

An Analysis of Quadrotor Flight in the Urban Canopy Layer

by

John Ware

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016 [September 2016]

© Massachusetts Institute of Technology 2016. All rights reserved.

Signature redacted

Author

.....

Department of Aeronautics and Astronautics
May 19, 2016

Signature redacted

Certified by

.....


Nicholas Roy
Associate Professor
Thesis Supervisor

Signature redacted

Accepted by

.....

Paulo C. Lozano
Chair, Graduate Program Committee



An Analysis of Quadrotor Flight in the Urban Canopy Layer

by

John Ware

Submitted to the Department of Aeronautics and Astronautics
on May 19, 2016, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

This thesis presents two distinct bodies of work concerning quadrotor flight in urban wind fields. The first attempts to characterize a quadrotor's ability to exploit urban wind fields for improved flight performance. A computational fluid dynamics model is used to obtain a wind field estimate given a 3D model of the environment and a prevailing wind estimate. Minimum-energy trajectories are then found through the environment using an empirically derived power consumption model for a specific quadrotor platform. It is shown in simulation that a minimum-energy planner aware of the wind field outperforms a naive, wind-unaware planner over metrics such as total energy consumption, time to goal, and failure rate. The second component of the work focuses on the development of an onboard wind sensor for quadrotors. Although it is not yet clear how to integrate these measurements into a global wind field estimate or use them in a planner, it is intuitive that on-board measurements could inform the local wind field estimate or validate the global one. Accordingly, an effort was made to integrate an existing microelectromechanical flow sensor into a quadrotor platform. Initial results from full-scale tests in the Wright Brothers wind tunnel at MIT demonstrate the the sensor's performance in flight and hover conditions.

Thesis Supervisor: Nicholas Roy

Title: Associate Professor

Acknowledgments

Thanks to my wife and family for supporting me through this. Living with an MIT student isn't always the easiest thing. Thanks also to all of the members of the RRG for being great collaborators and thinkers. Research can be an isolating endeavor and having the support of those around you makes all the difference. Above all, thanks to Nick for giving me this opportunity, celebrating the successes with me, and working through the challenges.

Dad, you are dearly missed and I can only hope that you can somehow share my joy in completing this phase of life.

Contents

1	Introduction	15
1.1	Motivation	15
1.2	Wind Fields	16
1.2.1	Mesoscale	17
1.2.2	Local Scale	18
1.2.3	Microscale	19
1.3	Wind Field Estimation	20
1.4	Minimum-Energy Trajectory Planning	24
1.5	Overview	25
2	Related Work	27
2.1	Wind Field Estimation	27
2.1.1	Local Measurements	28
2.1.2	Simulation	33
2.2	Minimum-Energy Planning	41
2.3	Wind Measurement	42
3	Onboard Wind Measurement	49
3.1	Wind Measurement Techniques	49
3.2	Operating Principle	51
3.3	Calibration and Validation	52
4	Wind Field Simulation	59

4.1	Simulation Environment	59
4.2	Prevailing Wind Conditions	60
4.2.1	Mixing and Reference Heights	61
4.2.2	Proximity to Obstacles	62
4.2.3	Speed and Heading Distributions	63
4.3	QUIC	65
4.3.1	QUIC-URB	65
4.3.2	QUIC-CFD	67
5	Planning Over MIT Campus with Steady Wind Fields	73
5.1	Problem Formulation	73
5.2	Power Consumption Model	75
5.3	Planning	81
5.4	Results	83
6	Conclusions	89
A	Wind Field Simulation Best Practices	91
A.1	Sources of Error	91
A.2	Simplification of Physical Complexity	92
A.2.1	Reynolds-Averaged Navier-Stokes (RANS)	94
A.2.2	Reynolds Stress Model (RSM)	94
A.2.3	Large Eddy Simulation (LES)	95
A.3	Geometric Representation	96
A.4	Domain Size	97
A.5	Boundary Conditions	98
A.5.1	Inflow	98
A.5.2	Wall Boundary Conditions	100
A.5.3	Top Boundary Conditions	101
A.5.4	Lateral Boundary Conditions	101
A.5.5	Outflow Boundary Conditions	101

A.6 Initial Conditions	101
A.7 Numerical Approximation	102
A.8 Spatial Discretization	102
A.9 Temporal Discretization	103
A.10 Iterative Convergence	103

List of Figures

1-1	Scales of urban wind fields	17
1-2	Wind Features	18
1-3	Example Urban Wind Field	20
1-4	Urban Flow Features	21
2-1	Wind profile estimation from Langelaan	29
2-2	Thermal soaring from Allen.	30
2-3	Wind field estimation with a Gaussian Process from Lawrance et al.	31
2-4	Fixed wind trajectories from Lawrance et al.	32
2-5	Dynamic wind field estimation from Chakrabarty and Langelaan	34
2-6	Wind field simulation from Galway et al.	35
2-7	Wind field library from Galway et al.	36
2-8	Wind field simulations from White et al.	38
2-9	Wind field simulation from Raza and Etele	39
2-10	Flight trajectory from Cybyk et al.	40
2-11	Flight trajectories from Chakrabarty et al.	42
2-12	Trajectory and vehicle from Muller et al.	46
2-13	Sensor layout from Yeo et al.	47
2-14	Trajectory from Yeo et al.	48
3-1	Wind sensor from Piotto et al.	52
3-2	Benchtop wind tunnel results from Bruschi et al.	53
3-3	Wright Brothers wind tunnel vehicle and sensor configuration	54
3-4	Wind sensor differential pressures	54

3-5	Wind sensor speed and heading results	56
3-6	Wind sensor pitch sample curve	57
3-7	Wind sensor pitch test speed and heading results	57
4-1	MIT campus satellite image	60
4-2	Green building weather station	63
4-3	Polar prevailing wind distributions	64
4-4	QUIC 3D sample model	66
4-5	QUIC-URB sample wind field	68
4-6	QUIC-CFD sample wind field	69
4-7	<i>In situ</i> measurement locations over MIT campus	70
4-8	<i>In situ</i> measurement results comparison	72
5-1	Quadrotor in Wright Brothers wind tunnel	78
5-2	Power consumption curve for quadrotor	79
5-3	Power consumption curve for quadrotor	80
5-4	MIT campus 3D model for QUIC-CFD	82
5-5	MIT wind field from QUIC-CFD	83
5-6	Planner energy consumption over MIT campus with QUIC-CFD . . .	84
5-7	Planner failure rates over MIT campus with QUIC-CFD	85
5-8	Naive planner failure locations over MIT campus with QUIC-CFD . .	86
5-9	Results from sample trajectory over MIT campus with QUIC-CFD .	86
5-10	Sample trajectory over MIT campus with QUIC-CFD	87
A-1	100

List of Tables

3.1 Wind sensor pitch test configurations 56

4.1 Measured and Model Wind Speed and Heading 71

Chapter 1

Introduction

1.1 Motivation

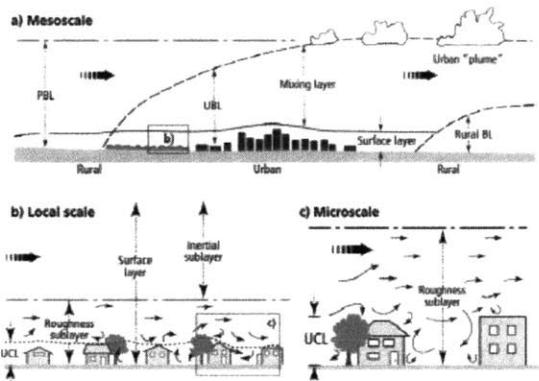
More than ever before, human civilization has become condensed into modern, dense, and complex urban developments. A relatively recent socio-economic shift has brought the fraction of the world's population living in urban areas to 54% with an annual growth rate of 1.84% [2]. Combined with the explosive growth in the commercial unmanned aerial vehicle (UAV) market, the constant improvements in UAV capabilities, ever increasing energy density of batteries, and gradual loosening of regulations makes the eventual operation of UAVs in urban environments seem all but certain. This trend has not been overlooked by companies such as Google and Amazon who seek to establish UAV delivery services in the consumer market [1, 3]. Others continue to push applications such as law enforcement, reconnaissance, surveillance, aerial photography, inspection, mapping, and first response. Although only a small number of companies have obtained permission to proceed with product development in these applications, the consumer UAV market is experiencing explosive growth. This growth has continued despite FAA regulations put in place to mitigate the slew of safety and regulatory challenges brought by this new technology, and UAV manufacturers have continued to evolve their products to improve flight performance and reduce costs.

While UAVs have traditionally been large fixed-wing or single-rotor vehicles, the vehicles that have come to dominate the commercial market and are the most likely


candidates for urban flight are smaller, multi-rotor Micro Air Vehicles (MAVs). As demonstrated by Orr et al. [80], Kathari et al. [56], and Galway et al. [41], the reduced mass and dimensions of MAVs make them more susceptible to the external forces imposed on them by the wind. Their size also limits their payload and thus constrains the amount of energy available for a given flight. These challenges are just a few of those introduced by the unique structure and terrain of urban environments, but reinforce the importance of considering the wind field during navigation. Despite the clear need to estimate and incorporate complex urban wind fields into trajectory planning, it has remained an open problem. The primary contribution of this work is to begin to address these challenges by analyzing the potential benefits of obtaining high-fidelity wind field estimates over any urban environment and explicitly consider their impact on a quadrotor’s flight performance over a set of minimum-energy trajectories.

