approximately 15,000 kNm that are absent or strongly underrepresented when the blocks are counted independently. These cycles correspond to the slow load excursions identified in the previous section. Their turning points occur in different 10-minute blocks and can only form complete cycles once the residues are combined over longer periods.

The recovered high-range cycles are rare compared with the shorter cycles. Their counts are several orders of magnitude lower than those in the main part of the distribution. However, their fatigue contribution is amplified by the Wöhler relation,

$$N_i = \left(\frac{S_0}{S_i} \right)^m, \quad (5.2)$$

as described in section 2.2. For example for $m = 10$, increasing the load range from $S_1 = 4,000$ kNm to $S_2 = 20,000$ kNm, which corresponds to the values of some of the recovered cycles, gives

$$\frac{N_2}{N_1} = \left(\frac{S_1}{S_2} \right)^{10} = \frac{1}{5^{10}} \approx 1.0 \times 10^{-7}. \quad (5.3)$$

The fivefold increase in load range therefore reduces the number of cycles to failure by approximately 10 million. Therefore, only a limited number of recovered cycles is required to produce the substantial DEL difference observed in subsection 5.3.1.

The cycle distribution brings the previous findings together. Independent processing captures most of the cycles by number, but misses part of the small high-range tail that develops over periods longer than 10 minutes. The segmentation effect is therefore not caused by a general increase in cycle counts, but by the loss of a small number of slow, large-range cycles to which the fatigue calculation is particularly sensitive.

5.4. Comparison between structural response channels

The flapwise results show that independent 10-minute processing can miss a small number of large cycles and produce a substantial reduction in DEL. The same analysis is now applied to the blade root edgewise moment and the two tower-bottom bending moments to determine whether this behaviour is common to all structural responses.

The comparison reveals a strong dependence on the physical loading mechanism of each channel. Responses that follow slow changes in aerodynamic loading are affected considerably, whereas responses dominated by cycles that close repeatedly within each rotor revolution show little sensitivity to segmentation.

5.4.1. Blade root edgewise bending moment

Figure 5.9 shows the cumulative DEL ratio for the blade root edgewise channel at $m = 10, 11$ and 12.




Figure 5.9: Cumulative DEL ratio $\text{DEL}_{\text{STA}}/\text{DEL}_{\text{LFFD}}$ for the blade root edgewise channel at $m = 10, 11$ and 12. Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.

For all three Wöhler exponents, the ratio remains above approximately 0.998 throughout the measurement period. Conventional processing therefore produces an edgewise DEL less than 0.2% below the residue-aware result. This difference is negligible compared with the 16–20% underestimation observed for the flapwise channel.

The corresponding cycle-range distributions are shown in Figure 5.10.




Figure 5.10: Cycle-count distribution for the blade root edgewise channel before and after residue recovery.

The two distributions are relatively similar across the complete range spectrum. A redistribution is nevertheless visible in several moderate-range bins. When each 10-minute block is processed independently, unresolved turning points at the block boundaries are treated only within that isolated segment. After the residues from consecutive blocks are combined, some of these turning points become part of larger cycles extending across several blocks. As a result, cycles that previously appeared as separate moderate-range excursions are replaced by fewer cycles at larger ranges.

This explains why the residue-aware distribution contains lower counts in some intermediate bins and higher counts in others. However, the redistribution does not create the very pronounced high-range tail observed for the flapwise channel. The changes in cycle range therefore largely balance at the fatigue level, resulting in only a negligible difference in the final edgewise DEL.

The limited segmentation effect is consistent with the physical origin of the edgewise response. Edgewise blade loading is strongly influenced by gravity as the blade rotates. The resulting cycles repeat once per revolution and close many times within each 10-minute block. In contrast to the flapwise response, the dominant edgewise cycle ranges are therefore already captured by conventional processing.

The edgewise channel provides an important comparison case. It shows that residue propaga-

tion does not automatically increase the fatigue estimate. A substantial correction occurs only when the response contains damaging cycles whose turning points are separated by more than the selected processing interval.

5.4.2. Tower-bottom bending moments

The tower-bottom measurements are recorded in the fixed geographical coordinate system introduced in section 3.1, following conventions as strain sensors are typically placed at these positions in the turbine. The Y axis is aligned with true North, while the X axis points East. The coordinate convention, as well as the wind climate at the site location, are shown in Figure 5.11.



Figure 5.11: Tower coordinate convention and site wind climate.

The prevailing wind direction is approximately south-westerly, giving the aerodynamic thrust vector substantial components in both fixed directions. Similar fatigue behaviour in the X and Y channels is therefore to be expected. However, their exact responses may still differ because the load projection changes with wind direction, nacelle orientation and the structural characteristics of the turbine and support structure.

Figures 5.12 and 5.13 show the cumulative DEL ratios for the X and Y tower directions respectively. Since the tower is made of steel, the Wöhler exponents are now much lower than before, now ranging from $m = 3 - 5$.




Figure 5.12: Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the tower-bottom X direction at $m = 3, 4$ and 5 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.




Figure 5.13: Cumulative DEL ratio DEL_{STA}/DEL_{LFFD} for the tower-bottom Y direction at $m = 3, 4$ and 5 . Shaded bands indicate plus or minus one standard error across overlapping windows of the same duration.

The two directions show very similar behaviour. By the end of the six-month period, conven-

tional processing produces a DEL approximately 3% below the residue-aware result at $m = 3$, 9% below it at $m = 4$, and 20% below it at $m = 5$. The 3% difference is small and may be comparable to sampling uncertainty in the available six-month record, whereas the larger differences at $m = 4$ and $m = 5$ provide clearer evidence of sensitivity to cycles extending beyond the 10-minute interval. As for the flapwise channel, the difference increases with the Wöhler exponent, indicating that the recovered cycles are concentrated at the larger end of the range distribution. At $m = 3$, these cycles have a relatively limited effect on the equivalent load. At $m = 5$, their increased weighting makes the segmentation error comparable with that observed for the blade flapwise response.

For the lower Wöhler exponents, the ratios approach a nearly constant value relatively quickly. This is because the fatigue contribution is distributed more evenly across the cycle population, so the result is less sensitive to the occasional appearance of very large cycles. As m increases, the largest ranges receive progressively greater weight and the cumulative ratio becomes more dependent on a smaller number of infrequent events. The $m = 5$ ratios therefore continue to vary near the end of the measurement period. This indicates that the result has not stabilised to the same extent as for $m = 3$ and $m = 4$, and remains influenced by relatively rare large excursions. These may be associated with multi-day changes in wind speed, shifts in wind direction or transitions between turbine operating states. A longer record would be required to determine whether the $m = 5$ ratio eventually converges to a stable value over a complete seasonal cycle.

As before, the associated cycle-range distributions are shown in Figures 5.14 and 5.15.




Figure 5.14: Cycle-count distribution for the tower-bottom X direction before and after residue recovery.




Figure 5.15: Cycle-count distribution for the tower-bottom Y direction before and after residue recovery.

For both tower directions, the two distributions overlap closely at the lower cycle ranges, where most of the counted cycles are located. The difference becomes increasingly visible towards the upper part of the spectrum. Residue-aware processing produces higher counts across several moderate-to-high range bins and extends the distribution beyond the maximum ranges identified by conventional processing. In the Y direction, for example, conventional counting contains almost no cycles above approximately 90,000 kNm, whereas the LFFD result extends to around 125,000 kNm.

The effect is therefore not limited to the recovery of one isolated extreme cycle. Combining the ordered residues changes how several unresolved excursions are paired, shifting part of the cycle population towards larger ranges and creating an extended high-range tail. Although the recovered cycles have very low counts compared with the dominant low-range population, their contribution to fatigue becomes important because of their much larger ranges. This explains why the tower DEL changes substantially, particularly at $m = 5$, despite the strong overlap between the distributions at lower ranges.

The similar behaviour of the X and Y channels shows that the segmentation effect is present in both fixed geographical directions and to a similar degree, as was to be expected from the wind rose distribution shown in Figure 5.11b. It would also have been valuable to examine the tower response in the fore-aft and side-side coordinate system, since this would separate loading aligned with and perpendicular to the rotor axis. However, as discussed in section 3.1, the nacelle yaw position was not available for the measurement period, so the required coordinate transformation could not be performed. The fore-aft and side-side responses are instead examined using the aeroelastic simulations in the simulation-results chapter.

Summary of channel dependence

The approximate DEL ratios obtained after six months are summarised in Table 5.3.

Table 5.3: Approximate six-month DEL ratios for the measured load channels.

Channel	Wöhler exponent	DEL_{STA}/DEL_{LFFD}	Observed effect
Blade flapwise	$m = 10-12$	0.80–0.84	Large
Blade edgewise	$m = 10-12$	0.998–0.999	Negligible
Tower X	$m = 3-5$	0.80–0.97	Small to large
Tower Y	$m = 3-5$	0.80–0.97	Small to large

The results confirm that the influence of 10-minute segmentation cannot be represented by one correction factor for the complete turbine. The blade edgewise response is almost unaffected because its dominant gravity-driven cycles close repeatedly within each block. In contrast, blade flapwise and tower bending are sensitive to slower changes in aerodynamic loading and turbine operation, allowing large cycles to develop across several block boundaries.

The Wöhler exponent further determines how strongly these recovered cycles affect the final DEL, so choosing the correct one is vital when estimating fatigue damage. A response may contain only a small number of additional high-range cycles, but their importance increases rapidly when a larger value of m is applied. The need for residue-aware processing must therefore be assessed using both the physical behaviour of the load channel and the fatigue properties of the material of the component being evaluated.

5.5. Operational origin of the recovered cycles

The preceding sections show that the missing fatigue contribution is concentrated in large cycles that extend across multiple 10-minute blocks. The next question is whether these cycles develop during one operating regime or while the turbine moves between different regimes. The SCADA classification introduced in section 3.7 is therefore combined with the flapwise rainflow results to identify the operating conditions associated with the recovered cycles. This channel is chosen, with an assumed $m = 10$, as it shows a clear segmentation effect and is strongly influenced by changes in aerodynamic loading and turbine operation. The procedure taken to assign each time block to a certain operational regime can be found in section A.3.

