



Comparison of Two Independent LIDAR-Based Pitch Control Designs

F. Dunne
*University of Colorado
Boulder, Colorado*

D. Schlipf
*University of Stuttgart
Stuttgart, Germany*

L.Y. Pao
*University of Colorado
Boulder, Colorado*

NREL Technical Monitor: Alan Wright

NREL is a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy, operated by the Alliance for Sustainable Energy, LLC.

Subcontract Report
NREL/SR-5000-55544
August 2012

Contract No. DE-AC36-08GO28308

Comparison of Two Independent LIDAR-Based Pitch Control Designs

F. Dunne
*University of Colorado
Boulder, Colorado*

D. Schlipf
*University of Stuttgart
Baden-Württemberg, Germany*

L.Y. Pao
*University of Colorado
Boulder, Colorado*

NREL Technical Monitor: Alan Wright
Prepared under Subcontract No. XEE-1-11300-01

NREL is a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy, operated by the Alliance for Sustainable Energy, LLC.

This publication received minimal editorial review at NREL.

NOTICE

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

Available electronically at <http://www.osti.gov/bridge>

Available for a processing fee to U.S. Department of Energy
and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: <mailto:reports@adonis.osti.gov>

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: <http://www.ntis.gov/help/ordermethods.aspx>

Cover Photos: (left to right) PIX 16416, PIX 17423, PIX 16560, PIX 17613, PIX 17436, PIX 17721



Printed on paper containing at least 50% wastepaper, including 10% post consumer waste.

Acknowledgments

This work was supported, in part, by the National Renewable Energy Laboratory (NREL) and the German Federal Ministry for the Environment, Nature Conservation, and Nuclear Safety.

Nomenclature

CP	collective pitch
IP	individual pitch
1P	once-per-revolution (of the rotor) (frequency) (0.2 Hz)
CW	continuous wave
DEL	damage equivalent load
DOF	degree of freedom
FIR	finite impulse response
LIDAR	Light Detection and Ranging
LPF	lowpass filter
MBC	multiblade coordinate
PI	proportional-integral
RMS	root mean square
	horizontal average wind speed
	horizontal linear wind shear
	vertical linear wind shear
	wind speed at blade i , $i = 1, 2$, or 3
u_0	average feedforward signal
u_H	horizontal MBC component of feedforward signal
u_V	vertical MBC component of feedforward signal
u_i	feedforward signal for blade i , $i = 1, 2$, or 3
	MBC transformation matrix
	inverse MBC transformation matrix

Executive Summary

Two different Light Detection and Ranging (LIDAR)-based feedforward controllers have previously been designed for the National Renewable Energy Laboratory's (NREL's) 5-MW reference wind turbine model under separate studies. Feedforward controller A uses a finite-impulse-response design, with 5 seconds of preview, and three rotating LIDAR measurements. Feedforward controller B uses a static-gain design, with the preview time defined by the pitch actuator dynamics, a simulation of a real nacelle-based scanning LIDAR system, and a lowpass filter defined by the LIDAR configuration. These controllers are now directly compared under the same LIDAR configuration, in terms of fatigue load reduction, rotor speed regulation, and power capture. The various differences in design choices are discussed and compared in this report. The authors also compare frequency plots of individual pitch feedforward and collective pitch feedforward load reductions, and results reveal that individual pitch feedforward is effective mainly at the once-per-revolution and twice-per-revolution frequencies. The authors also explain how to determine the required preview time by breaking it down into separate parts, and then comparing it to the expected preview time available.

Table of Contents

1	Introduction.....	6
2	Simulated Turbine and Turbulent Inflow	7
2.1	5-Megawatt Turbine Model and Baseline Control	7
2.2	Stochastic Turbulent Wind Field Simulator	7
3	Controller Comparison.....	9
4	Individual Pitch Feedforward Control	22
5	Preview Times	28
5.1	Actuator.....	28
5.2	Lowpass Filter.....	29
5.3	Turbine.....	30
6	Conclusions and Future Work	34
	References.....	35

List of Figures

Figure 1.	Feedforward control added to feedback control.....	6
Figure 2.	Impulse response of feedforward controller A.....	10
Figure 3.	Lookup table from steady-state wind speed to steady-state pitch angle	11
Figure 4.	Turbine loads method AAA	12
Figure 5.	LIDAR configuration A continuous-wave LIDAR scanning pattern.....	13
Figure 6.	LIDAR configuration B pulsed LIDAR scanning pattern.....	13
Figure 7.	Power spectral densities for method BBB and _B_, original simulation conditions..	16
Figure 8.	Turbine loads method AAA versus BAA, wind field AR2_s13.	17
Figure 9.	Turbine loads method AAA versus BAA, wind field 14 m/s.....	18
Figure 10.	Bode plot of feedforward controllers A and B, excluding scheduled gains.....	19
Figure 11.	Time delay of feedforward controllers A and B.....	20
Figure 12.	Tower top fore-aft (pitching) moment as a function of frequency..	23
Figure 13.	Blade root bending moment as a function of frequency.....	23
Figure 14.	Fore-aft nacelle acceleration as a function of frequency.....	24
Figure 15.	Blade pitch angle as a function of frequency	25
Figure 16.	Generator torque as a function of frequency	26
Figure 17.	Generator speed as a function of frequency	26
Figure 18.	Cumulative required preview times for the turbine, actuator, and lowpass filter.	28
Figure 19.	Phase of second-order actuator model.....	29
Figure 20.	Time delay of second-order actuator model.....	29
Figure 21.	Estimated preview required to cancel the delay of lowpass filtering.....	30
Figure 22.	Turbine preview, all available degrees of freedom	31
Figure 23.	Turbine preview, 5 degrees of freedom.....	32
Figure 24.	Turbine preview, spectra estimation method.....	33

List of Tables

Table 1. TurbSim Boundary Conditions	8
Table 2. Method Notation	9
Table 3. Features of Feedforward Controllers A and B	9
Table 4. Features of LIDAR Configurations A and B	10

1 Introduction

Wind turbine fatigue loads can be reduced with well-designed blade pitch control at above-rated wind speeds. The control is typically feedback-only, and relies on turbine sensors. However, LIDAR (Light Detection and Ranging) can be used to measure wind speeds that are upstream of a wind turbine. This preview of wind speeds can be used to improve blade pitch control, thereby reducing loads that otherwise would not be sensed until they have already affected the turbine.

Two different methods [1, 2] for implementing LIDAR-based pitch control were previously designed under separate studies. Both are feedforward controllers that are intended to be added on to a standard feedback controller as shown in Figure 1. Many design features differ between the two methods, and the main goal of this study is to determine which features are the most important in achieving good load reduction, and which choices should be made for each feature. Two additional goals are to find the frequencies at which individual pitch feedforward is effective, and to explain the contributing factors of required preview time.

In this report, Section 2 describes the turbine model, baseline controller, and simulation conditions. Section 3 describes and compares the two controllers. Section 4 explores details on individual pitch feedforward control, and Section 5 investigates the choice of preview time. Finally, Section 6 outlines the conclusions and future work.

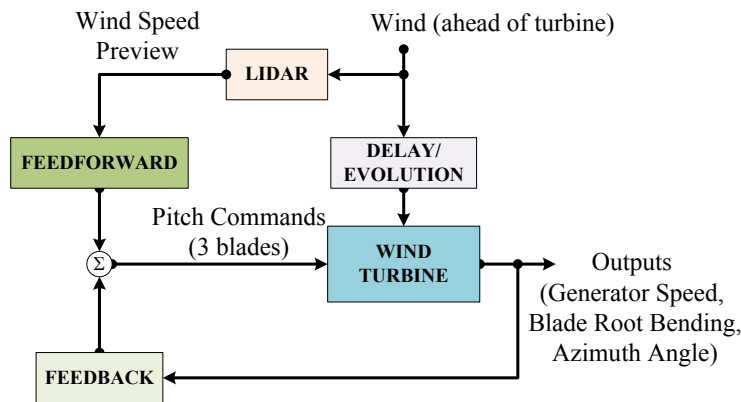


Figure 1. Feedforward control added to feedback control

2 Simulated Turbine and Turbulent Inflow

2.1 5-MW Turbine Model and Baseline Control

Simulations are performed using a full non-linear turbine model provided by the Fatigue, Aerodynamics, Structures, and Turbulence (FAST) [3] software code developed at the National Renewable Energy Laboratory (NREL). The particular model used in FAST is the NREL 5-MW reference turbine [4]. All available degrees of freedom (DOF) are turned on in simulations. For all simulations, a second-order pitch actuator model has been added to the 5-megawatt (MW) turbine model, with a natural frequency of 1 hertz and a damping ratio of 0.7.