1.2 Wind Fields

Wind field characteristics vary significantly with respect to the horizontal and vertical environmental scale under consideration. To capture this notion, meteorologists have defined three horizontal scales and a series of vertical boundary layers. The horizontal scales are the microscale, local scale, and mesoscale. Each of these scales, seen in 1-1a, has characteristic wind fields associated with it and either plays an important role in defining the urban wind fields we seek to exploit or motivates the approaches taken in related work. Similarly, the vertical boundary layers range between the urban canopy layer (UCL) and the troposphere and primarily aid in the definition of the mean horizontal velocity profile. The mean horizontal velocity profile is defined as the spatially and temporally averaged horizontal wind velocity as a function of altitude and provides an approximation of the wind speeds a UAV might encounter at a given altitude for a particular environment. As discussed later, the mean horizontal velocity profile can also be used to define the boundary conditions for a computational fluid dynamics (CFD) simulation. This section presents a brief summary of the relevant



(a) Environmental Scales



(b) Horizontal Velocity Profile

Figure 1-1: Illustrations of the horizontal and vertical environmental scales that define the urban boundary layers and the resulting mean horizontal velocity profile. Note that the horizontal velocity profile shown in figure 1-1b is averaged both temporally and spatially averaged in order to remove the effects of turbulence generated by the environment. Figures used with permission from Oke [79].

environmental scales, boundary layers, mean horizontal velocity profile, and their relevance to our problem.

1.2.1 Mesoscale

The mesoscale, seen in figure 1-1a, is defined on the order of tens of kilometers and tends to capture large scale shifts in weather patterns and the meteorological effects of entire cities. The mesoscale’s relevant vertical scales are defined within the troposphere by nested boundary layers called the free atmosphere and the planetary boundary layer (PBL). In the free atmosphere, geostrophic winds, resulting from the balance of the Coriolis effect and pressure gradients, act on synoptic scales of a kilometer or more and tend to have very little temporal or spatial variation. At a lower altitude, the PBL defines the highest extent of the terrain’s effect on the flow through frictional drag, solar heating, and evaporative effects.

Because of its irregularity, commercial passenger jets and full-scale UAVs, such as the Northrop Grumman Global Hawk, spend the majority of their flight time in the regions of the atmosphere far above the extent of the PBL. Due to their reliance on the lift provided by flow features created by the terrain, manned and unmanned




Figure 1-2: Illustrations of three common wind features used in avian and fixed-wing soaring to extend flight duration and range while conserving energy.

gliders must operate at lower altitudes within the PBL at the local scale.

1.2.2 Local Scale

The local scale, seen in figure 1-1a, is defined between one and several kilometers and considers the effects of topography, but mostly excludes effects due to specific environmental features or structures. Within the PBL, the terrain's effect on the flow increases with decreasing altitude and disturbs the uniform wind fields found at higher altitudes. The surface layer, also seen in figure 1-1a, is the region below which the effect of specific terrain features can be identified, and tends to have a height of several hundred meters. Within this layer, phenomena such as Lee waves, thermals, ridge lift, and gusts begin to appear. An example of each of these features can be seen in figure 1-2.

Ridge lift is a region of vertical flow caused when air is forced upwards as it moves over rising terrain. Lee waves are standing waves that are usually started by a region of ridge lift, but continue to propagate over a significant distance. Thermals are a region of rising air created by uneven heating of the surface. Although these features are not directly relevant to urban flight, they are the focus of much of the soaring literature discussed later, and corollaries exist within urban wind fields.

Critical for determining the prevailing wind conditions that drive the flow within an urban environment, the inertial and roughness sublayers are defined within the local scale and compose the surface layer. The upper boundary of the roughness sublayer (RSL) is defined by the highest extent of significant vertical mixing in the

flow. Above this is the inertial sublayer, where the flow is considered to have mixed and the effects of individual terrain features cannot be detected.

As shown in figure 1-1b, the mean horizontal velocity has a maximum value defined at the upper boundary of the surface layer and is usually called the free stream velocity. From this value, the mean horizontal velocity logarithmically decreases with decreasing altitude throughout the majority of the inertial sublayer and over part of the roughness sublayer. This straightforward relationship makes it common practice to take prevailing wind measurements at a given altitude and extrapolate to obtain the full velocity profile. However, the inherent complexity of the wind fields generated by the structures increases the risk of inaccurate measurements as the measurement altitude approaches the mean structure height.

1.2.3 Microscale

The microscale, seen in figure 1-1a, considers the effect of individual objects and structures and is defined between one meter and several hundred meters. Its vertical extent is entirely contained within the roughness sublayer, and includes the urban canopy layer (UCL). The UCL extends up to the mean feature height and because urban environments are dominated by the structures within them, the mean feature height tends to also be the mean structure height. However, as the name suggests, the concept originated with the study of wind flow within forest canopies. The flow within the urban canopy layer is driven by the flow in the upper sublayers, but is strongly governed by its dense structure. Just as the terrain's interaction with the flow creates complex features at higher altitudes, the urban canopy layer contains many flow features that can have a significant effect on flight performance. The most straightforward example of a flow feature exploitable by a UAV is the tendency for the wind to be directed down an aligned urban canyon and over a perpendicular one. Examples of both of these effects can be seen in figure 1-3. As shown, the x component of the wind velocity within the urban canyons is relatively smooth and continuous, while there is a significant reduction in magnitude directly behind the structures at this altitude.

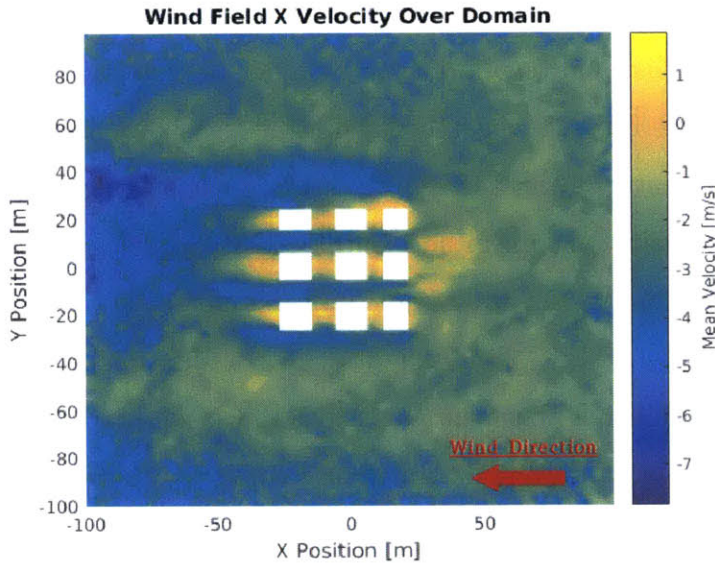



Figure 1-3: Example wind field estimate generated by a computational fluid dynamics solver with the prevailing wind flowing from the right to the left. Note the lower wind velocities found in the wake regions of the buildings and the accelerated flow down the aligned canyons.


Although those examples are intuitive, the urban environment also contains more complex features such as the frontal vortices and trailing wake regions found near single buildings, as described by Blocken et al. [14] and shown in figure 1-4a. As discussed by Hoydysh et al. [50] and shown in figure 1-4b, other wind features can be found in the vortices created by the interaction of the flow around several, closely packed structures. Kittiyoungkun [55] provides further discussion on the range of wind features found in the urban environment and their relevance to UAV flight. As all of these flow features demonstrate, the complex flow within the urban canopy layer is a challenge to estimate and care must be taken when selecting from the wide range of wind field estimation techniques.

1.3 Wind Field Estimation

Although there are many common wind field estimation techniques, no single approach is well suited to all applications and environments and very few are able to capture the complexity found within urban wind fields. This section contains a brief



(a) Single Structure Flow Features [14]



(b) Structure Group Flow Features [50]

Figure 1-4: Illustrations of the flow features found around a single structure and structure group. Note that the scales of each of these features varies with structure dimensions and some may cease to exist for certain configurations. Figures used with permission from authors.

summary of several common wind field estimation methods and their applicability to urban environments.

The simplest approach to wind field estimation is to assume a uniform wind velocity with intermittent gusts. The 1-cosine gust model derives from the Military Specification MIL-F-8785C and is applied to each velocity component individually. The inputs are the gust amplitude in meters per second, gust length in meters, and start time. Although this might be useful for testing a vehicle or structure's response to a smoothed step-input, it does nothing to encode a wind field's natural spatial or temporal variation. A more sophisticated approach is found in the Dryden and von Karman turbulence models. Both of these use characteristic length scales and intensities to define the power spectral densities over the linear and angular components of the wind velocity. The output is a static vector field with a degree of spatial variation corresponding to the input parameters. Although these are commonly used to assess an aircraft's response to turbulent wind fields, the temporal variation is only created by the aircraft's movement through the environment and does not exist for a stationary vehicle. Although the von-Karman model is the preferred continuous gust model of both the Federal Aviation Administration and Department of Defense, neither it nor the other methods discussed thus far are applicable to low altitude wind

fields because they have no way of incorporating the terrain's effect on the flow. An alternative approach to capturing the interaction between the flow and terrain is to forgo simple models and rely solely on a set of wind measurements throughout the domain.

