

Wind tunnel tests with combined pitch and free-floating flap control: Data-driven iterative feedforward controller tuning

Sachin T. Navalkar¹, Lars O. Bernhammer², Jurij Sodja³, Edwin van Solingen¹, Gijs A. M. van Kuik², and Jan-Willem van Wingerden¹

¹Delft Center for Systems and Control, Faculty of Mechanical, Maritime and Materials Engineering, Delft University of Technology, 2628 CD Delft, the Netherlands.

²Wind Energy Group, Faculty of Aerospace Engineering, Delft University of Technology, 2629 HS Delft, the Netherlands.

³Aerospace Structures and Computational Mechanics, Faculty of Aerospace Engineering, Delft University of Technology, 2629 HS Delft, the Netherlands.

Correspondence to: Sachin Navalkar (S.T.Navalkar@tudelft.nl)

Abstract. Wind turbine load alleviation has traditionally been addressed in literature using either full-span pitch control, which has limited bandwidth, or trailing-edge flap control, which typically shows low control authority due to actuation constraints. This paper combines both methods, and demonstrates the feasibility and advantages of such a combined control strategy on a scaled prototype in a series of wind tunnel tests. The pitchable blades of the test turbine are instrumented with free-floating flaps close to the tip, designed such that they aerodynamically magnify the low stroke of high-bandwidth actuators. The additional degree of freedom leads to aeroelastic coupling with the blade flexible modes. The inertia of the flaps was tuned such that instability occurs just beyond the operational envelope of the wind turbine; the system can however be stabilised using collocated closed-loop control. A feedforward controller is shown capable of significant reduction of the deterministic loads of the turbine. Iterative feedforward tuning, in combination with a stabilising feedback controller, is used to optimise the controller online in an automated manner, to maximise load reduction. Since the system is non-linear, the controller gains vary with wind speed; this paper also shows that iterative feedforward tuning is capable of generating the optimal gain schedule online.

1 Introduction

The increasing size and flexibility of wind turbines demand that attention be devoted towards the active and passive control of rotor loads in order to limit the costs related to both the construction as well as maintenance of the turbine blades and the support structure. The dominant dynamic loading of turbine components occurs at 1P (rotor speed), and its harmonics. One of the most interesting and readily accessible methods of blade load control is Individual Pitch Control (IPC) Bossanyi (2003), whereby each blade is pitched along its longitudinal axis independently to counteract the variation in wind loading. Numerous references can be found in literature which prove the efficacy of IPC in load control of wind turbines, both in the simulation environment Selvam et al. (2009), Bottasso et al. (2013), as well as experimentally Bossanyi et al. (2013). In these references, with ordinary levels of turbulence, it has been observed that IPC can achieve upto 30% reduction in the

standard deviation of blade loads. Previous experimental studies conducted by the authors Navalkar et al. (2015) show that in a controlled, wind tunnel environment, wind turbine blade load reductions upto 70% can be reached, since the blade loading under these circumstances is almost entirely deterministic. However, in all the references mentioned, the target of IPC has been to reduce low-frequency loads, primarily around the 1P (rotor frequency). While IPC can herewith address a large part of load spectrum, the emphasis on low frequencies is also a product of the low bandwidth that can be achieved with the full-span pitch control, which involves actuation of the large torsional inertia of the blades around their axes. As expected, IPC also leads to a substantial increase in pitch activity.

In an effort to reduce pitch actuator duty, target higher frequencies in the load spectrum, and address localised disturbances in the wind loading, recent literature has explored the concept of the ‘smart’ rotor Lackner and Van Kuik (2010): a rotor where the blades are instrumented with sensors and flow-modifying actuators at various radial locations. Reviews of such rotors Barlas and Van Kuik (2010) Bernhammer et al. (2012) invariably conclude that trailing-edge flaps give the best control authority for load alleviation. The load alleviation potential has been demonstrated in simulations Andersen et al. (2006) Bernhammer et al. (Article in Press) and experimentally in a wind tunnel Van Wingerden et al. (2011). Further, field tests of this concept have also been conducted Castaignet et al. (2013), although such a system is still not considered mature enough for incorporation in a commercial wind turbine. While these tests used conventional actuators, many references recommend the usage of smart actuators, such as piezoelectrics, in order to enhance bandwidth and achieve a high power-to-weight ratio. Such actuators unfortunately show low stroke and hence reduced control authority.

The concept of the free-floating flap Heinze and Karpel (2006) combines a trailing-edge flap that is free to rotate about its axis, with a small tab located on the flap, that can be actuated at a high speed to dynamically change flap camber. This concept was developed specifically to take advantage of aerodynamic leveraging to increase the low stroke of smart actuators. For a fixed wing instrumented with such a free-floating flap, it was experimentally shown Bernhammer et al. (2013) that it is possible to achieve enhanced control authority. Further, this study also demonstrated that such a flap could be completely autonomous in terms of energy consumption, and can be used as a plug-and-play device. This modularity shows promise for the construction and maintenance of future smart blade. However, this concept has not yet been demonstrated experimentally on a wind turbine.

Numerical and experimental investigations of the free-floating flap concept Heinze and Karpel (2006) Bernhammer et al. (2013) have shown that the additional degree of freedom adds a rigid-body mode to the system, the dynamics of which are strongly dependent on the total air speed at operation. Aeroelastic coupling of this mode with the flexible blade mode induces flutter at low wind speeds, an instability that can lead to dangerously high vibrations and even structural failure. However, it has also been shown in the references that closed-loop control of the tab can ensure safe operation of a fixed wing, well into the unstable regime. A pitchable wind turbine blade instrumented with free-floating flaps thus poses several control challenges. Firstly, the nature of the flap implies that its dynamic response is not constant but varies strongly with the wind speed. Such a system cannot be described by a linear time-invariant (LTI) state-space realisation, but can possibly be expressed as a linear-parameter-varying (LPV) system, where the time-varying parameter depends on the wind speed. Further, the presence of a stabilising closed-loop controller is mandatory. Finally, the uncertainties in flow and structure modelling imply that a robust controller may be unable to achieve the maximum possible load reductions. The advantage of using a data-driven control

strategy would be that input-output data could be used to (locally) optimise a simple user-defined criterion. Further, if such a strategy is used to tune a feedforward controller, then the optimised controller cannot, in the steady state, destabilise the plant, and in the best case, it may be able to achieve load reductions that may not be attainable by a conservative, globally robust controller.

5 Data-driven control of wind turbine loads has been demonstrated experimentally in Navalkar et al. (2015), where online recursive system identification was combined with online controller synthesis for minimising the periodic turbine loads. However, such a controller would be required to retune itself at every instant the ambient wind conditions change. An alternative methodology for the data-driven alleviation of wind loads has been described in Navalkar and Van Wingerden (2015) which employs the iterative feedback tuning (IFT) Hjalmarsson (2002) methodology to tune the gains of a fixed structure controller, hereby optimising a (convex) performance criterion. While IFT controllers have been used in the industry, they have typically been implemented to converge to linear time invariant controller structures Gevers (2002). The use of IFT for tuning the gains of time-varying controllers, as required for the current application, has been described in literature Navalkar and Van Wingerden (2015) but not yet demonstrated in practice.