The standard collective-pitch feedback controller for this turbine is a gain-scheduled proportional-integral (PI) control [4]. An individual pitch (IP) feedback-only controller was also designed, [5] and used as the baseline controller for this study. In addition to generator speed, the controller inputs are the three out-of-plane blade root bending moments, and the rotor azimuth, which is used for the multi-blade coordinate [6] (or d-q axis) transformation (for the IP feedback controller). This transformation converts the three blade signals from a rotating frame into a stationary frame. In the fixed stationary frame, the three signals are called the collective (average), horizontal, and vertical components. The horizontal and vertical components are controlled with PI controllers, and the collective component is controlled with the same PI feedback control as the standard collective-pitch controller.

This individual pitch baseline feedback controller is used in all simulations described in this study, sometimes with and sometimes without added feedforward control. This allows the team to investigate the possible improvements in load mitigation by adding a feedforward control element to a typical feedback controller.

2.2 Stochastic Turbulent Wind Field Simulator

The NREL TurbSim [7, 8] stochastic full-field inflow simulator is used to provide realistic wind fields for the turbine simulations. Most of the simulations described below are based on extensive observations taken in the high-plains environment of Southeast Colorado that now has a large operating wind farm. The Great Plains (GP-LLJ) spectral model available in TurbSim is used to simulate the wind conditions that are present at this site.

The boundary conditions for the TurbSim simulator are shown in Table 1. These values are derived from the averages of subpopulations (e.g., AR1) of actual measured wind conditions that are associated with 13 meters per second (m/s), (above-rated) hub-height mean wind speeds. The (Y-Z) grid encompassing the turbine rotor disk contains 31×31 points of three orthogonal wind components, with a sample rate of 20 hertz (Hz) and a total record length of 630 seconds. The first 99 seconds of data from each simulation are discarded before calculating any performance measures to allow initial transients to settle out. Thirty-one different realizations of each subpopulation were created, each 630 seconds long.

In addition to the wind fields shown here, the authors also mention below the use of a 14 m/s average, Class A wind field. This field is also 630 seconds long and was created using TurbSim, but has a greater turbulence intensity (18%), and is the International Electrotechnical Commission's Kaimal NTM spectral model with a Class A turbulence level.

**Table 1. TurbSim Boundary Conditions for 90-m Hub Height, 5-MW Turbine, Great Plains (Lamar)
Inflow Simulations**

Ensemble Simulation ID	u_{hub}^a (meters/second)	Ri_{TL}^b	α_D^c	u_*^d (m/s)	Coherent Structures?	Jet Height (m)
AR1	13	-0.10	0.077	0.514	N	no jet - pwr law
AR2	13	0.02	0.139	0.422	N	no jet - pwr law
AR3	13	0.20	0.363	0.135	N	no jet - pwr law
AR4	13	0.02	n.a. ^e	0.289	Y	90
AR5	13	0.20	n.a.	0.160	Y	90

^a hub-height mean wind speed

^b vertical stability parameter

^c vertical power law shear exponent

^d mean friction velocity (shearing stress) over the rotor disk

^e jet instead of power law

3 Controller Comparison

In this study, the authors compare two previously developed feedforward control methods, referred to as feedforward controller A [1] and feedforward controller B [2]. Each feedforward controller was designed for use in combination with its own individual pitch feedback controller. Both individual pitch feedback controllers are very similar, and both follow the description above in Section 2.1, but there are slight differences in implementation between the two, because they were implemented independently in Simulink by two different researchers in separate studies. For example, the feedback gains are different, and while they use the same pitch rate limit, they differ slightly in the way the rate limit is implemented (referred to as feedback controllers A [1] and B [2]). The two feedforward controllers were also each developed with their own LIDAR (Light Detection and Ranging) configuration, which are referred to as LIDAR configuration A [1, 9] and LIDAR configuration B [2].

To refer to different combinations of feedforward controller, feedback controller, and LIDAR configuration, ‘method xyz’ will be used, as shown in Table 2. A summary of the differences between feedforward controllers is shown in Table 3, and a summary of the differences between LIDAR configurations is shown in Table 4.

Table 2. Method Notation

Method xyz	
x	Feedforward controller
y	Feedback controller
z	LIDAR configuration

Notes: x, y, and z can each be either A, B, or __ (none). For instance, __A__ means feedback controller A only (without any feedforward controller nor any LIDAR configuration).


Table 3. Features of Feedforward Controllers A and B

Feature	Feedforward Controller A	Feedforward Controller B
Controller Design	Finite impulse response (FIR) + Wind-to-pitch lookup	Lowpass filter (LPF) + Wind-to-pitch lookup
Individual Pitch Control?	Yes (individual)	Yes (cyclic)

Table 4. Features of LIDAR Configurations A and B

Feature	LIDAR Configuration A	LIDAR Configuration B
LIDAR Type	Continuous wave	Pulsed
LIDAR Sample Rate [hertz (Hz)]	80	5
Measurement Locations	Three points, one ahead of each blade at 75% span	Five circles, 12 points each
Measurement Distance meters (m)	65	58 to 174
Preview Time Used seconds (s)	5	Varies
Convert Measurements to Mean, Horizontal, and Vertical Shear?	No	Yes
Assume Perfect Alignment w/Wind?	Yes	Yes

Feedforward controller A [1] uses a finite impulse response (FIR) design, with 5 seconds of preview. Its FIR filter coefficients were originally chosen heuristically. They were then optimized by using a genetic algorithm, trying thousands of variations, and converging on the set of coefficients with the best performance. Performance was based on fatigue load reduction (blade and tower damage equivalent loads and nacelle accelerations), root-mean-square pitch rate, peak rotor speed, and average power achieved in a simulation of the nonlinear turbine, with all degrees of freedom, in above-rated wind conditions. The impulse response of feedforward controller A is shown in Figure 2.

**Figure 2. Impulse response of feedforward controller A**

For input to feedforward controller A, three rotating, hub-mounted continuous-wave (CW) LIDAR measurements [9] are taken 65 meters (m) ahead of the turbine, one measurement ahead

of each blade, at about 75% span, as shown in Figure 5. The LIDAR model includes the effects of line-of-sight, volume-averaged wind speed measurements. The researchers refer to this wind measurement scheme as LIDAR configuration A. LIDAR configuration A does not model measurement obstruction by moving blades because it is assumed to be hub-mounted. The three blades' feedforward signals are controlled separately, each using its own wind speed measurement. Wind speed measurements are sampled at 80 Hz, matching the simulation rate for convenience. However, the higher frequencies are not used, since the FIR filter acts like a lowpass filter. Each measurement is separately filtered by the FIR filter described above, which has a DC gain of one. In series with each FIR filter is a lookup table from steady-state wind speed to blade pitch, shown in Figure 3. Results of this design (method AAA) are shown in Figure 4. Results are also shown here for a collective pitch feedforward variation, where the three wind speed measurements are averaged together, and the same average is fed into each blade's identical feedforward control channel.




Figure 3. Lookup table from steady-state wind speed to steady-state pitch angle, used for gain scheduling




Figure 4. Turbine loads using individual pitch baseline control alone [5] (method _A_), and with added feedforward controller A (method AAA), in individual pitch and collective pitch versions [1]

Notes: Percentages displayed are the average of the eight bars. Feedforward controller A reduces overall loads by 4.9% in the individual pitch version. The results are averages from FAST simulations across 155 wind files (representing AR1 through AR5).

Feedforward controller B uses a pulsed LIDAR model, sweeping a circle in 2.4 seconds, with 12 points at each of five different distances, as shown in Figure 6. This trajectory has been implemented with a LIDAR system developed and installed on the nacelle of a 5-megawatt turbine [10]. To obtain realistic measurements in the simulation, effects such as obstruction of the laser beam by the blades, line-of-sight volume-averaged measurements, and mechanical constraints of the scanner from real experimental data were considered. For instance, the same loss of about 30% of points could be observed in the simulation and in the measurements, caused by obstruction from the moving blades.



