

Adaptive Torque Control of Variable Speed Wind Turbines

Kathryn E. Johnson



NREL

National Renewable Energy Laboratory
1617 Cole Boulevard, Golden, Colorado 80401-3393
303-275-3000 • www.nrel.gov

Operated for the U.S. Department of Energy
Office of Energy Efficiency and Renewable Energy
by Midwest Research Institute • Battelle

Contract No. DE-AC36-99-GO10337

Adaptive Torque Control of Variable Speed Wind Turbines

Kathryn E. Johnson

Prepared under Task No. WER43302



NREL

National Renewable Energy Laboratory
1617 Cole Boulevard, Golden, Colorado 80401-3393
303-275-3000 • www.nrel.gov

Operated for the U.S. Department of Energy
Office of Energy Efficiency and Renewable Energy
by Midwest Research Institute • Battelle

Contract No. DE-AC36-99-GO10337

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/ordering.htm>



Executive Summary

Wind is a clean, renewable resource that has become more popular in recent years as a result of numerous advances in technology and public awareness. Wind energy is quickly becoming cost competitive with fossil fuels, but further reductions in the cost of wind energy are necessary before it can grow into a fully mature technology. One reason for the higher-than-necessary cost of wind energy is uncertainty in the aerodynamic parameters, which leads to inefficient controllers. This report explores an adaptive control technique designed to reduce the negative effects of this uncertainty.

The primary focus of this work is a new adaptive controller that is designed to resemble the standard non-adaptive controller used by the wind industry for variable speed wind turbines below rated power. This adaptive controller uses a simple, highly intuitive gain adaptation law designed to seek out the optimal gain for maximizing the turbine's energy capture. It is designed to work even in real, time-varying winds.

The adaptive controller has been tested both in simulation and on a real turbine, with numerous experimental results provided in this work. Simulations have considered the effects of erroneous wind measurements and time-varying turbine parameters, both of which are concerns on the real turbine. The adaptive controller has been found to operate as desired under realistic operating conditions, and energy capture has increased on the real turbine as a result.

Theoretical analyses of the standard and adaptive controllers were performed, as well, providing additional insight into the system. Finally, a few extensions were made with the intent of making the adaptive control idea even more appealing in the commercial wind turbine market.

Acknowledgements

I would like to acknowledge first my faculty co-advisors at the University of Colorado, Dr. Lucy Pao and Dr. Mark Balas, for all of their guidance and support in this work. Second, I am grateful to the many researchers at the National Wind Technology Center at the U.S. Department of Energy's National Renewable Energy Laboratory (NREL) for their support and assistance. In particular, I would like to thank Lee Jay Fingersh, Maureen Hand, Mike Robinson, and Karl Stol. Additionally, I appreciate the assistance of Vishwesh Kulkarni, who provided some valuable theoretical assistance for Appendix B, and Dale Lawrence.

Finally, I would like to acknowledge the financial support of the American Society for Engineering Education, the Department of Energy and NREL, and the University of Colorado.

Contents

CHAPTER 1: INTRODUCTION	1
1.1 WIND TURBINE OPERATION	1
1.2 REGIONS OF CONTROL	3
1.3 PREVIOUS CONTROL STRATEGIES	4
1.4 SCOPE AND OUTLINE	5
CHAPTER 2: REGION 2 CONTROL STRATEGIES	8
2.1 STANDARD REGION 2 CONTROL	8
2.2 STANDARD CONTROL LAW PROBLEMS	11
2.3 ADAPTIVE CONTROL	13
CHAPTER 3: MODEL DEVELOPMENT	17
3.1 MODEL OVERVIEW	17
3.2 MODEL VALIDATION	20
CHAPTER 4: SIMULATION RESULTS	23
4.1 OPTIMAL OPERATING POINT	23
4.2 SIMINT SIMULATION RESULTS	27
4.3 WIND MEASUREMENT PROBLEMS	32
4.4 TIME-VARYING PARAMETERS	33
4.5 SIMULATIONS USING SYMDYN	37
4.6 SIMULATED STRESSES ON TURBINE COMPONENTS	38
CHAPTER 5: REAL CART DATA	42
5.1 ADAPTATION HISTORY	42
5.2 ENERGY CAPTURE USING ADAPTIVE CONTROL	43
5.3 MEASURED STRESSES ON TURBINE COMPONENTS	47
5.4 REGION 2–3 TRANSITION	54

CHAPTER 6: THEORETICAL ANALYSIS	58
6.1 ASYMPTOTIC STABILITY OF $\omega = 0$	59
6.2 ASYMPTOTIC STABILITY OF ROTOR SPEED WITH CONSTANT, POSITIVE WIND INPUT	60
6.3 INPUT-OUTPUT STABILITY	62
6.4 CONVERGENCE CHARACTERISTICS OF THE GAIN ADAPTATION ALGORITHM	63
6.4.1 <i>Types of Instability</i>	63
6.4.2 <i>Bounds on $\gamma_{\Delta M}$ Gain for Stability</i>	65
6.4.3 <i>Asymmetric P_{favg} vs. \tilde{M} Curves</i>	67
CHAPTER 7: EXTENSIONS TO ORIGINAL GAIN ADAPTATION LAW	69
7.1 USE OF NACELLE ANEMOMETER	69
7.1.1 <i>CART Adaptations Using the Nacelle Anemometer</i>	70
7.1.2 <i>Effects of Tip Speed Ratio on Axial Induction Factor</i>	71
7.1.3 <i>Simulations of Axial Induction Factor</i>	74
7.2 FILTERED GAIN ADAPTATION LAW	78
CHAPTER 8: CONCLUSIONS AND FUTURE RECOMMENDATIONS	81
REFERENCES	83
APPENDIX A: MEAN EQUIVALENT WIND SPEED	86
APPENDIX B: STABILITY PROOF FOR GAIN ADAPTATION LAW	89

List of Common Symbols

A	rotor swept area
C_p	rotor power coefficient
C_{pmax}	maximum rotor power coefficient
C_q	rotor torque coefficient
C_T	thrust coefficient
J	rotor inertia
K	standard torque control gain
M	adaptive torque control gain
M^+	simulation-derived estimate of optimal torque control gain
M^*	turbine's true optimal torque control gain (possibly unknown)
P	turbine (rotor) power
P_{cap}	captured power
P_{avg}	average captured power divided by average wind power over a given time period
P_{wind}	power available in the wind
P_{wy}	power available in the wind, with approximate yaw error factor included
R	rotor radius
f_s	sampling frequency
a	axial induction factor
\underline{b}	vector of coefficients in filtered gain adaptation law: $\underline{b} = [b_0, b_1]$
\underline{c}	vector of coefficients in filtered gain adaptation law: $\underline{c} = [c_1, c_2]$

k	discrete time index
n	number of steps in adaptation period
v	wind speed
v_∞	wind speed (“free stream,” or far upwind; i.e., not affected by the turbine)
v_n	wind speed (measured with nacelle anemometer)
v_r	wind speed (at the rotor plane)
v_u	wind speed (measured with upwind anemometer)
β	blade pitch
$\gamma_{\Delta M}$	positive gain on adaptation law
λ^*	tip speed ratio (TSR) corresponding to $C_{p_{\max}}$
ρ	air density
τ_{aero}	aerodynamic torque
τ_c	generator (control) torque
ψ	yaw error
ω	rotor angular speed

List of Tables

TABLE 4-1. SIMULATION PARAMETERS	27
TABLE 4-2. SUMMARY OF PARAMETER CHANGES USED IN FIGURE 4-9 SIMULATION	35
TABLE 4-3. SELECTED DETAILS RELEVANT TO FIGURE 4-11 THROUGH FIGURE 4-13.....	39
TABLE 5-1. MEAN EQUIVALENT WIND SPEED AND MEAN GRID POWER DATA FOR STANDARD AND ADAPTIVE CONTROLLERS ($0.40M^+ < M < 0.52M^+$).....	44
TABLE 5-2. INCREASE IN ANNUAL ENERGY CAPTURE OVER STANDARD CONTROLLER FOR ADAPTIVE CONTROLLER WITH $0.40M^+ < M < 0.52M^+$	46
TABLE 5-3. SELECTED DETAILS RELEVANT TO FIGURE 5-5 THROUGH FIGURE 5-8.....	48
TABLE 7-1. RELEVANT DATA FOR FIGURE 7-5	74
TABLE A-1: MEAN EQUIVALENT WIND SPEED AND MEAN WIND SPEED EXAMPLE	88

List of Figures

FIGURE 1-1: TWO-BLADED UPWIND HORIZONTAL AXIS WIND TURBINE COMPONENTS	2
FIGURE 1-2: STEADY-STATE POWER CURVES	4
FIGURE 2-1: C_p VS. TSR AND PITCH FOR CONTROLS ADVANCED RESEARCH TURBINE (CART)	9
FIGURE 2-2: CART C_p VS. TSR AND CUBIC FUNCTION	11
FIGURE 2-3: SUB-OPTIMAL OPERATION DUE TO INACCURATE C_p -TSR CURVE	12
FIGURE 2-4: FRACTIONAL AVERAGE POWER P_{FAVG} VS. NORMALIZED M	14
FIGURE 3-1: BLOCK DIAGRAM OF THE SIMINT SIMULINK MODEL	17
FIGURE 3-2: PROP- AND SYMDYN-GENERATED TORQUE COEFFICIENT (C_Q) SURFACES	18
FIGURE 3-3: UPWIND HAWT OPERATING WITH NON-ZERO YAW ERROR ψ (TOP VIEW)	19
FIGURE 3-4: DECREASE IN AVAILABLE WIND POWER VS. YAW ERROR	20
FIGURE 3-5: VALIDATION STUDY OF SIMINT MODEL USING CART DATA	21
FIGURE 3-6: VALIDATION STUDY OF SIMINT MODEL USING SYMDYN	22
FIGURE 4-1: AVERAGE POWER CAPTURE FOR NORMAL AND LOW INERTIA	24
FIGURE 4-2: OPTIMUM TORQUE GAIN VS. TURBULENCE INTENSITY	25
FIGURE 4-3: SIMULATION RESULTS WITH CONSTANT WIND INPUT	28
FIGURE 4-4: SIMULATION RESULTS WITH REAL WIND DATA AS INPUT	29
FIGURE 4-5: FRACTIONAL AVERAGE POWER P_{FAVG} VS. NORMALIZED M (M/M^+) FOR SIMULATION DATA: (A) ALL DATA; (B) ZOOMED IN AROUND PEAK	30
FIGURE 4-6: ADAPTATIONS OF M FOR VARYING AMOUNTS OF YAW ERROR	31
FIGURE 4-7: SIMULATIONS WITH ERRONEOUS WIND MEASUREMENTS	32
FIGURE 4-8: SIMULATED CHANGE IN PARAMETERS WITH CONSTANT WIND SPEED INPUT	34
FIGURE 4-9: SIMULATED CHANGES IN PARAMETERS WITH VARIABLE WIND SPEED INPUT	35
FIGURE 4-10: COMPARISON OF SIMINT AND SYMDYN OUTPUTS	37
FIGURE 4-11: DRIVE SHAFT TWIST ANGLE FOR TWO DIFFERENT ADAPTATIONS	39
FIGURE 4-12: BLADE 1 AND 2 FLAP ANGLE FOR TWO DIFFERENT ADAPTATIONS	40
FIGURE 4-13: TOWER SIDE-TO-SIDE ANGLE FOR TWO DIFFERENT ADAPTATIONS	41
FIGURE 5-1: ADAPTATIONS OF M ON THE CART	43

FIGURE 5-2: GRID POWER VS. EQUIVALENT WIND SPEED FOR STANDARD AND ADAPTIVE CONTROL: (A) 10-MINUTE AVERAGES; (B) BINNED DATA WITH ERROR BARS	45
FIGURE 5-3: EXPECTED INCREMENTAL AND TOTAL ENERGY FOR STANDARD AND ADAPTIVE CONTROL ON THE CART	46
FIGURE 5-4: FRACTIONAL AVERAGE POWER P_{FAVG} VS. NORMALIZED M	47
FIGURE 5-5: DRIVE SHAFT TWIST ANGLE FOR TWO DIFFERENT ADAPTATIONS (REAL CART DATA)	49
FIGURE 5-6: BLADE FLAP BENDING MOMENT (REAL CART DATA)	49
FIGURE 5-7: BLADE EDGE BENDING MOMENT (REAL CART DATA)	50
FIGURE 5-8: TOWER SIDE-TO-SIDE BENDING MOMENT (REAL CART DATA)	51
FIGURE 5-9: TOWER FORE-AFT BENDING MOMENT (REAL CART DATA)	51
FIGURE 5-10: X, Y, AND Z ACCELEROMETER DATA	52
FIGURE 5-11: X, Y, AND Z ACCELEROMETER DATA VERSUS MEAN ROTOR SPEED DATA	53
FIGURE 5-12: BAD REGION 2 - REGION 3 TRANSITION ON THE CART	54
FIGURE 5-13: GOOD REGION 2 - REGION 3 TRANSITION ON THE CART	55
FIGURE 5-14: WIND SPEED FOR BAD AND GOOD TRANSITION EXAMPLES	56
FIGURE 5-15: STANDARD DEVIATIONS OF TURBINE LOADS VS. NORMALIZED TORQUE GAIN M DURING REGION 2 TO REGION 3 TRANSITIONS	57
FIGURE 6-1: SIMPLIFIED BLOCK DIAGRAMS (A) RELATING AERODYNAMIC TORQUE AND ROTOR SPEED, AND (B) GAIN ADAPTATION LAW.	58
FIGURE 6-2: FRACTIONAL AVERAGE POWER P_{FAVG} VS. ERROR IN M	59
FIGURE 6-3: CART C_p AND CUBIC FOR DIFFERENT M	61
FIGURE 6-4: INSTABILITY EXAMPLE 2 FOR GAIN ADAPTATION ALGORITHM	65
FIGURE 6-5: FINDING THE CRITICAL GAIN	66
FIGURE 6-6: RATIO OF MEAN CAPTURED POWER TO MEAN WIND POWER VS. ERROR IN TORQUE CONTROL GAIN M	68
FIGURE 7-1: CART (A) OPERATING AT THE NWTC WITH UPWIND MET TOWER SHOWN, AND (B) CLOSE-UP VIEW SHOWING THE NACELLE ANEMOMETER	70
FIGURE 7-2: ADAPTATIONS OF M ON THE CART USING UPWIND AND NACELLE ANEMOMETERS ...	71
FIGURE 7-3: FRACTIONAL AVERAGE POWER P_{FAVG} VS. NORMALIZED M	72
FIGURE 7-4: RELATIONSHIP BETWEEN NACELLE AND UPWIND WIND SPEED MEASUREMENTS	73

FIGURE 7-5: 10-MINUTE AVERAGES OF NACELLE VS. UPWIND WIND SPEED FOR VARIOUS TIP SPEED RATIOS	74
FIGURE 7-6: PROP-DERIVED CART MODEL SHOWING (A) THRUST COEFFICIENT VS. TIP SPEED RATIO, AND (B) AXIAL INDUCTION FACTOR VS. TIP SPEED RATIO	75
FIGURE 7-7: NORMALIZED M AND FRACTIONAL AVERAGE POWER P_{FAVG} FOR AXIAL INDUCTION FACTOR SIMULATION.....	77
FIGURE 7-8: AXIAL INDUCTION FACTOR A AND TIP SPEED RATIO λ OBTAINED IN SIMULATION.....	78
FIGURE 7-9: SIMULATION RESULTS FOR NEW FILTERED GAIN ADAPTATION LAW WITH CONSTANT WIND SPEED INPUT	79
FIGURE 7-10: SIMULATION RESULTS FOR NEW FILTERED GAIN ADAPTATION LAW WITH REAL (TIME-VARYING) WIND SPEED INPUT	80
FIGURE A-1: MEAN EQUIVALENT WIND SPEED EXAMPLE.....	86
FIGURE A-2: MEAN LOW SPEED SHAFT POWER VERSUS MEAN WIND SPEED AND MEAN EQUIVALENT WIND SPEED.....	87
FIGURE B-1: FEEDBACK SYSTEMS CONSIDERED BY SAFONOV: (A) GENERAL CASE, AND (B) SPECIFIC CASE WITH DISTURBANCES ENTERING ADDITIVELY	89
FIGURE B-2: STABILITY OF EXAMPLE FEEDBACK SYSTEM USING SAFONOV'S SECTOR METHOD	90
FIGURE B-3: ADAPTIVE CONTROL FEEDBACK SYSTEM.....	91
FIGURE B-4: RELATIONSHIP BETWEEN U AND Y WITHIN THE SUBSYSTEM S_{2I}	92
FIGURE B-5: STABILITY OF THE CLOSED-LOOP SYSTEM IN FIGURE B-3 USING SAFONOV'S SECTOR METHOD	94

CHAPTER 1: Introduction

Wind power was recognized as a valuable resource thousands of years before internal combustion engines and modern power plants were developed. Long ago, for example, people built windmills for milling grain and installed sails on their boats in order to reduce human labor demands. In recent decades, wind has experienced a resurgence in popularity by helping to meet the increasing energy consumption of a growing world population through the use of a renewable, non-polluting energy source.

To be competitive in the electricity market, wind turbines must produce electricity at a cost comparable to that of fossil fuel technology. Since the fuel (wind) is free, the major cost of wind energy comes from the initial capital outlay for the machine itself—roughly \$1000 per kW of installed capacity—and , and the supporting infrastructure. Relatively small costs are associated with operating and maintaining the turbine throughout its lifetime. Therefore, technology advances that reduce the initial cost of the turbine or allow it to capture more energy during operation result in a cost of energy (COE) decrease. Current wind technology produces power at 3¢ to 5¢ per kWh and is rapidly becoming cost competitive with fossil fuel energy.

All aspects of turbine design, construction, and operation are scrutinized in an attempt to decrease the COE. For example, turbine designs have improved through the use of advanced materials for increased strength and durability. Construction methods now incorporate strategies such as self-erecting tall towers and offshore ocean deployment. Finally, computer simulation tools used to model aerodynamic loads and dynamic responses are becoming more adept at modeling real turbines, which allows for the development of more sophisticated control algorithms. These sophisticated control algorithms aim to decrease the COE using a variety of techniques. The focus of this report is the design of an adaptive controller to maximize energy capture at lower wind speeds.

1.1 Wind Turbine Operation

There are many different types of wind turbines in use around the world, used in diverse applications from powering small water pumps on remote ranches to producing bulk electric power for modern cities. Capacity ratings vary from less than one kilowatt up to several megawatts; for example, a typical utility-scale turbine might generate up to 1.5 MW and sit atop a 80-meter tower. There are two principal turbine architectures: horizontal axis and vertical axis. Horizontal axis wind turbines (HAWT), which have an axis of rotation that is parallel with the wind direction, are the most common. The much less common vertical axis wind turbine (VAWT) spins like a top with an axis of rotation perpendicular to the wind. HAWTs have been shown to be more cost effective than VAWTs because the ability to place the HAWT rotor on a very tall tower allows the turbine to operate in a better wind resource due to vertical wind shear. In contrast, VAWTs are constrained to operate with one end of the rotor attached to the generator on the ground. Modern HAWTs typically have two or three blades and can be either upwind




Figure 1-1: Two-bladed upwind horizontal axis wind turbine components

machines, with the rotor spinning on the upwind side of the tower, or downwind machines. Figure 1-1 shows a two-bladed upwind HAWT and its major components.

A typical modern wind turbine has one of two basic operating modes: constant or variable speed. The rotor of a constant speed turbine turns at a constant angular speed, often selected so that the generator has the same frequency as the power grid to which it is connected, regardless of wind speed fluctuations. This approach can eliminate the need for expensive power electronics, but it constrains the rotor speed so that the turbine cannot operate at its peak aerodynamic efficiency in all wind speeds. Thus, a constant speed turbine usually produces less power at low wind speeds than does a variable speed turbine, which is designed to operate at a rotor speed proportional to the wind speed (below its rated wind speed). However, in order for a variable speed turbine to achieve its maximum power capture using standard controllers, complex aerodynamic properties must be well known; in practice, these uncertainties can easily lead to a variable speed turbine capturing even less power than a constant speed one in low winds. Adaptive control can solve this problem, and the research presented in this paper develops an adaptive control approach.

Most utility-scale turbines have three main active control actuation systems. The first, a yaw drive, turns the nacelle and rotor so that the wind is normal to the rotor plane. The yaw drive is intentionally very slow for large turbines in order to reduce gyroscopic loading. Variable speed turbines can also use generator torque control, which effectively tells the generator how much torque to demand from the rotor in order to optimize the rotor speed. Finally, most modern

turbines have active pitch control, in which either the entire blade or some section thereof can be rotated on demand. Pitch control changes the aerodynamic loading by changing the angle of attack. It can be either independent, in which each blade can be pitched separately from the others, or collective. Both generator torque control and pitch control can be relatively fast and are thus suitable as actuators in modern control schemes.

1.2 Regions of Control

Variable speed wind turbines have three main regions of operation. Figure 1-2 plots the power curves for the wind, an ideal (lossless) turbine, and an example variable speed wind turbine with a 43.3 m rotor diameter. The three major control regions are labeled on the turbine's power curve. In this example, the turbine produces maximum power (800 kW) at a rated wind speed of about 14 m/s, and it has a maximum power coefficient $C_{pmax} = 0.4$. Power coefficients are explained in more detail in Chapter 2.

Region 1, the section of the solid curve located to the left of 5 m/s wind speed, includes the times when the turbine is not operating and when it is starting up. In general, the region 1 control strategy involves monitoring the wind speed to determine whether it lies within the specifications for turbine operation and, if so, performing the necessary routines used in starting up the turbine. The use of modern control strategies is not usually critical in region 1.

Region 2, shown by the cubic section of the solid curve (approximately 5 to 14 m/s wind speeds), is an operational mode in which it is desirable to capture as much power as possible from the wind. Aerodynamic losses prevent the turbine from achieving its maximum theoretical power extraction from the wind, called the Betz Limit, but the goal is to approach the Betz Limit curve as closely as possible. All three control strategies (yaw drive, generator torque, and blade pitch) may be used in this region; however, it is common to use only generator torque and yaw control for most of the time in region 2, keeping the blade pitch constant at an optimal value for peak energy extraction. In general, the loads experienced by turbine components in region 2 are lower than those in region 3, and minimizing loads in this region is often not essential.

Region 3 operation occurs above rated wind speed, i.e., the wind speed above which maximum peak power is produced. The turbine must limit the fraction of the wind power captured so as not to exceed electrical and mechanical design loads. Thus, the solid curve flattens to the turbine's rated power, 800 kW, above the rated wind speed (approximately 14 m/s). In region 3, a variable speed turbine often maintains a constant speed and constant, rated power, pitching its blades in order to shed additional power. The industry standard control method is proportional-integral-derivative (PID) control on the blade pitch, a simple but effective means for limiting speed and power. Yaw control, generator torque, and blade pitch strategies can all be used to shed excess power and limit the turbine's energy capture as well as to achieve other control objectives. In this region, loads on components such as blades, drive shaft, and tower may become excessive and active control can be used to reduce vibrations. Load reductions decrease the amount of material required, thereby reducing initial capital costs.




Figure 1-2: Steady-state power curves

1.3 Previous Control Strategies

Control of modern variable speed wind turbines presents numerous interesting challenges to design engineers. Many researchers have developed methods for limiting power and speed in region 3 (e.g., Stol and Balas 2002), reducing loads (e.g., Eggers et al. 1998, Hand 2003, Wright 2003), and using adaptive control to compensate for unknown and time-varying parameters (e.g., Bhowmik et al. 1999, Freeman and Balas 1999, Song et al. 2000).

The design of advanced controls for wind turbines can be approached in the beginning from either a more theoretical or a more practical standpoint. On the practical side, it is common for the turbine itself to be more accurately modeled but the controls results to be less advanced. The theoretical standpoint, in contrast, may provide advanced control theory results but use a turbine model with additional simplifications. Of course, these two approaches can come together, providing advanced controls results while being faithful to the highly nonlinear, time-varying characteristics of wind turbines. Most research begins with one or the other of these two approaches, however.

The theoretical approach to variable speed wind turbine control design is a valuable contribution to both the controls community and the wind research community. Despite theoretical advances, however, the broader wind industry has for the most part continued to use the same simple

control schemes that have been in use for decades. Most of the published advanced control research relating to variable speed turbines begins with the theoretical approach and ends with simulation results. Sometimes the theoretical-to-practical approach leads to good results (Bossanyi 1987, 1989). Often, however, researchers who attempt implementation on a real turbine find that the turbine is so nonlinear and has so many unknowns that the simplifications used in the theoretical approach and simulations cause the theoretically designed controller to fail. Most of these failures, of course, are not published.

Classical techniques such as PID control of blade pitch are typically used to limit power and speed for turbines operating in region 3 (Bossanyi 1987, Hand and Balas 2000), and some sort of generator torque control is used in region 2 (Fingersh and Carlin 1998). While the specific techniques used to control the most modern turbines are proprietary and typically unpublished, only recently have any turbine manufacturers incorporated aspects of the decades of advanced control research into their commercial turbine designs. In part, this technology gap exists due to the risk adverse commercial practices of turbine developers and reflects the lack of theoretically derived controllers having been successfully tested and validated on real turbines operating in the field.

Enslin and Van Wyk (1992) have tested a simple direct-speed control method on a real variable speed turbine with some promising results. This method used a simple controller that incremented or decremented speed in response to changes in turbine power. This research has since been enhanced by the use of adaptive control to account for turbine uncertainties (Bhowmik et al. 1999), although the adaptive controller has not been tested on a real field turbine. Each of these controllers attempted to maximize the turbine's power and efficiency.

Unfortunately, not all turbines can be speed-controlled at the rapid rates required for the strategies in Bhowmik et al. (1999) and Enslin and Van Wyk (1992) to work. In particular, turbines with large rotational inertias require relatively long time periods for changes in speed in comparison to the extremely short times required for changes in wind speed. Because angular acceleration is proportional to net torque, the torque applied to the drive shaft during direct speed control of a large turbine may be higher than the design load for the shaft. In summary, any controller that relies on a small turbine's ability to accelerate or decelerate at very high rates in order to track wind speed changes might not be suitable for a large-scale turbine because the latter has a rotor inertia that is too large to allow those acceleration rates while remaining within safe shaft torque ranges. Thus, a different type of adaptive control is necessary in order to achieve the control objectives in Bhowmik et al. (1999) and Enslin and Van Wyk (1992) for a large turbine.

1.4 Scope and Outline

The research presented in this report begins with a practical approach to variable speed wind turbine control. Only after an intuitive idea is developed and tested both in simulation and on a real turbine are the theoretical aspects of the problem (stability and convergence) examined for a simplified system. The research was motivated by the fact that the aerodynamic properties of wind turbines are poorly understood and the fact that this lack of understanding contributes to a significant increase in the cost of wind energy. While several of the aforementioned papers

described adaptive control methods to deal with various turbine uncertainties, this report combines, for the first time, all of the following objectives:

1. Develop an adaptive strategy for region 2 control to address the effects of aerodynamic uncertainty.
2. Build from a control strategy that has been well-tested and used within the wind industry.
3. Verify the operation of the adaptive controller using both simulations and experimentation on a real, mid-sized, variable speed turbine.
4. Analyze both the original standard control and the new adaptive control from a theoretical stability viewpoint.

The typical effect of aerodynamic uncertainty is to cause sub-optimal energy capture in region 2, but adaptive control can reduce or eliminate the effects of this uncertainty. The cost of changing the control strategy on an existing turbine is very small when compared with the cost of the turbine, and increasing energy capture without adding expense will directly reduce the cost of wind energy.

The fact that this research builds from a control strategy that has been well-tested and used within the industry is important because many of the implementation issues have already been worked out. For example, the required actuators and many of the sensors are already in place on most commercial turbines, and their limitations are well known. Also, this familiarity will lend the new control strategy credibility within the commercial turbine industry.

In order for any new controller to be accepted, it must be verified using both simulations and experimentation on real turbines. This research uses a mid-sized (600 kW) turbine, which is smaller than the most modern utility-scale turbines being produced but similar in size to many of those already in operation. The size of the turbine is important because certain control objectives are much harder to meet on larger turbines, and every year the size of a “typical” utility scale turbine increases.

Chapter 2 begins with a discussion of the standard region 2 control used for variable speed wind turbines. It then describes some of the inherent problems in this control region and presents a basic adaptive control strategy designed to diminish those problems.

Chapter 3 describes the development of the simulation tool, SimInt, that will be used to analyze the new adaptive controller. This chapter includes a verification study in which the simulation results are compared with data from a real turbine. Chapter 4 uses this simulation tool and another, SymDyn (Stol and Bir 2003), to demonstrate the proper operation of the adaptive controller.

Field data collected on a real variable speed wind turbine are presented and analyzed in Chapter 5. The field data are intended to show two objectives: first, that the goal of increasing energy capture is met, and second, that no additional harm is caused to the turbine by the adaptive

strategy. Some of the results presented in this chapter have been published by Johnson et al. (2004a) and accepted for publication (Johnson et al. 2004b).

A theoretical analysis of the original and new control strategies is presented in Chapter 6, providing further insight into each controller. Many of these results, along with the proof in Appendix B, have been submitted for publication, but a reference is not available at the time of publication of this report. Chapter 7 goes on to extend the new adaptive controller to make it even more suitable for use on commercial turbines. Conclusions and recommendations for future work are presented in Chapter 8.

