Marine Hydrokinetic Turbine Array Performance and Wake Characteristics

Nickolas Stelzenmuller

a thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

University of Washington

2013

Committee

Alberto Aliseda, Chair

Brian Polagye

Phillip Malte

James Riley

Program Authorized to Offer Degree:

Mechanical Engineering

University of Washington

Abstract

Marine Hydrokinetic Turbine Array Performance and Wake Characteristics

Nickolas Stelzenmuller

Chair of the Supervisory Committee:
Associate Professor Alberto Aliseda
Mechanical Engineering Department

Three scale model horizontal axis MHK turbines (1:45) were tested in a flume at various array spacings. The scale rotors are based on the full-scale Department of Energy Reference Model 1, modified to reproduce the hydrodynamic performance of the full-scale turbine (20 m diameter) at the reduced experimental Reynolds number (10^5 vs 10^6, based on chord length). Flow incident on the turbines and in the wakes was characterized by Particle Image Velocimetry (PIV) and Acoustic Doppler Velocimetry (ADV) measurements. Tip speed ratio (TSR) similarity of the turbines is achieved by controlling the torque applied by magnetic particle brakes. Single turbines were characterized over a range of mean freestream velocities to explore the effect of Reynolds number on turbine performance. Measured turbine efficiencies of approximately 40% are similar to efficiencies predicted from full-scale simulations, indicating similar power extraction efficiency at scale. Wake characteristics and turbine efficiencies have been investigated at a range of TSR's, with the goal of determining array spacing and operating conditions that maximize overall array efficiency.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

The author wishes to thank all those that made this project possible. Thanks to Alberto Aliseda and Brian Polagye, for their constant guidance, the staff at the Bamfield Marine Science Centre, specifically Eric Clelland, for going above and beyond in his help with the experiments, Anthony Poggioli, Alejandro Fernandez Solana, and Rob Cavagnaro for spending weeks of their time in Canada helping with the experiments, Eamon McQuaide and Kevin Soderlund for their tireless help in the machine shop, Zoë Parsons for the generous use of her car for weeks in Canada, Jim Thompson for the loan of equipment, Bill Kuykendall for his help with the instrumentation, and to Teymour Javaherchi, Adam Niblick, James Joslin, Amirhossein Amini, Samantha Adamski, Danny Sale, Mike Barbour, Ryan Keedy, Chin Ng, and Colin Bateson for advice, discussion, and encouragement. This project was funded by the Department of Energy under grant DE-EE0003283

DEDICATION

to Elizabeth Rambus, D.N.P., who made it all possible

Chapter 1

INTRODUCTION

1.1 *What is Tidal Energy?*

1.1.1 Fundamentals of tidal energy.

The gravitational field due to the sun and the moon interact with the oceans and force changes in sea level. These changes in sea level represent potential energy extracted from the Sun-Earth-Moon gravitational system. This potential energy is converted into kinetic energy of the moving water and, eventually, dissipated into heat. In the presence of some particular combinations of coastal geometry and bathymetry, the moving water forms strong currents that concentrate the potential energy into very high kinetic energy density flows. When the sea level rises at the mouth of a long shallow inlet, the sea floods into the inlet, and when the sea level falls the sea ebbs out, forming strong alternating currents at the mouth of the basin. Likewise for long island chains separating large bodies of water, such as the Orkneys between the North Atlantic and the North Sea. The tidal currents created by these dynamics contain high kinetic energy fluxes per unit cross sectional area, which makes them ideal for converting into electricity this kinetic energy by way of Marine Hydrokinetic turbines placed across these tidal channels. These Marine Hydrokinetic (MHK) turbines, alternatively known also as tidal turbines or marine current turbines, act analogously to wind turbines and share many of their dynamics. This thesis focuses on the study of the characteristics of flow around these turbines, with particular attention to the performance of a small array and their near field effects on the flow.

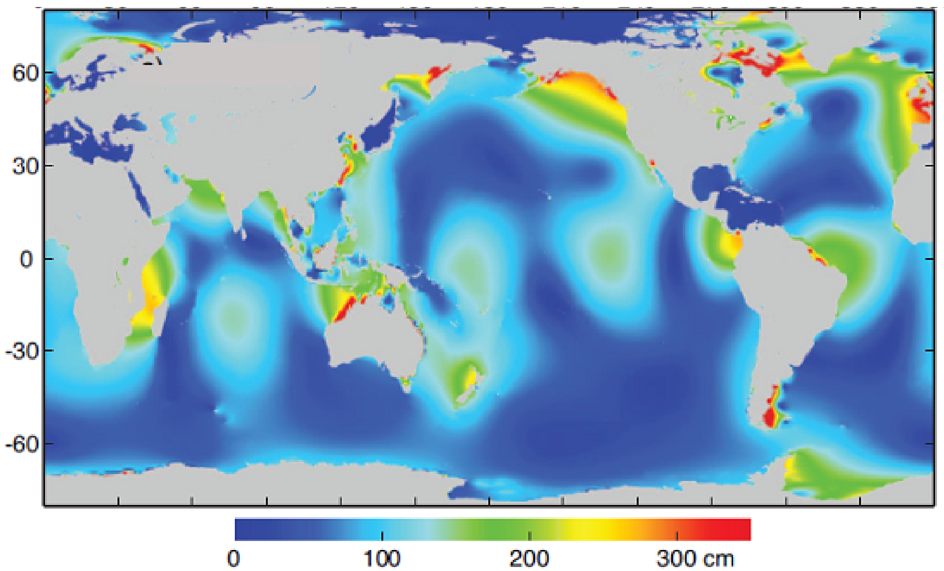

Figure 1.1: Average daily maximum tidal range.

Source: www.pacificstormsclimatology.org/images/glossary/tides.png

1.1.2 The role and scope of tidal energy

Climate change, increasing demand for energy, energy security concerns, and the diminishing reserves of conventional sources of energy have driven interest in developing sources of renewable energy. The solar and wind energy sectors have been growing rapidly, but both solar and wind energy suffer from unpredictability and intermittency—cloud cover and calm weather can halt energy conversion, sometimes unexpectedly. In contrast, the high degree of predictability of the positions of the sun and moon in relation to the earth allow tidal currents to be predicted with great accuracy. The predictability of tidal energy is a key advantage to utilities that need to provide a constant supply of power to the electrical grid.

Tidal currents occur in coastal areas world-wide, but there exist significant limitations to fully exploiting this energy resource. Among them:

1. Technical and economical limitations:

   Tidal currents contain kinetic power proportional to the cube of their velocity, therefore only relatively fast currents have enough power for viable commercial extraction. Mean kinetic energy densities less than 1 kW/m^2 are generally not considered to be commercially viable, although that threshold depends on the specific economic details of a project. Other key constraints to development include the proximity of the resource to electrical demand and particulates suspended in the flow such as ice, debris, and sediment that may damage the turbines.

2. Usage conflicts:

   Many area of coastal waterways are already in heavy use from shipping traffic, fishing, recreation, etc. Although tidal energy projects can coexist with other uses in many areas, usage conflicts will necessarily restrict development.

3. Environmental concerns:

   The interaction between tidal turbines and the marine ecosystem is largely unknown, and is the subject of many current research efforts. Concerns exist regarding the effect that tidal energy development will have on marine mammals and fish, as well as possible large-scale effects on estuarine dynamics [16]. Environmental regulators have adopted a precautionary approach to permitting tidal energy development projects, and further development will hinge on the results of environmental studies associated with early demonstration projects.

Global estimates of the tidal energy resource available for development are highly uncertain, ranging from 10s of GW up to 2 TW, and depend on how the above limitations are interpreted. Clearly, more research is required to assess the potential of this renewable energy source to satisfy future demand in an economically viable and environmentally responsible manner.

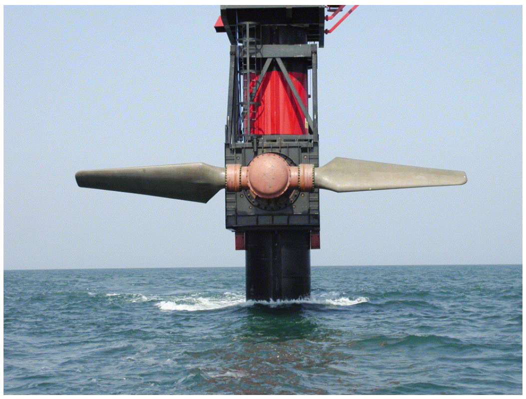

(a) Marine Current Turbine

(b) OpenHydro

(c) Verdant Power

(d) Ocean Renewable Power Company

Figure 1.2: Photographs of four MHK turbines currently being deployed in commercial tidal energy projects (photographs obtained from company websites).

1.2 Current State of Development of Tidal Energy

1.2.1 Tidal energy devices

The tidal energy industry is in an early stage of development. Few pilot-scale projects exist, and device development is still ongoing. As a consequence, there is a diversity of commercial and research tidal energy extraction devices being pursued.

Four tidal energy devices in commercial development currently are shown in Figure 1.2. The diversity of these devices reflects various approaches to the considerable engineering challenges faced by developers. For example, Figure 1.2 (a) shows the Siemens SeaGen turbine, which is designed to be easily moved vertically along the surface-piercing piling for maintenance. Figure 1.2 (b) shows the OpenHydro turbine, which has the electrical generator along the rim of the device, thereby avoiding central shaft and shaft seals. A more classical three-bladed design can be seen in Figure 1.2 (c). Although this Verdant Power prototype resembles the typical utility-scale wind turbine, the rotor is located downstream of the nacelle and the water-tight generator and gear box, which has to withstand water pressures up to 5 bar without significant maintenance for upwards of 2 years, make this a completely new engineering accomplishment. Figure 1.2 (d) highlights the Ocean Renewable Power Company cross flow turbine, which operates in a similar manner to a vertical axis wind turbine. While vertical axis turbines were not a commercial success in the wind arena, the advantages of insensitivity to flow direction in reversing tidal currents, potential ease of stacking for high coverage of tidal channel cross-sectional area, and the ability to share a common generator among multiple coaxial turbines, make this technology a promising candidate for analysis and development. Future maturation of the tidal energy industry may lead to some device design convergence, as happened in the wind industry.

Chapter 2

PROJECT BACKGROUND

This section provides some background on the project that is the subject of this thesis, outlines the project goals, and briefly reviews the published literature on tidal turbine laboratory experiments.

2.1 Project motivation

Tidal energy devices extract power from tidal currents to generate electricity. One of the key differences between wind energy and tidal current energy conversion is in the concentration of the resource. High winds are found along expansive areas, both onshore and offshore. Therefore, there is limited incentive to create dense arrays of wind turbines. In contrast, tidal energy resources are highly concentrated within narrow channels at the mouth of esturies or connecting large bodies of water. Only in these narrow passages is the kinetic energy density high enough to make energy conversion economically viable. This concentration of the resource makes understanding the minimum spacing between turbines critical to the success of this technology as a viable renewable energy source. Figure 2.1 illustrates the concentration of the tidal energy resource in Puget Sound, Washington. The power available in the flow is proportional to the cube of the velocity, so the available power in the areas of 2 m/s flow (red) is eight times that found in areas of 1 m/s flow (blue). A similar concentration in available power to that shown in Figure 2.1 is common in various tidal energy sites. This concentration implies that dense and optimized arrays of turbines

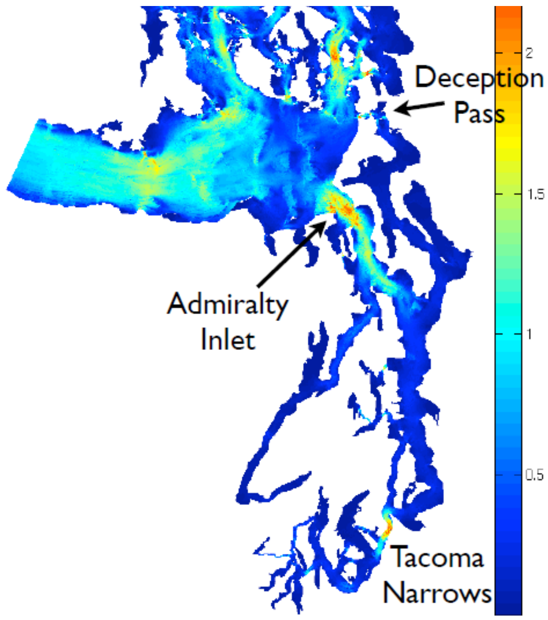

Figure 2.1: Tidal current speeds in the Puget Sound as predicted by numerical simulations from Kawase and Thyng[7]

will be required to create grid connected, utility-scale generation deployments for this natural resource.

Tidal turbine array optimization will require knowledge of turbine wakes, turbine wake interaction, the performance of turbines operating in the wakes of upstream turbines, and the effect of channel and free surface confinement on wake development and turbine performance. A complete understanding of phenomenon does not yet exist. The motivation of this thesis is to contribute to the fundamental understanding of these topics from the analysis of laboratory-scale measurements.

2.2 Thesis Objectives

The primary goal of this thesis is to improve the understanding of the physics of MHK turbine performance and the effect of turbine operation on the flow field through experimental study of scale-model turbines in a laboratory flume. We aim to extract general trends relating to turbine performance, wake development, and array optimization; trends that are not associated with a specific turbine design. Differences between experimental and full-scale conditions, including lower Reynolds number, flow confinement (blockage), and inflow properties, preclude direct extrapolations of these results to full-scale turbines and turbine arrays, but the experimental behavior of the turbine and the flow in the near field that has been observed provides insights into the hydrodynamics of MHK turbines. These measurements and analysis will be useful to understand full-scale deployments. For example, flow confinement causes the bypass flow around a turbine to accelerate (relative to the unconfined case), thereby reducing the wake recovery distance due to increased shear between the wake and the bypass flow. Consequently, direct use of the relationship between turbine spacing and performance for the experimental case may not be possible in the unconfined full-scale case. However, the trend of performance decay and recovery as a function of downstream turbine spacing is a result that will be useful for engineering of full-scale arrays.

The second goal of this thesis is to generate a large, high quality dataset that can be used to validate numerical models at laboratory-scale Reynolds numbers. This dataset includes turbine performance data, as well as data relating to wake development, for a single turbine and arrays of turbines. These numerical models, developed to match the scale, geometry, and boundary conditions of the experiments presented here [5] can be applied to the full-scale turbine prototypes and arrays. As Computational Fluid Dynamics (CFD) becomes a fast turn-around engineering tool, thanks to improvements in computational power and better numerical models, validated codes can be part of the predesign and prototyping process, quickly shifting through hundreds or thousands of design alternatives, and honing in on a few promising prototypes to test experimentally. High fidelity models can also be used to reduce the size of experimental testing matrices, further reducing the cost and turn-around time of the design iterations.

2.3 Previous work in laboratory-scale tidal turbine experiments

Marine hydrokinetic turbines are a relatively new field of research and this section provides an overview of the published literature and summary of some of the key results relevant to this project. This review is focused on experimental work, though the majority of the research in the field has been on numerical simulation of turbine performance and wake development. The recent review article by Ng et. al. [13] may be of interest for a broader look at recent research efforts.

2.3.1 Turbine performance characterization

Investigating turbine performance—the efficiency at which turbine extracts power from the flow—is one of the key motivations in laboratory-scale testing of tidal turbines. These results can be used to validate numerical models and extract general trends that can inform full-scale turbine design.

Bahaj et. al.[2] reported on the performance of a three-bladed scale-model turbine tested in both a cavitation tunnel and a towing tank. For a blockage ratio of 17% the peak performance was found to be ∼45%, which matches quite closely the same group's numerical predictions [1]. These experiments also explored the effect of rotor yaw and the blockage ratio (an effect confounded with the distance between the tip of the rotor and the free surface), both of which were found to have significant effect on turbine performance.

O'Doherty et. al. [14] performed tests on a similar scale-model turbine with a similar blockage ratio in a recirculating flume. Peak performance for this experiments agrees with the findings of Bahaj et. al., and is predicted reasonably well by the scale CFD models developed by this group. Large scatter is reported in the angular velocity measurements, which the authors speculate may be partially due to interference by the turbine support post.

Maganga et. al. [10] tested a similar turbine in a large recirculating flume with a very low blockage ratio (5%) and various inflow conditions. They found that a moderate shear in the velocity profile (8% difference from blade tip to blade tip) had a negligible effect on turbine performance, but very high inflow turbulence intensities (25%) significantly lowered turbine performance.

Common to the three experiments mentioned above is a surface piercing support post, and a electrical motor above the surface that connects to the rotor via a shaft and gear box to provide shaft loading.

2.3.2 Turbine wake characterization

Experimental characterization of a turbine wake at laboratory-scale—how it spreads, how quickly it recovers, how it responds to various turbine operating conditions, etc.—provides important information relevant to the design of turbine arrays and the optimization of array spacing. These results can also be used to validate numerical simulations, which, once experimentally validated, can provide a much less expensive design tool.

Mycek et. al. [11] tested two co-axially spaced turbines and measured their wakes with a laser dopper velocimeter point measurement. Similar to Maganga et. al. [10], they varied the inflow turbulence intensity between 5% and 25%. They found a much faster wake recovery for the high TI inflow condition. Specifically, the velocity deficit reached 10% (90% recovered) at only 6 rotor diameters downstream when the inflow TI=25%, but the velocity deficit was still $\sim$17% at 10 rotor diameters downstream for the inflow TI=5% case. They also noticed qualitative changes in the distributions of turbulence intensity between the upstream turbine wake and the downstream turbine wake.

Tedds et. al. [19] made a thorough survey of the near wake of a laboratory-scale turbine with an acoustic doppler velocimeter(ADV). From these data, the velocity deficit, turbulence intensity, turbulent kinetic energy, and Reynolds stresses are calculated. The results show a strong degree of anisotropy in the near wake of the turbine, and suggest the typical assumptions of isotropic turbulence may be problematic.

Stallard et. al. [18] describe turbine wake studies focusing on the wake of multiple turbines arrayed with various lateral spacings. The wakes interact as they expand downstream, and act to confine wake expansion. They report that wake interaction occurs when the turbines are spaced less that three rotor diameters laterally, and at that a merged group wake begins to form at a certain distance downstream that depends on this lateral spacing.

Recently, Chamarro et. al. [3] have published detailed velocity field measurements in the near wake of a tidal turbine. These results were obtained with a 3-D particle image velocimetry system, and represent some of the first turbine wake studies that have a high degree of spacial resolution. They report strong coherent tip vortices shedding from the three-bladed turbine that begin to interact with each other at one turbine diameter downstream of the rotor plane. This tip vortex interaction results in instability in the helical vortex structure, which they show to lose its structure by two diameters downstream. No hub vortex/tip vortex interaction was observed, although this interaction is clearly geometry dependent.

2.3.3 Turbine array optimization

Very little experimental work has been published on the topic of tidal turbine array optimization, except those described above. Jonsson et. al. [6] did some numerical work on the effect of turbulence intensity and length scale on wake development, which was validated by porous disk experiments. The authors claim that high turbulence intensity causes the wake to expand more quickly, and thus recover more quickly, and that quick wake recover also occurs if the integral length scale of the turbulence is greater than the wake width.

Myers et. al. [12] have also used porous disks to simulate turbine momentum extraction, with the goal of exploring array optimization. They report that zones of highly accelerated flow are observed when lateral spacing is small, and estimate that placing a turbine downstream in this accelerated flow could increase power extraction by 11% compared to free stream operation.

2.3.4 Laboratory-scale experimental considerations

Laboratory-scale experiments on MHK turbine performance and wake development are necessarily limited to relatively large test facilities by the competing concerns of Reynolds number effect on turbine performance and the blockage ratio effect. This balance between Reynolds number effect and blockage effect is explored in Whelan et. al. [20]. The authors show that the Reynolds number effect on performance of the foil sections that make up turbine blades is significant at typical laboratory length scales (based on chord length). This Reynolds number effect is due to the laminar separation bubble dynamics described in the review by Lissaman [9]. This review shows that the reattachment of a laminar separation bubble is largely a function of chord-based Reynolds number, and that the lack of reattachment has a severe impact of foil performance. As a consequence, there is a critical Reynolds number for a foil, above which the dependence of foil performance is relatively weak, but below which, foil performance dramatically decreases. Whelan et. al. [20] argued that, given this strong influence, the laboratory-scale rotor should be designed with foils that have good performance at low Reynolds numbers, instead of simply geometrically scaling down full-scale designs. This enables a closer matching of efficiency between full-scale and laboratory-scale at the same range of tip speed ratios, allowing for a better matching of wake development.

Chapter 3

DESIGN OF LABORATORY-SCALE MHK TURBINES

This chapter describes the process followed in this thesis for the design and manufacturing of the three laboratory-scale rotors, the mechanical design of the nacelles, and all the onboard instrumentation systems.

3.1 Rotor design

3.1.1 Department of Energy Reference Model 1

The Department of Energy has proposed a two-bladed horizontal axis MHK turbine geometry to be used as a common "open source" model for numerical and experimental research efforts, with the idea that a common geometry would allow various researchers to compare their results directly and accelerate the basic understanding and development process. The turbine rotor consists of two blades formed from NACA 63-424 foils, and has a diameter of 20 meters. This geometry has been used in previous numerical simulations [5, 8] at full-scale. The first series of experiments for this project used a geometrically-similar scaled down rotor based on the Department of Energy Reference Model 1 (DOE RM1), such that the laboratory-scale turbine had a diameter of 0.45 m (1:45 scaling). Results from initial testing of this turbine are shown in Figure 3.2. Experimental performance is much lower than predicted by blade-element momentum theory (BEMT).

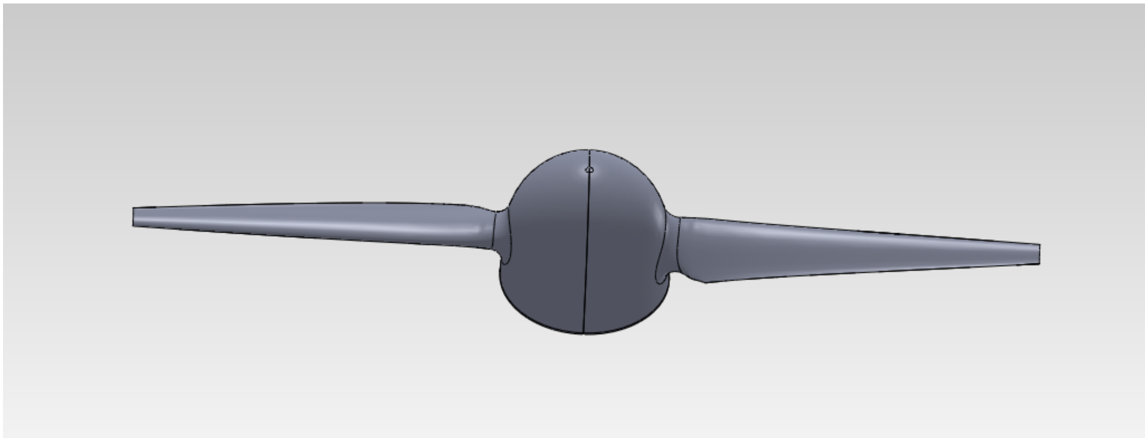

Figure 3.1: Solid model rendering of the DOE RM1 geometry

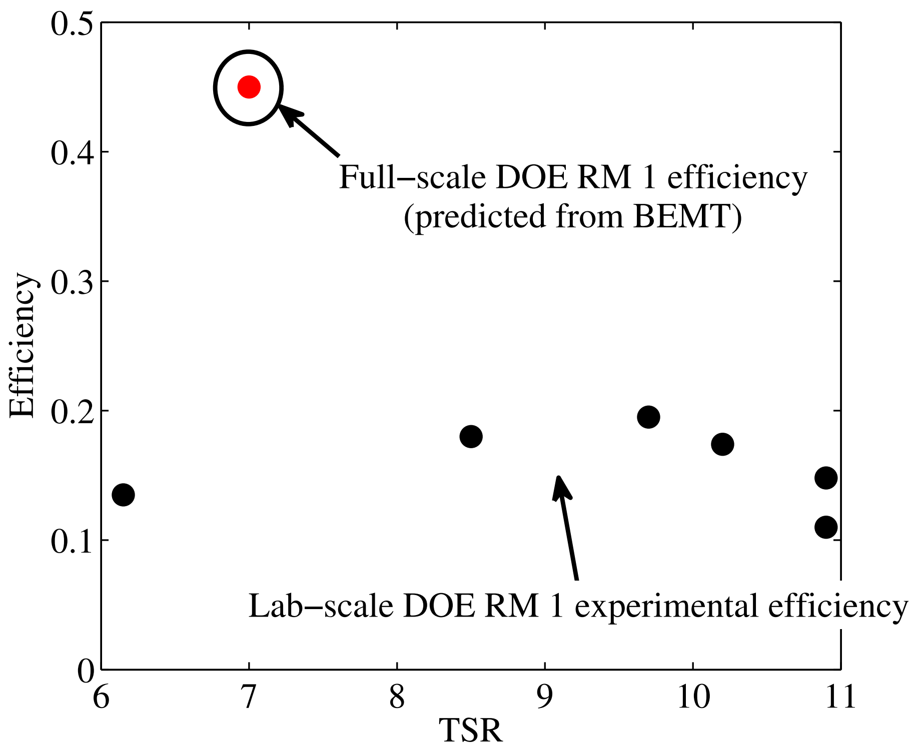

Figure 3.2: Experimental performance of the laboratory-scale DOE RM1 in 0.65 m/s flow, plotted with blade-element momentum theory prediction of the full-scale DOE RM1 performance at its optimum tips speed ratio.

A Reynold number effect was suspected, due to the relatively low chord-based Reynolds number of $Re \leq 80,000$ when the rotor was operating in flow speeds of 0.65 m/s, at Tip Speed Ratio values around 11. The laminar separation bubble dynamics described by Lissaman [9] result in a sharp decrease in performance below a critical Reynolds number. This critical Reynolds number varies by foil, and experimental foil performance data has not been published for the NACA 63-424 for the Reynolds number seen in testing ($Re \approx 70,000$), so there is some uncertainty as to what the critical Reynolds number for this foil is. Turbulence in the flow, 3-D flow effects, and flow curvature due to rotation may also play a role in the laminar separation bubble dynamics, and influence the critical Reynolds number for a particular foil.

Experimental constraints preclude increasing the diameter of the laboratory-scale rotor due to concerns with high blockage ratios. Based of the results of Whelan and Stallard [20], the decision was made to redesign the rotor in order to attempt to match, at the lab scale, the performance and wake properties of the full-scale DOE RM1, rather than merely scaling that geometry.

3.1.2 Redesigned rotor

A new rotor geometry was developed to avoid the low performance associated with the geometric scaling of the DOE RM1. Although the exact geometry of the DOE RM1 was not used in this new rotor, an attempt was made to match the performance and optimum tip speed ratio of the DOE RM1 as closely as possible. The two primary design goals associated with the redesign of the rotor were to increase the chord-based Reynolds number as much as possible in order to minimize any low-Reynolds number effects on performance, and to find a foil with good performance characteristics at low Reynolds numbers. The following steps were taken to achieve these design goals.

Greater braking capability

Initial tests of the DOE RM1 were performed with a free stream flow speed of 0.65 m/s. The initial tests could not be performed at the full speed speed of the flume— 1.1 m/s—due to the inability of the braking mechanism to apply sufficient torque to control the rotor. The nacelle instrumentation was redesigned in order to accommodate a magnetic particle brake that could apply a higher braking load on the turbine, the details of which are described in Appendix B.1. The ability to apply a high shaft load enables the turbines to be controlled at higher freestream velocities, which increases the operating Reynolds number.

Foil selection

Foil selection for the redesign rotor was made according to the following criteria:

1. The foil performance should be high, i. e. high lift coefficient, high ratio of lift to drag, at the relatively low Reynolds numbers used $Re \approx 100,000$.
2. The foil performance should have relatively low dependence on Reynolds number.
3. Experimental data should exist to verify conditions 1 and 2 near the operating Reynolds number.
4. The airfoil must be thick enough to form a structurally sound blade.

Condition 3 is important because potential flow solvers have high accuracy in predicting lift for low angles of attack at very large Reynolds numbers, but are unable to capture the low Reynolds number dynamics described by Lissaman [9], and should therefore not be used to predict airfoil performance near "transitional" Reynolds numbers. The range of Reynolds numbers seen in rotor operation ($70 \times 10^3 - 120 \times 10^3$) are relatively low for most foil applications, and consequently few results are available for foil performance at this range of Reynolds numbers. A number of foils have been designed for high performance at low Reynolds numbers, but these often have quite thin cross-sections, and were determined to have insufficient structural integrity for this application. The above considerations resulted in selection of the NACA 4415 foil. The NACA 4415 meets all of the above conditions, with high performance at low Reynolds numbers, relatively low dependence on Reynolds number, experimental data is available from wind tunnel testing, and the foil section is thick enough to support blade loading. This foil selection process was to some extent subjective, and other possible foil selections that meet the above criteria would have been possible.

Rotor design

Harp_opt, a rotor design code developed at the National Renewable Energy Laboratory [17], was used to maximize the blade chord while maintaining an optimum tip speed ratio equal to that of the DOE RM1. Harp_opt generates a range of initial rotor geometries, i. e. rotors with various chord and twist distributions, then uses the blade-element momentum code WT_perf [15] to rank their performance. WT_perf calculates local velocity and angle of attack for each section of the discretized blade using blade element moment theory, then calculates lift and drag forces on each section from coefficient of lift and coefficient of drag lookup tables. These lookup tables can be provided from potential flow codes such as Xfoil [4] or experimental wind tunnel data. Tip loss and hub loss is accounted for in WT_perf by empirical corrections. The results from WT_perf are used by Harp_opt to optimize the chord and twist distribution of the rotor. Harp_opt was used to create a turbine geometry using the NACA 4415 profile that maximizes performance and chord distribution, and has an optimum tip speed ratio equal to that of the DOE RM1 (optimum TSR≈7).

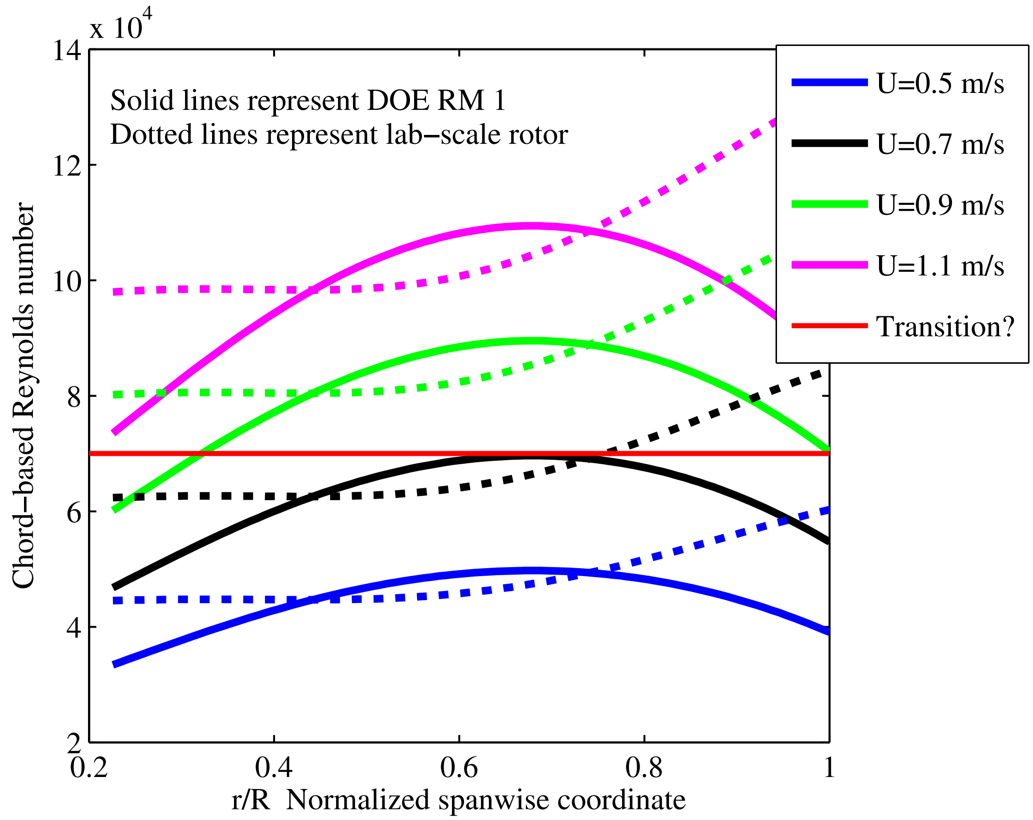

Figure 3.3: Chord-based Reynolds numbers for the 1:45 scale DOE RM1 and the redesigned turbine rotor for various free stream flow speeds at a TSR of 7.