Meteorology was one of the earliest fields to investigate wind field estimation using Kriging, also called Gaussian process regression. Gaussian process regression, discussed more thoroughly by Rasmussen et al. [86], takes a series of measurements as inputs and uses a function to define the covariance between any two points. This covariance function, or kernel function, often has a set of hyperparameters that are estimated given a set of training data. The resulting kernel function defines a distribution over functions that generates a continuous function over the space that can be queried with a test point to get the corresponding mean and covariance estimate. This technique has much to offer when a large number of measurements are evenly distributed throughout the domain and the resulting regression is well supported over the area of interest. While these assumptions are true for wind field estimation using data from satellite scatterometry, where mesoscale wind velocities are inferred from changes in reflected radiation from the earth's surface, they do not hold in urban environments due to the sparsity of measurements provided by existing weather stations or any quickly deployable system. A further complication is that the kernel functions used by the regression assume a spatial smoothness that is violated by any obstacles in the domain. These discontinuities cause the wind field estimate to incorrectly predict a strong correlation between wind velocities on either side of an obstacle, when the physical separation clearly breaks any physical constraints. This failure mode constrains the approach to only being used far from obstacles that interfere with the flow behavior. More recently, Lawrance et al. [61, 62] used a variant of this approach with a time history of wind measurements along a vehicle's flight path. Their approach relies on a spatio-temporal kernel to provide variable sampling density over the map and sparsification of the training set for online estimation. Despite these clever features, this approach suffers from the same drawbacks inherent in spatially smooth kernel functions in environments with obstacles.

While the complexity of the urban environment and its wind fields defeats these approaches, the dense and static structures actually provide some predictability given known prevailing wind conditions. A naive approach would be to use the publicly available measurements from weather stations as inputs to a simple regression. Unfortunately, the complex geometry of a typical urban environment would require weather stations at nearly every street corner to properly capture the resulting flow fields. Because weather stations tend to be distributed at the scale of neighborhoods, they cannot sufficiently capture the flow fields within the urban canyons, but are sufficiently dense to provide an accurate estimate of the local prevailing wind conditions in the upper sublayers. Combining these estimates with a model for the horizontal velocity profile and a map of the structures in the environment defines both the domain and input conditions for a computational fluid dynamics solver that can simulate the interaction of the flow with the environment. Although computationally taxing, the accuracy of this approach and its prior use in urban planning, structural engineering, and studies of fixed-wing flight performance in urban environments motivates its use for this application.

In general, CFD solvers use one of three main turbulence models to solve the Navier-Stokes equations. In order of increasing sophistication, realism, and solution complexity, these are Reynolds-averaged Navier-Stokes (RANS), large-eddy simulation (LES), and direct numerical simulation (DNS). RANS provides a time-averaged or steady solution and therefore generates a static wind field estimate. Due to its relative computational simplicity, RANS is the standard in the engineering industry, and the $k - \epsilon$ model is the most frequently used version. This two parameter turbulence model resolves the two transport equations for the turbulent kinetic energy, k , and the turbulent dissipation rate, ϵ . Despite the popularity of RANS models, they incorrectly assume that the turbulent behavior is equivalent at all scales. Because this is not true, LES directly solves for the flow in the larger, transient, and geometry dependent eddies, but resolves the more universal and isotropic smaller scale eddies with an approximate subgrid-scale model. Since LES is used in an unsteady simulation, it provides a time series of wind fields by explicitly computing the time evolution

of the flow. DNS, the most computationally intensive method, directly solves for the flow at all scales. Currently, this is only feasible for very small domains and will not be considered here. Further details on CFD simulation, turbulence models, and the significant of the various approaches will be discussed later.

1.4 Minimum-Energy Trajectory Planning

As established in the literature and discussed in greater depth later, a natural way to drive the exploitation of a wind field is to consider the amount of energy being expended by the agent as it navigates through the environment. For a vehicle without wing lift, it is natural that traveling upwind should have a higher action cost than traveling downwind due to the increased form drag. A winged vehicle adds some additional complexity because the vehicle's lift must be considered along with its airspeed. A vehicle's airspeed is defined as the wind speed minus the vehicle's ground speed. The study of this problem, minimum-energy trajectories for winged vehicles, is generally called soaring and is an important reference for the work presented here.

A common practice among birds as well as manned and unmanned gliders, soaring is the act of extracting energy from the wind field in order to improve some aspect of flight performance. In the earliest known description [5], Lord Rayleigh stated that when a bird does not work its wings while maintaining altitude, we must conclude that either the wind is not horizontal, or the wind is not uniform. These two alternatives are now known as static and dynamic soaring, respectively. Specifically, static soaring is the extraction of energy, usually in the form of potential energy, from a relatively uniform wind feature that provides lift through the vertical component of the wind velocity. Dynamic soaring is the extraction of energy from large gradients in the wind field such as those created by waves or hills and frequently exploited by albatrosses [95]. Both types share the common objective of extending flight range or duration. Despite its frequent use in nature as well as manned, fixed-wing flight, the effective use of soaring has yet to be demonstrated in complex wind fields such as those found in urban environments. Although the fixed surfaces of a quadrotor cannot

generate much lift, the minimum-energy trajectory planning approaches developed in the soaring literature offer an appropriate foundation for the work presented here.

1.5 Overview

This document begins with Chapter 2 by discussing related work in wind field estimation for UAV flight, trajectory planning for fixed-wing UAV soaring, and past approaches to on-board wind velocity measurement. The remainder of the work presents results from two distinct efforts to enable quadrotors to efficiently and safely navigate the complex wind fields found in urban environments. The first component focuses on wind field estimation and begins in Chapter 3 with the design, development, and testing of a low-cost wind sensor for two-dimensional wind velocity measurement on-board a quadrotor. This leads into Chapter 4 with a discussion of urban wind field simulation using a CFD solver. Chapter 5 presents the second component which analyzes a quadrotor's ability to exploit urban wind fields for improved flight performance. This includes a formal planning problem statement, the development of an empirically derived power consumption model, and discusses the simulation results of minimum-energy trajectory planning over portions of the MIT campus with known wind fields. Finally, Chapter 6 closes with conclusions.

Chapter 2

Related Work

The related work presented here draws from the research areas of fixed-wing UAV soaring, analyses of UAV flight performance in urban wind fields, and on-board wind estimation. Although a few of these efforts have focused on single and multi-rotor UAVs, the majority of the supporting literature comes from research on fixed-wing platforms. For clarity, the relevant components of each work have been arranged into the sections on wind field estimation, trajectory planning, and on-board wind estimation.

2.1 Wind Field Estimation

A prerequisite for exploiting a wind field is having an accurate estimate of that wind field. The simplest approach, demonstrated by Ceccarelli et al. [22] and Osborne et al. [81], is to assume that it is both static and uniform. Although this might be a decent approximation for an instantaneous snapshot of a synoptic wind field at high altitude, it ignores the dynamic nature of weather patterns and could become arbitrarily wrong as the prevailing wind shifts. Providing an incremental improvement, Rysdyk et al. [94] and Langelaan et al. [58] assume a uniform wind field given a prevailing wind vector with additive Gaussian noise. A series of air and ground speed measurements are then used to estimate each component of the prevailing wind velocity. Although this approach allows for slowly evolving, uniform wind fields, it forces

the user to choose between robustness and accuracy in the complex wind fields found at lower altitudes where the terrain’s influence on the wind field creates spatial and temporal variation. Despite its drawbacks, the method of using a history of local measurements to build a global wind field estimate is one frequently taken in the soaring literature.

2.1.1 Local Measurements

Using a set of local measurements can be a flexible and powerful approach to wind field estimation, but imposes strong constraints on the quality and distribution of those measurements. As might be expected, given an expressive representation and a sufficient number of wind measurements appropriately distributed throughout the domain, the global estimate obtained from the measurement set could closely approximate the true wind field. However, an overly sparse or poorly distributed set of samples over the domain could artificially and erroneously bias the resulting estimate.

One example of using a history of local measurements to estimate the global wind field was presented by Langelaan et al. [59], and assumed a polynomial parameterization of the wind velocity over altitude. Given wind velocity measurements from a Pitot tube mounted on a small, powered glider, a Kalman filter was used to estimate the polynomial coefficients of the horizontal wind speed profile. By repeatedly flying vertical loops, the estimate of the velocity profile was updated with the new measurements and propagated with a constant motion model between time steps. The evolving velocity profile can be seen in figure 2-1. A similar approach was taken by Sydney et al. [100] for a quadrotor flying near a single structure through a turbulent, spatially varying wind field while estimating the values of a parametric wind model using a recursive Bayesian filter. Although these techniques allow both temporal and spatial variation in the wind field, they solve both of these problems in overly simplified ways such that the generated estimates are not expressive enough to be used in more general environments.