The contribution of this paper is thus threefold: firstly, scaled wind turbine blades instrumented with outboard free-floating flaps are designed and manufactured for wind tunnel testing. Secondly, the load alleviation potential of the free-floating flaps in combination with individual pitch control is demonstrated for the first time in an experimental sense. The load alleviation potential is investigated in both the stable and unstable (post-flutter) modes of operation, and the importance of collocated control will be highlighted. Finally, the setup will serve as a test bench for a novel iterative feedback tuning algorithm that automatically tunes a controller gain schedule for load alleviation in real-time variable wind speed operation. The novelties found in the paper as explicitly stated as below:

- This is the first experimental demonstration of combined pitch and flap control.
- This is the first experimental demonstration of free-floating flaps applied to rotating wind turbine blades. This is also the first time their potential has been demonstrated for load reduction in wind turbines experimentally. Further, this is the first time that free-floating flaps have been shown to induce flutter on wind turbine blades experimentally.
- 25 – This is the first experiment where IFT has been devised and implemented for adaptively tuning the gain schedule on a non-linear (LPV) plant.

The remainder of the paper is organised as follows: Section II describes the design and manufacturing process for the wind turbine blades with free-floating flaps. Section III gives a brief description of the testing environment. The aeroelastic behaviour of the blades is studied in Section IV. The control algorithm used for load alleviation is formulated in Section V. The results of the testing are laid out in Section VI, and conclusions are drawn from these results in the final section.

2 Blade design and manufacturing

Since this paper reports on the first wind turbine implementation of free-floating flaps, first, the design of the experimental setup is discussed, and details regarding the materials, method of manufacture and assembly are provided. Primarily, the destabilising effect of the free-floating flap is studied in detail, and the parameters are tuned such that the blade is close to its
5 flutter point in order that maximal control authority is achieved.

The design of the blades formed the most important part of the design process of the scaled turbine, since it had to form a reasonable approximation of a full-scale wind turbine blade while adhering to the constraints set by the wind tunnel capabilities. The primary scaling that was aimed to be achieved was maintaining the ratio of blade first eigenfrequency to rotor speed (1P), as is done in Van Wingerden et al. (2011). This ratio is typically around 3.5 for the modern turbines, Bak (2013).

10 2.1 Blade design

The overall aerodynamic and structural design of the blades follows the procedure described in Van Wingerden et al. (2011), since the blades were designed for similar wind tunnel testing conditions. Aerodynamic and structural details of the blade design can be found in Hulskamp et al. (2011). However, as the wind tunnel experiments will also incorporate blade pitch control, the torsional inertia of the blades was reduced by scaling down the root chord by 30%. The root chord thus measures
15 200 mm, tapering to a tip chord value of 120 mm over a blade length of 750 mm, with a total blade twist of 12° .

Out of structural considerations, it was deemed necessary to minimise the weight of the blades, while ensuring adequate structural integrity to withstand the centrifugal and out-of-plane loading that the blade will be subject to during operation. An accurate aerodynamic shape of the blade was ensured by 3d printing the blade and then further reinforced with unidirectional carbon fibre spar caps, as shown in Figure 1. Small wind turbine blades have previously been manufactured in this manner by
20 the University of Stuttgart Bauer et al. (2014), and a comparison of different additive manufacturing techniques can be found in Karutz (2015). These references specifically investigate 3D printing of blades in a set of sections that are bonded together. In order to avoid solid plastic-plastic joints, it was decided that the blades in the current case would be printed in one piece.




Figure 1. Blade CAD model



Figure 2. Photograph of Blade

Three different materials (ABS M30, PC-ABS and nylon), that can be used for 3d printing, were evaluated qua their ability to bond with carbon fibre. For each material, a rectangular sample of size 200 mm x 30 mm, of thickness 3 mm, was 3d printed. Subsequently, each sample was bonded on the top and bottom with a single layer of unidirectional carbon fibres of thickness 0.14 mm, impregnated with epoxy resin. A four-point bending test to failure was then conducted with each of the samples. The distance between the supports was 140 mm, while the points of force application were 40 mm apart. The results of the test can be seen in Figures 3 and 4. In Figure 3, the behaviour to failure in bending can be observed. For small loads, the response is linear. At higher loads, small kinks can be observed in each of the curves, these physically represent the snapping of individual carbon fibres in compression. Finally, there is a large drop in strength when delamination occurs in the materials ABS M30 and nylon. For the material PC-ABS, brittle fracture occurs before delamination, as such the bond between this material and the carbon fibre spar is the best for this material. Further, it also holds its strength over a larger range of deformation than the other materials. Since PC-ABS also shows good mechanical workability, the choice was made to 3d print the scaled blade using this material.




Figure 3. Structural behaviour of the bond between 3d printed substrate and carbon fibre spar




Figure 4. The grey rectangle on the top and the two black rectangles below it are the 3d printed samples post fracture, placed on a sandstone-coloured desktop that forms the background. Top: ABS M30, Middle: PC-ABS, Bottom: Nylon

The blade was printed as a 3 mm-thick shell, with an internal spar structure, using stereolithography techniques. In order to add structural stiffness to the blades, a spanwise slot was engraved at the spar cap location on both the top and bottom of the blade. This slot was filled with a 0.14 mm thick layer of unidirectional carbon fibre tow impregnated with epoxy resin. The slot was then aerodynamically faired using crushed glass fibre epoxy filler, which was then sanded down for a smooth finish. A Computer-Aided Design (CAD) model of the blade and a photograph of the finished blade are shown in Figures 1 and 2. The CAD software Solidworks was used for designing the blade, with the blade material considered to be homogeneous and isotropic. The metal connection to the hub and the carbon fibre are modelled to be bonded to the blade ideally such that delamination is not possible. An ultimate loading case is simulated for a wind speed of 10 m/s, rotor speed of 400 rpm and

a thrust coefficient of 1. For this extreme case, the stresses in the plastic material are calculated to be less than the flexural strength of the material by a factor of safety of 1.3.

The designed static force-deflection curve, compared with the measured structural behaviour, is seen in Figure 5. It is interesting to note that the predicted stiffening effect of the carbon fibre layer is nearly identical. The tip deflection was calculated to be 17.2% lower with carbon fibre spars, while it was measured to be 16.6% lower after stiffening. A flexible mode analysis of the blade yields its first natural frequency as 16.43 Hz.

Post manufacture, the blades are instrumented with piezoelectric strain sensors on the top and bottom, at the root of each blade. These sensors provide a measure of the blade loads that are sought to be minimised.




Figure 5. Calculated stiffness characteristics compared with measured stiffness characteristics

2.2 Free-floating flap (FFF) design

The CAD design of the free-floating flaps is depicted in Figure 6. The leading edge is a continuation of the inboard portion of the blade, including the slot meant for carbon fibre stiffening. The hinge axis of the flap is mounted using bearings on an aluminium bracket just behind the spar; apart from the negligible bearing friction, it is entirely free to rotate. A T-section is connected to this axle, such that its interference with the mounting bracket provides limit stops for the rotation of the flap. The flap can hereby rotate freely through a maximum upward and downward deflection angle of 30°.

A metal plate (spring steel) of thickness 0.2 mm is sandwiched between the axle and the T-section. Two piezobenders (Macrofibre composite, type M8557-P1) are affixed rigidly to the top and the bottom, respectively, of this metal plate. The benders are electrically connected together in an antiparallel manner such that their piezoelectric effect reinforces each other and they produce the same magnitude but an opposite direction of strain in the substrate. A maximum voltage of +/- 500 V can be applied to the benders in order to emulate the behaviour of the trim tab from Heinze and Karpel (2006) and Bernhammer et al. (2013). Finally, an appropriate aerodynamic shape of the flap was achieved by embedding the instrumented metal plate into a highly compliant foam which was shaped according to the aerofoil geometry. The entire flap, from the angle-limiting

T-section to the foam spacers, is covered with a fairing shroud. A contactless angle encoder is embedded into the tip section, which provides feedback on the flap angular position.