Rotor manufacturing

The airfoil sections were lofted together in the a CAD software package according to the chord and twist distributions obtained from Harp_opt. This geometry was used to create a milling program, and that program was used by a three-axis CNC mill to machine the rotor from a solid billet of 6061 aluminum. This milling operation required flipping the stock over to access both sides of the rotor. A hole to accept the shaft was cut in the same operation to ensure a balanced rotor. The rotor was lightly sanded by hand, then anodized to prevent corrosion and maintain surface finish.

3.2 Nacelle mechanical design

3.2.1 Nacelle design constraints

The nacelle mechanical design was primarily driven by the constraints imposed by the instrumentation system. The primary design constraints were as follows.

Foundation fixed to flume bottom

The majority of the tidal turbine experiments found in the literature hold the turbines from a support post that pierces through the free surface. This offers several advantages, the primary one being that the system that applies the load on the shaft to the rotor can be placed above the surface. Holding the turbine from above the free surface also simplifies moving the turbines within the test facility. There are two disadvantages to holding the turbine from the surface. The first is the local deformation of the free surface by the post. One of the project goals was to measure the free surface deformation due to power extraction by the turbine. This would not be possible if the free surface was additionally deformed by the post. The second disadvantage is that the disturbance to the flow is greater when an immersed body pierces the free surface, and the drag and vibration imposed on the post is correspondingly greater, possibly contaminating the results. For these reasons, a design that supported the nacelle from the bottom of the flume was chosen for this project.

Waterproof nacelle body

A nacelle that is fixed to the bottom of the flume requires all of the instrumentation and the brake to be contained within the nacelle. Initial designs specified a flooded nacelle with individually waterproofed instruments. This proved possible with the torque sensor and encoder, but a satisfactory brake that could apply sufficient shaft torque, meet the other requirements necessary for the brake, and also be waterproofed was not commercially available. A nacelle that provided a waterproof cavity for the instrumentation and brake was therefore determined to be a design constraint.

Very low parasitic shaft drag not captured by the measurement system

The experiments described in this thesis require very sensitive and repeatable torque measurements. The maximum torque developed by these turbines as tested is approximately 2 N-m, so even parasitic shaft drag that is considered very low for most applications could be a significant fraction of the total measured torque. The sources of parasitic shaft drag in this system are the bearings and, given the requirements for a waterproof cavity outlined above, a waterproof shaft seal. If significant shaft drag cannot be avoided, as in the case of most shaft seals, then this shaft drag must be measurable by the torque sensor.

Streamlined low drag shape

The purpose of this experiment is to measure the performance and wake development of the rotor—as much as possible the nacelle should not affect these measurements. The body of the nacelle and the post should therefore present minimum blockage to the flow.

Manufacturability and accessibility

The nacelle needed to be manufacturable in-house, and with materials and processes that are compatible with the limited budget available. Due to the frequent changes and adjustments associated with a prototype, the body of the nacelle was required to be easily accessible. For example, the waterproofing could not be permanent glue, and the seals had to be reusable.

3.2.2 Nacelle design realization

This section describes some of the key decisions in the design of the nacelle.

Foundation system

The bottom of the flume used for these experiments is made of glass approximately 3 cm thick. Given the design constraint requiring the nacelle foundation to be fixed to the flume bottom, a method to attach the turbine post to this surface was needed. The large drag forces on the turbine during operation precluded a gravity foundation. Instead, an industrial vacuum pad was used as a base to which the support post was attached. The vacuum pad can be seen in the solid model of the turbine shown in Figure 3.4. The vacuum pad has a very high attachment force, does not damage the glass surface on which it is placed, and has a low profile for flow blockage. The vacuum pads were actuated by a vacuum pump in line with a water trap.

Nacelle waterproofing

The body of the nacelle was required to be waterproof as discussed above. There are two static seals and one dynamic shaft seal in the nacelle . The two static seals, a cable gland around the data cable and O-rings around the nacelle housing, are standard design components; these static seals are shown in Figure 3.4. The dynamic shaft seal is more complex, and required three subsystems to achieve the design goals. A solid model of the seal system is shown in Figure 3.5.

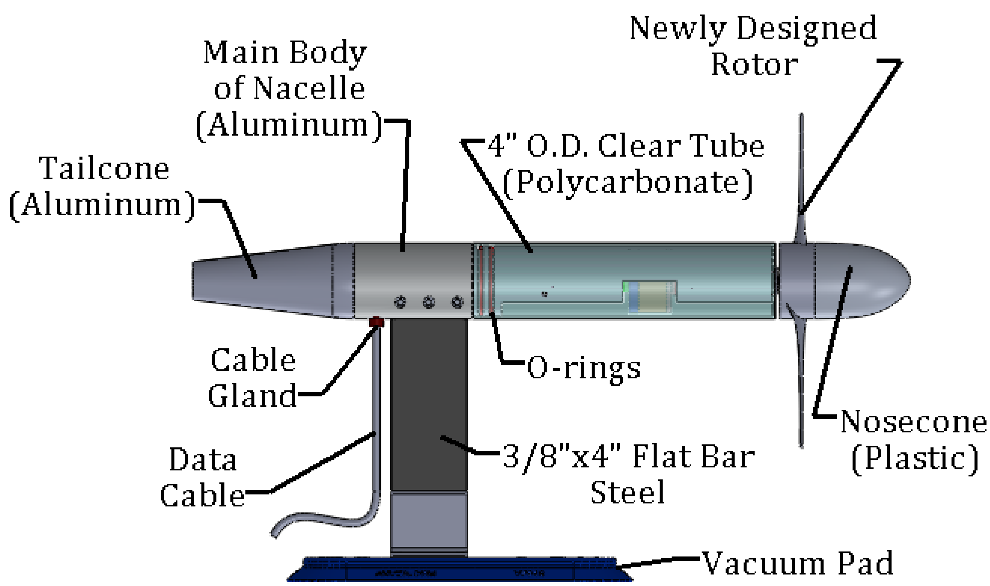

Figure 3.4: Solid model of the laboratory-scale turbine, as built.

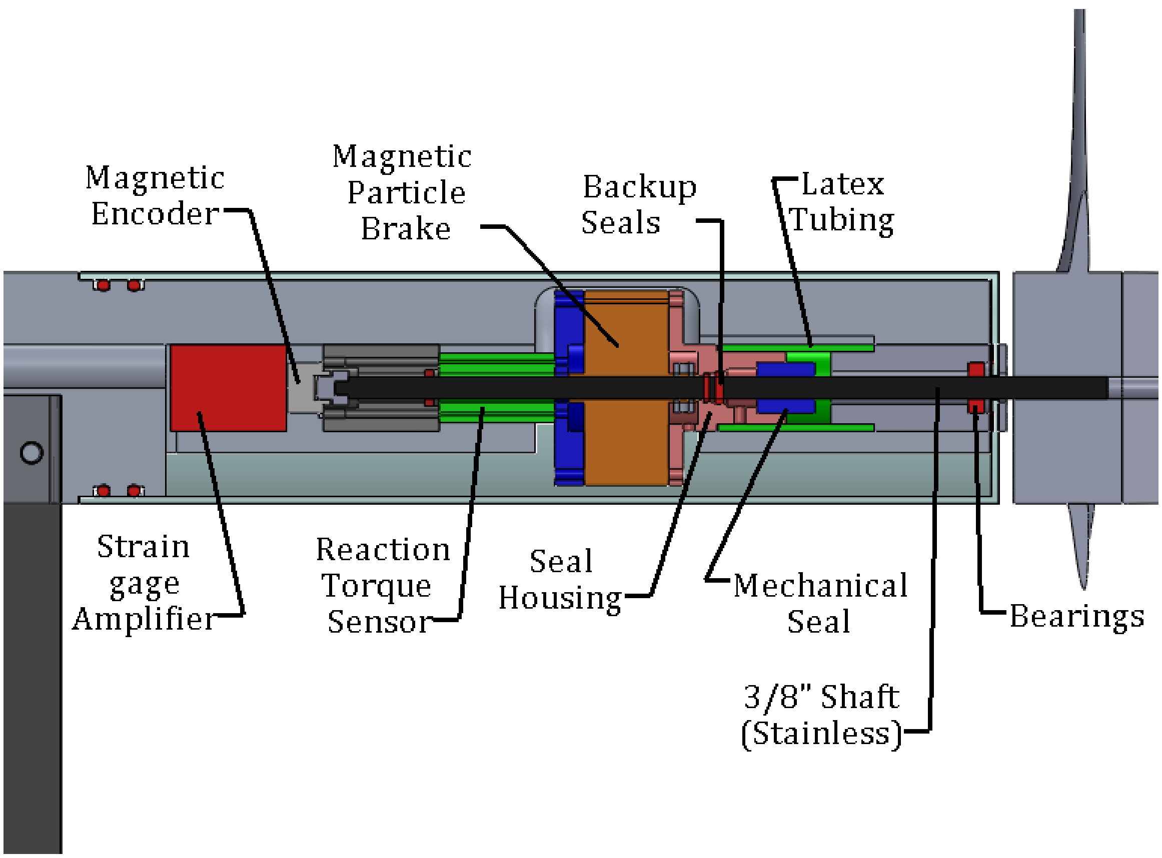

Figure 3.5: Cross-section of a solid model of the laboratory-scale turbine focusing on the instrument cavity, as built.

The three dynamic seal subsystems are a mechanical face seal, two PTFE lip seals, and a positive pressure system. The mechanical face seal and the two PTFE seals are shown in Figure 3.5. The key concept with regard to these shaft seals is that they are mounted on a seal housing attached directly to the magnetic particle brake. The seal housing is attached to the main body of the nacelle by a section of very flexible latex tubing. The shaft seals create significant shaft drag, but because the seals have a hard connection to the brake and a flexible connection to the main body of the nacelle all the shaft drag created by the seals is measured by the torque sensor. Almost all shaft seals have some leakage; to mitigate this, positive air pressure is pumped into the seal housing and a return line allows any water that has leaked through the seals to be forced out of the nacelle. The positive air pressure and return line are bundled with the data cable and routed out of the flume. The air pressure line is attached to an air compressor, and the air pressure is adjusted in operation such that the pressure forces a very small amount of air out of the front bearing.

Low parasitic shaft drag

The shaft drag associated with the shaft seals was measured by the torque cell as described above. This system results in a small amount of uncontrolled friction loading on the rotor, even when the brake is in free-spinning mode. As the brake is applied, the brake and the shaft seals create the measured shaft loading on the rotor. The only sources of unmeasurable parasitic torque in the system is the bearings and the hydrodynamic drag associated with the rotating shaft and hub. The bearings were chosen to be small, high precision, low drag bearings. The front bearing is flooded, but was regularly oiled to minimize drag. The hydrodynamic drag on the shaft and hub was judged to be minimal, although it was not fully characterized.

Streamlined nacelle and support structure

The design of the hub and nacelle was chosen to be as streamlined as possible to minimize the effect of the wake of the nacelle on the rotor testing. A streamlined shape was chosen for the hub, and the taper of the nacelle tail-cone was chosen to be 7° to prevent flow separation. Manufacturing limitations prevented the tail-cone extending to a sharp point, and was truncated as seen in Figure 3.4. The separation and recirculation associated with this truncated tail-cone can be appreciated in the velocity profiles measured by PIV, and presented in the Results chapter. Future experiments could include a tail-cone extension to prevent this separation. The support post was a simple steel flat bar, 5/8" x 4". This post was not streamlined, and the separated wake of the post does have some effect on the general flow field. The effect that this lack of a streamlined post has on the experiment is somewhat mitigated by the large distance from the rotor plane to the post, the relatively low frontal area of the post, and the flow field interrogation in the upper half of the water column. Future testing should consider a streamlined post as a natural improvement of the experiments described here.

Manufacturability and accessibility

The primary components of the nacelle require tight tolerances due to shaft alignment concerns. All components were designed to be manufacturable with 3-axis CNC machines, and the majority of components were cut out of aluminum for ease of manufacture. Accessibility was achieved by a removable acrylic sheath, which also allowed visual inspection for leakage during testing. This sheath was sealed by O-rings and silicon caulking, allowing easy removal.

3.3 Nacelle Instrumentation

In order to measure the efficiency and tip speed ratio of the turbines, the nacelles were instrumented with a torque sensor and a rotational position encoder. To apply a controlled shaft loading, a magnetic particle brake was used. The key design considerations for each of these instruments are described below, and a detailed description is given in Appendix B.

3.3.1 Torque measurement

A strain-gage-based reaction torque sensor was chosen to measure the hydrodynamic torque created by the rotor. The torque sensor was fixed to the main body of the nacelle and the brake, and thus measures all of the torque developed by the brake and shaft seals. The functional mechanism for measurements in this type of torque sensor consists of a four foil-backed strain gages connected in a Wheatstone bridge configuration and bonded to an aluminum body. Due to the long cable lengths necessary for this project ($\sim$12 m) a strain gage amplifier in the body of the nacelle was necessary to prevent signal-to-noise problems.

3.3.2 Rotational position measurement

The rotational position of the rotor is measured by a non-contact magnetic encoder. The encoder has a resolution of 256 pulses/revolution, and is able to resolve rotational direction. The position data generated by this system is used to derive rotational speed. More information on the derivation of rotational speed from position, which is non-trivial for this application due to quantization errors in the angular position, is given in Appendix A.

3.3.3 Applied torque

Braking torque is applied to the shaft by a magnetic particle brake. The magnetic particle brake is designed to generate a shaft torque proportional to input current and be independent of rotational speed. Other options for applying torque include a friction brake, a generator with a variable load, or a motor that drives the rotational speed independent of torque. These other options could offer advantages over the magnetic particle brake, and should be explored further in future work.

Chapter 4

EXPERIMENTAL DESIGN AND FLUME CHARACTERIZATION

4.1 Testing Facility

Laboratory-scale tidal turbine experiments are performed in testing facilities that can be divided into two types: towing tanks and flumes. The criteria for judging the suitability of a flume or towing tank are the following: cross-sectional area, flow speed, flow quality, length of test section, and optical access. As discussed in Section 2.3.4, the design of these experiments is necessarily a compromise between maximizing the Reynolds number (i. e. maximizing the rotor size and flow speed), and minimizing the blockage ratio (i. e. minimizing the rotor size and maximizing the cross-sectional area of the test section). For example, if the cross-section of the flume is small, the rotor diameter must be small to keep the blockage ratio at a reasonable level (where 5% is considered a low blockage ratio and 30% is considered high). But if the rotor diameter is small, a high freestream velocity is necessary to achieve chord-based Reynolds numbers high enough to prevent a large Reynolds number effect on performance. Only facilities with relatively high flow speeds and large cross-sectional areas are able to achieve this compromise.

The test facilities used for all the experiments described here is a recirculating flume located at the Bamfield Marine Science Centre (BMSC) in British Columbia, Canada. A photograph of this flume is shown in Figure 4.1.

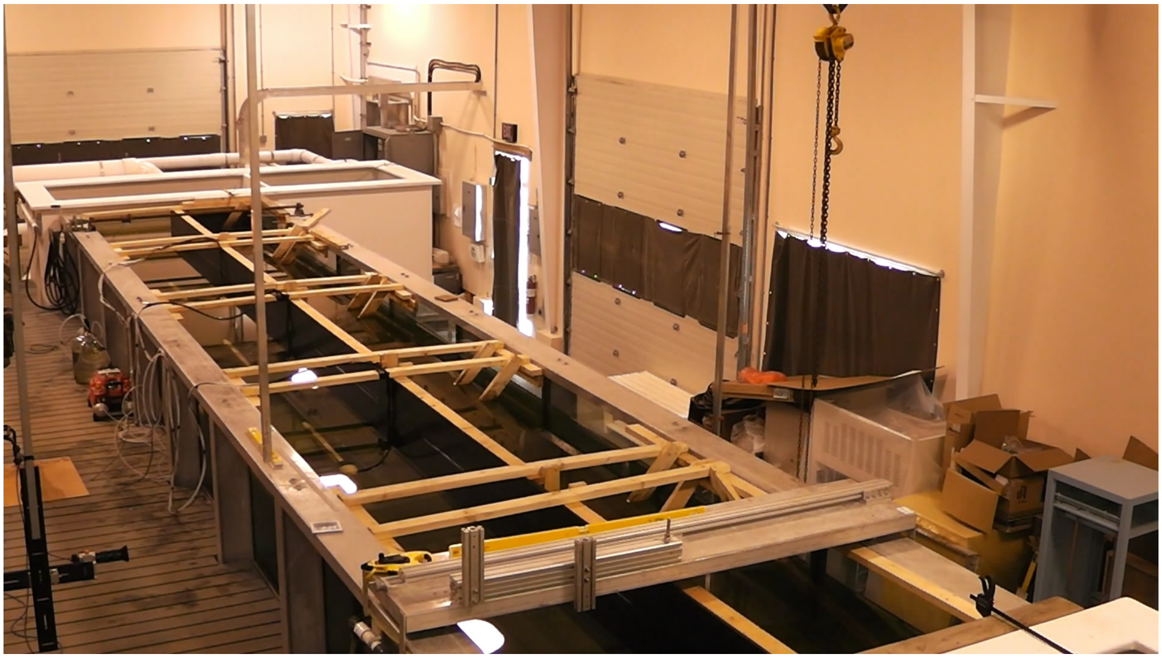

Figure 4.1: Photograph of the Bamfield Marine Science Centre flume

4.1.1 Flume dimensions and specifications

The BMSC flume has a width of 2 m, a depth of up to 1 m, and 12.3 m test section length with full optical access. The pumps that drive the flow are capable of a volumetric flow rate of approximately 1 m^3/s, which results in a freestream flow speed of 0.5 m/s. This flow speed was judged to be too slow to produce adequate Reynolds numbers, which are shown to be $\mathrm{Re} \sim 60,000$ at this flow speed in Figure 3.3. To increase the flow speed, and thus the Reynolds number, a partition was constructed in the flume, which can be seen in Figure 4.1. This partition halved the flume width from 2 m to 1 m, doubled the maximum flow speed, and doubled the blockage ratio.

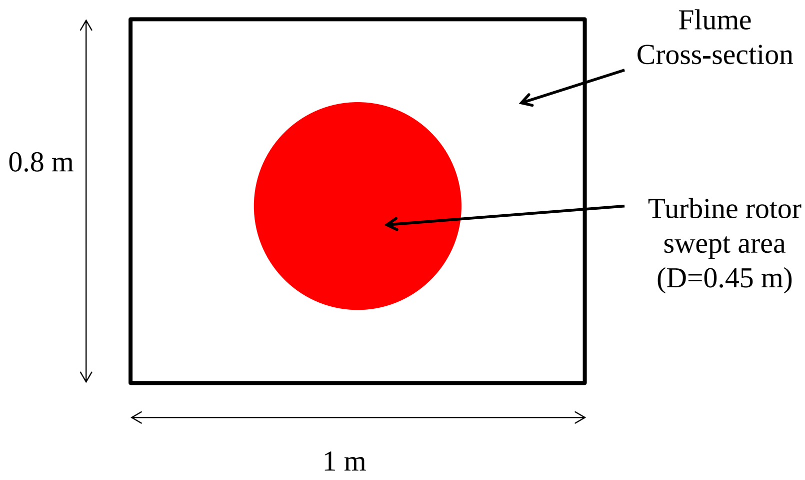

Figure 4.2: Blockage ratio schematic (to scale).

.

4.1.2 Flow confinement and blockage ratio

The blockage ratio for the modified flume was 20%, and is shown schematically in Figure 4.2.

This blockage ratio is high enough to affect the results in the following ways: the efficency of the turbines is increased, the TSR at which maximum efficiency is acheived ($TSR_{optimum}$) is increased, and the wake expansion is confined. Various blockage corrections have been proposed that use theoretical and empirical corrections to efficiency and $TSR_{optimum}$. These corrections attempt to predict the performance and $TSR_{optimum}$ of a turbine operating in an unconfined channel from the results obtained from a turbine tested in a confined channel. These corrections are not applied to the results presented here for three reasons. The first is that numerical simulations have been undertaken in a separate line of research at the University of Wahsington to model the experiment to scale and include the confinement, which enables a direct comparison with the experimental results. Second, the blockage corrections have not been settled or fully experimentally validated, and there is still large uncertainty about how to correctly apply them. The third reason is that blockage ratios have typically been developed for a single turbine, and it is unclear how the effect of confinement on wake expansion and recovery will affect the efficiency, $TSR_{optimum}$, and wake development of downstream turbines in an experimental turbine array.

4.1.3 Flow characteristics of the BMSC flume

Flumes designed for engineering fluid mechanics experiments typically have a large section upstream of of the test section to allow the flow to settle after it is discharged from the recirculation pumps. This flow then passes through a gentle constriction that is designed to accelerate the flow and introduce it to the test section in such a way that there is minimal turbulence, and little variation in flow speed across the cross-section. The BMSC flume lacks this settling section and constriction; the flow is forced around a corner, through a flow straightener, and into the test section. As a consequence, the flow has a relatively high turbulence intensity of approximately 5-10%, and both vertical and horizontal velocity shear in the mean flow. Details of the flow characterization of the flume can be found is Appendix D.

4.2 Outline of experiments

This section describes the experimental plan, turbine array configurations, and operating conditions. A detailed description of the testing procedure can be found in Appendix C.

4.2.1 Experimental plan

The goal of these tests was to collect torque and rotational position data while simultaneously measuring the flow with a particle image velocimetry (PIV) system. The

PIV system measures a 2-D 20 cm x 30 cm rectangular interrogation window, so a thorough characterization of the wake of the turbine entails sampling several of these areas at various streamwise stations along the wake. The turbines can be operated at various tip speed ratios by adjusting the brake on each turbine.

The primary variables in the experimental matrix are position of the PIV interrogation window and the TSR of the turbine. When testing multiple turbines it is possible to select a unique TSR for each turbine. This large experimental matrix was reduced to an experimental plan by making the following choices:

1. PIV interrogation windows were measured only along the centerline of the turbine, parallel to the flow, and vertically from a height of 40 cm from the bottom (mid-water column and hub height) to 70 cm from the bottom (10 cm below the free surface). Interrogation windows were taken at the following streamwise locations relative to the rotor plane: 2 rotor diameters (D) upstream, and 2, 3, 5, and 7 D downstream. These measurements result in a 2-D vertical slice of the wake sampled every two rotor diameters. Each interrogation window was measured for a duration of one minute, during which time the turbine brakes were not adjusted.
2. Various combinations of TSR were selected when testing multiple turbine arrays. For example, in a three-turbine array the upstream turbine could be operating at TSR=5, the midstream turbine operating at TSR=7, and the downstream turbine operating at TSR=9. Operating the upstream turbine at TSR=5 will produce a different wake that if it is operating at TSR=10, and because the wake of the upstream turbine provides the incident flow on the turbines downstream their performance may be affected. Typically eight combinations of TSR were chosen for the experimental plan.

To illustrate, a three-turbine co-axial array spaced five rotor diameters apart can be taken as an example. The PIV system is set up to measure an interrogation window two rotor diameters upstream of the upstream rotor, along the rotor centerline and vertically from the turbine axis of rotation to 10 cm from the free surface. The upstream turbine operates at TSR=5, the midstream turbine at TSR=5, and the downstream turbine at TSR=5. The PIV system takes data for one minute, during which time data from the torque sensors and rotational encoders in each turbine are being recorded. After one minute, the TSR of the upstream turbine is changed from 5 to 6, and the test is repeated. This procedure repeats until all eight combinations of TSR are tested, then the PIV system is moved to three rotor diameters downstream of the upstream turbine, and the same sweep through the combinations of TSR is performed. All eight combinations of TSR are tested for each location of the PIV system, until the wake measurements have been performed on all three turbines.

4.2.2 Turbine arrays

Single turbines, arrays of two turbines, and arrays of three turbines were tested, and the arrangement of these arrays are described below.

Single turbine

A single turbine was tested over a full range of TSR (from the stalled operating condition to the no-load operating condition), and for each operating condition the flow upstream and downstream of the rotor plane was measured with PIV. Each of the turbines were tested individually to ensure that all had similar performance curves. Finally, a single turbine was characterized at various freestream velocities in order to determine the effect of Reynolds number on performance. All of these tests were performed with the turbine at mid-channel, and the flume had a water depth of 0.8

m. The freestream velocity was approximately 1.1 m/s (as measured by PIV at 2 D upstream of the rotor plane and at hub height), unless stated otherwise.

Two turbine arrays

Two turbines were tested at various streamwise spacings. The two turbines were arranged co-axially, i. e. they were both placed in the centerline of the flume such that the rotors shared an axis of rotation. Four streamwise spacings— 5, 8, 11, and 14 rotor diameters—of the two turbines were tested. As with the single turbine tests, all tests had operating conditions of 1.1 m/s flowspeed and 0.8 m water depth.

Three turbine arrays

Four arrangements of three turbine arrays were tested, two with co-axial spacings and two with lateral spacings. The two co-axially spaced arrays differed only in

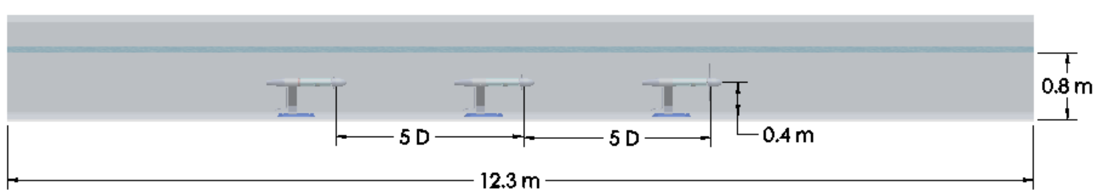

Figure 4.3: Solid model rendering of a three turbine co-axially spaced turbine array, separated by 5 rotor diameters

.

the distance that the turbines were separated. Both were arranged such that all of the rotor axes were on the centerline of the flume, and both had equal spacing between the three turbines. Figure 4.3 shows a side view of this configuration for a streamwise spacing of 5 rotor diameters. Three-turbine co-axial arrays with 5 and 7 rotor diameters spacing between turbines were tested.

Similarly two three-turbine arrays with lateral offsets were tested; a front view of one of these arrays is shown in Figure 4.4. All of the four three-turbine arrays have

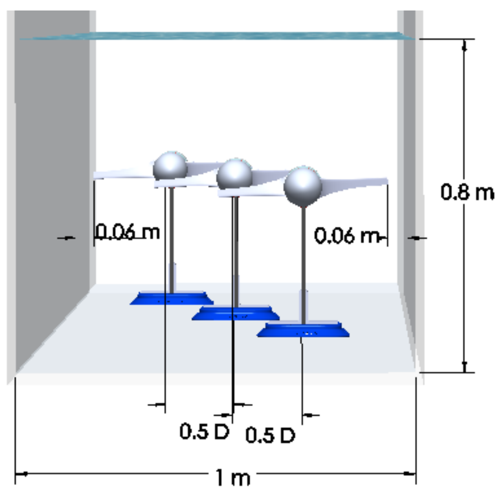

Figure 4.4: Front view of a solid model rendering of a three turbine array with 0.5 rotor diameter lateral offset.

a common position for the midstream turbine. For the laterally offset arrays, the upstream turbine was shifted 0.25 rotor diameters the right of the centerline of the flume, and the downstream turbine was shifted 0.25 rotor diameters to the left (from the perspective of looking downstream). These laterally offset configurations have similar transverse spacing to the other three-turbine arrays: one with a separation distance of 5 rotor diameters and the other with a separation of 7 rotor diameters. All three-turbine arrays were tested with the flume operating at 1.1 m/s and a water depth of 0.8 m.

4.3 Turbine data acquisition and control

During testing the signals from the three torque sensors and the three rotational encoders are sampled by the data acquisition system at a rate of 1000 Hz. These data are then minimally processed and streamed to disk. The rotational position data is transformed to rotational speed in real-time for the purpose of monitoring the tip speed ratio of each turbine during testing. This information is used to manually control the power supplies for the particle brakes for each turbine and therefore set the desired tip speed ratio of each turbine. There is variability in the tip speed ratios, especially for the downstream turbines in the turbine arrays and at low TSR. Design considerations regarding the turbine instrumentation can be found in section 3.3, and details regarding the instrumentation, braking system, and data acquisition system can be found in Appendix B.

4.4 Particle image velocimetry

The particle image velocimetry (PIV) system at the BMSC flume was used to investigate the flow field around the turbines. The PIV system consisted of a LaVision double exposure camera capable of 5 image pairs/sec, a control/timing server, and a YAG laser and optics. The laser was positioned beneath the flume and generated a vertical laser plane that was parallel to the direction of flow. This laser plane provided lighting for the camera which was arranged perpendicular to the flow. This PIV system configuration enabled interrogation windows 30 cm high and 20 cm in the streamwise direction. Although the PIV system was separate from the turbine data acquisition system the system clocks were synchronized, enabling time-series comparisons of the data. More information on the PIV system and the image processing of the PIV data can be found in Appendix B.5.

Chapter 5

PERFORMANCE AND WAKE CHARACTERIZATION OF A SINGLE TURBINE AND TURBINE ARRAYS

5.1 Definitions and conventions in the presentation of results

The primary results presented here are the coefficients of performance as a function of tip speed ratio and the measurements derived from the particle image velocimetry, including mean streamwise velocity profiles along the vertical direction and turbulence intensity profiles. The definitions and experimental derivations for these quantities are as follows:

5.1.1 Coefficient of performance and tip speed ratio

Coefficient of performance (C_p) is defined as the ratio between power extracted from the flow by the turbine and the power in the unperturbed flow through an area equal to that of the swept area of the rotor. Specifically, the definition of coefficient of performance is:

$$C_p \equiv \frac{T\omega}{\frac{1}{2}\rho\pi r^2 U^3} \tag{5.1}$$

and the definition of tip speed ratio is:

$$TSR \equiv \frac{\omega r}{U} \tag{5.2}$$

Where:

T = Torque, measured directly from the toque sensor, sampled at 1000 Hz

ω = Rotational speed, derived from encoder rotational position data, see Appendix A for details.

ρ = Fluid density, a nominal value of 1000 kg/m^3 has been used for all the results presented here.

r = Rotor radius, 0.225 m

U= Freestream velocity. The value of U is taken from the PIV data taken at two rotor diameters upstream of the most upstream rotor plane (if in an array). Not all tests were performed with simultaneous PIV measurements at the 2D upstream location, so a mean value from all of the 2D upstream centerline PIV measurements for each test was used for all calculations of C_p and TSR for that test. Where a characteristic value of the freestream velocity is given in the text, however, the channel centerline velocity was used as a simple representation of the freestream velocity.

All of the coefficients of performance and tip speed ratios reported here are calculated with the freestream velocity defined above. This decision was made in order to clarify the results but constitutes an abuse of notation, as the metric C_p is properly regarded as a ratio between the instantaneous power the rotor extracts from the flow and the instantaneous power in the flow available to the rotor. Since the downstream turbines in a turbine array are operating in the wakes of the upstream turbines, the power in the available flow is not equal to the power in the flow at freestream velocities. However, the interaction between the developing wake of the upstream turbine and the induction zone directly upstream of the downstream turbine is complex, and it is not clear where the flow speed should be measured when calculating the power available to the downstream turbine. A local freestream velocity for each turbine could be used to renormalize the Cp results, but insight into the overall power extraction of the turbine array would not be gained with such a metric. Therefore, the coefficients of performance and tip speed ratios of the downstream turbines reported here should be regarded as nominal values, and the values of these quantities based on local incoming velocity for each turbine would be, in general, higher than the nominal values. The use of this convention eliminates ambiguities with regard to where the flow speed associated with C_p and TSR of the downstream turbines is measured, and provides a more useful metric when evaluating the overall efficiency of turbine arrays. To gain insight into how the placement in the array influences the actual efficiency of a turbine, we can define an intrinsic efficiency as the ratio of power produced divided by kinetic energy flux at that turbine, 2 diameters upstream of its rotor plane, and the array component of the efficiency. This second component is defined as the ratio of actual kinetic energy flux at the location of the turbine inside the array (2D upstream of its rotor disk), over the kinetic energy flux in the undisturbed free stream (2D upstream of the most upstream turbines rotor disk). That way, the efficiency of the turbine, in its usual definition, as used here is the product of both efficiencies: its intrinsic times the array contribution (which is always smaller than 1).

5.1.2 Particle image velocimetry results

The details of the PIV system are given in section 4.4. The conventions used for the measurements presented below are as follows. Vertical profiles are derived from an average in time, over all the velocity fields obtained from PIV of images under the same conditions, and in the streamwise direction, over a row of the vector field at each vertical position. Unless explicitly stated otherwise, the PIV results were all collected from vertical planes at the channel centerline, oriented parallel to the flow, and extending 30 cm vertically and 20 cm in the streamwise direction. All measurements were taken over one minute intervals. The plots that display these results in general indicate the position of the rotor tip and the free surface. The cylinder defined by the circle traced by the rotor tips and extending downstream in the streamwise direction will be referred to as the rotor cylinder in the discussion.

5.2 Singe turbine characterization

A single turbine was tested in the flume to generate the following results:

- A performance curve at typical flume operating conditions (1.1 m/s).
- A set of performance curves at a range of free-stream velocities to explore the Reynolds number dependence of turbine performance.
- The wake properties of the turbine operating at a range of TSR at various streamwise locations.
- The performance of each of the three turbines tested separately, to confirm the similarity of the experimental models.