Working with a powered motor-glider, Allen [7] used a history of wind speed and lift rate measurements to estimate thermal location, size, and strength. Each of these




Figure 2-1: Langelaa et al. [59] show the evolution of a wind field profile estimate over time, in red, as the UAV makes additional loops through the environment to gather data. The dotted red lines show the 2 sigma bounds of the profile estimate. The blue lines show a curve fit of approximate ground truth data obtained from a weather balloon. Figures used with permission from authors.

parameters was found by using gradient descent to fit a spatial, squared exponential model of the wind velocity over the weighted history of measurements. A total of 23 thermals were detected and utilized over the course of 17 flights with an average altitude gain of 173 m. The flight path data from one of these flights can be seen in figure 2-2. Although this approach allowed for the successful discovery and tracking of thermals, it achieved this success by placing strong assumptions on the form of the wind field and then only flying near these wind features. In order to be successful outside of this narrow application, the method must allow for wind fields of a more general form.

In [61, 62], Lawrence et al. presented a wind field estimation technique that relies on Gaussian process regression over a history of local measurements and demonstrated its ability to build an accurate estimate of a static wind field with a single, stationary thermal. This nonparametric approach is powerful because it can use a set of wind measurements over the map to build a continuous estimate of an arbitrary wind field. Accordingly, this method captures spatial variation quite well, but does not

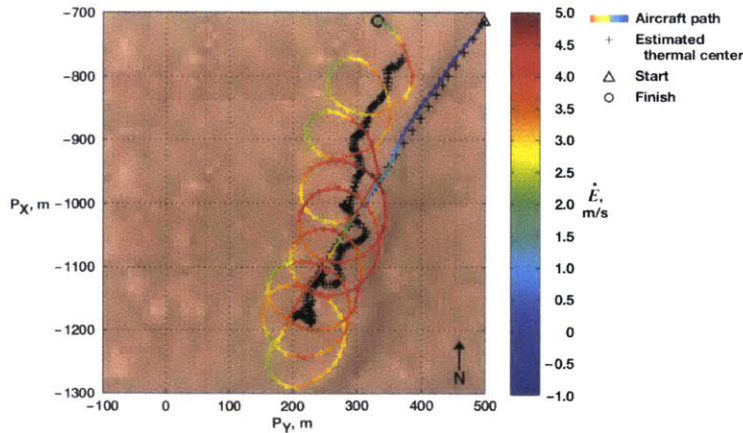


Figure 2-2: A figure from Allen [7] showing the flight path and tracking of a thermal. Note that the vehicle initially flies a straight path until it finds the thermal and then attempts to orbit the estimated thermal position for the remainder of the flight. Figures used with permission from authors.

elegantly handle temporal variation. For example, a temporal drift in a static wind field would require heavy resampling to get an accurate estimate of the shifted, but otherwise unchanged, wind field. Another drawback of this approach is its cubic complexity with respect to the number of measurements over every optimization step while learning the hyperparameters that define the covariance kernel function. Given the innate need to collect online measurements, the size of the measurement history must be carefully managed and places an upper bound on the resulting estimate's maximum domain size and resolution. An example of a wind field estimate from this approach can be seen in figure 2-3. This estimate is from a simulation of a glider exploring a $400 \text{ m} \times 200 \text{ m} \times 100 \text{ m}$ domain with no initial knowledge of the wind field. Accordingly, there is an initial period while the domain is being explored before the estimate closely approximates the true wind field. Note that the domain size is relatively small relative to the scales of high altitude wind features being exploited in typical fixed-wind soaring.

Also note that the simulation in figure 2-3 exaggerates the Gaussian process's ability to estimate complex spatial variation by including the sinusoidal wind feature across the domain. In general, it is challenging to infer the state of the wind field at a distance from the current measurement location, but including the rarely found,

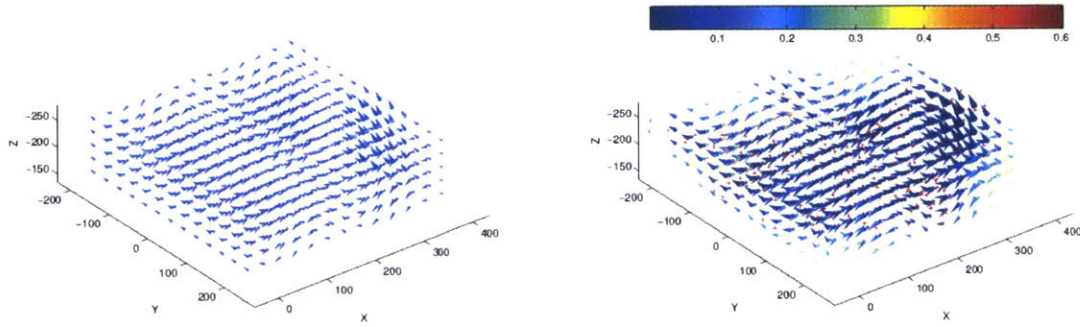


Figure 2-3: An illustration from Lawrance et al. [61] of the ground truth, shown on the left, and final mean estimate, shown on the right, of the wind field in a domain with a Lee wave and thermal. The cone’s size and orientation represent the wind speed and direction, respectively, and are color coded by the variance in $(m/s)^2$. The red crosses indicate the vehicle flight path through the environment. Figures used with permission from authors.

periodic Lee wave somewhat artificially enables local measurements to be accurately propagated over one dimension of the domain. If the wind field were composed of a dense set of more spatially irregular features, significantly more measurements would be required to get the same accuracy. This case would further complicate the challenge of tuning the length scale parameters of the squared exponential kernel function. Without a regular set of features with known length scales, it would be easy to overfit to a particular feature and smooth over higher frequency variation present in another.

In an attempt to address some of these issues, Lawrance et al. modify their approach in [63] to allow for time varying wind fields and feature drift. The spatio-temporal kernel proposed in their work also gives a relatively rigorous method for rejecting the least informative measurements in the set. As previously mentioned, this is a critical factor in managing the complexity of updating the estimate. A time series of wind field estimates and their corresponding glider trajectory for a single simulation can be seen in figure 2-4. Again, despite the more sophisticated kernel function, the size of the domain is relatively small compared to typical glider flight paths. By continuing to rely on the periodic Lee wave as the central wind feature, this work also underestimates the number of measurements needed for an accurate

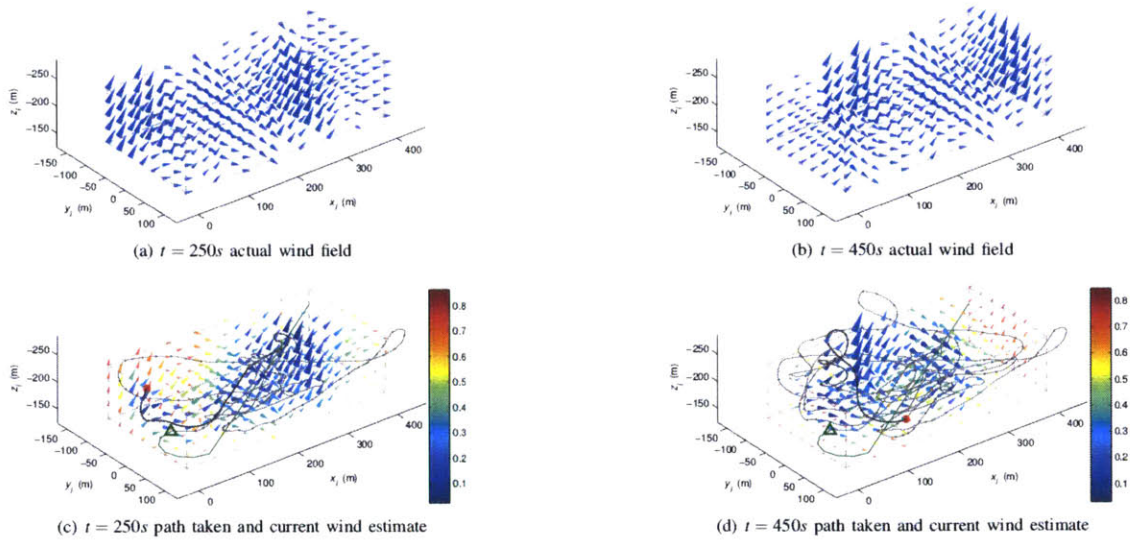


Figure 2-4: A series of snapshots of the wind field estimate and flight path from a simulation done by Lawrance et al. [63]. The path is indicated by the grey line and the autonomous soaring begins at the green triangle after an initial period of exploration and ends at the red circle. The cones represent the wind velocity magnitude and heading and are color coded by the variance in $(m/s)^2$. Figures used with permission from authors.

estimate in a more irregular wind field.

More fundamentally, Lawrance et al. [63] illustrate why the paradigm of using local measurements can be successful at high altitudes, but is much less effective at lower altitudes where the flow in the environment is dominated by its interaction with obstacles. As indicated by the authors, their approach was limited by the growth in computational complexity of updating the wind field estimate as the number of measurements increased. The relatively large amount of spatial variation in a similarly sized urban wind field would significantly increase the number of measurements required to accurately capture the underlying wind field. The temporal variation at any point would also prevent a single measurement from being representative of the expected wind velocity and would require repeated visits or extended periods of time to be spent at every measurement location. Finally, there is an inherent assumption of smoothness when using any common kernel function such as the squared-exponential. Urban wind fields violate this assumption because the struc-

tures in the urban environment introduce sharp discontinuities within the domain. For example, a measurement on one side of a building corner is only loosely correlated with the conditions on the opposite side. A measurement upstream of a small obstacle compared with one downstream also has little correlation, despite its close spatial proximity. With these limitations in mind, it becomes clear that any successful approach to urban wind field estimation will explicitly incorporate a model of the environment or similar environments in some way.