Figure 5. LIDAR configuration A continuous-wave LIDAR scanning pattern provides a preview (65 meters ahead of the turbine) LIDAR measurement of the horizontal wind speed at 75% span for each blade




Figure 6. LIDAR configuration B pulsed LIDAR scanning pattern

These LIDAR measurements are then reduced to three components:

(1)

where

= the horizontal average wind speed

and = the horizontal and vertical shear, respectively.

These components are found by using a least squares method on the past 12 measurements (the past full circle). This wind measurement scheme is referred to as LIDAR configuration B.

LIDAR configuration B is more realistic than LIDAR configuration A because configuration B is modeled on a system that has actually been implemented with a single LIDAR. Configuration A, however, would require three LIDARs, and it assumes that it is known where the blades will be a few seconds in advance. LIDAR configuration B also appears to have a higher bandwidth for providing accurate, rotor-effective wind speed measurements. Measurements from LIDAR configuration B match the rotor-effective wind speed felt by the turbine up to a spatial cutoff frequency of 0.06 radians per meter (rad/m). (Frequency in time is divided by wind speed in meters per second (m/s) to obtain spatial frequency). Using LIDAR configuration A, this spatial cutoff frequency is reduced to 0.04 rad/m. LIDAR configuration B likely provides a more accurate wind speed preview because measurements are taken at five different radii instead of just one. This was originally done with pulsed LIDAR, with all five measurements taken at the same time, as shown in Figure 6. A continuous-wave (CW) LIDAR could produce equivalent measurements by refocusing between the five points, one after the other, if the CW LIDAR's sample rate was at least five times that of the pulsed LIDAR.

LIDAR configuration A detects the blade-effective wind speeds that are needed by feedforward controller A. On the contrary, LIDAR configuration B is trying to capture the rotor-effective wind characteristics that are required by feedforward controller B. In this report, “blade-effective wind speeds” means one value for each blade, depending only on the sections of the flow that will hit the blades. In contrast, “rotor-effective wind characteristics” means the average wind speed and vertical and horizontal wind shears that are obtained from measurements covering the entire rotor plane, regardless of where the blades are. The correlation of one single measurement ahead of each blade with each blade's effective wind speed should be greater than the correlation of those three point measurements with the effective wind characteristics of the whole rotor disk. Therefore, feedforward controller A, combined with LIDAR configuration A, can compensate loads above the once-per-revolution (1P) frequency, if a good correlation with the blade-effective wind speed can be obtained at frequencies above 1P. To capture a rotor-effective wind speed correlation comparable to LIDAR configuration B, more measurement points are necessary—possibly up to five times more—since configuration B is designed with five measurements per azimuth and configuration A is designed with just one measurement per azimuth. LIDAR configuration B was designed, through choice of measurement points and processing, to improve the correlation with the rotor-effective wind characteristics. Nonetheless, using LIDAR configuration A also shows a good correlation for the rotor-effective wind characteristics.

LIDAR configuration A has the advantage that it retains more information on how the wind speed varies with azimuth. Blade loads can be caused by both wind speed differences over azimuth and wind speed differences in the x-direction (up/downwind). For changes in the x-direction, the authors estimate that the bandwidth for good measurements is around 0.1 or 0.2 Hz (at 13 m/s). For changes over azimuth, however, the bandwidth depends on the number of LIDAR measurements per LIDAR revolution. This could easily translate to higher than 0.1 or

0.2 Hz as the blade sees it. For example, a low-level jet can cause a high wind speed at hub height and lower wind speeds at both the top and bottom of the rotor plane. As the blades spin through this, they see a twice-per-revolution (2P), or 0.4 Hz load. Using LIDAR configuration A, as shown below, the researchers do use 2P measurements to reduce these 2P loads. LIDAR configuration B takes enough measurements to capture this low-level jet, but the information is lost when these measurements are simplified into the 1P components [6], which consist of the average wind speed and vertical and horizontal linear shears, as in [1].

Using only the 1P components of wind measurements, as in LIDAR configuration B, can be called cyclic feedforward. The advantage of cyclic feedforward is that it can be used with different LIDAR types and scan patterns, and is more independent of the preview time/scan distance. It should work for all wind speeds and can easily be adapted in real applications. To combine the advantages of configurations A and B, cyclic feedforward could be modified to reduce 2P loads in addition to 1P loads. This would involve transforming the ring of measurements into 2P components in addition to 1P components.¹

After capturing the wind measurements, feedforward controller B [2] first delays them so that the remaining preview time is just enough to compensate for the phase delay of the actuator, lowpass filter, and the turbine. This ensures that the blade pitch is actuated with the correct timing. The measurements are then sent to the feedforward controller, which is simply a lowpass filter followed by a set of static gains. The filter cutoff frequency is determined based on the correlation between the LIDAR measurements and the turbine-effective wind speed. Then for , there is a lookup table from steady-state wind speed to steady-state pitch to get u_0 , the average feedforward signal. The horizontal and vertical wind shears, and , are each multiplied by a scalar that is optimized to yield the best u_H and u_V blade pitch control components that cancel the wind disturbance. The feedforward blade pitch components u_0 , u_H , and u_V are added to the feedback controller's components in the stationary frame, and the sum is transformed back to the rotating frame through an inverse multiblade coordinate (MBC) transformation. This simpler design has the benefit of being easily tunable, which is very important when dealing with modeling uncertainty.

Originally, simulation for method BBB was done with a stochastic full-field wind (23×23 grid, $\Delta t = 0.25$ s) with a mean velocity of 16 m/s and a turbulence intensity of 18%. The results, which are provided in Figure 7, show greatly reduced tower and blade loads. When originally studied separately, with different wind fields, different performance metrics, and slightly different baseline feedback controllers, method BBB versus method _B_ appeared to have significantly greater load reduction than method AAA versus method _A_, which is shown in Figure 4 and Figure 8.

¹Both configurations A and B have enough measurement points to capture twice-per-revolution (2P) components. If a ring of measurements is used, as in configuration B, the Shannon sampling theorem requires measurements at more than four different azimuth angles to capture 2P components. If the blade locations are being tracked, as in configuration A, only three measurement points, one for each blade, are necessary to capture the 2P components, along with all other frequencies up to half the sampling rate. However, yaw errors, which are assumed to be zero in this study, may make it difficult to accurately capture 2P components, regardless of the number of measurements taken.