5.2.1 Performance curve for a single turbine

Figure 5.1 shows the performance curve of a single turbine, compared against the performance curve obtained from the rotor design code Harp_opt [17]. A description of the use of this design code in the rotor design can be found in Chapter 3. These results show that Harp_opt overpredicts peak performance by approximately 17%. Figure 5.1 also shows that Harp_opt predicts a optimum TSR of approximately 7.8, compared to the experimental optimum TSR of approximately 6.8. The slope of the Harp_opt performance curve from TSR 5.5-7 is relatively pronounced, but the experimental performance curve in this region is fairly flat. The overprediction of performance from the Harp_opt results is expected due to the idealized nature of the blade-element-momentum theory that is used by this code. More importantly, the peak performance of the experimental turbine of ~40% is close to what we expect

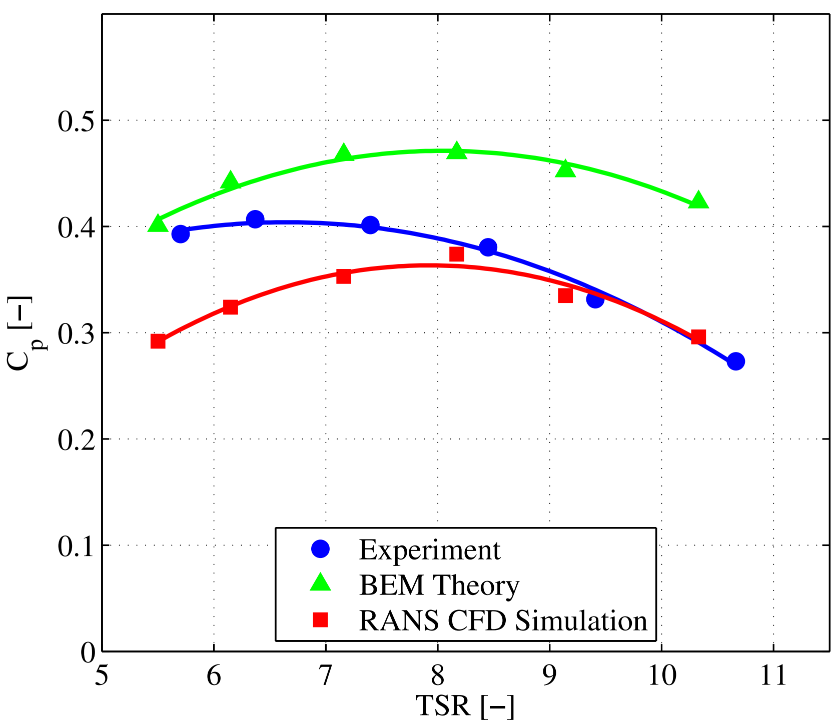

Figure 5.1: Performance curves from experimental data and from blade-element-momentum code predictions.

from a full-scale turbine. This similarity in the physics of the power extraction allows the trends uncovered by the experiments reported in this thesis, in performance and wake development, to be directly relevant to full-scale applications. In cases where the laboratory experiments do not have this physical similarity, as is the case in experiments conducted with a geometrically similar rotor at small scale, where the performance was in the single digits due to the poor performance of the 63-xxx airfoils at low Reynolds numbers ($\approx 50,000$), the wake development and the wake-turbine interactions is controlled by different interactions and the results would not be representative of full-scale behaviour.

Harp_opt predicts a lower performance at TSR of 5.5 than at a TSR of 7 due to some portion of the blade near the root being in stall conditions, i.e. the high angle of attack at blade sections of the root are higher than the angles of attack that cause stall in 2D airfoil tests. The portion of the performance curve between TSR 5.5-7 that is flatter than predicted by Harp_opt may be explained by a stall delay effect due to the high amplitude fluctuations in rotational speed at low tip speed ratios. Figure 5.2 shows the rotational speed variability of the rotor at low TSR and high TSR. The rapidly changing angles of attack along the blade associated with a highly variable rotational speed, in addition to the strongly 3-D nature of the flow around these blade sections, may be delaying stall on blade sections that would be in stall with steady 2-D flow. This effect would increase performance at low TSR, and thus flatten the performance curve at low TSR.

5.2.2 Performance curves for a single turbine at low Reynolds numbers

Figure 5.3 shows a family of performance curves measured at freestream velocities from 0.52–0.9 m/s. These results show that the performance curves for speeds down to 0.71 m/s collapse onto a single curve, within experimental error, but a progressive deterioration in performance is evident as the freestream velocity decreases below

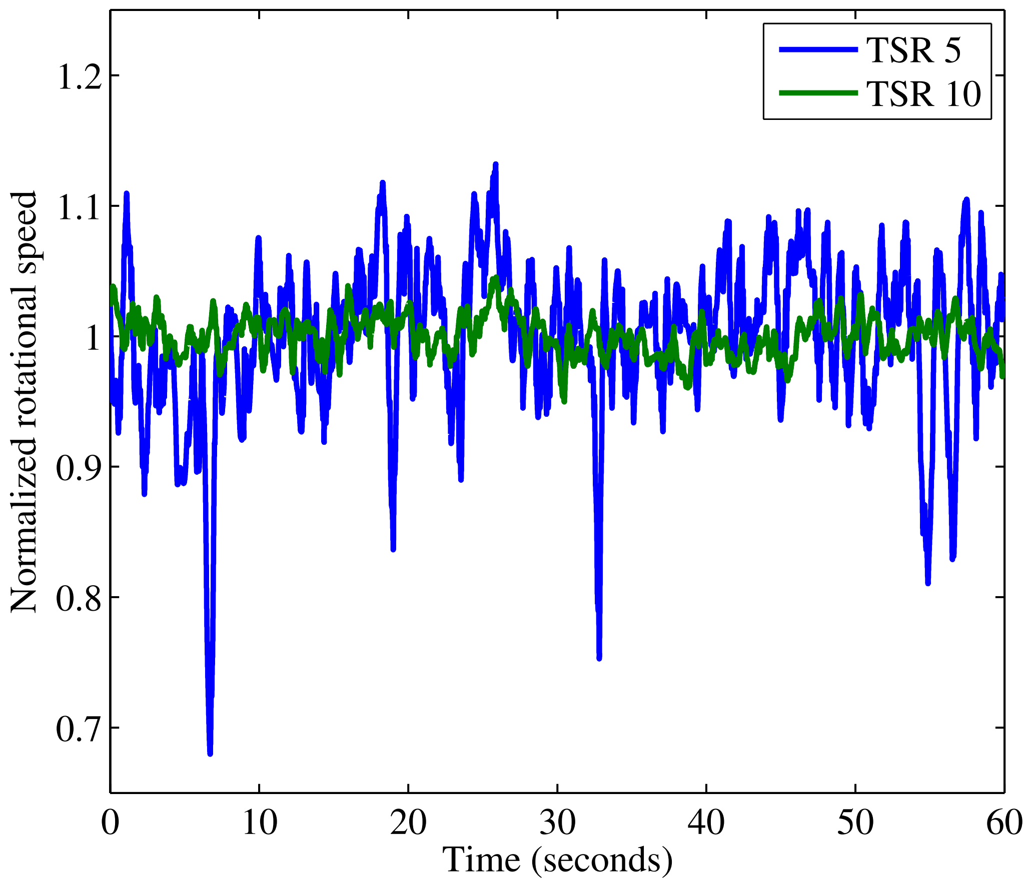

Figure 5.2: Rotational speed normalized by mean rotational speed over one minute for a high TSR operating condition and a low TSR operating condition.

that velocity value which corresponds to the chord-based Reynolds number of the blade dipping below 70,000 at significant portions of the span. The effect of Reynolds number was discussed in section 3.1, and the conclusion was that foil performance generally has a weak dependence on Reynolds number except near some transition Reynolds number that is specific to each particular foil and inflow conditions. Near this transition Reynolds number, the foil performance has a large dependence on Reynolds number, due to the reattachment (or lack thereof) of the laminar separation bubble. This effect is hinted at in Figure 5.3, as no change in performance is seen for

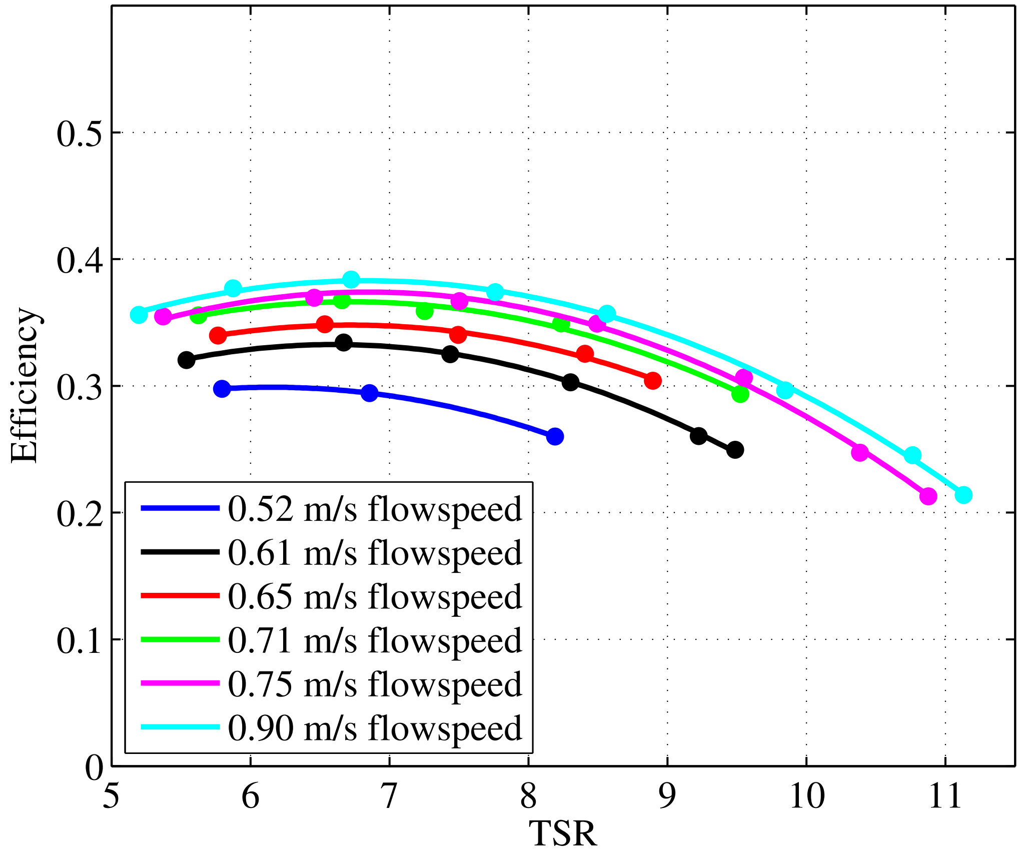

Figure 5.3: Performance curves from experimental data at various flowspeeds

freestream velocities between 0.9 m/s and 0.71 m/s, but a large drop in performance is seen between turbine performance for free-stream velocities lower than 0.71 m/s. This indicates that the Reynolds number associated with 0.71 m/s freestream velocity is a transition point below which airfoil hydrodynamics shows a strong dependence on Reynolds number and above which it does not.

5.2.3 Wake characteristics

Particle image velocimetry (PIV) was used to investigate the flow upstream and downstream of the turbines. A full description of the PIV system is given in section 4.4, and a brief outline of the conventions for the presentation of these results is given in section 5.1.2. The mean velocity profiles upstream and in the wake of a single turbine operating at a tip speed ratio of 7 are shown in Figure 5.4.

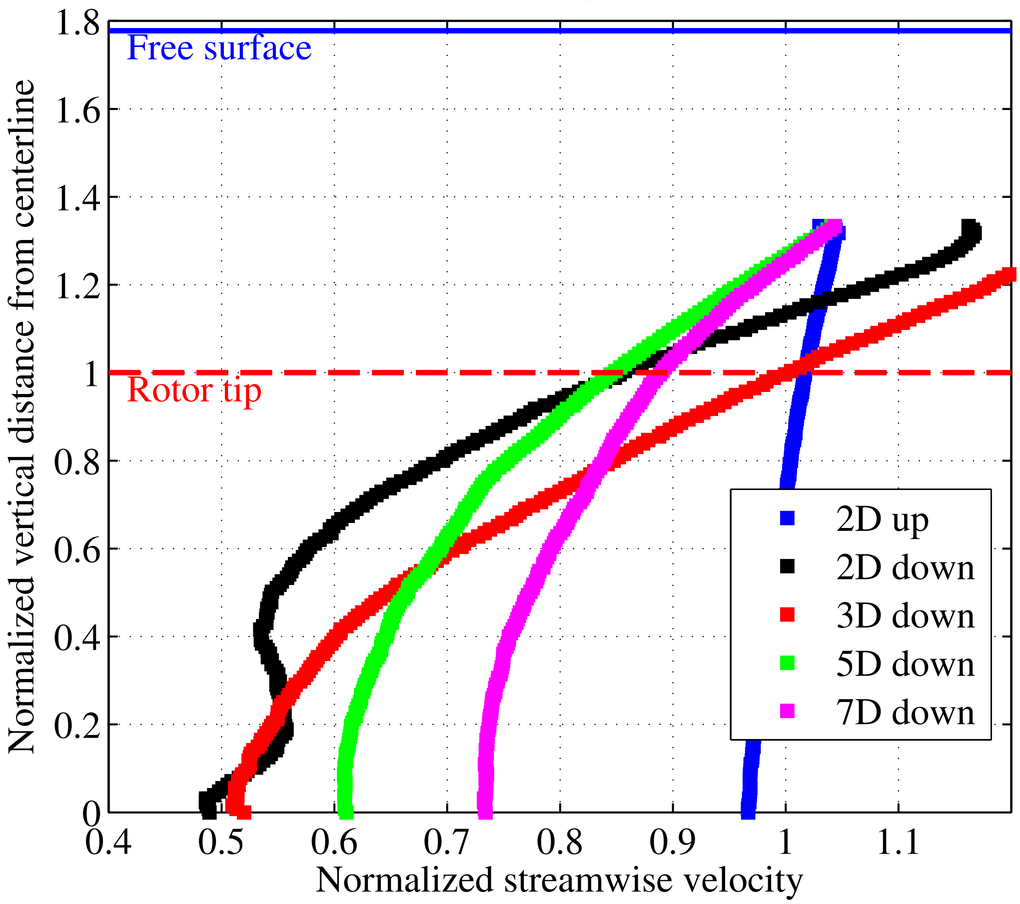

Figure 5.4: Mean streamwise velocity profiles 2 rotor diameters upstream and 2–7 rotor diameters downstream of the rotor plane for a single turbine operating at TSR=7.

Figure 5.4 shows a slightly sheared inflow velocity profile, a velocity deficit that has a maximum at 2 diameters downstream, and a recovering wake that has a centerline velocity ∼73% of the free-stream velocity at 7 diameters downstream of the rotor plane. The measurement position two diameters downstream is just downstream of the back edge of the nacelle, and the velocity profile at that location presents a kink near the centerline due to the recirculation zone directly downstream of the blunt nacelle edge. Velocity profiles at 2D and 3D downstream show an accelerated zone outboard of the rotor tip, and a highly sheared zone across the boundary of the rotor cylinder.

Figure 5.5 shows mean streamwise velocity profiles for a range of TSR; each subplot is a streamwise location in the wake. The velocity profiles at 2 diameters downstream of the rotor plane are variable with TSR, but farther downstream the velocity profiles collapse onto a single profile for all TSR. This result shows that, while the near wake is strongly influenced by blade rotation and therefore by the tip speed ratio, the mean velocity of the wake more than 5 diameters downstream of the rotor plane does not depend strongly on upstream operating conditions, at least for in the range TSR 5-10, where the efficiency of the turbine is fairly invariant. Our hypothesis is that the lower power extraction for values of non-optimal TSR is compensated with higher turbulent dissipation in the blade tip vortices (for high TSR) and in the separated flow near the root (for low TSR) such that the overall reduction in the flow kinetic energy flux has little dependency on operating TSR. This result does not support the working hypothesis that the overall efficiency of an array can be optimized by operating the upstream turbines at off-optimum conditions, allowing more kinetic energy flux to reach the downstream turbines, and thus balance and optimize the power extraction across the array. Figure 5.5 clearly show that the mean velocity in the wake, i.e. the incoming flow available to the downstream turbines in an array, does not depend strongly on tip speed ratio after only 5 diameters downstream, and thus operating the upstream turbine at sub-optimal tip speed ratios will produce less power, and more power will be dissipated in the near wake, but more energy flux will not reach the downstream turbines, as desired.

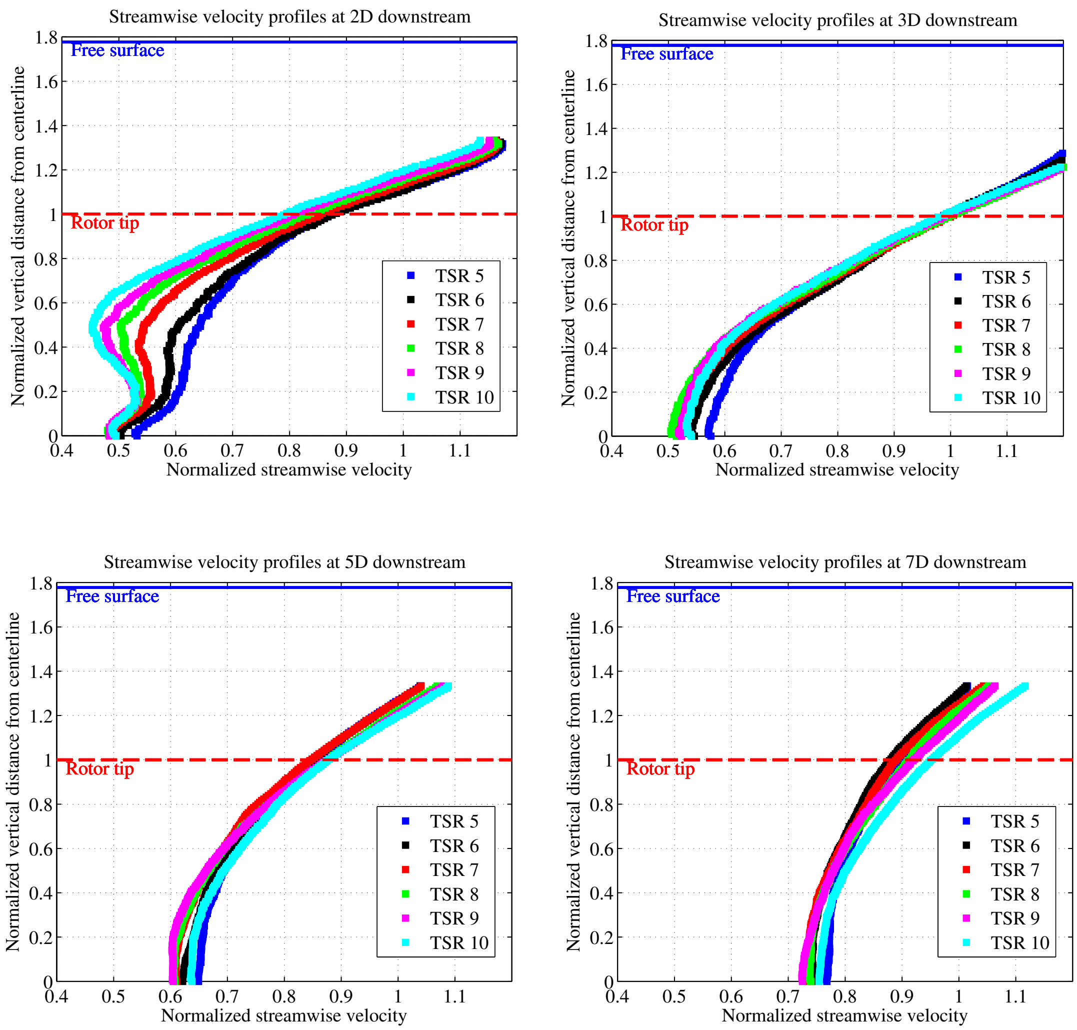

Figure 5.5: Mean streamwise velocity profiles at various locations in the wake of a single turbine for a range of tip speed ratios.

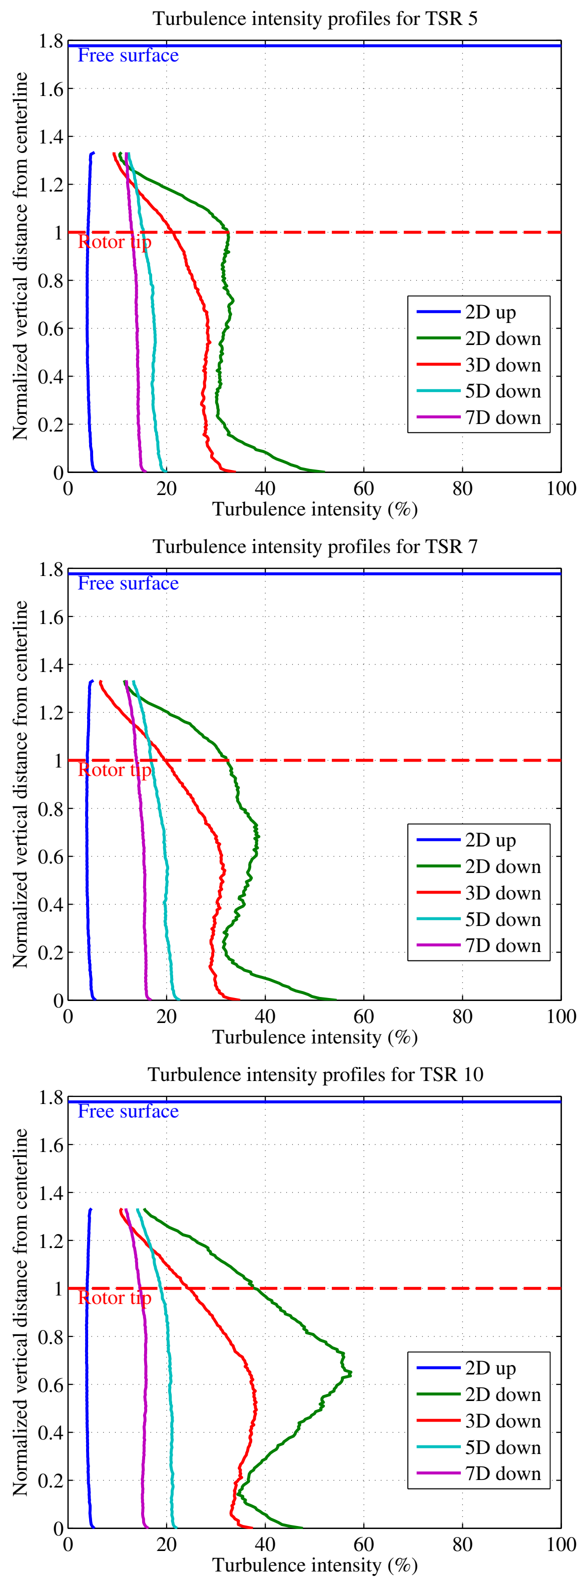

Figure 5.6: Mean turbulence intensity profiles at various locations in the wake of a single turbine for a range of tip speed ratios.

Figure 5.6 show a similar trend in turbulence intensity profiles—significant differences between operating conditions in the near wake that lessen and finally disappear in the far wake. In the turbulence intensity profiles at 2 diameters downstream there is clear evidence of the high turbulent production in the shear layers near the edge of the rotor cylinder and in the highly turbulent recirculation zone directly downstream of the nacelle near the centerline. The well-defined peak of turbulence intensity associated with the TSR=10 operating condition is due to the more densely spaced tip vortices, as a rapidly spinning rotor will produce a helical tip vortex structure with a greater pitch than a more slowly spinning rotor. These differences in the near wakes are not present at five diameters downstream, and the turbulence intensity is well distributed axially in the wake.

5.3 Co-axially spaced turbine arrays

The results from testing three different configurations of co-axially spaced turbine arrays are reported in this section. The first configuration consists of two turbines separated by several co-axial distances. The second configuration is an array of three turbines separated by a distance of five rotor diameters between each two turbines, and the third configuration is an array of three turbines with five and seven rotor diameters between consecutive turbines and a 0.25 D lateral offset.

5.3.1 Two co-axially spaced turbines at various spacings

Two turbines with a variety of co-axial distances were tested in order to elucidate the effect of spacing on performance and wake development on the downstream turbine. The performance curves for the upstream and downstream turbines are presented in Figure 5.7. These data show that the turbine located five diameters downstream (of the front turbine) has less than half the efficiency of the upstream turbine. As

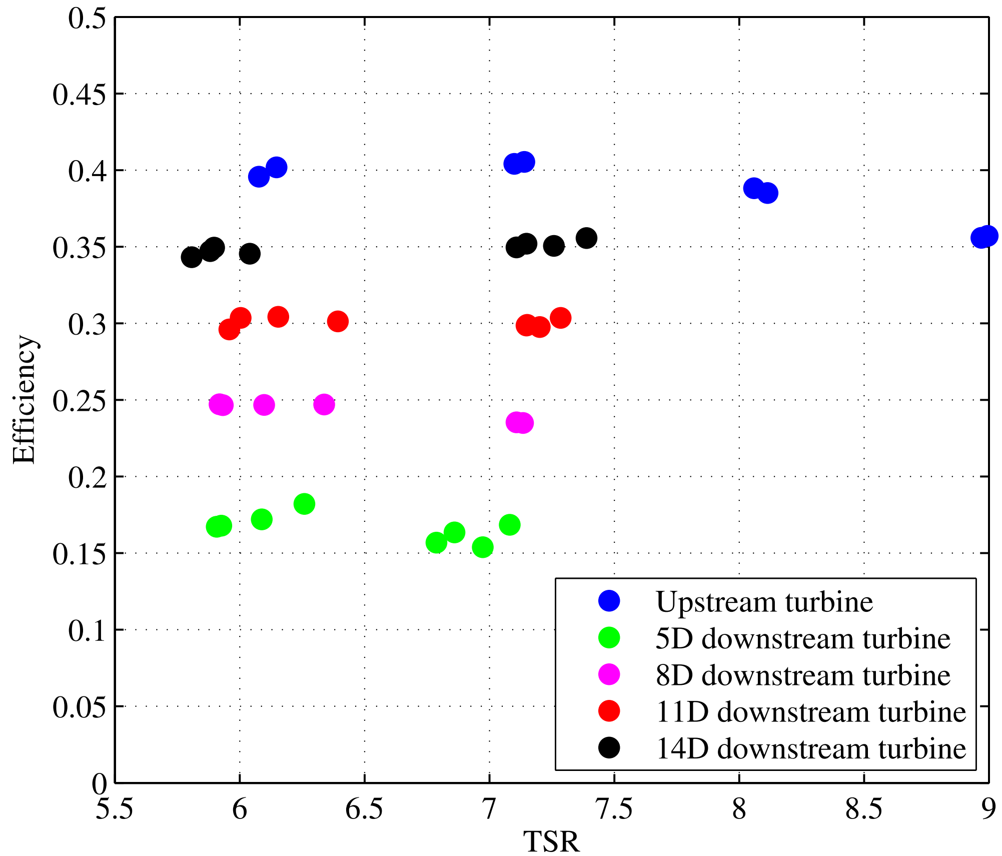

Figure 5.7: Performance curves for an upstream turbine and a downstream turbine at various co-axial distances.

the turbine spacing is increased, a steady increase in performance is observed for the downstream turbine and, at the largest separation distance tested, 14 diameters, the performance of the downstream turbine has recovered to within 10% of the performance of the upstream turbine.

Figure 5.8 shows the recovery of the downstream turbine performance as the distance between the two turbines is increased. Extrapolating the data with a linear trend predicts that the downstream turbine performance will recover to 100% of the upstream turbine performance at a spacing distance of 16D.

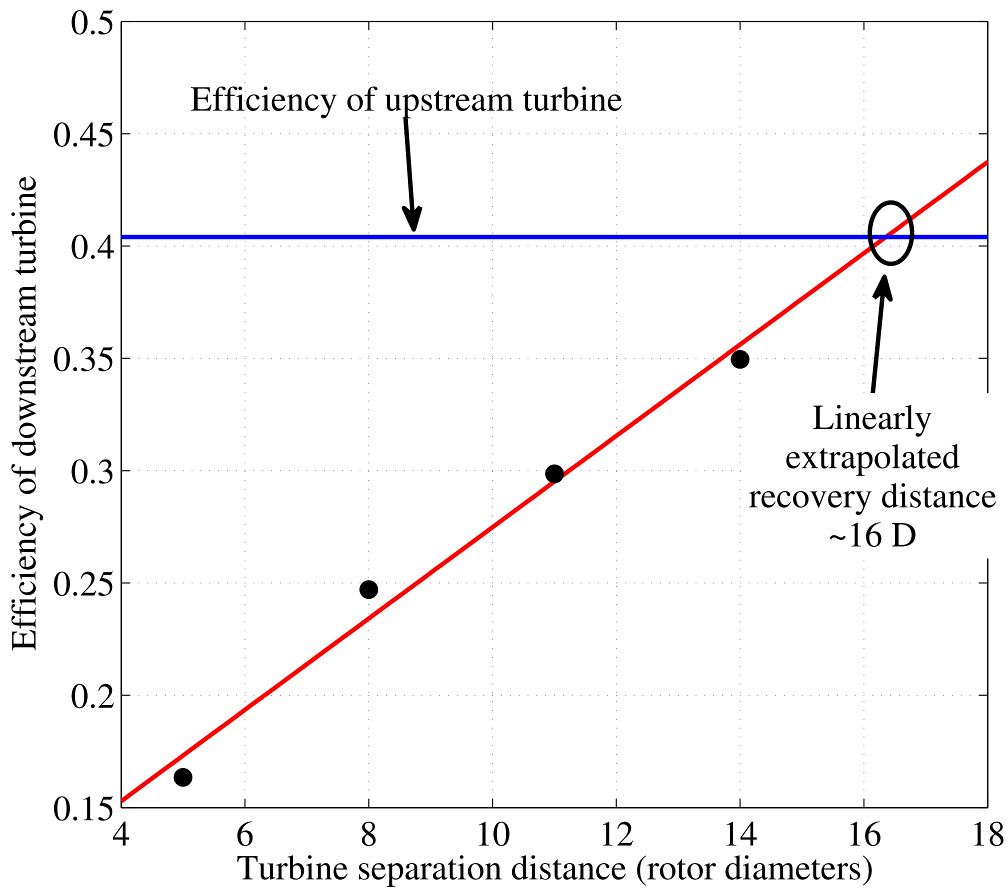

Figure 5.8: Mean downstream turbine peak performance as a function of turbine spacing in a two-turbine co-axially spaced array

Figures 5.9, 5.10, and 5.11 show the mean streamwise velocity profiles in the wake of the upstream and downstream turbines for the two-turbine co-axial array. All three configurations show that the wake recovery of the second turbine occurs more quickly than the front turbines.