2.1.2 Simulation

A more direct approach to wind field estimation uses CFD and a model of the environment to simulate the wind fields generated by the current prevailing wind and weather conditions. Recently, several authors have made significant headway in applying CFD solvers to wind field estimation for UAV flight. The following section outlines their results and investigates the strengths and weaknesses of their approaches in regard to urban wind field estimation.

Chakrabarty and Langelaan [24] demonstrated the use of CFD for wind field estimation for a fixed-wing, unpowered glider path planning over a mountain range in central Pennsylvania during a 12 hour period on October 7, 2007. The simulated glider was a RnR Products SB-XC with a mass of 10 kg and a roughly 4 m wing span. The objective was to extend the vehicle's range as it attempted to reach a goal by exploiting regions of lift within the wind field. The time series of wind field estimates was computed by Young et al. [108] with the numerical weather prediction package Weather Research and Forecasting (WRF). The simulation domain had a large extent of $70 \text{ km} \times 70 \text{ km} \times 5 \text{ km}$ with a discretization of 0.44 km horizontally and a descending density with increasing altitude. Unlike using a history of local measurements, CFD solvers directly compute global wind fields that explicitly consider the effect of the terrain. Although CFD solvers are computationally expensive, the results are more likely to be accurate, and assuming offline computation, significantly larger domains can be considered. To extend the use of a CFD solver to a real-time application with arbitrary prevailing wind conditions, a library of wind fields would

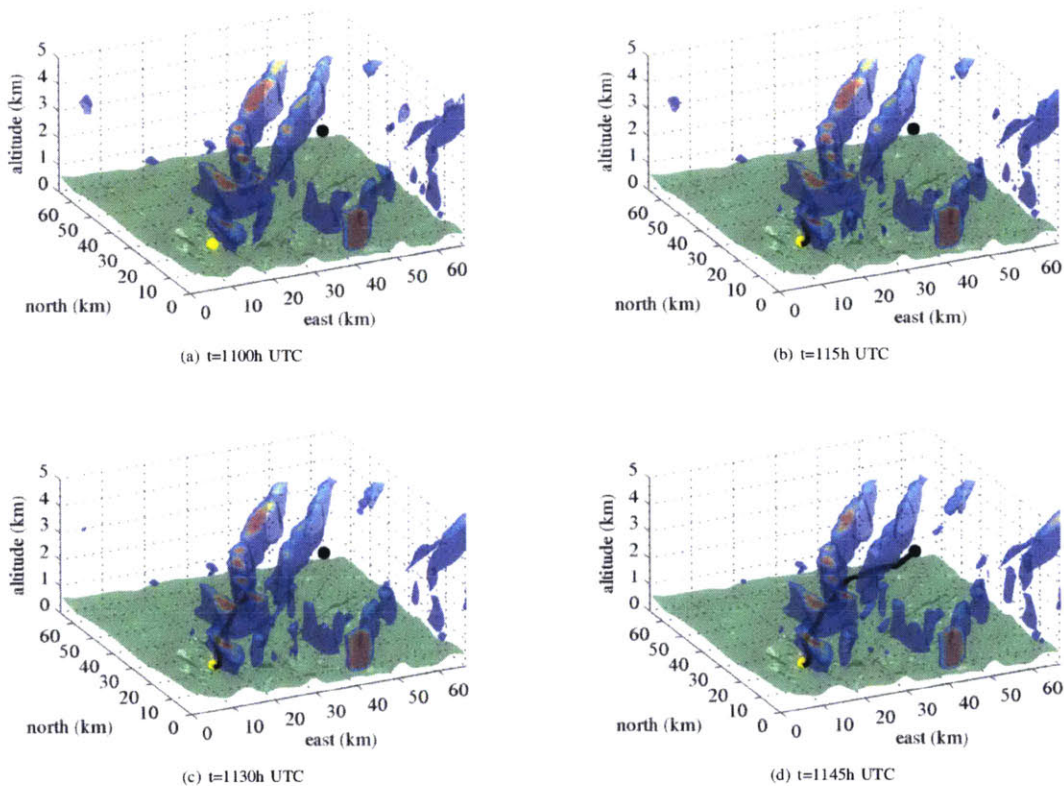


Figure 2-5: Simulation results from Chakrabarty and Langelaan[24] showing four snapshots in time over the course of the flight. Note that the vehicle follows the color coded lift regions as they evolve over time. Figures used with permission from authors.

need to be precomputed and stored for a spanning set of prevailing wind speeds and headings. Although the use of a numerical solver for wind field estimation that explicit considers the terrain offers promise for obtaining urban wind field estimates, the WRF cannot handle the dense structure of urban environments. Fortunately, other CFD packages offer suitable alternatives.

Although no previous results have addressed planning over urban wind fields, several efforts have been made to generate wind fields using CFD solvers to characterize their effect on UAV flight performance. Galway et al. [43, 40, 41, 42] investigated the effect of wind fields within a sparse urban environment on the path following performance and control effort of both a single-rotor and fixed-wing UAV. In an effort to avoid recomputing the wind field for every new environment, a library of 3D wind field primitives was built using the ANSYS CFX CFD package. The library was composed




Figure 2-6: Two 2D vector plots showing wind fields on a horizontal plane at half the structure height from the 3D simulation environment used in Galway et al. [42]. The freestream wind speed was 8.5 m/s flow from bottom to top. Flow direction is shown in the direction of the arrow heads with white being higher speeds and black being lower speeds. Note the interaction between the two structures on the right that acts to speed up the flow through the canyon. Figures used with permission from authors.

of wind fields for small groups of structures with one or two buildings over a set of prevailing wind conditions. The objective of the work was to combine these wind field primitives to generate an accurate estimate over a more complex environment with many structures. An example of these wind field solutions can be seen in figure 2-6.

Each of these wind fields was computed using an unsteady LES simulation with a 0.1 s time step over a total of 20 s. The resulting time series of wind fields was then post processed to find a bounding wake region that contained the time varying portion of the wind field. Once this wake region was identified, it could be directly




Figure 2-7: An illustration of the 3D simulation environment used in Galway et al. [42]. The prevailing wind is approximately 8 kts from the East. Note the wake boundary regions outlined on each structure. Figures used with permission from authors.

combined with other elements of the library as long as their wake regions did not overlap. Further work was done to combine wind fields with overlapping wake regions in special cases. An example of a wind field composed of many library elements, some of which overlap, can be seen in figure 2-7.

By simulating the vehicle’s flight along a given target trajectory through a pre-computed wind field unknown to the vehicle, Galway et al. further demonstrated that the forces generated by the vehicle’s interaction with the wind field had a significant effect on the control effort and trajectory following performance of both the single-rotor and fixed-wing UAV. Although Galway’s approach allows for approximate wind fields to be generated for this specific subset of sparse urban environments with building groups no larger than one or two structures, these constraints on structure density and geometry prevent the approach from capturing wind fields in more complex and dense environments found at ground level in a typical city. Although no attempt at trajectory planning was made in this work, it is also unclear how a raw unsteady wind field simulation would be used in a planner because even a wind field with a dominantly periodic evolution would need to be synchronized with the current conditions.

Orr et al. [80] and Gross et al. [47] made several efforts investigating UAV flight in

urban environments through an Air Force Research Laboratory (AFRL) project called the Cooperative Operations in Urban Terrain (COUNTER). This work focused on the effect of wind fields over a small group of buildings on a fixed-wing UAV's ability to reach a set of waypoints. The fixed-wing UAV used in simulation for this work had a mass of 9.5 kg, a wing span of 1.8 m, and a cruising speed of 13 m/s. The simulation domain was a small grouping of 15 structures that composed a town square. The tallest of the structures was three stories. Steady wind field simulation was done with a CFD package called the Air Vehicles Unstructured Solver (AVUS). The simulated wind field had a prevailing wind speed of 4.6 m/s and southerly heading and although no detailed analysis was presented on the effects of the wind field on the vehicle's path following performance, it was noted that the addition of the wind field prevented the UAV from reaching several of the waypoints in its path. In this case, the addition of the wind field into the simulation caused the vehicle to exceed its maximum airspeed along the target trajectory and forced it to abandon the unreachable waypoints.

Similar to the work of Chakrabarty et al [24] and Galway et al. [43, 40, 41, 42], the approach taken by Orr et al. [80] and Gross et al. [47] has the benefit of explicitly considering the effect of the terrain on the wind field. However, relative to our application, it improved upon these approaches by doing so for a dense urban environment without the need to combine wind field primitives. A potential drawback lies in the questionable ability of a steady CFD solver to accurately capture the unsteadiness present in an urban environment. Although it might be able to capture the mean flow velocities, it does not provide any uncertainty estimate over the map.