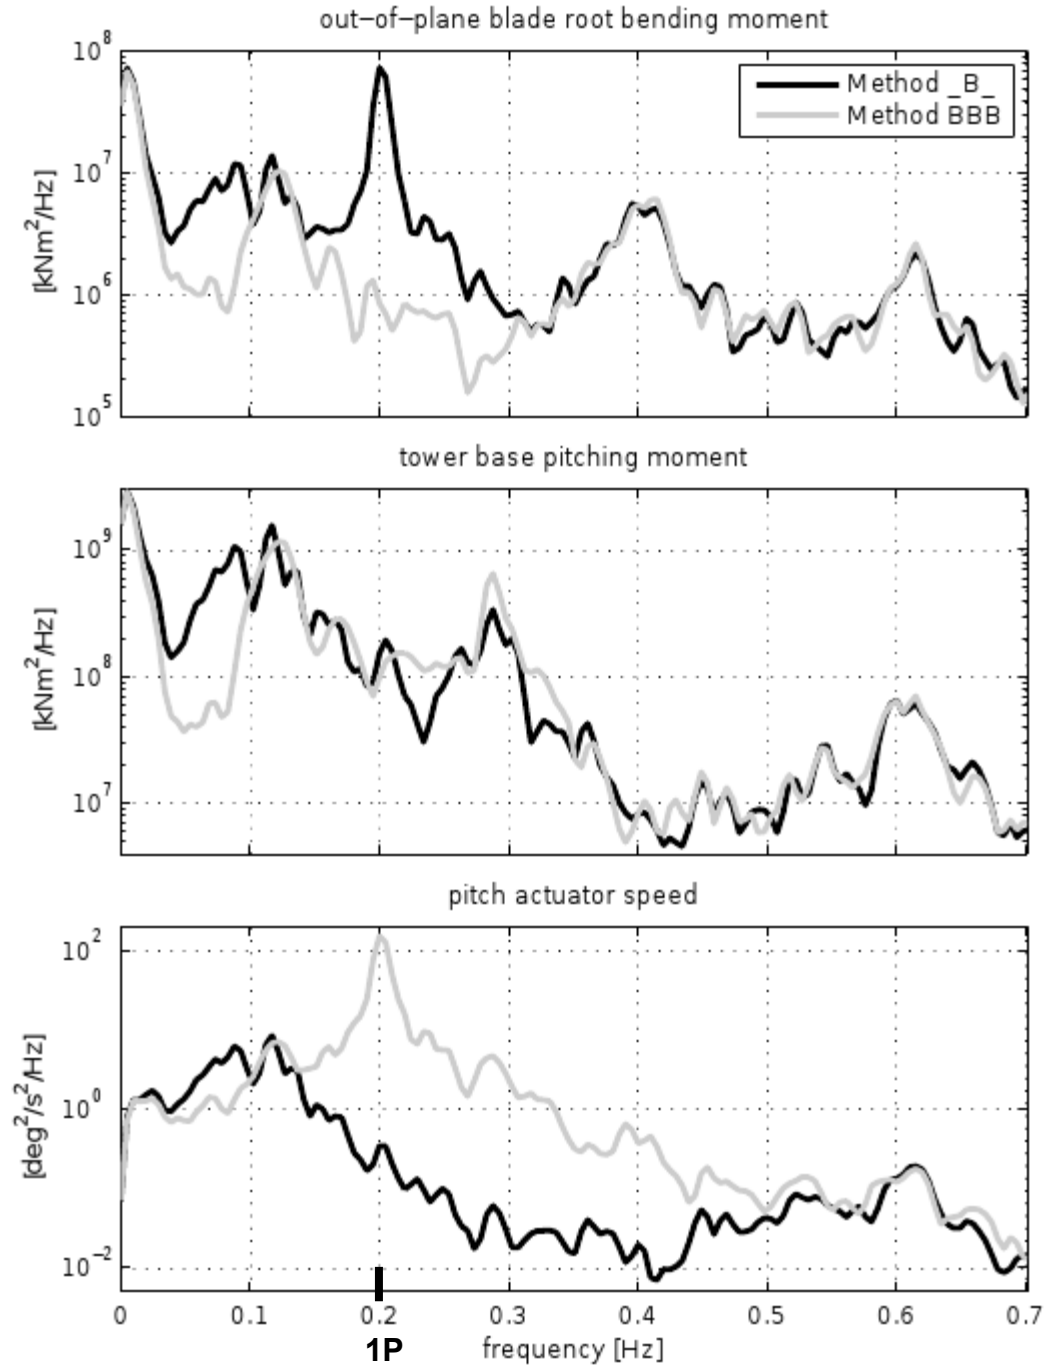


Figure 7. Power spectral densities for method BBB and _B_ under original simulation conditions. The once-per-revolution (1P) frequency is 0.2 Hz.

To compare the two feedforward controllers only, the researchers simulated method AAA and method BAA in the same wind fields, and using the same performance metrics. Results are shown in Figure 8 for a 13 m/s, less turbulent wind field, and in Figure 9 for a 14 m/s more

turbulent wind field. Both feedforward controllers perform similarly, with feedforward controller A showing slightly more load reduction in the 14 m/s wind, and feedforward controller B showing slightly more load reduction in the 13 m/s wind. The more turbulent 14 m/s wind allows for more load reduction, with both controllers averaging about 9% overall, versus 5% in the 13 m/s wind. Blade root and tower base load reduction were the original performance metrics of feedforward controller B, and when looked at individually, these two measures have consistently greater load reduction than the average of all eight bars. Because performance for feedforward controller A was originally measured using this eight-bar average, this accounts for some of the originally perceived differences in the performance of the two controllers. Both controllers improve rotor speed regulation compared to the baseline, and, in the more turbulent 14 m/s wind, they also improve power capture, as shown in Figure 9. Feedforward controller B has a lower root-mean-square (rms) pitch rate than A but does not regulate rotor speed quite as well as A, as shown in Figure 8 and Figure 9. Feedforward controller A reduces tower top damage-equivalent load (DEL) and nacelle fore-aft acceleration (y-accel) more than feedforward controller B. As will be shown in Figure 12 below, much of feedforward controller A's tower top DEL reduction is at 0.6 Hz, which translates to 0.4 Hz in the rotating (blade) coordinate system. Therefore, the better tower top load reduction by feedforward controller A may be due to the increase in its magnitude at 0.4 Hz, which is shown in Figure 10 and discussed further below.




Figure 8. Turbine loads using individual pitch baseline control (A) alone, [5] with added individual pitch feedforward controller A, and with added individual pitch feedforward controller B. The rated rotor speed is 12.1 revolutions per minute (rpm), and the rated power is 5,000.0 kilowatts (kW). X-accel, Y-accel, and Z-accel are respectively fore-aft, side-to-side, and up-down nacelle accelerations. The wind field used is AR2_s13, from the Great Plains set, at 13 m/s average wind speed.




Figure 9. Turbine loads using individual pitch baseline control (A) alone, [5] with added individual pitch feedforward controller A, and with added individual pitch feedforward controller B. The rated rotor speed is 12.1 rpm, and the rated power is 5,000.0 kW. X-accel, Y-accel, and Z-accel are respectively fore-aft, side-to-side, and up-down nacelle accelerations. The wind field used is the IEC Kaimal NTM spectral model with a Class A turbulence level, at 14 m/s average wind speed.

This same pitch-rate/rotor-speed tradeoff that is appearing for feedforward control also appeared when comparing feedback controllers A and B. B was designed to use reduced feedback gains when combined with a feedforward controller. This greatly reduced rms pitch rate, and somewhat increased peak rotor speed. In some cases, the reduced feedback gains also led to reduced loads. Overall, this feedback gain reduction appears beneficial because it reduces pitching action and allows the feedforward controller more control authority.

Figure 10 shows a Bode plot of feedforward controllers A and B. Feedforward controller A was designed for 13 m/s average wind speed. Feedforward controller B was designed for the full range of region three wind speeds, and its cutoff frequency and preview time vary with wind speed. Here, the design for the 13 m/s average wind speed is shown. The two controllers, while designed using very different methods, look strikingly similar. Both designs drop to 50% magnitude at about 0.1 Hz. This has been roughly estimated to be the highest frequency at which it can be assumed that the wind measured with LIDAR matches the wind speed that is seen at the turbine, at an average wind speed of 13 m/s [11, 12]. In addition to the frequency responses

shown in Figure 10, each controller also contains a lookup table with gains that vary depending on wind speed. Both tables are based on steady-state wind-to-pitch gains.




Figure 10. Bode plot of feedforward controllers A and B, excluding scheduled gains

Figure 11 shows the controllers' phase converted to time delay. The phase and time delay matter mostly when the magnitude Bode plot is high, and here the controllers differ somewhat. Below 0.1 Hz, controller A has a time delay of about 5 seconds. Controller A is also designed to receive its wind preview input 5 seconds in advance, so all of the available preview is used by the controller. Controller B, on the other hand, has a time delay of about 2.9 seconds, at frequencies below 0.1 Hz. Controller B is designed to receive its wind preview input 3.9 seconds in advance, so an extra 1.0 seconds is still available to compensate for the additional delays caused by the actuator and turbine. This time value is also tunable in B.




Figure 11. Time delay of feedforward controllers A and B
time delay = $-(\text{phase from Figure 10})/(360 \text{ degrees})/(\text{frequency})$

Feedforward controllers A and B also differ in the coordinate system in which they are applied. Controller A is applied in individual blade coordinates as follows:

Here, u_i is the pitch feedforward command to blade i , where $i=1, 2$, or 3 . A represents the transfer function plotted in Figure 10, and u_i is the wind speed at blade i . Controller B is instead applied in MBC coordinates as follows:

Here, u is the average, u_h is the horizontal component, and u_v is the vertical component of feedforward pitch command to the blades. B represents the transfer function plotted in Figure 10, and u is the average, u_h is the horizontal shear component, and u_v is the vertical shear component of wind speed at the blades. These two methods (applying the controller in individual versus MBC coordinates) are not equivalent for the vertical and horizontal components of the

signal, because frequencies change between the rotating and nonrotating frames. The two methods are equivalent only for the collective component (average) of the signal. This is shown in the equations below,

where

, T = the MBC transformation matrix
 , and T^{-1} = the inverse MBC transformation matrix.

if C_{MBC} (controller applied in MBC coordinates)

this does not imply that C_{ind} (controller applied in individual coordinates)

however, it does imply that $\text{mean}(C_{MBC}) = A \text{mean}(C_{ind})$

because $\text{mean}(C_{MBC}) = \frac{1}{N} \sum C_{MBC} = \frac{1}{N} \sum A C_{ind} = A \text{mean}(C_{ind})$.