Chapter 1 has provided a background on wind turbines and a summary of the relevant previous control research in addition to defining the scope and outline of this research. Chapter 2 will now explore two region 2 control strategies: the standard control and the adaptive control.

CHAPTER 2: Region 2 Control Strategies

The primary objective in region 2 control is to maximize energy capture. In general, the lower wind speeds in this region cause less stress on turbine components than the higher wind speeds of region 3, so it is generally not as important to use active control to reduce loads. This chapter will begin with an explanation of the standard region 2 controller and its shortcomings and then introduce an adaptive controller designed to alleviate those shortcomings.

2.1 Standard Region 2 Control

Region 2 control of a variable speed wind turbine is often accomplished by setting the control torque (i.e., generator torque) equal to a gain times the rotor speed squared:

$$\tau_c = K\omega^2, \quad (2.1)$$

where

$$K = \frac{1}{2} \rho A R^3 \frac{C_{p_{\max}}}{\lambda_*^3}. \quad (2.2)$$

Tip speed ratio (TSR) is defined as

$$\lambda = \frac{\omega R}{v}, \quad (2.3)$$

and C_p is a measure of the ratio of the rotor power to the power available in the wind:

$$C_p = \frac{P}{P_{wind}}, \quad (2.4)$$

where

$$P_{wind} = \frac{1}{2} \rho A v^3. \quad (2.5)$$

In general, the power P in Equation 2.4 can be defined in different ways, which can result in slightly different interpretations of C_p ; however, the most common definition considers P to be the aerodynamic rotor power:

$$P = \tau_{aero} \omega, \quad (2.6)$$




Figure 2-1: C_p vs. TSR and pitch for Controls Advanced Research Turbine (CART)

where τ_{aero} is the aerodynamic torque applied to the rotor by the wind and ω is the rotor angular speed. The theoretical maximum power coefficient C_p attainable by an ideal turbine is called the Betz Limit and can easily be derived using the Actuator Disc Concept (see, for example, Burton et al. 2001). The control law (Equation 2.1) is intended to keep the turbine operating at the peak of its C_p -TSR-Pitch surface. Equation 2.4 shows that rotor power P increases with C_p , so operation at C_{pmax} is clearly desirable. One turbine's estimated C_p -TSR-Pitch surface, derived from a simulation, is shown in Figure 2-1.

Figure 2-1 was created with the modeling software PROP (Walker and Wilson 1976), which uses blade-element-momentum theory (Burton et al. 2001). The PROP simulation was performed to obtain the operating parameters for the Controls Advanced Research Turbine (CART). The CART is a 600 kW, two-bladed, upwind turbine at the National Renewable Energy Laboratory's (NREL) National Wind Technology Center (NWTC) and is the turbine test bed used in this research and for testing other advanced control algorithms. It is equipped with many more sensors than would be found on a commercial turbine, including a meteorological (met) tower about two rotor diameters upwind in the prevailing wind direction. This met tower has a number of weather-related sensors, including anemometers at hub height and at the heights of the blade tips when the rotor is oriented vertically. As shown in Figure 2-1, it is possible—though undesirable for normal operation—for C_p to be negative; negative C_p corresponds to motoring operation.

To study the dynamics of the system, consider a simple rigid body model of a wind turbine. The equation relating net torque and angular acceleration is

$$\dot{\omega} = \frac{1}{J}(\tau_{aero} - \tau_c), \quad (2.7)$$

where

$$\tau_{aero} = \frac{1}{2} \rho A R C_q(\lambda, \beta) v^2 \quad (2.8)$$

and

$$C_q(\lambda, \beta) = \frac{C_p(\lambda, \beta)}{\lambda}. \quad (2.9)$$

While very simple, this rigid body model is a fair approximation for the CART, which has a very rigid rotor. Now, substitute Equations 2.8 and 2.1 into 2.7:

$$\dot{\omega} = \frac{1}{J} \left(\frac{1}{2} \rho A R C_q(\lambda, \beta) v^2 - \frac{1}{2} \rho A R^3 \frac{C_{p_{max}}}{\lambda_*^3} \omega^2 \right). \quad (2.10)$$

Then, simplify Equation 2.10 further with Equations 2.9 and 2.3:

$$\dot{\omega} = \frac{1}{2J} \rho A R^3 \omega^2 \left(\frac{C_p(\lambda, \beta)}{\lambda^3} - \frac{C_{p_{max}}}{\lambda_*^3} \right). \quad (2.11)$$

Since J , ρ , A , R , and ω^2 are all nonnegative, the sign of $\dot{\omega}$ depends on the sign of the difference within the parentheses. By definition, $C_p \leq C_{p_{max}}$. Thus, when $\lambda > \lambda_*$, $\dot{\omega}$ will be negative, and the rotor will decelerate toward $\lambda = \lambda_*$ (since wind is constant in this example). On the other hand, if $\lambda < \lambda_*$, $\dot{\omega}$ will be positive when

$$C_p \geq \frac{C_{p_{max}}}{\lambda_*^3} \lambda^3. \quad (2.12)$$

The curve

$$C_p = \frac{C_{p_{max}}}{\lambda_*^3} \lambda^3 = F \lambda^3 \quad (2.13)$$

is plotted as the dotted line in Figure 2-2, and the CART's PROP-derived C_p - λ curve for a fixed pitch of -1° is the solid line. Note that the solid line is simply a two-dimensional slice of Figure 2-1. Thus, the inequality in Equation 2.12 is satisfied for tip speed ratios ranging from about 3.3 to 7.5. So, as long as the CART has a tip speed ratio of 3.3 or greater, its standard control law (Equation 2.1) will cause its speed to approach the optimum tip speed ratio.

When $\lambda < 3.3$, the requirement given by Equation 2.12 is no longer satisfied, and $\dot{\omega} < 0$. In this case, the rotor speed would actually slow toward zero, and, by Equation 2.11, the turbine would remain stopped once $\omega = 0$. However, most turbines have other control mechanisms in place to




Figure 2-2: CART C_p vs. TSR and cubic function

prevent this from occurring (otherwise, the turbines would never start in the first place). The CART, for example, uses blade pitch to change the aerodynamic torque at low tip speed ratios. While the specific tip speed ratio numbers and control mechanisms are different for different turbines, the general idea of the dynamics presented here hold for all variable speed turbines.

2.2 Standard Control Law Problems

Unfortunately, modeling tools such as PROP are not perfectly accurate, and fixed controller designs based on these modeling tools are generally still sub-optimal. Figure 2-3 presents a graphical depiction of how sub-optimal performance results in energy loss.

If the dotted curve in Figure 2-3 is the actual C_p surface but the solid curve is the best guess made by the control engineer based on the modeling tools, the control torque τ_c will be designed to drive steady-state operation to the asterisk at the peak of the solid curve. Since τ_c is directly proportional to C_{pmax} as given in Equations 2.1 and 2.2, the net result is that the control torque will be higher than warranted by the dotted (real) curve. Then, considering Equation 2.7 and for a given aerodynamic torque τ_{aero} , a higher control torque τ_c will result in a smaller acceleration. All else being equal, the end result of the discrepancy between the two curves will be that the turbine will operate at a slightly lower speed (and thus tip speed ratio) than if the solid curve

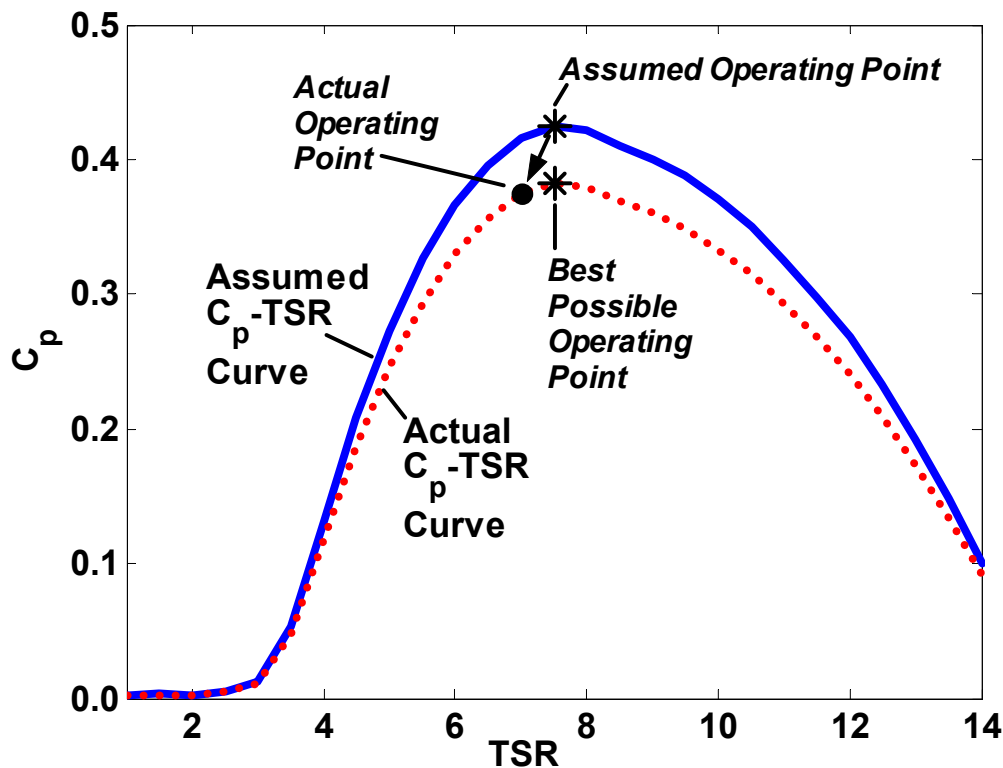


Figure 2-3: Sub-optimal operation due to inaccurate C_p -TSR curve

were correct. For this example, steady-state operation will occur at the black circle on the dotted curve, which is lower than the real system's best possible operating point, the asterisk on the dotted curve. Note that true steady-state operation, as described above, occurs only with a constant wind, which is never the case in the field.

Even if it could be assumed that the initially chosen gain K is optimal, wind turbine blades will change over time as a result of occurrences such as bug build-up and blade erosion, with the same net result as a sub-optimally chosen initial K . Fingersh and Carlin (1998) studied the effect of error in λ^* and $C_{p_{max}}$ on energy loss for a C_p curve derived from the NREL Variable Speed Test Bed turbine data. The NREL Variable Speed Test Bed turbine was a 20 kW, direct-drive turbine not directly related to the CART, but similar results would be expected on other variable speed turbines with similar C_p curves. The study showed that a very common 5% error in the optimal tip speed ratio λ^* alone can cause an energy loss of around 1% - 3% in region 2—a significant loss in this industry—for this particular turbine. For example, a wind farm rated at 100 MW and operating with a reasonable 35% capacity factor can produce about 307 GWh of energy in a given year. If the cost of energy is \$0.04 per kWh, each GWh is worth about \$40,000, meaning that a 1% loss of energy on this wind farm is equivalent to a loss of \$123,000 per year.

2.3 Adaptive Control

A simple, highly intuitive adaptive controller has been devised to address the problems with the standard control. In region 2, the adaptive control is very similar to the non-adaptive case presented in Equations 2.1 and 2.2:

$$\tau_c = \begin{cases} 0, & \omega < 0 \\ \rho M \omega^2, & \omega \geq 0 \end{cases}. \quad (2.14)$$

The adaptive gain M incorporates all of the terms in the non-adaptive torque control gain K except the air density ρ ; ρ is kept separate because it is slowly time-varying (with changes in the weather) and completely uncontrollable. M is adapted after a certain number n of time steps of operation in region 2; n is selected to be large enough to average out high frequency wind variations and the slowness of the turbine response under certain conditions. Testing on the CART thus far indicates that the averaging period defined by n will need to be on the order of tens of minutes to hours. The control law is split between positive and negative regions of ω because it is not desirable to apply control when the turbine is operating in reverse. In fact, most turbines have separate control mechanisms to prevent reverse operation, and, except where specifically noted, this research assumes positive speed operation only.

A simulation was conducted using a rigid body model relating net torque and angular acceleration as in Equation 2.7 and using Equation 2.14 for the control torque. This simulation was run for 200 seconds with each of 26 different values of the gain M , and the turbine's behavior for each of the 26 gain values was averaged in order to produce Figure 2-4. The resulting curve led to the development of the gain adaptation law described next. In Figure 2-4 and throughout the report, M^+ is the assumed optimal operating point based on the standard non-adaptive torque control law coefficient K (Equation 2.2) and the PROP-derived C_p surface in Figure 2-1. The notation M^* will be used to denote the true (often unknown) optimal gain for a given turbine.

The performance characteristic “fractional average power,” or P_{favg} , is similar to the turbine's power coefficient C_p and is defined as the ratio of the mean power captured to the mean wind power. Depending on the choice of instruments used in the calculation (in particular, the choice of anemometer), P_{favg} can be greater than the Betz Limit, so the C_p nomenclature has been avoided in order to prevent confusion. P_{favg} is computed as:

$$P_{favg}(k) = \frac{\frac{1}{n} \sum_{i=1}^n P_{cap}(k-n+i)}{\frac{1}{n} \sum_{i=1}^n P_{wy}(k-n+i)}, \quad (2.15)$$

where

$$P_{wy} = \frac{1}{2} \rho A v^3 \cos^3(\psi) \quad (2.16)$$

and




Figure 2-4: Fractional average power P_{favg} vs. normalized M

$$P_{cap} = \tau_c \omega + J \omega \dot{\omega}. \quad (2.17)$$

The first term in P_{cap} is the generator power and the second is the kinetic power (i.e., the time derivative of the kinetic energy) of the rotor. The yaw error factor ($\cos^3(\psi)$) in Equation 2.16 is a necessary component of the available power calculation because a yaw error can significantly reduce the amount of power that the turbine is physically able to capture. While active control of the turbine's yaw position is used in order to minimize the steady-state yaw error, that yaw control is not part of this research. The reason for incorporating the \cos^3 factor into Equation 2.16 is so that slowness in the yaw control does not compromise the adaptive torque control by providing the gain adaptation law with an inaccurate measure of available wind power. Verification of the \cos^3 factor is discussed in Section 3.1. In addition, P_{cap} is used in the calculation of P_{favg} rather than the turbine power P given in Equation 2.6 because the sensor requirements are more reasonable, given the instrumentation normally available on an industrial turbine. The two definitions of the turbine's captured power are closely related, differing only by the inevitable mechanical losses in the turbine's gearbox that make $P_{cap} < P$ by a small amount. Given that fact and the fact that $P_{wind} \geq P_{wys}$, it is impossible to state definitively whether $P_{favg} < C_p$ as defined in (2.4) or vice versa at any given instant. However, it can be stated that they are closely related.

The controller begins by changing M by some value ΔM . At the end of the adaptation period, the controller evaluates the turbine's performance. If the fraction of the average power P_{favg} is greater than the fractional mean power in the preceding adaptation period, the controller selects a new ΔM of the same sign as the previous one. This process continues in the same manner until the fractional power is less than that of the preceding adaptation period. At that point, the new ΔM is calculated to have the opposite sign of the previous ΔM . Eventually, M should converge toward M^* , the turbine's true optimal gain.

The mathematical equations implementing this gain adaptation are

$$M(k) = M(k-n) + \Delta M(k), \quad (2.18)$$

$$\Delta M(k) = \gamma_{\Delta M} \operatorname{sgn}[\Delta M(k-n)] \operatorname{sgn}[\Delta P_{favg}(k)] |\Delta P_{favg}(k)|^{1/2}, \quad (2.19)$$

and

$$\Delta P_{favg}(k) = P_{favg}(k) - P_{favg}(k-n). \quad (2.20)$$

In Equation 2.19, the $|\Delta P_{favg}(k)|^{1/2}$ factor is an indicator of the closeness of M to the optimal operating point. When M is such that operation is near the peak of the curve shown in Figure 2-4, a given ΔM will cause a smaller $|\Delta P_{favg}|$ due to the flatter nature of the curve near its peak. Thus, ΔM will get smaller and smaller as the optimal point is approached. The exponent $1/2$ and the positive gain $\gamma_{\Delta M}$ were chosen based on empirical results in simulation. Selection of $\gamma_{\Delta M}$ is addressed further in Section 4.3.

Because the averaging period is on the order of hours, it may seem that this simple adaptation scheme is rather slow. However, in the decades-long life of a wind turbine, several hours, days, or even weeks spent in identifying the optimal operating point are not significant in terms of power lost or time wasted. In addition, the aerodynamic changes that occur to all turbine blades typically take place on an even slower time scale than that of the adaptive controller. Thus, an adaptation period on the order of hours should not be a significant problem for a commercial turbine designer.

One main difference between the adaptive control described in this research and some of the adaptive controllers developed in previous research is which parameters are assumed to be unknown. In this research, the controller attempts to have the turbine power track the wind power but assumes C_{pmax} and λ^* are unknown. In contrast, previous adaptive controllers, such as those presented by Freeman and Balas (1999) and Song et al. (2000), force the turbine to track a desired rotor speed in region 2. However, in order to develop a model for desired rotor speed, it is necessary to have some knowledge of the C_p surface, particularly C_{pmax} and λ^* . As discussed above, the fact that those two parameters are not well known is a major source of energy loss in region 2. An additional difference among the various adaptive controllers is the lengthy averaging period used in this research, compared to the very short time periods used in previous adaptive controllers.

Now that an adaptive controller has been developed, the focus of Chapter 3 is to design the simulation tools that will be used to perform the first tests of the controller.

CHAPTER 3: Model Development

Wind turbine modeling is an ever-developing field of research that has made great strides in recent years; however, there is much yet to be learned about complex turbine aerodynamics. Modeling tools such as SymDyn (Stol and Bir 2003) and FAST (Buhl et al. 2003) can be used to model numerous degrees of freedom (DOF) for a large variety of HAWTs and have undergone validation studies (Stol and Bir 2000, Buhl et al. 2001). For example, each can be configured to allow for infinitely stiff or flexible drive trains, blades, and towers on two- or three-bladed upwind or downwind turbines. Additional DOF include nacelle yaw position and rotor angular speed. While the ability to model numerous DOF is crucial to many aspects of advanced turbine control research, it is not really necessary for the purposes of this research. In fact, a simple model that can reasonably predict changes in turbine speed with respect to wind is sufficient for testing the gain adaptation law presented in Chapter 2. The advantage of creating a simple model rather than using one of the previously validated codes lies in the fact that the simple model is able to run more quickly. This chapter describes the model created for testing this research.

3.1 Model Overview

Matlab's® modeling software, Simulink®, was used to model the turbine and controller system. The basic components of the system include the turbine model, the controller, and the wind input. A simplified diagram of the Simulink model, hereafter referred to as the SimInt (Simple Integrator) model, is presented in Figure 3-1.

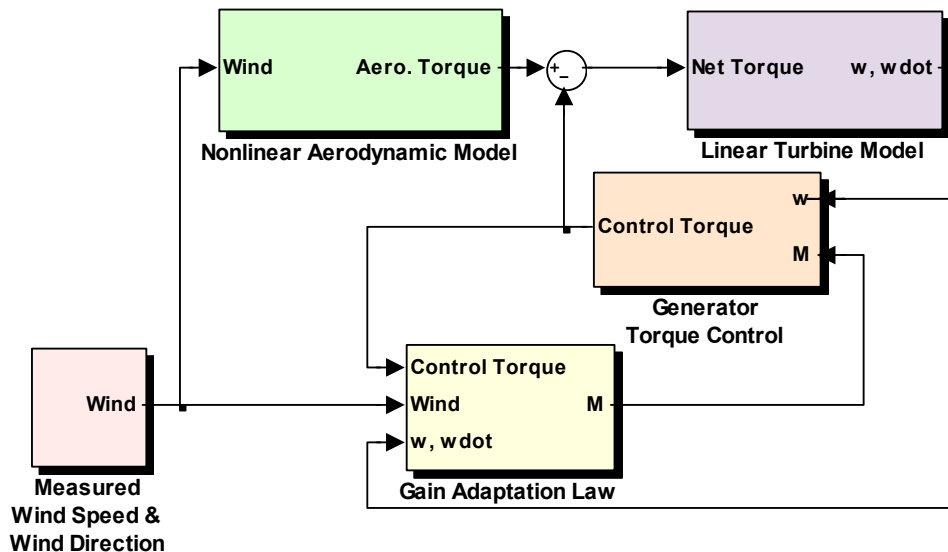


Figure 3-1: Block diagram of the SimInt Simulink model

The turbine controller is made up of two blocks. The first, labeled *Generator Torque Control*, implements Equation 2.14. The second, labeled *Gain Adaptation Law*, performs all of the calculations necessary to adapt the gain M (Equations 2.15–2.20).

The turbine model also consists of two blocks. The *Linear Turbine Model* block is simply an implementation of Equation 2.7 relating net torque to angular acceleration and then integrating acceleration to obtain angular velocity. For a large turbine operating fairly efficiently, the effects of the rotor inertia are far more important from a modeling perspective than the effects of damping, so damping terms were neglected in this model. The other part of the turbine model is the *Nonlinear Aerodynamic Model*, which uses a highly nonlinear look-up table to model the relationship between wind speed and the aerodynamic torque applied to the rotor, as given by Equation 2.8. This look-up table is the major source of the discrepancies between the model’s rotor speed output and the CART’s measured speed for a given wind input, since it is the uncertainty in this table that has made this research necessary. In fact, two different look-up tables have been used in these experiments. The surfaces generated from these tables are plotted in Figure 3-2. While the two surfaces are similar in shape, they clearly have distinct maxima and minima. It is not known which surface more accurately represents the torque coefficient (C_q) surface for the real CART, which is why both have been used in various experiments.

Either simulated or measured wind data can be used as input for the SimInt simulations, which operates at 100 Hz. This frequency was selected to mimic the controller frequency on the real CART. Although pitch control is not a part of this research, there is some pitch control implemented within the *Nonlinear Aerodynamic Model* subsystem in order to cause the model to operate as closely as possible to the real CART. This block also contains the \cos^3 yaw error factor derived in Section 2.3 and obtained from the wind direction input.

Yaw control is not a component of this research, which utilizes the existing yaw controller for the real CART experiments. However, it is important to note that a turbine operating with non-zero steady-state yaw error will generally capture less energy than one operating with zero yaw error. When a turbine has non-zero yaw error, the velocity component of the wind normal to the

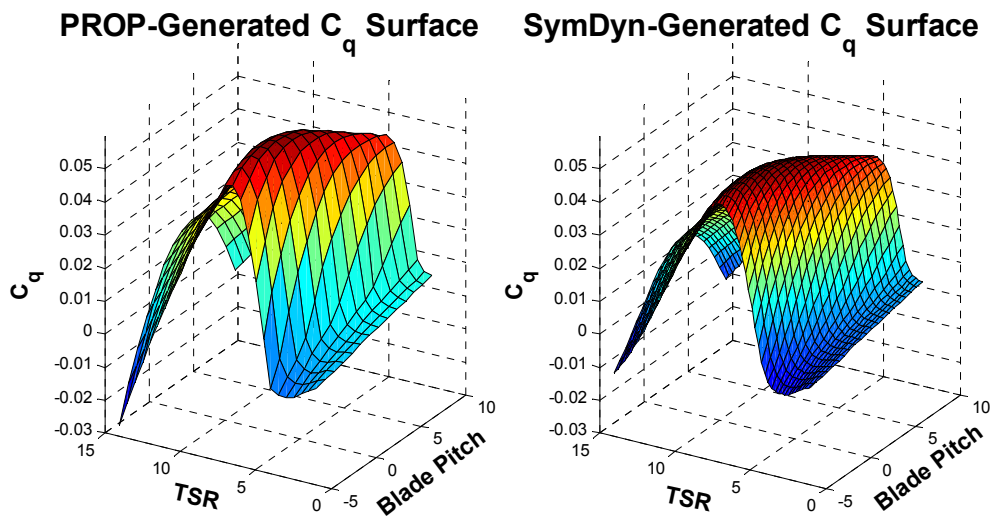


Figure 3-2: PROP- and SymDyn-generated torque coefficient (C_q) surfaces

rotor plane is less than the speed of the wind in the wind direction, which means that the actual wind available to power the turbine is less than the measured non-directional wind speed. If the rotor acted as a solid disk and wind components parallel to the rotor plane did not have any effect on the rotor, the difference in power would be a factor of $\cos^3(\psi)$, where ψ is the yaw error, as shown in Figure 3-3. This \cos^3 factor is derived from the fact that the component of the wind velocity normal to the rotor plane has a magnitude that is scaled by the cosine of the yaw error angle from the magnitude of the wind speed, as shown in Figure 3-3. The cube is a result of the fact that the power in the wind is a function of the cube of the wind speed.

In reality, wind components parallel to the rotor plane do have an effect on the rotor, so the \cos^3 simplification is not perfectly accurate. However, some experiments have been performed to assess its accuracy, and good agreement has been found for small yaw errors up to about 20° . Figure 3-4 plots the decrease in available power for a given yaw error for two cases: the \cos^3 approximation and real data collected from NREL's Unsteady Aerodynamics experiment (Fingersh et al. 2001). It should be noted that the Unsteady Aerodynamics experimental data, collected in a wind tunnel, give virtually identical results to a SymDyn simulation run with the CART's parameters. Thus, those simulation results are not plotted.

Given that a properly operating turbine will have zero mean yaw error and that fluctuations in instantaneous wind direction (and thus yaw error) will be very similar from one adaptation period to the next, these small percentages will cause no more than a very small fraction of a percent difference in the calculation of P_{wy} for any two adaptation periods. Thus, using the \cos^3 approximation to the yaw error factor, rather than a simulation-generated curve requiring an additional look-up table, is not expected to cause problems for the gain adaptation law. Although inserting a look-up table into a simulation would not be a problem, it is desirable to minimize the use of such tables in the real CART's controller. In addition, since the CART is too large for its aerodynamic parameters to have been tested in a wind tunnel, there is no real




Figure 3-3: Upwind HAWT operating with non-zero yaw error ψ (top view)




Figure 3-4: Decrease in available wind power vs. yaw error

guarantee that using a table of simulation values or data collected from an experiment with another turbine will more accurately represent the CART's aerodynamic behavior than the simple \cos^3 approximation.

3.2 Model Validation

Models can be validated in either of two ways: by comparing the model's output to the output of an already validated model, or by comparing the model's output to real turbine data. Various degrees of accuracy are required for various applications. For example, a model that will be used in the design of controls to dampen vibrations must be able to represent a turbine's structural modes fairly accurately. In the case of the SimInt model presented in this research, the model need not represent most CART DOF since it is only attempting to improve upon an existing region 2 controller that uses only the rotor speed DOF, ω . In fact, even if the rotor speed in response to wind is not perfectly modeled in comparison to real data, the simulations can still be valid tests of the adaptation law. This validity is due to the fact that the entire motivation behind the adaptive controller is that the turbine's power and torque coefficients are not well determined, and a failure to model speed precisely is a result of those uncertainties. Figure 3-5 plots the rotor speed output of two simulations using real wind speed as an input as well as the

CART's rotor speed data collected during the same time period. The two simulations use the two torque coefficient surfaces in Figure 3-2.

The rotor speed trends shown in Figure 3-5 are clearly similar between the SimInt simulations and data, though the specific magnitudes of the peaks and valleys are different. Part of the problem lies in the uncertainty of the PROP- and SymDyn-generated C_q surfaces discussed before, and part lies in the difficulty of obtaining an accurate wind measurement. In fact, the wind input data used to create Figure 3-5 was measured by an anemometer located at hub height (the height of the CART's axis of rotation, 36.6 m) 2.0 rotor diameters upwind of the CART. Not only is there a poor correlation between the wind at the small anemometer and the wind over the whole rotor plane, but the wind also requires a finite nonzero time to travel the 86.6 m between the anemometer and the rotor. The time values of the CART data in Figure 3-5 were shifted to account for the average time delay during the simulation, but it is still apparent that some of the peaks and valleys on the simulation curves occur at slightly different times than those on the CART data curve.

Another validation test was performed by comparing the output of the SimInt model with the SymDyn-generated C_q surface to the output of a CART model in SymDyn itself. Since the C_q surfaces are the same for these two models and there is no confusion over the wind measurement, it is expected that the outputs will be very close. As shown in Figure 3-6, despite the different initial speeds provided to the two simulations, this is clearly the case.