5.5.1. Fatigue attribution by operating regime

Figure 5.16 compares the flapwise fatigue contributions of the different operating regimes before and after residue recovery. This complements Figure A.7, which shows only the proportion of time spent in each regime and does not indicate its contribution to the total fatigue damage. The shares are normalised separately for each processing method and therefore sum to 100%.




Figure 5.16: Fatigue share associated with each operating regime for the blade root flapwise channel at $m = 10$, before and after residue recovery.

Under conventional processing, approximately 88% of the fatigue sum is assigned to Operational intervals and around 12% to Idling, with Fault accounting for a negligible amount. This broadly follows the time distribution the turbine spends in each regime, because the cycles are counted separately inside each 10-minute block and are assigned to the state of that block.

However, after residue recovery, the Operational share decreases to approximately 69%, while the share associated with Idling rises to around 30%. The Fault contribution remains small, which is consistent with the very limited number of intervals classified in this state.

The larger share attributed to Idling does not mean that steady idling is suddenly more damaging than normal production. Instead, it can be interpreted in the following way. Many of the recovered cycles connect a low-load idling period with a higher-load Operational period, for example during startup or shutdown. When these load levels occur in different 10-minute blocks, conventional processing treats them separately, while residue-aware processing combines them into one transition cycle. The increased Idling share therefore reflects fatigue associated with changes between operating states rather than additional loading during steady idling.

5.5.2. Damage matrices and transition cycles

The operating origin of the recovered cycles can be shown more directly by the damage matrices in Figure 5.17. Each matrix entry represents the fraction of the fatigue sum associated with a cycle whose two turning points occur in the indicated operating regimes. Diagonal entries correspond to cycles that begin and end within the same regime. Off-diagonal entries correspond to cycles connecting two different regimes. Since the matrices are folded, transitions in opposite directions are combined into one entry.




Figure 5.17: Folded operating-regime damage matrices for the blade root flapwise channel at $m = 10$. Left: conventional blockwise processing. Right: residue-aware processing.

The conventional matrix is diagonal by construction. Every 10-minute block receives one operating-state label and is counted independently. Consequently, every cycle found within that block is recorded as beginning and ending in the same state. This procedure cannot represent a cycle whose turning points occur on opposite sides of an operating-state transition.

The residue-aware matrix is markedly different. Only approximately 39% of the fatigue remains on the diagonal, while around 61% is associated with cycles connecting different regimes. The dominant contribution is the Idling–Operational combination, which accounts for approximately 59% of the complete LFFD-corrected fatigue sum. At the same time, the Operational–Operational contribution decreases from approximately 88% under conventional processing to around 37% after residue recovery.

This does not mean that the cycles previously classified as Operational simply change label, rather, combining the residues changes which turning points form complete rainflow cycles. Turning points that appeared as separate open excursions in operational and idling blocks can become the endpoints of one larger transition cycle. The resulting cycle replaces part of the smaller within-state contribution and is assigned to the off-diagonal Idling–Operational entry.

The result identifies transitions between idling and production as the main operational source of the missing flapwise fatigue in this dataset. During such transitions, rotor speed, blade pitch, aerodynamic thrust and the mean flapwise load can all change substantially. The full load excursion may extend across several blocks, particularly during gradual startups, shutdowns, temporary interruptions or repeated changes in availability.

This interpretation is consistent with the previous results. The aggregation analysis showed that much of the missing fatigue develops over periods of hours or days, while the cycle-range distribution showed that the recovered cycles are concentrated at comparatively large ranges. Operating-state transitions provide a physical mechanism for producing cycles with both characteristics: their endpoints may be widely separated in time and may represent substantially different mean loading levels.

The dominance of the Idling–Operational entry is specific to the events contained in the available measurement period. Since representative high-wind shutdowns are absent and Fault intervals are scarce, no general conclusion can be drawn about their lifetime contribution. The result nevertheless demonstrates that processing each operating interval independently can omit an important part of the fatigue generated as the turbine moves between operating conditions.

These findings motivate the simulation cases examined later in chapter 6. In addition to longer periods of normal production, the BHawC study considers continuous sequences containing idling, startup, production and shutdown. This makes it possible to isolate the fatigue generated by the transitions themselves and to compare it with the result obtained when the same operating states are processed separately.

Overall, the measurement results show that independent 10-minute processing can underestimate fatigue when important load cycles extend across block boundaries. The effect depends strongly on the load channel and Wöhler exponent: it is substantial for blade flapwise and tower bending, but negligible for blade edgewise bending. The recovered fatigue is associated mainly with a small number of high-range cycles, many of which connect idling and operational periods. These findings motivate the simulation study in chapter 6, where the effects of longer production histories and continuous operating-state transitions are examined under controlled conditions.

6

Simulation-based results

This chapter presents the results of the BHawC simulation campaign introduced in chapter 4. The analysis begins with the baseline design load cases, for which conventional 600 s simulations are compared with directly simulated 3600 s records to isolate the effect of simulation duration while wind speed and operating conditions remain unchanged. The subsequent analysis preserves continuity across several operating states, allowing the influence of startup and shutdown transitions to be examined separately.

Results are presented in detail for the blade root flapwise bending moment, since the measurement analysis identified this response as particularly sensitive to longer-duration load variations. The corresponding results for the blade root edgewise and tower-bottom channels are included in the final channel comparison, but more detailed plots can be found in section B.1.

6.1. Baseline results

The baseline campaign compares the fatigue estimates obtained from 600 s and 3600 s simulations for the same design load cases. All results are calculated using consistent Wöhler exponents, reference cycle counts and DEL normalisation. Consequently, the longer simulations will not produce a larger DEL simply because they contain six times as many cycles. Instead, both durations provide estimates of the damage rate under the same nominal condition, and any difference arises from the load ranges sampled within the available records.

As a first step, the simulation comparison is restricted to DLC 1.2, normal power production under turbulent inflow. This provides the most direct basis for assessing the effect of extending the simulation duration, while keeping the turbine operating conditions unchanged. The remaining DLCs are considered later, once normal-production behaviour has been established.

6.1.1. Normal-production comparison

Figure 6.1 compares paired 600 s and 3600 s DLC 1.2 blade root flapwise bending-moment signals at 12 m/s. Since both simulations use the same turbulence seed, their first 600 s are identical, after which only the 3600 s case continues.

This example shows that a longer record can contain load behaviour absent from the initial 600 s. In the one-hour signal, the load level decreases noticeably around 1200 s. Although the cause cannot be identified from the load signal alone, such additional excursions may affect the DEL depending on their range.




Figure 6.1: Blade root flapwise bending-moment histories from paired 600 s and 3600 s DLC 1.2 simulations at 12 m/s using the same turbulence seed.

The same rainflow-counting procedure is applied to both simulation durations using the standard `FLS.rainflow`, as the duration is short enough to not have to use the LFFD approach to recover long cycles. Figure 6.2 presents the resulting DLC 1.2 DELs across the operational wind-speed range for the root blade flap moment. The solid curves show the mean DEL across the available seeds, while the shaded regions show the minimum-to-maximum seed range.




Figure 6.2: Mean blade root flapwise DEL for the 600 s and 3600 s DLC 1.2 simulations. The shaded regions indicate the minimum-to-maximum range across turbulence seeds.

The flapwise DEL generally increases with wind speed. Below rated wind speed, increasing aerodynamic loading and rotor speed produce progressively larger fatigue loads. A less regular response is visible around the transition to rated operation, around 13 m/s, where the controller moves between torque-controlled and pitch-regulated behaviour. Above rated wind speed, pitch regulation limits the increase in mean aerodynamic thrust, but turbulence and pitch activity continue to generate substantial flapwise variation. The DEL therefore remains high over much of the above-rated range and eventually flattens out, rather than increasing indefinitely with wind speed.

The seed envelopes overlap across most of the operating range, showing that the duration effect is generally small compared with the variation between turbulent realisations. Since the envelopes represent the full seed range, they are used only as a descriptive comparison. Above approximately 25 m/s, however, the mean DEL from the 3600 s simulations remains consistently higher. This may reflect broader turbulent excursions in the longer records or changes in turbine control within the HWRT region. Identifying the cause would require further analysis of blade pitch, rotor speed, generated power and mean flapwise load, which is outside the scope of this study.

Furthermore, the 600 s simulations also show a wider seed envelope over much of the wind-speed range. Since each result is calculated from a shorter record containing fewer cycles, it is more sensitive to the particular turbulent structures occurring within each seed. A single unusually large or mild excursion can therefore have a greater influence on the resulting DEL. The 3600 s records contain a larger cycle population, reducing this seed-to-seed sampling variability and producing a narrower envelope.

Although the wind-speed-dependent comparison provides useful insight, the turbine does not operate at every wind speed for the same amount of time. The contribution of each bin to the long-term fatigue estimate must therefore also account for its probability of occurrence. This is introduced in Figure 6.3, using the Weibull wind-speed distribution at the turbine's site.




Figure 6.3: Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the blade root flapwise moment.

The Weibull distribution is concentrated at comparatively low wind speeds, whereas the largest weighted fatigue contributions occur mainly between approximately 20 and 25 m/s. This shows that the most frequently occurring operating conditions are not necessarily those governing the flapwise fatigue result.

Although the probability of these higher wind speeds is lower, the associated load cycles are sufficiently large to dominate once their ranges are raised to the applied Wöhler exponent. The resulting lifetime contribution is therefore concentrated in a limited part of the above-rated operating range, where the balance between occurrence probability and fatigue severity is greatest. Across most wind speeds, the 600 s and 3600 s simulations contribute similar shares of the weighted damage. Each histogram is normalised to sum to 100% over the full wind-speed range. Above approximately 25 m/s, however, the 3600 s simulations again show a consistently larger contribution.

To show more clearly how these weighted damage contributions influence the final fatigue result, Figure 6.4 presents the same comparison on an equivalent-load scale.