This configuration causes a step change in the chordwise profile just aft of the spar, that produces undesirable aerodynamic behaviour, which is a well-known trade-off against the increase in the deformability of the trailing edge. In this experiment, to achieve a proof of concept for free-floating flaps, aerodynamic accuracy is sacrificed for control authority in the design of the flap.




Figure 6. Flap cross-section: trim tabs replaced by chordwise piezobenders




Figure 7. Photograph of flap

3 Aeroelastic blade analysis

While in the previous section it was ensured that the behaviour of the blade under the ultimate static structural load was acceptable, an aeroelastic analysis is required to determine the change in its structural response with increasing wind speed.

The rigid-body mode of the free-floating flap is expected to couple with the first flexible mode of the system, giving rise to a low wind speed form of flutter.

In order to analyse the aeroelastic behaviour of the blade instrumented with a free-floating flap, the blade is modelled in MSC/NASTRAN as a cantilever beam of non-uniform cross-section (CBAR elements). The various cross-sections of the modelled beam were taken at ten equidistant spanwise stations along the blade. Each element is rigidly connected to a flat-plate aerodynamic panel of the corresponding chordwise length. The flap is modelled in a similar manner. First, a structural modal analysis of the blade is carried out, at zero wind speed. The calculated modes of the blade are given in Table 1. The corresponding modal frequencies predicted by Solidworks are:

- 1st flapwise frequency: 18.97 Hz (Solidworks), 19.44 Hz (Nastran)
- 1st edgewise frequency: 78.37 Hz (Solidworks), 76.67 Hz (Nastran)
- 2nd flapwise frequency: 84.8 Hz (Solidworks), 87.88 Hz (Nastran).

It should be noted that these frequencies differ from the actual modal frequencies measured experimentally. This discrepancy arises because the experimental frequencies correspond to rotor modes, and hence also include the flexibility in the blade connection pieces, the motors and the hub and are hence necessarily lower than the blade modal frequencies. This difference between the calculated blade frequencies and the measured rotor frequencies can be considered modelling uncertainty, and forms powerful motivation for a data-driven controller that tunes itself in accordance with the true system parameters.

It is most interesting to note that the lowest flexible mode is the flapwise mode, with modal frequency 19.44 Hz. This is the mode most likely to couple unstably with the rigid body flap mode. The blade is significantly stiffer in both the lead-lag and torsional directions; these modes are hence unaffected by aerodynamic coupling. An actual turbine blade Bak (2013) is relatively softer in these directions, however even for such a blade, the flapwise mode is the most relevant one for load analysis and also possesses the lowest frequency. The current scaled blade design, with high lead-lag and torsional stiffness, allows us to study the low-speed flutter phenomenon with limited complexity.

Table 1. Structural modes of the blade at zero total air speed

Mode description	Modal frequency	Mode description	Modal frequency
Rigid-body flap mode	0 Hz	1 st Flapwise mode	19.44 Hz
1 st Lead-lag mode	76.67 Hz	2 nd Flapwise mode	87.88 Hz
3 rd Flapwise mode	223.9 Hz	2 nd Lead-lag mode	291.3 Hz
1 st Torsional mode	361.6 Hz	4 th Flapwise mode	449.6 Hz

The low-speed flutter phenomenon, as predicted by NASTRAN, can be seen in Figures 8 and 9. Here, the abscissæ correspond to total air speed, which is defined as the resultant of the inflow wind speed and turbine rotational speed at the blade tip. It should be noted that the speed regulation trajectory of the wind turbine is linear, such that rotor speed increases linearly with wind speed at the rate of 51.1 rpm/(m/s).

For the purpose of aeroelastic analysis, the blade has been considered to be held stationary, with inflow wind speed equal to the total air speed defined above. This assumption is not strictly valid, since the incident wind speed is lower at the inboard sections of the blade. However, since these sections undergo lower structural deformations, it is expected that the impact on the aeroelastic behaviour of the blade is also lower. Further, the blade is twisted such that the angle of attack along the span remains more or less constant. Since the blade is non-rotating and subject to constant wind flow, the aerodynamic panels attached to each section maintain a constant angle along the blade span. With these simplifying assumptions, a first-order approximation of the flutter behaviour of the turbine blade is synthesised.

It can be directly observed that the frequency of the rigid body flap mode rises linearly with total air speed. Due to coupling of this mode with the first blade flexible mode, the blade mode becomes unstable at the onset of flutter. For the given physical

configuration of the blade, flutter occurs at a total air speed of 36 m/s, which corresponds to a turbine rotor speed of 340 rpm, thus at a speed beyond the designed operational speed of the wind turbine (230 rpm).¹




Figure 8. Variation in modal frequency with total incident air speed.




Figure 9. Variation in modal damping with total incident air speed.

The aeroelastic analysis served as a guideline for designing the kinematic parameters of the free-floating flap. A sensitivity analysis showed that the flutter speed depends strongly on the inertia of the flaps about the hinge axis. As seen in Fig. 10, an increase in flap inertia increases the flutter speed. Since an increase in flutter speed is also associated with a decrease in control authority, the flap inertia is chosen such that flutter occurs at a speed just beyond the operational regime of the wind turbine. As such, the flap inertia is chosen to be 270 g-cm², so that the flutter speed is 36 m/s, as described before.

This aeroelastic analysis also served as a guideline for the design of the blades and for identifying the range of operation permissible in the experiments described in the sequel. Experimentally, it was observed that the onset of flutter occurred at 315 rpm. However, since this mode involves exponentially diverging vibrations in the blades, which cannot be physically limited, open-loop experiments in this unstable regime were not conducted out of safety considerations.

4 Testing Environment

The blades designed and analysed as above were mounted on the test turbine setup used previously in Navalkar et al. (2015). As described in this reference, the blades are connected to the hub through pitch servomotors.

The hub is mounted on a shaft supported by two main bearings located in the nacelle. The electrical connections of the hub are transferred to the stationary part of the wind turbine via slip rings, rated at 500 V, which is also the maximum voltage that can be fed to the piezobenders located outboard on the blades. Further, the shaft is instrumented with a torque transducer and

¹In principle, a higher (pre-flutter) rotor operational speed of upto 300 rpm could have been chosen; the tip speed ratio is in both cases virtually identical. However, the speed of 230 rpm gives the best ratio of forcing frequency to blade eigenfrequency. Further, the current tower design yields a very lightly damped tower torsional mode at 280 rpm, which is deemed necessary to avoid out of practical considerations




Figure 10. Variation of flutter speed with flap inertia.

speed encoder, and connected mechanically with the generator. The turbine is direct-drive; the rotor speed is the same as the generator speed. The generator is in turn connected electrically in series with an adjustable dump load amenable to resistance control. Thus, in principle this setup can also provide torque control. However, in this series of tests, the resistance of the dump load is kept constant. This implies that the wind turbine is in constant torque operation, and its rotor speed rises linearly with the incoming wind speed. This form of control deviates from classical variable-speed variable-pitch turbine control, which utilises collective pitch to ensure constant speed regulation above rated wind speed. However, the variable-speed constant-load operation of the scaled turbine serves three purposes: overspeed behaviour can be investigated, which may induce flutter, below-rated turbine behaviour can be emulated, and the use of adaptive control can be evaluated in terms of its ability to retune itself to adapt to changed operating conditions.