Figure 5.12 compares the velocity profiles at 3D (a) and 6D (b) downstream of each turbine for all of the array separation distances. For the array separated by 5D, there is a distinctly faster wake recovery, but there is little difference in the velocity

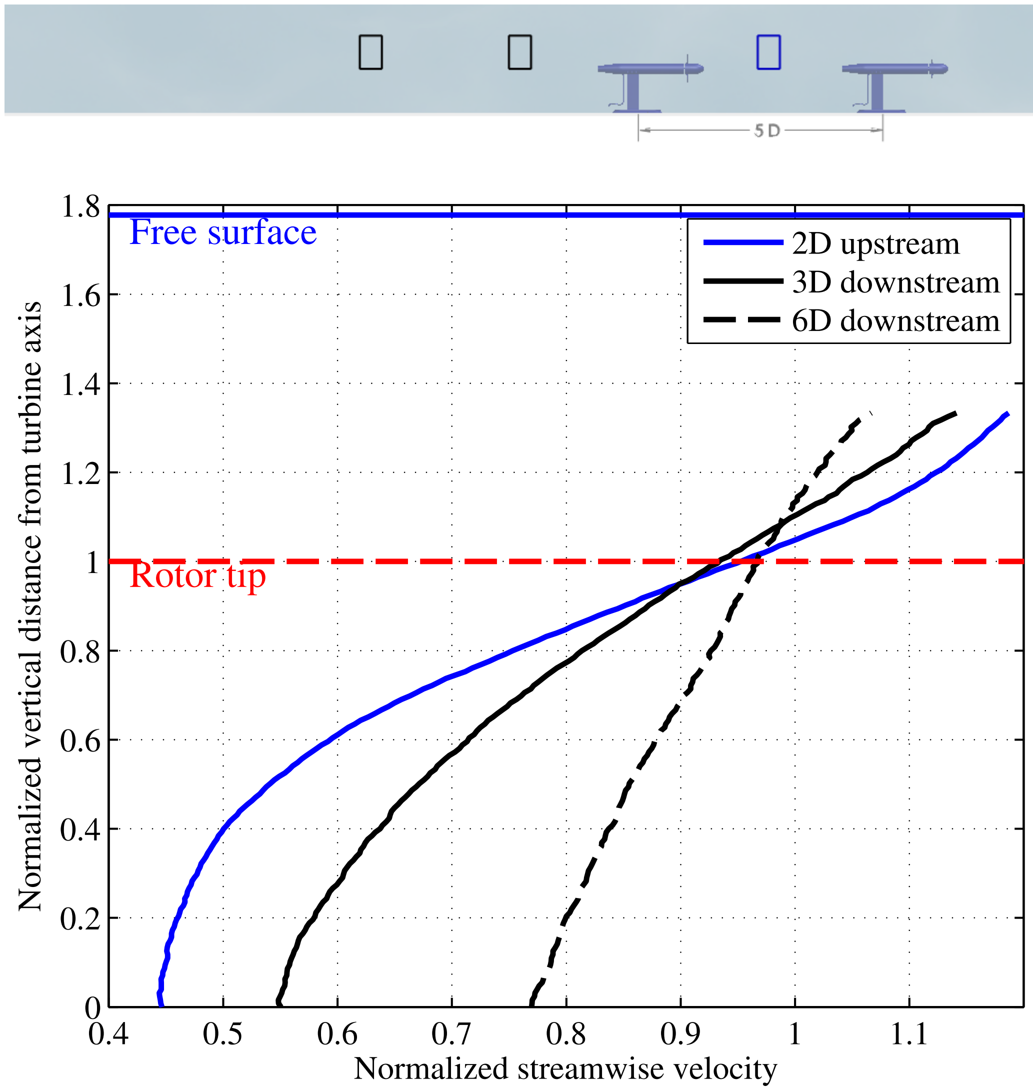

Figure 5.9: Mean streamwise velocity profiles associated with a co-axial array of two turbines separated by 5D. Velocity profiles are shown 2D upstream of downstream turbine rotor plane (3D downstream of the upstream turbine), 3D downstream of the downstream turbine, and 6D downstream of the downstream turbine.

profiles for the arrays separated by 8D and 11D. This indicates that the difference in performance between the 8D and 11D separated arrays seen in Figure 5.7 is due to a factor other than mean incident velocity.

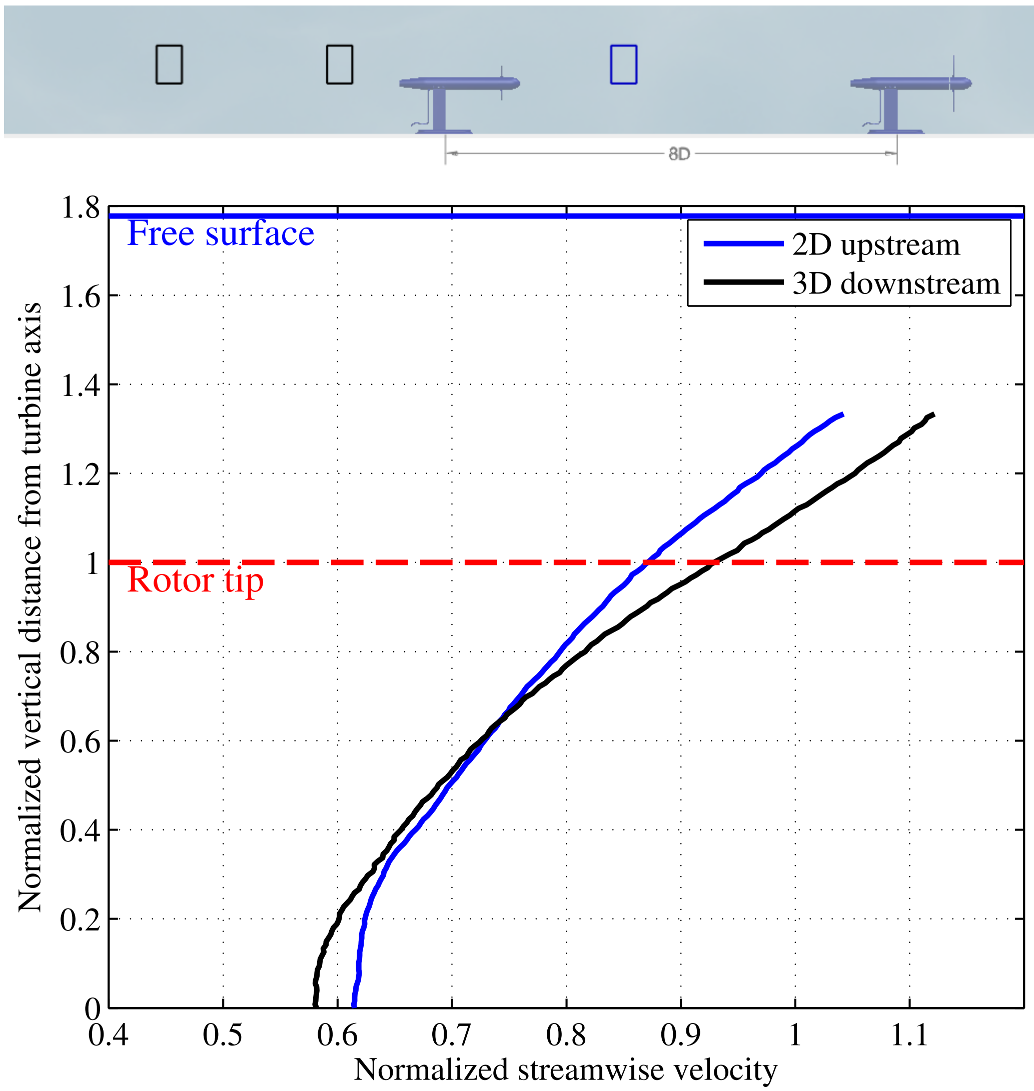

Figure 5.10: Mean streamwise velocity profiles associated with a co-axial array of two turbines separated by 8D diameters. Velocity profiles are shown 2D upstream of the downstream rotor plane (6D downstream of the upstream turbine), and 3D downstream of the downstream turbine.

Figures 5.13, 5.14, 5.15, and 5.16 are structured similarly to Figures 5.9, 5.10, 5.11, and 5.12 and show the mean tubulence intensity profiles in the wakes of two co-axially spaced turbines. Contrary to the results for the mean velocity profile, the

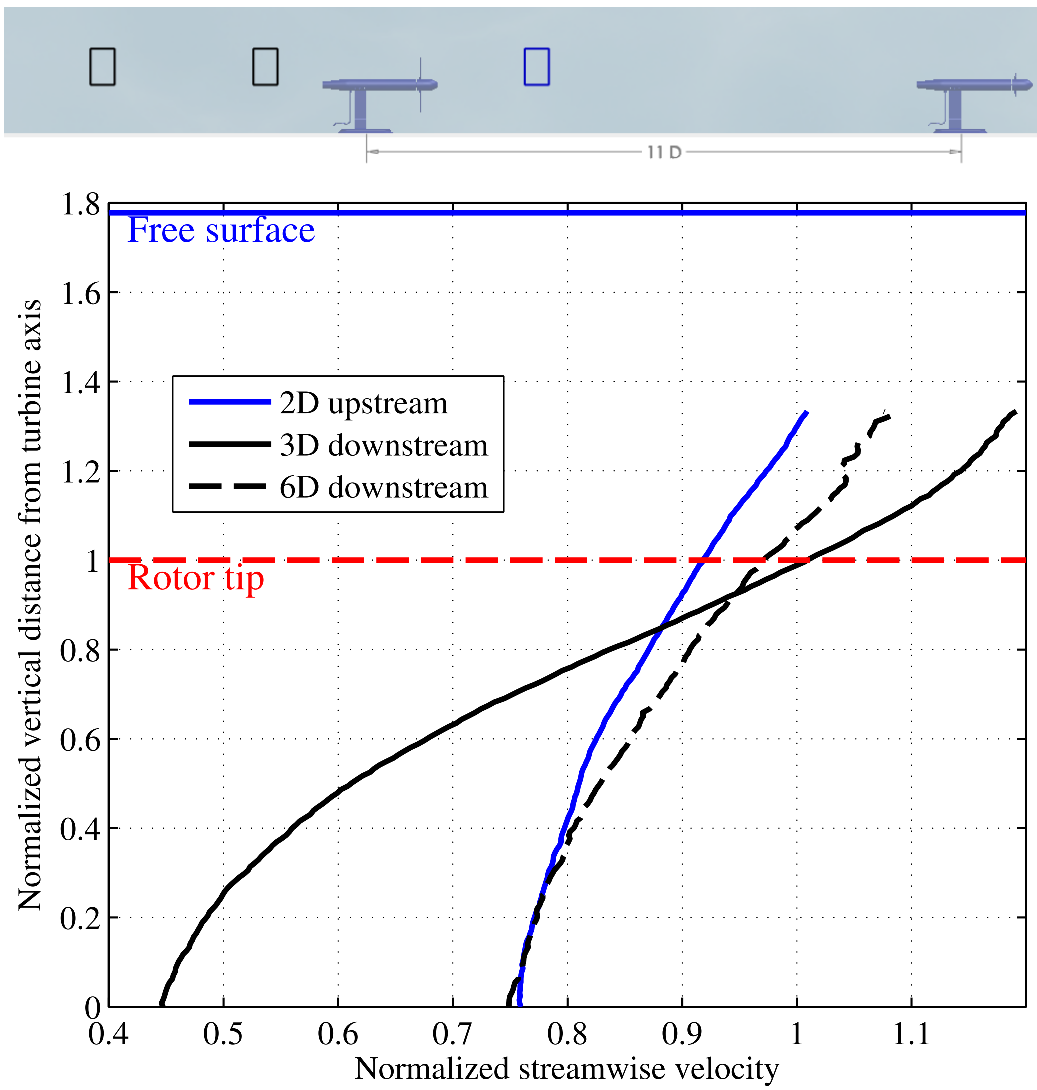

Figure 5.11: Mean streamwise velocity profiles associated with a co-axial array of two turbines separated by 11D Velocity profiles are shown 2D upstream of the downstream turbine (9D downstream of the upstream turbine), 3D downstream of the downstream turbine, and 6D downstream of the downstream turbine.

turbulence intensity in the wake of the second turbine is higher and takes much longer to relax back to freestream conditions than the TI in the wake of the front turbine. Turbulence production is clearly highest near the rotor tip (as the blade tip vortices

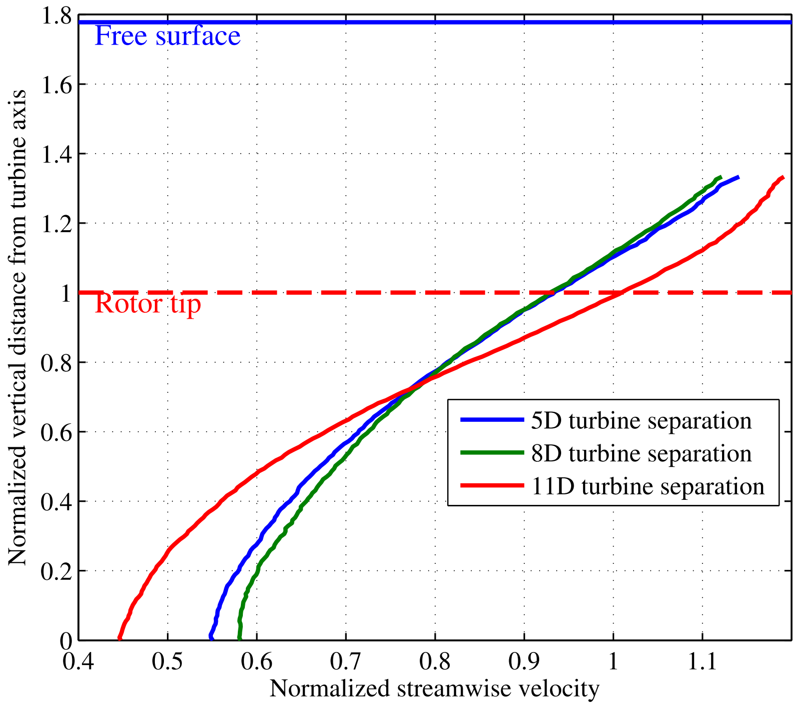

(a) 3D downstream

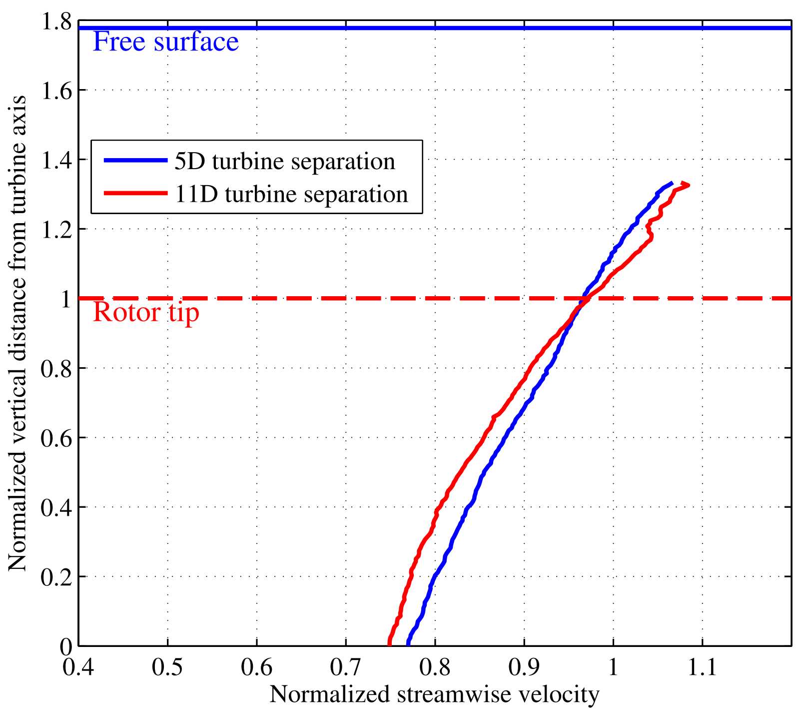

(b) 6D downstream

Figure 5.12: Comparison of the streamwise velocity profiles at 3D (a) and 6D (b) downstream of the downstream turbine in variously spaced co-axial arrays.

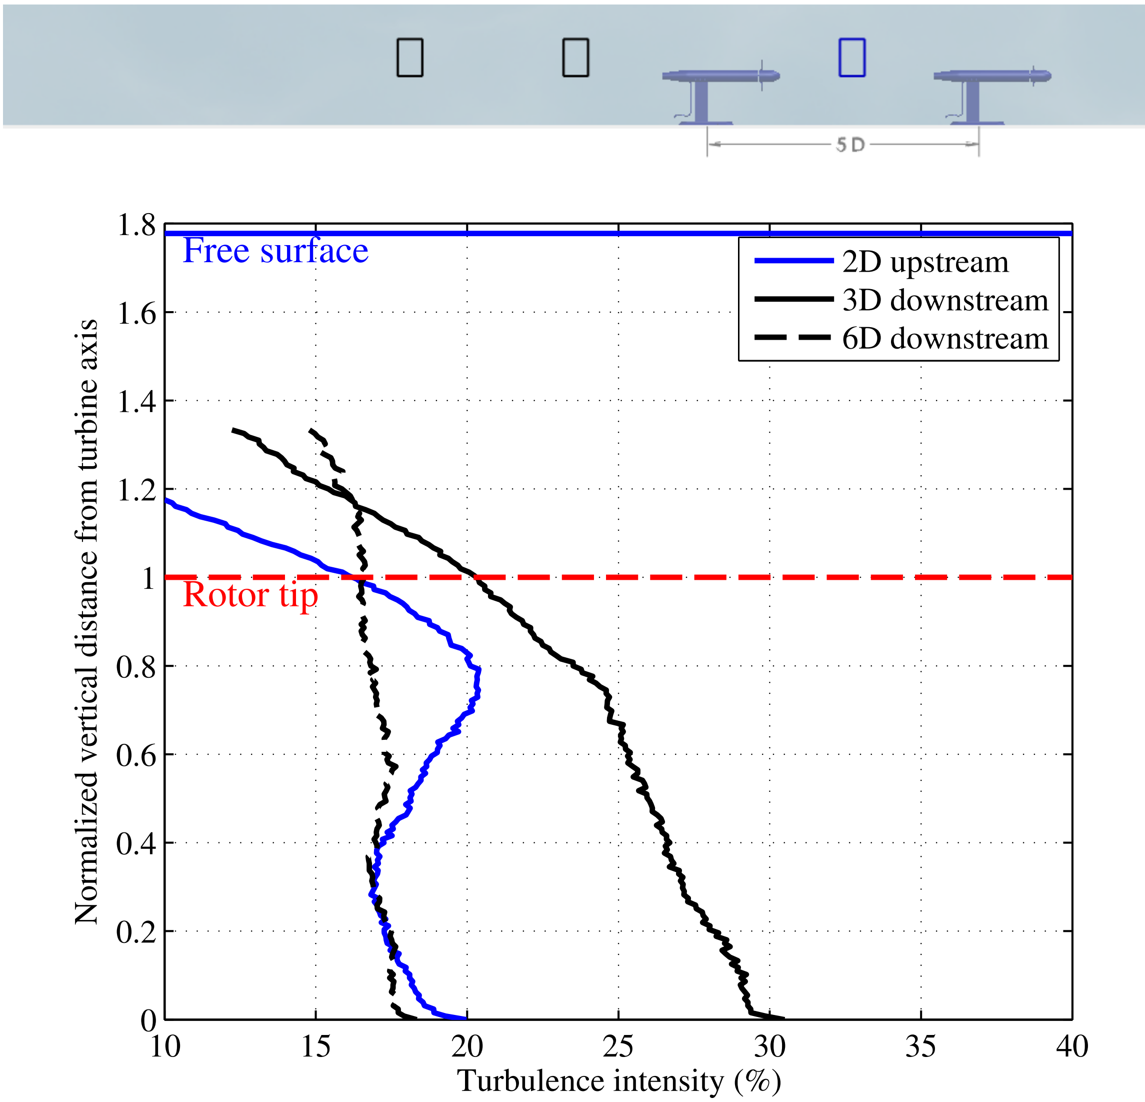

Figure 5.13: Mean turbulence intensity (TI) profiles associated with a co-axial array of two turbines separated by 5D. TI profiles are shown 2D upstream of the upstream turbine rotor plane (3D downstream of the upstream turbine), 3D downstream of the downstream turbine, and 6D downstream of the downstream turbine.

are included in the calculation of the turbulence fluctuations from the PIV data) and near the nacelle, where the flow separation and recirculation sheds vortices into the

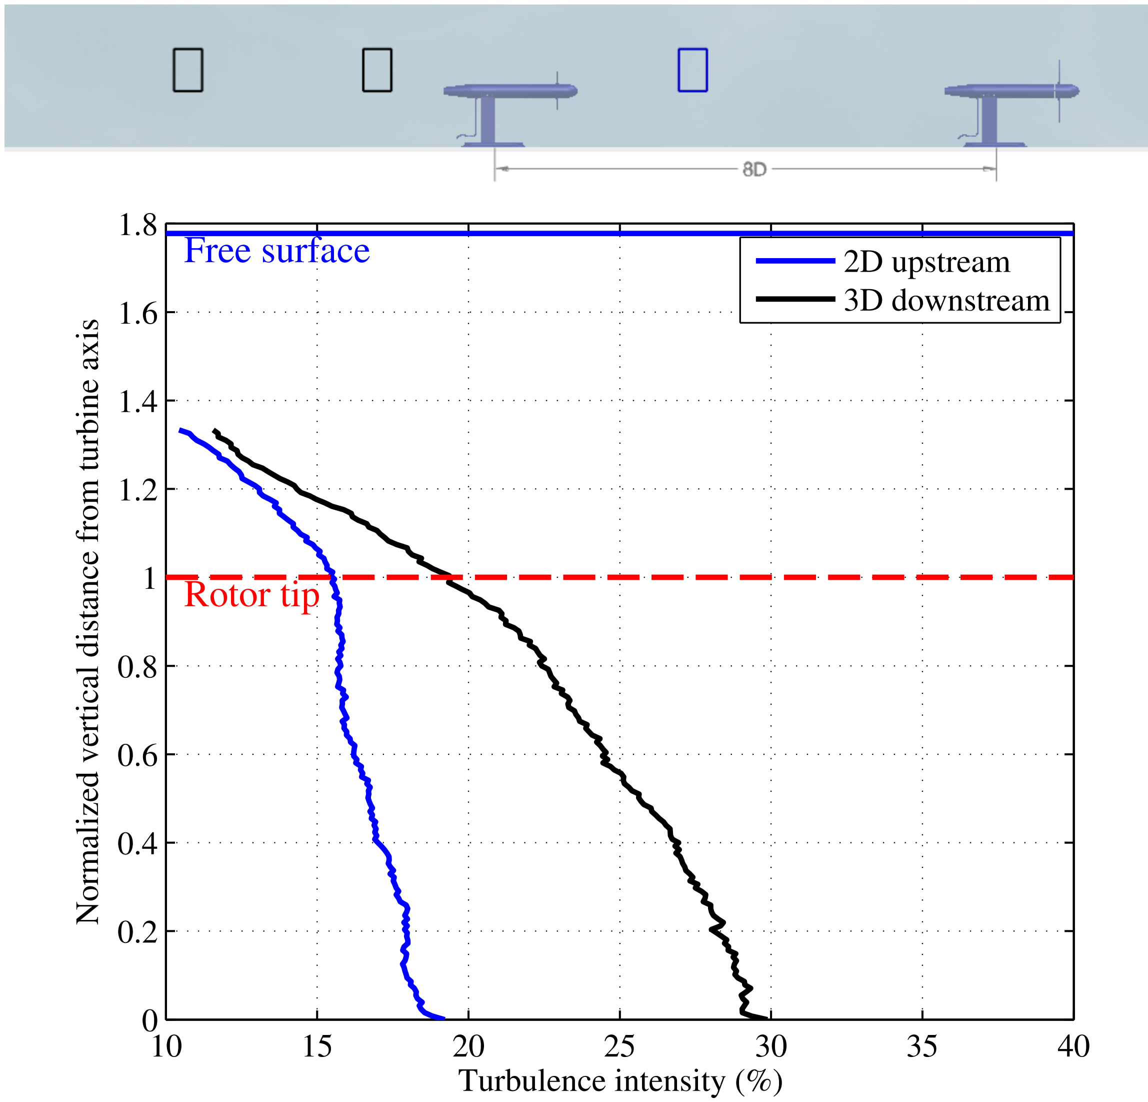

Figure 5.14: Mean turbulence intensity (TI) profiles associated with a co-axial array of two turbines separated by 8D. TI profiles are shown 2D upstream of the upstream turbine rotor plane (6D downstream of the upstream turbine), and 3D downstream of the downstream turbine

wake. These velocity fluctuations mix across the wake cross section, represented in the figures by the vertical profiles. When the TI profiles are compared at a given

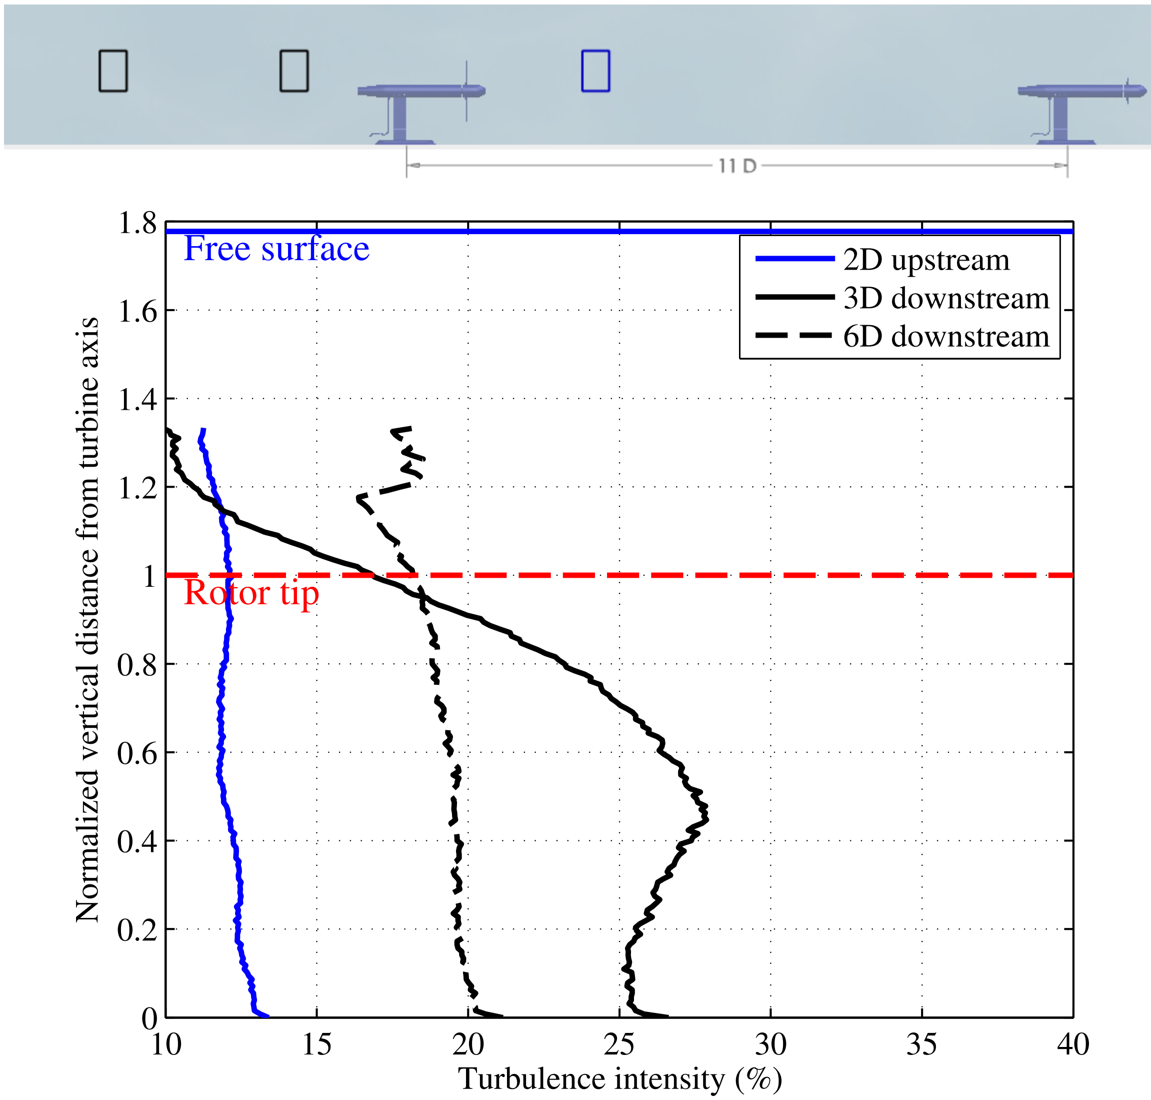

Figure 5.15: Mean turbulence intensity (TI) profiles associated with a co-axial array of two turbines separated by 11D. TI profiles are shown 2D upstream of the upstream turbine rotor plane (9D downstream of the upstream turbine), 3D downstream of the downstream turbine, and 6D downstream of the downstream turbine.

station (3D or 6D downstream of the second turbine), we observe that the TI reaches a “fully developed” condition for turbine separation equal or larger than 8D. The TI

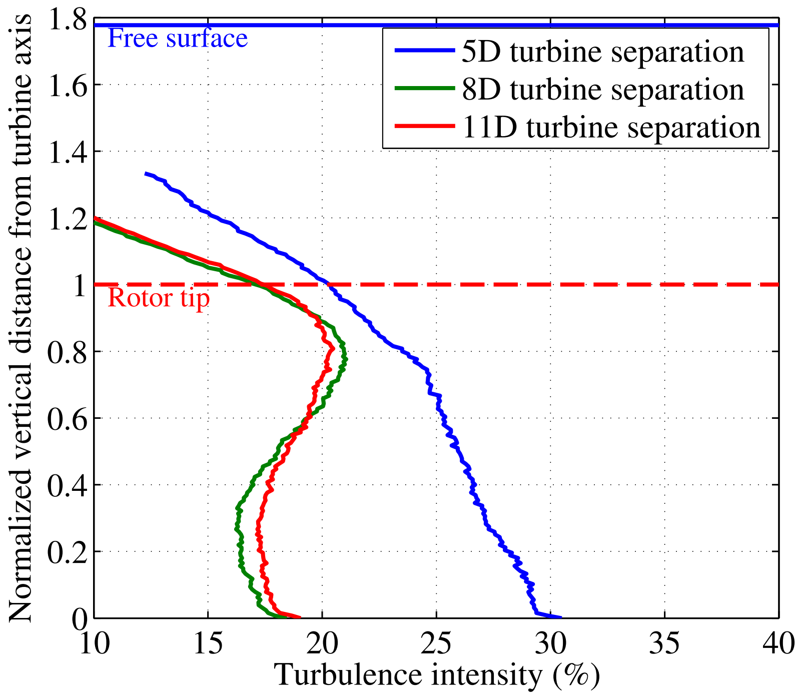

(a) 3D downstream

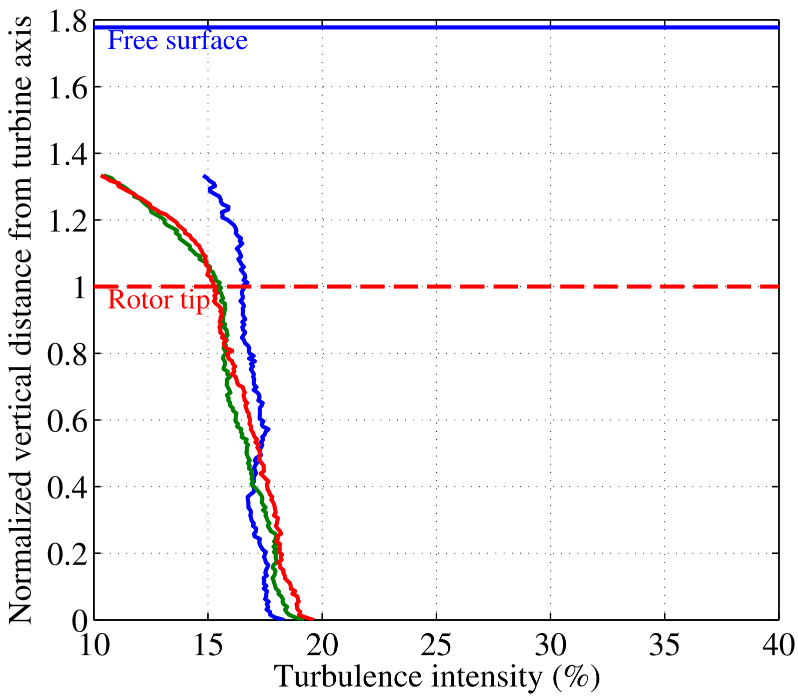

(b) 6D downstream

Figure 5.16: Comparison of the turbulence intensity profiles at 3D (a) and 6D (b) downstream of the downstream turbine in variously spaced co-axial arrays.

profiles for the turbine separations of 5D and 8D are very different, much higher for the 5D separation where it has not had enough time to decay before it encounters the second turbine, but it is approximately equal for 8D and 11D separations, showing that the decay slows down and that separations further than 8D may be wasteful if a compact array is the objective.

5.3.2 Three co-axially spaced turbines separated by seven rotor diameters

The performance curve results from the testing of three co-axially spaced turbines with a separation of seven rotor diameters are presented in Figure 5.17.