White et al. [103] used steady and unsteady CFD simulation as well as wind tunnel and in-situ measurements to investigate the ability of a small fixed-wing glider to soar in the rear wake region of a structure. The wind vectors over a vertical 2D plan can be seen in figure 2-8 and show good agreement between the measured data and simulation data. Notably, the steady $k - \epsilon$ RANS solution has strong agreement with the wind tunnel measurements, and the LES simulation required the inclusion of inflow turbulence to match these results. Given the fixed-wing UAV's lift characteristics found in a separate wind tunnel experiment, it was determined that this UAV

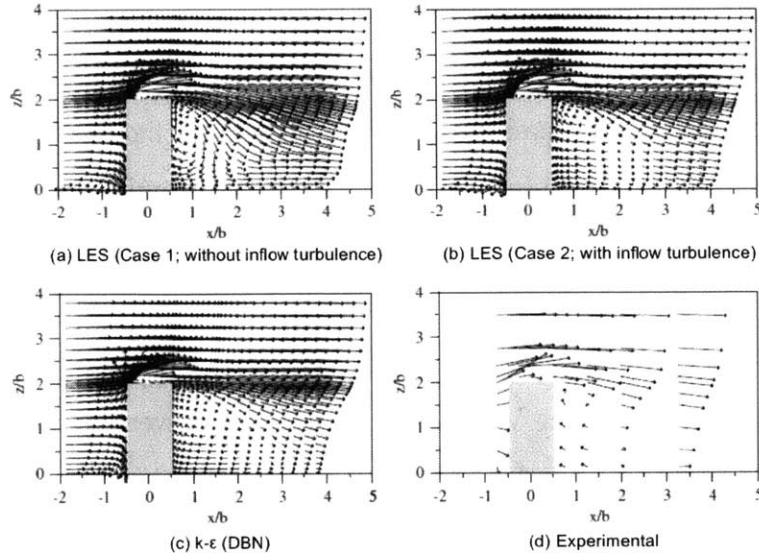


Figure 2-8: White et al. [103] presents the velocity vectors from numerical simulations comparing the RANS and LES turbulence models (a-c) and the in-situ experimental data (d). Note the similarity between the RANS simulation and the LES simulation with the modeled inflow turbulence. Figures used with permission from authors.

would be able to effectively soar in the wake region of this building. Unfortunately, no planning or flight simulations were performed to justify this conclusion, but this work remains an example of the ability of CFD simulation to compute accurate estimates of complex wind fields close to structures at the resolutions necessary for UAV flight.

Sutherland [99] also studied UAV flight performance in the wake region of a single structure, but focused on quadrotor flight performance while investigating the effective differences between using a RANS and LES turbulence model. This work used the OpenFOAM CFD solver to generate a set of wind fields around a single structure with dimensions $20 \times 20 \times 120$ m. In an effort to ensure the accuracy of the wind field estimates, much of this work was directed towards a series of sensitivity analyses on spatial and temporal resolution as well as verification studies. As in Galway et al.[42], a wind field database was constructed in order to decouple the 30 s runtime of the flight simulation and the 0.97 and 17.7 hour wind field computation time for the RANS and LES models, respectively. Once the wind field library was built, it was used to analyze the performance of a quadrotor position controller in various locations around the structure for both CFD methods. In summary, the LES simu-




Figure 2-9: Raza and Etele [87] show a vector field from a wind simulation of the flow around a single structure and the trajectory waypoints in their simulation. Figures used with permission from authors.

lations did significantly better resolving the innately unsteady turbulent fluctuations and therefore its wind fields had a more significant effect on the quadrotor’s flight performance. Sutherland also notes that the additional complexity of LES means that it should primarily be used in the pursuit of designing and testing autonomous control algorithms for multirotor UAVs on the order of 0.5 m in size and 2 kg in mass. Building further on these results, Raza and Etele [87] used the same wind field database to characterize the performance of several quadrotor position controllers while flying in the structure’s wake region. Again, these authors concluded that LES was superior to RANS in its ability to accurately model the smaller-scale perturbations placed on the quadrotor due to turbulence, but also resulted in significantly increased solution times.

Cybyk et al. [31] performed a similar investigation into the effects of wind fields on fixed-wing flight performance over a multikilometer portion of Baghdad, Iraq. The terrain geometry was resolved to 1 m and the wind field to 6 m resolution. The unsteady, LES wind fields were generated with the FAST3D-CT CFD package. Harms et al. [49] presented a validation study of FAST3D-CT against the data from a large field trial in Oklahoma City summarized in work by Allwine et al. [10, 8, 9, 16].




Figure 2-10: An illustration from Cybyk et al. [31] of the Procerus Unicorn UAS and the simulated mission route through downtown Baghdad with a prevailing wind heading of 20 degrees and prevailing wind speed of 3 m/s. Figures used with permission from authors.

The UAV used in simulation was the Procerus Unicorn, a powered glider with a 1 m wing span that weighs approximately 2 kg. An example trajectory from a single simulation along with an image of the UAV is shown in figure 2-10. This trajectory was simulated without wind, with a 3 m/s prevailing wind speed using only translational wind velocity, and with a 3 m/s prevailing wind speed with both translational and rotational wind velocities. Although no planning was done and a detailed analysis of the UAV's flight performance was not performed, these few trajectories generated significant differences in the vehicle's control effort, position, and attitude.

In a slightly different application, the Navy has used both RANS and LES CFD solvers to estimate wind fields near structures for flight simulation and the analysis of helicopter and fixed-wing flight performance. Bogstad et al. [15], Zan [109], and Forrest et al. [37] used CFD to create an extensive ship-airwake database for a ship-specific flight simulator to train helicopter pilots for landing and takeoff from various points on a ship. Similarly, Crozon et al. [30] used CFD to analyze the interaction

between the rotor and the frigate deck in a single simulation.

2.2 Minimum-Energy Planning

A critical component of the quadrotor’s ability to exploit an urban wind field for improved flight performance is a planner that will choose the appropriate path through the environment. The soaring community has thoroughly addressed this problem using minimum-energy trajectory planning for its ability to encode the desire not only to avoid flying into headwinds and catch tailwinds, but to provide vertical lift and exploit shear layers in the wind field. Although a quadrotor cannot generate lift as efficiently as a fixed-wing vehicle, the same inherent connection exists between a minimum-energy trajectory and the ability to efficiently fly through wind fields. The following section presents the relevant details of closely related work from the soaring community that informed our planning algorithm.

Chakrabarty et al. [23] developed a graph-based planning method that uses A* search to plan over the domain using a feasible action set. The edge cost is the energy consumption along that edge in the graph. A weighting between the Euclidean distance heuristic and energy consumption is proposed to tune the behavior of the planner. Figure 2-11 shows the behavior of the planner for different weightings of the heuristic over a simple domain.

In an extension of this work, Chakrabarty et al. [24] considered time-varying, complex wind fields over a mountain range. Although the wind field estimation component of this work was previously discussed, this is also one of the few examples of UAV path planning in time-varying wind fields. To circumvent the complexity of planning over time varying wind fields, a kinematic tree algorithm, originally presented in Langelaan [57], is extended to use an explicit representation of time in a receding horizon framework. This planning approach minimizes the distance to the goal while maximizing the total energy.

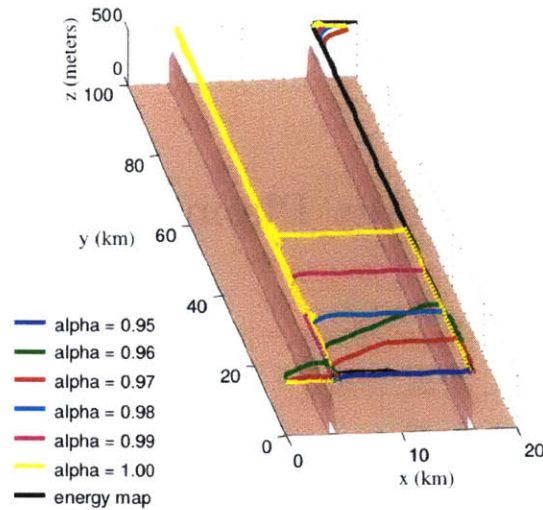


Figure 2-11: A comparison from Chakrabarty et al. [23] of A* paths with varying weights on minimum-distance and energy consumption along with the ideal path through this environment and cost map. Figures used with permission from authors.

2.3 Wind Measurement

Although it is not yet clear how to incorporate on-board wind velocity measurements into a wind field estimate, it is intuitive that an accurate local measurement could provide useful information for a trajectory planner or wind field estimation algorithm. This is reinforced by the work of Alexander and Vogel [6] which investigated the ability of several species of birds to align themselves with their intended direction of travel and mitigate the effects of wind gusts using measurements of the wind velocity along their flight trajectory. Additional work was done by Gewecke and Woike [44] and Brown and Fedde [19] demonstrating the connection between the airflow over avian feathers and the bird’s steering impulses as well as the ability to predict stall and measure airspeed. Local wind measurements could also provide a check on the accuracy of the current global wind field estimate within that region of the map. The following section presents the relevant details of closely related work from the UAV community on on-board wind velocity estimation and measurement. A few of these efforts have focused on single-rotor MAVs, but the majority of the work considers fixed-wing vehicles.

Although larger vehicles can afford to carry complex and heavy wind sensors such as LIDAR, SONAR, and sophisticated ultrasonic anemometers, MAVs have strict weight and power budgets that restrict their sensor payloads. Accordingly, much work over the past decade has focused on eliminating the need to carry an additional sensor specifically for wind velocity measurement. Mayer et al. [69] and Reuder et al. [89] demonstrated a straightforward example of this by constraining a powered glider to a circular flight path with a constant thrust and angle of attack in order to use the GPS position measurements to estimate wind velocity. Although successful in both simulation and a field trial, the constraints imposed on the vehicle's trajectory prevent it from being more broadly applicable.