So the fact that feedforward controllers A and B are applied in different coordinate systems makes a difference to the individual pitch command variations, but it does not make a difference to the average pitch command.

To summarize this overall comparison, feedforward controllers A and B, while designed using very different methods, end up with very similar Bode plots and simulation results. Feedforward controller B is simpler to design and more tunable, which is beneficial because there will always be modeling uncertainty. It is a good idea to explore using feedforward at frequencies higher than 1P, as was attempted in controller A, to lead to even better performance. This higher frequency operation may be the reason for controller A's greater load reduction in the tower top and fore-aft acceleration. Reducing feedback gains when feedforward is present is a good method to reduce pitch rate, but there is some tradeoff in the form of increasing rotor speed error.

LIDAR configuration B is more realistic and provides more accurate wind speed estimates, because it takes more measurements that cover more of the rotor plane. Fitting measurements to average speed, vertical shear, and horizontal shear is convenient in converting a large set of LIDAR measurements to usable controller inputs, but it loses some useful spatial information as in the example of the low-level jet. In more turbulent wind fields, the addition of feedforward control provides a greater percent load reduction than in less turbulent wind. Finally, blade root and tower base load reduction is often greater than load reduction in other areas of the turbine structure, so one should be careful to compare controllers under the same performance metrics.

4 Individual Pitch Feedforward Control

When wind speeds are constant in time but varying over the rotor plane, loads appear on the blades only at the once-per-revolution frequency (1P) (0.2 Hz, which corresponds to a rotor speed of 12.1 rpm) and at its harmonics [twice-per-revolution (2P), three-times-per-revolution (3P), and so on]. It was therefore theorized that individual pitch (IP) control is only useful at these distinct frequencies. To test this theory, the researchers look at loads as a function of frequency for both collective pitch (CP) and IP configurations of method AAA described above.

Figure 4 shows that feedforward controller A reduces the blade root damage-equivalent load (DEL), tower top DEL, and root-mean-square x-acceleration (nacelle fore-aft acceleration) in the IP configuration when compared to the CP configuration. These three loads are plotted as a function of frequency in Figure 12, Figure 13, and Figure 14. On the non-rotating parts of the structure, the 1P, 2P, and 3P blade loads described above respectively transform to loads at 0P (DC), 3P, and none. Figure 12 shows that tower top fore-aft loads are reduced by the IP feedforward at below 0.1 Hz, increased from 0.15 to 0.3 Hz, and reduced again between 0.4 and 0.7 Hz, when compared to CP feedforward. This does fit with load reduction near 0P and 3P, as expected. Figure 13 shows that the IP action slightly reduces blade root loads from 0.23 to 0.45 Hz. This includes 2P, but not 1P and 3P. The IP feedback controller is designed to work at 1P. This explains why 1P loads are lower than 2P loads. It also may explain why IP feedforward does not help at 1P: the IP feedback may have done it so well that there is nothing left for IP feedforward to do. Finally, Figure 14 shows that IP feedforward makes fore-aft nacelle acceleration worse from 0.2 to 0.3 Hz, and very slightly better from 0.32 to 0.65 Hz. This follows a similar pattern to tower top loads, except for the part near DC, since a DC fore-aft nacelle acceleration does not make much sense for a turbine that does not move away from its foundation.

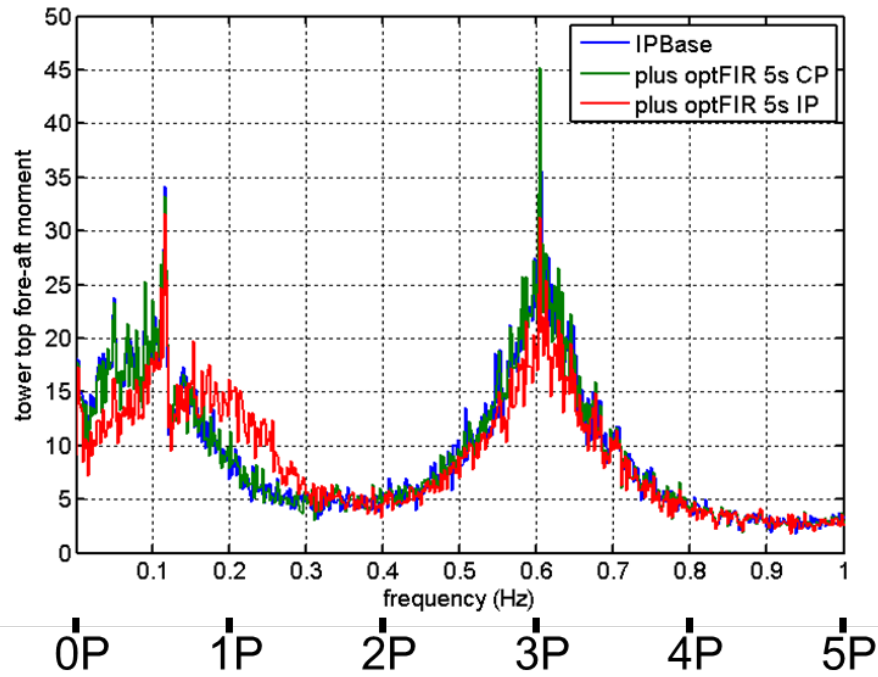


Figure 12. Tower top fore-aft (pitching) moment as a function of frequency. Average fast Fourier transform (FFT) magnitude over AR1_s12 through AR1_s27.

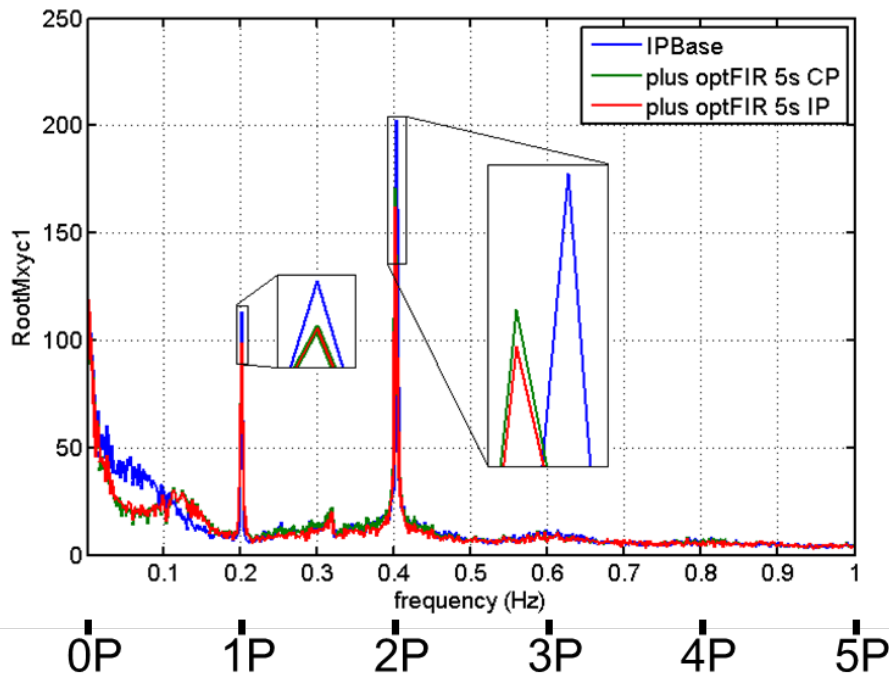


Figure 13. Blade root bending moment as a function of frequency. If M_{xyc1} is the root bending moment in-plane and M_{xyc2} is out-of-plane, then the plot is M_{xyc1} . Average FFT magnitude over AR1_s12 through AR1_s27.




Figure 14. Fore-aft nacelle acceleration as a function of frequency. Average FFT magnitude over AR1_s12 through AR1_s27.