The midstream turbine in this three-turbine configuration (7D behind the first) was observed to have similar efficiency to the second turbine, when it was spaced eight diameters downstream, in the two-turbine configuration discussed above and shown in Figure 5.7. The downstream-most turbine is shown to have significantly better efficiency than the midstream turbine. An explanation for this is found, at least partially, by comparing the velocity profiles 2D upstream of the midstream rotor plane (shown in Figure 5.20) against the velocity profile 2D upstream of the downstream rotor plane (shown in Figure 5.21). This two figures show the faster wake recovery in the combined wake of the upstream and midstream turbines compared to the wake recovery of the upstream turbine alone. There are two causes of this faster wake recovery: the midstream turbine extracts less momentum from the flow than the upstream turbine, and it produces high turbulent fluctuations and mixing, resulting in the combined wake of the upstream and downstream turbine having a much higher turbulence intensity, which promotes momentum diffusion radially across the wake. The difference in performance between the midstream turbine performance and the downstream performance has implications for the optimization of array spacing. For example, from the perspective of optimizing overall array performance the optimal position of the midstream turbine may be closer to the downstream turbine, as the

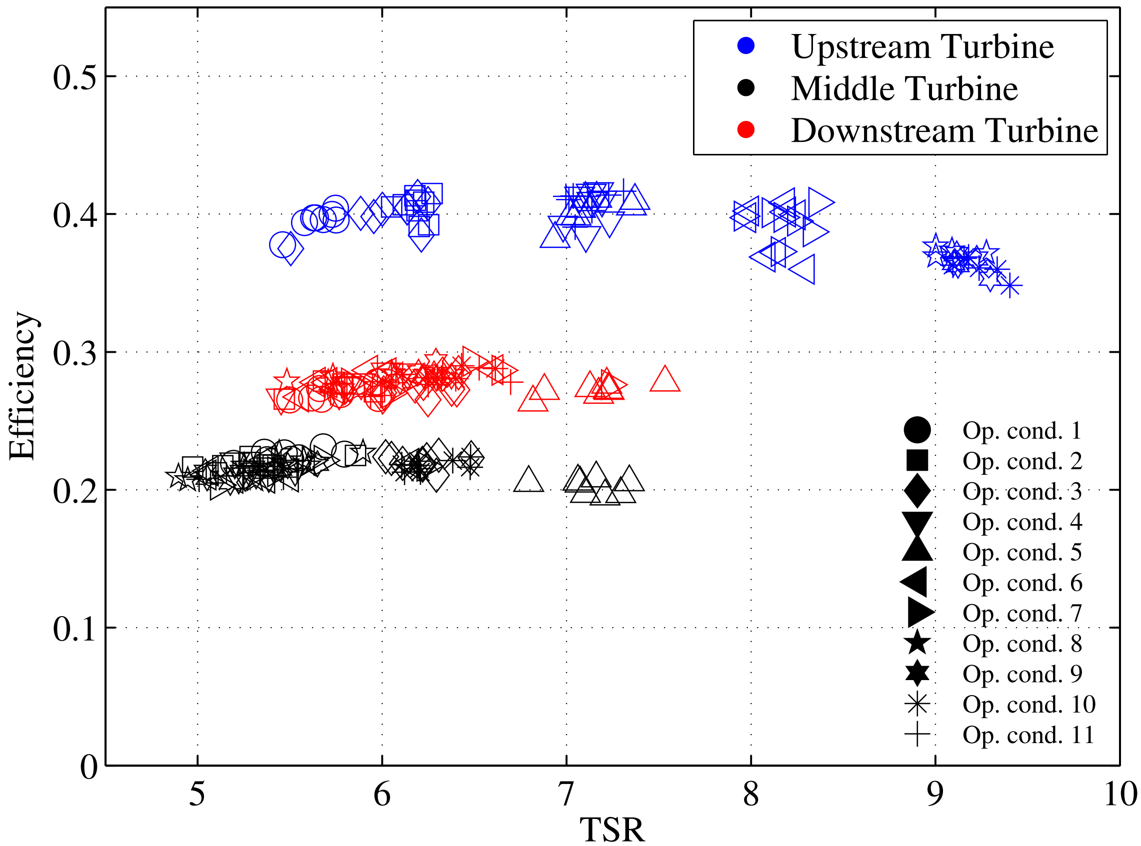

Figure 5.17: Performance curves for three co-axially spaced turbines separated by seven rotor diameters. Different markers represent different simultaneous tests at various combinations of tip speed ratio.

faster wake recovery of the combined wake could allow a smaller spacing than the spacing between the upstream and midstream turbines. If, however, higher power density (defined as power produced per area of the sea-floor occupied by the array) is desired, the two back rows could be moved closer to the front turbine with the corresponding decrease in array footprint but very little loss of power produced (by the midstream turbine) and unchanged power in the first and third rows.

5.3.3 Three co-axially spaced turbines separated by five rotor diameters

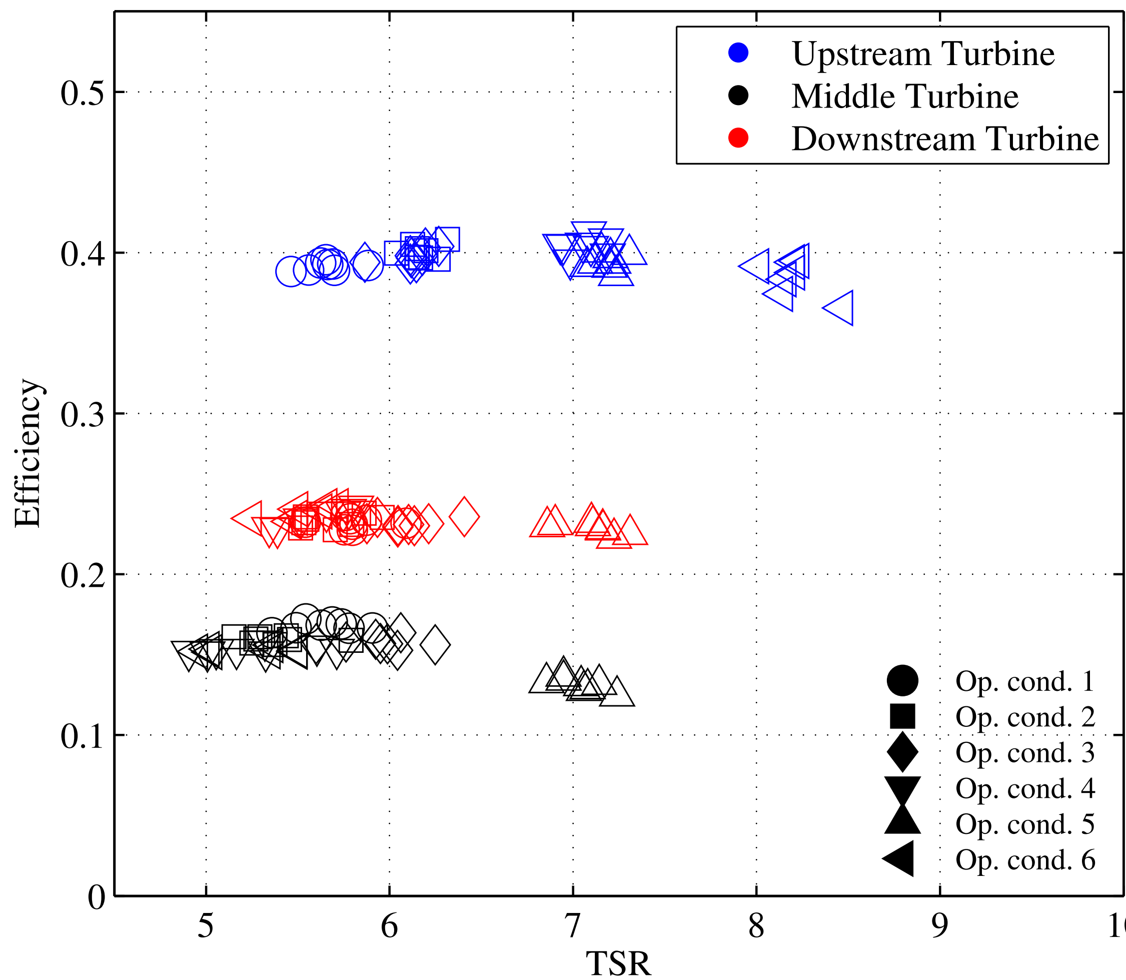

Figure 5.18: Performance curves for three co-axially spaced turbines separated by five rotor diameters. Different markers represent different simultaneous tests at various combinations of tip speed ratio.

Similar trends to the previous configuration (three turbines spaced by 7 rotor diameters coaxially) are observed in the performance curves for three co-axially spaced turbines with a separation distance of five rotor diameters, shown in Figure 5.18.

As expected, the midstream turbine in this configuration has a lower efficiency than the midstream turbine in the previous (seven diameter spacing) configuration, and similar performance to the downstream turbine separated by five diameters in the two-turbine configuration, shown in Figure 5.7. The downstream turbine also has higher efficiency than the midstream turbine, as observed in the configuration with 7D spacing discussed above and presented in Figure 5.17.

Figures 5.17 and 5.18 show the mean coefficients of performance for all three turbines measured simultaneously. Little variation is shown in the performance of the midstream and downstream turbines as the TSR of the upstream turbine varies from TSR=5.5–9.5. This independence of midstream and downstream turbine performance from upstream turbine operating condition has implications for turbine array optimization. As mentioned above with regard to the velocity profiles in the wake of a turbine for different TSR, the tuning of the array by operating the upstream turbine in a suboptimal TSR is not feasible, as it does not improve the power production of the turbines located downstream. The measurements of power production for 3 turbine arrays, shown in figures 5.17 and 5.18, confirm this hypothesis resulting from the observation of the wake of a single turbine.

It is interesting to note that the velocity profile upstream and downstream of the last turbine in the array quickly becomes "fully developed". Figure 5.21 shows that the flow 5D downstream of this third turbine has already recovered back to the incoming flow for this turbine. This is confirmed by the data in figure 5.24 where we can observe that the flow velocity profiles at 5D downstream of the second (midstream) and third (downstream) turbines are essentially equal. This shows that the array is already in "fully developed" mode and that any turbines placed in further rows downstream at the correct spacing (in this case 5D) will operate similarly to turbine number 3 (the downstream-most) in this small array. This is important information in the effort to extrapolate the findings of these experiments, and possibly of simulations, where

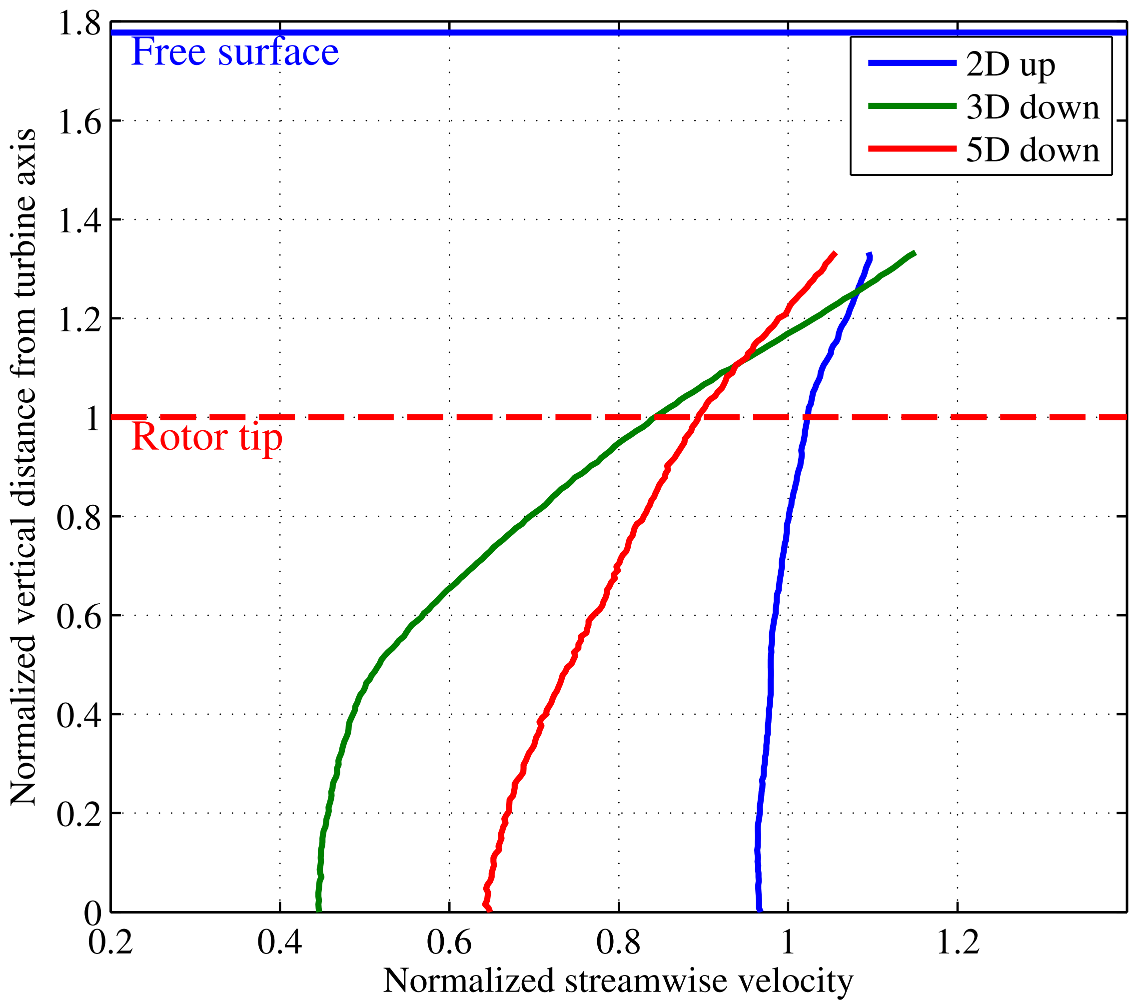

Figure 5.19: Mean streamwise velocity profiles in the wake of the upstream turbine in a three turbine co-axially spaced array separated by 7D.

the number of turbines is a limitation, to full-scale arrays with tens or hundreds of turbines.

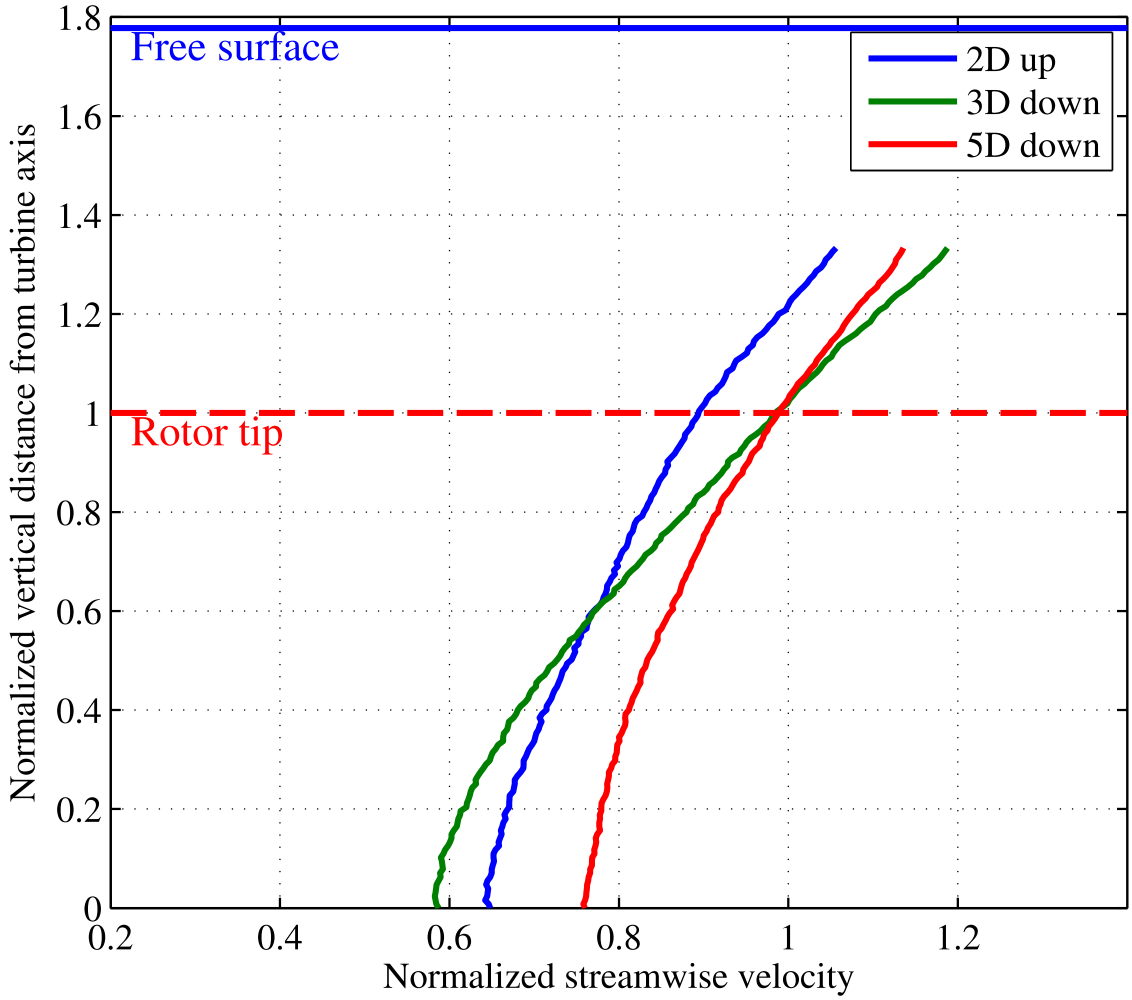

Figure 5.20: Mean streamwise velocity profiles in the wake of the midstream turbine in a three turbine co-axially spaced array separated by 7D.

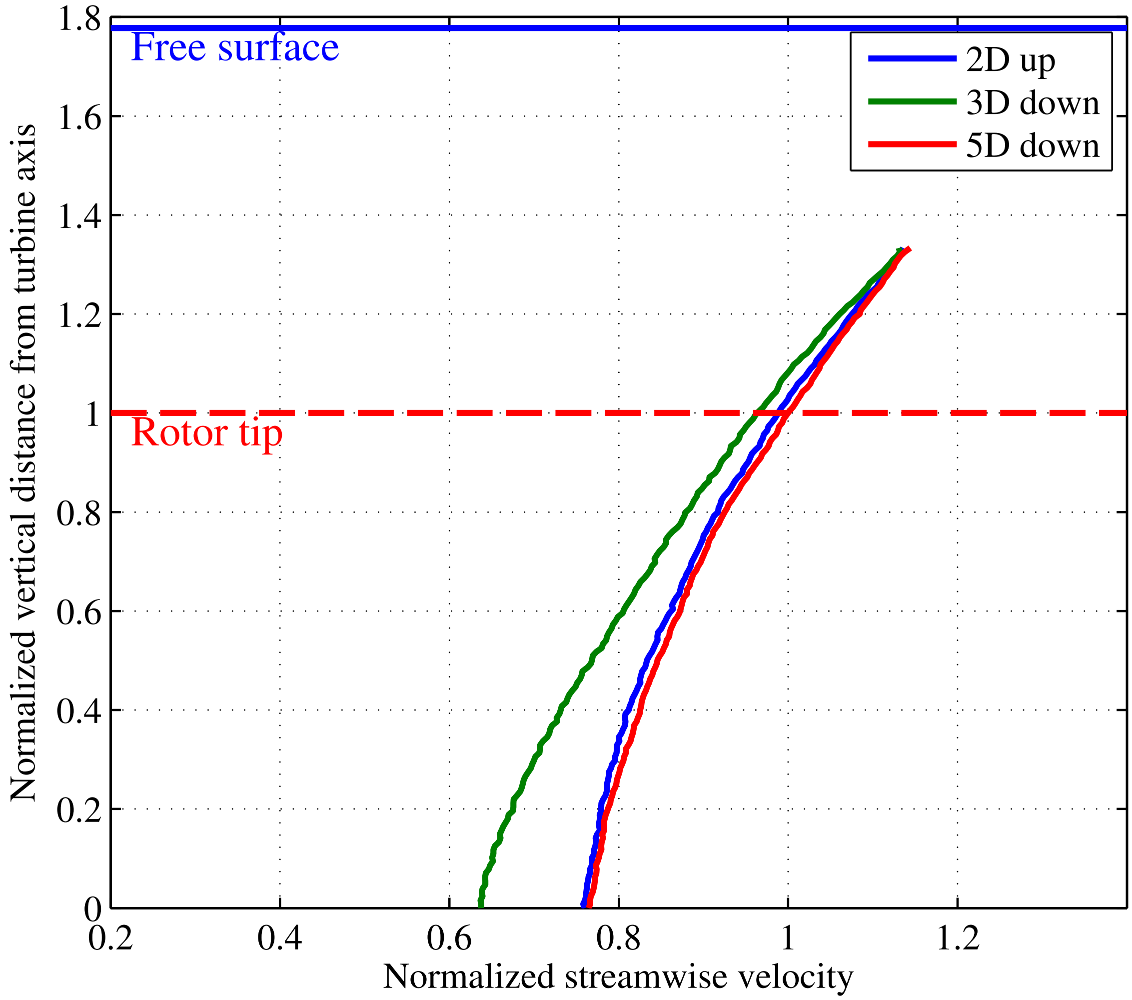

Figure 5.21: Mean streamwise velocity profiles in the wake of the downstream turbine in a three turbine co-axially spaced array separated by 7D.

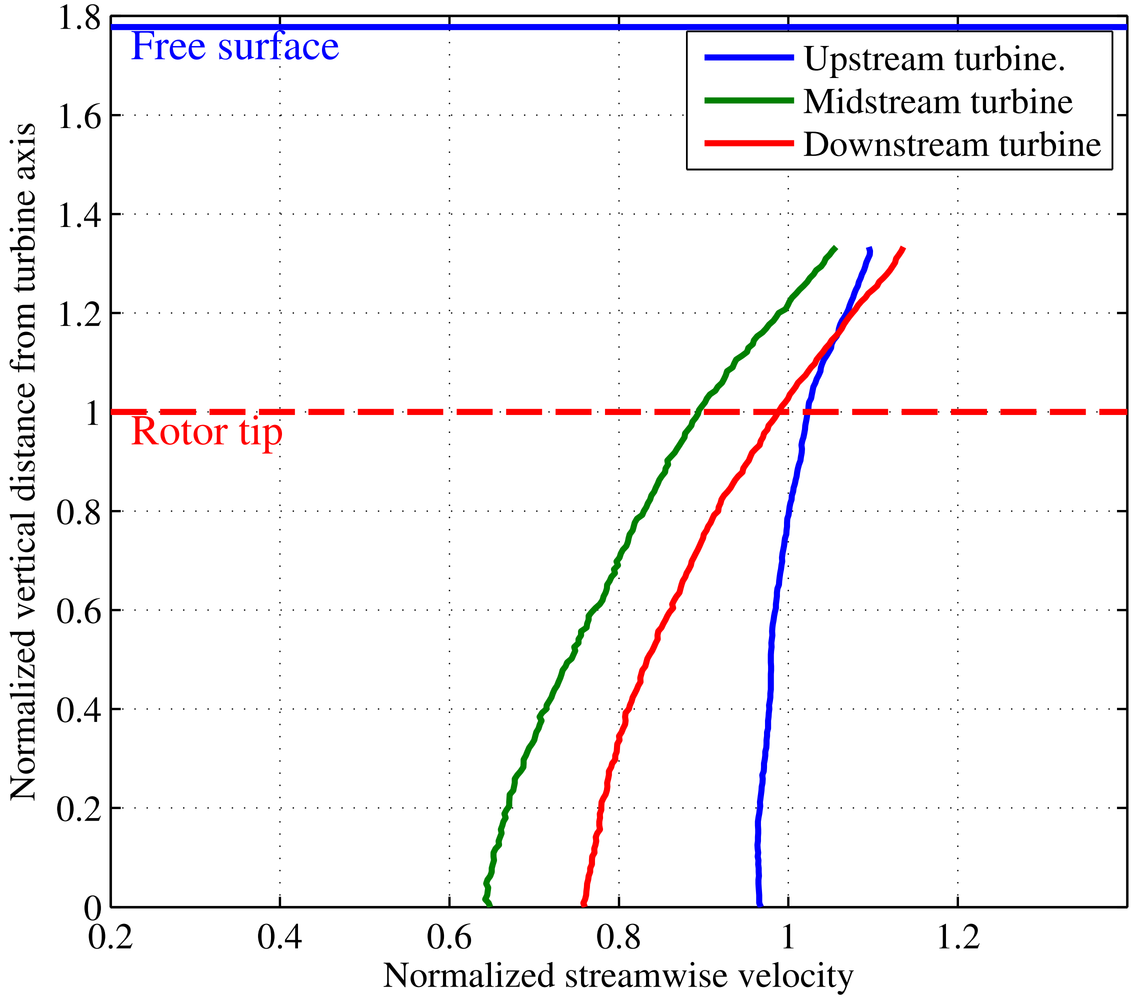

Figure 5.22: Comparison of the mean streamwise velocity profiles 2D upstream of each turbine in a three turbine co-axial array separated by 7D.

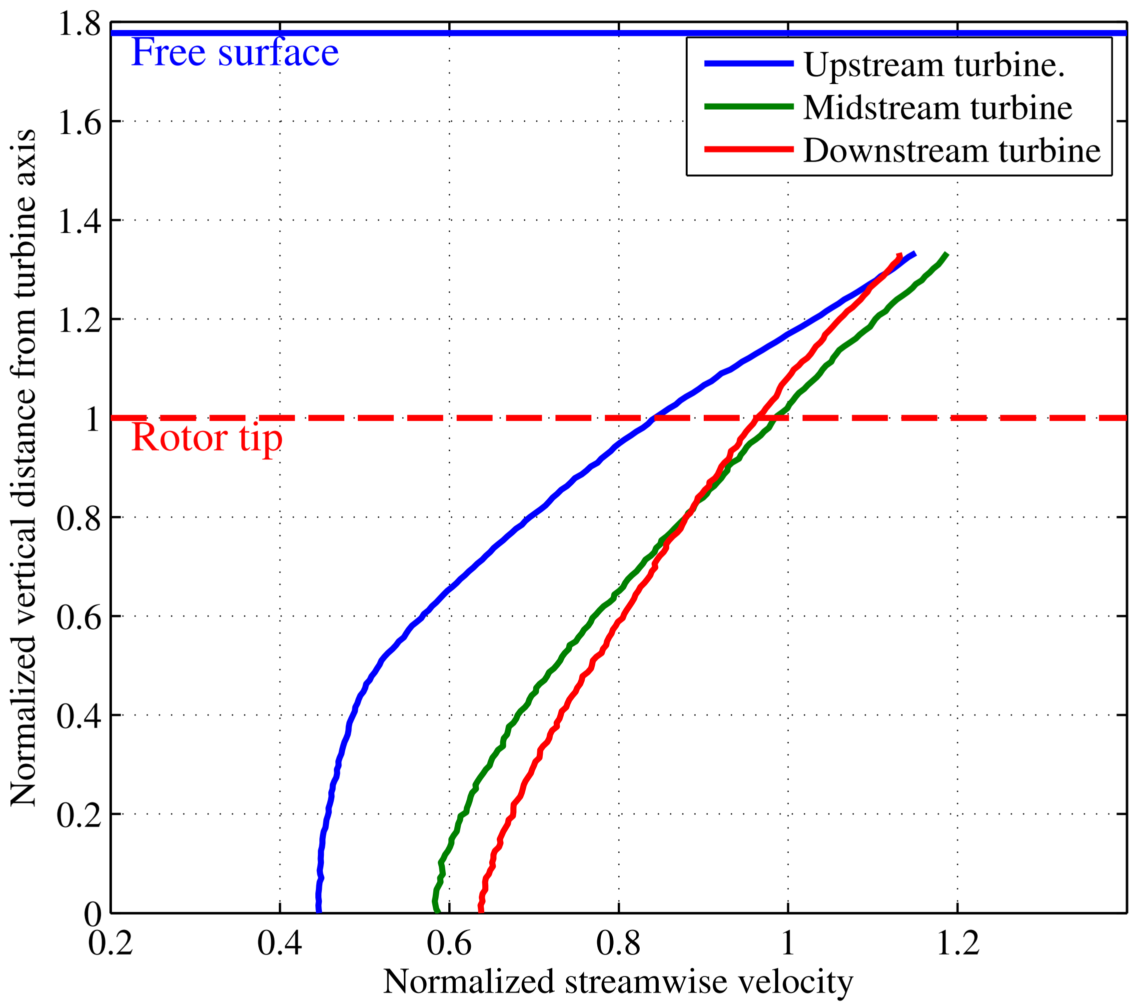

Figure 5.23: Comparison of the mean streamwise velocity profiles 3D downstream of each turbine in a three turbine co-axial array separated by 7D.

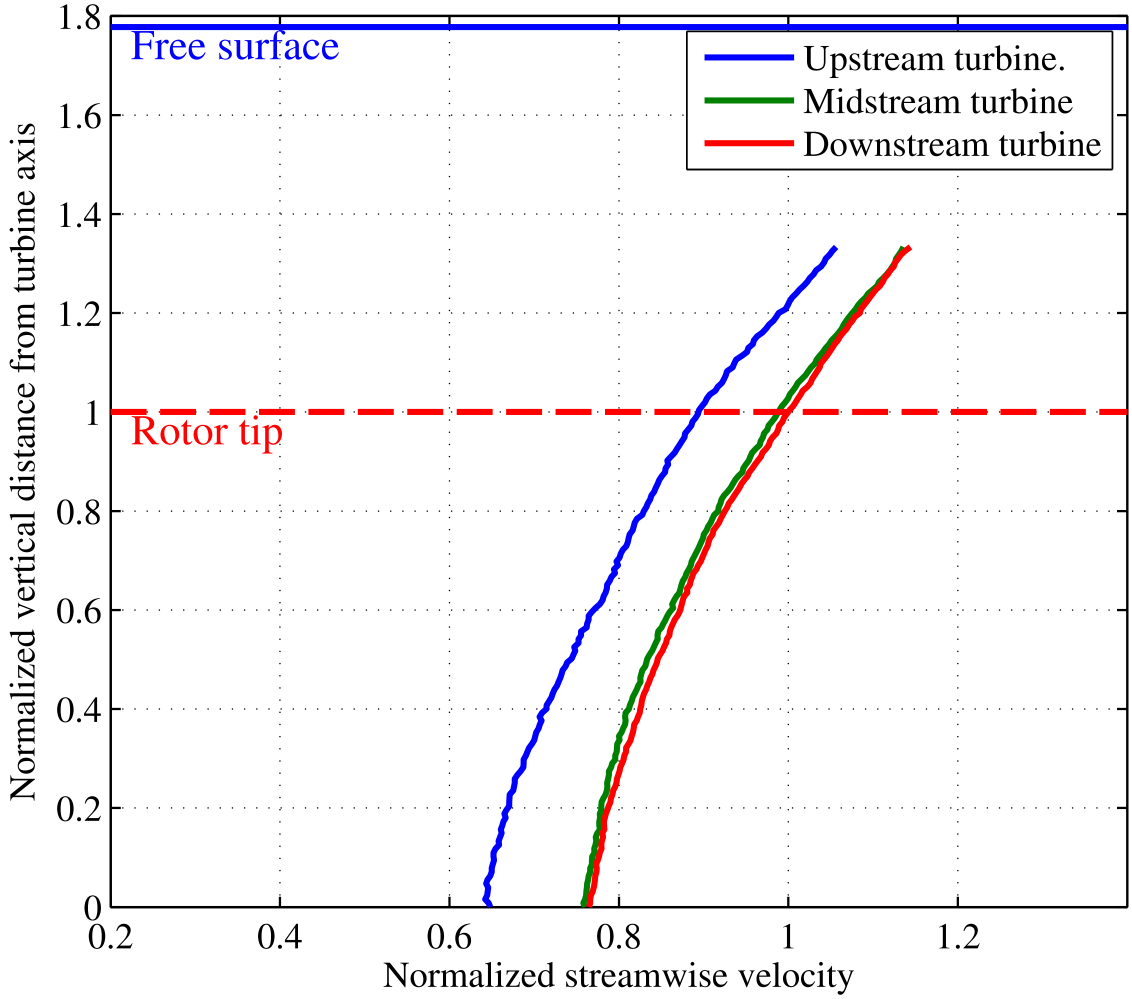

Figure 5.24: Comparison of the mean streamwise velocity profiles 5D downstream of each turbine in a three turbine co-axial array separated by 7D.

5.4 Laterally-offset turbine arrays

Two configurations of laterally-offset three-turbine arrays were tested; a front view of these configurations is shown in Figure 4.4. The layout of these configurations is described in section 4.2.2. Figures 5.25 and 5.26 show the performance of the arrays with spacings of 0.25 D lateral/5D streamwise and 0.25D lateral/7D streamwise, respectively.