Figure 6.4: Weibull-weighted equivalent-load contribution for the blade root flapwise moment.

Expressing the weighted contributions as equivalent loads makes the wind-speed regions governing the final result easier to identify. The contribution increases through the below-rated and rated regions, reaches its maximum between approximately 17 and 25 m/s, and then decreases towards the upper end of the operating range. At lower wind speeds, the high probability of occurrence is not sufficient to compensate for the comparatively small load ranges. At the highest wind speeds, the loads remain substantial, but the low probability of occurrence limits their contribution to the overall result.

The 600 s and 3600 s curves remain relatively close throughout the wind-speed range contributing most strongly to the weighted fatigue estimate. The one-hour simulations generally produce slightly larger contributions in this region, but the difference remains small compared

with the variation between seeds. The final power production DEL is therefore determined by the accumulated contribution of several wind-speed bins, rather than by the result at one individual wind speed.

To examine the effect of simulation duration more clearly, Figure 6.5 presents the ratio between the 3600 s and 600 s DELs for each wind-speed bin. A value above one indicates that the longer simulation produces a higher DEL, while a value below one indicates a lower result.




Figure 6.5: Ratio of the 3600 s and 600 s blade root flapwise DELs for each DLC 1.2 wind-speed bin. The shaded envelope spans the minimum and maximum seed-wise ratios.

The ratio varies around one and includes both increases and decreases, showing that extending the simulation does not automatically increase the DEL in every wind-speed bin. The seed envelope is widest at several low wind speeds, where the absolute DEL is small and a limited number of larger excursions can strongly affect the ratio. Division by a small 600 s reference value further increases the apparent variation. The envelope generally narrows over the central operating range, indicating more consistent results between seeds, but widens again at some high wind speeds where the response appears more sensitive to the turbulent realisation.

Across the wind-speed range contributing most strongly to the weighted fatigue estimate, the ratios are generally closer to one, although the 3600 s simulations tend to produce moderately higher values. After all wind-speed bins are combined at the damage level using their Weibull probabilities, the blade root flapwise DLC 1.2 DEL increases by 3.76% when the duration is extended to 3600 s, as reported in Table 6.1. This is considerably smaller than the largest bin-specific differences and demonstrates why the individual ratios should not be averaged directly. A large percentage difference at a wind speed with a small fatigue contribution may have little influence on the final weighted result.

The variation between individual turbulent realisations can be examined more directly in Fig-

ure 6.6. Each point compares the DEL obtained from the 600 s simulation with the DEL obtained from the extended 3600 s history using the same turbulence seed and nominal wind speed. The dashed diagonal represents equal DEL for the two durations.




Figure 6.6: Seed-matched comparison of the 600 s and 3600 s blade root flapwise DELs. Colour indicates the nominal wind speed, while the dashed line represents equal DEL for both durations.

In general, the points follow the equality line reasonably closely, confirming that the two durations produce similar DELs for most combinations of wind speed and turbulence seed. However, points occur on both sides of the line because the additional 3000 s can contain load excursions that are either more or less severe than those occurring during the initial 600 s. The duration effect for an individual simulation therefore depends on the particular turbulent loading sampled after the end of the shorter record.

Looking closely, a larger proportion of the high-wind cases lies above the equality line, which is consistent with the generally higher mean 3600 s DEL observed previously. The spread between seeds also increases at higher wind speeds, where larger turbulent fluctuations produce a wider range of load excursions. Consequently, this seed-to-seed variability is much greater than the change in the final Weibull-weighted DEL. The longer simulations therefore do not produce a uniform increase in fatigue loading, but provide a broader and more representative sample of the possible response at each wind speed.

To conclude, the normal-production comparison shows that increasing the simulation duration from 600 s to 3600 s produces a modest increase in the weighted blade root flapwise DEL. The effect varies substantially between individual wind speeds and seeds, but becomes considerably smaller once the complete wind-speed distribution is combined.

Other normal-production channels

The corresponding DLC 1.2 results for the remaining load channels are provided in section B.1. In summary, the blade root edgewise moment is almost unaffected by the increased duration, with the weighted DEL increasing by only 0.07%. The 600 s and 3600 s results remain closely aligned across the wind-speed range, as shown in Figures B.1 and B.10. This is consistent with the edgewise response being dominated by gravity-driven cycles that repeat many times within a 600 s simulation.

On the other hand, for the tower channels, the fore-aft DEL increases by 1.90%, whereas the side-side DEL decreases by 0.38%. The fore-aft response is aligned with the mean wind direction and is therefore governed largely by variations in aerodynamic thrust. Extending the simulation provides more opportunity to capture slower changes in thrust and a wider range of turbulent load excursions, which can lead to a modest increase in the resulting DEL. The side-side response, by contrast, has a smaller mean aerodynamic component under zero-yaw conditions and is dominated more strongly by oscillatory lateral and rotor-related loading. These cycles repeat frequently within a 600 s simulation, so increasing the duration has little systematic influence on the estimated fatigue load. The small negative difference should therefore not be interpreted as a meaningful reduction caused by the longer duration, but as part of the stochastic variation between the finite records. This interpretation is supported by the appendix comparisons, where the duration-related differences remain small relative to the seed-to-seed spread for both tower channels, as shown in Figures B.14 and B.15.

Overall, the normal-production results indicate that the effect of extending a stationary simulation to one hour is modest and depends on the dominant loading mechanism of each response. The analysis is therefore extended next to the remaining chosen design load cases, where differences in operating state and transient loading may produce a different sensitivity to simulation duration.

6.1.2. Combined baseline DLCs

The preceding comparison considered normal power production in isolation. However, the long-term fatigue estimate also includes periods of idling, together with discrete startup and shutdown events. The remaining baseline design load cases are therefore included to determine whether their duration sensitivity changes the combined result. As introduced in Table 4.4, these include DLC 3.1 startup, DLC 4.1m manual shutdown, DLC 6.4l low-wind idling, DLC 6.4h high-wind idling and DLC 7.2 idling following loss of grid.

Each DLC is first converted into an annual fatigue contribution using the weighting procedure introduced in chapter 4. Continuous operating states are weighted according to their wind-speed distribution and expected duration, whereas startup, shutdown and idling faults are weighted by their assumed number of annual occurrences. The individual contributions are then combined at the damage level before being converted into a total equivalent load based on an expected 25 year lifetime.

Figure 6.7 shows the weighted equivalent-load contribution of each design load case for the blade root flapwise moment for both the 600 s and 3600 s simulations. The error bars represent the variation across turbulence seeds. No error bars are shown for startup and shutdown because these simulations do not use turbulent inflow and therefore have no seed-to-seed variation. DLC 1.2 produces the largest component DEL for both durations. Its slightly higher 3600 s value is consistent with the normal-production results discussed previously, as the longer simulation samples a broader range of turbulent load excursions, including occasional larger and more damaging cycles that may not occur within the initial 600 s.




Figure 6.7: Weighted equivalent-load contribution of each baseline DLC for the blade root flapwise moment. Error bars indicate turbulent seed variability.

Furthermore, startup and shutdown produce noticeable component DELs because both events involve substantial changes in rotor speed, blade pitch and aerodynamic thrust as the turbine moves between idling and normal operation. The idling cases generally produce lower flapwise loads because the turbine is not generating power and the mean aerodynamic loading is reduced by design. Nevertheless, the clearest duration-related difference among these cases occurs for DLC 7.2, as shown in Figure 6.8.




Figure 6.8: Blade root flapwise bending-moment histories from paired 600 s and 3600 s DLC 7.2 simulations at 12 m/s using the same turbulence seed.

The blade root flapwise response during idling following loss of grid contains slow, large-amplitude variations that extend over several hundred seconds. The initial 600 s interval captures only part of this behaviour, whereas the 3600 s record contains several reversals between large positive and negative bending moments. These additional reversals introduce more high-range rainflow cycles and explain why the longer simulation produces a higher DEL for this condition. Nevertheless, its absolute contribution remains well below that of normal production, so the relatively large percentage change should not be interpreted as a dominant effect on the total fatigue result.

Although the startup and shutdown component DELs appear relatively close to the power production value in Figure 6.7, this does not directly reflect their importance to the total fatigue estimate. This figure represents the severity of the loading associated with each DLC after applying the relevant wind-speed weighting, but it does not yet account for how frequently each operating condition occurs.

The occurrence weighting is included in Figure 6.9. The proportion of time spent in each operating state and the frequency of transient events follow Table 4.6. Since normal production represents the largest share of turbine operation, DLC 1.2 contributes approximately 94% of the total annual flapwise damage. Startup, shutdown and idling can produce noticeable loads when they occur, but their lower occurrence greatly reduces their contribution to the combined result. The error bars show that the estimated damage shares can vary considerably between turbulence seeds, with a noticeably larger spread for the 600 s simulations than for the 3600 s simulations once again.




Figure 6.9: Contribution of each baseline DLC to the total annual blade root flapwise fatigue damage. Error bars indicate turbulent seed variability.

The dominance of DLC 1.2 means that the combined flapwise DEL is governed mainly by the duration effect identified during normal production. DLC 7.2 shows a much larger relative increase between the 600 s and 3600 s cases, but its contribution to the annual damage is too small for this difference to control the final result. Duration-related changes must therefore be

interpreted together with the occurrence and absolute fatigue contribution of the corresponding design load case. For the other load channels being evaluated, similar plots to Figures 6.7 and 6.9 can be found in section B.1.

Table 6.1 summarises the percentage change from the 600 s to the 3600 s results for all design load cases and response channels. Positive values in red indicate an increase in equivalent load, making the current 600 s approach non-conservative, while negative values in green indicate a decrease.

Table 6.1: Percentage change in DEL from the 600 s to the 3600 s simulations for each load channel and DLC.