In addition to frequencies of loads, it is also useful to look at the frequencies of actuator operation. Figure 15 shows that IP feedforward, when compared to CP feedforward, only increases pitching action at 2P and 4P (0.4 Hz at 0.8 Hz). This does not imply that IP feedforward is not doing anything at other frequencies; the blades can switch to pitching independently without necessarily increasing the amount that they pitch. Figure 15 also shows that, compared to feedback alone, both CP and IP have reduced pitching action below 0.1 Hz. Generator speed error and blade loads, shown in Figure 17 and Figure 13, are also greatly reduced below 0.1 Hz, so it is encouraging that, in this frequency range, preview allows for less blade pitch action, less generator speed error, and lower blade loads all at the same time.




Figure 15. Blade pitch angle as a function of frequency. Average FFT magnitude over AR1_s12 through AR1_s27.

Figure 16 shows generator torque versus frequency. Both CP and IP feedforward greatly reduce generator torque action below 0.11 Hz (because they are reducing generator speed error here). From 0.2 Hz up to 1 Hz, IP is slightly lower than both CP and feedback, except at the spike at 0.32 Hz, where both CP and IP are higher than feedback only. The plot of generator speed, shown in Figure 17, looks almost exactly the same as the generator torque plot, except in different units. This is because the region three generator torque controller is set up for constant power control: its torque command is inversely proportional to generator speed (after lowpass filtering with 0.25 Hz cutoff). However, upon zooming in on the generator speed plot, there is no difference between CP and IP feedforward from 0.2 to 1 Hz, except that exactly at the 0.32 Hz spike, IP is lower than CP.

The spike shown here at 0.32 Hz corresponds with the first tower fore-aft mode and the first tower side-to-side mode. This spike appears in the plots of generator torque, generator speed, drivetrain torsion, tower bending fore-aft and side-to-side, and nacelle motion in x, y, and z directions. It is important because this one frequency makes up the overwhelming majority of nacelle side-to-side motion, which is often a problem, showing increased loads when feedforward control is added, because it is very lightly damped (1% damping ratio). In this set of data, IP does slightly better than CP at reducing nacelle side-to-side velocity at 0.32 Hz, but both are worse than feedback alone.

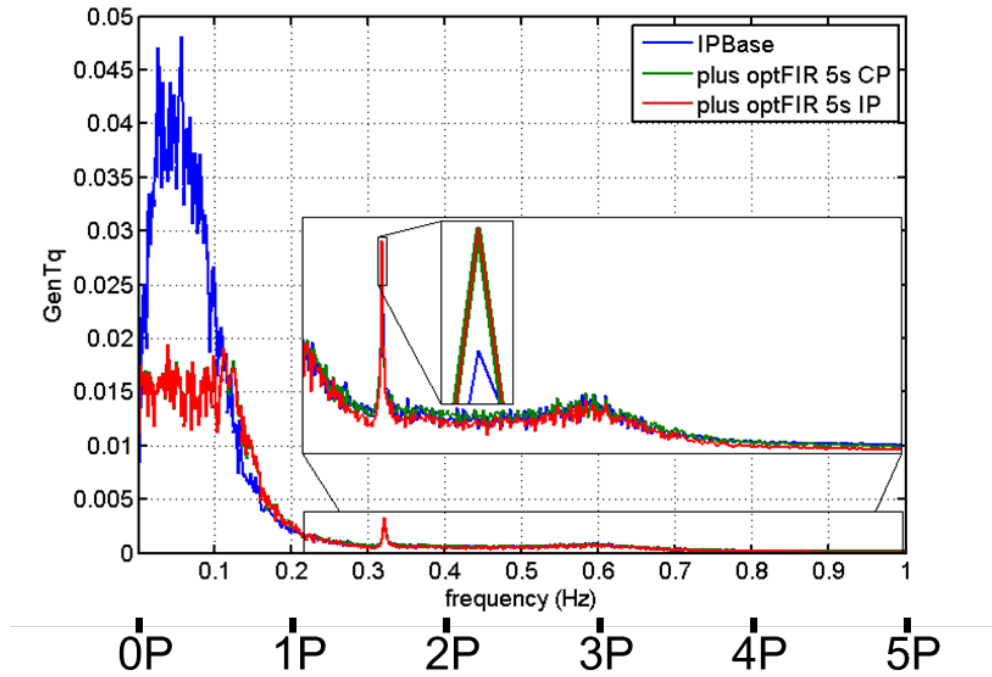


Figure 16. Generator torque as a function of frequency. Average FFT magnitude over AR1_s12 through AR1_s27.

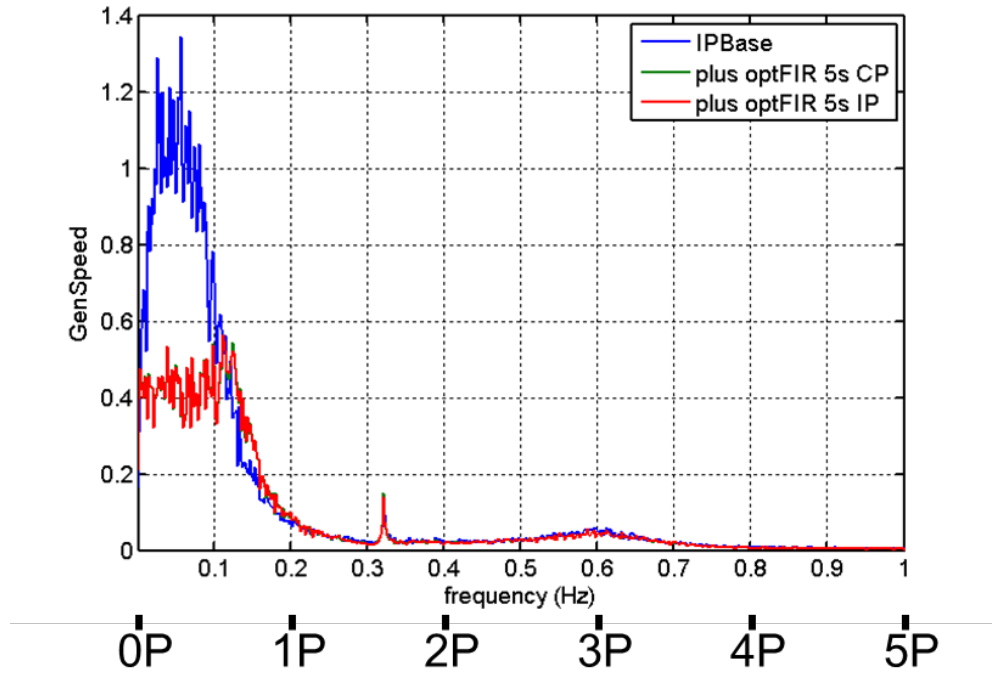


Figure 17. Generator speed as a function of frequency. Average FFT magnitude over AR1_s12 through AR1_s27. (This subset of the wind fields was chosen arbitrarily due to computer memory limitations).

These plots of turbine response parameters versus frequency have confirmed that IP feedforward load reductions all happen at or near the $1P_{rotating} = 0P_{nonrotating}$ or $2P_{rotating} = 3P_{nonrotating}$ frequencies. Nonrotating loads are reduced at both frequencies, and blade root loads are only reduced at $2P$ because IP feedback already does a good job of reducing them at $1P$. This implies that it might be possible to add IP feedback to the other loads to eliminate the need for individual pitch feedforward. The researchers have also gained some insight into the cause of the increased nacelle side-to-side motion that often occurs when feedforward is added. Finally, it has been shown that below 0.1 Hz, feedforward control simultaneously reduces blade pitch actuation, generator speed error, and blade loads.

5 Preview Times

The preview time required for any fully effective, feedforward controller for fatigue load reduction is the sum of three values, which are described in further detail in the following subsections. The first is determined by the actuator. The second is determined by the lowpass filter used. The third is determined by the turbine itself. Figure 18 shows estimates for these three values, along with the expected preview time available from the Light Detection and Ranging (LIDAR). There is slightly more preview time available than is required, and this is desirable because some processing time is also necessary. The amount of preview time available from LIDAR depends on focus distance and wind speed. For this 5-megawatt (MW) turbine, it has been estimated that the minimum measurement error is achieved when the LIDAR is focused about 60 to 100 meters (m) upwind, for a 13 meters per second (m/s) wind speed, when accounting for the effects of measurement angle errors, LIDAR spatial averaging, and wind evolution [12]. The ideal distance was expected to increase somewhat with wind speed. For the plot shown here, researchers assume that the LIDAR is focused 70 m ahead at 12 m/s and 120 m ahead at 24 m/s. The preview time estimates shown for the actuator, lowpass filter, and turbine are explained in the following subsections.