10 The nacelle is connected rigidly to the top of a tower, mounted on bearings on its base. The tower (and hence the entire wind turbine) can yaw freely around its base. For this set of experiments, the tower is kept fixed such that the plane of the rotor is always perpendicular to the incoming wind speed.

The entire assembly is mounted inside the Open Jet Facility of the Delft University of Technology, which is an open jet wind tunnel of 6 m test cross-section and 2.85 m equivalent open jet diameter. A photograph of the turbine can be seen in Figure 11. While wind speeds up to 35 m/s can be achieved in this wind tunnel, the operation of the wind turbine under the current settings requires no more than 6 m/s, with a rated wind speed of 4.5 m/s (and thus a tip speed ratio of 5.35).

Data acquisition and online control is furnished at a sampling frequency of 2 kHz by a real-time PC, on which the controller is compiled using Matlab-Simulink xPCTarget. There are two primary sensing elements: the load sensors at the blade roots and the free-floating flap angle sensors. Further, there are two primary actuators: the piezobenders on the flaps and the pitch motors. The objective of the experiments is to use these sensing and actuating elements to achieve load control of the scaled wind turbine.




Figure 11. Photograph of the assembled turbine with pitch and flap control.

5 Iterative Feedforward Tuning for Combined Pitch and Flap Control

For a wind turbine in the field, the blade loads arise mainly out of wind shear, tower shadow, turbulence and its rotational sampling. As such, the blade load spectrum for a typical turbine shows dominant peaks at the rotor speed (1P) and its harmonics: for a two-bladed turbine at 2P, 4P, and so on, while for a three-bladed machine at 3P, 6P etc. The presence of turbulence broadens these peaks and adds energy in the high frequency region of the spectrum.

In the wind tunnel environment, the levels of turbulence are low. The main cause of the blade loads is the tower passage, which leads to sharp peaks at 1P and its harmonics. The objective of the experiments is to demonstrate that these peaks can be attenuated by pitch and flap control, which by extension implies that a significant portion of the load spectrum of an in-field turbine can be addressed by these actuation and control methods.

For achieving load control, Iterative Feedforward Tuning (IFT) of the pitch and flap controllers will be implemented. This technique specifically targets deterministic disturbances, as seen in the blade load spectrum of the turbine, with minimal control action. As long as there exists a nominally stabilising controller in the loop to avoid the unstable flutter region, the controllers tuned using IFT will not render the plant unstable. Further, IFT ensures that data-driven tuning of the controllers makes them converge to an optimal control action over a number of iterations.

It should be noted that this optimality refers to the local optimum of the user-defined cost criterion, and is unrelated to global controller optimality. There are, at present, no global optimality proofs for IFT. Indeed, if a feedback controller is tuned using

IFT for a poorly-designed performance criterion, it may yield an unstable closed loop. However, since this paper considers IFT for feedforward control, this issue is not relevant. Further, if the step size in the gradient descent algorithms is too large, the parameter tuning process may become unstable. These issues have been dealt with by Hjalmarsson (2002).

The optimal controller parameters depend strongly on the incoming wind speed and hence demand a linear parameter-varying (LPV) controller. LPV controller tuning using IFT has been explored and shown to work in the simulation environment Navalkar and Van Wingerden (2015). However, the computational burden and number of experiments required for tuning imply that this methodology is required to be modified to meet the demands of real-time control in the wind tunnel. Hence, a quasi-LPV approach will be followed in this section. Accordingly, while the plant remains LPV at all times, when the wind speed varies slowly in the wind tunnel and the plant is approximated as LTI for the duration of each set of IFT experiments.

As a consequence of this assumption of constant dynamics, IFT tunes controller parameters that are optimal only for one specific operating point, while being suboptimal for the rest of the operating range. It is for this reason that the ordinary IFT process has to be repeated for different constant wind speeds, or an IFT gain schedule has to be generated for a varying wind speed.

Firstly, the notation for this section will be introduced, then the three IFT experiments will be described, and finally, the method for creating a gain schedule for controller wind speed adaptivity is described.

5.1 Preliminaries and Notation

An LPV formulation is set up for describing the wind turbine system, either in open-loop, or in closed-loop with a nominally stabilising controller that allows the system to be operated in a post-flutter regime:

$$x_{k+1} = A_k x_k + B_k u_k, \quad (1)$$

$$y_k = C_k x_k + D_k u_k + v_k. \quad (2)$$

In these equations, $x_k \in \mathbb{R}^{n_x}$ is the state vector of unknown size, $u_k \in \mathbb{R}^4$ are the control inputs, including the two pitch signals and the two flap signals, while $y_k \in \mathbb{R}^2$ are the blade load signals as measured by the sensors located at the roots of the two blades. The signal $v_k \in \mathbb{R}^2$ is the external forcing signal, produced in this case by tower passage; it is a superposition of a periodic signal with zero mean white noise. The state-space matrices A_k , B_k , C_k and D_k are considered unknown and have the appropriate dimensions. Although the system matrices are unknown, they are assumed to admit a specific LPV structure, as per Van Wingerden (2008). These matrices, as well as the disturbance v_k , are considered to be affine functions of the ambient wind speed, $V_k \in \mathbb{R}$:

$$A_k = A^{[0]} + V_k A^{[1]}, \quad (3)$$

and similarly for the other matrices. Here, $A^{[0]}$ and $A^{[1]}$ signify the unknown components of A_k , one which remains constant over time, and one which varies linearly with the wind speed, respectively. As per Van Wingerden (2008), there should be one more term that varies with V_k^2 , but the influence of this term is limited and it is neglected in this paper.

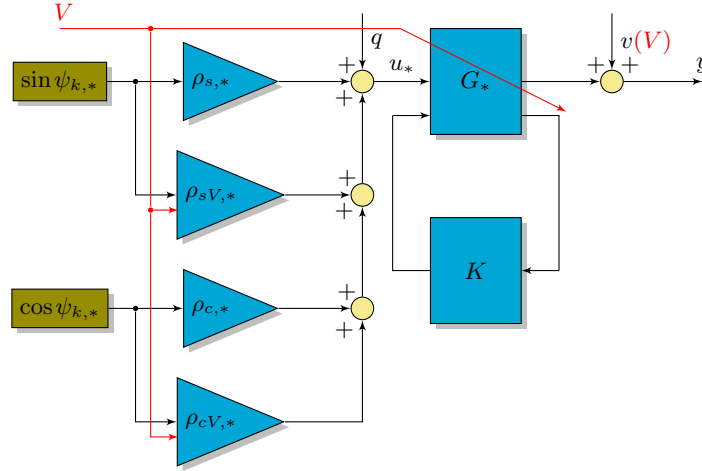


Figure 12. IFT implementation: wind turbine load control. * stands for IPC or IFC.

Typically for a wind turbine controller, a measure of the wind speed is not directly available. It is expected that advanced wind measuring sensors like the Lidar will in the future be able to provide accurate wind speed measurements Dunne et al. (2011). Meanwhile, it is possible to use the collective pitch angle in the above-rated region, and the generator speed in the below-rated region to approximate the value of ambient wind speed; in the current experiments the latter approach is used. In the theoretical framework, it is hence assumed that the controller possesses perfect knowledge of the wind speed.