Load channel	DLC 1.2	DLC 3.1	DLC 4.1m	DLC 6.4h	DLC 6.4l	DLC 7.2	Total
Change from 600 s to 3600 s [%]							
Root blade flap	3.76	0.50	-1.76	0.59	6.71	8.35	3.62
Root blade edge	0.07	-0.13	-0.04	–	-0.43	–	0.07
Tower bottom fore-aft	1.90	0.14	0.46	-4.58	9.92	-0.47	1.76
Tower bottom side-side	-0.38	0.64	0.01	0.32	-4.87	2.98	-0.36

For the blade root flapwise moment, the combined DEL increases by 3.62%. This remains close to the 3.76% increase obtained for DLC 1.2 alone because normal production accounts for the majority of the annual flapwise damage. The remaining design load cases introduce both positive and negative changes, but their lower occurrence means that they only slightly modify the final result. The larger relative increases for DLC 6.4l and DLC 7.2 therefore have limited influence on the combined DEL.

The blade root edgewise moment changes by only 0.07%, confirming that the duration effect is negligible for this response, in the same way it was negligible in the measurement data. The dominant edgewise loading is associated with gravity-driven cycles that repeat once per rotor revolution and are therefore sampled many times within a 600 s simulation. Extending the record mainly adds further repetitions of an already stable cycle population rather than introducing substantially different load ranges.

For the tower-bottom fore-aft moment, the combined DEL increases by 1.76%, which is slightly lower than the 1.90% increase obtained for normal production alone. Although low-wind idling shows a comparatively large positive difference, its annual fatigue contribution is small. This increase is partly offset by the negative changes in high-wind idling and DLC 7.2, leaving the combined result close to the DLC 1.2 value. The fore-aft response therefore remains somewhat sensitive to the longer turbulent record, but the inclusion of the remaining operating conditions does not substantially increase the overall duration effect.

The tower-bottom side-side DEL decreases by 0.36%, again remaining close to the -0.38% difference obtained for normal production. The positive and negative changes among the remaining DLCs are either small or associated with conditions that contribute little to the annual damage. Their influence on the combined result is therefore limited.

Overall, extending the stationary simulations from 600 s to 3600 s produces only modest changes in the combined fatigue estimates. The largest increase is obtained for blade root flapwise bending, at 3.62%, followed by tower-bottom fore-aft bending at 1.76%, while the blade edgewise and tower side-side responses remain effectively unchanged. The longer records provide a broader sample of the turbulent loading within each operating condition, but one hour is still short relative to many of the low-frequency variations observed in the measurements.

Simulations extending beyond 3600 s would therefore be expected to capture a wider range of slow load variations and could produce larger differences, consistent with the increasing fatigue contribution observed when longer aggregation periods were applied to the measurement data in subsection 5.3.3. However, increasing the duration in this manner rapidly becomes computationally demanding, since longer simulations would be required for every wind-speed bin, turbulence seed, load channel and design load case. The associated increases in simulation time, data storage and post-processing effort limit the practical duration that can be investigated using a full design-load-case campaign.

Furthermore, even substantially longer stationary simulations would not capture cycles caused by transitions between operating states unless these transitions are represented explicitly. The following section therefore proposes a different approach and considers a continuous sequence containing idling, startup, normal production and shutdown, allowing the fatigue contribution of the transitions themselves to be examined.

6.2. Continuous sequence of design load cases

The same preprocessed 3600 s load signal, containing several operating regimes, is analysed in two ways. First, the complete signal is rainflow counted continuously. Second, it is divided at the operating-state boundaries, each segment is counted independently, and the resulting damage is summed before calculating the DEL. All simulation and processing inputs are therefore identical, only the treatment of cycles crossing operating-state boundaries differs.

Figure 6.10 shows a representative blade root flapwise load history and the operational-state boundaries used in the comparison, for a wind speed of 13 m/s. The turbine follows the operational states described in Table 4.5. As shown, startup introduces a substantial change in the flapwise load as the turbine accelerates and thrust increases, going from idling to normal operational conditions. A further large and inverse change occurs during shutdown, as the turbine leaves normal power production and returns towards an idling condition.




Figure 6.10: Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for the blade root flapwise moment.

Although startup and shutdown occupy only a small fraction of the complete simulation, the associated changes in load level can be important for fatigue. As mentioned previously, rain-flow damage is governed not only by the number of cycles but also by their ranges. A small number of cycles spanning the transition between two substantially different load levels may therefore have a noticeable influence on the resulting DEL.

Figure 6.11 compares the mean blade root flapwise DEL obtained using the continuous and DLC-segmented approaches. Both curves follow the same overall wind-speed-dependent trend because they are calculated from the same load histories. The DEL increases rapidly through the below-rated range, reaches its highest value close to rated operation, and reduces slowly throughout the above-rated region as the blades are pitched-out. The wind-speed range considered extends from 3 to 23 m/s in 2 m/s increments, as this was deemed to represent realistic conditions under which the turbine could initiate startup.

Nevertheless, the most important result is that the continuously processed DEL is higher at every investigated wind speed. The difference is small at 3 m/s, but becomes clearly visible from 5 m/s onwards and persists throughout the rest of the operating range. The difference between the mean curves is generally larger than the minimum-to-maximum seed envelopes. This indicates a systematic processing effect rather than a difference caused only by one or two unusual turbulence realisations.




Figure 6.11: Mean blade root flapwise DEL from continuous and DLC-segmented processing. The shaded regions indicate the minimum-to-maximum range across turbulence seeds.

The magnitude of this effect is shown more directly in Figure 6.12. The ratio is close to one at 3 m/s, showing that the two approaches produce almost the same result under this condition. It then increases to reach a local maxima of almost 1.30 at 7 and 15 m/s. Over much of the remaining operating range, continuous processing produces a flapwise DEL approximately 15–25% higher than DLC-segmented processing. The ratio decreases again towards the highest investigated wind speed, but remains above one.

This result differs considerably from the stationary-duration comparison presented previously. When the 600 s and 3600 s stationary simulations were compared, the individual ratios oc-

curred on both sides of 1.0 because the longer record could contain either more or less severe turbulent loading than the initial 600 s. In the present comparison, however, the underlying signal is identical and the continuous result remains consistently larger. The difference can therefore be attributed directly to the treatment of the operational-state boundaries.




Figure 6.12: Ratio of blade root flapwise DEL from continuous and DLC-segmented processing. The shaded regions indicate the minimum-to-maximum range across turbulence seeds.

Further insight is provided by Figure 6.13, which shows the contribution of each operational state to the damage obtained when the segments are counted separately. For each wind-speed bin, the displayed shares are normalised and sum to 100%. At the lowest wind speeds, the separately counted damage is dominated almost fully by shutdown. Between approximately 7 and 15 m/s, startup and shutdown together account for most of the flapwise damage, despite their comparatively short durations relative to production. This confirms that the transient operating states can be highly fatigue-relevant because of the large changes in aerodynamic loading that occur while the rotor accelerates, the blades pitch and the turbine enters or leaves power production.

Nevertheless, the contribution from normal production increases with wind speed and becomes dominant above approximately 19 m/s. This development corresponds closely with the DEL-ratio result in Figure 6.12. The largest continuous-to-segmented ratios generally occur where substantial damage is distributed between more than one operating state. Under these conditions, the complete load history contains significant turning points on both sides of an operational boundary, allowing the continuous rainflow count to identify cycles that connect the different states. When the damage becomes concentrated predominantly within normal production at the highest wind speeds, a greater proportion of the important cycles can already be identified within that segment and the difference between the two processing approaches decreases.

However, the correspondence should not be strictly interpreted as meaning that the separately counted damage alone determines the ratio. The ranges and ordering of the turning points at each boundary are also important. Nevertheless, the result provides a physically consistent explanation for why segmentation has its greatest influence when several operational states

contribute meaningfully to the fatigue history.




Figure 6.13: Mean contribution of each operational state to the total blade root flapwise fatigue damage when the simulation is divided into separate DLC segments.

In conclusion, after combining the wind-speed bins using the adopted Weibull probabilities, continuous processing increases the blade root flapwise DEL by 19.19% relative to DLC-segmented processing. This is substantially larger than the 3.62% increase obtained by extending the independently simulated baseline design load cases from 600 s to 3600 s. The result therefore shows that explicitly representing realistic transitions and retaining the continuity of the load history can have a much greater influence than extending the duration of otherwise stationary operating conditions.

Table 6.2 presents the corresponding result for all four load channels. Since the DEL change is positive for all, it indicates that processing the full operational sequence continuously produces a higher Weibull-weighted DEL than counting the operational states separately. Detailed plots for the other channels can be found in section B.2.

Table 6.2: Percentage increase in the Weibull-weighted DEL obtained from continuous processing relative to DLC-segmented processing.

Load component	DEL change [%]
Blade root flapwise moment	19.19
Blade root edgewise moment	0.04
Tower-bottom fore-aft moment	9.74
Tower-bottom side-side moment	1.49

The results show a clear dependence on the dominant loading mechanism of each response. The blade root edgewise DEL changes by only 0.04%, meaning that continuous and seg-

mented processing produce effectively identical results. As shown in Figure B.31, nearly all of the separately counted edgewise damage is produced during normal power production for all wind speeds. Once again, the edgewise response is dominated by gravity-driven loading that repeats once per rotor revolution, so a large number of the relevant cycles are already completed within the production segment. Entering or leaving power production therefore introduces very little additional edgewise fatigue when the boundaries are counted continuously. This is further supported by the fact that for edgewise loading, the mean load does not jump significantly from one regime to the other as it did for flapwise, as can be seen in Figure B.22.

The tower-bottom fore-aft DEL increases by 9.74%, making it the second-most sensitive channel after blade root flapwise bending. The corresponding operational-state contributions in Figure B.32 show that the fore-aft damage is distributed relatively evenly between idling, startup, production and shutdown over much of the wind-speed range. The fore-aft response is strongly influenced by aerodynamic thrust, which changes substantially as the turbine moves between non-producing and producing states, as shown in Figure B.23. Continuous processing therefore has several opportunities to connect load reversals occurring on opposite sides of the operational boundaries. The resulting increase is lower than for blade flapwise bending, but remains large enough to be important and is much bigger than the 1.76% difference shown in Table 6.1.