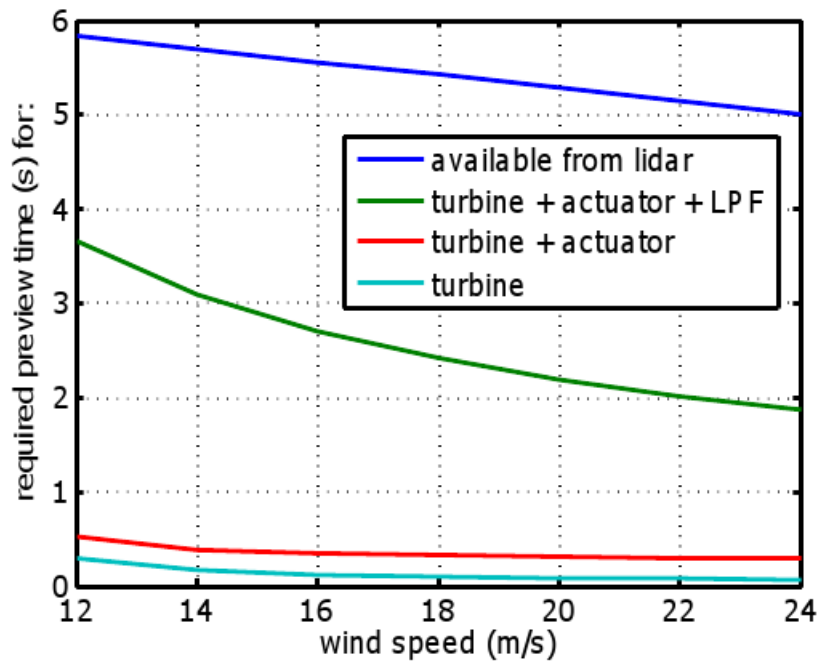


Figure 18. Cumulative required preview times for the turbine, actuator, and lowpass filter (LPF). Available preview time from the LIDAR is also shown.

5.1 Actuator

The actuator is modeled by a second-order lowpass filter with a natural frequency of 1 hertz (Hz) and a damping ratio of 0.7. Figure 19 and Figure 20 show how phase and time delay vary with frequency for this model. At frequencies below 0.2 Hz, the time delay remains at about

0.225 seconds. Therefore, 0.225 seconds of preview time is required to compensate for this actuator delay.




Figure 19. Phase of second-order actuator model (natural frequency of 1 Hz and damping ratio 0.7)




Figure 20. Time delay of second-order actuator model (natural frequency 1 Hz, damping ratio 0.7)
time delay = $-(\text{phase from Figure 19}) / (360 \text{ degrees}) / (\text{frequency})$

5.2 Lowpass Filter

LIDAR wind preview measurements are not accurate at high frequencies. As a result, some lowpass filtering must be included in any feedforward controller that relies on these measurements. The LIDAR itself, because of its volume averaging, acts as a lowpass filter with no phase delay, but more lowpass filtering is still necessary. The more that is filtered out in the controller, the more phase delay that is added, and therefore the more preview time is necessary. Doubling the order of the lowpass filter doubles the required preview time, because it doubles the phase delay. Doubling the frequency of the lowpass filter poles cuts the required preview time in half because the majority of the phase delay shifts to higher frequencies. For the 13 m/s wind speed, Figure 10 shows two configurations of the feedforward controller. Both provide similar amounts of lowpass filtering. One of these, feedforward controller B, is a second-order Butterworth lowpass filter with a cutoff of 0.08 Hz. Converting the phase delay of this filter from degrees to time in seconds, as in Figure 11, shows that it remains fairly constant at 2.9 seconds of delay at all frequencies below 0.1 Hz. (Using a Bessel filter instead of a Butterworth filter would make this delay vary less, but it does not make much difference since the filter order is only two).

Therefore, an effective lowpass filter at 13 m/s will require about 2.9 seconds of preview. The cutoff frequency of the lowpass filter is 0.04 radians per meter in spatial frequency. To find the required lowpass filter preview time at other wind speeds, the researchers assume that this spatial frequency will remain approximately constant as the wind speed increases. Figure 21 shows an estimate of the required preview time as a function of wind speed due to lowpass filtering.




Figure 21. Estimated preview required to cancel the delay of lowpass filtering (LPF), as a function of wind speed

5.3 Turbine

The third value of required preview time appears because for some output y , there is often more phase delay in the blade-pitch-to- y transfer function than there is in the wind-speed-to- y transfer function of the wind turbine. (Here y is some subset of the full vector of turbine outputs and can be a vector or scalar). This difference in phase delay is equal to the phase advance of the ideal model-inverse feedforward controller

where

- = the closed-loop (feedback included) transfer function from blade pitch to output
- = the closed-loop transfer function from wind speed to output.

For ease of calculations in this section, the researchers will let y be the generator speed output, and will use collective pitch feedforward control. Then, the equation above simplifies to

where

P = the open-loop plant
and the feedback controllers do not matter.

This is because the torque and collective pitch feedback controllers both take their input from generator speed, and since the feedforward is designed to reduce generator speed error to zero, these torque and collective pitch feedback controllers drop out of the equation for F . The individual pitch feedback controller also does not matter because it does not affect the average blade pitch, which is what is being added to with this collective pitch feedforward control.

To find the required preview time, given the above configuration, the researchers first linearized the open-loop turbine model at seven different wind speeds, from 12 m/s to 24 m/s, and solved for F at each wind speed. For each F , the researchers then took the phase advance at various frequencies and converted it to time in seconds. The results are shown in Figure 22 and Figure 23. The magnitude of each F is approximately flat in the frequency range of interest. Therefore, a simple feedforward generator speed controller can be made using a lookup table to find the required magnitude and time advance as a function of wind speed.




Figure 22. Estimated required preview time due to the turbine itself. All available degrees of freedom turned on in linearization. Above 0.5 Hz, time values begin to vary much more widely, but the researchers do not intend to use frequencies any higher for feedforward control.




Figure 23. Estimated required preview time due to the turbine itself. Five degrees of freedom turned on in linearization: first flapwise blade mode ($\times 3$ blades), generator, and drivetrain rotational-flexibility.

Figure 22 shows the results when all available degrees of freedom are turned on in linearization. The preview time required varies with wind speed, but stays fairly constant over the frequencies of interest.

Figure 23 shows results when only 5 degrees of freedom are turned on in linearization: first flapwise blade mode ($\times 3$ blades), generator, and drivetrain rotational-flexibility. Because the two plots are very similar, it can be concluded that the required preview time depends mainly on some or all of only these 5 degrees of freedom. A third set of linearizations was also created with only two degrees of freedom: generator and drivetrain rotational-flexibility. The resulting required preview time was almost exactly zero at all frequencies and at all wind speeds. This shows that the first flapwise blade mode is the main reason why the phase delay from blade-pitch-to-generator-speed is different than the phase delay from wind-speed-to-generator-speed. In other words, the amount of preview time required by the turbine depends mainly on blade flexibility.

It is also interesting to note that when the blade flap degree of freedom is turned on, a nonminimum phase zero appears in the blade-pitch-to-generator-speed transfer function, but not in the wind-speed-to-generator-speed transfer function. It is unlikely to be a coincidence that this nonminimum phase zero and the difference in phase delay appear only under the same circumstances. Preview time can also be used for model-inverse-based control when the plant has nonminimum phase (unstable) zeros. [1] Nonminimum phase zeros often appear in plants with noncollocated sensors and actuators, and they appear in the linearized 5-MW model when blade flexibility degrees of freedom are turned on. In this case, the ideal model-inverse feedforward controller is unstable, but it can be approximated by a stable, non-causal controller. Various approximation methods exist, including the noncausal series expansion [13, 14].

To confirm the results above, the team also found the required turbine preview time using a different method: spectra estimation. In this method, the researchers first find the frequency components of the rotor-effective wind speed of a class A wind field. Passing these frequency components through the closed-loop transfer function for the system (turbine plus controllers), an estimate for the spectra of the rotor speed was found. The standard variation in the rotor speed is just the square root of the integrated spectra. Researchers choose the preview time that minimizes this variation in rotor speed. Figure 24 shows the results of spectra estimation. These results match Figure 22 at the lower frequencies, as expected.