Chapter 3 has discussed the development and validation of the simulation tool SimInt. The




Figure 3-5: Validation study of SimInt model using CART data

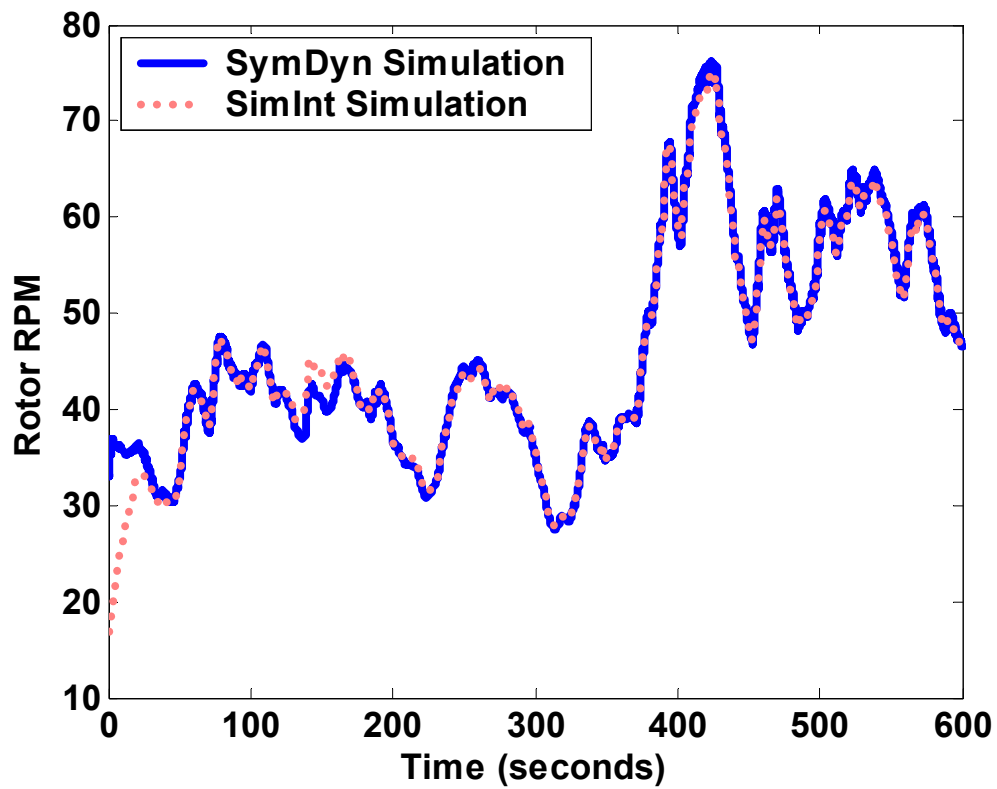


Figure 3-6: Validation study of SimInt model using SymDyn

simulation results presented in Figure 3-5 and Figure 3-6 demonstrate that the SimInt model represents the CART's behavior accurately enough to begin simulations of the gain adaptation law. Those simulations are described in Chapter 4.

CHAPTER 4: Simulation Results

Chapter 4 presents a compilation of simulation results obtained while studying various aspects of the gain adaptation law described in Chapter 2. Both the SimInt and SymDyn simulators were used in these experiments. Section 4.1 addresses the question of what result is expected if the gain adaptation law is working properly. Section 4.2 then shows the outputs of the baseline SimInt simulations, followed in Sections 4.3 and 4.4 by the results of simulations meant to test the turbine's response to real-life problems such as inaccurate wind measurements and time-varying parameters. Finally, tests performed on the adaptive controller using SymDyn are described in Sections 4.5 and 4.6.

4.1 Optimal Operating Point

As explained in Chapter 1, the need for an adaptive controller for region 2 wind turbine control stems from the fact that the optimal operating point is unknown. Since this uncertainty is mostly due to the fact that the C_p and C_q surfaces are not well known, and since these surfaces are well defined in simulations, it would seem that finding the optimal operating point for a simulation would be a simple task. The control engineer simply plugs the known values of ρ , A , R^3 , C_{pmax} , and λ_* into Equation 2.2 to find K , and the torque control law (Equation 2.1) causes the turbine to capture as much energy as possible. Unfortunately, it has recently been determined (Johnson et al. 2004a) that a given turbine does not have one single optimal torque control gain for all wind conditions, so it is not quite that simple to achieve maximum energy capture, even in simulation.

The following simulation results show that using a gain K that is 1% to 20% smaller than that given by the standard control law (Equation 2.2) actually results in improved energy capture, depending on the turbulence in the wind. The reason for this result is the fact that a turbine with high rotor inertia spends much of its region 2 operational time trying to regain the optimum tip speed ratio lost as a result of wind gusts and lulls; it simply cannot track wind speed perfectly. Since rotor inertia is included in the simulation, this result applies equally to simulations and real turbines. Because the power in the wind is proportional to the cube of the wind speed, it is more important to track wind gusts than lulls. This concept is explained further throughout this section.

As shown in Figure 4-1, a turbine with a very low rotor inertia ($J = 1000 \text{ kg}\cdot\text{m}^2$) has a much different optimal operating point than one like the CART, with $J = 388,500 \text{ kg}\cdot\text{m}^2$. In Figure 4-1, "Nominal Power Capture" is the maximum power captured by the turbine with very low rotor inertia at $K/M^+ = 1.0$, all other physical parameters used in the simulation being the same. In this plot, three different 100-Hz sampled wind data sets were scaled to have the same means (7.533 m/s) and used as inputs in simulations. The curves were each created by simulating the CART's behavior 24 times, with each simulation lasting one hour and having a different (fixed) value of the torque control gain, K . The nominal value, M^+ , is given by




Figure 4-1: Average power capture for normal and low inertia

$$M^+ = \frac{1}{2} AR^3 \frac{C_{p \max}}{\lambda_*^3}, \quad (4.1)$$

which is simply the equation for K (Equation 2.2) with the time-varying air density ρ removed. For the low rotor inertia ($J = 1000$) case, the power capture is virtually identical no matter which wind is used as input—the lines are indistinguishable on the scale of Figure 4-1. However, there are noticeable differences in the three plots showing the normal inertia ($J = 388,500$) case. In these simulations, optimum power capture (indicated by the stars) is achieved at anywhere from 89% to 93% of the nominal optimal control torque gain M^+ . According to these three curves, it is clear that captured power for this turbine model could be increased by 0.5% simply by reducing the gain K by an average of 10% below the nominal value. While 0.5% may not sound like an impressive number, the control modification required is so simple, costing nothing, that there is no question that this change is worthwhile in this case.

Why does the optimum torque gain differ for different wind inputs? Since the three curves from the low inertia simulations lie almost on top of each other and only the high inertia curves differ by wind input, it seems likely that the rotor inertia is the cause of the change in optimum gain from one wind speed to another. A higher inertia will result in slower tracking of the optimal tip speed ratio λ_* in response to wind gusts and lulls. When characterizing the wind, turbulence

intensity is one measure used to describe the amount of “gustiness” in the wind. It is usually defined as the standard deviation of the wind speed divided by the mean wind speed over 10 minutes. The following simulations, presented in Figure 4-2, show that the optimum operating point is closely related to a slightly modified measure of the wind’s turbulence intensity. This modified one-hour turbulence intensity uses 60 minutes of data rather than 10 in the mean and standard deviation calculations.

Figure 4-2 was created in a manner similar to that of Figure 4-1. Two of the wind data sets used to create Figure 4-1 were slightly modified to create 15 different standard deviations per wind input, resulting in 15 different one-hour turbulence intensities. The simulations were then run for each of these 15 one-hour turbulence intensities per wind data set, holding the mean wind speed constant throughout the experiment. Each simulation was one hour long and used the normal inertia and other CART parameters. As shown by the “Wind Data Set 2” curve, for high one-hour turbulence intensities, the torque gain resulting in the greatest power capture can be nearly 20% lower than the nominal value given by Equation 4.1.

Similar simulations performed while varying the mean wind speed and holding the one-hour turbulence intensity constant showed that the mean wind speed does not significantly affect the optimum torque control gain. Thus, it can be concluded that it is the standard deviation of the wind speed that forms the correlation between turbulence intensity and optimal torque control gain. The fact that the two curves in Figure 4-2 do not lie on top of each other indicates that other wind characteristics also play a role. For example, the Richardson number, a measure of




Figure 4-2: Optimum torque gain vs. turbulence intensity

atmospheric stability, and other factors may also affect the optimum gain, but these are not investigated in this paper.

Why is the optimal torque control gain lower for higher turbulence intensities? Consider this example: imagine that the constant wind has a speed of 10 m/s, and the CART is operating in steady state with $\omega = 3.46$ rad/s (for a tip speed ratio of $\lambda^* = 7.5$). At some time, there is a step change in wind speed to 11 m/s, which causes the instantaneous tip speed ratio to drop to $\lambda = 6.82$. The torque control law (Equation 2.1) will cause the rotor to accelerate toward $\lambda = 7.5$, as described in Section 2.1, but in the meantime, the power loss is proportional to the cube of the wind speed, or 1331 (m/s)³, by Equation 2.5. Similarly, if there is a step change in wind speed down to 9 m/s, the power loss is proportional to $9^3 = 729$. Compared to the 1000 (m/s)³ baseline, the step increase in wind speed results in a 33% gain in power available, whereas the step decrease results in only a 27% loss in available power. Because of the cubic law relating wind speed and power (Equation 2.5), it is more important to regain the optimal tip speed ratio quickly following wind gusts than following lulls. Since it is safe to assume that, in general, a given wind input will have a similar number of gusts and lulls, using a torque control gain that is slightly smaller than the one given by the standard gain K (Equation 2.2) is more likely to make the rotor acceleration (Equation 2.7) positive, which means the rotor will accelerate more easily in response to a gust.

It is clear from Figure 4-2 that turbulence intensity has a significant effect on the optimal torque control gain K . However, this experiment did not consider the effect of changes in mean wind speed on the optimal torque control gain. Additional simulations have been performed for two cases: in the first, the turbulence intensity was held constant while the mean wind speed was varied, and in the second, the standard deviation of the wind speed was held constant while the mean wind speed was varied. To clarify, in the first case, the standard deviation of wind speed was varied, while in the second, since standard deviation was constant, turbulence intensity was varied by definition. In the first case, the optimal torque control gain changed only slightly (about 2% on average) over a 5 m/s to 12 m/s range of mean wind speeds (approximately the region 2 operational range for the CART). In the second case, the optimal torque control gain varied by about 11% on average over the same range of mean wind speeds.

Several researchers, including Casanova Alcalde and Freris (1982) and Pierce (1999), have published research describing the effect of mean wind speed on a turbine's time constant. The time constant is inversely proportional to the mean wind speed, so higher mean wind speeds result in smaller time constants. Thus, under near constant wind conditions, a turbine that is disturbed by a 1 m/s step change in wind speed should regain its optimal tip speed ratio more quickly for a higher mean wind speed than for a lower. As a result, one might argue that mean wind speed should have a significant effect on the optimal torque control gain. However, a closer examination reveals that this analysis is consistent with the preceding analysis on the effects of turbulence intensity, since mean wind speed is a factor in turbulence intensity. In fact, that 1 m/s step change would provide a higher turbulence intensity for the lower wind speeds than for the higher. In effect, our simulation results show a "near cancellation effect" among the mean wind speed, standard deviation, and time constant as far as the optimal torque control gain is concerned. The precise cause for this effect is beyond the scope of this research, but one possible explanation is that even though the optimal tip speed ratio is regained more quickly at higher wind speeds for a gust of a given size, proportionally sized gusts (i.e., turbulence intensity

held constant for all mean wind speeds) will result in larger deviations from optimal tip speed ratio and greater potential for energy loss due to the additional energy available at these higher wind speeds. The 1%-3% deviation in optimal torque control gain K for different wind speeds with constant turbulence intensity is an area that deserves further study. However, it is small compared to the 13%-18% deviation apparent in Figure 4-2 for changes in turbulence intensity and is not considered in the remainder of this paper.

Since the adaptive controller can begin with any initial choice of M , it is not necessary to perform an independent study of turbulence intensity in order to achieve optimal operation. The adaptive controller is able to take the turbulence intensity into account as the gain adaptation law adapts M over time, eventually seeking out the optimal point M^* , assuming the turbulence intensity is somewhat consistent between adaptation periods. However, in simulation results for which very turbulent winds are used as input, M is expected to adapt to a value 10% to 20% less than the standard value M^+ given by Equation 4.1 when the C_q surface used in the simulation is the same as the one used to determine the parameters in Equation 4.1, i.e., the PROP-derived C_q surface. When the simulation uses another C_q surface in the calculation of aerodynamic torque (Equation 2.8), M may adapt to an entirely different value, depending on the new surface's parameters.

4.2 SimInt Simulation Results

The first simulations use a constant wind speed and constant, zero yaw error as the wind input. In this case, the turbulence intensity equals zero and the PROP-derived C_q surface is used, so M is expected to adapt to M^+ exactly. One of the simulation outputs, which used a constant wind input, is plotted in Figure 4-3, with simulation parameters listed in Table 4-1.

Table 4-1. Simulation Parameters

Parameter	Value [†]
J	388,500 kg•m ²
R	21.65 m
M^+	174.48 m ⁵ /rad ³
γ_{AM}	23.17 rad ³ /m ⁵
ρ	1.0 kg/m ³
f_s	100 Hz

[†] These values are referenced to the high-speed shaft. On the rotor side of the gearbox, $M^+ = 7530.41 \text{ m}^5/\text{rad}^3$. and $\gamma_{AM} = 1000 \text{ rad}^3/\text{m}^5$.

Figure 4-3 uses an adaptation period of 20 minutes, meaning that n , the number of steps between adaptations, is equal to 120,000. Clearly, the bulk of the changes in this particular case occur in the first dozen adaptation periods. After that, the value of M oscillates near the expected value, M^+ . Because the adaptation scheme uses only the most recent past value in addition to the current value of the adaptation parameters, these oscillations are not unexpected. Chapter 7 will investigate incorporating additional past values to reduce the oscillations. In Figure 4-3, both M and ΔM are normalized by M^+ (Equation 4.1), which, according to Figure 2-1, is nominally the ideal value.

When it comes to wind turbine control, showing that a controller works for a constant wind input provides very little information about whether it will work under a time-varying wind input, though a controller that fails under constant wind is very unlikely to succeed under varying winds. Thus, the next step is to experiment with time-varying wind inputs.

Figure 4-4 was created by a SimInt simulation that used measured wind speed data as the wind input to the simulation and the PROP-derived C_q surface. After an initial step in the wrong direction—caused by the initial conditions provided to the simulation—it is clear in the top plot of Figure 4-4 that M converges toward a value of around 85% to 90% of M^+ . In the process, the P_{avg} curve appears to be seeking its maximum, discounting the first value, which was the unattainable initial guess provided by the control engineer. The turbulence at the NWTC during the winter wind season tends to be higher than the average at a typical wind site; thus, it is expected that using measured wind speeds will result in M converging to some value




Figure 4-3: Simulation results with constant wind input

approximately 15% lower than M^+ . Further, it is clear that convergence under real wind input does not occur as quickly as it does with constant wind, though it seems that the system is stable in some sense. It is also important to note that Figure 4-4 required the use of a much larger n than Figure 4-3 did, as a result of the time-varying wind input. In the case of Figure 4-4, $n = 4,680,000$, which translates to 13 hours. All other parameters shown in Table 4-1 are the same for both simulations. One final clarifying point remains to be made regarding Figure 4-4 and similar ones that follow. Each value on the Fractional Mean Power curve is the P_{fav} corresponding to the normalized M from the *preceding* time period. Thus, the second value of P_{fav} , 0.37, is the fractional mean power collected by the turbine during the first adaptation period when $M = 1.5M^+$, which is computed at the end of the first time period.

The important things to realize regarding the simulation outputs are that the final value of M appears to be close to the predicted optimal torque control gain value for the given turbulence intensity and that M appears to be stable if not convergent. These two results suggest that the adaptive controller is working as desired. The same data plotted as time-series in Figure 4-4 can be plotted against each other, as in Figure 4-5, rather than against time to help explain the oscillations that occur from about 150 to about 350 hours.

Figure 4-5(a) contains all the P_{fav} and M data used in the creation of Figure 4-4 and includes a best-fit quadratic and its equation. The correlation in the data is high, with $R^2 = 0.997$. At this level, it appears that a 13-hour adaptation period is long enough to eliminate most of the negative effects of noise. This noise is believed to be mostly a result of the difficulties in measuring the




Figure 4-4: Simulation results with real wind data as input




Figure 4-5: Fractional average power P_{favg} vs. normalized M (M/M^+) for simulation data: (a) all data; (b) zoomed in around peak

wind hitting the rotor. Figure 4-5(b) is simply a zoomed-in version of (a) in which the noise effects are much more obvious. Though the data still follow the trend line, some data points are much farther away from it than others. Based on numerous simulations, it has become clear that the longer the adaptation period, the more zoom is necessary in order to observe the effects of noise in the data. Thus, when the adaptive gain M gets closer to its optimal value M^* , longer adaptation periods might be necessary in order to achieve nicely convergent results.

For simplicity and computational ease, the simulations used to create Figure 4-3 through Figure 4-5 were performed assuming a zero yaw error. However, experimentation on the real CART will not allow that simplification, so it becomes important to study the effects of yaw error on the simulation outputs. Figure 4-6 shows a plot of simulation outputs for three cases: one with large yaw errors provided to the simulation, one with zero yaw errors, and one halfway between the two extremes. For this simulation, the SymDyn-generated C_q surface was used, but M^+ is still calculated from the PROP-generated surface. The difference in C_q surfaces explains the lower normalized M values apparent in the upper plot.

Figure 4-6 was created using three 44.5 hour simulations having the same wind speed input and the same initial conditions but different yaw error inputs. The “Full Yaw Error” simulation involved very large yaw errors up to 79° , with a mean absolute value of about 10° . Positive and

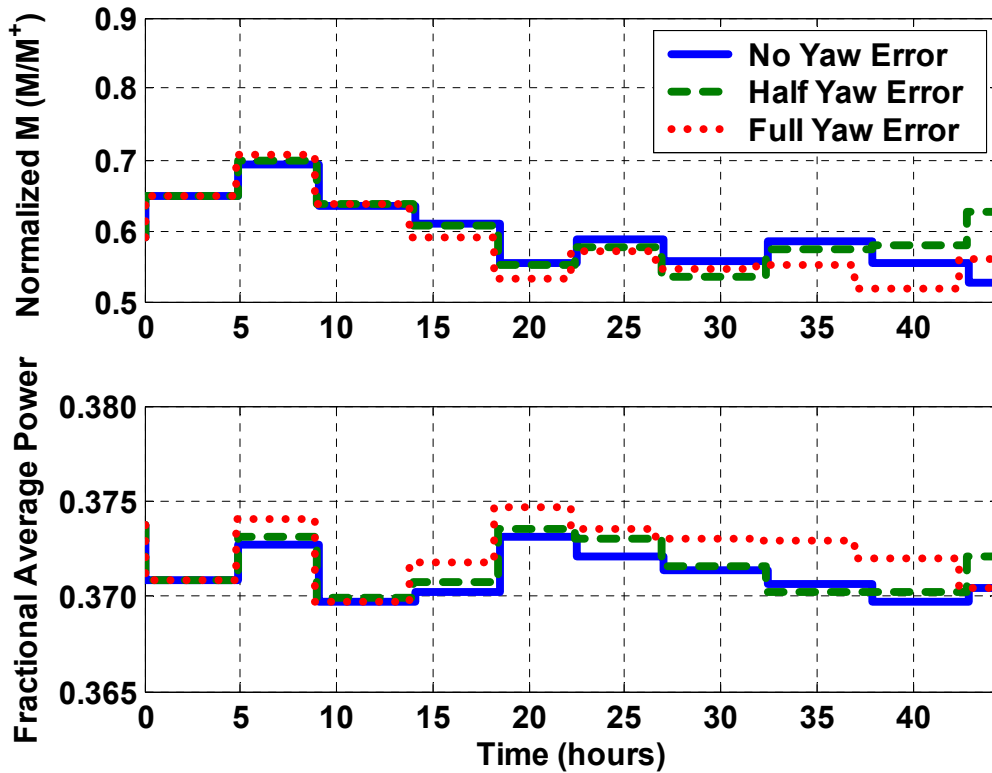


Figure 4-6: Adaptations of M for varying amounts of yaw error

negative yaw errors of the same magnitude were assumed to be equivalent. The adaptation period used for the simulations was four hours, or $n = 1,440,000$. Note that the adaptation period consists only of times during which the turbine is operating in region 2; thus, when region 1 or 3 operation occurs, the adaptation algorithm pauses and waits for the winds to increase or subside so region 2 operation can continue. During the first five hours of the simulation, it is apparent that about an hour of region 1 or 3 operation occurred, since M adapted approximately five hours after the beginning of the simulation rather than four.

As explained in Section 3.1 and shown in Figure 3-4, the \cos^3 yaw error factor approximation is only very good up to about 20° of yaw error. Approximately 11% of the yaw error data used in the “Full Yaw Error” simulation exceeds this value, which means that the P_{wy} calculations given by Equation 2.16 will be less accurate for this simulation. The “Half Yaw Error” simulation was run with the “Full Yaw Error” yaw error data scaled by one half, and only 0.4% of the yaw errors in this simulation exceed 20° . As shown in Figure 4-6, the adaptations of M for the “Half Yaw Error” case closely resemble those in the “No Yaw Error” case, whereas those in the “Full Yaw Error Case” diverge more quickly. From these simulations, it becomes clear that a real turbine using this gain adaptation law will require a fairly good yaw controller in order to prevent the yaw error from affecting the gain adaptation law. However, it is not unreasonable to assume that the average magnitude of the yaw error will be less than 10° for a properly operating turbine.

In summary, simulations have shown that the gain adaptation law proposed in Chapter 2 will cause the adaptive gain to trend toward convergence, thereby maximizing the fractional average power captured by the turbine, assuming that the adaptation period is long enough and yaw control is adequate. These simulations have also assumed that the wind measurement is accurate and the turbine parameters are constant. The following sections present simulations to address the outcomes when these last two assumptions are invalid.

4.3 Wind Measurement Problems

Section 4.2 provides simulation results demonstrating that the gain adaptation law performed fairly well under the condition that the wind was measured correctly. Since there is no assurance that this would be true for a real turbine, it is important to simulate the effects of inaccurate wind measurements.

The wind measurement at any point will not be equal to the wind acting on the rotor blades in space or time. Since it is impossible to obtain precise wind measurements, the question becomes: “What types of imprecisions will cause the adaptive controller to fail?” Figure 4-7 shows some simulation results that help to answer that question. The upper plot shows the normalized adaptive gain M over time, while the lower plot shows the error between the true wind power and the erroneous wind power measurement for each case. The root mean square (rms) wind power in each 44.5 hour simulation is 768 kW, and the power error in the “Offset Error” case was chosen to be equal to the rms value of the error in the “Noise Error” case, 54.3

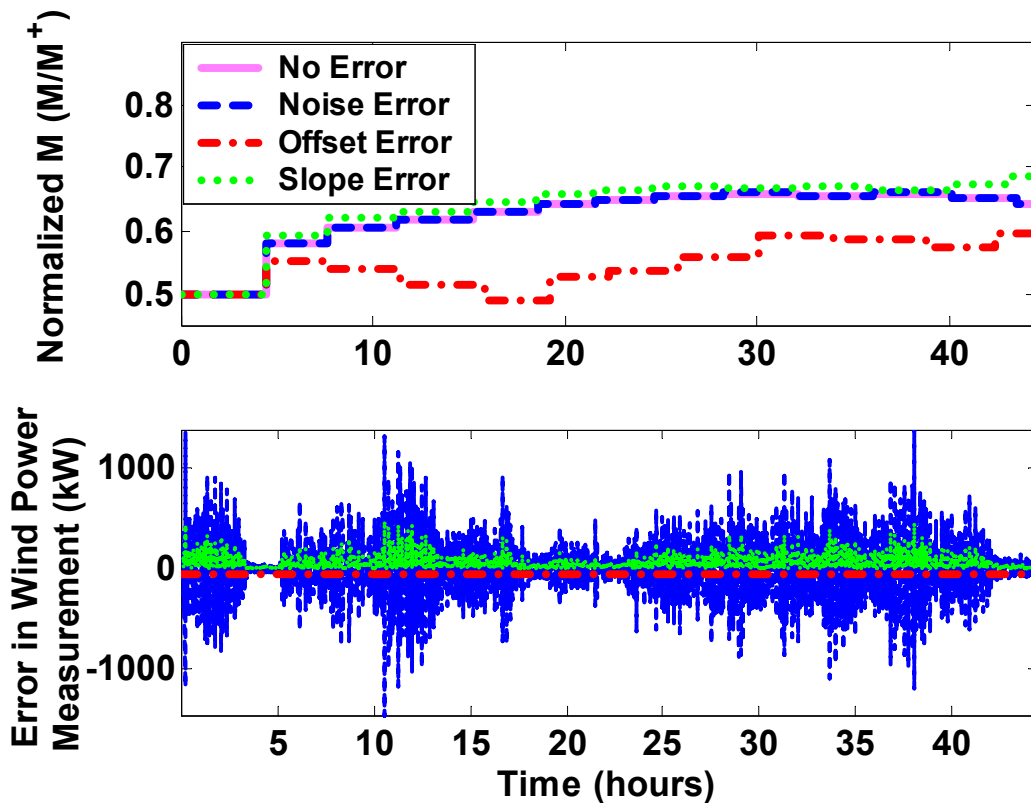


Figure 4-7: Simulations with erroneous wind measurements

kW, or 7.1% of the rms wind power. The rms error in the “Slope Error” case is also 54.3 kW, which is achieved by setting the measured wind power equal to 92.9% of the actual wind power.

In the interest of reducing computational time, the adaptation period used to run the simulations plotted in Figure 4-7 is three hours ($n = 1,080,000$) rather than 13 hours. The gain adaptation law is not expected to cause convergence of M as well or as quickly when this shorter adaptation period is used, but it is sufficient for the purposes of this experiment. In this case, the result of interest is the difference in adaptations of M among the four cases.

These simulations were performed by corrupting the wind power calculation P_{wy} given by Equation 2.16. In each case, the same wind speed input was provided to the aerodynamic torque calculation (Equation 2.8), but the calculated wind power was corrupted by band-limited white noise, by multiplying by a constant scaling factor, or by the addition of a constant offset, respectively. In other words, the v used to calculate τ_{aero} (Equation 2.8) was the true wind speed, but the P_{wy} calculation (Equation 2.16) was corrupted. The point of the experiment is to test whether M adapts to the same value when the wind measurement is corrupted.

The important result of these experiments is that, if there is white noise in the wind speed measurement, the gain adaptation law performs as well as if there is no corruption of the wind measurement. In fact, the two plots lie virtually on top of each other. This result is to be expected: in the case of white noise, it is reasonable to expect that the mean value of the noise will be approximately zero over the length of the adaptation period, so each P_{favg} (Equation 2.15) measurement will be similarly corrupted. A similar explanation also holds for the case of the scaled wind speed measurement; in this case, the step size ΔM given by Equation 2.19 will simply be scaled by the same value as the wind speed measurement, leading to a slight mismatch in step size when compared with the “No Error” case but still converging to nearly the same value. However, if there is an offset in the wind speed measurement, such as might occur with a miscalibrated anemometer, the gain adaptation law performs less well, with a noticeable difference occurring as early as the first adaptation. Since an offset in the anemometer calibration would be relatively easy to detect (i.e., by checking to see if the instrument is reading some nonzero value in the absence of wind), it is reasonable to assume that this type of wind measurement corruption can be avoided. Scaling and noise errors are much more likely to occur in the field.

4.4 Time-Varying Parameters

One other real-life problem that challenges turbine control design is time-varying aerodynamic parameters, specifically C_p , C_q , and λ . The adaptive controller has been designed not only to determine the optimal operating point initially, but also to track the changes in these parameters over time and account for them in the control law. Once again, several simulations were conducted in order to test the ability of the gain adaptation law to deal with these changes. The first of these, plotted in Figure 4-8, uses the PROP-derived C_q surface with a constant wind input and assumes that there is a sudden step decrease in the turbine’s power coefficient C_p . This sort of sudden change could occur if, for example, a blade were struck and damaged. The simulation assumes that only the magnitude of C_p is affected but that the peak still occurs at λ^* .




Figure 4-8: Simulated change in parameters with constant wind speed input

It is clear from Figure 4-8 that, for a step change in the turbine's power coefficient and a constant wind speed input, the gain adaptation law works quite well. In fact, the top plot shows that, after an initially very large step caused by the initial conditions, M adapts toward its optimal value, which is exactly M^+ for this constant-wind simulation, which uses the PROP-derived C_q surface. Then, M oscillates around M^+ with small steps until the parameter change occurs at 12.5 hours, after which it seeks the new value. It is apparent that the 5% reduction in C_p leads directly to a 5% reduction in M , which is expected because, for the zero turbulence intensity case, the optimal torque control gain is given by M^+ , which is directly proportional to C_p (Equation 4.1).

Next, a variable wind speed input simulation was conducted for both step and ramp changes in the turbine's power coefficient, this time using the SymDyn-generated C_q surface. Figure 4-9 is a plot of this simulation's output, which shows a total of four different parameter changes over more than 1200 hours of simulation time. The nature of the changes is summarized in Table 4-2, followed by a more thorough explanation of the performance of the adaptive controller following each change.




Figure 4-9: Simulated changes in parameters with variable wind speed input

Table 4-2. Summary of Parameter Changes Used in Figure 4-9 Simulation

Change	Time (Hours)	Description
#1	150.17	Step change in τ_{aero} from 100% to 90%
#2	356	Begin ramp change in τ_{aero} starting at 90% and with a slope of -0.11% per hour (-984.25% per year)
#3	578.5	Begin ramp change in τ_{aero} starting at 75% and with a slope of $+0.023\%$ per hour ($+196.85\%$ per year)
#4	890	Begin ramp change in τ_{aero} starting at 82% and with a slope of -0.00013% per hour (-1.18% per year)

In Figure 4-9, change #1 is a step decrease in aerodynamic torque just like the one used in the Figure 4-8 constant wind input simulation. In the variable wind speed input case, note how the adaptive gain M was oscillating by small steps around its optimal value prior to the change, while the fractional average power also seemed to be seeking its maximum. Then, the step decrease occurred, which caught the gain adaptation law by surprise, causing a significant step upward in the wrong direction. This step in the wrong direction is a result of timing and chance: the previous adaptation was in the downward direction, so when the C_p surface was decreased suddenly in the middle of the adaptation period, the adaptive controller reacted as if the previous downward step were to blame and adjusted upward accordingly. Of course, the significant

change in energy capture was caused by the external change, and it did not take long for the gain adaptation law to track the change and begin oscillating near the new optimum value. It is also clear from the lower plot that the fractional average power is seeking a value 10% less than it was before the step change occurred.