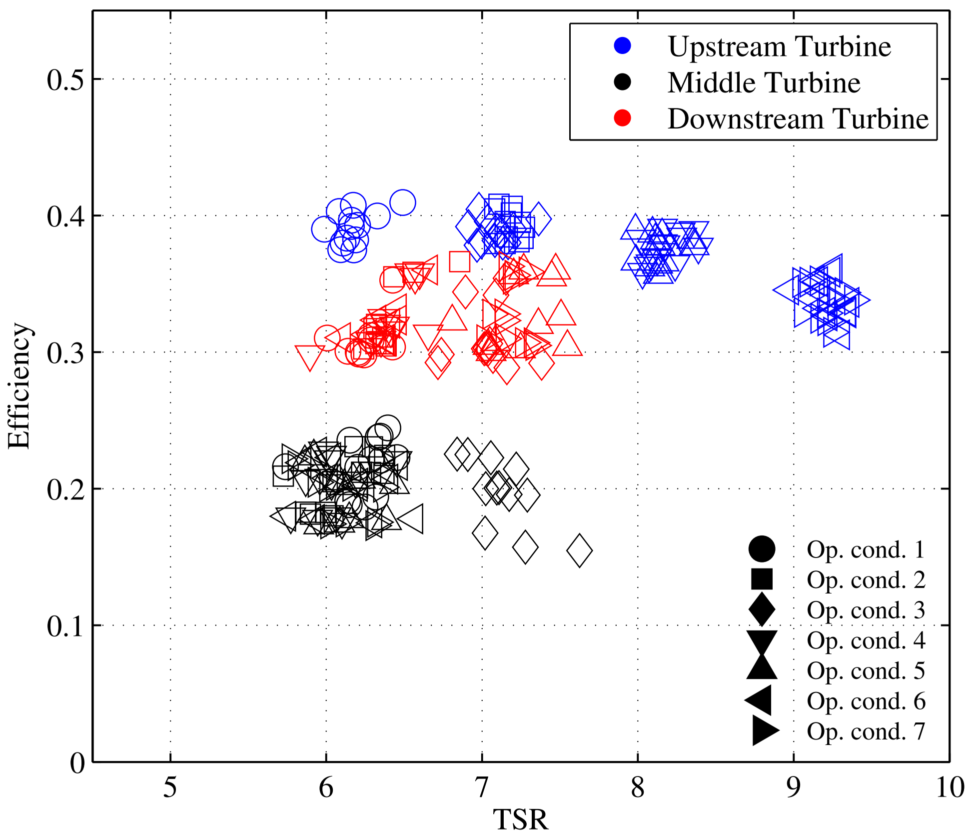

Figure 5.25: Performance curves for three turbines separated by 5 rotor diameters in the streamwise direction and 0.25 rotor diameters in the transverse direction. Different markers represent different simultaneous tests at various combinations of tip speed ratio.

Turbine arrays with lateral offsets, in the context of a high blockage ratio flume, direct the bypass flow towards the downstream turbines. As a consequence the midstream turbine and especially the downstream turbine are operating more efficiently than co-axial arrays with similar streamwise separation distances. More variation is seen in turbine performance in offset arrays than in co-axially arrays, possibly reflecting the greater horizontal shear in velocity.

The closeness to the walls and the impact of blade-tip vortices on downstream turbines cause a much higher variability in the performance of the second (midstream) and third (downstream-most) turbines, compared to the coaxial array cases. Generally, the turbine array is more sensitive to highly unsteady flow near the flume walls and on the shear layer around the turbine cylinder. It is also subject to a higher effective blockage ratio that makes the power production more sensitive to TSR values than shown in the coaxial cases.

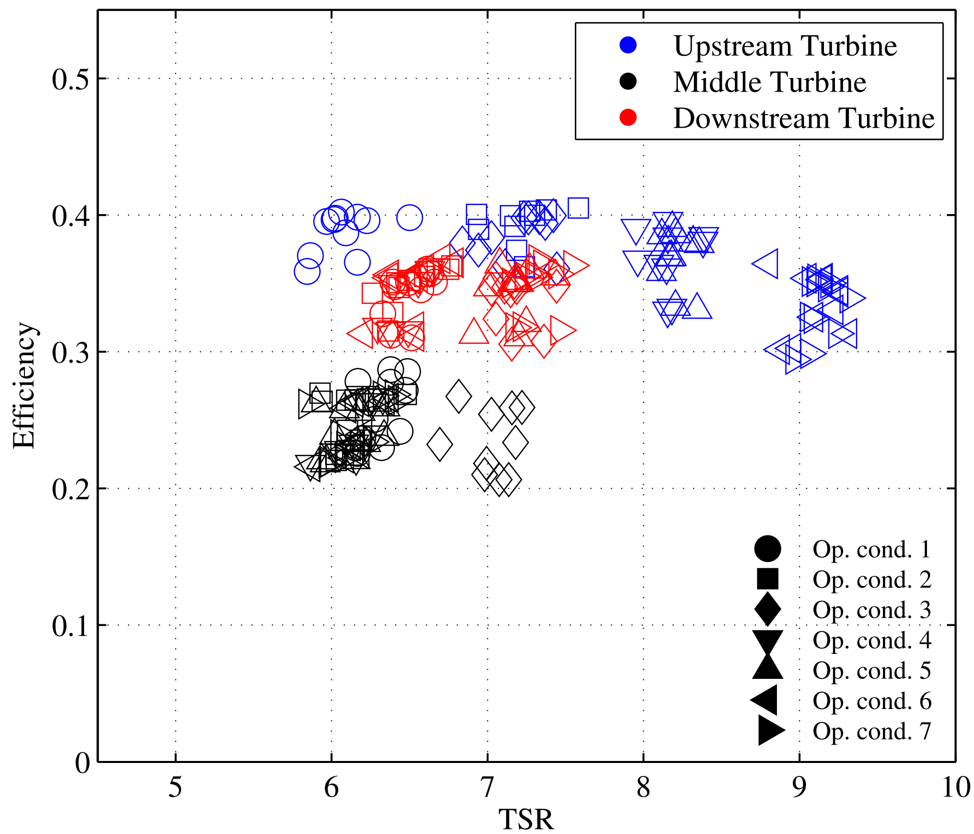

Figure 5.26: Performance curves for three turbines separated by 7 rotor diameters in the streamwise direction and 0.25 rotor diameters in the transverse direction. Different markers represent different simultaneous tests at various combinations of tip speed ratio.

Chapter 6

CONCLUSIONS

This section presents the conclusions of this thesis with regard to the experimental investigation of laboratory-scale tidal turbines. Trends in the experimental results are briefly summarized, and directions for future work are identified.

6.1 Laboratory-scale turbine design

Three scale-model tidal turbines were designed and manufactured. These turbines performed well, and are available for future experiments.

6.1.1 Rotor design

A laboratory-scale rotor was designed with Harp_opt [17] for the purpose of matching the efficiency and $TSR_{optimum}$ of the full-scale Department of Energy Reference Model 1. The performance of this rotor as tested was similar (though lower) than that predicted by Harp_opt, illustrating the efficacy of Harp_opt as a design tool for this application. Comparison of the experimental and numerical curves indicates that there may be a stall-delay effect on the rotor for which Harp_opt does not account. Quantifying this stall-delay may be important for full-scale turbine design optimization.

6.1.2 Nacelle design

An instrumented nacelle was designed and manufactured, and acted as an effective test bed for the laboratory-scale rotor. This nacelle has many advantages, including a method of attaching to the bottom of the flume to avoid piercing the free surface, a system of seals that do not add to the unmeasured parasitic drag, a waterproof housing that allows the in-situ inspection of the nacelle instrumentation, and a streamlined shape. The blunt back end of the nacelle, where the 7° tapering was truncated proved to create a not-insignificant separation and contributed to the wake mixing.

6.2 Trends in results

6.2.1 Turbine performance

The performance of a single turbine was characterized over a range of tip speed ratios. The peak performance of the laboratory-scale turbine is similar to that of a full-scale turbine, indicating that laboratory-scale experimental results can be a significant tool in elucidating trends relevant to full-scale tidal turbine performance and wake development. The experimental performance curve was found to have a higher than expected performance at low TSR, possibly due to a stall-delay effect in the experiment created by the unsteady rotational speed.

A Reynolds number dependence of the turbine performance was shown, and the results agree well with the literature regarding the transition Reynolds number phenomenon associated with laminar separation bubble reattachment. The Reynolds number effect was observed for freestream velocities (up to 0.7 m/s) well below typical flume operating conditions for the experiments reported here (1.1 m/s), indicating that the Reynolds number effect on foil performance does not play a role for these experiments.

The performance of two co-axially spaced turbines was measured, and the performance of the downstream turbine was found to be approximated well by a linear function of the separation distance between the turbines. The tip speed ratio of the upstream turbine had no observable effect on the performance of the downstream turbine.

Four different configurations of three-turbine arrays were tested, and for all configurations a lower performance was observed for the midstream turbine than the downstream turbine. This suggests that an optimized array may require an uneven turbine spacing scheme, e. g. moving the midstream turbine closer to the downstream turbine in order to equalize their performances. The performance curves for the downstream turbines were flatter than for the upstream turbine, i. e. the efficiency was not as dependent on tip speed ratio. The reason for this is not clear, although it may be related to the highly turbulent incident flow and the higher amplitude fluctuations in rotational speed of the downstream turbines.

6.2.2 Wake development

Wake development was characterized by mean streamwise velocity and turbulence intensity profiles. Differences in velocity and turbulence intensity profiles in the near wake were observed for different turbine tip speed ratios, but these differences largely disappeared as the wake developed, with the profiles collapsing onto a single curve at three rotor diameters downstream. The wake of downstream turbines recovers more quickly than the wake of a single turbine (or of the front turbine in the array), which was expected due to the effect of turbine-enhanced turbulence on momentum diffusion. This increase in wake recovery does not have a linear dependence on turbine spacing in the way that downstream performance does, for example the wake recovery of the downstream turbine is very similar for different spacings of two-turbine array.

A comparison of the velocity profiles in a three-turbine array showed that the velocity profiles in the wake of the midstream turbine were very similar to those of the downstream turbine. This may have implications for large array optimizations, as it suggests a useful relationship between wake recovery and turbine spacing.

6.3 Comparisons with numerical simulations

Work is currently ongoing comparing the results presented here with results from numerical simulations.

It is the intention of the author to make the experimental dataset that was generated by this project to be freely available for the validation of numerical models. To that end, the work of organizing and presenting this dataset in an web-accesible database is ongoing.

6.4 Directions for future work

6.4.1 Further analysis

Many more insights than the ones presented in this thesis can be extracted from the experimental dataset. Mean values of velocity and turbulence intensity are only two of the flow properties that may be relevant for turbine performance. Turbulent length scales, turbulent kinetic energy, and Reynolds stresses can be derived from the PIV data will be investigated in future work. Time-series analysis can also be performed to understand the correlation between flow properties, turbine rotational speed, and turbine performance.

Measurements of free surface deformation were performed at various turbine rotor proximities to the free surface, although this was not discussed in this thesis. Future work will include analysis of these data, and comparing the results to numerical simulations.

6.4.2 Future experiments

The scale-model turbines created for these experiments are available for further testing. Future work should include testing at other facilities that have different cross-sectional areas and freestream flow characteristics. These tests would show the effect of blockage on turbine performance and wake development. They could also investigate how the turbine and wake properties differ in lower turbulence intensity inflow conditions.

Based on the results presented here, future tests can reduce the size of the test matrix by reducing the number of combinations of TSR that are tested. These results show that wake development and turbine array efficiency has little dependence on TSR. Instead of testing many combinations of TSR, future tests can make more detailed investigations of the flow field. Using PIV to explore the bottom half of the water column, off-axis vertical planes, and horizontal planes in the wake would deepen the understanding of wake development and wake interaction.

BIBLIOGRAPHY

[1] A.S. Bahaj, W.M.J. Batten, and G. McCann. Experimental verifications of numerical predictions for the hydrodynamic performance of horizontal axis marine current turbines. *Renewable Energy*, 32(15):2479 – 2490, 2007.

[2] A.S. Bahaj, A.F. Molland, J.R. Chaplin, and W.M.J. Batten. Power and thrust measurements of marine current turbines under various hydrodynamic flow conditions in a cavitation tunnel and a towing tank. *Renewable Energy*, 32(3):407 – 426, 2007.

[3] Leonardo P. Chamorro, Daniel R. Troolin, Seung-Jae Lee, R.E.A. Arndt, and Fotis Sotiropoulos. Three-dimensional flow visualization in the wake of a miniature axial-flow hydrokinetic turbine. *Experiments in Fluids*, 54(2):1–12, 2013.

[4] M. Drela and H. Youngren. Xfoil 6.9 user primer. Technical report, Massachusetts Institute of Technology.

[5] A. T. Javaherchi, N. Stelzenmuller, and A. Aliseda. Experimental and numerical analysis of the doe reference model 1 horizontal axis hydrokinetic turbine. In *Proceedings of the 1st Marine Energy Technology Symposium*, 2013.

[6] Christian Jonsson, Peter B Johnson, and Ian Eames. Energy extractors in turbulent flow: Wake decay and implications for farm layout. In *Proceedings of the 9th European Wave and Tidal Energy Conference (EWTEC)*, 2011.

[7] M. Kawase and K.M. Thyng. Three-dimensional hydrodynamic modelling of inland marine waters of washington state, united states, for tidal resource and environmental impact assessment. *Renewable Power Generation, IET*, 4(6):568–578, 2010.

[8] M. Lawson, Y. Li, and D. Sale. Development and verification of a computational fluid dynamics model of a horizontal-axis tidal current turbine. In *Proceedings of the 30th International Conference on Ocean, Offshore, and Arctic Engineering*, 2011.

[9] P.B.S. Lissaman. Low-reynolds-number airfoils.

[10] F. Maganga, G. Germain, J. King, G. Pinon, and E. Rivoalen. Experimental characterisation of flow effects on marine current turbine behaviour and on its wake properties. *Renewable Power Generation, IET*, 4(6):498–509, 2010.

[11] P. Mycek, Gaurier B., G. Pinon, and E. Rivoalen. Numerical and experimental study of the interaction between two marine current turbines. In *Proceedings of the 9th European Wave and Tidal Energy Conference*, 2011.

[12] LE Myers, B Keogh, and AS Bahaj. Layout optimisation of 1 st-generation tidal energy arrays. In *European Wave and Tidal Energy Conference*, 2011.

[13] Kai-Wern Ng, Wei-Haur Lam, and Khai-Ching Ng. 20022012: 10 years of research progress in horizontal-axis marine current turbines. *Energies*, 6(3):1497–1526, 2013.

[14] T. O'Doherty, A. Mason-Jones, D.M. ODoherty, C.B. Byrne, I. Owen, and Y. Wang. Experimental and computational analysis of a model horizontal axis tidal turbine. In *8th European Wave and Tidal Energy Conference (EWTEC), Uppsala, Sweden*, 2009.

[15] A. D. Platt and M.L. Jr. Buhl. Wt_perf user guide for version 3.05.00. Technical report, National Renewable Energy Laboratory.

[16] B. Polagye, B. Van Cleve, A. Copping, and K. Kirkendall. Environmental effects of tidal energy development. NOAA Tech. Memo NMFS F/SPO-116, U.S. Dept. Commerce, http://ir.library.oregonstate.edu/xmlui/handle/1957/21617, 2011.

[17] D. Sale. Harp_opt user's guide. Technical report, National Renewable Energy Laboratory.

[18] T. Stallard, R. Collings, T. Feng, and J.I. Whelan. Interactions between tidal turbine wakes: experimental study of a group of 3-bladed rotors. In *Proceedings of the 9th European Wave and Tidal Energy Conference*, 2011.

[19] S.C. Tedds, I. Owen, and R.J. Poole. Near-wake characteristics of a model horizontal axis tidal stream turbine. *Renewable Energy*, 63(0):222 – 235, 2014.

[20] J.I. Whelan and T. Stallard. Arguments for modifying the geometry of a scale model rotor. In *Proceedings of the 9th European Wave and Tidal Energy Conference*, 2011.

Appendix A

DERIVATION OF ROTATIONAL SPEED

A.1 Background

We would like to measure the angular velocity of a rotating shaft. The instrument used to measure this angular velocity is an incremental encoder connected to a data acquisition system (DAQ). The encoder sends a voltage pulse to the DAQ every time the encoder is rotated a certain fraction of a revolution. For example, the resolution of the encoders used in the MHK turbine testing was 256 pulses/revolution, so each time the encoder rotates 256/360 degrees the encoder sends a pulse to the DAQ. The counter of the DAQ reads the pulses into Labview, which translates them into position data (i.e. 0-360 degrees). Labview samples the DAQ at a certain frequency, which in the case of the MHK experiments was 1000 Hz.

So we have 1000 position data points per second (in degrees). A snippet of the data is shown below.

Timestamp	**Angular Position (degrees)**
07/13/2013 09:58:52.036 AM	348.3984375
07/13/2013 09:58:52.037 AM	350.15625
07/13/2013 09:58:52.038 AM	351.9140625
07/13/2013 09:58:52.039 AM	354.0234375
07/13/2013 09:58:52.040 AM	355.78125
07/13/2013 09:58:52.041 AM	357.890625
07/13/2013 09:58:52.042 AM	359.6484375
07/13/2013 09:58:52.043 AM	1.40625
07/13/2013 09:58:52.044 AM	3.1640625
07/13/2013 09:58:52.045 AM	5.2734375
07/13/2013 09:58:52.046 AM	7.03125
07/13/2013 09:58:52.047 AM	8.7890625

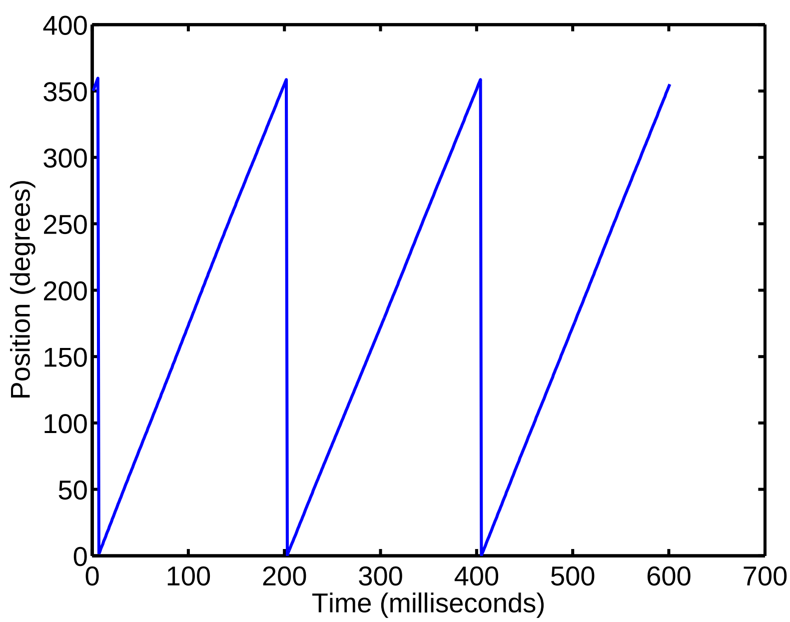

Figure A.1: Small snippet of raw position data

A simple numerical derivative of this position data will produce rotational velocity

$$\frac{Position_{n+1} - Position_n}{\Delta time}$$

Because the angular position data is modulus 360 the data processing code needs to recognize and effectively process the rotational velocity points that are calculated across the zero point, for example:

$$\frac{1.1deg - 359.2deg}{\Delta t} = \frac{-358.1deg}{\Delta t}$$

which is clearly incorrect. Currently this is handled in the data processing code by replacing all negative velocity data with NaN's, on the assumption that the rotor is always turning in the same direction.

Figure A.2 shows an example of the raw velocity data that was calulated in this way. The raw velocity data shows three distinct bands—a primary band at approximately 470 rpm and two secondary bands at 410 rpm and 530 rpm. A closer look at the relationship between the position data and the sampling rate will help to explain this banding behavior.

For example, imagine a rotating shaft with an average rotational speed of 475 rpm and fluctuations of ±50 rpm. An encoder with a resolution of 256 pulse/rev that was measuring this rotational speed will produce a 2025±213 pulse/sec signal. This signal is then sampled at 1000 Hz.

Figures A.3-A.5 show that we can reproduce this banding effect by simulating a simple sinusoidal velocity signal, numerically integrating that signal to simulate the position signal, discretizing and sampling that position signal in a way that mimics the encoder and DAQ, and using the forward difference method to calculate the velocity. This demonstrates that the banding shown in Figure A.2 is an artifact of using the forward difference method to derive velocity from a position signal that has been discretized

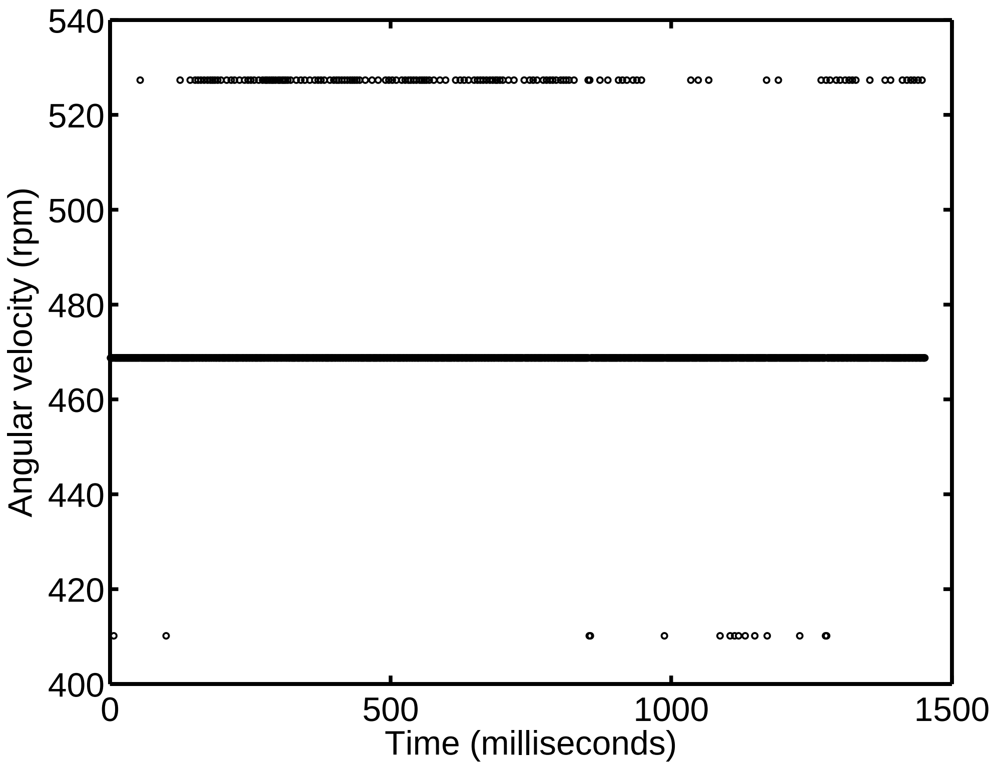

Figure A.2: Small snippet of raw velocity data

and sampled in this manner. Clearly, alternatives to the forward difference method are needed.

Methods of deriving velocity from discretized and sampled position data

The following methods were considered for deriving rotational speed from the position data.

1. A first order forward difference method, as described above.

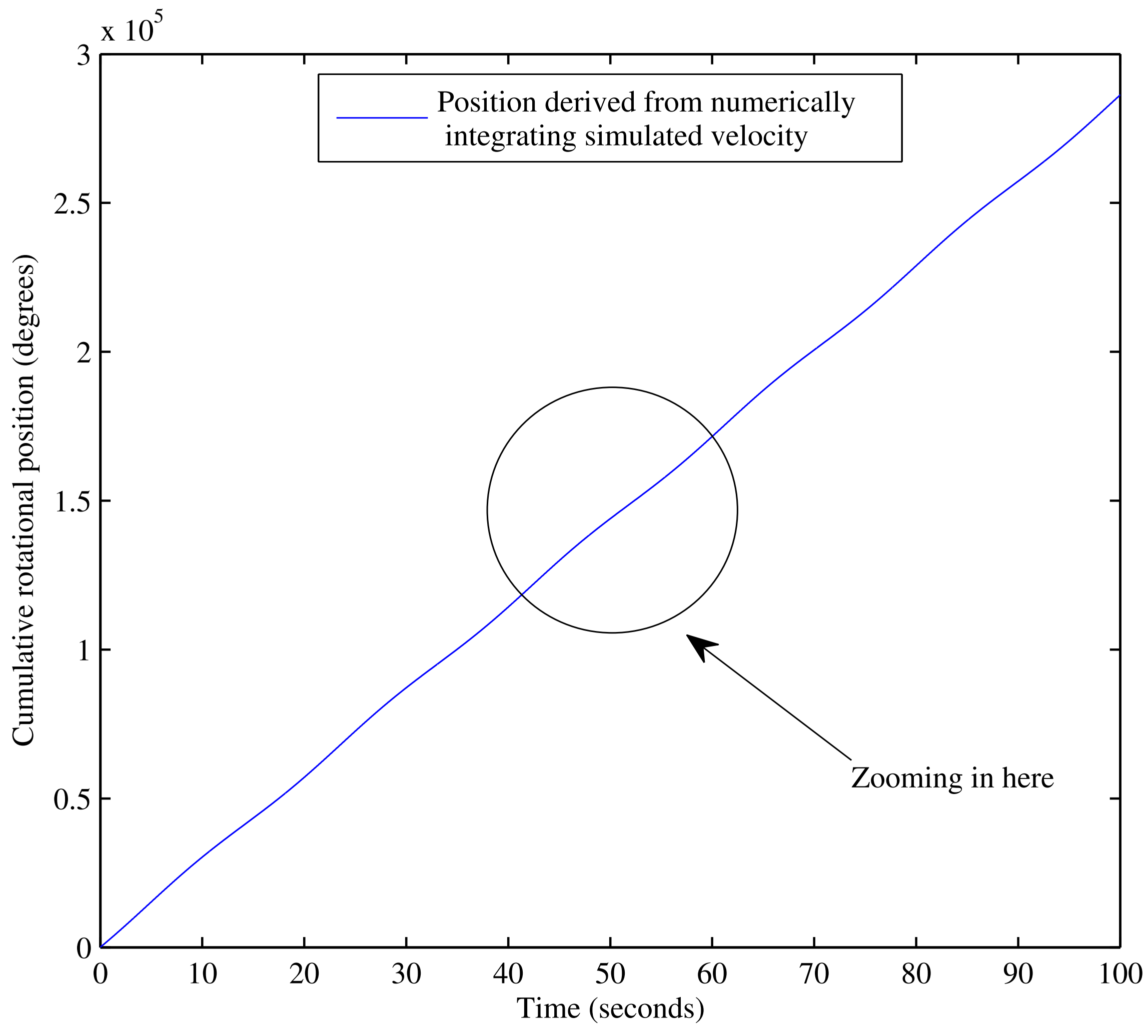

Figure A.3: Example position data taken from sinusoidal velocity

2. A down-sampled forward difference method, i. e. using the forward difference method over a greater time interval.

$$velocity = \frac{pos_{n+10} - pos_n}{10\Delta t}$$

3. A central difference method over N points, e. g. if N=5:

$$\frac{f_{n-2} - 8f_{n-1} + 8f_{n+1} - f_{n+2}}{12\Delta t}$$

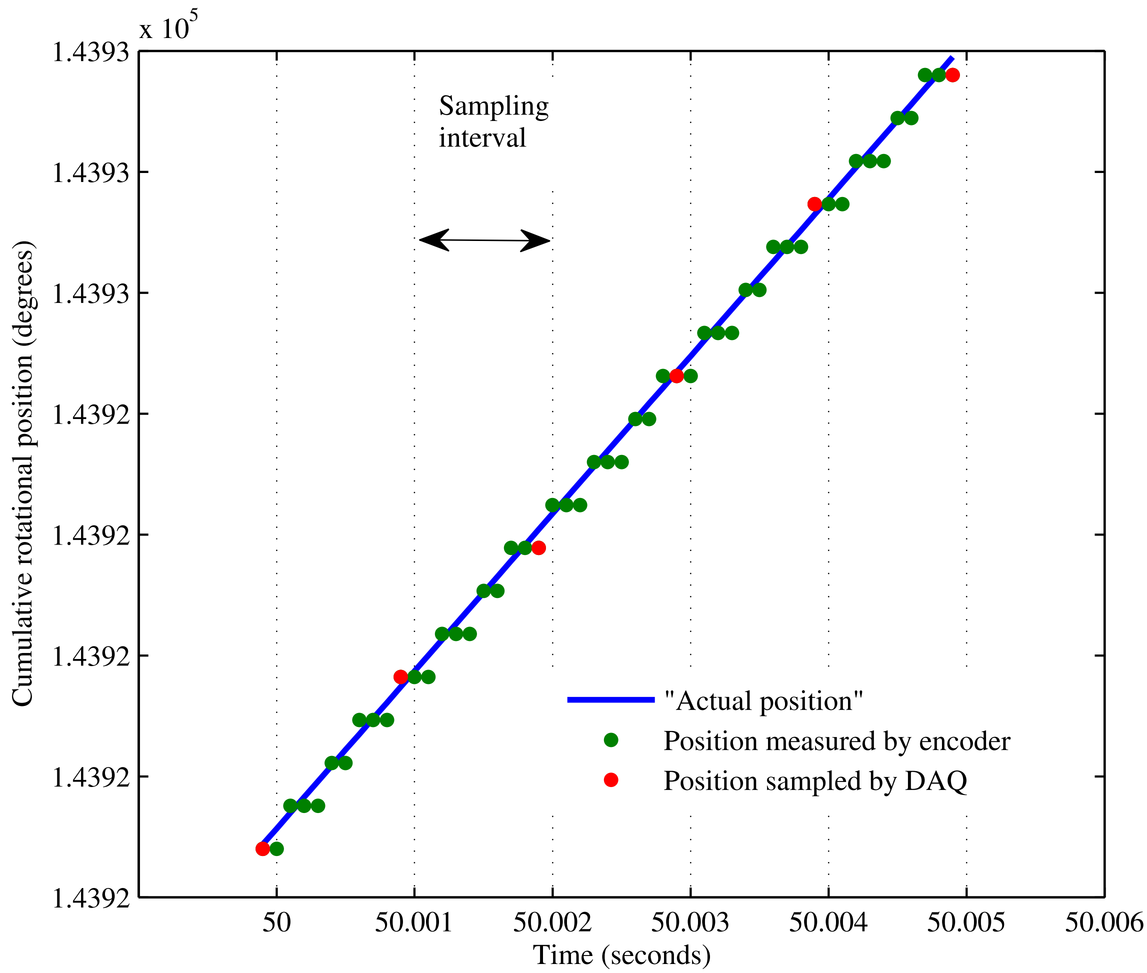

Figure A.4: Small snippet of position data, shown with the example position as measured by the encoder.

4. Using a moving average window to smooth the forward difference method result.
5. Fitting the position data with a B-spline of various sizes, and analytically taking the derivative of the spline function.

We do not have access to the true velocity data, so to compare the various methods of deriving velocity data from position data we can use a test function. The general testing procedure was:

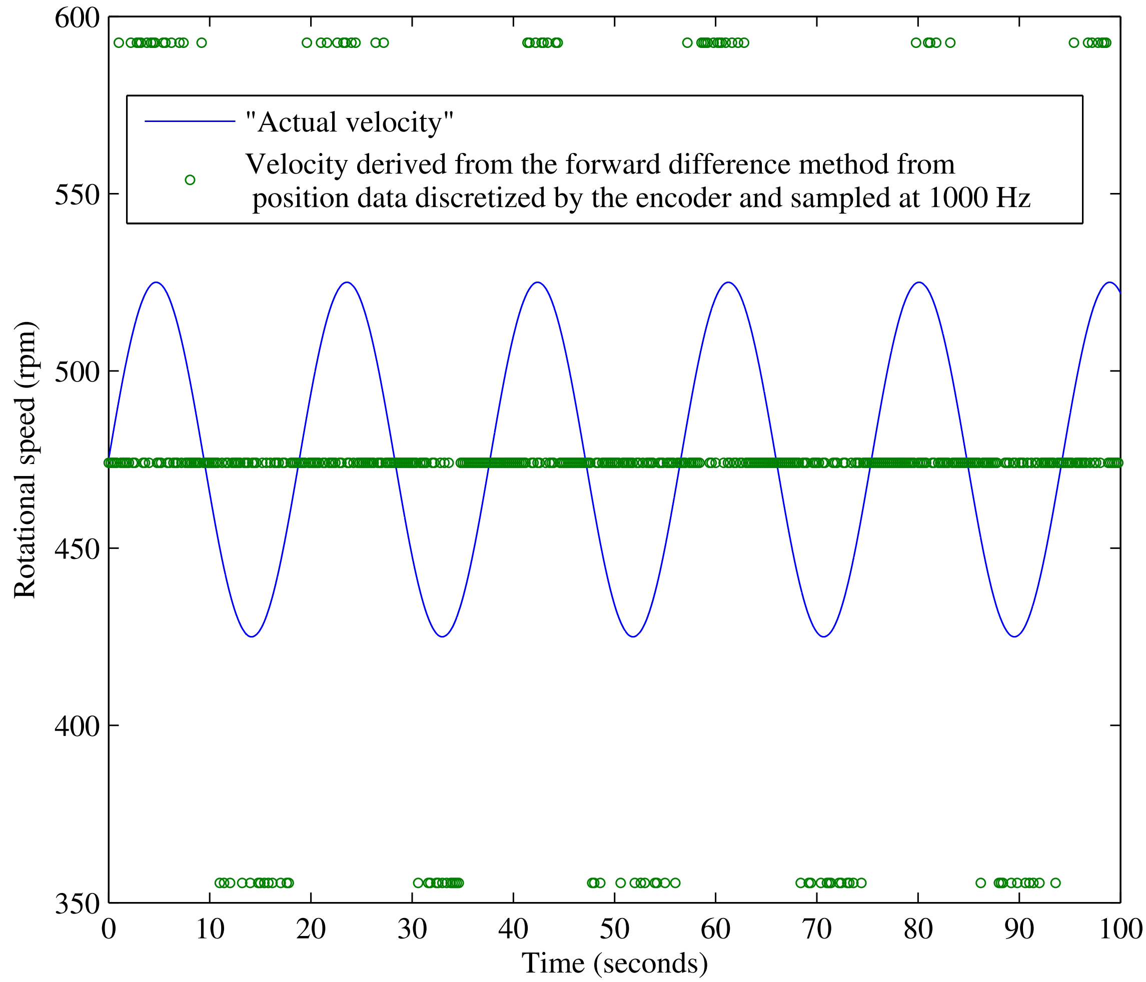

Figure A.5: Example velocity data shown with velocity data that is a result of measuring the example postion with the encoder, sampling that data at 1000 Hz, and taking the first order forward difference (velocity=$\frac{pos_{n+1}-pos_n}{\Delta t}$)

1. Use some combination of analytical curves to simulate velocity data.
2. Numerically integrate this “true” velocity to obtain position data.
3. Simulate the encoder and DAQ by discretizing and sampling the position data as described above.
4. Apply the various methods to the discretized and sampled position.