The feedforward disturbance attenuating controller to be designed is considered to be a full LPV controller parameterised as follows:

$$\xi_{k+1} = A_{c,k}(\rho)\xi_k + B_{c,k}(\rho)r_k, \quad (4)$$

$$u_k = C_{c,k}(\rho)\xi_k + D_{c,k}(\rho)r_k - q_k. \quad (5)$$

Here, the controller is considered to be a fixed structure controller, such as a PID controller, with state $\xi \in \mathbb{R}^{n_c}$ of fixed dimension. The reference signal for this controller is taken to be a set of azimuth-locked basis functions, as in Navalkar and Van Wingerden (2015), thereby rendering this method a form of adaptive cyclic pitch and flap control. This form of open-loop control is depicted in the block diagram in Figure 12. For the pitch controller, these are sinusoidal functions of frequency 1P, while for the flap controller, these are sinusoidal functions of frequency 2P. Thus, the pitch and the flap control are both decoupled in the frequency domain, and are expected to strictly attenuate loads at their respective frequencies, in order to mitigate the maximum amount of disturbance with a minimum of control effort. Thus, with two sinusoidal basis functions for each frequency 1P and 2P, the reference signal is $r_k \in \mathbb{R}^4$. The term q_k refers to an auxiliary input that will be used in the IFT experiments described in the next subsection.

The objective of IFT is to minimise the loads as measured by the load sensors; so the cost criterion is:

$$J = \frac{1}{2N}(y^T y), \quad (6)$$

where N is a sufficiently long prediction horizon. For attenuating periodic loads, N is taken to be a multiple of the fundamental period of these loads. The term $y \in \mathbb{R}^{2N}$ is the load signal stacked over this horizon: $y = [y_1^T, y_2^T, \dots, y_N^T]^T$. The other signals are stacked in a similar manner. The key element of the IFT methodology is the optimisation of the system performance with the help of an experimentally derived performance gradient with respect to the controller parameters. This performance gradient is give by:

$$\frac{\partial J}{\partial \rho} = \frac{1}{N} \frac{\partial y^T}{\partial \rho} y. \quad (7)$$

Since the gradient contains stacked signals, it is more convenient to cast the system equations into a lifted format. Thus, for instance, the lifted system matrix for the wind turbine plant is given by the Toeplitz-like matrix $T \in \mathbb{R}^{2N \times 4N}$ as:

$$T = \mathcal{T}(A_k, B_k, C_k, D_k) = \begin{bmatrix} D_1 & 0 & \cdots & 0 \\ C_2 B_1 & D_2 & \cdots & 0 \\ C_3 A_2 B_1 & C_3 B_2 & \cdots & 0 \\ C_4 A_3 A_2 B_1 & C_4 A_3 B_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ C_N A^{N-1} \dots B_1 & C_N A^{N-1} \dots B_2 & \cdots & D_N \end{bmatrix}. \quad (8)$$

Like this lifted plant matrix T , a similar lifted matrix T_c can be constructed for the controller. It can be observed that the system matrices are functions of the wind speed V_k , which is approximated to be constant for each set of IFT experiments, but changes over the course of different sets of IFT experiments (or IFT iterations). As such, for the case where the wind speed is held constant at V_* for a set of IFT experiments, the system matrices of the plant and the controller are $T(V_*)$ and $T_c(V_*)$ respectively. Now, the linear time-invariant (LTI) IFT set of experiments that yield the controller gradient for the fixed wind speed V_* are described next.

5.2 IFT Experiments

In traditional IFT Hjalmarsson (2002), designed for LTI systems, the controller parameters can be iteratively optimised by repeatedly conducting a set of three experiments for each controller parameter. If the wind speed is considered to be constant over this set of experiments, the wind turbine plant reduces to an LTI system, and the same approach can be followed for optimising controller parameters for that specific wind speed. This section recapitulates the IFT methodology with this perspective.

In the first IFT experiment, the auxiliary signal is set to zero ($q_I = 0$). It is assumed that, over the set of these experiments, the wind speed is constant at a value of V_* . Accordingly, the output data collected is related to the system matrices as:

$$y_I = T(V_*)T_c(V_*)r + v_I(V_*). \quad (9)$$

Considering Equation 7, in order to find the performance gradient, it is necessary to determine the gradient of the output with respect to the parameters, $\frac{\partial y}{\partial \rho}$. As such, the equation above is differentiated with respect to each controller parameter $\rho_{j\rho}$,

$j_\rho = 1, \dots, n_\rho$, where n_ρ is the number of controller parameters. This results in the following equality:

$$\frac{\partial y_I}{\partial \rho_{j_\rho}} = T(V_*) \frac{\partial T_c(V_*)}{\partial \rho_{j_\rho}} r. \quad (10)$$

It should be noted that in the above equation, the only unknown on the right hand side is the system matrix of the plant, $T(V_*)$. In order to estimate its filtering effect, the second experiment uses the following auxiliary input:

$$5 \quad q_{II} = \frac{\partial T_c(V_*)}{\partial \rho_{j_\rho}} r. \quad (11)$$

Herewith, the output in the second experiment becomes:

$$y_{II} = (T(V_*) - T(V_*) \frac{\partial T_c(V_*)}{\partial \rho_{j_\rho}}) r + v_{II}(V_*). \quad (12)$$

The required output gradient $\frac{\partial y_I}{\partial \rho_{j_\rho}}$ is now given by:

$$\frac{\partial y_I}{\partial \rho_{j_\rho}} = y_I - y_{II} + v_{II}(V_*) - v_I(V_*). \quad (13)$$

- 10 Since the disturbance signal v is a superposition of a periodic signal and random noise, and the stacking length N is a multiple of period of the noise, the term $v_{II}(V_*) - v_I(V_*)$ does not contain a periodic component, and is purely zero mean white noise. As such, the output gradient below is ergodically unbiased:

$$\frac{\partial \hat{y}_I}{\partial \rho_{j_\rho}} = y_I - y_{II}. \quad (14)$$

However, the performance gradient cannot in this case be constructed simply as:

$$15 \quad \frac{\partial J}{\partial \rho_{j_\rho}} \Big|_{V=V_*} = \frac{1}{N} (y_I - y_{II})^T y_I. \quad (15)$$

This is because the noise in the estimate of the output gradient $\frac{\partial y_I}{\partial \rho_{j_\rho}}$ is correlated with the disturbance components in y_I and the performance gradient estimate would hence be biased. So, a third experiment, replicating the deterministic conditions of the first experiment, is required to be conducted, in order to obtain the statistically uncorrelated output y_{III} . Finally, the performance gradient is given by:

$$20 \quad \frac{\partial J}{\partial \rho_{j_\rho}} \Big|_{V=V_*} = \frac{1}{N} (y_I - y_{II})^T y_{III}. \quad (16)$$

With this performance gradient estimated from data, an optimisation method, such as a steepest descent method, can now be employed to obtain the optimal value of the controller parameter. It is to be noted, however, that the controller parameter derived in such a manner is optimal only for the operating wind speed. The iterations for achieving such a conditionally optimal controller parameter can be denoted by:

$$25 \quad \rho_{j_\rho}^{i+1}(V_*) = \rho_{j_\rho}^i(V_*) - \gamma^i R^{-1} \frac{\partial J}{\partial \rho_{j_\rho}} \Big|_{V=V_*}. \quad (17)$$

Here, the term γ is an (iteration-dependent) scalar step size that can be tuned for achieving the desired convergence rate. It should be noted that a step size that is too large may lead to non-convergence. The term R represents a positive definite matrix, which is identity for the steepest descent method, but can be the Hessian or an estimate of the Hessian with respect to the controller parameter for increasing the rate of convergence.