Once M seemed to be settling near its new optimal operating point, the system was perturbed again with a ramp change in aerodynamic torque beginning at the time marked “#2.” The slope of this ramp is very large compared to what most turbines would be expected to endure—nearly 1000% per year—particularly for such a long time period. One situation that might cause such significant decreases in C_{pmax} would be ice build-up on the blades, but it is unlikely that an ice storm would continue for 200 hours or more. In this case, the gain adaptation law did not perform as well as it did for the step change case. During all of the adaptation periods, the turbine captured less energy than it had during the previous ones, which caused the adaptive gain to switch directions at every step. The adaptive gain does decrease on average, as it should, but there does not appear to be any trend toward convergence of M to a steady-state value. In fact, since the power coefficient C_p is rapidly decreasing, M should not be converging on one particular value but rather attempting to track the changes in C_p .

Change #3 occurs shortly before 600 hours and signifies the end of the decreasing power coefficient. At this time, a positive ramp was applied in the simulation with a smaller—but still relatively large—magnitude slope. This type of slope might occur as the ice slowly melts from the blades in near-freezing temperatures. In this case, because the power coefficient is constantly increasing, the gain adaptation law tends to keep adapting M in the same direction as the previous adaptation, erroneously responding as if the adaptation had caused the increased energy capture. The direction of the adaptation changes only when M adapts far enough over the peak of the C_p curve that it causes a decrease in energy capture despite the constantly increasing power coefficient. The Fractional Average Power plot reflects the steady increase in C_p .

While the performance of the gain adaptation law following changes #2 and #3 is not superb, real turbines are unlikely to endure ramp changes with slopes of these magnitudes over such long time periods. The more realistic scenario would be for the turbine to undergo a few hours of icing followed by melting, during which time the gain adaptation law might have trouble seeking its optimum value. However, once the ice has melted, changes such as bug build-up and blade erosion will occur on a much slower time scale, better suited for the gain adaptation law. The ramp change beginning at #4 is an example of a more realistic scenario in which the turbine undergoes a gradual decrease in power coefficient at a rate of about 1% per year. During this time period, the gain adaptation law causes M to seek its new optimal value, which is changing on a time scale so much longer than the 13-hour adaptation period that M is once again able to approximate convergence by the end of the simulation.

The variable speed wind input simulations demonstrate the gain adaptation law’s strengths and weaknesses. The gain adaptation law performs well for discrete (sudden) changes in the turbine’s aerodynamic parameters and for those that show significant changes only over very long time scales relative to the adaptation period. However, for situations in which the parameter changes are noticeable over a time scale similar to the adaptation period, the performance of the gain adaptation law is noticeably compromised. Fortunately, these

moderately fast changes typically don't go on for very long and don't cause M to become unstable.

4.5 Simulations Using SymDyn

All of the previously described simulations were performed using SimInt, which, as explained earlier, includes only the rotor angular speed DOF and is thus capable of running more quickly than more complex simulation tools such as SymDyn and FAST. However, for verification purposes, it is valuable to compare the gain adaptation law's operation in SimInt with its operation in another, more well-known simulation tool. The adaptive gain is not expected to adapt equally from one simulation tool to the next, since each has its own distinct turbine model. However, if each is modeling the same real turbine—in this case, the CART—the adaptive gain should adapt similarly in both SimInt and SymDyn. This similarity is evident in Figure 4-10.

Both the SimInt and SymDyn simulations were started at time 0 with the same initial values of M , ω , and most other parameters. However, although the two turbine models were set up to have the same true optimal torque control gain M^* , the C_p surfaces have peaks with different magnitudes; therefore, in order to cause the initial step to be approximately the same, the initial P_{favg} values were the same percentage of maximum rather than the same absolute value. The upper plot in Figure 4-10, showing normalized M , shows a reasonable amount of agreement. There is a time period between about 30 and 60 hours during which the adaptive gains M for

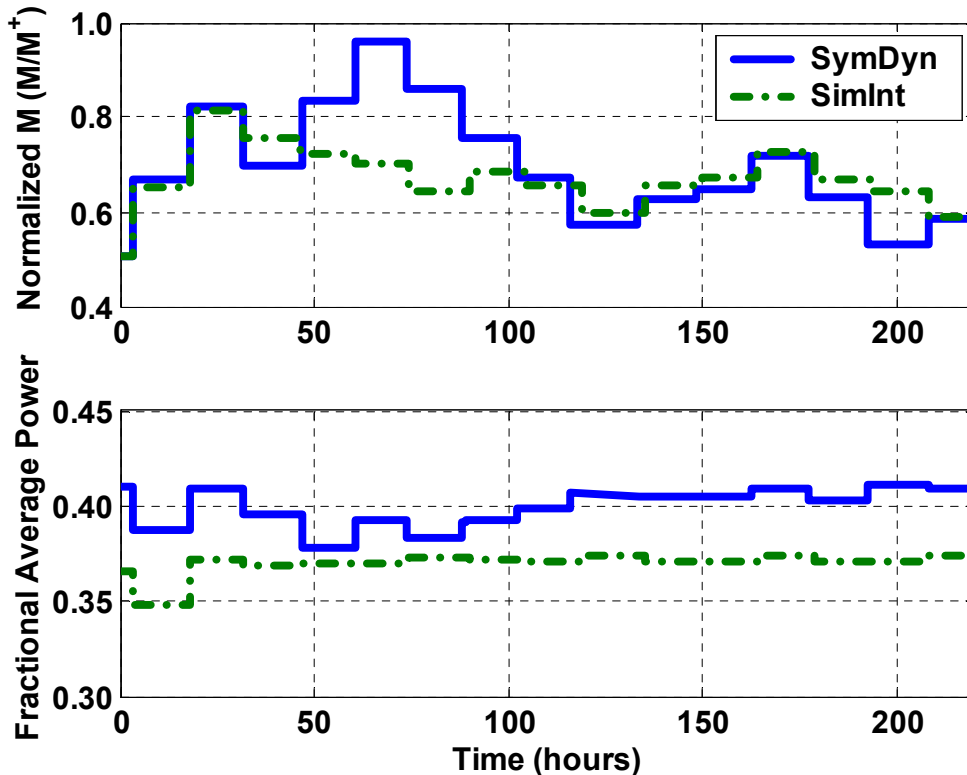


Figure 4-10: Comparison of SimInt and SymDyn outputs

each simulation appear to diverge, but after that time they once again approach each other and remain close together from about 100 hours through the end of the simulation. It is also clear from the lower plot that each simulation is adapting its gain M and seeking maximum power capture.

Since SimInt simulations run approximately five times faster than SymDyn ones, the SymDyn simulation presented in Figure 4-10 was the only one performed as a verification of the gain adaptation law; after this experiment showed reasonable agreement between the two simulation tools, SimInt remained the simulation tool of choice for all further experiments designed to test the gain adaptation law. However, unlike SimInt, SymDyn is capable of analyzing additional DOF besides generator speed. Thus, the SymDyn simulation data were analyzed to determine whether the gain adaptation law might be likely to cause additional stress on the turbine. This analysis is the topic of the next section.

4.6 Simulated Stresses on Turbine Components

In addition to generator speed, SymDyn is capable of simulating many other degrees of freedom, including drive shaft compliance angle, blade flap angles, tower twist angle, and tower side-to-side and fore-aft angles. The ability to model these DOF is a key to SymDyn's usefulness in control design and testing. In tests of the adaptive controller within SymDyn, the question to be answered is whether the adaptive generator torque controller (Equation 2.14) causes additional turbine transient stresses that are absent when the turbine is controlled by the standard non-adaptive generator torque controller (Equation 2.1). Long term, steady-state changes in loading are discussed further in Section 5.3 but are difficult to analyze over the time spans used in these simulations and will not be considered in this section. However, there is a chance that an abrupt change in generator torque could excite an unknown number of modes each time an adaptation occurs. Therefore, this section will focus on SymDyn data collected immediately before and after adaptations.

It appears from the SymDyn experiments that most of the DOF mentioned above do not undergo much additional stress as a result of the adaptation of the torque control gain M . The following three plots, Figure 4-11 through Figure 4-13, provide a sample of the experimental data consisting of three DOF (drive shaft torsion, blade flap bending, and tower side-to-side bending) for two different adaptations occurring during the SymDyn simulations plotted in Figure 4-10. Table 4-3 provides some information relevant to each case. Each adaptation involves an upward change in M , though the direction of the change did not appear to affect the nature of the stresses for other adaptations examined. In addition, both occurred at approximately the same rotor speed. The maximum rotor speed for the CART is 4.37 rad/s, so 2.9 to 3.0 rad/s is well within region 2 operation. Finally, Adaptation #4 occurs when the rotor is slowing down, while Adaptation #7 occurs while it is speeding up. Rotor acceleration is closely related to wind acceleration, and, in general, higher winds cause more stress on turbine components.

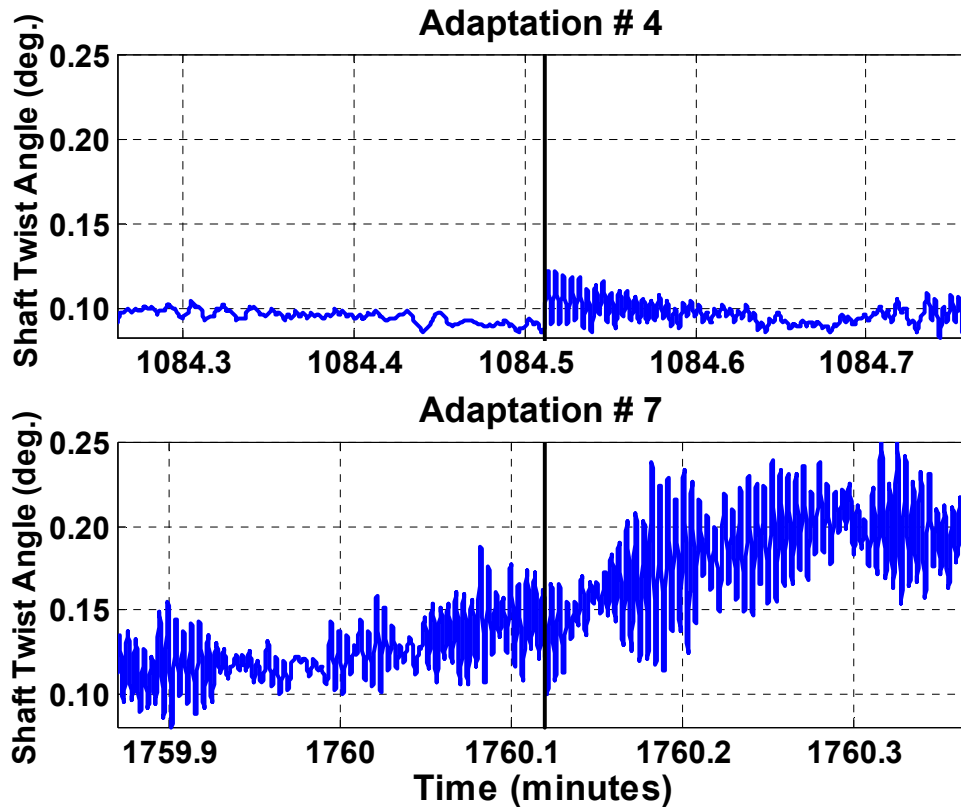


Figure 4-11: Drive shaft twist angle for two different adaptations

Table 4-3. Selected Details Relevant to Figure 4-11 through Figure 4-13

Adaptation	Normalized M (M/M^*)	Rotor Speed	Rotor Acceleration
#4	0.67 to 0.82	2.9 rad/s	< 0
#7	0.84 to 0.96	3.0 rad/s	> 0

Of the DOF that SymDyn is capable of modeling, the drive shaft twist angle is the one most closely tied to the generator, since it connects the generator to the rotor. Thus, the most significant effects of the torque control gain adaptation are expected to be seen in the shaft, and this is indeed the case. Figure 4-11 is a plot of the shaft twist angle for the two adaptations.

In the top plot of Figure 4-11, Adaptation #4, it is clear that the drive shaft experiences an undesirable response to the adaptation, which occurs at time 1084.51, marked by the vertical line. However, there are several reasons why this response is not as worrisome as it first appears. First, the oscillations induced by the adaptation decrease within a few seconds, indicating that this is likely a stable mode. Second, the maximum magnitude of the oscillations—about 0.3° —is still smaller than the oscillations apparent in the lower plot (Adaptation #7) even before the lower plot’s adaptation time, which indicates that the adaptation does not cause oscillations more severe than those produced by environmental variables. Finally, in the simulation, the adaptation is modeled as occurring within one time step, or 0.01 seconds. In the case of a real turbine, it is unlikely that it would be possible to change the generator torque that rapidly, so the real turbine’s

shaft would not see such a large step change at all. Thus, the shaft twist angle oscillations are something to watch closely when implementing the adaptive controller on a real turbine, but do not represent a reason to decide against implementation.

Figure 4-12 plots the flap angle of each blade for the 15 seconds before and after each adaptation. “Flap” is the out-of-plane blade bending direction; for a horizontal axis turbine, flap bending typically occurs in the direction of the wind as the wind exerts a lift force on the blade. As is the case with Figure 4-11, it appears that the environmental conditions occurring during the time period around Adaptation #7 were more turbulent than those occurring around Adaptation #4, since the blade flap angles are larger and undergo more variation. The blade flap angles plotted for Adaptation #7 do increase somewhat after the adaptation, but an increase in mean flap angle is to be expected for a faster spinning rotor. Based on these representative plots, a sudden adaptation is not expected to cause undue stress on the blade flap bending. It should be noted, however, that blade edge bending, which occurs in the rotor plane and is likely to be more sensitive to drive shaft twist angle oscillations, cannot be modeled in SymDyn.

The final plots of note for these two adaptations, Figure 4-13, show the tower side-to-side bending angle. It appears that the adaptation of the generator torque gain has no consistent effect on this mode, since in one case the oscillations continue to decrease in magnitude—a trend started before the adaptation—and in the other case they continue a trend of increasing in magnitude. Environmental variables and rotor speed are the major driving factors for tower bending, since the first tower natural frequency occurs at a rotor speed well within region 2




Figure 4-12: Blade 1 and 2 flap angle for two different adaptations

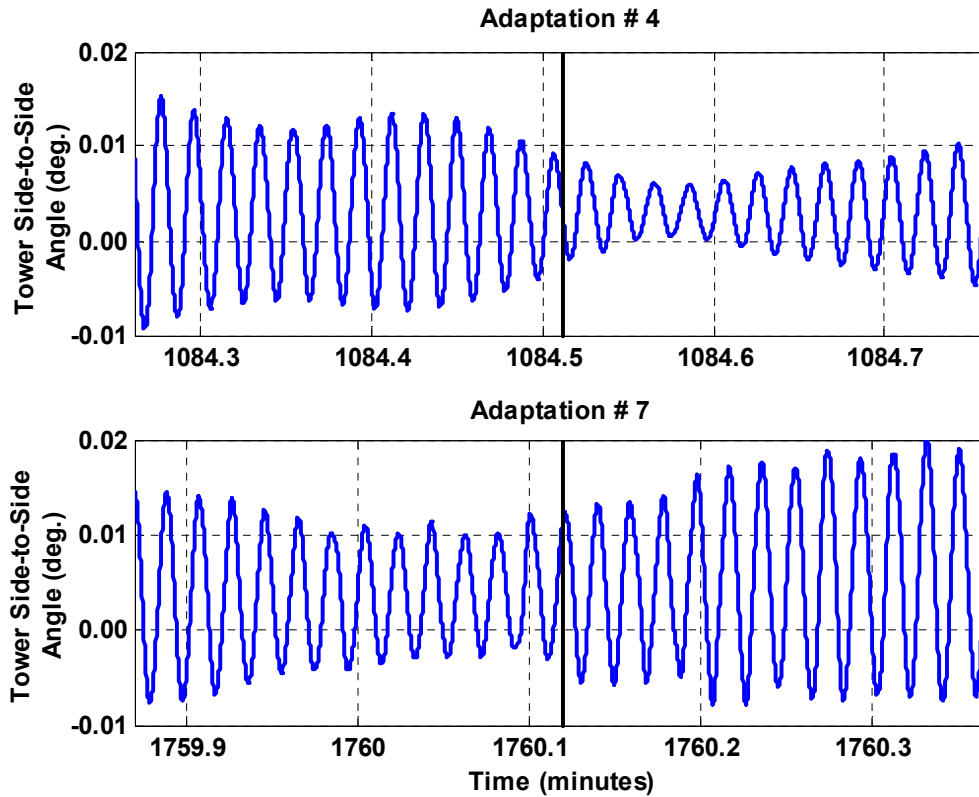


Figure 4-13: Tower side-to-side angle for two different adaptations

(about 2.7 rad/s). In fact, this rotor speed is reached during the last few seconds of the Adaptation #4 plot, which probably explains the increasing oscillations.

In summary, of the DOF modeled in the SymDyn simulations, only drive shaft torsion appears to be excited by the adaptation of the gain in the adaptive controller, and this excitation is likely to be much less pronounced on a real turbine. The other DOF modeled in simulation but not plotted (tower fore-aft angle and tower twist angle) show, like the tower side-to-side angle plots, no consistent response to the adaptations.

The gain adaptation law has been tested thoroughly in simulation, using both the SimInt and SymDyn modeling tools. The SimInt simulations have tested a number of areas of concern, including non-zero yaw error, inaccurate wind measurement, and changing parameters. Outputs of the SimInt and SymDyn simulations have been compared and found to be similar, and the SymDyn outputs were examined for any new stresses in the turbine components. The proposed gain adaptation law is now ready for testing on the real CART.

CHAPTER 5: Real CART Data

The simulation results presented in Chapter 4 provide promising evidence that the adaptive controller will increase a turbine's energy capture by adapting the torque control gain M toward its optimum value M^* . These are, however, only simple simulations of an extremely complex system. In order to validate the gain adaptation law, it is useful to test its operation on a real turbine, because issues that might not arise in a simulation must be resolved before industry would be able to fully trust a new controller. The CART is an invaluable tool for determining the benefit of this new adaptive controller to the turbine industry, and Chapter 5 analyzes the data collected while running the CART with the adaptive torque controller. Section 5.1 examines the adaptation history during these experiments. Next, Section 5.2 demonstrates the increase in energy capture as a result of the gain adaptation law. Finally, Sections 5.3 and 5.4 examine the turbine's dynamic response to the adaptation law (Equations 2.18 to 2.20) and the adaptive controller (Equation 2.14) compared with the standard, non-adaptive controller (Equation 2.1).

5.1 Adaptation History

Figure 5-1 shows data collected on the CART during approximately 10 months of operation under the adaptive controller (Equation 2.14). The adaptation period used to collect the data plotted in the upper plot of Figure 5-1 is 10 minutes for the first 11 adaptations, 20 minutes for the next nine, 30 minutes for the following 51, 60 minutes for the next 10, and 180 minutes ($n = 360,000$) for the final 13.

During the time covered by Figure 5-1, numerous changes were made to the adaptive controller as minor errors were discovered. In addition, several problems with sensors on the CART were discovered. For example, the drop in M between about nine and 15 hours is partly the result of an erroneous wind direction sensor. Each discontinuity in the upper plot of Figure 5-1 is the result of a restart in the gain adaptation law caused by either a change in the adaptation law or problems with faulty hardware. Finally, in the interest of clarity, the time scale in Figure 5-1 reflects only the time spent operating in region 2.

These experimental data show that even the one-hour adaptation period used between 33 and 43 hours is probably too short to result in convergence of M on a turbine as large as the CART. This result is consistent with the fact that a 13-hour adaptation period was used to obtain the degree of convergence evident in Figure 4-4. However, the results of the experiments using a three-hour adaptation period, enlarged in the lower plot of Figure 5-1, are somewhat more promising. Though the exact value of M^* for the CART is unknown (the $M^+ = 174.456$ used throughout this report is the value obtained via the PROP simulation plotted in Figure 2-1), the black line at $M = 0.47 M^+$ in Figure 5-1 is believed to be close to the CART's true optimal torque control gain M^* based on other (constant-speed) experimental data. This value and the data used to obtain it are discussed further by Fingersh and Johnson (2004). Note that once the adaptation period was lengthened to three hours, the adaptive gain was never more than one step away from




Figure 5-1: Adaptations of M on the CART

this approximate optimal value; instead, the adaptive gain oscillated around that value with relatively small steps. The fact that both the constant speed and adaptive control experiment data suggest a similar optimal torque gain M^* provides validity to each experiment's results.

Although some of the oscillations in Figure 5-1 are larger than others, particularly the one at about 76 hours, it is believed that a longer adaptation period reduces their magnitudes and provides better convergence. This belief has proven true in simulation. However, given the time constraints on the use of the CART, which is a test bed for several ongoing experiments at any given time, it was not possible to use a significantly longer adaptation period (such as 13 hours).

5.2 Energy Capture Using Adaptive Control

While data showing the adaptive gain M approaching a steady-state value is a promising reassurance that the gain adaptation law is working properly, it is actually secondary to the underlying problem: maximizing energy capture. In theory, the optimal M is the one that maximizes energy capture, so the question of whether M converges to its optimal value M^* and the question of whether energy capture is maximized are one and the same. Baseline power data from the CART running in variable speed mode with the standard controller (Equation 2.1) with gain $K = 0.91M^+$ (Fingersh and Johnson 2004) are compared to power data with the CART

running with the adaptive controller (Equation 2.14) where $0.40M^+ < M < 0.52M^+$ in Figure 5-2. The individual data points in Figure 5-2(a) are the averages of 10-minute data sets, while the curves in Figure 5-2(b) show the average values of the set of 10-minute averaged data contained in each 1 m/s wind speed bin. Superimposed on the curves in Figure 5-2(b) are the vertical and horizontal standard error bars obtained by dividing the standard deviation of the data within each 1 m/s bin by the square root of the number of points in each bin.

While there is insufficient data available to show an increase in energy capture for one particular adaptive gain M , it is clear from Figure 5-2 that controlling the CART with M within the range of gains $0.40M^+ < M < 0.52M^+$ does significantly increase energy capture for the CART, particularly in medium winds. In high winds, both schemes use the same region 3 controller, so the two curves are expected to converge. Table 5-1 summarizes the data plotted as solid lines in Figure 5-2. For mid-range wind speeds, it appears that the adaptive controller causes the CART to capture between about 5% and 14% more energy than the standard controller does.

Table 5-1. Mean Equivalent Wind Speed and Mean Grid Power Data for Standard and Adaptive Controllers ($0.40M^+ < M < 0.52M^+$)

Wind Speed Range (m/s)	Mean Equivalent Wind Speed (m/s)			Mean Grid Power (kW)		
	Standard Controller	Adaptive Controller	Percent Increase	Standard Controller	Adaptive Controller	Percent Increase
(4.5,5.5)	5.24	5.18	-0.98	20.42	16.56	-18.90
(5.5,6.5)	5.97	6.01	0.75	33.89	37.35	10.19
(6.5,7.5)	6.95	7.06	1.60	70.84	79.08	11.63
(7.5,8.5)	8.08	8.02	-0.79	109.99	122.45	11.33
(8.5,9.5)	9.00	9.04	0.45	161.43	181.11	12.19
(9.5,10.5)	9.90	9.92	0.19	216.97	247.92	14.27
(10.5,11.5)	10.98	10.87	-0.96	287.39	312.91	8.88
(11.5,12.5)	12.05	12.02	-0.27	388.43	407.80	4.99
(12.5,13.5)	12.98	13.13	1.10	446.12	469.78	5.30
(13.5,14.5)	13.89	13.94	0.38	488.93	514.05	5.14
(14.5,15.5)	15.18	14.94	-1.58	544.52	538.57	-1.09
(15.5,16.5)	15.93	16.16	1.41	558.80	564.02	0.93
(16.5,17.5)	17.30	16.97	-1.97	587.00	585.49	-0.26
(17.5,18.5)	17.87	18.20	1.81	586.65	583.02	-0.62

Because wind power is a function of the cube of the wind speed (Equation 2.5), the mean equivalent wind speed is used in Table 5-1 and Figure 5-2 rather than the mean wind speed. The mean equivalent wind speed, discussed further in Appendix A, is a measure of the constant wind speed that would provide the mean wind power calculated over the 10-minute averaging period. Since the mean equivalent wind speed in Table 5-1 for each controller was within about 2% over the wind speed range of 4.5 m/s to 18.5 m/s, it is reasonable to state that the wind power available to the standard and adaptive controllers was similar. For mean equivalent wind speeds ranging from 5.5 to 14.5 m/s, there are at least 10 data sets within each 1 m/s range.

Great care must be taken not to misinterpret the data provided in Table 5-1. The only way to be certain about the improvement in energy capture caused by the adaptive controller is to run each

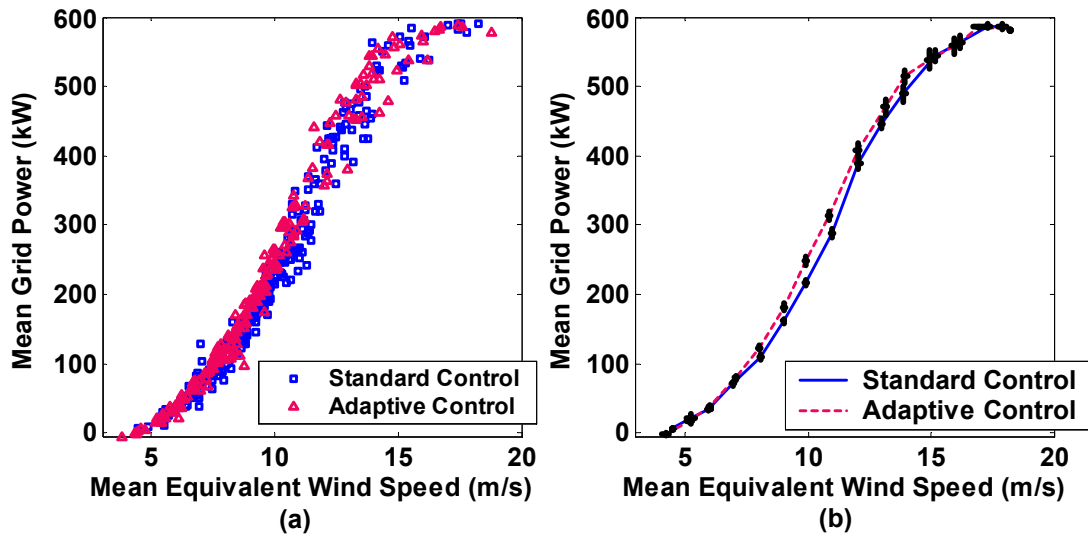


Figure 5-2: Grid power vs. equivalent wind speed for standard and adaptive control: (a) 10-minute averages; (b) binned data with error bars

controller on the turbine with identical wind input, obviously an impossibility on a real turbine outside of a wind tunnel. However, comparisons among turbines and controllers are frequently made by multiplying a turbine's power curve by the probability of a certain wind speed occurring and thus calculating the expected energy. This type of analysis was conducted for the CART's standard and adaptive power curves ($0.40M^+ < M < 0.52M^+$) and is shown in Figure 5-3. The power curves were specified as being zero below 5 m/s (i.e., the "cut-in" wind speed) and rated power (600 kW) above 18 m/s ("rated" wind speed), with a polynomial best fit line applied to the data between 5 and 18 m/s.

The two solid lines in Figure 5-3 show the expected hourly energy at each wind speed, assuming a Rayleigh wind distribution with a mean of 8 m/s. For example, this Rayleigh distribution assumes that the wind speed will be between 10 m/s and 10.25 m/s 1.75% of the time, or 1.05 minutes out of every hour. In order to form the solid curves showing expected incremental energy, this probability was multiplied by the value of the standard and adaptive control power curves at this wind speed, with the result being the amount of energy (kWh) expected by the turbine due to wind speeds of 10 m/s to 10.25 m/s during any given hour. The expected total annual energy, marked by the dotted curves, is simply the integral of the expected incremental energy and multiplied by the number of hours in a year.

Table 5-2 shows the percent increase in annual energy obtained by using the adaptive controller rather than the standard for six different wind speeds. Since 25 m/s is the cut-out wind speed, meaning that no additional energy is captured for wind speeds above 25 m/s, the conclusion is that the CART would capture 5.5% more energy under the adaptive control than under the standard control while operating at a site with a Raleigh wind distribution.