By comparison, the tower-bottom side-side DEL increases by only 1.49%. Its separately counted damage is primarily associated with startup and production, with smaller contributions from idling and shutdown, as shown in Figure B.33. Although several operating states contribute to this response, the small continuous-to-segmented difference indicates that the transitions do not form additional high-range cycles comparable with those found in the flapwise and fore-aft directions, as the mean loads in the different regimes are closer to each other than in the other cases, as seen in Figure B.24. This is consistent with the side-side moment being less directly influenced by changes in mean aerodynamic thrust under the zero-yaw conditions considered here.

The comparison between channels is therefore physically coherent. The largest effects occur for the blade flapwise and tower fore-aft responses, both of which are strongly affected by changes in aerodynamic loading as the turbine enters and leaves normal production. The blade edgewise response, whose dominant cyclic loading remains present and repetitive within the production segment, is almost unaffected. The side-side response lies between these cases but remains much closer to the segmented result, as it has done consistently throughout this investigation.

It should be noted that the large 19.19% flapwise increase represents the controlled sequence investigated here and should not be applied directly as a correction to the complete design-lifetime DEL. Each simulated sequence contains one startup and one shutdown within a one-hour period, whereas the actual importance of these transitions depends on their long-term frequency and ordering during turbine operation. The purpose of the comparison is instead to isolate the effect of maintaining continuity across operational-state boundaries.

Within this scope, the result is clear: summing the damage from independently counted DLC segments is not equivalent to rainflow counting the corresponding continuous operating sequence. The following section evaluates whether existing 600 s simulations can approximate a continuous one-hour load history through concatenation.

6.3. Concatenated simulation results

Lastly, as introduced in section 4.4, a short study is performed by concatenating existing 600 s simulations to form one-hour load histories. This provides a practical alternative to repeating the full aeroelastic campaign with directly simulated 3600 s cases. Figure 6.14 compares the concatenated blade root flapwise DEL with the mean DEL from the direct one-hour simulations at each wind speed. A ratio of one indicates equal DEL, while values below and above one indicate lower and higher concatenated results, respectively.




Figure 6.14: Ratio of blade root flapwise DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases. The shaded band indicates seed variance of the one-hour long simulations.

The DEL ratio varies around one, showing that concatenation can either underpredict or overpredict the directly simulated one-hour result. The largest differences occur between approximately 6 and 10 m/s, where the ratio falls to about 0.5 at some wind speeds. This region also shows a relatively wide shaded band, indicating greater seed-to-seed variability in the one-hour simulations. The largest overprediction occurs near rated wind speed, where the concatenated DEL is approximately 16% higher than the reference. Above rated wind speed, the ratio generally remains closer to one, although it decreases again towards the upper end of the operating range.

The irregular changes between neighbouring wind speeds show that concatenation does not introduce a consistent bias. Instead, the agreement depends on the particular load excursions contained in the selected 600 s segments and in the directly simulated one-hour histories. This sensitivity is strongest at low wind speeds, where the seed variability is also larger, and for the flapwise response, where a small number of large cycles can strongly influence the DEL because of the high Wöhler exponent.

The corresponding results for the remaining load channels are provided in section B.3. The blade root edgewise moment shows the closest agreement with the directly simulated reference, consistent with its repetitive gravity-driven loading. Greater variation is observed for the tower-bottom fore-aft response, while the tower-bottom side-side result generally remains

within a narrower range.

Overall, concatenation can approximate the one-hour DEL at individual wind speeds, but the agreement is not consistent across the operating range. Since the method joins independently initialised simulations, it cannot fully reproduce the continuous evolution of the turbine response and turbulent inflow over one hour, even when abrupt discontinuities at the boundaries are minimised. The resulting rainflow-cycle distribution may therefore differ from that of a genuine 3600 s simulation.

Concatenation should consequently be treated as a practical sensitivity study rather than as a direct replacement for longer simulations. Furthermore, the comparison includes stochastic differences between the selected turbulent realisations. A direct seed-matched concatenation study and event-by-event comparison with measurements would provide a stronger validation, but both are outside the present scope and are retained as future work.

Conclusions and recommendations

This thesis investigated how representing wind-turbine load histories as independent 10-minute records affects fatigue estimates compared with longer, continuous histories. Measurements from an operating offshore wind turbine were combined with controlled BHawC aeroelastic simulations to distinguish three related effects: the loss of rainflow cycles at processing boundaries, the limited duration of conventional simulations, and the separation of different turbine operating states.

The main conclusion is that the 10-minute format is not inherently inadequate. The important limitation arises when consecutive records are treated as unrelated histories and their unresolved rainflow residues are discarded. This can produce a substantial non-conservative fatigue estimate when the structural response contains large load excursions that develop across several records. The magnitude of the effect depends strongly on the physical loading mechanism, the Wöhler exponent and the operating conditions contained in the load history.

7.1. Conclusions

The measurement analysis first established that cycles crossing 10-minute boundaries can be recovered without loading and processing the complete high-frequency record at once. The custom four-point rainflow implementation reproduced the established continuous rainflow calculation to the reported numerical precision for all validation windows. The residue-aware LFFD workflow also reproduced the continuous result when the same histories were divided into 10-minute blocks. Storing the closed-cycle contributions together with the ordered unresolved residues therefore preserves the fatigue information required to reconstruct cycles spanning multiple records. In the present implementation, this also made it possible to analyse the complete six-month record, which could not be processed directly in one pass because of memory limitations.

The comparison between conventional and residue-aware processing showed that the segmentation effect develops rapidly over the first part of the record and then becomes less sensitive to the particular analysis window. For blade root flapwise bending, the final six-month ratio DEL_{STA}/DEL_{LFFD} was approximately 0.84 at $m = 10$ and decreased to around 0.80 at $m = 12$. Independent 10-minute processing therefore underestimated the residue-aware flapwise DEL by approximately 16% to 20%. The ratio did not return towards one as more data was included, indicating that the difference was not caused only by one isolated event. Its apparent stabilisation nevertheless applies only to the processing ratio within the available

record and does not demonstrate that six months are fully representative of lifetime operation.

The segmentation effect was strongly dependent on the structural response channel. The measured blade root edgewise DEL differed by less than approximately 0.2%, whereas the tower-bottom X and Y moments were underestimated by approximately 3% at $m = 3$, 9% at $m = 4$, and 20% at $m = 5$. A single correction factor therefore cannot represent the influence of segmentation across the complete turbine. Blade edgewise bending is dominated mainly by gravity-driven cycles that repeat with each rotor revolution and close many times within a 10-minute block. Blade flapwise and tower bending are more sensitive to slower changes in aerodynamic loading, wind direction and turbine operation, allowing important turning points to remain separated across several blocks.

The cycle-range analysis explains why a relatively small number of missed cycles can produce these differences. Residue recovery caused little change in the large population of low- and moderate-range cycles already captured within each block. The main difference was concentrated in the upper tail of the distribution. For the flapwise channel, LFFD recovered cycles above approximately 15,000 kNm that were absent or strongly underrepresented under independent processing. Although these cycles occurred infrequently, their fatigue contribution increased rapidly with load range through the Wöhler relation. Increasing the Wöhler exponent therefore made the result more sensitive to the loss of these high-range cycles.

The aggregation analysis provided a time-scale interpretation of this behaviour. Independent 10-minute processing recovered approximately 80% to 84% of the full-record flapwise DEL, depending on the Wöhler exponent. Allowing residues to interact within each day increased this share to approximately 95% to 96%, while weekly aggregation recovered about 99%. Monthly aggregation produced an almost exact match with the complete six-month result. Most of the missing flapwise DEL was therefore associated with cycles crossing one or more 10-minute boundaries but closing within hours or days. Only a small additional contribution required residue propagation over several days or weeks. These time scales are specific to the investigated record, but they show that most of the segmentation effect can arise well before a six-month cycle is completed.

The operational-state analysis identified one important physical origin of the recovered cycles. Under residue-aware processing, approximately 61% of the measured flapwise fatigue contribution was associated with cycles whose endpoints occurred in different operating regimes. The Idling–Operational combination alone accounted for approximately 59% of the total. This does not imply that steady idling was more damaging than normal production. Instead, it indicates that transitions between low-load non-producing conditions and higher-load power production created large excursions whose turning points occurred in separate 10-minute blocks. The numerical shares are specific to the simplified classification and the events present in the six-month record, but the result demonstrates that transitions can form a major part of the fatigue contribution lost through independent processing.

The stationary BHawC simulations showed a different and generally smaller effect. After applying the adopted wind-speed and operating-state weights, extending the baseline simulations from 600 to 3600 s increased the combined DEL by 3.62% for blade root flapwise bending and by 1.76% for tower-bottom fore-aft bending. The blade root edgewise DEL increased by only 0.07%, while the tower-bottom side-side DEL decreased by 0.36%. The differences varied considerably between wind speeds and turbulence seeds, and the longer simulation did not consistently produce a higher DEL in every case. Extending a stationary simulation therefore improves the temporal sampling of that condition and reduces sensitivity to individual turbulent realisations, but it does not produce a universal duration correction.

The modest effect of increasing the stationary simulation duration is consistent with the inflow model used. Each case remains centred on one fixed mean wind speed and operating condition. Extending the simulation from 600 to 3600 s therefore samples more turbulent fluctuations, but it does not introduce the slower changes in mean wind speed, wind direction or turbine state observed in the measurements. Increasing simulation duration and preserving continuity between operating conditions consequently address two different sources of fatigue variation.

Preserving continuity across operating-state transitions had a much larger effect. The combined one-hour idling–startup–production–shutdown signal was processed both as one continuous history and as separate operating-state segments. Since both approaches used the same aeroelastic load history, the difference was caused only by cycles crossing the segment boundaries. Continuous processing increased the Weibull-weighted DEL by 19.19% for blade root flapwise bending, 9.74% for tower-bottom fore-aft bending, 1.49% for tower-bottom side-side bending and 0.04% for blade root edgewise bending.