- 5 With this method, the optimal controller parameters for a specific wind speed can be iterated to. The next section details the synthesis of a gain schedule for adapting the parameters for the case with slowly varying wind speed.

5.3 Data-driven Gain schedule synthesis

The previous section details the manner in which, for a constant wind speed, an updated estimate of the ideal controller parameters for that wind speed can be obtained. In this section, it is assumed that, in each iteration i , the ideal parameters vary
10 as an affine function of the wind speed V^i in the following manner:

$$\rho_{j\rho}^*(V^i) = \rho_{j\rho}^{[0],*} + V^i \rho_{j\rho}^{[1],*}. \quad (18)$$

While the above equation indicates a linear relationship between the optimal parameter and the scheduling variable V , this may not in practice always be the case. However, the same equation can also be extended to an arbitrarily high degree of complexity, using either polynomial or any other suitable basis functions. The choice of the number of functions depends upon
15 the non-linearity of the scheduling, and on the signal to noise ratio achieved by the sensors, and may in practice be difficult to estimate a priori.

The objective of IFT is then to iterate to the optimal values of $\rho_{j\rho}^{[0]}$ and $\rho_{j\rho}^{[1]}$ based on the data inferred from the experiments described in the previous section. In the simple case of affine scheduling dependence described above, this can be achieved by recursive least squares estimation of the gain schedule. Thus, at every iteration that a pair $\rho_{j\rho}^{i+1}$ and V^{i+1} is obtained from
20 Equation 17, recursive linear regression is used to update the gain schedule.

This procedure is repeated until the gain schedule converges to the optimal gain schedule described in Equation 18. Hereby, IFT is able to synthesise an optimal gain-scheduled combined pitch and flap controller for the case where the wind speed varies slowly. The optimal tuning of such a controller will be demonstrated experimentally on the wind tunnel setup described in the previous sections.

25 6 Results

To recapitulate, the objective of the wind tunnel experiments was to achieve blade load control for the scaled wind turbine, using full-span pitch actuation and free-floating flap control, with Iterative Feedforward Tuning for optimal performance of the load controller. It should be noted that since experiments are conducted under constant load operation, the rotor speed varies linearly with wind speed. Thus, a rated wind speed of 4.5 m/s corresponds to a rotor speed of 230 rpm. The flutter speed of 6
30 m/s (total air speed 34 m/s) corresponds to a rotor speed of 315 rpm. In this section, operating conditions will be designated by the operating rotor speed.

6.1 System identification and stabilising controller

Initially, the response of the wind turbine blade to flap actuation is studied and compared with the simulations. Open-loop identification experiments are conducted in the pre-flutter regime, 200-300 rpm, with a zero-mean white noise (maximally +/- 500 V) imposed on the piezobenders. Predictor-based subspace identification (PBSID) Van der Veen et al. (2013) is performed using the acquired data to obtain the transfer function between the tab actuation and the flap angle and blade root load measurements. The transfer functions are depicted in Figures 13 and 14.

It can be observed that significant phase loss occurs in the transfer from the actuator to the blade root loads. This implies that stabilising the system using the measurements from the root loads poses a control challenge, and it may prove difficult in the case of uncertain systems to guarantee robust stability in the unstable post-flutter region. Further, it also motivates the use of local load sensors to enhance load attenuation capabilities. On the other hand, the phase loss in the transfer between the actuator and the flap angle measurement is minimal. This collocated sensor is hence ideal for system stabilisation in the post-flutter region. A simple classically tuned controller is used for stabilisation, it is not designed for load reduction, and is hence not optimal. It is described in continuous time as follows:

$$K = \underbrace{0.0001}_{\text{Static gain}} \underbrace{\frac{s/0.001 + 1}{s/10 + 1}}_{\text{High pass}} \underbrace{\frac{s^2 + 0.001s * 50/2\pi + (50/2\pi)^2}{s^2 + 0.1s * 50/2\pi + (50/2\pi)^2}}_{\text{notch for 50 Hz electrical back-coupling artefact}} \underbrace{\frac{1}{2\pi s/100 + 1}}_{\text{Low pass}}. \quad (19)$$

This controller is now used in closed-loop for studying system behaviour beyond flutter. Closed-loop identification experiments are performed in a similar manner and the transfer functions are obtained using PBSID, also shown in the Figures 13 and 14. In all identification experiments, the variance accounted for (VAF) Van der Veen et al. (2013) values exceed 60%. The dynamic behaviour can be seen to follow the predicted aeroelastic response from Figures 8 and 9. The frequency of the blade flexible mode remains more or less constant, Figure 13, however the damping goes on reducing until it is unstable at 340 rpm, as indicated by the sharp peak at 74 rad/s (11.8 Hz). On the other hand, Figure 14 shows that the frequency of the rigid-body mode increases with wind speed, along with the damping, as predicted in NASTRAN. Finally, system identification shows that the control authority of the flaps is low at low frequencies, but increases substantially at and beyond 2P (8Hz), making it suitable for reducing 2P loads and loads induced by turbulence.

6.2 Optimal IFT for constant wind speeds: Pre-flutter

The next step was to study the effect of the IFT load controllers for combined pitch and flap control. The block diagram for the load controllers is shown in Figure 12. Accordingly, the pitch and flap actuation signals were combinations of 1P and 2P sinusoidal basis functions, respectively. The basis functions are scheduled on the azimuth and are hence phase-locked. IFT was used to train the amplitudes of these basis functions; thus, with two basis functions for each frequency and each blade, for pitch and flap control both, a total of eight gains were required to be tuned.

The IFT process was first studied for a constant operational speed. Selected results, at an operational speed of 230 rpm, are presented here, although similar results were also observed throughout the operational range. The convergence of the controller

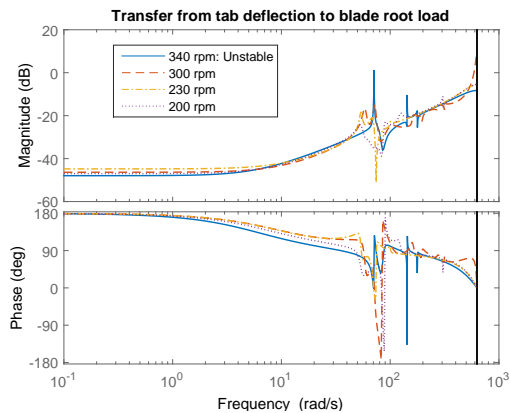


Figure 13. Transfer from piezobender actuators to blade root loads, different wind speeds.

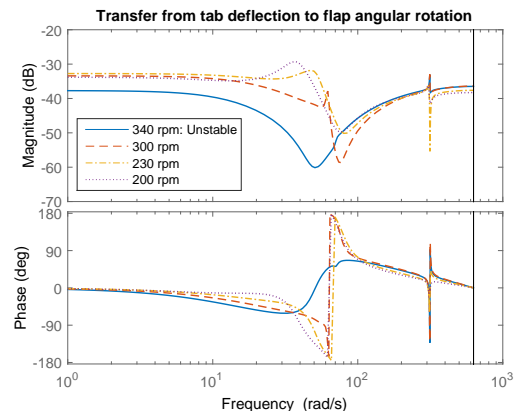


Figure 14. Transfer from piezobender actuators to free floating flap angle, different wind speeds.

gains and the IFT cost criterion can be seen in Figures 15 and 16. It can be seen that, within ten minutes, the controller gains converge to their optimal values. The performance of the controller after convergence can be visualised in Figures 17 and 18. The figures show that the actuation demanded, both pitch and flap, is purely sinusoidal, as constrained by the respective basis functions. Further, the load components in the blade load spectrum at the frequencies 1P and 2P are almost entirely eliminated by the pitch and flap action respectively. Thus, IFT is successful in tuning the controllers as required.