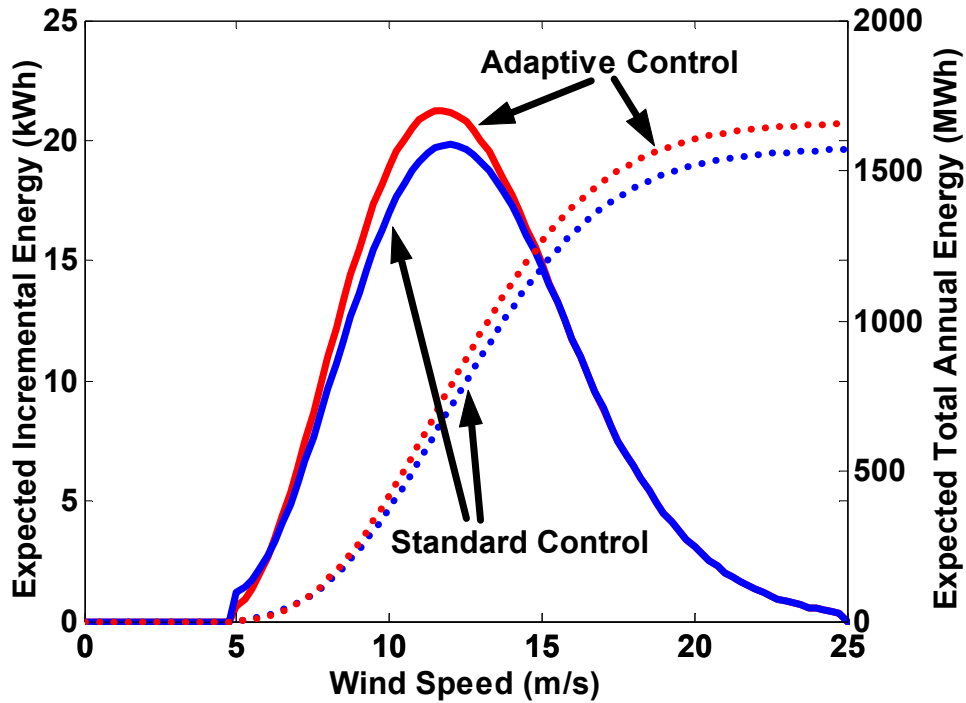


Figure 5-3: Expected incremental and total energy for standard and adaptive control on the CART

Table 5-2. Increase in Annual Energy Capture over Standard Controller for Adaptive Controller with $0.40M^+ < M < 0.52M^+$

Wind Speed (m/s)	Annual Energy: Standard Control (MWh)	Annual Energy: Adaptive Control (MWh)	Percent Increase
12	710.39	780.05	9.8
14	1039.5	1123.7	8.1
16	1289.9	1376.7	6.7
18	1440.6	1526.9	6.0
20	1517.5	1603.8	5.7
25	1569.3	1655.5	5.5

Since maximizing energy capture was the primary goal of this adaptive controller (Equation 2.14), it seems that the controller has achieved its goal. Unfortunately, it is impossible to tell exactly when a real turbine’s energy capture is maximized, as opposed to merely increased. It is possible, however, to plot the computed fractional average power P_{avg} (Equation 2.15) vs. the torque control gain M , fit a polynomial curve, and see where the peak of the curve occurs. This is done in Figure 5-4, which includes only the 180-minute adaptation period data (i.e., the data plotted in the lower plot in Figure 5-1).

Unfortunately, because of time constraints on using the CART for experimental testing and measurements, only these 13 data points were collected using the three-hour adaptation period,

and Figure 5-4 shows that there is still significant noise within the measurements. This noise is mostly the result of the short (180-minute) adaptation period, which makes it difficult to obtain a high correlation between measured wind and wind hitting the CART's rotor. The peak of the best-fit curve occurs at $M = 0.49M^+$, which is close to the value $M = 0.47M^+$ computed from the constant speed CART data and plotted as the straight line in the lower plot of Figure 5-1.

The data collected while running the CART with the adaptive controller (Equation 2.14) has shown an increase in energy capture. All else being equal, an increase in energy capture by a given turbine will lead to a decrease in cost of energy. However, any new control technique that requires a turbine to be structurally modified in a way that increases the cost of the turbine or shortens its life span may not reduce the cost of energy after all. The next section examines the measured stresses on the CART, particularly those immediately before and after an adaptation of the torque control gain.

5.3 Measured Stresses on Turbine Components

The CART is outfitted with strain gauges, accelerometers, and other instruments that can be used to determine whether it is experiencing high loads. In general, region 2 loads are smaller than region 3 loads because of the slower wind speed and turbine speed in region 2. Thus, design loads, or loads that drive the turbine's structural design, rarely occur completely within region 2 (though some do occur in the transition between region 2 and region 3). It is important to understand whether the adaptive controller (Equation 2.14) adds any stress to the turbine beyond that normally be expected for region 2 operation under the standard non-adaptive controller (Equation 2.1). These additional loads could be of two types: transient or steady-state. The




Figure 5-4: Fractional average power P_{fav} vs. normalized M

transient loads are those that occur over a short time period directly following and as a result of an adaptation of the adaptive gain M . The long-term steady-state loads are those that would result from a change in the turbine's operating point. For example, if a lower adaptive torque gain M causes the rotor to spin more quickly and thus undergo more revolutions in a given time, the damage equivalent fatigue loads for blade edge bending might be expected to be higher, since the blades would experience more gravity loading cycles.

Of these two types of loads, only the transient loads can be fairly attributed to the adaptive nature of the proposed adaptive controller (Equation 2.14). While it is true that the adaptive controller could lead to steady-state operation that causes an increase in damage equivalent fatigue loads compared with those experienced with the standard non-adaptive controller (Equation 2.1), this increase would be due only to the fact that the standard controller failed to use the optimal gain $K = M^*$. This section will assume that the turbine was designed to withstand damage equivalent fatigue loads consistent with the use of the standard controller with gain $K = M^*$, although, in practice, this design might not be possible since M^* is likely unknown. Thus, the loads of interest in this research are the transient loads that result directly from the adaptations of the adaptive controller. This analysis is very similar to that conducted on the SymDyn simulation data in Section 4.6 except that here real CART data are the subject of the examination. Once again, data collected 30 seconds before and after each of two adaptations are plotted, and a brief summary of the two selected adaptations appears in Table 5-3.

Table 5-3. Selected Details Relevant to Figure 5-5 through Figure 5-8

Adaptation	Normalized $M (M/M^*)$	Rotor Speed	Rotor Acceleration
#6	0.47 to 0.41	3.6 rad/s	< 0
#12	0.38 to 0.46	4.1 rad/s	< 0

Figure 4-11 shows that, in simulation, the effect of an adaptation on the drive shaft twist angle is noticeable and possibly detrimental to the drive shaft. Fortunately, the same effect is not seen in the real CART data, plotted in Figure 5-5 for two different adaptations. The independent axes in Figure 5-5 show time in seconds after the respective data sets began rather than from some universal zero time.

Based on the data plotted in Figure 5-5 and a review of other data immediately before and after other adaptations, it appears that the adaptations of the generator control torque gain do not have a detrimental effect on the drive shaft twist angle. This result is likely due to two facts. First, on the real CART, it is impossible to change the torque control gain instantaneously, and second, the torque control signal sent to the turbine is actually filtered with a one second time constant for the purpose of preventing undesirable transient loads. Thus, even though the gain adaptation law causes the value of M to change instantaneously, it takes several seconds for the power electronics to catch up to the change. This slow response can be credited for eliminating the ringing effect that was apparent in the simulation outputs (Figure 4-11).

The CART also has blade and tower strain gauges that may reveal any transient effects of the adaptation on those structures. Blade flap and edge bending moment are plotted in Figure 5-6 and Figure 5-7, respectively, and tower side-to-side and fore-aft bending moment follow in Figure 5-8 and Figure 5-9.

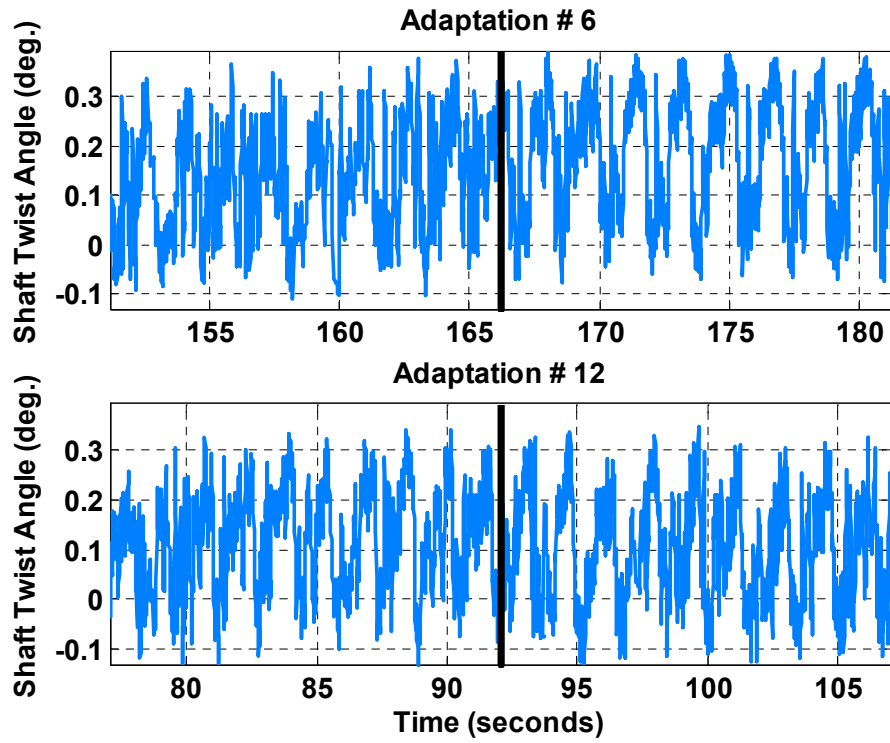


Figure 5-5: Drive shaft twist angle for two different adaptations (real CART data)

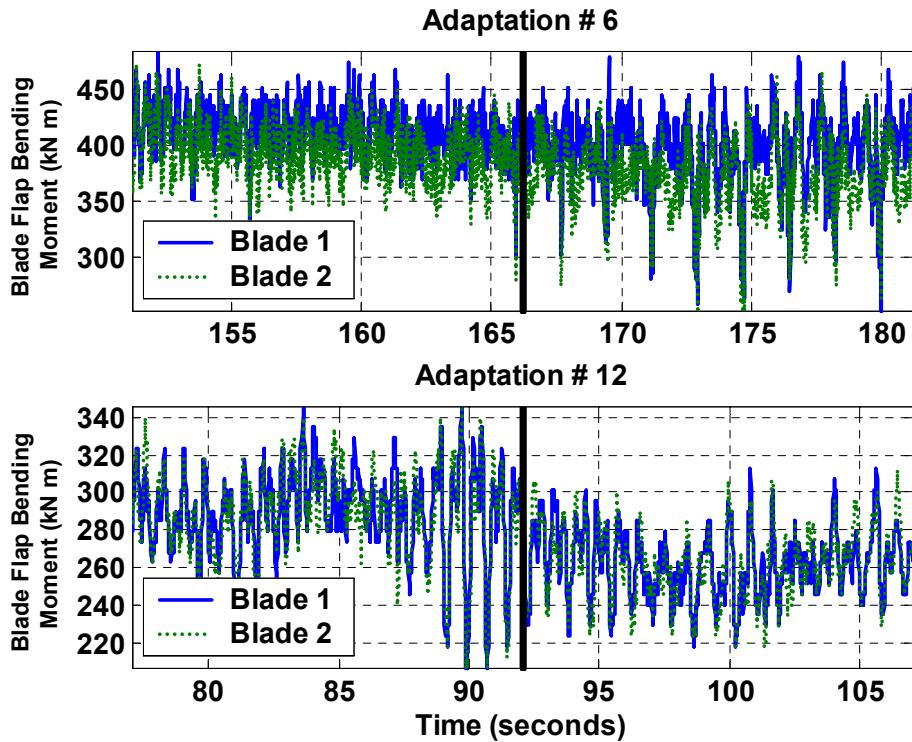


Figure 5-6: Blade flap bending moment (real CART data)

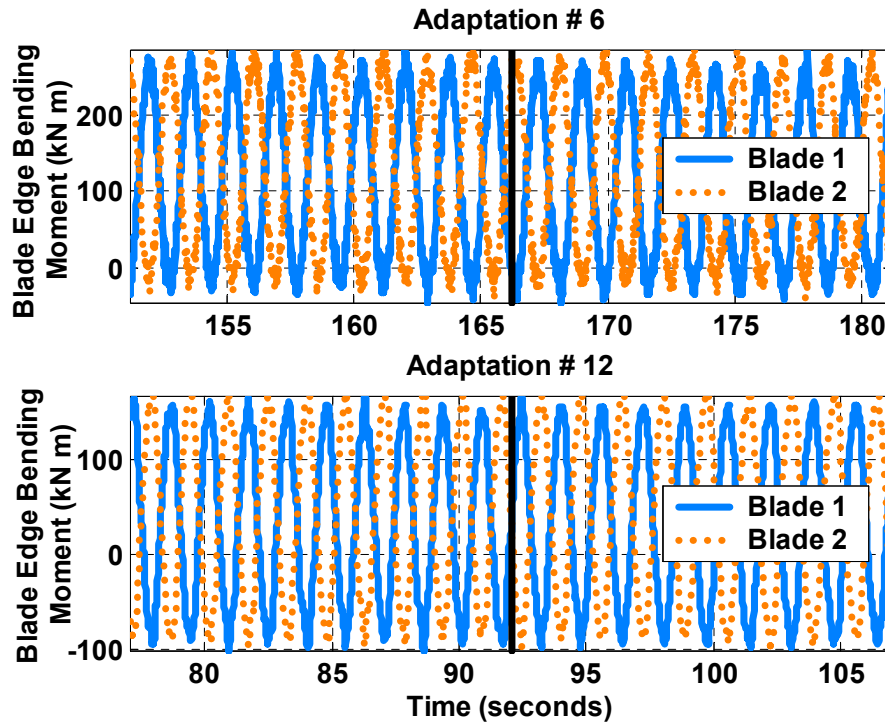


Figure 5-7: Blade edge bending moment (real CART data)

It is difficult to tell from Figure 5-6 alone whether the adaptation of the torque control gain has a transient effect on blade flap bending. Following Adaptation #6, it appears that there is an increase in the standard deviation of the bending moment of each blade; however, the deviation is reduced following Adaptation #12. Based on the inconsistency of these results and the fact that the wind speed increases several seconds after Adaptation #6 and decreases following Adaptation #12, it is most likely that the adaptation had little or no effect on the blade flap bending. Rather, the blades were responding to the increase and decrease in wind speed. In Figure 5-7, it is clear that the adaptation has no transient effect on the blade edge bending moments, which are driven mainly by gravity and the acceleration and deceleration of the rotor.

Data in Figure 5-8 and Figure 5-9 also suggest transient loads resulting from environmental effects rather than adaptation effects. In each case, there is a slight increase in tower bending moment following Adaptation #6 and a slight decrease following Adaptation #12. Since wind speed is known to affect tower bending, and since there are no consistent tower bending responses among all of the adaptations, it can be safely deduced that the adaptations of the torque gain will not have a detrimental effect on the tower bending moment.




Figure 5-8: Tower side-to-side bending moment (real CART data)




Figure 5-9: Tower fore-aft bending moment (real CART data)

The time series data plotted in Figure 5-6 through Figure 5-9 provide a partial picture of the transient effects of the adaptations on the CART. However, more than 75 adaptations have occurred under the gain adaptation law (2.14) and it is simply impossible to plot time-series data for all of them. Instead, the following figures plot some statistical data for the time periods immediately preceding and following an adaptation.

Figure 5-10 plots the ratio of the standard deviation of the accelerometer data in three Cartesian coordinates in the 30 seconds before and after 18 of the adaptations. These 18 adaptations were chosen from the 60 and 180 minute adaptation data, excluding those adaptations that occurred within the first or last 30 seconds of a data set. Figure 5-10(a) uses the adaptation number as the independent axis, while Figure 5-10(b) uses the ratio of the value of M after the adaptation to the value of M before the adaptation. From Figure 5-10(a) it is clear that, in general, the adaptation of the torque control gain does not have a significant short-term effect on the variation of the turbine's accelerometer outputs (as measured by the standard deviation), since most of the ratios cluster around 1.0 (i.e., similar variation before and after the adaptation). However, a trend becomes clear in Figure 5-10(b). In this case, it becomes clear that when the adaptive gain M is decreased ($M_{k+1}/M_k < 1.0$), the variation in the accelerometer outputs goes up, whereas when M is increased, the variation goes down. This result makes sense, since (all else being equal) a smaller M results in a smaller control torque (Equation 2.14) and thus a faster rotor speed (Equation 2.7), which tends to cause the turbine to shake more.

Indeed, Figure 5-11 shows the relationship between rotor speed and the standard deviation of accelerometer data. Clearly, when the rotor speed increases, as it does when M decreases, the accelerometers record larger variations in acceleration than prior to the increase in rotor speed. Thus, if M adapts downward from the initial guess M_0 , it is expected that the turbine will shake more than it would have had the standard non-adaptive controller been used throughout the turbine's life. However, this steady-state result is not the fault of the adaptive controller but rather of the initially inaccurate guess M_0 . If the primary goal of region 2 operation were to minimize turbine shaking rather than to maximize energy capture, the problem would be easily

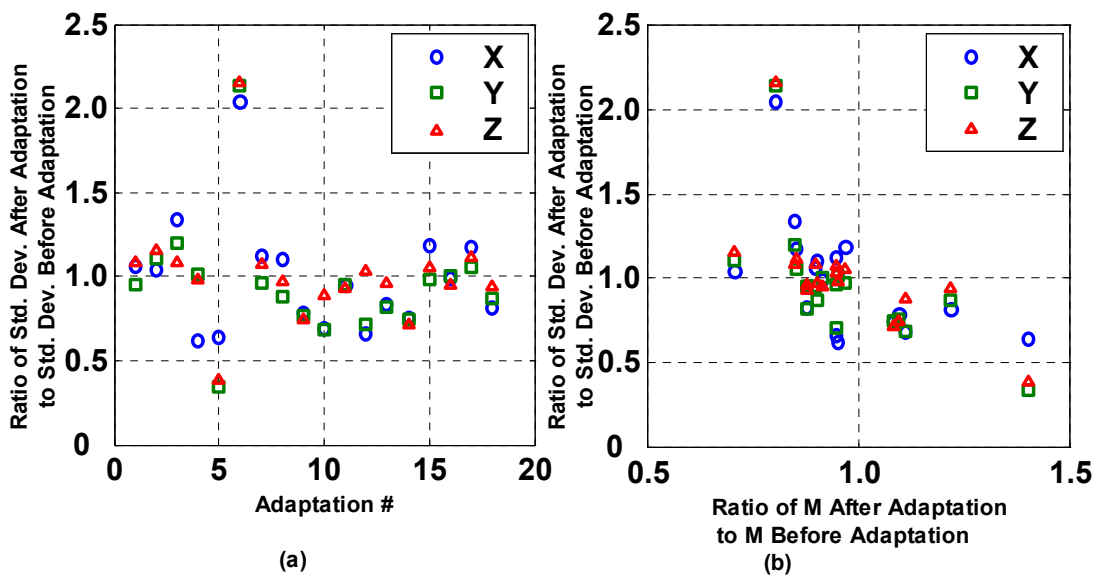


Figure 5-10: X, Y, and Z accelerometer data




Figure 5-11: X, Y, and Z accelerometer data versus mean rotor speed data

solved by stopping the turbine in this region.

Examination of other CART data—including low speed shaft and high speed shaft torque transducer output, nacelle roll, pitch, and yaw accelerations, and electrical outputs—have found no cause for worry that the adaptive controller’s abrupt adaptations will harm the turbine. In other words, the turbine’s loads under the adaptive controller (Equation 2.14) appear to be so similar to those under the standard non-adaptive controller (Equation 2.1) operating with K equal to the true optimal torque control gain M^* that there should be no need to adjust the turbine design in any way in order to maintain a similar life span. Thus, any increase in energy derived from the adaptive controller can be considered to be a direct decrease in the cost of wind energy, assuming the turbine was designed to withstand steady-state fatigue loads under the standard controller with $K = M^*$.

Now that it has been determined that the adaptive controller does not appear to increase the stresses on the CART beyond those experienced under the non-adaptive control in region 2, there is one more area of operation to examine. This area is the transition between region 2 and region 3 control, an area that has been known to cause significant loading on commercial turbines.

5.4 Region 2–3 Transition

The distinct region 2 and region 3 control objectives lead to many turbines using separate control strategies in region 2 and region 3. On many turbines, the strategy is as simple as switching from one controller to the other (e.g., constant-pitch, generator torque control in region 2 to constant-torque, variable pitch control in region 3). On the CART, both the pitch and generator torque control are technically “active” at all times, though they are designed so that the pitch reaches its constant saturated value in region 2, and the generator torque reaches its constant saturated value in region 3. The transition between controllers can cause significant loading on a turbine’s mechanical and electrical components, with overspeed and overpower transients commonly seen in the industry (Eggers et al. 2002, Leith and Leithead 1997). One example of a bad transition on the CART is plotted in Figure 5-12; an example of a good transition follows in Figure 5-13. For ease of comparison, both are plotted on the same scales. In the bad transition example of Figure 5-12, the torque control gain was constant throughout the transition at $0.72M^+$, while in the good transition example of Figure 5-13 this gain was almost the same at $0.71M^+$. Clearly, the torque control gain M is not the only factor affecting the smoothness of the transition.




Figure 5-12: Bad region 2 - region 3 transition on the CART

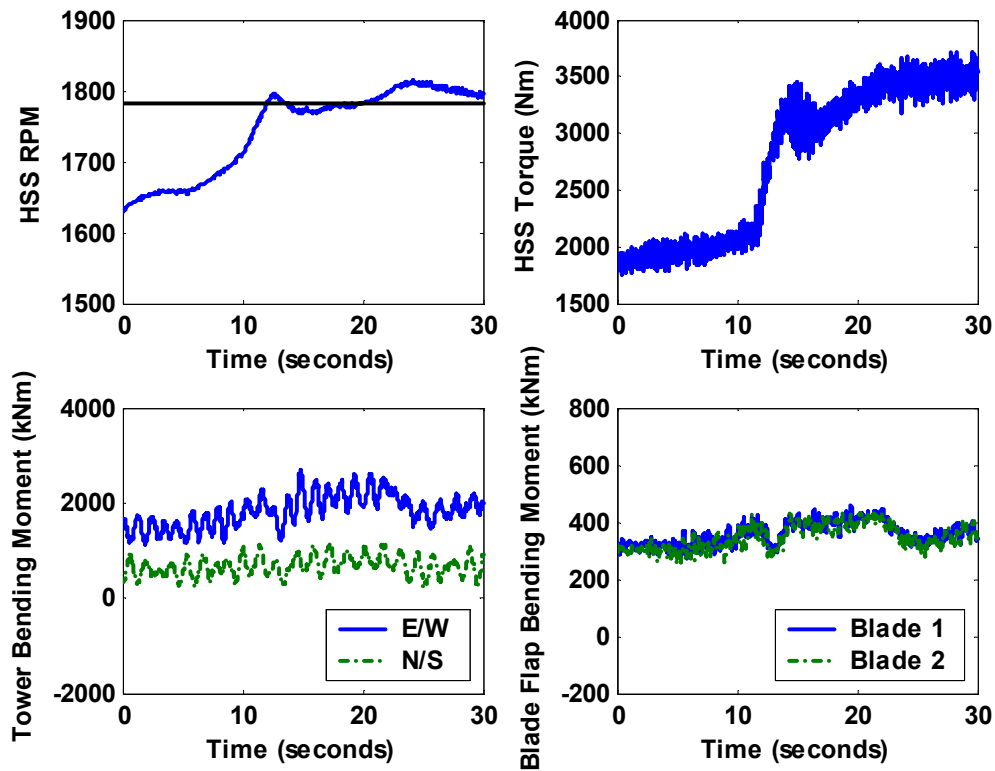


Figure 5-13: Good region 2 - region 3 transition on the CART

In Figure 5-12, there are large oscillations in many of the CART data channels, including the high speed shaft angular speed (rpm), the high speed shaft torque, the tower bending moments in two directions, and the two blade flap bending moments. The horizontal line plotted in the HSS RPM subplot marks the transition between regions; region 3 operation occurs when the speed is above this line. During the time period plotted, the turbine was switching between controllers as the speed moved above and below the transition speed. These data are from one of the more extreme transitions the CART experienced during the variable speed adaptive controller testing. Fortunately, most transitions occur with much lighter loading on the turbine components, like the example plotted in Figure 5-13.

Since the torque control gain M was almost identical between the two examples and the region 2 control scheme was the same, the difference between the transitions is likely the result of atmospheric factors. Indeed, the wind speed was very different between the two examples, as plotted in Figure 5-14. In the first 10 seconds of the bad transition example, the wind speed rose from 12 m/s to 19 m/s, then dropped back down to about 15 m/s. Such rapid fluctuations can cause high loading on the turbine. By contrast, the wind speed measured during the good transition ranged from 11 m/s to 15 m/s over the 30-second period plotted.

Note that the problem of the region 2 to region 3 transition on a variable speed turbine is not a standard control problem wherein a variable must move from one steady-state value to another




Figure 5-14: Wind Speed for Bad and Good Transition Examples

through a transient region. When the wind speed is fluctuating right around the region 2–region 3 transition region, it is desirable for the turbine to switch back and forth between control regions rather than for it to reach region 3 control and remain there. Thus, it would be desirable to have a means of alleviating the loading problems that are caused by the controller.

There are many possible solutions to the transition loading problem, and adaptive torque gain control is one option. However, in order to derive the most benefit from an adaptive torque controller in the transition region, the controller must be able to adapt on a time scale similar to or shorter than the time usually spent in transition (seconds). The adaptive torque gain controller proposed in this research, then, will not actively improve the transition. However, it is not well understood whether the identification of the optimal torque control gain performed by the adaptive controller will have an unintentional effect on the transition region. This is the question addressed in this section.

If the use of the optimal torque control gain M in the region 2 control law were to have an effect on the transition loading, then there should be a correlation between the magnitude of the oscillations of some of the turbine components and the value of M . One way to measure the magnitude of the oscillations is with the standard deviation. In this study, 39 30-second transition periods were identified within the CART data that contained at least two seconds in region 2 within the first 20% (6 seconds) of the period and 10 seconds in region 3 within the last 50% (15 seconds) of the period. During the intermediate 30% (9 seconds), the turbine could be operating within any control region. These specifications were chosen to ensure that the transitions being compared were similar in nature so that the statistical analysis would be valid. The values of the torque control gain varied from $0.31M^+$ to $1.00M^+$ among the 39 periods of the study.

It is clear from Figure 5-15 that there is no statistically significant correlation between the value of the torque control gain M and the standard deviations of the blade flap and edge bending moments, tower bending moments, and accelerometer outputs in any direction. Since these

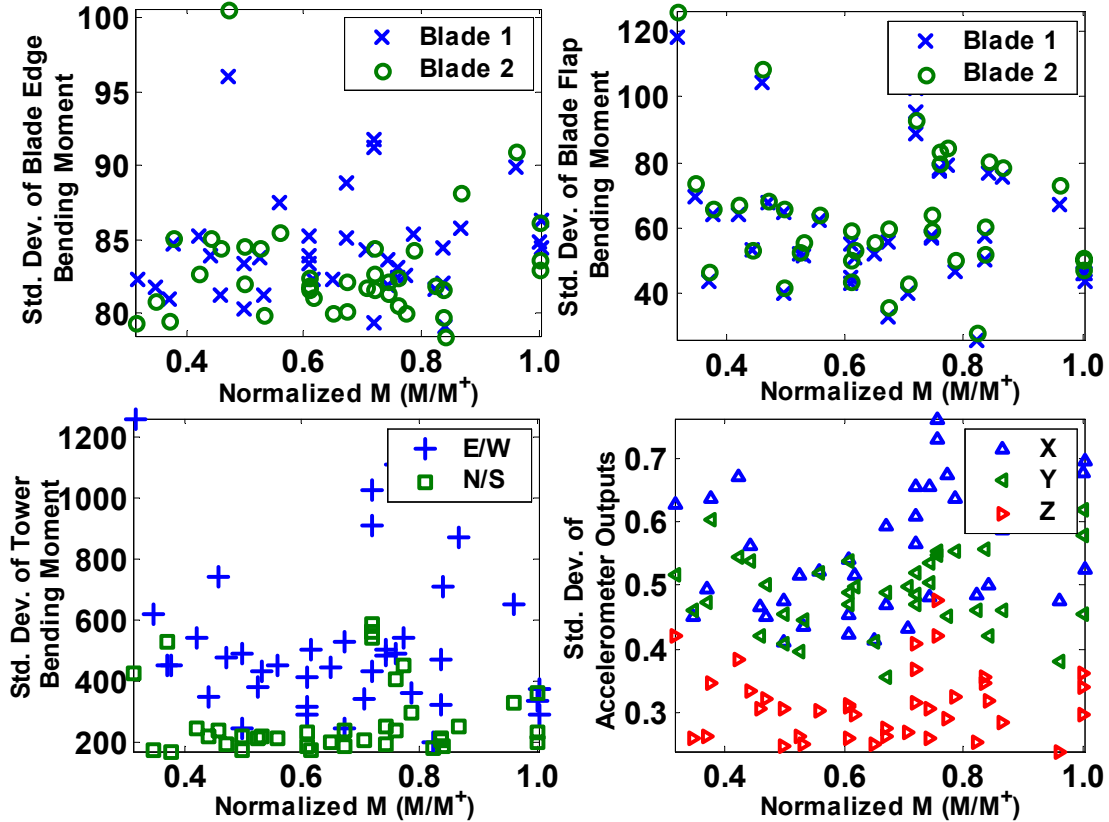


Figure 5-15: Standard deviations of turbine loads vs. normalized torque gain M during region 2 to region 3 transitions

measurements are likely candidates for loading problems to appear, it seems unlikely that this adaptive controller (Equation 2.14) will have a significant effect on turbine loading during the transition between regions.