5. Compare the resulting velocity with the "true" velocity.
6. Compare the errors across the various methods.

The first test function will be the relatively low frequency signal shown in Figure A.6

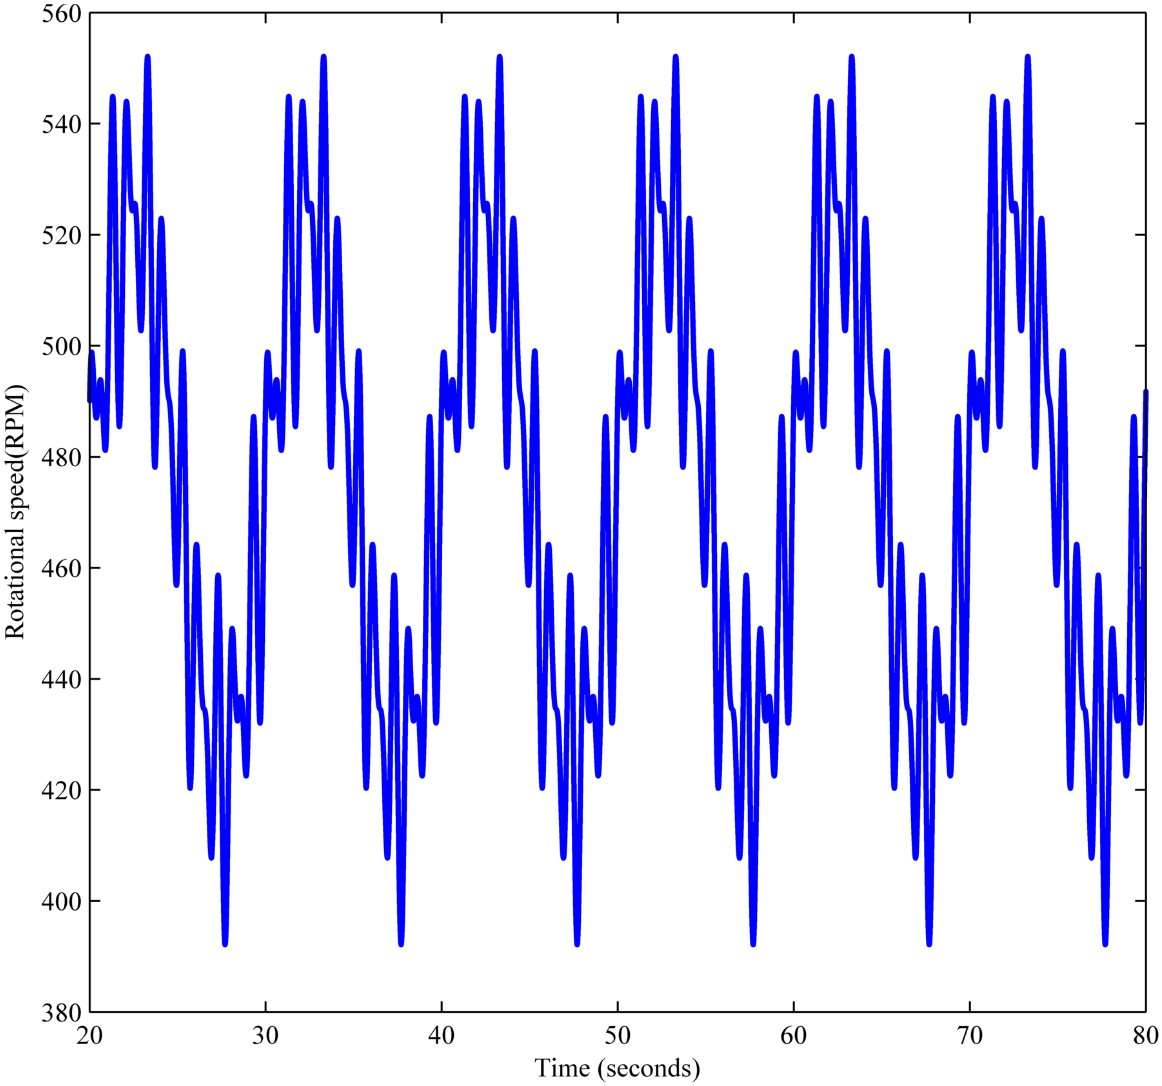

Figure A.6: Low-frequency velocity test function.

Applying the various velocity derivation methods to this test function we can gage how well each method performed. The resulting velocities and errors associated with these methods are plotted in Figures A.7–A.10.

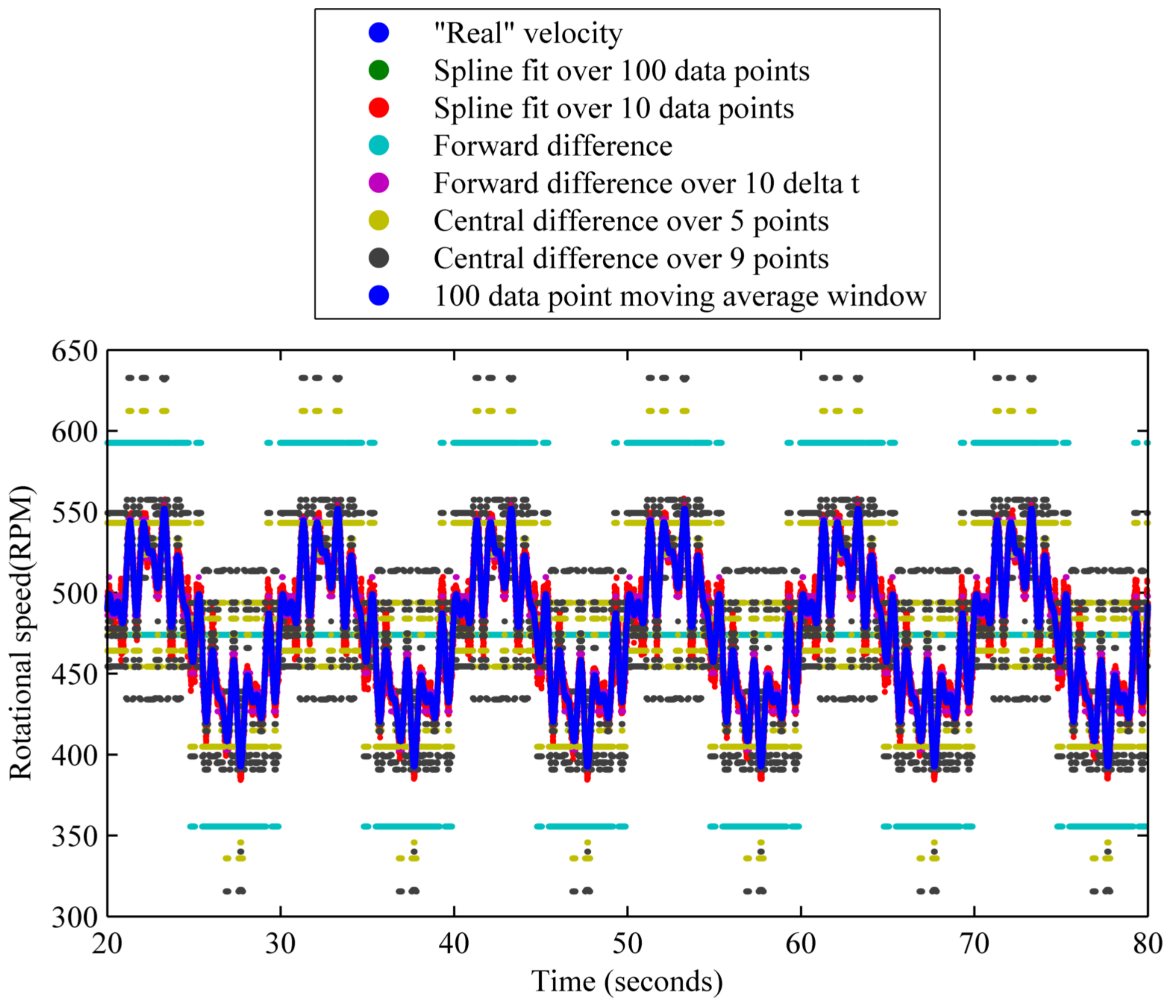

Figure A.7: Velocities derived from the low-frequency test function by various methods.

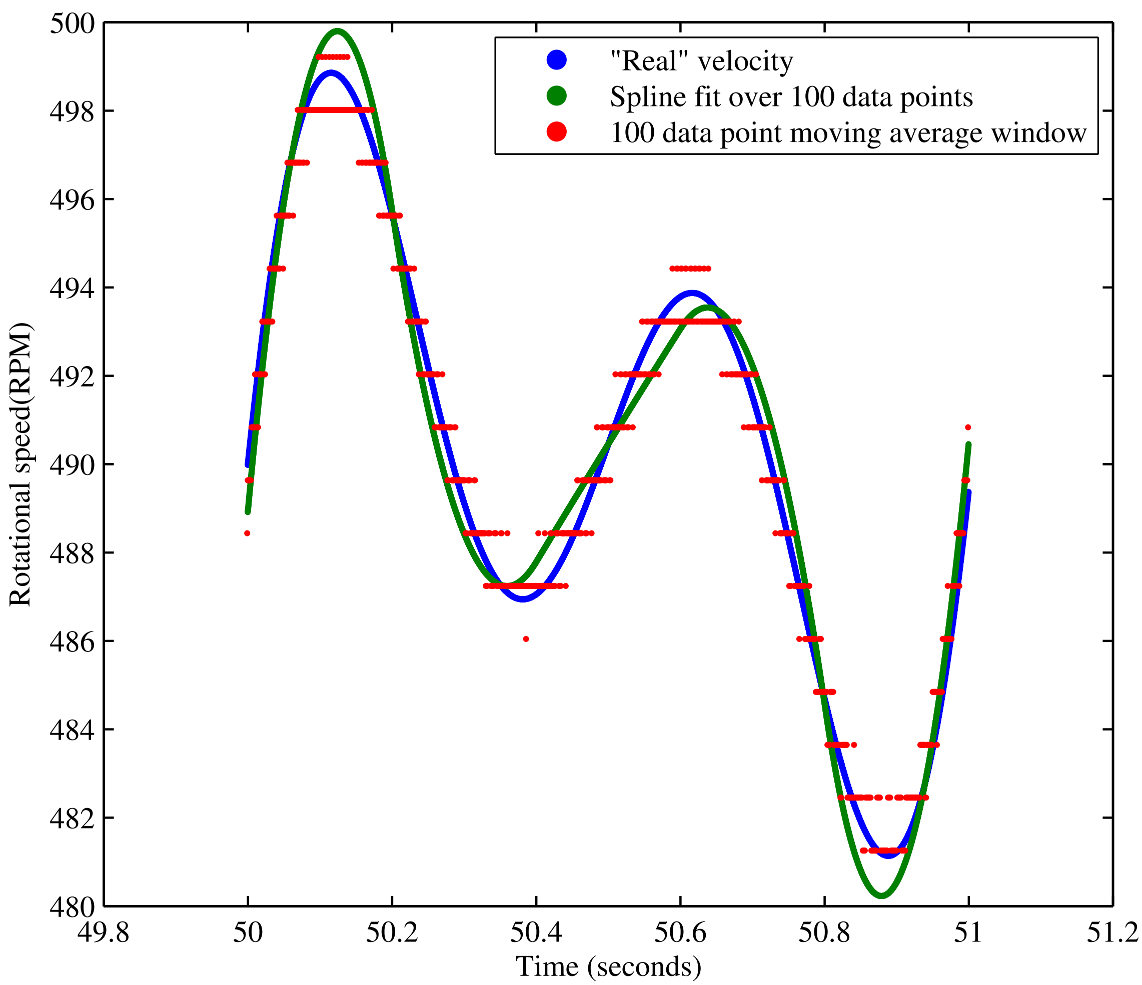

Figure A.8: Detailed view of two of the derived velocities shown in Figure A.7.

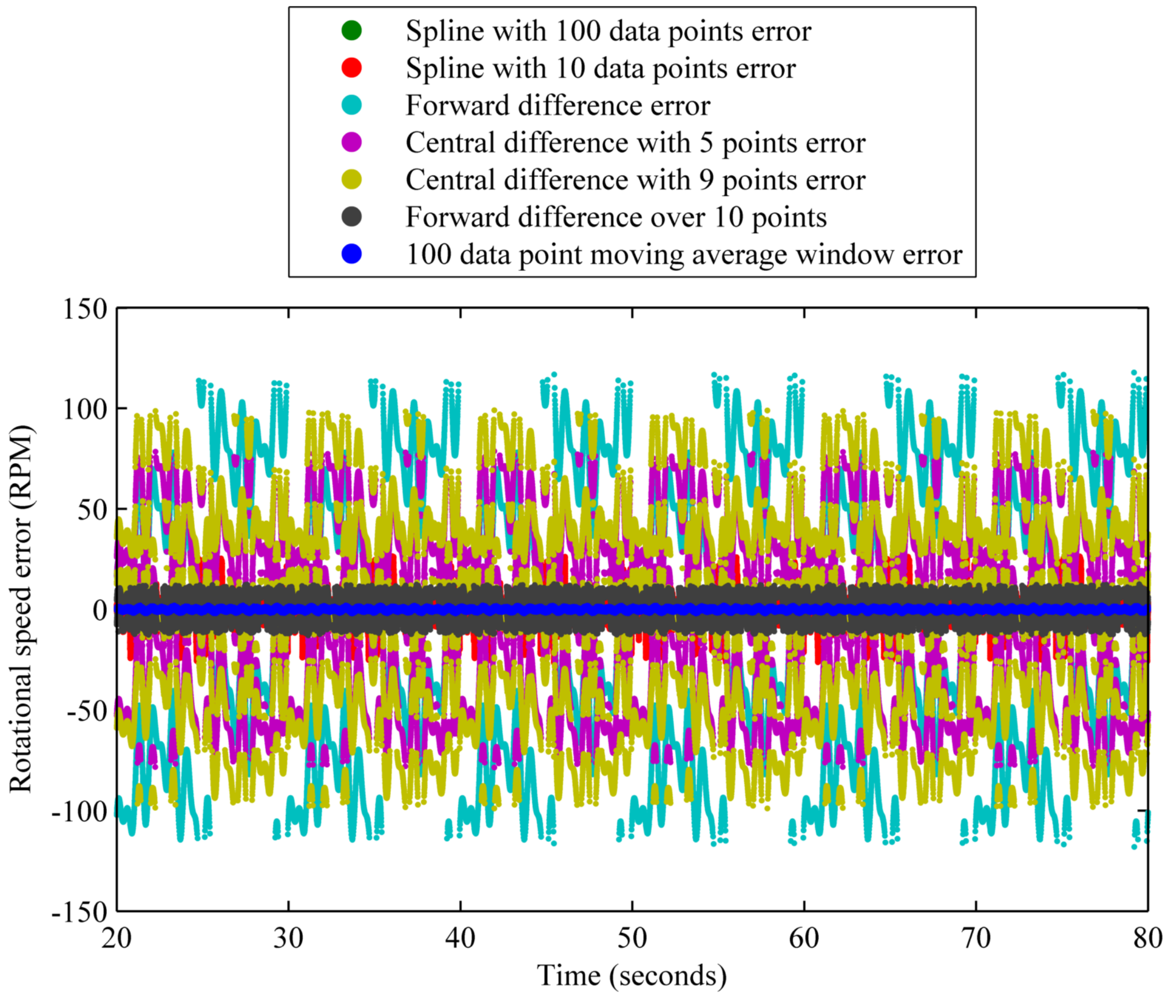

Figure A.9: Error from each of the velocity derivation methods from the low-frequency test signal.

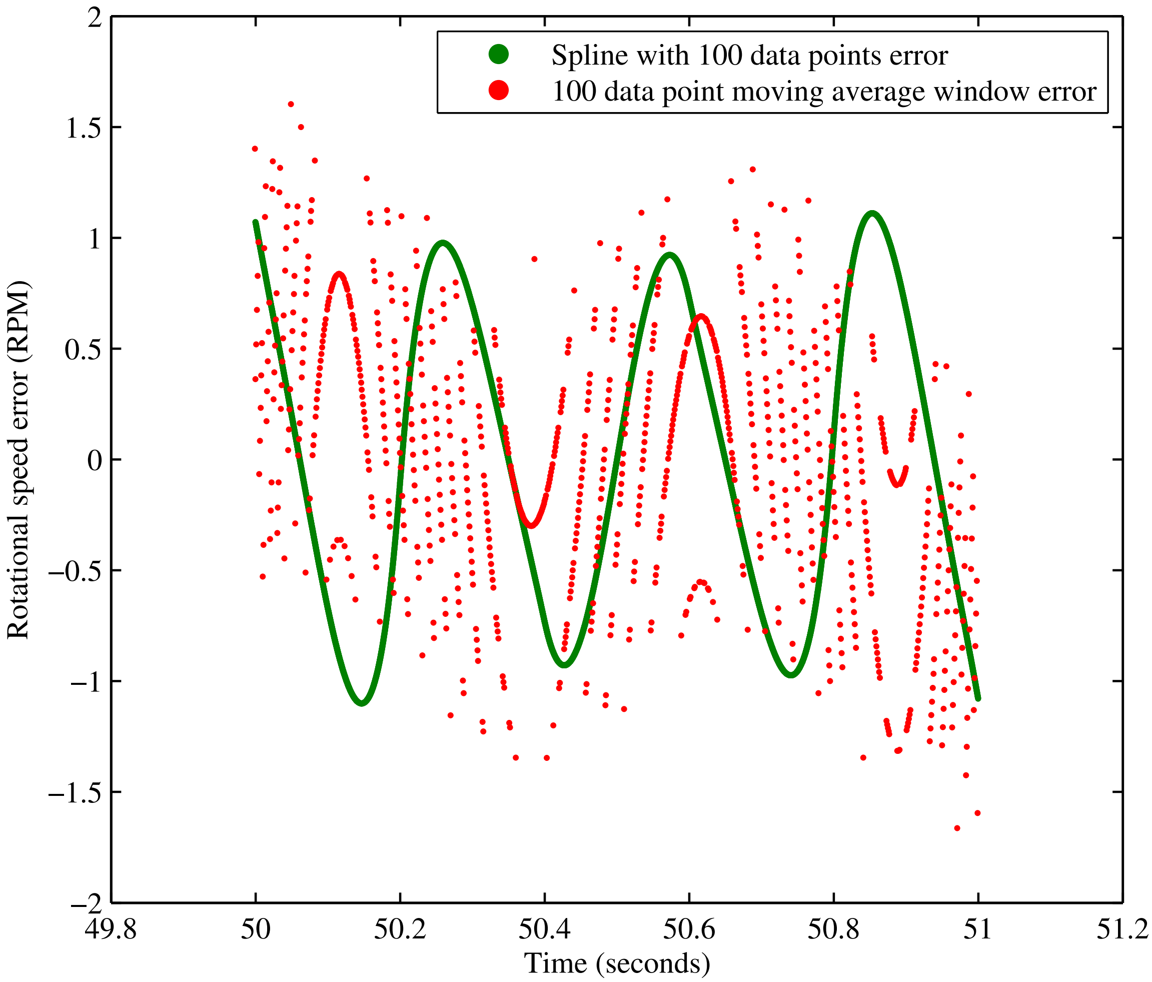

Figure A.10: Detailed view of the errors associated with two of the derived velocities shown in Figure A.9.

above The error associated with the 100-point moving average and the 100-point spline fit are the lowest error methods for this test function. Each shows an error of ~0.2%.

A high-frequency test function was also used to explore the dependence of the velocity derivation methods on signal frequency. Figure A.11 shows the high-frequency test function.

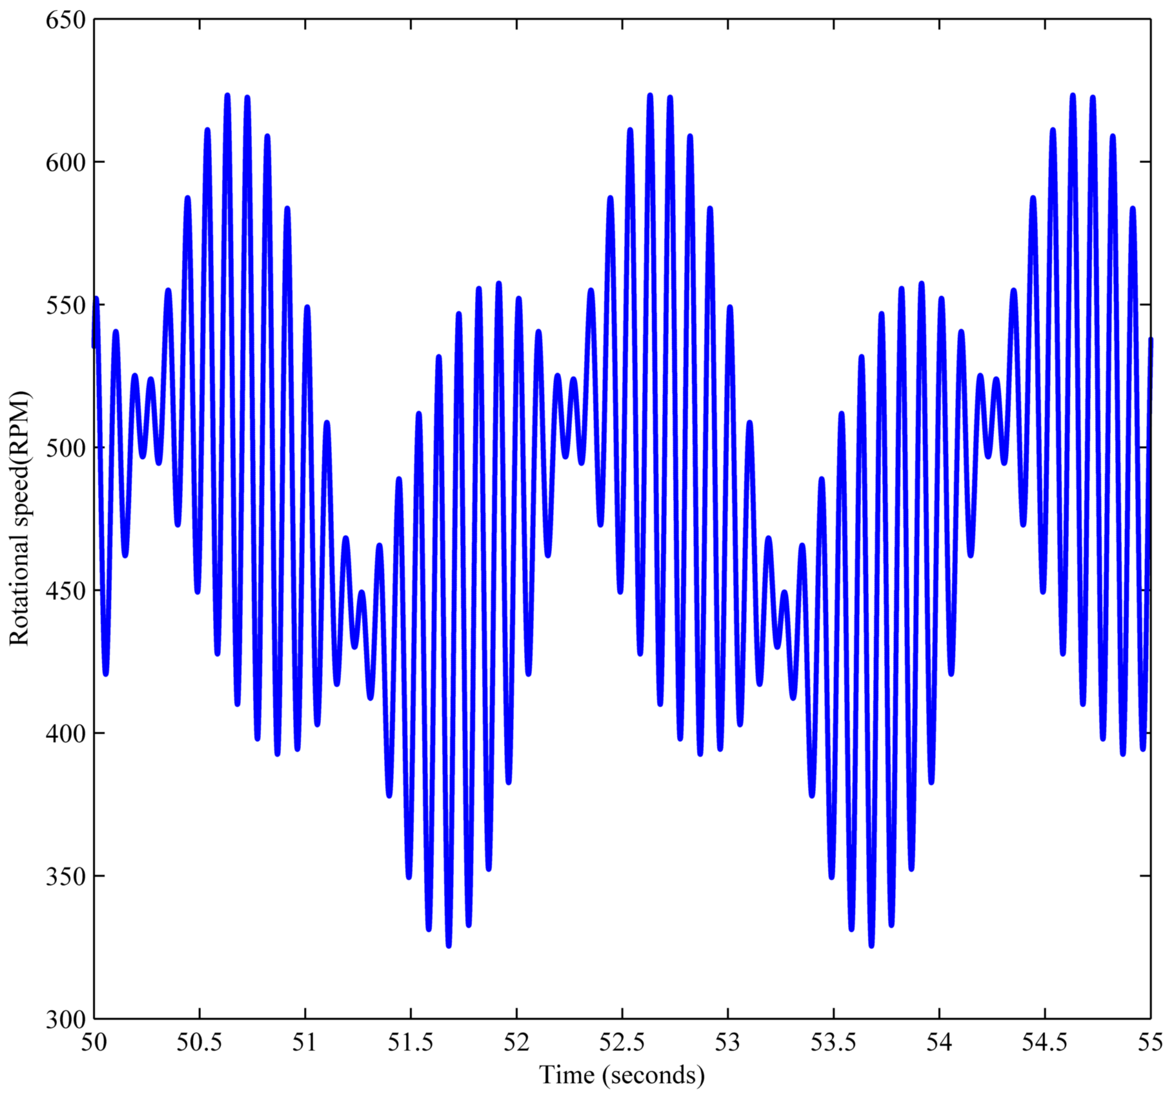

Figure A.11: High-frequency velocity test function.

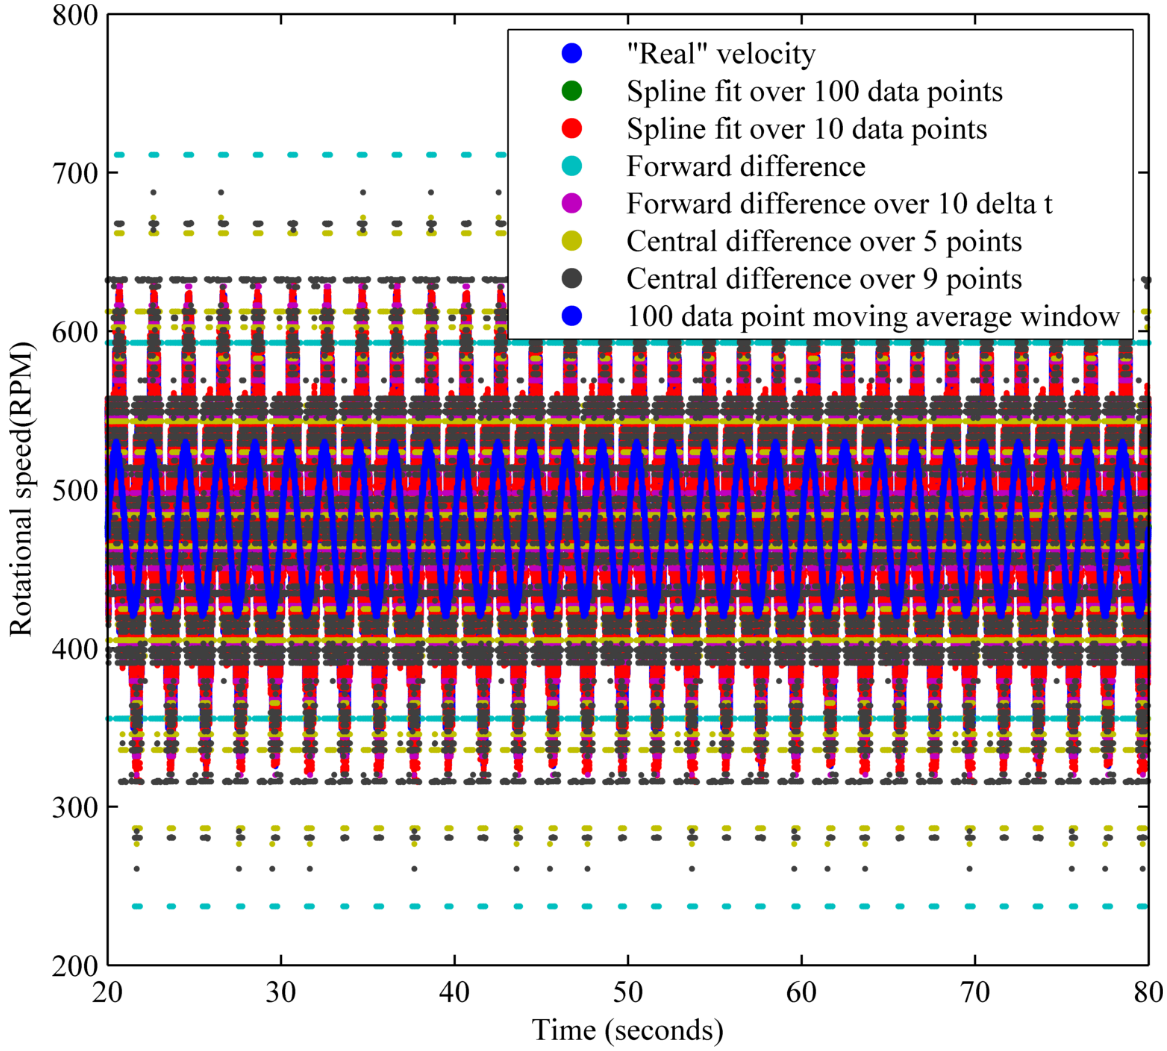

Figure A.12: Velocities derived from the high-frequency test function by various methods.

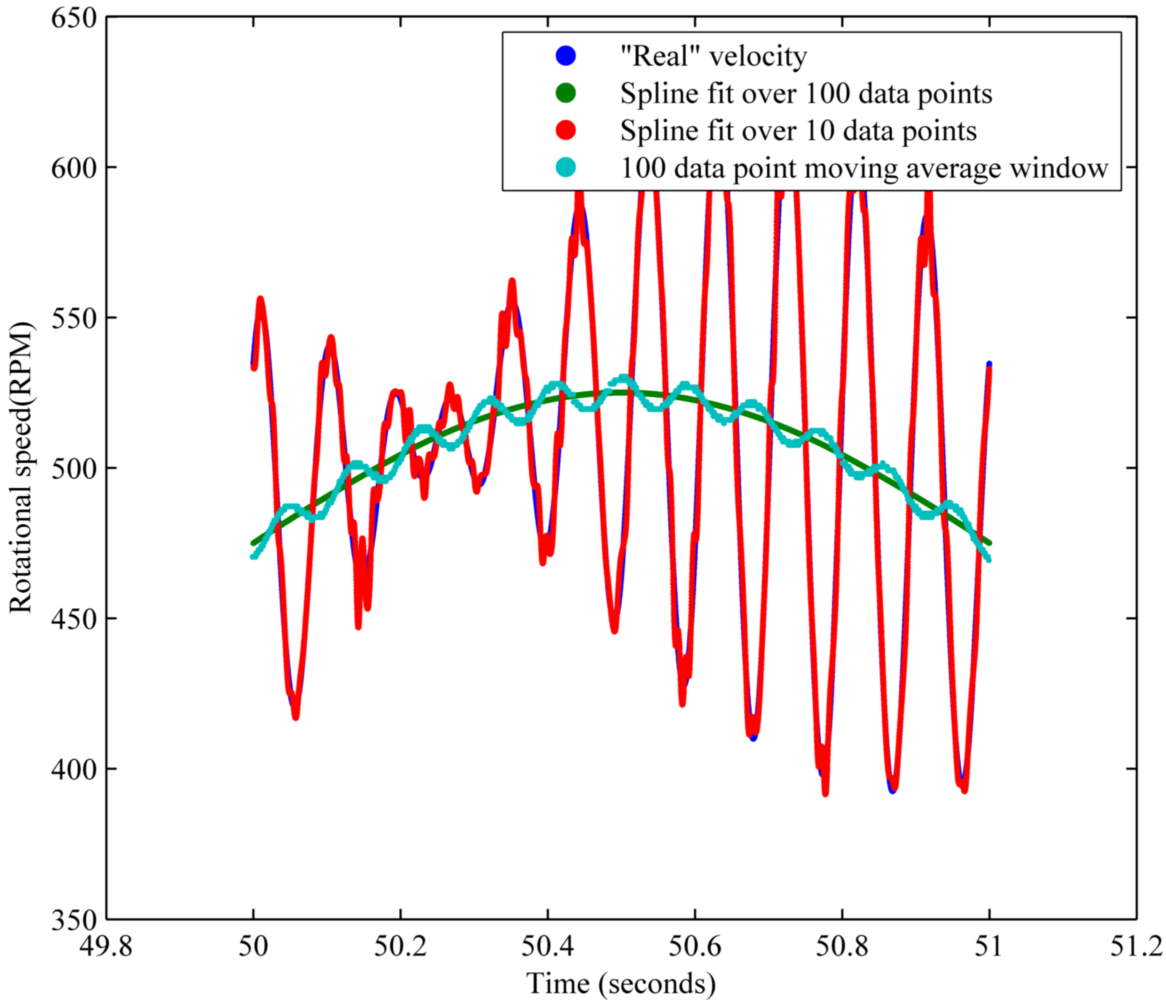

Figure A.13: Detailed view of three of the derived velocities shown in Figure A.12.