One final point of note is that the converged gains for the two blades are not exactly antisymmetric, this is especially pronounced for the flap actuation. The primary reason for this is a difference in the manufacture of the two blades. Specifically for the flap dynamics, for the scaled blade, a difference in the order of a few grammes in its weight distribution can strongly alter system dynamics and even prepone the onset of flutter.

Commercially manufactured blades are ideally expected to be identical; they would require antisymmetric pitch action and identical flap action for load attenuation, as produced by a conventional IPC controller Bossanyi (2003). Such a controller does not achieve optimal load reduction in the case there are discrepancies in blade manufacture or aging. The IFT controller designed above is thus shown capable of accounting for blade asymmetry and adjusting control action for maximising load reduction.

6.3 Optimal IFT for constant wind speeds: Post-flutter

Next, the free-floating flap is connected in series with the stabilising PID controller described above and the wind turbine is run at an operational speed of 330 rpm, in the post-flutter region. Since it is not optimally tuned, this controller does not maximise load reductions. Hence, IFT is used to tune the feedforward load reducing pitch and flap controller gains in a manner similar to the previous experiments; however in this case the underlying plant is the stabilised post-flutter wind turbine in closed-loop with the PID controller. From Figures 19 and 20 it can be seen that the optimised IFT controller gains are still




Figure 15. Convergence of controller gains over iterations

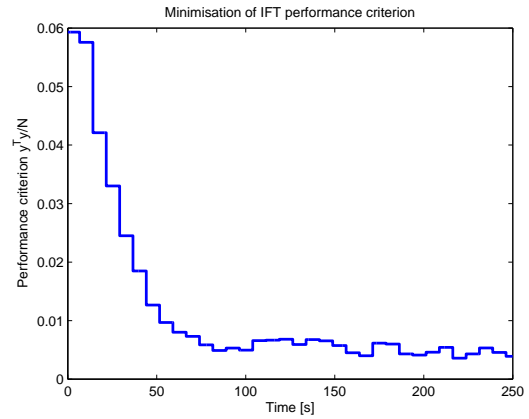


Figure 16. Minimisation of IFT cost criterion over iterations




Figure 17. Actuator duty cycles of optimised controller (Pre-flutter)




Figure 18. Load reductions achieved by optimised controller (Pre-flutter)

able to achieve load reduction even in this highly challenging unstable operational regime. The Figure 19 shows that the pitch controller no longer issues antisymmetric commands; a traditional IPC controller is no longer adequate in this region. Further, the flap command has already reached its maximum limits of +/- 500 V.

6.4 Optimal IFT gain-schedule for varying wind speeds

- For an in-field wind turbine, controller gain optimisation cannot be implemented considering the wind speed to be constant. Hence, the gain scheduling approach described in the previous section is followed, where, instead of the absolute values of the controller gains, the coefficients of a gain schedule, $\rho^{[0]}$ and $\rho^{[1]}$ are optimised based on the IFT experiments. This method is tested in the wind tunnel, with a varying operational speed profile as depicted in Figure 21. The convergence of the gain




Figure 19. Actuator duty cycles of optimised controller (Post-flutter)

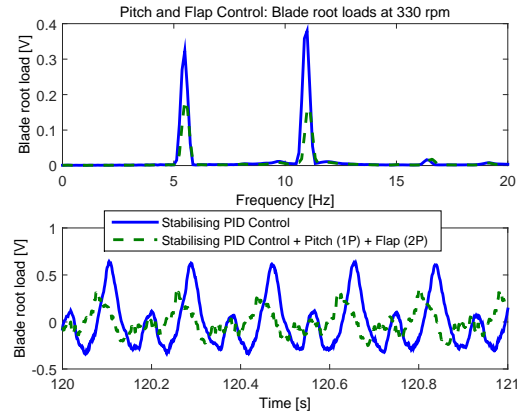


Figure 20. Load reductions achieved by optimised controller (Post-flutter)

schedule coefficients can be seen in Figures 22 and 23. It can be seen that in the first 100 seconds, since the wind speed is constant, a good gain schedule cannot be identified owing to a lack of persistency of excitation in the scheduling parameter. However, as the wind speed changes beyond this point in time, the gain schedule rapidly converges to an optimum. The gain schedule finally achieved is compared with the optimal controller gains obtained from the previous set of experiments in 5 Figures 24 and 25. It can be seen that for the pitch controller, the linear gain schedule obtained is a good fit to the optimal values obtained at constant wind speed. On the other hand, the flap controller optimal gains show a non-linear variation with wind speed and the linear gain schedule obtained achieves reduced goodness of fit.




Figure 21. Varying operational speed for optimisation of gain schedule.




Figure 22. Optimisation of gain schedule intercepts for varying wind speed conditions.




Figure 23. Optimisation of gain schedule slopes for varying wind speed conditions.




Figure 24. Gain schedule at varying wind speeds versus optimal gains at constant wind speed: Pitch Control.

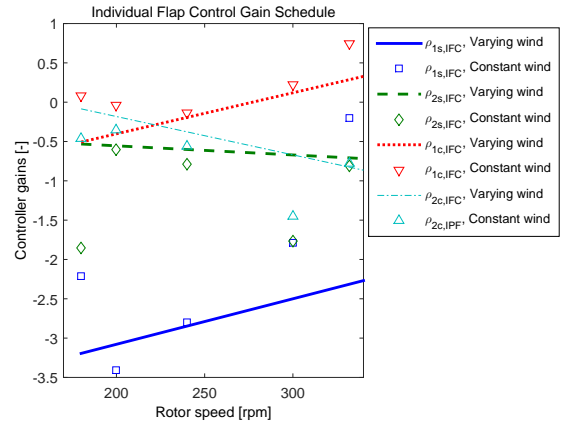


Figure 25. Gain schedule at varying wind speeds versus optimal gains at constant wind speed: Flap Control

Herewith, the combined pitch and flap controller has been shown to be able to reduce blade loads both in pre- and post-flutter conditions. Further, an optimal gain schedule for these controllers is automatically tuned online using IFT in varying wind conditions.

7 Conclusions

- 5 A successful experimental proof of concept has hereby been achieved for the first time of free-floating flaps for wind turbines and of combined pitch and flap control for blade load mitigation.

Free-floating flaps were designed for the first time for the application of wind turbine load control. Numerical aeroelastic analysis concluded that such flaps show significant control authority in the desired frequency band (2P and beyond). However the additional degree of freedom couples aerodynamically with the flapwise flexible mode of the blade and causes flutter at low wind speeds, just outside the design envelope. Using a feedback controller, the blade can be stabilised in the post-flutter region. Both of these results were validated experimentally in the wind tunnel.

Blades were manufactured using the novel combination of 3d-printing with carbon fibre layup, and instrumented with free-floating flaps close to the blade tip. The concept of Iterative Feedforward Tuning of the gains of phase-locked basis functions was used to achieve blade load reductions. The pitch control action was composed of a superposition of 1P sinusoidal basis functions, while for the flap control action, 2P sinusoidal basis functions were used. It was shown that, for a constant pre-flutter wind speed, ideal rejection of 1P and 2P loads in the blade load spectrum could be achieved with combined pitch and flap control. Further, at a post-flutter wind speed, the system was connected in closed-loop with a stabilising PID controller using collocated feedback. Iterative Feedforward Tuning was able to optimise, in this unstable regime, the load control gains of the pitch and flap control action for this closed-loop plant. Load rejection was also achieved in the challenging post-flutter regime, although the flap actuation duty reached close to its physical limits under these conditions. Finally, for the case of varying wind speed conditions, the IFT methodology was able to autonomously synthesise an optimal linear gain schedule, in real time, for the combined pitch and flap controller. Such a gain schedule was found to be near-optimal for a large portion of the range of operation.