Chapter 5 has examined the CART data obtained while running the adaptive controller (Equation 2.14) and found some promising results. The first of these results is that the gain adaptation law (Equations 2.18–2.20) causes the adaptive gain M to converge toward a steady-state value which is near its optimal value M^* as computed from other data. The second result is that the CART's energy capture is increased when the turbine runs with an adaptive gain M in a narrow range around the estimated M^* . Third, the actual adaptation of the adaptive gain M does not appear to have a detrimental effect on the turbine. All of these results are very promising for adaptive torque control of variable speed turbines. Then, in a result that is neither positive nor negative, it was determined that the gain adaptation law has little or no effect on the transition between region 2 and region 3, a well-known trouble area for variable speed turbine control.

CHAPTER 6: Theoretical Analysis

There are several stability questions to be addressed that can help to ensure safe and desired operation of a variable speed adaptive torque gain control law and also to provide further insight into the systems. The first three questions relate to the stability properties of the turbine's torque control law—a continuous-time problem on a rapid time scale. First, the simple problem of the asymptotic stability of the equilibrium point $\omega = 0$ in the absence of wind is addressed; the conclusion is that the standard region 2 controller causes the equilibrium point $\omega = 0$ to be asymptotically stable. A similar conclusion is reached in the case of a constant wind input. Next, the finite gain stability properties of the system from input (wind) to output (rotor speed) are examined. All three of these stability results are proved under the assumption that the adaptive control gain $M > 0$ is constant; this is a valid assumption because the gain adaptation takes place discretely and on a time scale several orders of magnitude slower than that of the wind and rotor speed changes (many hours vs. seconds). Simplified block diagrams for the systems are shown in Figure 6-1. In Figure 6-1(a), the linear plant is simply as stated in Equation 2.7, and the nonlinear controller is as stated in Equation 2.14.

The final stability question regards the convergence of the adaptive gain $M \rightarrow M^*$ given the proposed gain adaptation law. Figure 6-1(b) shows the simplified block diagram for this system, where the nonlinear plant is the turbine's P_{favg} vs. \tilde{M} relationship and the nonlinear controller is given by Equations 2.18–2.20. One example of a P_{favg} vs. \tilde{M} curve, shown in Figure 6-2, is based on SimInt simulations using the CART's PROP-generated C_p surface. The assumption in this section is that the P_{favg} vs. \tilde{M} relationship on most modern turbines resembles this curve. This is a valid assumption because the shape of this curve is qualitatively based on the shape of the turbine's C_p surface, and the shape of a C_p surface is qualitatively the same from turbine to turbine. In all of these proofs the air density, ρ , is assumed to be a constant greater than zero. In reality, changes in air density are small in magnitude.




Figure 6-1: Simplified block diagrams (a) relating aerodynamic torque and rotor speed, and (b) gain adaptation law.




Figure 6-2: Fractional average power P_{favg} vs. error in M

6.1 Asymptotic Stability of $\omega = 0$

Because wind turbines are designed to spin as freely as possible, the friction due to mechanical bearings, gear mesh, generator core losses, and air resistance during operation is very small. However, in the formal proof of the asymptotic stability of the equilibrium point $\omega = 0$, the equation for $\dot{\omega}$ (Equation 2.7) should be amended to include a frictional term b , where $b > 0$:

$$\dot{\omega} = \frac{1}{J}(\tau_{aero} - \tau_c - b\omega). \quad (6.1)$$

Now, given the equation for aerodynamic torque (Equation 2.8) and Equation 2.14, Equation 6.1 can be expanded to

$$\dot{\omega} = \begin{cases} \frac{1}{2J} \rho A R C_q v^2 - \frac{b}{J} \omega, & \omega < 0 \\ \frac{1}{2J} \rho A R C_q v^2 - \frac{\rho}{J} M \omega^2 - \frac{b}{J} \omega, & \omega \geq 0 \end{cases} \quad (6.2)$$

Theorem 1: The plant (Equation 6.1) and the nonlinear controller (Equation 2.14) have an asymptotically stable equilibrium point at $\omega = 0$ when $v = 0$ and M is a constant greater than zero.

Proof: When $v = 0$, the first term in Equation 6.2 becomes zero. In this case, the simple Lyapunov function candidate $V = \frac{1}{2}\omega^2$ has the derivative

$$\dot{V} = \begin{cases} -\frac{b}{J}\omega^2, & \omega < 0 \\ -\frac{\rho}{J}M\omega^3 - \frac{b}{J}\omega^2, & \omega \geq 0 \end{cases}$$

which is negative for all $\omega \neq 0$ and equals zero for $\omega = 0$. Thus, the equilibrium point $\omega = 0$ is globally asymptotically stable. \square

Note also that, when b is assumed to be zero, the equilibrium point is still stable in a global sense but no longer asymptotically stable for $\omega < 0$. The asymptotic stability still holds in a local sense for all $\omega \geq 0$.

6.2 Asymptotic Stability of Rotor Speed with Constant, Positive Wind Input

The next stability result to be addressed concerns whether or not the rotor speed ω converges to an equilibrium value under a constant, positive wind speed input. While it is unreasonable to assume a constant wind speed in the field, it is still desirable to understand the system response under these controlled conditions. Once again, the plant is given by (Equation 6.1) and the nonlinear controller is given by (Equation 2.14). This analysis is similar to the example provided in Section 2.1 explaining the operation of the standard, non-adaptive region 2 controller. However, unlike the standard controller (Equation 2.1), the adaptive controller (Equation 2.14) does not assume perfect knowledge of the aerodynamic parameters $C_{p_{max}}$ and λ^* . Thus, the cubic curve $F\lambda^3$ given in (Equation 2.13) and plotted in Figure 2-2 must be rewritten as follows:

$$C_p = \frac{\rho M \lambda^3 v + \lambda^2 b R}{\frac{1}{2} \rho A R^3 v} = G(\lambda, M, b, v). \quad (6.3)$$

In (Equation 6.3), b is several orders of magnitude smaller than M^* , so if it is assumed that M is within a neighborhood (say, an order of magnitude) of its true optimal value, the second term in the numerator of (Equation 6.3) is nearly zero in comparison. Although this fact is not required for the following analysis, it simplifies the re-drawing of Figure 2-2 because the wind speed v and air density ρ both cancel out and Figure 2-2 can be re-drawn in Figure 6-3 for various M values independently of v and ρ . When Figure 6-3 is plotted using representative values of ρ , v , and b , its qualitative nature does not change, and the new curves are indistinguishable from those plotted on the scale in use.

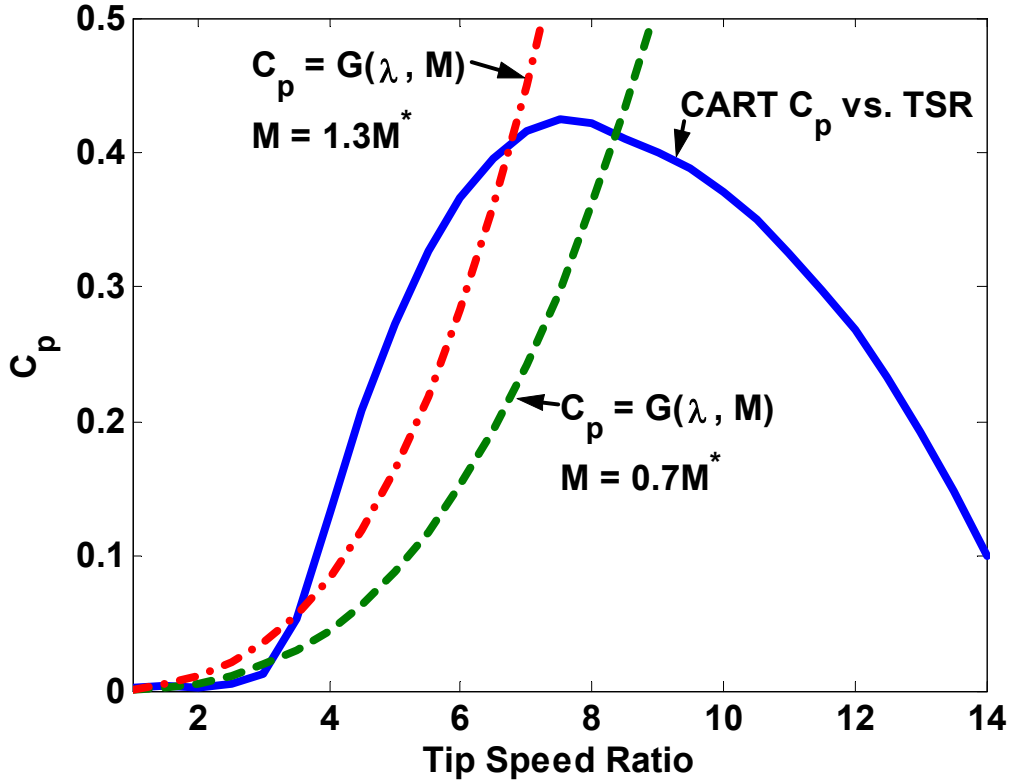


Figure 6-3: CART C_p and cubic for different M

The first item of note apparent in Figure 6-3 compared to Figure 2-2 is that the cubic function no longer intersects the CART's C_p curve at its peak when $M \neq M^*$. This is because optimal power capture cannot be achieved for $M \neq M^*$. Let λ_2 be defined as the intersection of $G(\lambda, M)$ with the CART's C_p vs. λ curve, $C_p(\lambda)$, such that $G(\lambda, M) > C_p(\lambda)$ for all $\lambda > \lambda_2$, i.e., the highest value of λ for which the two curves intersect. Let λ_1 be defined as the next highest intersection point, i.e., the λ for which $0 < \lambda_1 < \lambda_2$ and $G(\lambda, M) < C_p(\lambda)$ for all $\lambda_1 < \lambda < \lambda_2$ and $G(\lambda, M) > C_p(\lambda)$ for some $\lambda < \lambda_1$ within a neighborhood of λ_1 . For the dashed $M = 0.7M^*$ curve, these values correspond to $\lambda_1 = 3.1$ and $\lambda_2 = 8.4$. The following theorem states that, for a constant wind input, the tip speed ratio will converge to the intersection of the two curves, λ_2 , as long as the initial tip speed ratio is greater than λ_1 . Assume $\lambda_1 > 0$.

Theorem 2: The plant (Equation 6.1) and the nonlinear controller (Equation 2.14) have a locally asymptotically stable equilibrium point at $\lambda = \lambda_2$ when v and M are constants greater than zero. The domain of attraction is $\lambda_1 < \lambda < \infty$.

Proof: In the domain $0 < \lambda_1 < \lambda < \infty$, $\omega > 0$ holds (since $\omega = \lambda v/R$). Define $\tilde{\lambda} = \lambda_2 - \lambda$. Now choose the Lyapunov function $V = \frac{1}{2} \tilde{\lambda}^2$. For $\omega > 0$,

$$\dot{V} = (\lambda - \lambda_2) \left(\frac{1}{2J} \rho A R^2 C_q v - \frac{1}{JR} \rho M \lambda^2 v - \frac{b}{J} \lambda \right). \quad (6.4)$$

Substitution of C_p / λ for C_q in (Equation 6.4) and a little algebra provides the result that $\left(\frac{1}{2J} \rho A R^2 C_q v - \frac{1}{JR} \rho M \lambda^2 v - \frac{b}{J} \lambda \right) > 0$ when $\lambda_1 < \lambda < \lambda_2$ (i.e., when $G(\lambda, M) < C_p(\lambda)$) and $\left(\frac{1}{2J} \rho A R^2 C_q v - \frac{1}{JR} \rho M \lambda^2 v - \frac{b}{J} \lambda \right) < 0$ when $\lambda > \lambda_2$ causes $G(\lambda, M) > C_p(\lambda)$. Thus, $\dot{V} < 0$ for all $\lambda_1 < \lambda < \infty$ except $\lambda = \lambda_2$, for which $\dot{V} = 0$. This provides the result that the equilibrium point $\lambda = \lambda_2$ is locally asymptotically stable in the domain $\lambda_1 < \lambda < \infty$. \square

Note that this proof of the convergence of λ to a specific value is equivalent to the convergence of ω to a specific value for a specific wind speed because $\omega = \lambda v / R$. Also, note that when $M = M^*$, the curves $G(\lambda, M) > C_p(\lambda)$ intersect at $(\lambda^*, C_{p_{max}})$ and therefore optimal power capture is achieved for the constant wind input case.

While it is useful to understand the characteristics of the system when the wind input is constant, this unrealistic situation would only occur in a wind tunnel. Thus, the behavior of the system when the wind input is not constant is the next item studied.

6.3 Input-Output Stability

The next question to be answered is that of input-output stability. In reality, all wind turbines have a maximum safe operating speed, and some type of aerodynamic braking is usually used to prevent the turbine from operating at speeds that are above this maximum. These aerodynamic controls can include actuators for changing the blade pitch, special tip devices called tip brakes, blade flaps or ailerons, or some other method. However, it is still useful to examine whether the generator torque control would bound the turbine speed in some sense in the absence of these other controllers.

Theorem 2: If $C_q \leq 1$, the plant (Equation 6.1) and the nonlinear controller (Equation 2.14) is finite-gain \mathcal{L}_2 stable, where squared wind speed v^2 is the input and rotor speed ω is the output.

Proof: Passivity theory is used to answer this question. Consider the kinetic energy of the rotor: $\frac{1}{2} J \omega^2$. Now choose a positive definite function similar to the kinetic energy function:

$$V = \frac{1}{\rho A R} J \omega^2, \quad (6.5)$$

where $\rho > 0$ is once again considered to be a constant. Now, in order for the system to be output strictly passive, the time derivative of Equation 6.5 must be less than or equal to the product of the input and output minus some positive constant times the square of the output (Khalil 2002). In this particular case, with v^2 the input and ω the output, the requirement is

$$\dot{V} \leq v^2 \omega - \delta \omega^2 \quad (6.6)$$

for some $\delta > 0$. In fact, the time derivative of Equation 6.5 is

$$\dot{V} = \begin{cases} C_q v^2 \omega - \frac{b}{\frac{1}{2}\rho AR} \omega^2, & \omega < 0 \\ C_q v^2 \omega - \frac{b}{\frac{1}{2}\rho AR} \omega^2 - \frac{M}{\frac{1}{2}AR} \omega^3, & \omega \geq 0 \end{cases} \quad (6.7)$$

Let $\delta = \frac{b}{\frac{1}{2}\rho AR}$. Then, since $\frac{M}{\frac{1}{2}AR} \omega^3 \geq 0$ for $\omega \geq 0$, the inequality in Equation 6.6 is satisfied whenever $C_q \leq 1$. If $C_q \leq 1$, Lemma 6.5 in Khalil (2002) provides the conclusion that the wind turbine system, from v^2 to ω , is finite-gain \mathcal{L}_2 stable. \square

The condition $C_q \leq 1$ is nearly always satisfied for modern turbines. The CART, for example, has an estimated $C_{q_{max}}$ of approximately 0.06. However, the motivation for the adaptive controller is that the C_p and C_q curves are unknown, so it would be nice to have a stability result independent of the assumption on C_q . Indeed, since the Betz Limit states that the maximum possible C_p for any real turbine is $16/27$, and the two curves are related by Equation 2.9, it can be guaranteed that $C_q \leq 1$ for any $\lambda \geq 16/27$. Thus, the \mathcal{L}_2 stability result is valid $\forall \lambda \geq 16/27$. When $\lambda \leq 16/27$, the question is simple, since by the definition of λ in Equation 2.3, it is known that $\omega = \frac{\lambda v}{R} \leq \frac{16}{27} \frac{v}{R}$. Thus, for finite λ , \mathcal{L}_∞ stability is given.

6.4 Convergence Characteristics of the Gain Adaptation Algorithm

Since the gain adaptation law performs no calculations during $(k-n)nT_s < t < knT_s$, M_{k-1} can replace M_{k-n} without loss of generality. (The discrete time index k has been changed to a subscript for convenience.) A few assumptions are made:

Assumption 1: M^* is constant. Although the turbine parameters (and thus the optimal gain M^*) change with time, this is a valid simplification because the turbine's physical changes are typically noticeable only over months or years, while the gain adaptation law has an adaptation period of less than a day.

Assumption 2: The P_{fav} vs. \tilde{M} curve has a shape similar to the one in Figure 6-2, at least in some local region around the optimal operating point. This is generally assumed to be true for any modern turbine. Specifically, the curve has a maximum at $\tilde{M} = 0$, is continuously differentiable, and is strictly monotonically increasing on $\tilde{M} < 0$ and strictly monotonically decreasing on $\tilde{M} > 0$.

Assumption 3: The adaptive controller has been operating sufficiently long that the specifics of the initial conditions provided to the controller are no longer relevant. If the initial conditions provided are M_0 , P_{fav0} , ΔM_0 , and M_I , then $k > 2$ is the time frame of interest.

6.4.1 Types of Instability

Begin the analysis by using intuition and considering the possible ways that the system could go unstable (i.e., $|\tilde{M}| \rightarrow \infty$ as $k \rightarrow \infty$). Clearly, one possibility is for $|\tilde{M}_k| > |\tilde{M}_{k-1}|$ with either $\text{sgn}(\tilde{M}_k) = 1 \forall k > 2$ or $\text{sgn}(\tilde{M}_k) = -1 \forall k > 2$. However, it is simple to show that this scenario

cannot occur with the gain adaptation law presented in Equations 2.18–2.20. Indeed, the error \tilde{M} will never take more than one consecutive step in the wrong direction.

Theorem 3: Given *Assumption 1*, the feedback system in Figure 6-1(b) with nonlinear controller (Equations 2.18 - 2.20) and the nonlinear plant given by *Assumption 2*, will never result in $|\tilde{M}_{k+1}| > |\tilde{M}_k| > |\tilde{M}_{k-1}|$ for $\text{sgn}(\tilde{M}_{k+1}) = \text{sgn}(\tilde{M}_k) = \text{sgn}(\tilde{M}_{k-1})$ for any $k > 2$.

Proof: Before the proof is begun, it should be noted that, by Equation (2.15), $P_{fav_{g_{k+1}}}$ is the calculation performed at the end of the adaptation interval during which $M = M_k$; thus, $P_{fav_{g_{k+1}}}$ corresponds to \tilde{M}_k .

Assume for simplicity that $\text{sgn}(\tilde{M}_{k-1}) = \text{sgn}(\tilde{M}_k) = \text{sgn}(\tilde{M}_{k+1}) = 1$. (The case in which $\text{sgn}(\tilde{M}_{k-1}) = \text{sgn}(\tilde{M}_k) = \text{sgn}(\tilde{M}_{k+1}) = -1$ can be proven identically with the appropriate sign changes.) Under this assumption, $|\tilde{M}_{k+1}| > |\tilde{M}_k| > |\tilde{M}_{k-1}|$ is equivalent to $\tilde{M}_{k+1} > \tilde{M}_k > \tilde{M}_{k-1}$. This sequence, which represents two consecutive steps away from $\tilde{M} = 0$, gives the following results:

$$\tilde{M}_k - \tilde{M}_{k-1} = -\Delta M_k > 0 \Rightarrow \Delta M_k < 0. \quad (6.8)$$

$$\tilde{M}_{k+1} - \tilde{M}_k = -\Delta M_{k+1} > 0 \Rightarrow \Delta M_{k+1} < 0. \quad (6.9)$$

Further, since $\tilde{M}_k > \tilde{M}_{k-1}$ in this scenario, it can be concluded that $P_{fav_{g_{k+1}}} < P_{fav_{g_k}}$ due to the strictly monotonically decreasing nature of this side of the P_{fav} vs. \tilde{M} curve. Thus, by Equation 2.20,

$$\Delta P_{fav_{g_{k+1}}} < 0. \quad (6.10)$$

Now, from Equation 6.8 and Equation 6.10 it is clear that $\text{sgn}(\Delta M_k) = \text{sgn}(\Delta P_{fav_{g_{k+1}}}) = -1$. Thus, by Equation 2.19, $\text{sgn}(\Delta M_{k+1}) = 1$. This contradicts Equation 6.9, and the conclusion is that it is impossible for $|\tilde{M}_{k+1}| > |\tilde{M}_k| > |\tilde{M}_{k-1}|$ for $\text{sgn}(\tilde{M}_{k+1}) = \text{sgn}(\tilde{M}_k) = \text{sgn}(\tilde{M}_{k-1})$ for any $k > 2$. Thus, it is impossible for the adaptive gain M_k to take two or more consecutive steps in the wrong direction for $k > 2$. \square

Since Theorem 3 shows that it is impossible for the sign of the adaptation step to be incorrect for more than one consecutive step, intuition would state that the magnitude of the adaptation step—specifically the gain $\gamma_{\Delta M}$ —is the critical factor in determining whether the gain adaptation law is stable. Figure 6-4 gives an example of a situation in which the gain $\gamma_{\Delta M}$ was large enough to cause instability of M . In this example, $|\tilde{M}_{k+1}| > |\tilde{M}_{k-1}|$, $\forall k > 2$ but $|\tilde{M}_{k+1}| \not> |\tilde{M}_k|$, $\forall k > 2$.




Figure 6-4: Instability example 2 for gain adaptation algorithm

6.4.2 Bounds on $\gamma_{\Delta M}$ Gain for Stability

Since the type of instability shown in Figure 6-4 can occur whenever $|\Delta M_k| > |\tilde{M}_{k-1}|$, it is logical to consider

$$|\Delta M_k| = |\tilde{M}_{k-1}|, \quad \tilde{M}_{k-1} \neq 0 \quad (6.11)$$

to be the critical case, which may be referred to as the marginal stability case. What gain $\gamma_{\Delta M}$ makes Equation 6.11 true? For the symmetrical curve

$$y(\tilde{M}_{k-1}) = y_k = \alpha_2 \tilde{M}_{k-1}^2 + \alpha_0, \quad (6.12)$$

where $\alpha_2 < 0$ and α_0 is any real number, $\gamma_{\Delta M}$ is simple to find.

In the critical gain scenario of this example, the system is oscillating among the three points plotted in Figure 6-5. If $\Delta M_k = \tilde{M}_{k-1}$, then $\tilde{M}_k = 0$ by Equation 2.18. Also, substituting y_k for P_{favgk} in Equation 2.20 and considering the definition of ΔM_k (Equation 2.19), in this case the




Figure 6-5: Finding the critical gain

gain $\gamma_{\Delta M}$ is such that $\Delta M_{k+1} = \Delta M_k$, so $\tilde{M}_{k+1} = -\tilde{M}_{k-1}$. Following the equations through one more step shows that, in this critical situation, $\tilde{M}_{k+2} = 0$ and the adaptive gain will oscillate among these three points. Thus, an upper bound on the gain $\gamma_{\Delta M}$ for stability can be found by equating

$$\Delta M_k = \tilde{M}_{k-1} = -\tilde{M}_{k+1} \quad (6.13)$$

and solving for $\gamma_{\Delta M}$ in terms of α_2 . Combining Equations 6.13, 2.19, and 6.12 and $\tilde{M}_k = 0$ gives the requirement

$$\gamma_{\Delta M} = \pm \sqrt{\frac{1}{|\alpha_2|}} \quad (6.14)$$

Since $\gamma_{\Delta M} > 0$, the positive value of (6.14) is chosen. Thus, the gain adaptation law (Equations 2.18–2.20) will not cause instability of \tilde{M}_k on the curve (Equation 6.12) whenever

$0 < \gamma_{\Delta M} < |\alpha_2|^{-1/2}$. In fact, since $\gamma_{\Delta M} = |\alpha_2|^{-1/2}$ is the marginal stability case, $0 < \gamma_{\Delta M} < |\alpha_2|^{-1/2}$ will actually cause convergence of $\tilde{M} \rightarrow 0$.

The rate of convergence of $\tilde{M} \rightarrow 0$ is dependent on several factors, including the adaptation period, the size of $\gamma_{\Delta M}$, and the shape of the P_{favg} vs. \tilde{M} curve. A longer adaptation period will likely slow the convergence overall but may reduce the number of adaptation steps required for M to come within a certain percent error from M^* . Based on simulation results, a longer adaptation period will also reduce the magnitude of the oscillations around $\tilde{M} = 0$, as well. Assuming that $\gamma_{\Delta M}$ is small enough for the gain adaptation law to be stable, the size of $\gamma_{\Delta M}$ also affects the rate of convergence towards $\tilde{M} = 0$. A smaller $\gamma_{\Delta M}$ will result in a slower approach but smaller oscillations around $\tilde{M} = 0$. Similarly, a system with a flatter P_{favg} vs. \tilde{M} curve will experience a slower approach of $\tilde{M} \rightarrow 0$ because its step sizes will be smaller but will also undergo smaller oscillations around $\tilde{M} = 0$.

6.4.3 Asymmetric P_{favg} vs. \tilde{M} Curves

Since the requirement on $\gamma_{\Delta M}$ is dependent on the magnitude $|\alpha_2|$, any $\gamma_{\Delta M}$ chosen for a specific curve will also guarantee convergence of M_k on a shallower curve (i.e., one with a smaller $|\alpha_2|$) since that curve would have a larger allowable $\gamma_{\Delta M}$ by Equation 6.14. A similar result can be stated for an asymmetric curve: if the gain $\gamma_{\Delta M}$ is chosen to guarantee convergence based on the slope of the steeper side of the curve, it would guarantee convergence over the entire curve, since the step sizes ΔM_k calculated on the shallower side of the curve would be smaller in magnitude than those on the steeper side for a given ΔM_{k-1} . Thus, for any turbine P_{favg} vs. \tilde{M} curve, there exists a $\gamma_{\Delta M}$ that guarantees convergence of the adaptive gain M_k , and this $\gamma_{\Delta M}$ depends on the steepness of the P_{favg} vs. \tilde{M} curve.

Unfortunately, the P_{favg} vs. \tilde{M} curve is not explicitly known for any turbine. In practice it is necessary to approximate the shape and slope of the curve for a given turbine. A more conservative selection of $\gamma_{\Delta M}$ is then likely to result in stability (and convergence), but will also result in smaller step sizes and thus slower convergence.

An example of the selection of $\gamma_{\Delta M}$ for the CART is provided in Figure 6-6. The y curve is chosen to fit snugly inside the P_{favg} curve while satisfying $y < P_{favg}$; in this case, $\alpha_2 = -0.00001$. Thus, the maximum allowable $\gamma_{\Delta M}$ for “guaranteed” stability is 316. The gain used in testing on the actual CART (before this stability analysis was performed) was $\gamma_{\Delta M} = 100$. This $\gamma_{\Delta M} = 100$ value was determined empirically from simulations and early hardware testing. While actual turbine results indicate stable performance of the adaptive control law, this stability analysis provides further reassurance and guidelines in choosing $\gamma_{\Delta M}$.

A formal proof of the stability of the feedback system given by Figure 6-1(b) is provided in Appendix B with one modification: Equation 2.19 is replaced by

$$\begin{aligned} \Delta M(k) &= \gamma_{\Delta M} \operatorname{sgn}[\Delta M(k-n)] \Delta P_{favg}(k) \\ &= \gamma_{\Delta M} \operatorname{sgn}[\Delta M(k-n)] \operatorname{sgn}[\Delta P_{favg}(k)] |\Delta P_{favg}(k)|^2 \end{aligned} \quad (6.15)$$




Figure 6-6: Ratio of mean captured power to mean wind power vs. error in torque control gain M

which is identical except for the square root. An argument is presented in Remark 2 following the proof as to the why the modified proof is still relevant.

CHAPTER 7: Extensions to Original Gain Adaptation Law

Experimentally demonstrating that the adaptive control law (Equation 2.14) has increased energy capture compared to the original non-adaptive law (Equation 2.1) on the CART is very promising, but there are still some negative aspects of Equations 2.14 and 2.18–2.20 and the ways in which they were implemented that make them less desirable to a commercial turbine designer. The biggest of these is likely the fact that an upwind anemometer was used for the P_{wy} calculation (Equation 2.16), since this anemometer would likely be unavailable on a commercial machine. In addition, the fact that the proposed gain adaptation law (Equations 2.18–2.20) incorporates only one past value into its calculations makes it more oscillatory and more prone to measurement and sensor errors than, for example, a filtered law might be. This chapter addresses these two concerns by making extensions on the gain adaptation law (Equations 2.18–2.20).

7.1 Use of Nacelle Anemometer

The earliest adaptive control tests performed on the CART actually used the nacelle anemometer rather than the anemometer on the met tower (the “upwind anemometer”) for the calculation of P_{wy} (Equation 2.16). However, the sensor was changed after it was determined that P_{favg} (Equation 2.15) values were higher than the Betz Limit, which was known to be impossible. Although the aerodynamics behind the rotor and around the nacelle are very complicated while the turbine is operating, it was then believed that the major errors in the nacelle anemometer wind speed measurements are caused by the well known axial induction factor. The axial induction factor is a measure of the slowing of the wind from the upwind, or free-stream, velocity to the velocity at the rotor plane (Manwell et al. 2002). The effect of the axial induction factor is that, for a given upwind wind speed, a faster spinning rotor (i.e., higher tip speed ratio) results in slower wind downwind of the rotor. Thus, the measured P_{wy} actually appears lower than the true available wind power, and the value of P_{favg} is misleadingly larger. As the adaptive gain M decreased at each adaptation step, the average tip speed ratio increased; P_{favg} became even larger, as a result, it was thought, of the axial induction factor; and an undesirable positive feedback loop formed. For this reason, the use of the nacelle anemometer was abandoned early in the testing. Figure 7-1 shows the CART and the general locations of the anemometers, with the upwind anemometer on the met tower upwind of the turbine and the nacelle anemometer on top of the nacelle at the back end.