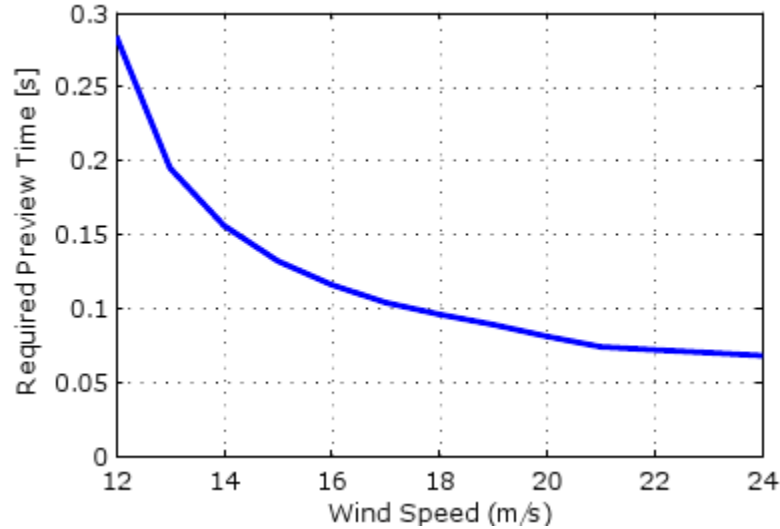


Figure 24. Spectra estimation method using a class A wind field. All degrees of freedom turned on in linearization. Matches Figure 22.

In simulations, a range of preview times was tried to see if they match the linearization results. The turbine preview time that resulted in the lowest root-mean-square generator speed error varied considerably, depending on the random seed used to generate the wind field, ranging from 0.2 to 0.6 seconds over just three seeds. These values are for controller B, in 14 m/s wind, with the actuator and lowpass filter preview times subtracted from the total preview time. The linearization value of 0.15 seconds falls just below this range of simulation results.

Required turbine preview time may increase greatly from the less than 0.3 seconds shown here if the focus is on other outputs besides generator speed. A controller that is designed only to regulate generator speed is not ideal for reducing loads, although it may help reduce them somewhat. Previous work done using preview control [1] suggests a turbine-required preview time of 2 seconds when only penalizing pitch rate and generator speed error, and 20 seconds when tower sway is also penalized. However, other work [2] shows that the preview time required to minimize loads depends on wind speed and is at most 1 second. Future work should further investigate the required turbine preview time when designing for blade root and tower loads as well as generator speed regulation.

6 Conclusions and Future Work

In this study, researchers compared two previously independently developed feedforward controllers and Light Detection and Ranging (LIDAR) configurations to better understand what properties of such controllers yield improved load mitigation. Further, the team evaluated the minimum preview times required for the implementation of feedforward controllers. Additional remarks on the findings are also provided below.

First, in comparing the two feedforward controllers and LIDAR configurations, researchers found that the two feedforward controllers have similar performance. Further, the simpler of the two feedforward controllers (feedforward controller B) can be more easily tuned for real-world use. One version of LIDAR measurement processing (A) does not discard information on twice-per-revolution (2P) wind speed variations, and is able, in simulation, to additionally reduce loads at that frequency. The LIDAR configuration with more measurement points, covering more of the rotor plane, (B) gives a better estimate of the rotor-effective wind speed. The LIDAR configuration with one measurement for each blade (A) measures only the relevant sections of wind inflow, and therefore gives a better estimate of the blade-effective wind speed. However, this requires an accurate prediction of where the blades will be a few seconds in advance. It was also concluded that a greater percent load reduction is possible in more turbulent wind, and that feedback gains should be reduced when feedforward control is used.

Second, researchers studied individual pitch feedforward load reductions as a function of frequency. It was found that individual pitching action resulted in load reductions near the DC (0P) and three-times-per-revolution (3P) frequencies in the stationary frame, caused by pitching at once-per-revolution (1P) and 2P in the rotating frame. Researchers also found that both collective pitch feedforward and individual pitch feedforward caused reductions in both pitch rate and blade root loads below 0.1 hertz (Hz).

Third and finally, the team explained the contributing sources of required preview time. The main source is the phase delay of the lowpass filter, which is needed to filter out high-frequency components of the wind measurements, since these higher frequency measurements are usually not accurate enough to be used to mitigate loads. This filter will likely require 2–3 seconds of preview, with both required preview time and available preview time decreasing as wind speed increases.

Future work is being carried out to gain a better analytic understanding of what is required of the LIDAR measurement accuracies and feedforward controller in order to enable improvements in the load mitigation control performance over feedback-only control.

References

1. Dunne, F.; Pao, L.Y.; Wright, A.D.; Jonkman, B.; Kelley, N.; Simley, E. "Adding Feedforward Blade Pitch Control for Load Mitigation in Wind Turbines: Non-Causal Series Expansion, Preview Control, and Optimized FIR Filter Methods," *Proc. AIAA Aerospace Sciences Meeting*, Orlando, FL, Jan. 2011.
2. Schlipf, D.; Schuler, S.; Grau, P.; Allgöwer, F.; Kühn, M. "Look-Ahead Cyclic Pitch Control Using LIDAR," *Proc. The Science of Making Torque from Wind*, Heraklion, Greece, June 2010.
3. Jonkman, J.; Buhl, M. L. *FAST User's Guide*, NREL/EL-500-38230, NREL Report, National Renewable Energy Laboratory, Golden, CO, August 2005.
4. Jonkman, J.; Butterfield, S.; Musial, W.; Scott, G. *Definition of a 5-MW Reference Wind Turbine for Offshore System Development*, NREL/TP-500-38060, NREL Report, National Renewable Energy Laboratory, Golden, CO, February 2009.
5. Bossanyi, E.A. "Individual Blade Pitch Control for Load Reduction," *Wind Energy*, Vol. 6, 2003, pp. 119–128.
6. Bir, G. "Multi-blade Coordinate Transformation and its Application to Wind Turbine Analysis," *Proc. AIAA Aerospace Sciences Meeting*, Reno, NV, January 2008.
7. Kelley, N.D.; Jonkman, B.J. *Overview of the TurbSim Stochastic Inflow Turbulence Simulator: Version 1.21* (Revised Feb. 1, 2007), NREL/TP-500-41137, NREL Report, National Renewable Energy Laboratory, Golden, CO, April 2007.
8. Jonkman, B.J., *TurbSim User's Guide*, NREL/TP-500-46198, NREL Report, National Renewable Energy Laboratory, Golden, CO, September 2009.
9. Simley, E.; Pao, L.Y.; Frehlich, R.; Jonkman, B.; Kelley, N. "Analysis of Wind Speed Measurements using Coherent LIDAR for Wind Preview Control," *Proc. AIAA Aerospace Sciences Meeting*, Orlando, FL, January 2011.
10. Rettenmeier, A.; Bischoff, O.; Hofsäß, M.; Schlipf, D.; Trujillo, J. J.; Kühn, M. "Wind Field Analyses Using a Nacelle-based LIDAR System," *Proc. European Wind Energy Conference*, 2010.
11. Schlipf, D.; Trabucchi, D.; Bischoff, O.; Hofsäß, M.; Mann, J.; Mikkelsen, T.; Rettenmeier, A.; Trujillo, J.; Kühn, M. "Testing of Frozen Turbulence Hypothesis for Wind Turbine Applications with a Scanning LIDAR System," *International Society of Acoustic Remote Sensing Symposium*, Paris, France, 2010.
12. Simley, E.; Pao, L.Y.; Kelley, N.; Jonkman, B.; Frehlich, R. "LIDAR Wind Speed Measurements of Evolving Wind Fields," *Proc. AIAA Aerospace Sciences Meeting*, Nashville, TN, January 2012.

13. Rigney, B.P.; Pao, L.Y.; Lawrence, D.A. “Nonminimum Phase Dynamic Inversion for Settle Time Applications,” *IEEE Trans. Control Systems Technology*, Vol. 17, No. 5, September 2009, pp. 989–1005.
14. Widrow, B.; Walach, E. *Adaptive Inverse Control - A Signal Processing Approach*, Wiley, Hoboken, NJ, 2008.