This channel dependence is physically consistent with the measurement results. Startup and shutdown involve substantial changes in rotor speed, blade pitch, aerodynamic thrust and mean structural loading. Blade flapwise and tower fore-aft bending respond directly to these changes, allowing large cycles to form between the non-producing and producing parts of the sequence. The edgewise response remains dominated by repeating gravity-driven cycles and is therefore almost unaffected by whether the state boundaries are counted continuously. The side-side response is less directly affected by changes in mean thrust under the zero-yaw conditions considered and consequently shows a smaller increase.

The 19.19% flapwise and 9.74% fore-aft increases should not be interpreted as direct lifetime correction factors. The controlled sequence contains one startup and one shutdown within one hour, whereas their lifetime importance depends on the actual frequency, order and surrounding conditions of comparable events. The comparison instead demonstrates that summing damage from individually counted operating states is not equivalent to counting the corresponding continuous load history. For the investigated channels, preserving the transition itself had a greater influence than extending otherwise stationary simulations from 600 to 3600 s.

The measurement and simulation results therefore agree at the level of the underlying mechanism. Both indicate that responses governed by changes in aerodynamic mean loading are more sensitive to segmentation than responses dominated by rapidly repeating cycles. Both also show that the largest differences arise when substantial load changes occur across the boundaries between producing and non-producing conditions. This agreement should not be interpreted as a direct validation of simulated load magnitudes against measurements, since the measured inflow was not reproduced and the measured tower moments were defined in fixed geographical directions rather than fore-aft and side-side coordinates.

Finally, concatenating independently initialised 600 s simulations did not provide a consistent substitute for directly simulated one-hour histories. The resulting DEL could be either above or below the one-hour reference, with particularly large differences at some below-rated wind speeds. Boundary matching reduced visible discontinuities but could not restore continuity in turbulence, rotor phase, structural motion or controller memory. Consequently, concatenation may provide a useful diagnostic of sensitivity to longer post-processed records, but it cannot be assumed to reproduce the cycle distribution of a physically continuous aeroelastic simulation.

Taken together, the results answer the main research question. Representing load histories as independent 10-minute records has little effect when the dominant fatigue cycles close repeat-

edly within each record. It can, however, substantially underestimate fatigue when damaging turning points are separated by longer aerodynamic variations or changes in operating state. The importance of this effect must therefore be evaluated separately for each response channel and material model. No universal low-frequency correction factor can be applied across components, turbines or operating histories.

7.2. Practical implications

For measurement-based fatigue monitoring, the results do not require the 10-minute storage format to be abandoned. Instead, the information retained from each block should be extended. Storing only a DEL or a cycle histogram removes the order of the unresolved turning points and makes cross-boundary cycles impossible to reconstruct. A more complete monitoring workflow should store the closed-cycle contribution, the ordered rainflow residue and, where operational attribution is required, the corresponding time indices. These compact outputs can then be combined over different periods without repeatedly loading the complete high-frequency signal.

The required residue-propagation period should be selected for the specific channel and dataset. For the measured flapwise response, daily aggregation recovered most of the missing DEL, weekly aggregation recovered approximately 99%, and monthly aggregation produced an almost exact match with the full six-month result. These values provide evidence that very long uninterrupted records are not always required to recover most of the effect. They should not, however, be treated as universal limits. Responses dominated by repeating rotor cycles may require almost no propagation, whereas channels affected by multi-day wind-direction changes, availability events or prolonged shutdowns may require longer periods.

For aeroelastic assessment, the results do not provide a strong basis for uniformly replacing every conventional 600 s simulation with a 3600 s simulation. Extending the stationary cases produced modest changes in the combined DELs and remained subject to turbulence-seed variability. Longer simulations can improve the sampling of an individual operating condition, but computational effort may be more effectively directed towards the channels and conditions expected to contain slow or transitional loading.

Continuous simulations should therefore be prioritised where the turbine moves between states with substantially different mean loads. This is particularly relevant for blade flapwise and tower fore-aft bending during startup, shutdown, temporary interruptions, curtailment or recovery from faults. Such sequences should be designed using realistic event durations and should ultimately be weighted by observed or expected transition frequencies. Their results should complement, rather than automatically replace, the conventional DLC framework.

Where continuity is important, directly simulated histories remain preferable to concatenated short records. Concatenation may be useful for screening existing simulation databases or identifying conditions deserving further investigation, but its accuracy should be checked against directly simulated histories before it is used in design, certification or lifetime-assessment decisions.

More generally, the findings show that the treatment of boundaries should be considered explicitly in fatigue workflows. Environmental and operating conditions may be represented separately for statistical weighting, but this does not guarantee that the fatigue cycles generated while moving between them have also been represented. The importance of these connecting cycles should be assessed whenever the mean structural response changes substantially between adjacent conditions.

7.3. Limitations and future work

The measurement results are based on approximately six months of data from one offshore turbine. Although several processing ratios became comparatively stable, the record does not cover a complete annual cycle and cannot establish lifetime representativeness. Fault periods were scarce, and representative high-wind shutdowns were absent. Extending the analysis to at least one complete year, and preferably to several years and turbines, would show whether the observed channel dependence and recovery time scales remain consistent across seasons, wind climates, sea states and availability conditions.

The operational classification was intentionally simplified to three regimes and assigned one state to each 10-minute interval. More detailed controller and status information would allow startup, shutdown, curtailment, derating, maintenance, grid faults and different idling conditions to be separated. Higher-resolution operational labels would also reduce ambiguity when assigning the endpoints of recovered cycles to particular events. Future monitoring campaigns should therefore retain detailed status signals together with blade pitch, rotor speed, power, wind direction and nacelle heading.

The lack of nacelle-heading measurements prevented the measured tower moments from being transformed from fixed geographical X and Y directions into fore-aft and side-side components. This limited the comparison with the BHawC results to the broader physical mechanism. Recording the nacelle position would allow the aerodynamic thrust direction to be separated more clearly and would support a direct investigation of whether recovered tower cycles arise mainly from changes in thrust magnitude, wind direction or both.

The BHawC campaign used six turbulence seeds per wind-speed bin and simulations of at most one hour. One hour remains short compared with the daily and weekly scales identified in the measurements. Future work should therefore use targeted longer simulations rather than necessarily extending every DLC. Priority should be given to representative wind speeds, channels and operating sequences for which the present results indicate high sensitivity. Increasing the seed count would also provide a stronger statistical basis for separating systematic duration effects from stochastic variation.

The stationary Mann turbulence fields reproduce the prescribed short-term turbulence statistics but do not include all forms of mesoscale atmospheric variability. Future simulations could introduce continuously changing mean wind speed, wind direction, shear or turbulence intensity, either through synthetic profiles, measured inflow histories or higher-fidelity atmospheric models. This would make it possible to separate cycles caused by slowly varying inflow from those caused by changes in turbine operating state.

The controlled transition sequence should also be expanded. The present sequence uses prescribed durations and one startup and shutdown per simulation. A larger event library containing different initial conditions, transition speeds, shutdown types, curtailment events and fault recovery sequences would show how sensitive the recovered cycles are to event timing and controller response. Combining these events using transition frequencies and ordering derived from operational data would provide a more defensible estimate of their contribution to lifetime fatigue.

A direct event-by-event or one-to-one comparison would provide the strongest link between the measurement and simulation investigations. Measured periods containing pronounced recovered cycles could be selected together with the corresponding wind and operational information. Reproducing these periods in the aeroelastic model would allow the timing, range and operating origin of the simulated cycles to be compared with the measurements, rather

than comparing only general channel trends. Site-specific LIDAR measurements would further strengthen this comparison by providing a more detailed description of the incoming wind field during each event, including changes in wind speed, direction and shear.

Finally, the fatigue calculations were intentionally based on load range, a power-law Wöhler relation and Palmgren–Miner linear accumulation. Component-specific mean-stress corrections, nonlinear damage accumulation and load-sequence effects were not included. Future work should examine whether these modelling choices alter the relative importance of the recovered cycles, particularly because many transition cycles connect operating states with substantially different mean loads. Such developments would be most useful once the environmental and operational representativeness of the input histories has first been improved.

7.4. Final statement

The conventional 10-minute framework remains a practical basis for wind-turbine fatigue assessment, but the boundaries between records should not automatically be treated as fatigue boundaries. For rapidly repeating responses, independent processing is sufficient. For aerodynamically driven responses and histories containing substantial operating transitions, discarding the chronological residue can remove a small but highly damaging part of the cycle population.

A more complete fatigue assessment therefore requires both an accurate representation of the loading within each condition and a defensible treatment of how consecutive conditions connect. Residue-aware processing provides an efficient solution for long measurement records, while targeted continuous aeroelastic simulations can represent the transitions that independent DLCs omit. Together, these approaches retain the practical advantages of short-duration records while reducing the risk of non-conservative fatigue estimates.