Later experimental results, however, necessitated a re-evaluation of this early conclusion. First, the 10-minute adaptation periods used in the earliest experiments were found to be not nearly long enough to provide a good correlation between measured wind speed at either anemometer and the wind hitting the rotor plane. Thus, some of the effects believed to be the result of the axial induction factor may have been merely the result of this poor correlation. Second, further data analysis revealed that the CART’s true optimal M was far below (around 47%, as discussed in Section 5.1) the originally estimated value M^+ , so the fact that M kept adapting down in early




Figure 7-1: CART (a) operating at the NWTC with upwind met tower shown, and (b) close-up view showing the nacelle anemometer

experiments may not have been attributable to the axial induction factor at all. In fact, the use of the nacelle anemometer for the wind speed measurement may have given better results than were recognized at the time. However, the fact remains that the wind is slowed by the rotor, and, as a result, using the nacelle anemometer in the P_{wy} calculation does cause the calculated P_{favg} to be higher than the Betz Limit. The question addressed in this section is whether this slowing factor causes problems for the gain adaptation law, or whether it is consistent enough over the operating range that it can simply be viewed as a constant multiplying factor in the calculations. If it is like a constant multiplying factor, then the “Slope Error” simulations performed in Section 4.3 suggest that it should not cause significant problems for the gain adaptation law.

7.1.1 CART Adaptations Using the Nacelle Anemometer

In fact, the adaptive gain M seems to converge nearly as well when the nacelle anemometer is used for the wind speed measurement as when the upwind anemometer is used, as shown in Figure 7-2. The CART’s estimated optimal value of M , M^* (Section 5.1), is plotted as the solid horizontal line, and the 180-minute adaptations using the upwind and nacelle anemometers are plotted for comparison. One difference between the two plots is that, at 21 hours, the nacelle anemometer M takes a second step in the wrong direction (assuming the horizontal solid line is the true M^*). The reason for this, and for the general oscillatory nature of both curves, is assumed to be the shorter-than-optimal adaptation period. However, both curves show general agreement on the approximate optimal value of M , and it is hoped that future testing with longer adaptation periods will show even better convergence properties.

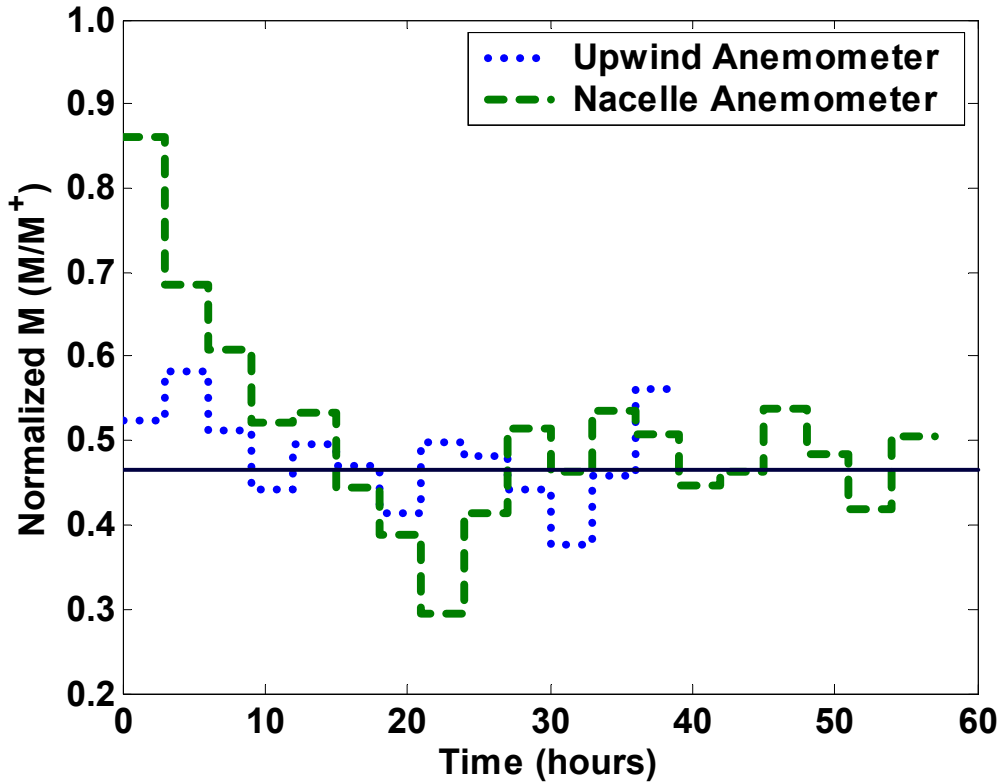


Figure 7-2: Adaptations of M on the CART using upwind and nacelle anemometers

Figure 5-4 shows a plot of the fractional average power P_{favg} vs. normalized M , including the best-fit quadratic, obtained from running the CART using the upwind anemometer. In this figure, the power values were mostly close to the C_{pmax} predicted from data analysis ($C_{pmax} \cong 0.41$). However, the presence of the slowing factor is clear in a similar plot, Figure 7-3, created using the nacelle anemometer data. In this plot, half of the data lie above the Betz Limit, which means that there must be an error in either the wind power P_{wy} (Equation 2.16) or the turbine captured power P_{cap} (Equation 2.17) measurement. Even the data that lie below the Betz Limit are unreasonably high for a real turbine, and certainly too high for the CART, given the knowledge obtained from much of the other data collected in various experiments. However, the peak of the best fit curve occurs at $M = 0.50M^+$, which is very close to the peak of the best fit curve of Figure 5-4 ($M = 0.49M^+$). Since the adaptive controller cares only where the peak of the P_{favg} vs. \tilde{M} curve occurs and not its actual magnitude, these data support the idea that the nacelle anemometer is a valid instrument to use in the CART's gain adaptation law calculations.

7.1.2 Effects of Tip Speed Ratio on Axial Induction Factor

Since it appears that the adaptive controller treats the CART's nacelle wind speed measurement as being some constant multiple of the upwind wind speed measurement, it seems likely that there would be a high correlation between the two measurements. Randall et al. (2002) examined experimental data and derived best-fit lines for the relationship between upwind and




Figure 7-3: Fractional average power P_{favg} vs. normalized M

nacelle wind speed on several different turbines. Smaïli and Masson (2003) also addressed this question for one particular turbine, deriving some relevant equations and finding correlations between the two measurements based on experimental data. They found that the relationship between the two wind speed measurements varies with blade pitch angle and is dependent on the shape of the nacelle. In fact, common sense dictates that the geometry of the nacelle and blades and the precise location of the nacelle anemometer all play an important role in determining this relationship. Thus, the specific equations in Smaïli and Masson (2003) are not expected to hold for the CART's upwind and nacelle wind speed relationship. Note especially that the instantaneous correlation between the two measurements will be poor, since there is a time delay of several seconds between the two sensors; averages over time periods significantly longer than that delay, however, should show a correlation.

Figure 7-4 shows a plot of about 350 10-minute wind speed averages. The best fit line is superimposed on the data and shows a high correlation ($R^2 = 0.98$) between the two. The equation of the best fit line is $v_n = 0.8543v_u - 0.04720$, where v_n is the wind speed measured by the nacelle anemometer and v_u is the wind speed measured by the upwind anemometer.

Figure 7-4 includes data for M having a range of $0.27M^+$ to $1.42M^+$. Thus, it does not provide any information regarding how the relationship between the two measurements of wind speeds varies with tip speed ratio or axial induction factor. However, it does provide support for the idea that the gain adaptation law works when the nacelle anemometer is used because the nacelle




Figure 7-4: Relationship between nacelle and upwind wind speed measurements

wind speed measurement is very close to a scaled version of the “true” (upwind) wind speed measurement. There is an offset of 0.05 m/s present in the equation, but this value is small (less than 1%) compared with normal operating wind speeds of 8-9 m/s. Since wind power P_{wy} (Equation 2.16) is a cubic rather than linear function of wind speed, the exact percentage error in P_{wy} caused by this offset error in the wind speed measurement is not consistent from one wind data set to another. However, for one 2-hour wind data set used as an input in this report, the 0.05 m/s offset in wind speed corresponds to a 1.3% error in the wind power measurement. For comparison, the offset error simulated in Section 4.3 was about 7% of the power calculation. Thus, the wind measurement offset error that results from using the nacelle anemometer may cause small problems for the gain adaptation law but it is small compared to the offset that caused such significant problems in Section 4.3. The scaling factor of 0.8543 is a much larger percentage of the wind speed, but this type of error (constant slope error) is handled much more easily by the gain adaptation law (Equations 2.18–2.20), as discussed in Section 4.3.

The question, then, becomes whether the “constant” scaling factor observed in Figure 7-4 is truly constant or whether it depends on the tip speed ratio. Figure 7-5 shows the same data presented in Figure 7-4 divided into six different plots based on the mean tip speed ratio λ for each data set. Once again, the best fit line is superimposed over the data in each plot. The y-intercept on each plot is very close to zero, and the slopes appear to be very similar. The relevant data for each subplot are listed in Table 7-1, where the coefficients a and b are from the equation $v_n = av_u + b$. From the table, it is apparent that the slopes of each best-fit line are close (within about 5%) and,

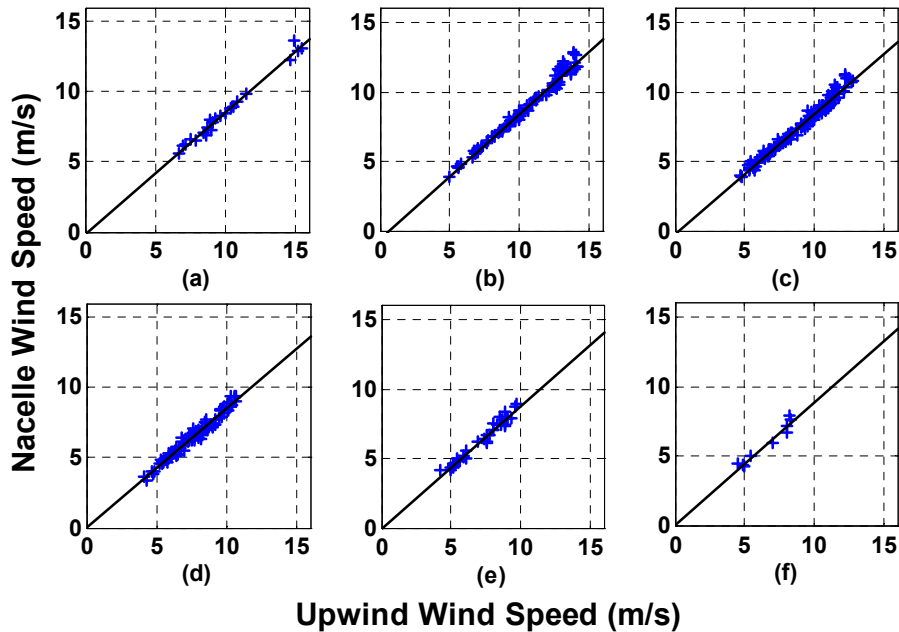


Figure 7-5: 10-minute averages of nacelle vs. upwind wind speed for various tip speed ratios

more importantly, that there is no clear trend relating the slopes and the tip speed ratio intervals to which they correspond. Thus, it seems that the axial induction factor does not play as large a role on the CART as originally believed, which may be because the CART's blades taper to ellipses near their roots rather than retaining a good airfoil shape. In any case, the nacelle anemometer should provide a valid measurement to use in the CART's gain adaptation law.

Table 7-1. Relevant Data for Figure 7-5

Plot	# of Points	Range of λ	a	b	R^2
(a)	22	(0.0,6.5]	0.86605	-0.13357	0.98646
(b)	90	(6.5,7.5]	0.89567	-0.51903	0.98214
(c)	129	(7.5,8.5]	0.85451	-0.13552	0.97966
(d)	85	(8.5,9.5]	0.85262	0.0004371	0.96954
(e)	35	(9.5,10.5]	0.88175	-0.020186	0.96728
(f)	9	(10.5, ∞)	0.87757	0.13417	0.95346

7.1.3 Simulations of Axial Induction Factor

Although the data presented above show that the axial induction factor does not prevent the use of the nacelle anemometer for this adaptive controller on the CART, it cannot be concluded that the same result holds for all turbines. Since no other real turbines were available for this test, simulations were performed to assess the effect of the axial induction factor on a turbine with a strong correlation between axial induction factor and tip speed ratio. In fact, this relationship was obtained from the earlier PROP analysis of the CART, but this research has already shown

that the PROP-derived data do not closely resemble data for the CART. Thus, these axial induction factor simulations can be considered to represent a separate turbine.

The axial induction factor causes the wind speed at the rotor to be slower than the upwind, or free-stream, wind. The relationship between the two wind speeds is given by

$$v_r = v_\infty(1 - a), \tag{7.1}$$

where v_r is the wind speed at the rotor and v_∞ is the free-stream wind speed. For the purposes of these simulations, v_r is assumed to be equal to the wind speed measured by the nacelle anemometer, v_n . If the axial induction factor a were constant for a given turbine, using v_n (equivalently, v_r in Equation 7.1) in the gain adaptation law's calculation of P_{wy} (Equation 2.16) would simply result in a slope error on the wind power measurement. As discussed in Section 4.3, this slope error is not likely to result in a change of the convergence properties of the gain adaptation law, though it might change the step size and thus the rate of convergence. However, in this simulation a is not constant but instead varies nonlinearly with tip speed ratio λ . Before performing the axial induction factor simulations, this relationship must be established.

The relationship between axial induction factor and tip speed ratio is nonlinear and poorly defined. The aerodynamics behind the relationship are very complex and are outside the scope of this research, but may be studied further in Manwell et al. (2002) and other wind energy textbooks. The PROP simulations provide an estimate of the CART's relationship between tip speed ratio and thrust coefficient, C_T ; this simulated relationship is plotted in Figure 7-6(a). The thrust coefficient describes the ratio of the thrust force on the turbine and the dynamic force applied, similarly in nature to the power and torque coefficients C_p and C_q discussed previously. Wind turbine aerodynamic theory (Manwell et al. 2002) can then be used to relate axial

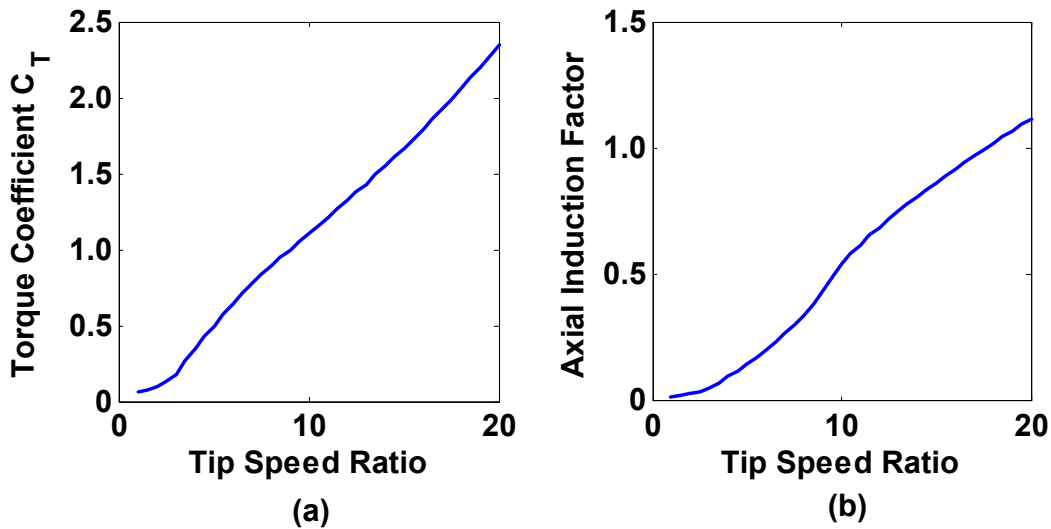


Figure 7-6: PROP-derived CART model showing (a) thrust coefficient vs. tip speed ratio, and (b) axial induction factor vs. tip speed ratio

induction factor a and thrust coefficient C_T as follows:

$$a = \begin{cases} \frac{1}{2} - \frac{1}{2} \sqrt{1 - C_T}, & C_T \leq 0.96 \\ \frac{1}{F} \left(0.143 + \sqrt{0.0203 - 0.6427(0.889 - C_T)} \right), & C_T > 0.96 \end{cases} \quad (7.2)$$

where F is a correction factor depending on the radial distance from the axis of rotation and other factors and is assumed to be equal to 1.0 for these simulations. The first term in a is derived using blade element momentum theory for a non-rotating wake and the second is commonly known as Glauert's empirical relationship. Figure 7-6(b) combines Equation 7.2 and the PROP-derived relationship of C_T and λ shown in Figure 7-6(a) to determine the $a - \lambda$ relationship used in the simulations in this section.

Now that the $a - \lambda$ relationship has been established, the simulations follow in a straightforward manner. The SimInt model used for the Chapter 4 simulations and shown in Figure 3-1 is modified slightly by Equation 7.1 so that the nacelle (rotor plane) wind speed is used in the P_{wy} (Equation 2.16) calculations, effectively resulting in a variable slope error on the measured wind speed. The PROP-derived C_q surface is used in the simulation, so a properly functioning gain adaptation law should cause M to converge to some value within the range $(0.8M^+, 1.0M^+)$, depending on the turbulence intensity of the wind input. The adaptation period is three hours rather than 13 in order to reduce the simulation time and because the longer time period is not necessary in order to observe the expected results. The simulation results are plotted in Figure 7-7 and Figure 7-8.

Figure 7-7 shows exactly the undesirable effects of axial induction factor that were expected. After the first step up due to initial conditions, the adaptive gain M proceeds to adapt down for the remainder of the simulation as the fractional average power P_{favg} rises. Unfortunately, this rise is not a real effect, but rather the effect of the variable slope error on the wind measurement. Note that the fractional average power is always greater than one in the lower plot. This impossible result—in which the turbine produces more energy than it extracts from the wind—is due to the slope error on the wind power measurement given by Equation 7.1.

Figure 7-8 shows the time series plots of axial induction factor and tip speed ratio during the simulation and the best fit lines for each. Clearly, both a and λ rise throughout the simulation. The reason for the increase in λ is the decrease in M , which for a given wind speed causes higher rotor accelerations by Equation 2.7 and thus higher tip speed ratios. As the tip speed ratio increases, the axial induction factor increases by Figure 7-6(b), which then causes the measured wind speed $v_n = v_r$ to decrease by Equation 7.1. This slower measured wind speed causes P_{wy} (Equation 2.16) to appear falsely lower, which then causes P_{favg} (Equation 2.15) to appear falsely larger at the next adaptation step, which in turn causes M to continue to adapt in the same downward direction by Equations 2.18–2.20. This positive feedback loop will continue indefinitely unless external controls are in place to prevent M from leaving a certain range, for example.

This section has provided conflicting information in the form of real turbine data showing that the axial induction factor does not cause improper adaptations when the nacelle anemometer is used for wind speed measurements and simulation data showing that it does cause improper




Figure 7-7: Normalized M and fractional average power P_{avg} for axial induction factor simulation

adaptations. These conflicting data lead to the conclusion that the decision on whether or not to use the nacelle anemometer or upwind anemometer for the wind measurement is turbine-dependent because of the differences in various turbines' blade and nacelle geometry and nacelle anemometer location.

The use of the nacelle anemometer rather than the upwind anemometer for measuring the wind speed would make it easier to implement the adaptive controller (Equation 2.14) on a commercial wind farm. However, it does little to address the highly oscillatory behavior of the gain adaptation algorithm or the lengthy adaptation period required for good results. Filtering the original gain adaptation law might provide improvements in these two areas, and a new filtered law is the topic of Section 7.2.

The use of the nacelle anemometer rather than the upwind anemometer for measuring the wind speed would make it easier to implement the adaptive controller (Equation 2.14) on a commercial wind farm. However, it does little to address the highly oscillatory behavior of the gain adaptation algorithm or the lengthy adaptation period required for good results. Filtering the original gain adaptation law might provide improvements in these two areas, and a new filtered law is the topic of Section 7.2.




Figure 7-8: Axial induction factor a and tip speed ratio λ obtained in simulation

7.2 Filtered Gain Adaptation Law

Filtering the very basic gain adaptation law (Equation 2.18) might help reduce the oscillations in M as it converges. The modified gain adaptation law is of the form

$$M(k) = c_1 M(k-1) + c_2 M(k-2) + b_0 \Delta M(k) + b_1 \Delta M(k-1), \quad (7.3)$$

where $\underline{c} = [c_1, c_2]$ and $\underline{b} = [b_0, b_1]$ are constant coefficients chosen for the desired transient and convergence properties. Taking the z-transform of Equation 7.3 gives the result

$$\frac{M(z)}{\Delta M(z)} = \frac{b_0 z^2 + b_1 z}{z^2 - c_1 z - c_2}, \quad (7.4)$$

which has poles at $\frac{1}{2}(c_1 \pm \sqrt{c_1^2 + 4c_2})$. If $c_1 + c_2 = 1$ is constrained (so that in a sense the overall magnitude of M remains unchanged from the original law (Equation 2.18), the poles of Equation 7.4 are located at $(c_1 - 1, 1)$. In general, locating a pole on the unit circle gives only marginal stability, but in this case, the goal is not to reject disturbances in ΔM but rather to reject disturbances in atmospheric factors such as wind speed and direction. Since the equations

relating M and these atmospheric factors are nonlinear, it is impossible to use the z-transform to place these poles as desired. In this case, intuition and simulations were used in order to choose the values of \underline{b} and \underline{c} .

The first simulations used constant wind speed as the input and Equation 7.3 as the gain adaptation law, with the rest of the system being identical to those used in the original gain adaptation law in Chapter 2. The SymDyn-generated C_q surface was used in the τ_{aero} calculation (Equation 2.8). These results are plotted in Figure 7-9, with each curve normalized to M^+ as usual.

Based on the simulation results plotted in Figure 7-9, it seems that the best filter, marked with stars, has coefficients $\underline{b} = [1.0, 0.0]$ and $\underline{c} = [0.75, 0.25]$. This curve has no overshoot—ignoring the initial step forced upon all six simulations—and converges to the same final value as the original curve (plotted as the solid line). However, intuitively it seems that it would be nice to have a filter on the ΔM term, as well (i.e., $b_1 \neq 0$), since this value is highly subject to wind disturbances. Of course, wind disturbances are not present in a constant-wind input simulation, so the value of the b_1 term is not apparent. In any case, the best of the three filters with nonzero b_1 is the curve marked with diamonds and having coefficients $\underline{b} = [0.9, 0.1]$ and $\underline{c} = [0.75, 0.25]$. The best filtered gain adaptation laws were next tested with real wind data as input and compared with the original. These results are plotted in Figure 7-10.

In the first 90 hours, there is very little difference between the curves showing the filtered gain




Figure 7-9: Simulation results for new filtered gain adaptation law with constant wind speed input

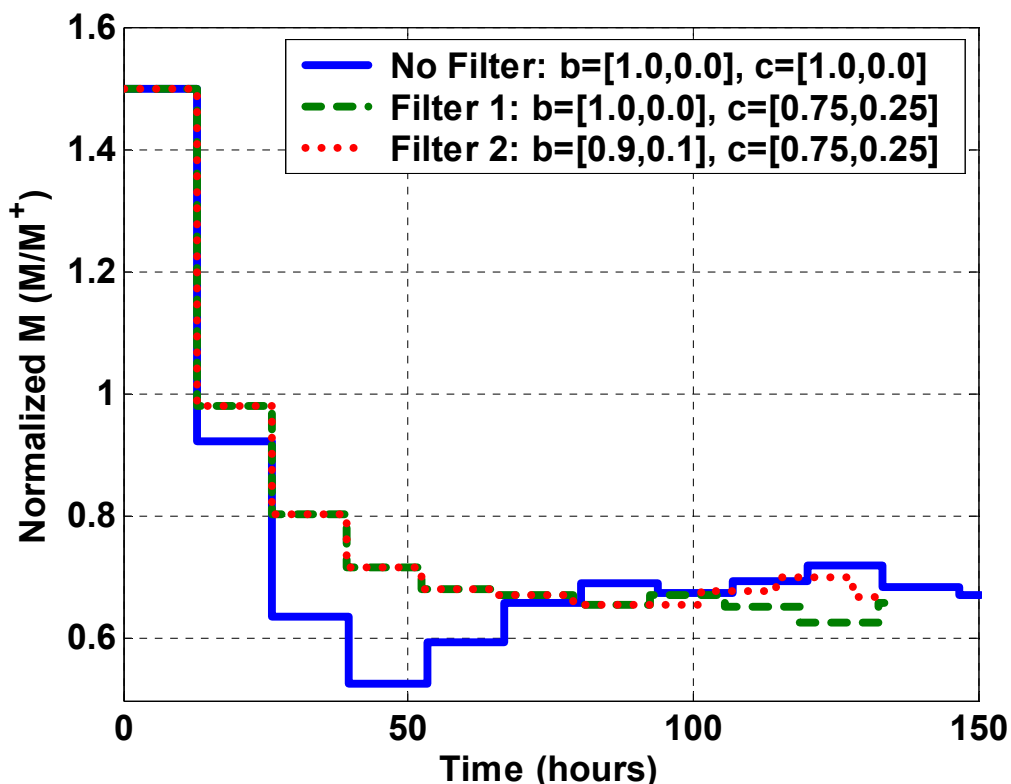


Figure 7-10: Simulation results for new filtered gain adaptation law with real (time-varying) wind speed input

adaptation law outputs, and it is promising that both approach nearly identical final values at the end of 130 hours of simulation. Also, neither experiences the same overshoot apparent in the no filter case between 26 and 65 hours. Because of the time constraints on CART testing, further simulation studies could not be performed before a decision had to be made about which set of filter coefficients to use on the CART. Therefore, the intuitively preferred “Filter 2” was chosen and testing commenced. Unfortunately, several sensor problems and component failures that occurred during the time the filtered gain adaptation law was being tested on the CART prevented the controller from obtaining enough accurate measurements that a conclusion could be drawn as to its effectiveness. Obtaining more data using this algorithm thus is one of the recommendations for future work.

Filtering the gain adaptation law with a well-chosen filter can provide superior results in terms of the convergence properties of M in simulation. Some of the filters tested in simulation were found to be superior to the original, unfiltered gain adaptation law. Unfortunately, no real CART data are available to support the simulation conclusions.

CHAPTER 8: Conclusions and Future Recommendations

A new adaptive control strategy has been proposed, tested, and analyzed in this research. The new controller is designed to reduce the energy loss that results from uncertainty in a turbine's aerodynamic parameters. Specifically, the new adaptive control strategy replaces the constant gain torque controller commonly used by the wind industry in generator torque control with a new adaptive gain, using measurements of turbine power and wind power as the basis for the gain adaptation law.

Because the new adaptive controller described in Chapters 2 through 6 keeps the same structure as the standard generator torque controller in region 2 or subrated operation, it is already familiar to the wind turbine industry. This familiarity should make it easier to implement the new adaptive controller commercially, since many of the controller implementation issues have already been worked out.

Two different simulators have been used to examine the operation of the new gain adaptation law: SimInt and SymDyn. SimInt is a new tool designed specifically for this research that has limited applicability to other problems. Specifically, SimInt is capable of simulating only the rotor angular speed and requires the production of a nonlinear C_q surface for the desired turbine via a separate simulator. Because of its simplicity, however, SimInt can run very quickly compared to other simulation tools. SymDyn, on the other hand, is capable of simulating 10-11 degrees of freedom (depending on the number of blades) but is about five times slower than SimInt. These two independent simulation tools have produced similar results when provided with the same initial conditions, lending credibility to each other and to the assertion that the new adaptive controller and gain adaptation law perform as desired. Simulations have been performed assuming perfect and imperfect wind input measurements and constant and time-varying turbine aerodynamic parameters. In general, the adaptive controller has been shown in simulation to operate as desired when provided with realistic operating conditions and without causing harm to the turbine.

In addition to these simulations, tests have been performed on a real, mid-sized turbine to verify the effectiveness of the adaptive controller. From these experiments it was determined that the aerodynamic properties of this research turbine were not well known, and a significant increase in energy capture was achieved by using the adaptive controller. This result lends credence to the idea that this adaptive control could provide significant benefits to the commercial turbine industry, since it is reasonable to assume that the aerodynamic properties of this research turbine have been studied more extensively than those of many commercial turbines. Additionally, despite the fact that the adaptation period of the real turbine's gain adaptation law was significantly shorter than those used in the simulations, the adaptive gain did converge toward what is believed to be the optimal value for the turbine.

Concurrently with the real turbine experiments and simulations, a stability analysis was performed on the standard and adaptive control laws. The analysis showed theoretically that the standard law does achieve some desired stability results. Also, the gain adaptation law does

appear to cause the adaptive gain to converge toward the optimal value assuming that the gain within the gain adaptation law is small enough, and a technique was presented for finding that small enough gain. A more formal proof of the stability of a very similar gain adaptation law was developed; it is presented in Appendix B.

Finally, two extensions to the adaptive controller were proposed and tested. The first, which entailed the use of the nacelle anemometer rather than the upwind meteorological tower anemometer in the wind speed measurement, was a simple change that makes it even easier to implement this adaptive controller on a commercial turbine. Secondly, a filtering rule was proposed and tested in simulation to reduce the overshoot and oscillatory behavior of the adaptive gain as it converges toward its optimal value.