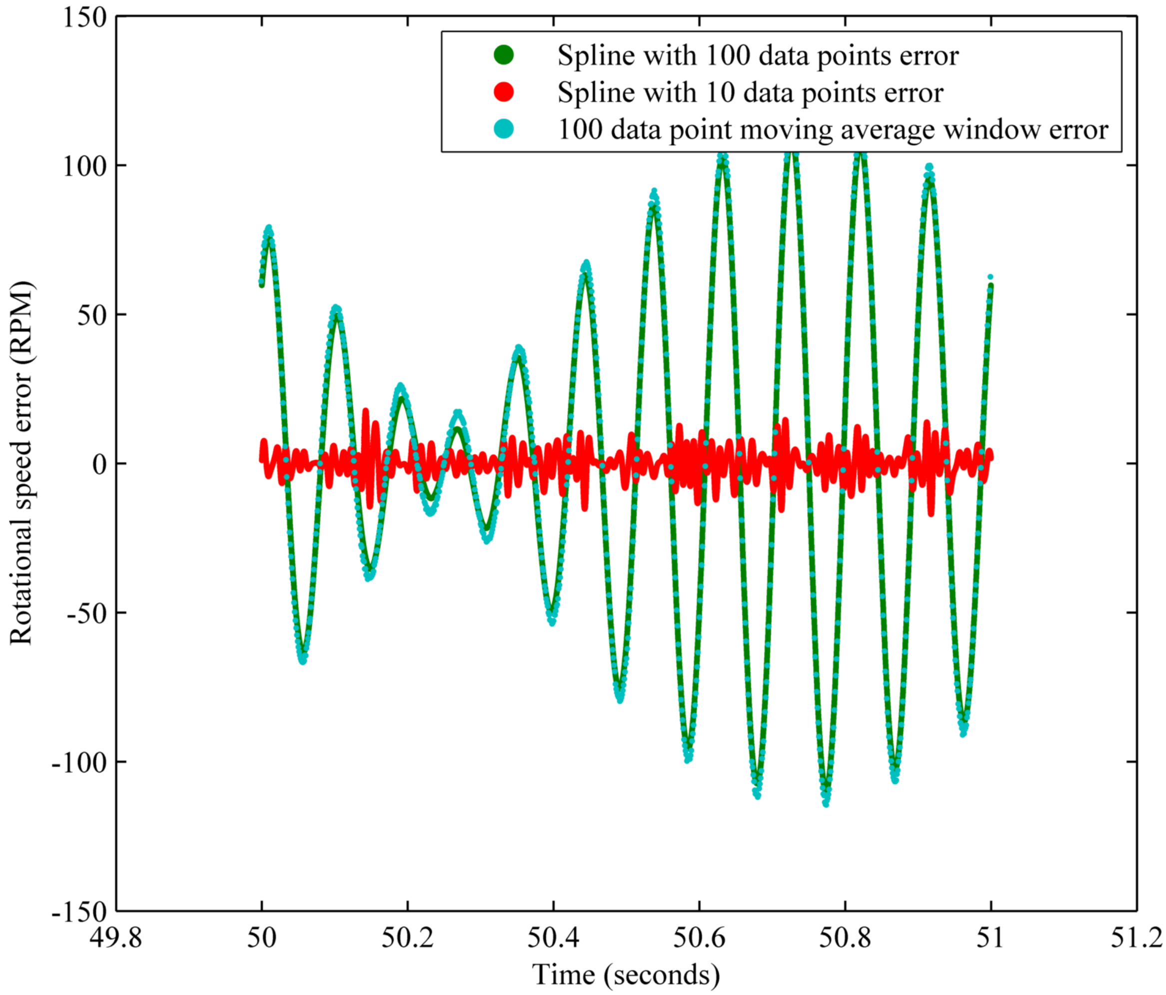

Figure A.14: Detailed view of the errors associated with two of the derived velocities of the high-frequency test function.

A.2 Adaptive spline method

Comparing the errors associated with velocity derivation methods shown in Figures A.9, A.10, and A.14 suggests that the method of fitting splines is the best method for tracking the true velocity. However, for the low-frequency test function the 10-point spline has high error, and the 100-point spline had low error, but for the high-frequency test function the reverse was true: the 10-point spline had low error and the 100-point spline had high error. The dependence of error on frequency for a spline of a given length suggested a method of adaptive spline lengths. This method follows the following algorithm:

1. The position data is put in monotonic form, i. e. adding 360° for each revolution, such that the position vector does not reset to zero when it becomes larger than 360°.
2. Large (e. g. 1,000-point) splines are fitted to the position data.
3. The splines are discretized and sampled in a way that mimics the encoder and DAQ.
4. This discretized and sampled data is compared to the actual position data, and the error is calculated.
5. The length of each spline is halved, and the above process is repeated.
6. The error associated with this reduced spline is compared to the error associated with the previous spline (normalized by length).
7. If the new error is larger, the old spline is replaced, otherwise the spline is halved again.
8. This process iterates until the minimum error per length is achieved.

The intention of this algorithm is to reconstruct the true rotational velocity by proposing a low-frequency model of the true position (the large spline), discretizing and sampling the model and comparing to the discretized and sampled true position, then progressively increasing the frequency of the model until the error is minimized. This algorithm attempts to minimize spurious high frequency models by proposing low-frequency models first, and only increasing the frequency if warranted.

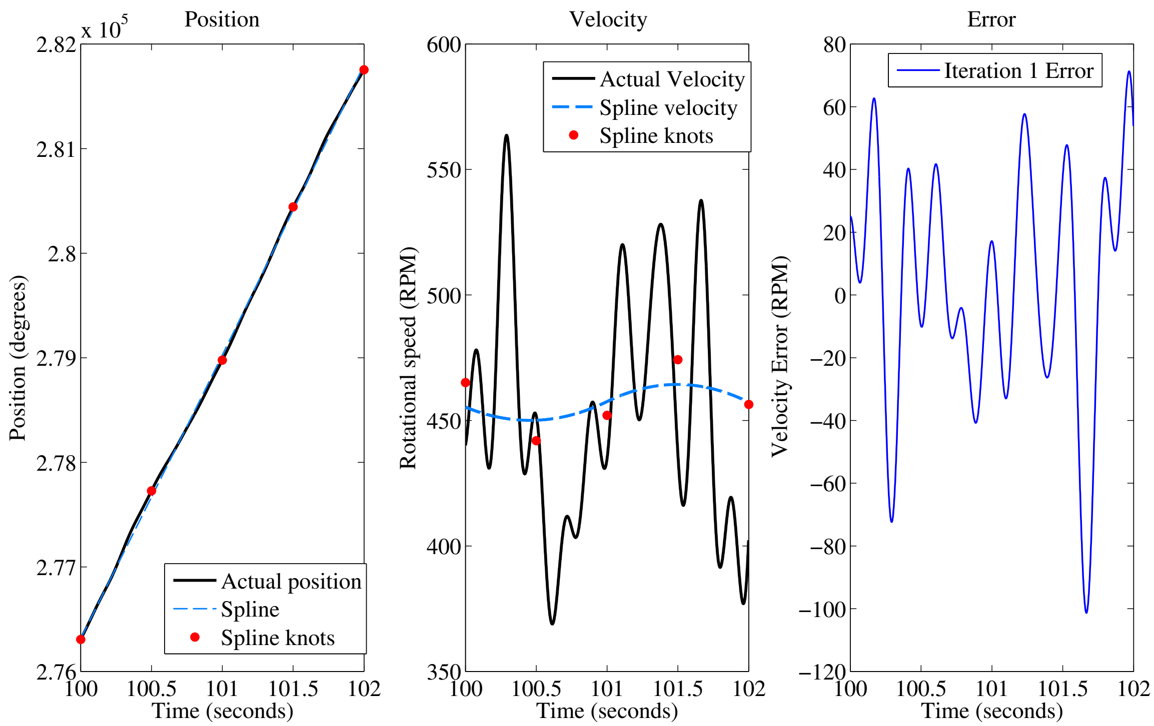

Figure A.15: Adaptive spine method over several iterations with a test function

Accuracy of the spline-fitting angular velocity

The adaptive spline method seems to give low error for a variety of test functions, but there is no direct way to validate this method with the measured turbine position data because we don't have access to the true turbine rotational speed. To get a sense of how well this method works we need to relate the position error (to which we have access), to the velocity error (for which we do not have access). To explore this relationship, we created a large number of test functions with a combination of sinusoidal functions with varying amplitudes and frequencies. We applied the adaptive spline algorithm to these test functions and found a global position error (the sum of all of position errors along the position vector). We then defined a global velocity error the same way: the sum of the errors between the test function velocity (derived analytically) and the spline velocity (derived analytically). This two metrics were calculated for each test function and plotted in Figure A.16. A sample of the test functions is given in Figure A.17.

Figure A.16 shows that a low global position error is a good predictor of a low global velocity error for tested functions. As the frequency of the test functions increases the position error and the velocity error increase, and the correlation between the two metrics becomes weaker.

A.3 Conclusion

The combination of the encoder resolution, the sampling rate, and the rotational speed created position data that produced banded rotational speed calculations when a simple forward difference method was used to derive velocity. Several more involved methods were tested to extract useful high-frequency information from the data. B-splines were found to work well for two test cases, but the error associated with this method was dependent on the frequency of the test function. An adaptive method was proposed that allows the size of the spline to vary along the position vector in a way that best fits the data. A simple global metric was proposed to evaluate the error associated with this method when the true velocity data is unavailable.

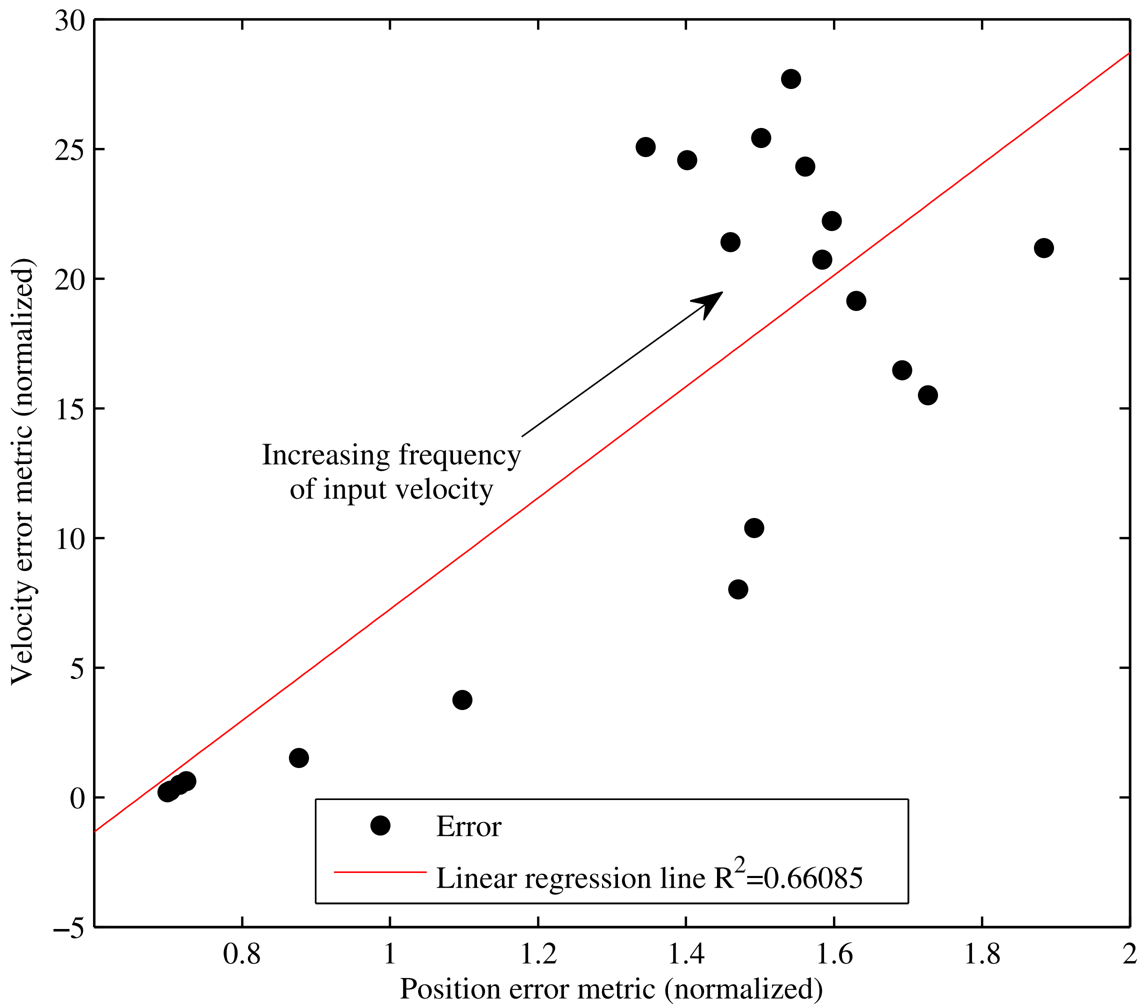

Figure A.16: Velocity error metric plotted against position error metric. The correlation becomes weaker with increasing frequency of the test velocity.

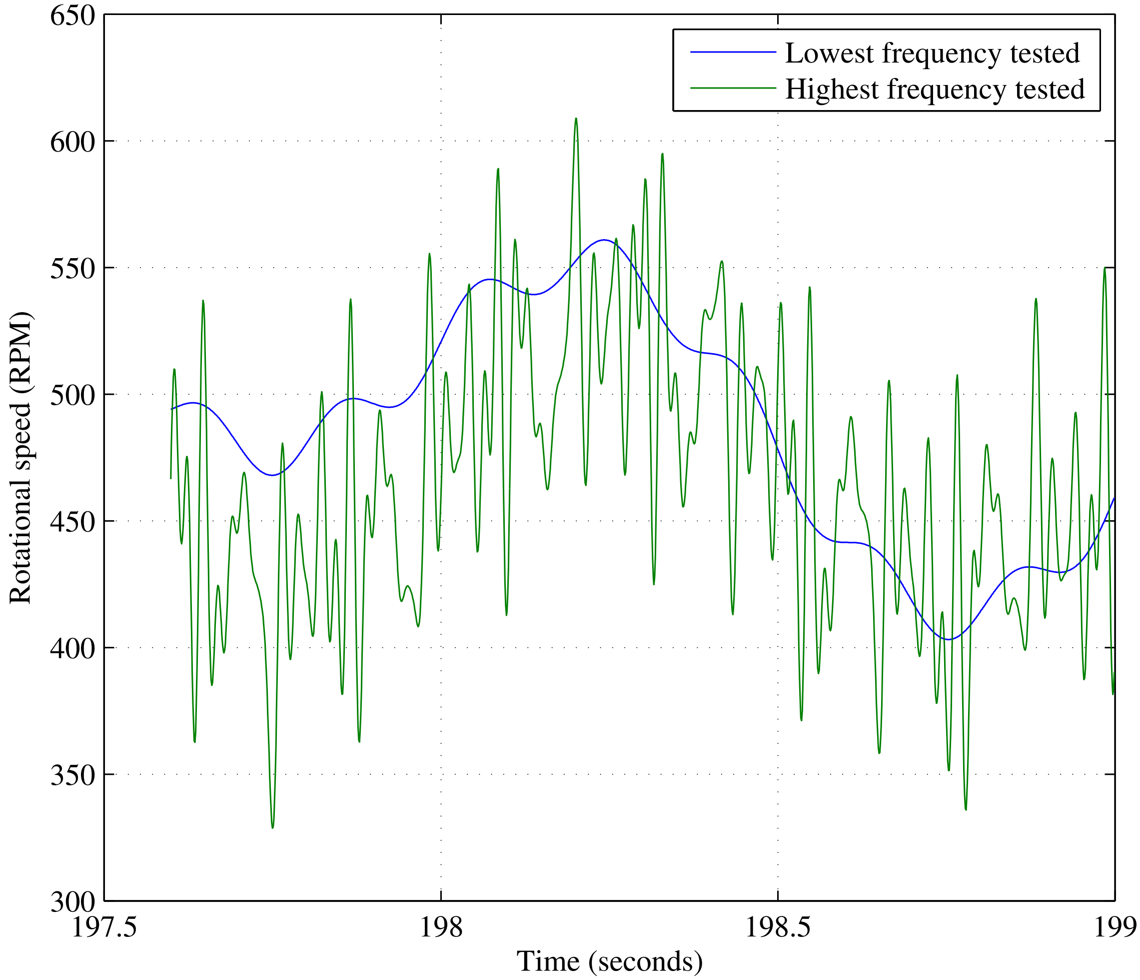

Figure A.17: Snippets of the lowest and highest frequency velocities on which error correlation analyis was performed.

Appendix B

INTRUMENTATION AND DATA ACQUISITION

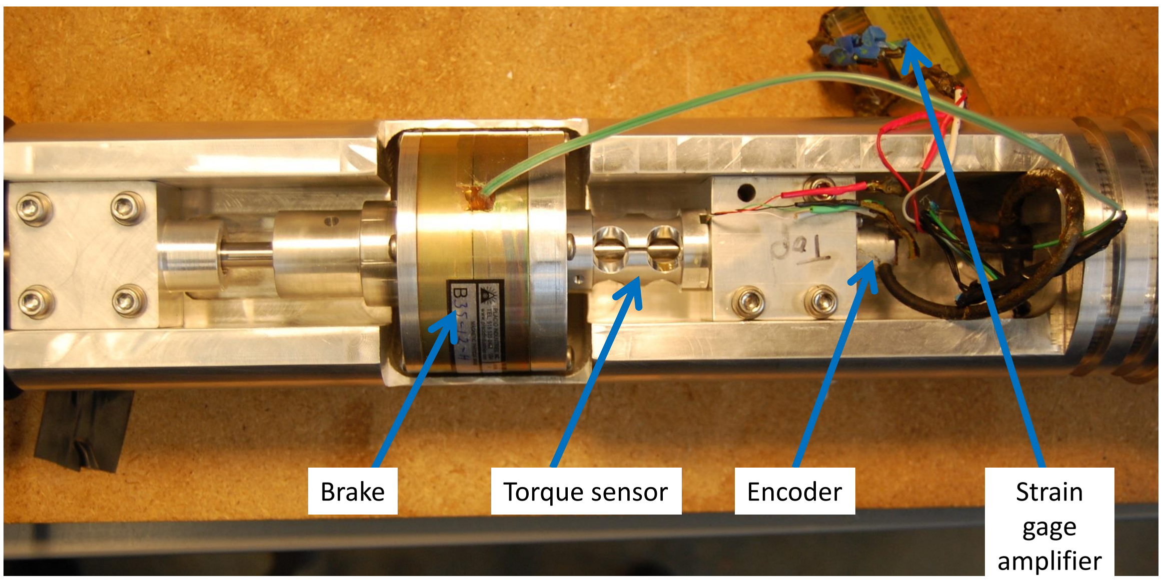

Figure B.1: Photograph of the nacelle instrumentation

B.1 Encoder details

Renishaw RLS RM22I non-contact rotary encoders were used to measure rotor rotational position. This encoder uses the magnetic Hall effect to measure rotational position with an nominal resolution of 1024 counts per revolution. The functional resolution when used with the National Instruments data acquisition system was 256 counts per revolution. The manufacturer states that this encoder is accurate to within

± 0.7 degrees. This encoder uses a reference signal ("Z") that resets the count every revolution.

B.2 Torque sensor details

Futek TFF325 reaction torque sensors (rated capacity 6 N-m) were used to measure the torque developed by the rotor. These torque sensors measure torque through the use of strain gages placed on an aluminum housing. The aluminum housing is designed to deflect under torque such that the strain measured by the strain gages is proportional to the torque applied to the aluminum housing. A Wheatstone bridge connects four strain gages and allows precise measurement of the changes in electrical resistance through the strain gages associated with the application of torque. The manufacturer's specifications report the accuracy of the zero balance as $\pm 1\%$ of rated output, the non-repeatability as $\pm 0.05\%$ of rated output, and the hysteresis as $\pm 0.2\%$ rated output.

The output voltage (2 mV/V) associated with the torque sensor was amplified at the nacelle due to signal-to-noise concerns associated with the long cable length between the torque sensor and the data acquisition system. Dataforth DI-8B38 strain gage input modules were used to provide this amplification. These strain gage input modules provide an excitation voltage to the strain gages on the torque sensor (10 V), amplify the output to ± 5V, and provide high frequency filtering (100 dB per decade of normal mode rejection for frequencies above 8 kHz).

B.3 Magnetic particle brake details

Placid Industries B35 hollow-shaft magnetic particle brakes were used to apply loads on the rotors. These brakes consist of a housing containing fine metallic particles, and a disk coupled to the shaft that spins within the housing. Coils contained in the housing create magnetic flux lines that increase the resistance of the particles to slip. This resistance is proportional to the current applied to the coil, and independent of the rotation rate. The resistance is dissipated as heat.

Placid Industries constant current power supplies were used to control the torque that the brake applies to the rotor. These power supplies allowed the adjustment of current by a manually controlled potentiometer.

B.4 Data acquisition system details

Torque and rotational position data were acquired and processed with a National Instruments PCIe-6341 X series data acquisition card, used in conjunction with LabVIEW software. These data were acquired at 1000 Hz. The LabVIEW Virtual Intrument (VI) used to acquire, process, save, and present the data was developed for this project. This VI processes the analog voltage from the torque sensor and the rotational position data from the encoder into meaningful results (torque in N-m and rotational speed in TSR) that are monitored during the experiment. The raw data (analog voltage and rotational position) is streamed to disk for later analysis. Timestamps of millisecond resolution from the computer clock were recorded simultaneously so that the PIV data (which was acquired on a separate computer) could be coordinated with the encoder and torque sensor data. The computer synchronized with a time server every five minutes to prevent drift in the on-board clock.

B.5 Particle image velocimetry details

The particle image velocimetry system consisted of a PIV dual frame-straddling camera, a timing unit, and a laser. All of these components were part of a LaVision PIV system that is the property of the Bamfield Marine Science Centre. The camera has an acquisition rate of 5 Hz, and a resolution of 1392 x 1024 pixels per image. The time between images was set to 1500 microseconds for all tests. All PIV measurements were performed for 1 minute, resulting in 300 image pairs. A simple calibration was performed by holding a ruler in the laser plane, taking an image of the ruler, and measuring the distance in pixels across a fixed distance on the ruler. The camera was at a right angle to the flume walls and streamwise direction, so no image distortion or parallax corrections were necessary.

The laser was placed beneath the flume on an optical table, and a laser plane was created and directed vertically by a cylindrical lens and mirror. A jig was constructed so that the optical table could be easily moved along the length of the flume while maintaining the laser plane parallel to the flow and on the centerline of the channel.

PIV image pairs were processed in MATLAB with the PIVLab toolbox. The images were preprocessed with a high-pass filter and intensity capping, and a three-pass PIV algorithm was used, resulting in a scaled vector field of size 173 rows x 124 columns (16x 16 pixel resolution). The resulting vector field was processed with a min/max filter, standard deviation filter, and a normalized median filter. Occasionally the PIV camera dropped the second image in an image pair; the resulting vector field was then discarded.

The timestamps associated with the PIV images and the timestamps associated with the torque and rotational speed data are equal to within approximately five milliseconds. To achieve this, the turbine data acquisition computer was synchronized to a time server every five minutes to prevent the drift of the on-board clock, as was done in the PIV data acquisition computer.

B.6 Acoustic doppler velocimeter details

A Nortek Vector ADV was used to make point measurements of the flow velocity. This instrument resolves all three components of the flow. An attempt was made to align the probe as perpendicular to the flume freestream as possible, such that the mean vertical and cross-channel flow speeds were zero. This ADV has an acquisition rate of 100 Hz.

Appendix C

TESTING PROCEDURE

A checklist was developed during the experimental data collection to uniformize the testing procedure. That checklist is included below to provide a sense of how the experiments were performed and as a guide for those that are performing similar experiments.

C.1 Pre-test Checklist

1. Confirm experimental plan, including turbine spacing and turbine operating conditions (TSR).
2. Check the filepaths and settings in the LabVIEW VI.
3. Check the LaVision imaging acquisition settings.
4. Check vacuum water trap, if it is 3/4 full then empty before test.

C.2 Test Protocol

1. Turn on the torque sensor amplifier power supply and the magnetic particle brake power supply.
2. Turn on the air compressor to pressurize the nacelles (set to 10 PSI)
3. Put turbines in flume.

4. Align turbines parallel to the flume and to correct streamwise locations with a jig.
5. Check nacelles for leaks.
6. Turn on vacuum pump.
7. Quickly set cable-holding suction cups.
8. Clamp cable bundles out of the way.
9. Let vacuum develop and make sure water has stopped entering water trap.
10. Start LabVIEW with naming convention: Directory=(Config_name), Filename=[(PIV_location) (test number)].
11. Wait at least 30 seconds.
12. Start pumps (one at a time).
13. Turn brakes down and hand start turbines.
14. Brake to correct TSR according to experimental plan, wait to stabilize.
15. Take a sample image with the PIV camera to check image quality.
16. Begin PIV acquisition with same naming convention as LabVIEW VI.
17. Restart LabVIEW after acquisition, repeat.
18. Periodically check for turbine leaks.
19. After experimental plan has been completed start shutting down pumps (keep LabVIEW running).
20. Stop LabVIEW 30 seconds after water is still.
21. Turn off vacuum pump.
22. Remove turbines from flume.

23. Put laser on standby .

24. Turn off air compressor and open tank drain.

25. Turn off pump breakers.

26. Back up Labview data.

Appendix D

FLOW CHARACTERIZATION OF THE BMSC FLUME

The mean flow and turbulence intensity in the flume at the Bamfield Marine Science Centre was characterized with an Acoustic Dopper Velocimeter (ADV) at 38 points along the test section. The ADV measures three components of the instantaneous velocity at a point; a time-series of ADV measurements were made at each of the 38 points, with approximately one minute in duration. These data were then filtered (de-spiked) and used to calculate mean velocity and turbulence intensity. The mean streamwise velocity in the test section is shown in Figure D.1, and the turbulence intensity is shown in Figure D.2.

These measurements show a range in mean streamwise velocity from 0.92-1.06 m/s, and a range of turbulence intensity from 4.2-11.8%. There is significant vertical and horizontal shear in the velocity profiles, especially near the flume entrance. Two meters from the flume entrance these measurements show a 7% horizontal variation in mean streamwise velocity, though this horizontal variation decreases farther from the flume entrance.

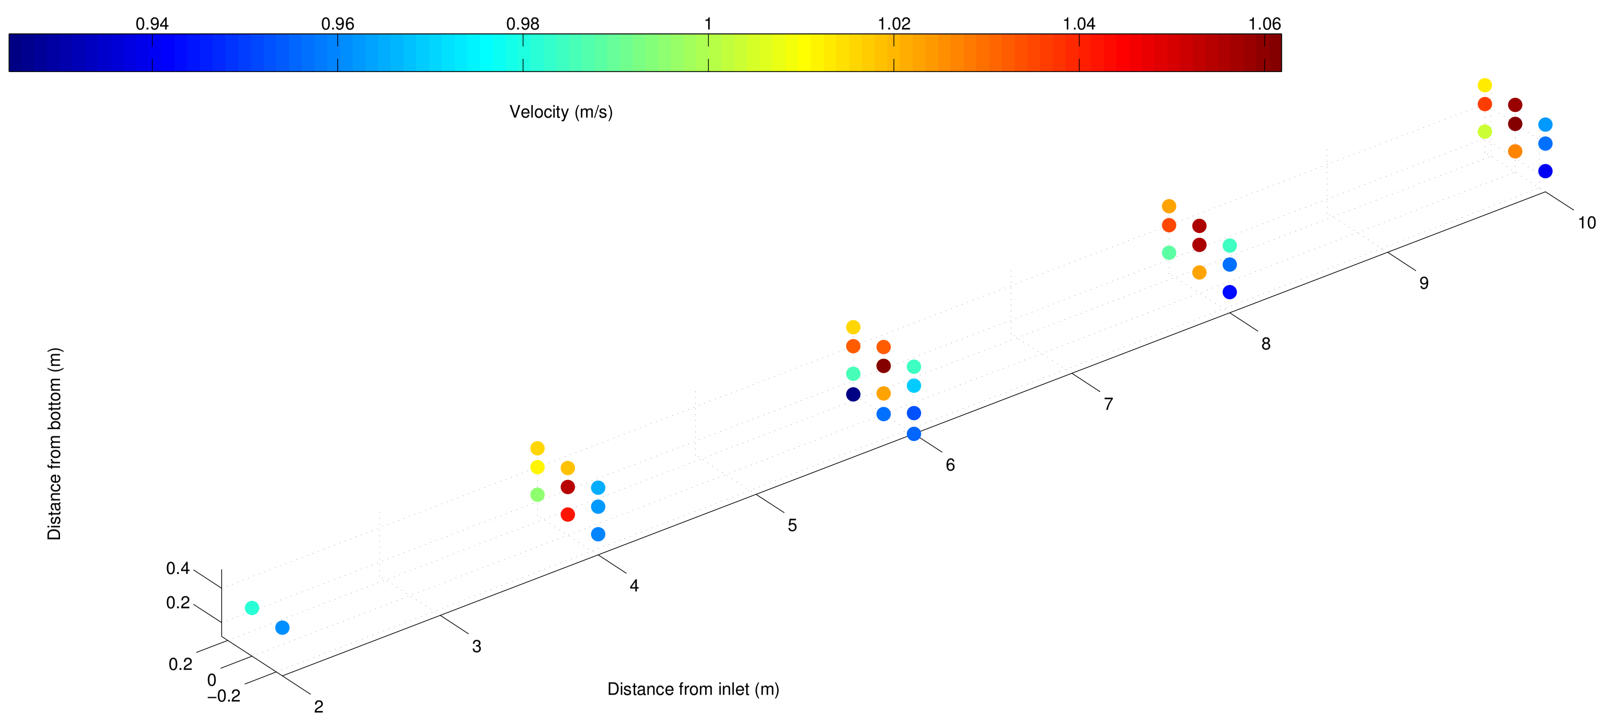

Figure D.1: Mean streamwise velocity in the BMSC flume test section measured by ADV.

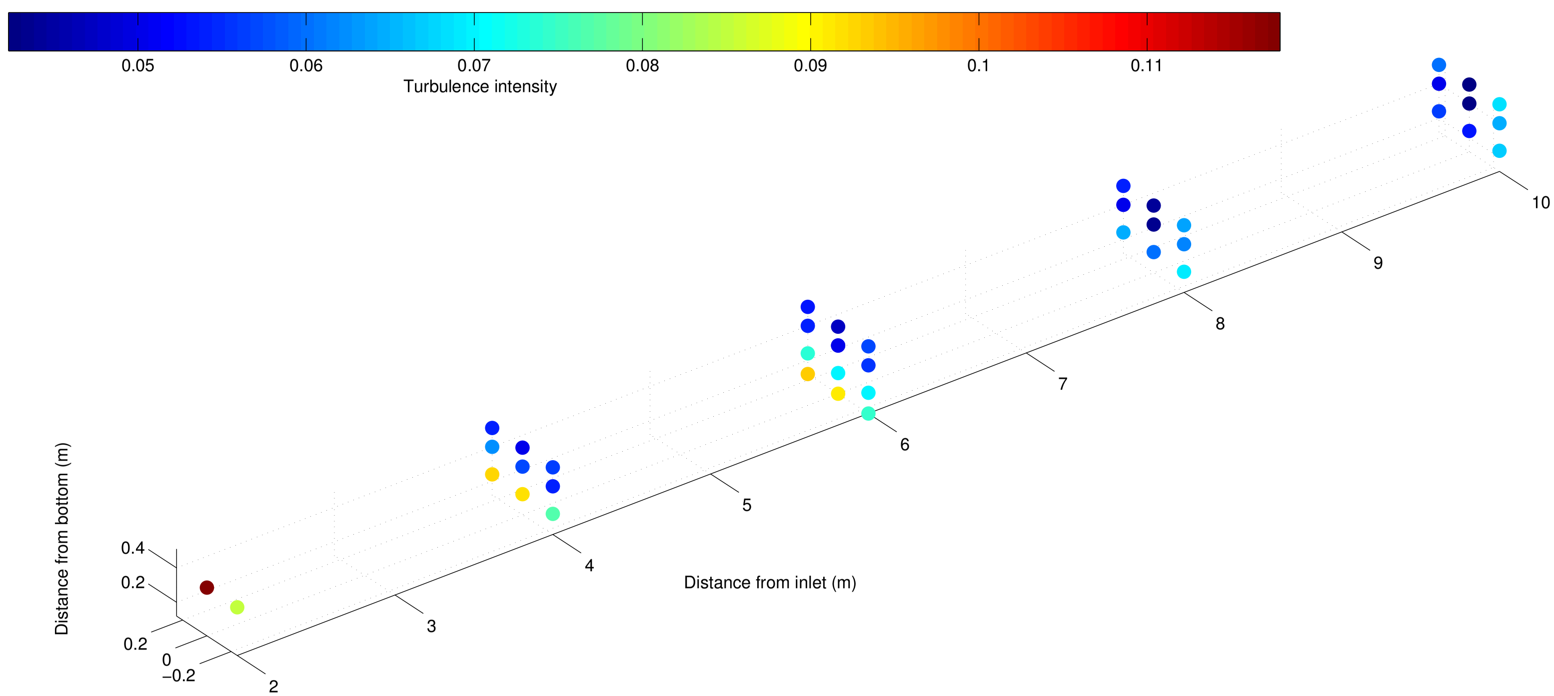

Figure D.2: Mean turbulence intensity in the BMSC flume test section measured by ADV.