References

- [1] Herbert J. Sutherland. *On the Fatigue Analysis of Wind Turbines*. Technical Report SAND99-0089. Albuquerque, NM, USA: Sandia National Laboratories, 1999. DOI: 10.2172/9460.
- [2] Negin Sadeghi et al. “Quantifying the effect of low-frequency fatigue dynamics on offshore wind turbine foundations: a comparative study”. In: *Wind Energy Science* 8 (2023), pp. 1839–1852. DOI: 10.5194/wes-8-1839-2023.
- [3] C. Hübler et al. “Reliability analysis of fatigue damage extrapolations of wind turbines using offshore strain measurements”. In: *Journal of Physics: Conference Series* 1037 (2018), p. 032035. DOI: 10.1088/1742-6596/1037/3/032035.
- [4] International Electrotechnical Commission. *IEC 61400-1:2019: Wind Energy Generation Systems – Part 1: Design Requirements*. International Standard IEC 61400-1:2019. Geneva, Switzerland: International Electrotechnical Commission, 2019. URL: <https://webstore.iec.ch/en/publication/26423>.
- [5] International Electrotechnical Commission. *IEC 61400-13:2015: Wind Turbines – Part 13: Measurement of Mechanical Loads*. International Standard IEC 61400-13:2015. Geneva, Switzerland: International Electrotechnical Commission, 2015. URL: <https://webstore.iec.ch/en/publication/23971>.
- [6] Isaac Van der Hoven. “Power Spectrum of Horizontal Wind Speed in the Frequency Range from 0.0007 to 900 Cycles per Hour”. In: *Journal of Meteorology* 14.2 (1957), pp. 160–164. DOI: 10.1175/1520-0469(1957)014<0160:PSOHS>2.0.CO;2.
- [7] Song-Lak Kang and Hoonill Won. “Spectral Structure of 5 Year Time Series of Horizontal Wind Speed at the Boulder Atmospheric Observatory”. In: *Journal of Geophysical Research: Atmospheres* 121.20 (2016), pp. 11946–11967. DOI: 10.1002/2016JD025289.
- [8] Emilio Bianchi, Edinson Andrés Solarte Casanova, and Tomás Manuel Guozden. “Spatiotemporal variability of the wind power resource in Argentina and Uruguay”. In: *Wind Energy* 22 (2019). Figure shows spectral density of wind speeds according to Van der Hoven, adapted from Dutta et al., pp. 1086–1100. DOI: 10.1002/we.2342.
- [9] Nirav Dangi et al. “The Effect of Turbulent Coherent Structures in Atmospheric Flow on Wind Turbine Loads”. In: *Renewable Energy* 241 (2025), p. 122248. DOI: 10.1016/j.renene.2024.122248.
- [10] Herman Frederik Veldkamp. “Chances in Wind Energy: A Probabilistic Approach to Wind Turbine Fatigue Design”. PhD thesis. Delft, The Netherlands: Delft University of Technology, 2006. ISBN: 978-90-76468-12-9.
- [11] Gabriel Marsh et al. “Review and application of rainflow residue processing techniques for accurate fatigue damage estimation”. In: *International Journal of Fatigue* 82 (2016), pp. 757–765. DOI: 10.1016/j.ijfatigue.2015.10.007.
- [12] Negin Sadeghi et al. “Fatigue damage calculation of offshore wind turbines’ long-term data considering the low-frequency fatigue dynamics”. In: *Journal of Physics: Conference Series* 2265 (2022), p. 032063. DOI: 10.1088/1742-6596/2265/3/032063.


- [13] M. A. Miner. "Cumulative Damage in Fatigue". In: *Journal of Applied Mechanics* 12 (1945), A159–A164. DOI: 10.1115/1.4009458.
- [14] Jenni Rinker. *Fatigue Damage and Wöhler Curves*. Teaching material, lecture slides. Slides originally prepared by Taeseong Kim. DTU Wind Energy.
- [15] Herbert J. Sutherland. "A Summary of the Fatigue Properties of Wind Turbine Materials". In: *Wind Energy* 3.1 (2000), pp. 1–34. DOI: 10.1002/1099-1824(200001/03)3:1<1::AID-WE28>3.0.CO;2-2.
- [16] ASTM International. *Standard Practices for Cycle Counting in Fatigue Analysis*. Standard ASTM E1049-85(2023). West Conshohocken, PA, USA: ASTM International, 2023. DOI: 10.1520/E1049-85R23. URL: <https://www.astm.org/e1049-85r23.html>.
- [17] C. Amzallag et al. "Standardization of the rainflow counting method for fatigue analysis". In: *International Journal of Fatigue* 16.4 (1994), pp. 287–293. DOI: 10.1016/0142-1123(94)90343-3.
- [18] M. Matsuishi and T. Endo. "Fatigue of Metals Subjected to Varying Stress". In: *Proceedings of the Kyushu District Meeting of the Japan Society of Mechanical Engineers*. Fukuoka, Japan, 1968.
- [19] Darrell F. Socie. "Rainflow Cycle Counting: A Historical Perspective". In: *The Rainflow Method in Fatigue: The Tatsuo Endo Memorial Volume*. Ed. by Y. Murakami. Oxford, UK: Butterworth-Heinemann, 1992, pp. 3–10. DOI: 10.1016/B978-0-7506-0504-5.50009-6.
- [20] *Wikimedia Commons image*. Wikimedia Commons. Accessed 2026-03-03. URL: <https://commons.wikimedia.org/w/index.php?curid=45235675>.
- [21] Kenneth Thomsen. *The Statistical Variation of Wind Turbine Fatigue Loads*. Risø-R 1063(EN). Roskilde, Denmark: Risø National Laboratory, 1998. ISBN: 87-550-2410-6. URL: <https://www.osti.gov/etdeweb/servlets/purl/292704>.
- [22] J. C. Kaimal et al. "Spectral Characteristics of Surface-Layer Turbulence". In: *Quarterly Journal of the Royal Meteorological Society* 98.417 (1972), pp. 563–589. DOI: 10.1002/qj.49709841707.
- [23] Jakob Mann. "Wind Field Simulation". In: *Probabilistic Engineering Mechanics* 13.4 (1998), pp. 269–282. DOI: 10.1016/S0266-8920(97)00036-2.
- [24] Shadan Mozafari et al. "Effects of Finite Sampling on Fatigue Damage Estimation of Wind Turbine Components: A Statistical Study". In: *Wind Engineering* 47.4 (2023), pp. 799–820. DOI: 10.1177/0309524X231163825.
- [25] Shadan Mozafari et al. "Sensitivity of Fatigue Reliability in Wind Turbines: Effects of Design Turbulence and the Wöhler Exponent". In: *Wind Energy Science* 9 (2024), pp. 799–820. DOI: 10.5194/wes-9-799-2024.
- [26] Laurent M. M. van den Bos et al. "Fatigue Design Load Calculations of the Offshore NREL 5 MW Benchmark Turbine Using Quadrature Rule Techniques". In: *Wind Energy* 23.5 (2020), pp. 1181–1195. DOI: 10.1002/we.2470.
- [27] Rad Haghi, Cassidy Stagg, and Curran Crawford. "Wind Turbine Damage Equivalent Load Assessment Using Gaussian Process Regression Combining Measurement and Synthetic Data". In: *Energies* 17.2 (2024), p. 346. DOI: 10.3390/en17020346.
- [28] Bruno Rodrigues Faria et al. "Inclusion of low-frequency cycles on tower fatigue lifetime assessment through relevant environmental and operational conditions". In: *Journal of Physics: Conference Series* 2767 (2024). Proceedings of TORQUE 2024, p. 042021. DOI: 10.1088/1742-6596/2767/4/042021.

- [29] Srinivas Guntur et al. “A validation and code-to-code verification of FAST for a megawatt-scale wind turbine with aeroelastically tailored blades”. In: *Wind Energy Science* 2 (2017), pp. 443–468. DOI: 10.5194/wes-2-443-2017.
- [30] Etienne Muller et al. “Field-data-based validation of an aero-servo-elastic solver for high-fidelity large-eddy simulations of industrial wind turbines”. In: *Wind Energy Science* 9 (2024), pp. 25–48. DOI: 10.5194/wes-9-25-2024.
- [31] Neil N. Davis et al. “The Global Wind Atlas: A high-resolution dataset of climatologies and associated web-based application”. In: *Bulletin of the American Meteorological Society* 104.8 (2023), E1507–E1525. DOI: 10.1175/BAMS-D-21-0075.1.

Use of Artificial Intelligence

Artificial Intelligence (AI) was used as a supporting tool during parts of the preparation of this thesis. Its use included assistance with the development and debugging of selected post-processing scripts, the creation and improvement of figures used to explain the methodology and results more clearly, and the rephrasing of selected paragraphs to improve readability and clarity.

All AI-generated suggestions were critically reviewed, adapted and verified by the author before being included. The author remains fully responsible for the methodology, implementation, interpretation of the results, and all conclusions presented in this thesis.



Measurement data processing

A.1. Measurement data filtering

Measurements were obtained from a 7 MW offshore direct-drive wind turbine with a rotor diameter of 154 m. The selected period extended from 1 May to 1 November 2022, providing six months of continuous data. This period was selected because it was sufficiently long to provide a representative record, while exhibiting the lowest incidence of sensor faults and missing data within the available measurement campaign.

The measurement data pre-processing procedure is summarized in Figure A.1.




Figure A.1: Simplified workflow used for measurement data pre-processing.

Four structural load channels were considered:

Table A.1: Measurement channels used in the fatigue analysis.

Database channel	Description
BladeA_FlapMom	Blade root flapwise bending moment
BladeA_EdgeMom	Blade root edgewise bending moment
Tower_MomX_Bottom	Tower bottom bending moment about the geographical East direction
Tower_MomY_Bottom	Tower bottom bending moment about the geographical North direction

The four load channels were not originally stored at the same sampling frequency. Before the fatigue analysis, each channel was therefore resampled onto a common and uniformly spaced time grid at 20 Hz. Linear interpolation was used during this step so that the channels shared the same timestamps and could be divided into equally sized analysis intervals.

The complete six-month record was too large to load into memory at once. For the long-duration analysis, the measurements were therefore read in consecutive one-day chunks. These daily chunks were used only for data handling and did not represent boundaries in the fatigue calculation. When valid measurements continued from one day into the next, the LFFD residue was carried across the daily boundary so that cycles extending over midnight could still be recovered.

After resampling, the signals were divided into consecutive, non-overlapping 600 s blocks. At a sampling frequency of 20 Hz, each block contained 12,000 samples. Only complete 10-minute blocks were retained, ensuring that all blocks had the same duration and could be compared using the same DEL normalisation.

The amount of valid data was then checked separately for every block. Blocks containing no missing samples were retained without modification and remained connected to the neighbouring valid blocks. When at least 95% of the samples were available, the short gaps were filled using linear interpolation. These blocks were retained, but were treated as isolated intervals so that interpolated values could not connect fatigue cycles between neighbouring blocks.

Blocks containing less than 95% valid data were rejected. The LFFD residue chain was also terminated at these locations because the load history across the missing interval was unknown. Joining the residues on either side of a substantial data gap could otherwise create an artificial large cycle that was not present in the measured signal.