In summary, the research described in this report has achieved the following:

- Development of a new adaptive control law to increase region 2 energy capture on variable speed turbines
- Verification of the new adaptive controller in simulation and testing to demonstrate effectiveness and safety
- Theoretical analysis of the standard non-adaptive controller and the new adaptive controller to support simulation and testing results and provide further insight into the problem
- Extension of the new adaptive controller to make it more appealing to the commercial wind turbine industry.

Recommendations for future work along this line of research lie mainly within the area of additional real turbine testing. Not only should the adaptation period be extended significantly in an attempt to improve the convergence properties of the adaptive gain, but the adaptive controller should also be run for a long enough time period that its effectiveness at tracking aerodynamic properties over long time periods can be evaluated. Unfortunately, testing time on real wind turbines is limited, and the latter experiment could take years to perform.

In addition to the real turbine testing, more work will be done on the proof of stability and convergence of the gain adaptation law. A more formal, thorough proof will help to make the adaptive controller more acceptable within the wind industry.

References

- Bhowmik, S., Spée, R., and Enslin, J. 1999. "Performance Optimization for Doubly-Fed Wind Power Generation Systems." *IEEE Transactions on Industry Applications*, 35, No. 4, pp. 949-958.
- Bossanyi, E. 1989. "Adaptive Control of the MS2 Wind Turbine. Practical Results." *Wind Engineering*, Vol. 13, No. 5, pp. 259-273.
- Bossanyi, E. 1987. "Adaptive Pitch Control for a 250 kW Wind Turbine." *Proceedings of the British Wind Energy Conference*, pp. 85-92.
- Buhl, M., Jonkman, J., Wright, A., Wilson, R., Walker, S., and Heh, P. 2003. *FAST User's Guide*. NREL/EL-500-29798, Golden, CO: National Renewable Energy Laboratory.
- Buhl, M., Wright, A., and Pierce, K. 2001. "FAST_AD Code Verification: A Comparison to ADAMS." *Proceedings of the 20th ASME Wind Energy Symposium*, Reno, NV, pp. 368-377.
- Burton, T., Sharpe, D., Jenkins, N., and Bossanyi, E. 2001. *Wind Energy Handbook*. New York: John Wiley & Sons, Ltd.
- Casanova Alcalde, V. and Freris, L., 1982, "Control Policies for Maximising Energy Extraction from Wind Turbines," *Proceedings of the 4th International Symposium on Wind Energy Systems*, Stockholm, Sweden, pp. 233-246.
- Eggers, A., Ashley, H., Chaney, K., Rock, S., and Digumarthi, R. 1998. "Effects of Coupled Rotor-Tower Motions on Aerodynamic Control of Fluctuating Loads on Light-Weight HAWTs." *Proceedings of the 17th ASME Wind Energy Symposium*. Reno, NV, pp. 113-122.
- Eggers, A., Moriarty, P., Chaney, K., Digumarthi, R., and Holley, W. 2002. "Influence of Transition Modes and Gravity Loads on Rotor Fatigue and Power Control." *Proceedings of the 21st ASME Wind Energy Symposium*, Reno, NV, pp. 237-245.
- Enslin, J., and Van Wyk, J. 1992. "A Study of a Wind Power Converter with Micro-Computer Based Maximal Power Control Utilising an Over-Synchronous Electronic Scherbius Cascade." *Renewable Energy*, Vol. 2, No. 6, pp. 551-562.
- Fingersh, L., and Carlin, P. 1998. "Results from the NREL Variable-Speed Test Bed." *Proceedings of the 17th ASME Wind Energy Symposium*, Reno, NV, pp. 233-237.
- Fingersh, L., Simms, D., Hand, M., Jager, D., Cotrell, J., Robinson, M., Schreck, S., and Larwood, S. 2001. "Wind Tunnel Testing of NREL's Unsteady Aerodynamics Experiment." *Proceedings of the 20th ASME Wind Energy Symposium*, Reno, NV, pp. 194-200.

Fingersh, L., and Johnson, K. 2004. "Baseline Results and Future Plans for the NREL Controls Advanced Research Turbine." *Proceedings of the 23rd ASME Wind Energy Symposium*, Reno, NV, pp. 87-93.

Freeman, J., and Balas, M. 1999. "An Investigation of Variable Speed Horizontal-Axis Wind Turbines Using Direct Model-Reference Adaptive Control." *Proceedings of the 18th ASME Wind Energy Symposium*, Reno, NV, pp. 66-76.

Hand, M. 2003. *Mitigation of Wind Turbine/Vortex Interaction Using Disturbance Accommodating Control*. Ph.D. thesis, Boulder, CO: University of Colorado.

Hand, M., and Balas, M. 2000. "Systematic controller design methodology for variable-speed wind turbines." *Wind Engineering*, Vol. 24, No. 3, pp. 169-187.

Johnson, K., Fingersh, L., Balas, M., and Pao, L. 2004a. "Methods for Increasing Region 2 Power Capture on a Variable Speed HAWT." *Proceedings of the 23rd ASME Wind Energy Symposium*, Reno, NV, pp. 103-113.

Johnson, K., Fingersh, L., Balas, M., and Pao, L. 2004b. "Methods for Increasing Region 2 Power Capture on a Variable Speed HAWT." Accepted for publication in the *Journal of Solar Energy Engineering*.

Johnson, K., Pao, L., Balas, M., Kulkarni, V., and Fingersh, L. "Stability Analysis of an Adaptive Torque Controller for Variable Speed Wind Turbines." Submitted for publication in the proceedings of the 2004 Conference on Decision and Control.

Khalil, H. K. 2002. *Nonlinear Systems*. Third Edition. Upper Saddle River, NJ: Prentice Hall, p. 242.

Leith, D., and Leithead, W. 1997. "Implementation of Wind Turbine Controllers." *International Journal of Control*, Vol. 66, No. 3, pp. 349-380.

Manwell, J., McGowan, J., and Rogers, A. 2002. *Wind Energy Explained: Theory, Design, and Application*. New York: John Wiley & Sons, Inc., p. 86.

Pierce, K., 1999, "Control Method for Improved Energy Capture Below Rated Power," *Proceedings of the 3rd ASME/JSME Joint Fluids Engineering Conference*, Paper No. FEDSM99-S295-12, San Francisco, CA.

Randall, G., McCoy, T., Smith, B, and Link, H. 2002. "Applicability of Nacelle Anemometer Measurements for Use in Turbine Power Performance Tests." *Proceedings of the AWEA Windpower Conference*, Portland, OR.

Safonov, M. G. 1980. *Stability and Robustness of Multivariable Feedback Systems*, Cambridge, MA: The MIT Press.

- Smaïli, A., and Masson, C. 2003. "Numerical Investigations of Nacelle Anemometry for Horizontal Axis Wind Turbines." *Proceedings of the 22nd ASME Wind Energy Symposium*, Reno, NV, pp. 389-398.
- Stol, K., and Balas, M. 2002. "Periodic Disturbance Accommodating Control for Speed Regulation of Wind Turbines." *Proceedings of the 21st ASME Wind Energy Symposium*, Reno, NV, pp. 310-320.
- Stol, K., and Bir, G. 2003. *SymDyn User's Guide*, NREL/EL-500-33845. Golden, CO: National Renewable Energy Laboratory.
- Stol, K., and Bir, G. 2000. "Validation of a Symbolic Wind Turbine Structural Dynamics Model." *Proceedings of the 19th ASME Wind Energy Symposium*, Reno, NV, pp. 41-48.
- Song, Y., Dhinakaran, B., and Bao, X. 2000. "Variable Speed Control of Wind Turbines Using Nonlinear and Adaptive Algorithms." *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 85, pp. 293-308.
- Walker, S., and Wilson, R. 1976. *Performance Analysis Program for Propeller Type Wind Turbines*. Corvallis, OR: Oregon State University.
- Wright, A. 2003. *Modern Control Design for Flexible Wind Turbines*. Ph.D. thesis. Boulder, CO: University of Colorado.
- Zames, G. 1966a. "On the input-output stability of time-varying nonlinear feedback systems—Part I: Conditions derived using concepts of loop gain, conicity, and positivity." *IEEE Transactions on Automatic Control*, Vol. 11, No. 2, pp. 228-238.
- Zames, G. 1966b. "On the input-output stability of time-varying nonlinear feedback systems—Part II: Conditions involving circles in the frequency plane and sector nonlinearities." *IEEE Transactions on Automatic Control*, Vol. 11, No. 3, pp. 465-476.

APPENDIX A: Mean Equivalent Wind Speed

It is common practice within the wind industry to create plots showing mean power versus mean wind speed over some averaging period in order to assess the performance of a turbine. However, this practice can be somewhat misleading due to the fact that the power available in the wind is a function of the cube of the wind speed, and the mean of a cubic function is not equivalent to the cube of the mean. Thus, comparing two turbines or controllers based on mean wind speed may not provide a truly accurate picture of the situation. Consider the following example, detailed in Table A-1 and plotted in Figure A-1. Air density ρ is set equal to 1.0 for simplicity. The wind speed is assumed to be constant over the ranges (0,3], (3,7], and (7,10] minutes with the mean values listed.

If a given turbine with a constant power coefficient C_p were run for each of these two 10-minute periods, it would capture more energy during the period called Case 1 simply because more energy was available. However, the two periods have identical mean wind speeds, so an analysis considering only mean wind speed might infer that the power coefficient C_p must have changed from one period to the next. Thus, this work plots power versus mean equivalent wind speed rather than mean wind speed. Mean equivalent wind speed is defined as

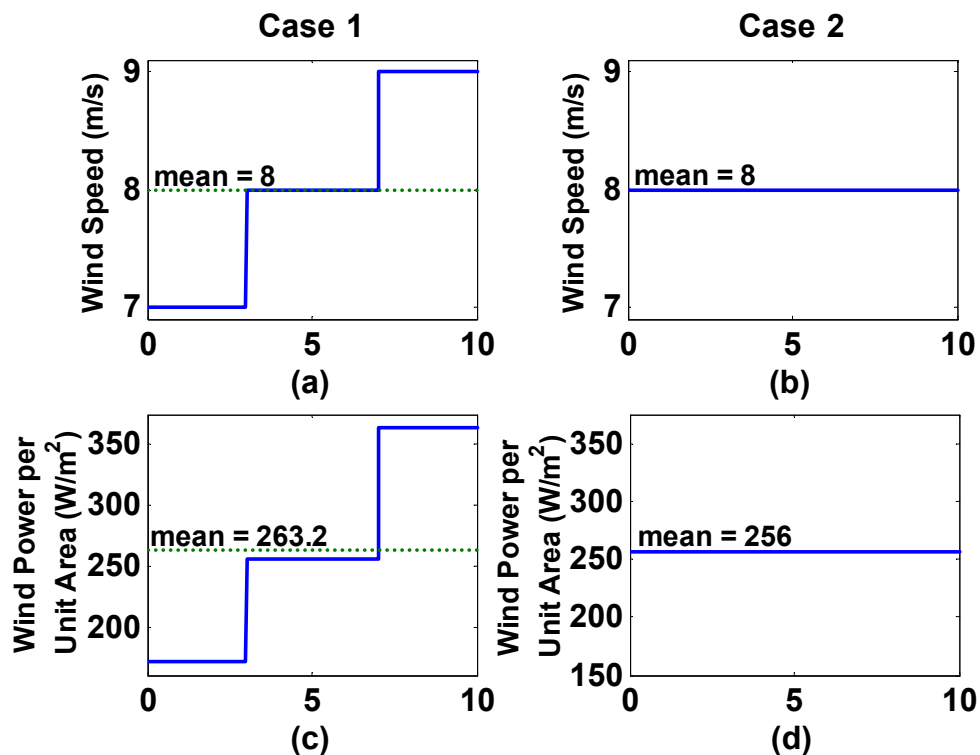


Figure A-1: Mean equivalent wind speed example

$$\bar{v}_{equiv} = \sqrt[3]{\frac{2}{A} \frac{1}{n} \sum_n P_{wind}(n)}. \quad (A.1)$$

In the example described above, there is about a 3% difference between the wind power per unit area from Case 1 to Case 2. Whether or not this is a significant difference depends on the application. In some cases, it may be perfectly reasonable to use mean wind speed rather than mean equivalent wind speed in an analysis. The CART's adaptive control low speed shaft power data for $0.40M^+ < M < 0.52M^+$ is plotted versus mean wind speed and mean equivalent wind speed in Figure A-2.

Since none of the 10-minute data sets collected while running the CART experienced constant winds, the mean equivalent wind speeds are expected to be slightly higher than the mean wind speeds, which is clearly the case in Figure A-2.

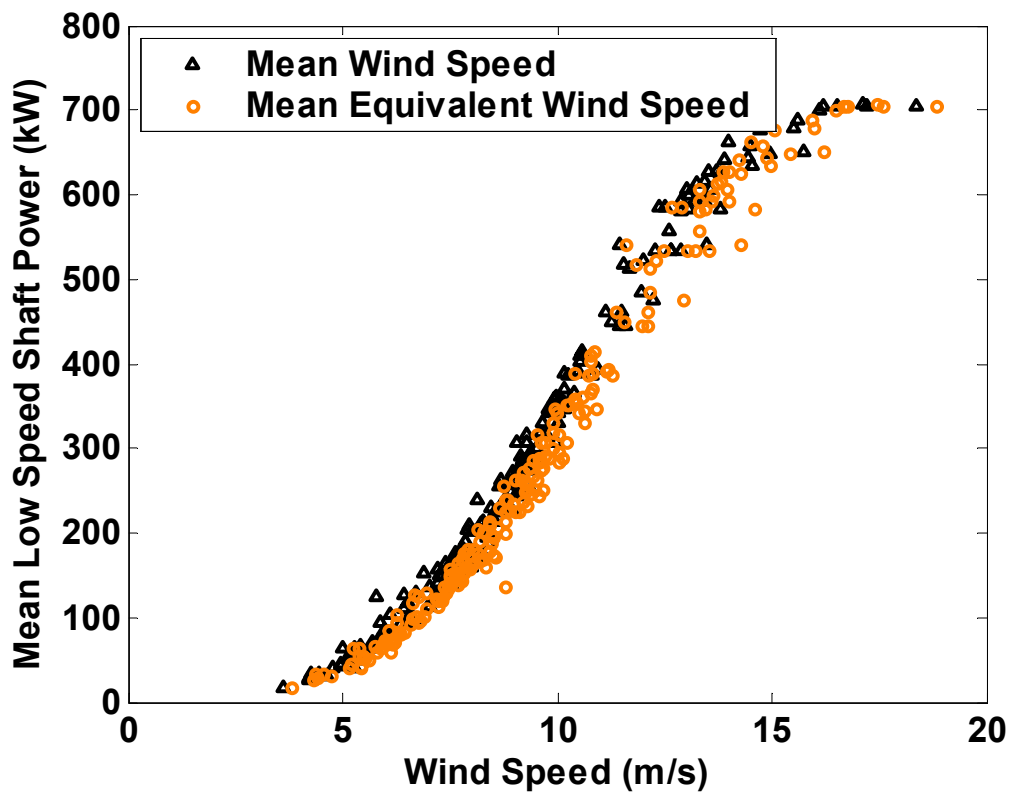


Figure A-2: Mean low speed shaft power versus mean wind speed and mean equivalent wind speed

Table A-1. Mean Equivalent Wind Speed and Mean Wind Speed Example

Time Range (minutes)	Case 1		Case 2	
	Mean Wind Speed (m/s)	Mean Power (2.5) per Unit Area: P_{wind}/A (kW/m ²)	Mean Wind Speed (m/s)	Mean Power (2.5) per Unit Area: P_{wind}/A (kW/m ²)
(0,3]	7	171.5	8	256.0
(3,7]	8	256.0	8	256.0
(7,10]	9	364.5	8	256.0
(0,10]	8	263.2	8	256.0

APPENDIX B: Stability Analysis for Gain Adaptation Law

The question of the convergence of $M \rightarrow M^*$ was addressed in Section 6.4 but not proven formally. This appendix provides a formal stability (though not convergence) analysis for a similar system in which the equation for ΔM (Equation 2.19) is replaced by (Equation 6.15).

Many stability results for time-varying nonlinear feedback systems were first proposed by Zames (1966a, 1966b). These were later expanded by Safonov (1980). The proof in this appendix is based on the sector stability criterion given by Theorem 2.2 in (Safonov 1980), hereafter referred to as Safonov's Theorem 2.2. This theorem applies to the two subsystem feedback system given in Figure B-1(a), where d_1 and d_2 are disturbance inputs to each subsystem. The disturbances in this proof can be considered to enter the subsystems additively, as shown in Figure B-1(b).

Safonov's Theorem 2.2 incorporates a functional F operating on the signals $x \in X_e$ and $y \in Y_e$, where X_e and Y_e are extended normed spaces. F defines an inner product on x and y as follows:

$$F(x, y, t) \equiv \langle F_{11}y + F_{12}x, F_{21}y + F_{22}x \rangle_t, \quad (\text{B.1})$$

where t denotes a truncation defined by the usual truncation operator; $F_{11}0 = F_{12}0 = F_{21}0 = F_{22}0 = 0$; $F_{11}, F_{21}: Y_e \rightarrow \mathcal{L}_e$; and $F_{12}, F_{22}: X_e \rightarrow Y_e$. In this definition, \mathcal{L}_e is an extended inner product space. For Safonov's Theorem 2.2, the sector of F is defined using the inner product (Equation B.1):

$$\text{sector}(F) \equiv \{(x, y) \in X_e \times Y_e \mid F(x, y, t) \leq 0 \forall t \in T\}. \quad (\text{B.2})$$




Figure B-1: Feedback systems considered by Safonov: (a) general case, and (b) specific case with disturbances entering additively

In Equation B.2, for the purposes of this analysis, T is considered to be the time $T = [0, \infty)$. Now, Safonov's Theorem 2.2 requires the following:

1. F is as given in Equation B.1 and F_{11} , F_{12} , F_{21} , and F_{22} have finite incremental gain
2. the mappings d_1 into $S_1(d_1)$ and d_2 into $S_2(d_2)$ are bounded about $S_1(0)$ and $S_2(0)$
3. $S_1^{-1}(0)$ is strictly inside $\text{sector}(F)$ and S_2 is outside $\text{sector}(F)$

Safonov's Theorem 2.2 concludes that, given (1. – 3.), the system given by Figure B-1(a) is closed-loop bounded. In simple terms, Safonov's requirement for stability of the closed loop system given in Figure B-1(a) is that the inverse of the graph of S_1 be strictly inside a given $\text{sector}(F)$ and the graph of S_2 be outside $\text{sector}(F)$. A two-dimensional interpretation of this requirement is provided in Figure B-2. In this example, the graphs of two fictitious operators S_1 and S_2 are shown along with a $\text{sector}(F)$. In this example, the graph of S_2 lies outside of $\text{sector}(F)$ and the graph of the inverse of S_1 lies strictly inside $\text{sector}(F)$. Thus, the closed-loop system in this example is stable.

Now Zames's and Safonov's results can be applied to the wind turbine's modified adaptive controller with gain adaptation law given by Equation 6.15. First, decompose the system given by Equations 2.18, 2.20, and 6.15 into the block diagram in Figure B-3.

The nonlinearity N_2 captures the relationship between \tilde{M} and fractional mean power P_{fav} .

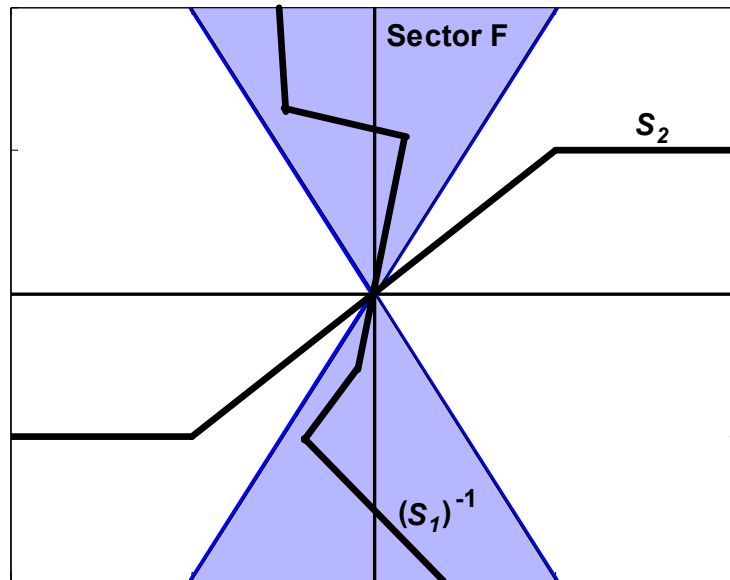


Figure B-2: Stability of example feedback system using Safonov's sector method

Denote

$$N_1(x) \equiv \gamma_{\Delta M} \operatorname{sgn}(x) |x| = \gamma_{\Delta M} x$$

$$N_3(x,y) \equiv xy$$

$$N_4(x) \equiv \operatorname{sgn}(x).$$

Let u_k be the input to S_{21} . The output of S_{21} is given by $y_k = \bar{N}_2(u_k)$. Let p^* denote the minimal slope of N_2 in the domain of interest and let p^* denote its maximal slope; without loss of generality, let $p^* = -p^*$. Effectively, the nonlinearity defined by S_{21} is the slope of a line segment connecting two points on the nonlinearity N_2 . Let the output of the linear block H_2 , which corresponds to \tilde{M} , be called v , and the output of the nonlinear block N_2 , or P_{favg} , be called w . The following relationships can be easily verified:

$$v_k = v_{k-1} - u_k, = - \sum_{i=0}^k u_i$$

$$w_k = N_2(v_k) = N_2(v_{k-1} - u_k),$$

$$y_k = w_k - w_{k-1},$$

where y , the output of S_{21} , corresponds to ΔP_{favg} , and u , the input to S_{21} , corresponds to ΔM . Thus,

$$y_k = \bar{N}_2(u_k) = N_2(v_{k-1} - u_k) - N_2(v_{k-1}). \quad (\text{B.3})$$




Figure B-3: Adaptive control feedback system

describes the relationship between u and y in the nonlinearity S_{21} . For any given v_{k-1} in the domain of interest, the relationship given by Equation B.3 is simply the slope of the line segment connecting the points $(v_{k-1}, N_2(v_{k-1}))$ and $(v_{k-1} - u_k, N_2(v_{k-1} - u_k))$, as shown in Figure B-4. While the magnitude and sign of this slope for a given input u_k vary depending on v_{k-1} , the slope nevertheless is bounded by the maximal and minimal slope of the nonlinearity N_2 as long as the points $(v_{k-1}, N_2(v_{k-1}))$ and $(v_{k-1} - u_k, N_2(v_{k-1} - u_k))$ are within the domain of interest. Note that when $u_k \in L_{\infty e}$, so is v_k since the gain of H_2 is bounded (finite) when operating on an extended normed space.

Now, let x be the input to the subsystem S_{22} , or equivalently the input to S_2 . Then, the relationship between x and u is given by

$$u_k = (x_k)(\text{sgn}(u_{k-1})),$$

and the relationship from x to y by

$$y_k = N_2\left(-x_k \text{sgn}(u_{k-1}) - \sum_{i=1}^{k-1} x_i \text{sgn}(u_{i-1})\right) - N_2\left(-\sum_{i=1}^{k-1} x_i \text{sgn}(u_{i-1})\right), \quad (\text{B.4})$$

which is qualitatively similar to Equation B.3 and can be pictured in the same manner with the appropriate replacements of signal names. Now, in order to use Safonov's Theorem 2.2 to prove stability of this system, the functional F must be selected to give the proper sector results. In this case, let X_e and Y_e be the extended normed space $L_{\infty e}$ and let the components of F be the gains




Figure B-4: Relationship between u and y within the subsystem S_{21}

$F_{11} = F_{21} = 1$, $F_{12} = p^*$, and $F_{22} = -p^*$. Thus, requirement (1.) is satisfied. Similarly, if the

disturbance inputs are additive as in Figure B-1(b), (2.) is satisfied automatically. Then, only requirement (3.) remains to be shown.

Define the inner product in the usual way for discrete time systems, i.e.,

$$\langle x, y \rangle_t = \sum_{i=1}^{t-1} x_i y_i .$$

Also, let $z_k = \sum_{i=1}^{k-1} x_i \operatorname{sgn}(u_{i-1})$. Then, the operator F gives the following inner product for this system:

$$\begin{aligned} \langle y + p^* x, y - p^* x \rangle_t &= \sum_{k=1}^{t-1} \left\{ \begin{aligned} &[N_2(-x_k \operatorname{sgn}(u_{k-1}) - z_k) - N_2(-z_k) + p^* x_k] \cdot \\ &[N_2(-x_k \operatorname{sgn}(u_{k-1}) - z_k) - N_2(-z_k) - p^* x_k] \end{aligned} \right\} . \\ &= \sum_{k=1}^{t-1} \left\{ [N_2(-x_k \operatorname{sgn}(u_{k-1}) - z_k) - N_2(-z_k)]^2 - (p^* x_k)^2 \right\} \end{aligned} \quad (\text{B.5})$$

Now, given the maximal and minimal slope of the nonlinearity N_2 , it is known that

$$\left| \frac{N_2(-x_k \operatorname{sgn}(u_{k-1}) - z_k) - N_2(-z_k)}{|x_k|} \right| < p^* ,$$

where $|x_k|$ can replace $-x_k \operatorname{sgn}(u_{k-1})$ for simplicity in the denominator because the slope of the line segment connecting the two points is bounded by the same value (p^*) regardless of whether the input x is positive or negative. Thus,

$$[N_2(-x_k \operatorname{sgn}(u_{k-1}) - z_k) - N_2(-z_k)]^2 - (x_k p^*)^2 < 0 ,$$

for all k , with the result that the sum given by Equation B.5 is less than or equal to zero for all t . This result proves, by definition, that the subsystem S_2 lies inside the sector $([-p^*, p^*])$. To use the notation in Safonov's Theorem 2.2, let $\operatorname{sector}(F)$ be the *complement* of $\operatorname{sector}([-p^*, p^*])$, which may be written $\operatorname{sector}([p^*, -p^*])$, as shown in Figure B-5. Thus, S_2 lies *outside* the $\operatorname{sector}(F)$.

Also, note that the subsystem S_1 is a positive memoryless linear operator with gain γ_{AM} ; the graph of S_1 is simply a line through the origin having slope equal to γ_{AM} . Thus, the graph of the inverse of S_1 is a line through the origin having slope equal to $1/\gamma_{AM}$. By requirement (3.) of Safonov's Theorem 2.2, the closed-loop system is stable if this line with slope $1/\gamma_{AM}$ lies *strictly inside* the $\operatorname{sector}(F)$, or equivalently the $\operatorname{sector}([p^*, -p^*])$. In other words, γ_{AM} must be chosen such that $\gamma_{AM} < 1/p^*$; if this inequality holds, then the system given in Figure B-3 is closed-loop bounded.




Figure B-5: Stability of the closed-loop system in Figure B-3 using Safonov's sector method

Remark 1: A stable system may still be obtained by replacing the term $|\Delta P(k)|$ with the more aggressive term $|\Delta P(k)|^{1/2}$. Note that, with such an update, the subsystem S_1 is no longer a memoryless linear operator with gain $\gamma_{\Delta M}$ but rather a memoryless monotonically increasing positive nonlinearity with graph bounded by the sector($[0, \infty)$). If it can be shown that the magnitude of the open loop gain is less than unity, stability of the system is established by Theorem 1 in (Zames 1966a); equivalently, the circle criterion and related multiplier theory techniques can also be used to establish the stability. Future work will include a rigorous proof, but, roughly speaking, the proof will show that the gain of S_2 is arbitrarily small around the operating point $(\Delta P, \Delta M) = (0, 0)$ since the gain of S_1 is arbitrarily large in its neighborhood. The gain of S_2 may be allowed to be relatively high in the region where $|\Delta P|$ is large since the gain of S_1 is very small for large $|\Delta P|$.

Remark 2: Effectively, the stability analysis problem was posed, for the $|\Delta P|$ update case, as the stability analysis of a feedback interconnection having a *linear time invariant* (LTI) system in the feedforward path and a sector nonlinearity in the feedback path. A family of stable controller gain adaptations may be obtained by using the standard positivity preserving multipliers in conjunction with the first order accumulator and difference operators. In this manner it is also possible to compensate for plant uncertainties effectively.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services and Communications Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE (DD-MM-YYYY) August 2004		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To) March 2002- March 2004	
4. TITLE AND SUBTITLE Adaptive Torque Control of Variable Speed Wind Turbines			5a. CONTRACT NUMBER DE-AC36-99-GO10337		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Kathryn E. Johnson			5d. PROJECT NUMBER NREL/TP-500-36265		
			5e. TASK NUMBER WER43302		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) National Renewable Energy Laboratory 1617 Cole Blvd. Golden, CO 80401-3393				8. PERFORMING ORGANIZATION REPORT NUMBER NREL/TP-500-36265	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S) NREL	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER	
12. DISTRIBUTION AVAILABILITY STATEMENT National Technical Information Service U.S. Department of Commerce 5285 Port Royal Road Springfield, VA 22161					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT (Maximum 200 Words) The primary focus of this work is a new adaptive controller that is designed to resemble the standard non-adaptive controller used by the wind industry for variable speed wind turbines below rated power. This adaptive controller uses a simple, highly intuitive gain adaptation law designed to seek out the optimal gain for maximizing the turbine's energy capture. It is designed to work even in real, time-varying winds.					
15. SUBJECT TERMS wind energy; adaptive controller; wind turbines					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code)

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18