Another technique to circumvent the need for an additional sensor is to use a model of the vehicle dynamics or actuators. Neumann and Bartholmai [77] used measurements from a quadrotor's IMU and GPS along with a drag model to estimate the wind speed. Validation flights were done using a ground truth anemometer over the course of a 20 minute flight within a single $30\text{ m} \times 30\text{ m}$ area. Although both the wind speed and heading roughly tracked the ground truth values, errors of more than 60° in heading and approximately 50% in speed were present. The RMSE of the wind speed was 0.6 m/s while hovering and 0.36 m/s while flying. The RMSE for the heading was 14.02° while hovering and 14.77° while flying. Similarly, Cho et al. [26] used airspeed and GPS measurements on a powered glider in a Kalman filter framework to estimate wind speed. Along these same lines, Myschik et al. [75] used GPS position, magnetometer heading, IMU measurements, and control surface deflection to calculate the angle of attack, sideslip angle, and resulting wind speed of a fixed-wing UAV.

Other work in this area focused on wind estimation without additional sensors on-board single and multi-rotor vehicles. Moyano Cano and Javier [72] developed a model of a quadrotor's pitch response to horizontal wind and used it to characterize the horizontal velocity profile in offshore wind farms. Marino et al. [68] attempted a more minimalist approach to estimate oncoming flow velocity magnitude and direction using power consumption measurements from each of a quadrotor's four rotors.

Although it was found that quadrotor endurance was improved by an approximate minimum of 50% with updraft velocities of greater than 5.5 m/s, the oncoming flow vector could only be resolved under limited conditions. Finally, Divitiis [35] applied the same method of using a dynamics and thrust model to a single-rotor MAV with the addition of a neural-network to estimate wind velocity from these inputs.

Although many of these approaches were successful in estimating the wind velocity in their respective domains, they all suffer from the need to either restrict vehicle motion or to build and maintain an accurate dynamics model. It is impractical to tightly constrain the motion of a UAV in a dense urban environment and the relatively laborious task of updating a dynamics model would need to be repeated for each new vehicle and as the UAV payload or configuration changed.

Additional efforts have been made to leverage the Pitot tube airspeed sensor frequently found on fixed-wing gliders. Chan et al. [25] demonstrate three real-time, filtering approaches to estimating the local wind velocity with the Spoonbill UAV system, a powered glider developed by the Remotely Piloted Vehicle and Micro-Satellite Research Laboratory (RMLRL). The onboard sensors consisted of a GPS, AHRS, airspeed sensor, angle of attack sensor, and sideslip angle sensor. A test flight was performed with the UAV flying a 1 km \times 1 km box pattern while ground truth wind data was collected by an ultrasonic anemometer. Strong agreement was found between the estimated and measured speed and headings with a maximum of 10% error in speed and only a few degrees in heading. Other efforts were made by Patcher et al. [82] and Zhang [110] using GPS inertial and air data measurements to estimate horizontal components of wind speed and vehicle heading angle. Langelaan et al. [58] improved on this approach by estimating wind velocity, its rate of change, and spatial gradient. While these approaches work well for fixed-wing gliders, they cannot be applied to any vehicle that does not have a minimum forward airspeed since a Pitot tube needs the resulting dynamic pressure.


Given the sensor's inherent constraints, several have attempted to replace or improve upon the Pitot tube as an airspeed sensor. One alternative to the standard Pitot tube is a multi-port pressure probe that provides velocity and turbulence information

in several dimensions. Kroonenberg and Martin [102], Reuder and Jonassen [90], and Spiess et al. [98] used 5 and 9-hole pressure probes to successfully measure atmospheric wind velocity and turbulence on the M^2AV and SUMO platforms. Both are fixed-wing UAVs with sensor suites that also included a GPS and IMU. A similar approach was taken by Mansour et al. [67] with a 7-hole pressure probe to characterize the atmospheric boundary layer near wind turbines. Their custom pressure probe consisted of a 20 mm hemispherical probe head and allowed for a relatively small dynamic pressures and large incidence angles of up to 70 degrees. Although exceedingly accurate, these pressure probes still require a minimum forward airspeed and restrict the wind's incidence angle. Both of these constraints are problematic for a quadrotor measurement platform.


In a relatively extreme effort to improve measurement accuracy, Bange and Rainer [11] used a full-scale helicopter towing a sensor array to measure wind velocity and turbulence in an urban environment. Although the system performed well, the size and cost of the system prohibits its use on most UAVs and all MAVs.

In an effort to remove the dependence on both the Pitot tube and magnetometer, both Rodriguez et al. [93] and Moore [71] used a downward facing camera and optical flow algorithm to detect the slip angle of a fixed-wing UAV. Combining this with ground speed estimates provided by GPS allowed the wind speed to be calculated. The approach was demonstrated both in simulation and in a single test flight. In both cases, the system was only used to estimate a global wind speed and heading.

In an effort closely resembling the approach taken here, Müller et al. [73] developed a 2D, low-power flow sensor based on MEMS differential pressure sensors. The indoor robotic blimp and flow sensor used by Müller et al. [73] are shown in figure 2-12b. Because the blimp platform was designed for indoor flight, the lack of wind meant that any airflow measured by the vehicle would be due to the motion of the vehicle or the flow induced by its rotors. Given the assumption of a static wind field, the experiments done by Müller et al. [73] demonstrated that an accurate flow sensor could be used to recover the vehicle velocity. Due to the blimp's limited payload capacity, a custom flow sensor was developed based on the SDP600 commercial differ-



(a) Validation Test



(b) Vehicle and Sensor Platform

Figure 2-12: Work by Muller et al. [73] shown in panel 2-12a the validation test and the robotic indoor blimp with the mounted MEMS flow sensor in panel 2-12b. Figures used with permission from authors.

ential pressure sensor manufactured by Sensirion. The factory housing was removed and tubes were mounted on the sensor face in order to focus the sensor’s sensitivity on a single axis. Two of these units were mounted perpendicular to each other and positioned both longitudinally and laterally with respect to the vehicle body frame. A model was built to perform probabilistic velocity estimation from the sensor measurements. The resulting successful velocity estimation from the algorithm as well as a naive estimator and ground truth data over the course of a test flight can be seen in figure 2-12a. Although this approach works well for indoor environments, it is not clear how well it would extend to outdoor flight and how the induced flow from a quadrotors propellers would affect the solution.

In a similar effort, Yeo et al. [107, 105, 106] demonstrated a unique multi-port pressure probe system distributed over a quadrotor platform such that it could measure both horizontal and vertical flow velocities. This work’s primary objectives were to both estimate vehicle velocity in an indoor environment and estimate the strength and position of downwash from other nearby quadrotors. As seen in figure 2-13, this was achieved by mounting a total of six flow sensors with four vertical flow sensors mounted on each of quadrotor’s arms and two horizontal flow sensors mounted perpendicular to each other and directly above the center of the vehicle. The system was demonstrated in a series of indoor flight tests with a single rotor mounted on a sliding carriage to simulate the presence of another quadrotor while using motion capture




Figure 2-13: A diagram from Yeo et al. [106] showing the four vertical flow sensors and two horizontal sensors mounted on the DJI quadrotor. Figures used with permission from authors.

as ground truth. The comparison of the resulting velocity estimates can be seen in figure 2-14. Although this system was successful at estimating a quadrotor's ground velocity in an indoor environment, its maximum measurement velocity of 5 m/s limits its usefulness in larger outdoor environments with strong winds. The effect of large pitch angles on the accuracy of the estimate is also unclear.

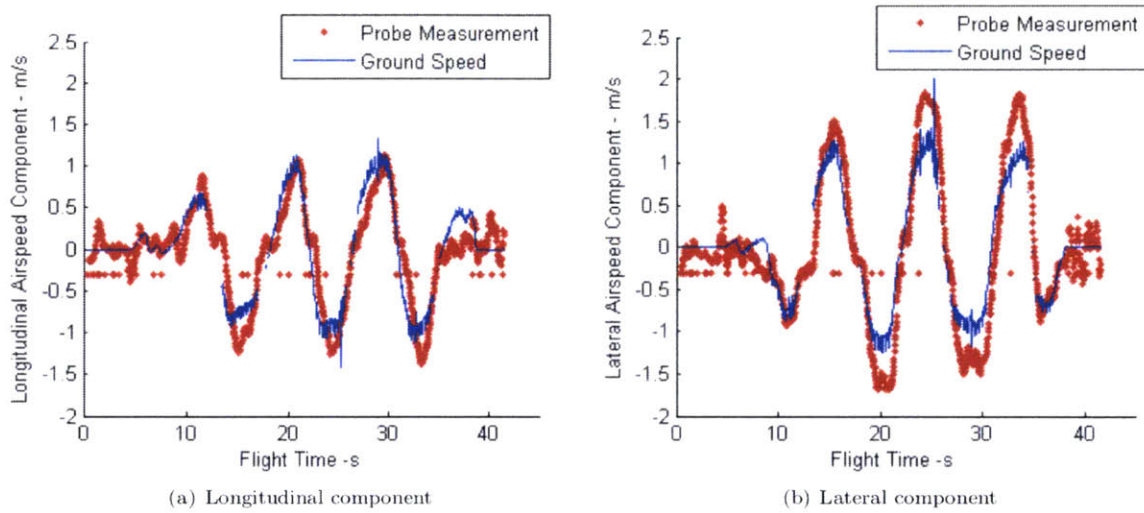


Figure 2-14: Flight test results from Yeo et al. [106] showing the performance of the velocity estimation against a ground truth reference provided by motion capture. Figures used with permission from authors.