Before rainflow counting, consecutive repeated values were removed and each accepted signal was reduced to its turning points. These turning points consisted of the local maxima and minima required by the rainflow algorithm. No additional filtering, detrending, clipping or amplitude correction was applied. In particular, the slow variations were retained because they form part of the low-frequency loading investigated in this thesis.

For the conventional and LFFD analyses, rainflow processing was applied to the accepted 10-minute blocks according to the methods described in section 3.2. The DEL normalisation was based on the total duration of the accepted blocks. Rejected intervals were therefore excluded from the calculation rather than being interpreted as periods of zero structural loading.

A slightly different loading procedure was used for the continuous-reference validation. Each selected validation interval was loaded and resampled as one uninterrupted signal, after which

rainflow counting was applied to the complete interval without introducing 10-minute or daily counting boundaries. The same start time, end time and valid-data duration were used for the continuous, conventional and LFFD calculations. Their differences could therefore be attributed to the fatigue-processing method, rather than to differences in the analysed data.

A.2. Visualising the LFFD process

This section provides a visual example of the residue-aware low-frequency fatigue dynamics (LFFD) process described in section 3.3. The example uses one hour of blade root flapwise bending moment data divided into six consecutive 10-minute blocks. The purpose is to illustrate how unresolved turning points from the individual blocks are retained and used to recover cycles extending across the original block boundaries.

Figure A.2 shows the complete one-hour load signal. The vertical dashed lines indicate the boundaries between the six 10-minute blocks. In conventional processing, each block would be rainflow counted independently. Cycles closing within a block would be counted normally, while any turning points that cannot form a complete cycle within that block would remain in its residue.




Figure A.2: One-hour blade root flapwise bending moment divided into six consecutive 10-minute blocks.

The residue remaining from each block is shown together with the original signal in Figure A.3. Each colour represents the ordered residue obtained from a different 10-minute block, while the black markers indicate the retained turning points. Most of the high-frequency fluctuations have already formed closed cycles within their respective blocks and are therefore not included in the residue. Only the unresolved load excursions needed for the later recovery stage are retained.




Figure A.3: Rainflow residues from the individual 10-minute blocks overlaid on the original load signal.

The individual residues are then joined in chronological order, as shown in Figure A.4. Duplicate or non-alternating turning points introduced at the block boundaries are removed so that the combined series remains a valid turning-point sequence. Although the sequence contains far fewer points than the original high-frequency signal, it preserves the larger load changes that remained unresolved within the individual blocks.




Figure A.4: Residues from the six 10-minute blocks joined in chronological order.

Rainflow counting is then applied once more to the combined residue, allowing turning points originating from different 10-minute blocks to form complete cycles. The resulting cycle ranges are shown in Figure A.5. In this example, 20 additional full cycles are identified, reducing the

number of unresolved half cycles from 48 to 8. These cycles could not have been recovered when the six blocks were processed independently.

The example considers only six consecutive 10-minute blocks. The same procedure can be applied over much longer periods by continuing to combine the ordered residues from subsequent blocks. As more data is included, some of the eight half cycles remaining in this example may close and form additional full cycles.




Figure A.5: Cycle ranges obtained by rainflow counting the combined residue.

The complete LFFD result is obtained by combining three contributions:

1. The cycles that close within the individual 10-minute blocks
2. The full cycles recovered by rainflow counting the combined residue
3. The half cycles remaining in the final residue at the end of the analysed period.

In this way, the method preserves the practical 10-minute data structure while recovering fatigue cycles that extend across the original processing boundaries [11, 12, 17].

A.3. Distribution of operating regimes from measurement data

This section presents the distribution of the simplified operating regimes and relates their occurrence to wind speed and the measured turbine response. The available measurement channels used for the operational classification are rotor speed, blade pitch and grid power, together with the measured 10-minute mean wind speed. As described in section 3.7, the combined values of these signals are used to assign each 10-minute interval to one of three simplified operating regimes: Fault, Idling or Operational.

Figure A.6 summarises the measured turbine response and the resulting classification. The upper panels show the 10-minute mean blade pitch, rotor speed, grid power and flapwise bending moment as functions of wind speed, with trend lines added to highlight the overall behaviour. The lower panels show the assigned regime of each interval and the relative regime share within each wind-speed bin.




Figure A.6: Measured blade pitch, rotor speed, grid power and blade root flapwise bending moment as functions of 10-minute mean wind speed, together with the classified operating regimes.

From Figure A.6, it can be seen that rotor speed and grid power generally increase through the below-rated region, while blade pitch remains constant and only begins to rise as rated operation is approached. The flapwise bending moment also increases with wind speed until rated wind speed is reached, after which it starts decreasing as the blades begin to pitch out to reduce loading and keep power production constant. The scatter around these trends reflects the range of operating conditions present in the six-month record. Some points lie systematically away from the main trends, most likely because of curtailed or otherwise reduced-power

operation.

In terms of operating regime, Idling dominates at the lowest wind speeds, mainly below cut-in, while Operational intervals account for most observations across the normal production range. A smaller number of Idling and Fault intervals also appears at higher wind speeds, indicating temporary interruptions or periods of unavailability. Moreover, the total fraction of the six-month record spent in each regime is shown in Figure A.7 as a percentage of the complete measurement period.




Figure A.7: Time share of the simplified operating regimes over the measurement period.

Approximately 82% of the 10-minute intervals are classified as Operational, around 18% as Idling and less than 1% as Fault. However, these fractions describe how often each regime occurs, but not how much it contributes to fatigue. A state that takes up only a limited part of the record can still be associated with important fatigue cycles if entering or leaving that state produces a large change in structural loading.

B

Additional load-channel simulation results

This appendix presents the supplementary simulation results for the blade root edgewise, tower bottom fore–aft and tower bottom side–side load channels, complementing the blade root flapwise results discussed in chapter 6. Section B.1 contains the baseline comparisons between the 600s and 3600s simulations. Section B.2 then presents the results from the continuous operational sequences and compares continuous processing with processing the operating states separately. Finally, section B.3 compares concatenated 600s histories with the directly simulated 3600s cases.

B.1. Baseline simulations




Figure B.1: DEL as a function of wind speed for DLC 1.2 for the root blade edge moment.




Figure B.2: DEL as a function of wind speed for DLC 1.2 for the tower-bottom fore-aft moment.




Figure B.3: DEL as a function of wind speed for DLC 1.2 for the tower-bottom side-side moment.




Figure B.4: Wind-speed-dependent damage-equivalent loads and Weibull weighting for the root blade edge moment.




Figure B.5: Wind-speed-dependent damage-equivalent loads and Weibull weighting for the tower-bottom fore-aft moment.




Figure B.6: Wind-speed-dependent damage-equivalent loads and Weibull weighting for the tower-bottom side-side moment.




Figure B.7: Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the root blade edge moment.




Figure B.8: Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the tower-bottom fore-aft moment.




Figure B.9: Weibull occurrence probability and the contribution of each wind-speed bin to the total weighted DLC 1.2 fatigue damage of the tower-bottom side-side moment.




Figure B.10: Ratio of the 3600 s and 600 s blade root edge weighted DELs for each DLC 1.2 wind-speed bin.




Figure B.11: Ratio of the 3600 s and 600 s tower-bottom fore-aft weighted DELs for each DLC 1.2 wind-speed bin.




Figure B.12: Ratio of the 3600 s and 600 s tower-bottom side-side weighted DELs for each DLC 1.2 wind-speed bin.




Figure B.13: Comparison of 600s vs. 3600s DEL per simulation seed for root blade edge moment.




Figure B.14: Comparison of 600s vs. 3600s DEL per simulation seed for tower-bottom fore-aft moment.




Figure B.15: Comparison of 600s vs. 3600s DEL per simulation seed for tower-bottom side-side moment.




Figure B.16: Equivalent-load contribution of each design load case for the root blade edge moment.




Figure B.17: Equivalent-load contribution of each design load case for the tower-bottom fore-aft moment.




Figure B.18: Equivalent-load contribution of each design load case for the tower-bottom side-side moment.




Figure B.19: Contribution of each design load case to the total fatigue damage of the root blade edge moment.




Figure B.20: Contribution of each design load case to the total fatigue damage of the tower-bottom fore-aft moment.




Figure B.21: Contribution of each design load case to the total fatigue damage of the tower-bottom side-side moment.

B.2. DLC combination simulations




Figure B.22: Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for blade-root edgewise moment.




Figure B.23: Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for tower-bottom fore-aft moment.




Figure B.24: Combined simulation showing four operational states: idling (blue), startup (orange), normal power production (green), and shutdown (purple) for tower-bottom side-side moment.




Figure B.25: Mean blade-root edgewise DEL from continuous and DLC-segmented processing.




Figure B.26: Mean tower-bottom fore-aft DEL from continuous and DLC-segmented processing.




Figure B.27: Mean tower-bottom side-side DEL from continuous and DLC-segmented processing.




Figure B.28: Ratio of blade-root edgewise DEL from continuous and DLC-segmented processing.




Figure B.29: Ratio of tower-bottom fore-aft DEL from continuous and DLC-segmented processing.




Figure B.30: Ratio of tower-bottom side-side DEL from continuous and DLC-segmented processing.




Figure B.31: Mean contribution of each operational state to the total blade-root edgewise fatigue damage when the simulation is divided into separate DLC segments.




Figure B.32: Mean contribution of each operational state to the total tower-bottom fore-aft fatigue damage when the simulation is divided into separate DLC segments.




Figure B.33: Mean contribution of each operational state to the total tower-bottom side-side fatigue damage when the simulation is divided into separate DLC segments.

B.3. Concatenation simulations




Figure B.34: Ratio of blade root edgewise DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases.




Figure B.35: Ratio of tower bottom fore-aft DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases.




Figure B.36: Ratio of tower bottom side-side DEL from concatenated 600 s histories to the mean DEL from directly simulated 3600 s DLC 1.2 cases.