The shortcomings of the IFT control approach are as follows:

- For the case where LTI controllers are devised for constant operating points, the control action for an intermediate speed is obtained by interpolating between the gains of the closest wind speeds. This is the approach followed by most industrial gain-scheduled controllers. For a highly non-linear plant, this approach is no longer optimal for the intermediate wind speeds, in common with such conventional gain-scheduled controllers.
- For the case where a gain-schedule is automatically tuned by IFT for varying wind speeds, the control is not optimal at any operating point, but it is globally optimal for a range of operating points. However, for the case (as with IFC), where the desired gain schedule is not linear, the controller may possibly behave poorly across the entire wind speed region. This case has more parallels with an LTI robust control design, which optimises globally but may be severely suboptimal for local operating points. An LPV robust control design may be superior in general, but accuracy of modelling is critical for such an approach.

With the inclusion of flap control, the individual pitch controller can focus purely on 1P load attenuation; this reduces pitch activity, especially in the high-frequency region of the spectrum, and can enhance the longevity of the pitch actuation mechanism.

With free-floating flap control with variable pitch wind turbines hereby validated, future work will primarily concern the LPV (linear parameter-varying) modelling and validation of the augmented turbine blade, and the synthesis of an optimal LPV

controller. Also, turbulent and gust load mitigation with local flap control would form the next step in the study of wind turbines with free-floating flaps.

Acknowledgements. This work was supported by the INNWIND.EU Project, an EU Consortium with Academic and Industrial Partnership for Innovations in Wind Energy.

- 5 We would like to thank ir. Kees Slinkman and ing. Will van Geest (Delft Centre for Systems and Control) for their help in the design and assembly of the experimental setup.

References

- Andersen PB, Gaunaa M, Bak C, Buhl T (2006). Load alleviation on wind turbine blades using variable aerofoil geometry. *Proc. of the EWEC, Brussels, Belgium*.
- Bak C (2013). Description of the DTU 10 MW reference turbine. *DTU Wind Energy Tech. Rep. I-0092*.
- 5 Barlas TK, Van Kuik GAM (2010). Review of the state of the art in smart rotor control research for wind turbines. *Progress in Aerospace Sciences* 46:1-27.
- Bauer S, Czakalla M, Göppele F, Grandy M, Hammer A, Hofmann L, Jenter T, Kaus S, Pullman T, Rahn A, Samol D, Schlesiger J, Schwenk J, Beyer F and Cheng PW(2014). Abschlussbericht Windenergieprojekt WS 13/14 - Windenergie aus dem 3D-Drucker. *Universität Stuttgart*.
- Bernhammer LO, Van Kuik GAM, De Breuker R (2012). How far is smart rotor research and what steps need to be taken to build a full-scale
10 prototype? *Journal of Physics: Conference Series* 555 (1), 012008.
- Bernhammer LO, De Breuker R, Karpel M, Van der Veen GJ (2013). Aeroelastic control using distributed floating flaps activated by piezo-electric tabs. *Journal of Aircraft* 50 (3): 732-740.
- Bernhammer LO, Sodja J, Navalkar ST, De Breuker R, Karpel M (2015). Experimental and numerical investigation of an autonomous flap for load alleviation. *Journal of Aircraft*, Article in Press.
- 15 Bernhammer LO, De Breuker R, Van Kuik GAM. Fatigue and extreme load reduction of wind turbine components using smart rotors. *Journal of Wind Engineering and Industrial Aerodynamics*, Article in Press.
- Bossanyi EA (2003). Individual blade pitch control for load reduction. *Wind Energy* 6 (2): 119-128.
- Bossanyi EA, Fleming PA, Wright AD (2013). Validation of individual blade pitch control by field tests on two- and three-bladed wind turbines. *IEEE Trans. Control Syst. Technol.* 21 (4):1067-1078.
- 20 Bottasso CL, Croce A, Riboldi CED, Nam Y (2013). Multi-layer control architecture for the reduction of deterministic and non-deterministic loads on wind turbines. *Renew. Energy* 51:159-169.
- Castaignet D, Barlas TK, Buhl T, Poulsen NK, Wedel-Heinen JJ, Olesen NA, Bak C, Kim T (2013). Full-scale test of trailing edge flaps on a Vestas V27 wind turbine: active load reduction and system identification. *Wind Energy* 17 (4):549-564.
- Dunne F, Pao LY, Wright AD, Jonkman B, Kelley N (2011). Adding feedforward blade pitch control to standard feedback controllers for
25 load mitigation in wind turbines. *Mechatronics* 682-690.
- Gevers M (2002). A decade of progress in iterative process control design: from theory to practice. *Journal of Process Control* 12:519-531.
- Heinze S, Karpel M (2006). Analysis and wind tunnel testing of a piezoelectric tab for aeroelastic control applications. *Journal of Aircraft* 43 (6):1799-1804.
- Hjalmarsson H (2002). Iterative feedback tuning - an overview. *Journal of Adaptive Control and Signal Processing* 16:373-395.
- 30 Hulskamp AW, Van Wingerden JW, Barlas TK, Champlaud H, Van Kuik GAM, Bersee HEN, Verhaegen M (2011). Design of a scaled wind turbine with a smart rotor for dynamic load control experiments. *Wind Energy* 14:339-354.
- Karutz R. Untersuchung Generativer Fertigungsverfahren im Rotorblattbau von Kleinwindkraftanlagen. *Universität Stuttgart Bachelor Thesis*.
- Lackner MA, Van Kuik GAM (2010). A comparison of smart rotor control approaches using trailing edge flaps and individual pitch control.
35 *Wind Energy* 13:117-134.
- Navalkar ST, Van Wingerden JW (2015). Iterative feedback tuning of an LPV feedforward controller for wind turbine load alleviation. *The 1st IFAC LPV Systems Conference, Grenoble, France*.

- Navalkar ST, Van Solingen E, Van Wingerden JW (2015). Subspace predictive repetitive control for variable pitch wind turbines. *IEEE Trans. Control Syst. Technol.* 23 (6):2101-2116.
- Selvam K, Kanev S, Van Wingerden JW, Van Engelen T, Verhaegen M (2009). Feedback-feedforward individual pitch control for wind turbine load reduction. *Int. J. Robust Nonlinear Control* 19 (1): 72-91.
- 5 Van der Veen GJ, Van Wingerden JW, Bergamasco M, Lovera M (2013). Closed-loop subspace identification methods: an overview. *Control Theory and Applications IET* 7 (10): 1339-1358.
- Van Wingerden JW, Hulskamp AW, Barlas TK, Houtzager I, Bersee HEN, Van Kuik GAM, Verhaegen M (2011). Two-degree-of-freedom active vibration control of a prototyped ‘smart’ rotor. *IEEE Trans. Control Syst. Technol.* 19:284-296.
- Van Wingerden JW (2008). Control of wind turbines with ‘smart’ rotors: proof of concept and LPV subspace identification. *PhD Thesis,*
10 *Delft University of Technology.*