Chapter 3

Onboard Wind Measurement

Although it is not yet clear how to properly incorporate on-board wind measurements into a wind field estimation algorithm, it is intuitive that on-board measurements could provide useful information about the local wind conditions as well as the validity of the current global wind field estimate. With this in mind, the following section presents initial results on the integration of an existing MEMS-based wind sensor into the quadrotor.

3.1 Wind Measurement Techniques

There are a range of common sensing modalities that can be used to measure wind speed, but most of these are unsuitable for use on-board a quadrotor. The simplest wind sensor is the ubiquitous vane anemometer. These have either cups or a windmill to measure wind speed and are commonly used to measure wind conditions on personal weather stations. Although their simplicity provides robustness, the inertia of the rotating components and bearing friction within the assemblies limit the dynamic response of these sensors and restricts their use to hovering in wind that varies slowly relative to the sensor dynamics. An additional drawback is that the vehicle's movement imparts forces on the sensor which induce artificial changes in the wind heading. Although this design's simplicity is attractive, these inherent weaknesses prevent it from being a suitable candidate for on-board wind sensing in dynamic urban wind

fields.

The pitot tube is another wind sensor commonly used in both hobbyist and commercial aviation. These sensors measure the dynamic pressure between a forward facing pressure port and a perpendicular, static pressure port and work well on fixed wing aircraft. The constant forward airspeed of a fixed wing vehicle creates the necessary dynamic pressure to obtain accurate wind speed measurements. Another drawback is the sensor's sensitivity to the wind's angle of attack with respect to the forward pressure port. Most Pitot tubes only work with a relative angle to the flow of no more than 15 degrees. These requirements preclude their use on quadrotors that might hover for long periods of time and can assume nearly any orientation with respect to the wind heading.

Hot-wire anemometers provide extremely fast response times by measuring the change in current required to maintain a constant power loss across a heated element exposed to the wind. Hot-wire anemometers can be purchased with varying degrees of complexity and measurement capability. At one end of the spectrum are single element sensors that only measure scalar wind speed and are commonly found in the hobbyist, commercial, and research markets at a range of sampling frequencies and resolutions. At the other end of the spectrum are more complex versions of these sensors created almost exclusively for the research. These sensors have multi-element arrays whose collective outputs are used to estimate the wind velocity in two or three dimensions. Although these more sophisticated devices are exceedingly accurate and have very low response times, they cost several thousand dollars and are too delicate to be flown on a quadrotor.

Ultrasonic anemometers are a relatively modern wind measurement technique. The ultrasonic anemometer has become the new standard in wind measurement for both personal and commercial weather stations but has not been adapted for use on a quadrotor for several reasons. These sensors use pairs of ultrasonic emitters and receivers to calculate the wind velocity by measuring changes in the speed of sound caused by varying the air velocity between the sensing pair. Arrays of these pairs can be used to calculate the wind velocity in two or three dimensions. Although 2D

ultrasonic anemometers can weigh as little 100 grams in the case of the LCJ Capteurs CV7, most weigh significantly more. Similarly, there are no readily available, 3D anemometers with the appropriate size, power consumption, and weight required by a MAV. Unfortunately, 2D ultrasonic anemometers only work within a very narrow range of pitch angles with respect to the air velocity. This limitation prevents them from being suitable as an on-board sensor for a MAV that regularly operates at significant pitch and roll angles. Given the deficiencies of each of these anemometer designs with respect to use on-board a quadrotor, there is a clear need for another approach.

3.2 Operating Principle

The sensor used in this work was originally developed by Piotta et al. [84] and was adapted for use on a quadrotor in an ongoing collaboration with MIT. Although MEMS-based anemometers have been considered a viable alternative by Robadey et al. [91], Kim et al. [54], and Kaltsas et al. [53], their delicate sensing arrays cannot be exposed to the elements without risking damage. Due to their fragility, MEMS-based sensors without protective housings are mostly confined to lab environments. To address this issue, Piotta et al. developed a wind sensor, shown in the right panel of figure 3-1, that uses a novel sensor housing to calculate the wind speed and heading based on the pressure differential across a cylinder as measured by two MEMS differential pressure sensors. The key to the design is in its use of a series of channels, shown in the left panel of figure 3-1, whose entrances are carefully placed around the diameter of the cylinder and feed each of the differential pressure sensor ports. The correct placement of the channels makes the diametric pressure across the cylinder a monotonic function with respect to heading over a 180 degree range. Once this has been achieved, two identical disks, rotated 90 degrees relative to each other, can be used to measure the diametric pressure along two axes. These two pressure values can then be used to compute the flow speed and heading with respect to the sensor frame of reference.

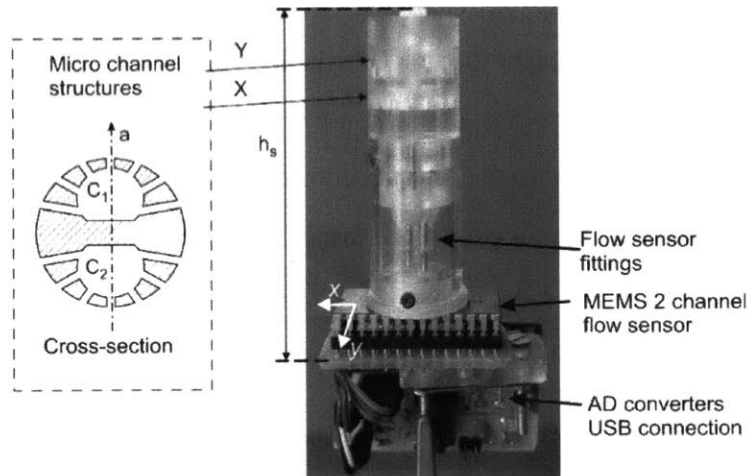


Figure 3-1: The left panel shows a cross section of the sensor and illustrates the channel structure that generates the necessary differential pressure measurements. The right panel is a photograph of the sensor used on the quadrotor. The upper, translucent portion is the plastic sensor housing that is built in several stacking segments. The MEMS flow sensor designed and fabricated by Piotto et al. [84] is at the base of the sensor housing in a dual in-line integrated circuit package.

3.3 Calibration and Validation

Previous work by Bruschi et al. [20] outlines the process to map the sensor's output voltage to a differential pressure. This calibration is insensitive to the sensor's operating conditions and did not have to be repeated for use on a quadrotor. A second calibration that maps from the two measured differential pressures to a corresponding speed and heading was first performed by Bruschi et al. [20] in a small benchtop wind tunnel shown in figure 3-2 and later supplemented by data from the full scale Wright Brothers wind tunnel at MIT. The data and fit from the small benchtop wind tunnel can be seen in figure 3-2.

The tests in the full scale Wright Brothers wind tunnel were performed with the sensor mounted on-board the quadrotor and the quadrotor rigidly mounted on a tripod. As shown in figure 3-3 the sensor was elevated 22 cm above the rotor plane to isolate it from turbulent effects. Wind speed was recorded with an off-board ultrasonic anemometer at 4 Hz and all other sensor data was recorded on the on-board dual core i7 computer.




Figure 3-2: Results from tests performed in a benchtop wind tunnel with a 10 cm diameter cross section. The top panel shows measured versus actual wind speed over the full set of headings between 0 and 360 degrees. The bottom panel shows the measured versus actual angle. Both plots show standard error.

Two sets of trials were conducted over a range of wind speeds and headings of 0 and 45 degrees. Each wind speed was tested with the propellers on and off to capture their effect on the measurement accuracy. The first test sought to characterize the sensor while the vehicle was hovering and received an intermittent gust and thus maintained a 0 degree pitch and roll angle. Measurements were taken over wind speeds of 6, 8, 10, 12, 14.7, 16, 20, and 25 m/s for each of the two headings. The raw differential pressures generated within the sensor housing during these tests can be seen in figure 3-4. Note that the magnitude of the differential pressure increases for both axes when the vehicle has a yaw of 45 degrees, but only for one of the axes while the sensor axis is aligned with the wind heading.

The measurements from the small wind tunnel are densely sampled over both speed and heading, but do not exceed 8 m/s, while the measurements from the Wright Brothers wind tunnel were more sparsely sampled with a max speed of 25 m/s. With their relative strengths in mind, both datasets were used to create a new calibration between the two differential pressures and the corresponding wind speed and heading.





Figure 3-3: Panel (a) shows a diagram showing the sensor mounting position on-board the quadrotor. The sensor height above the rotor plane, h , was 22 cm. The propellers center distance, d , was 20 cm. The propeller diameter, ϕ_d , was 25.5 cm. Panel (b) shows the quad mounted on the test fixture in the Wright Brothers wind tunnel.




Figure 3-4: The measured differential pressures given a range of wind speeds tested in the Wright Brothers wind tunnel. The pressures for the X and Y axis sensors, and their standard errors, are reported in the left and right panels, respectively. The plot for each axis shows pressure values with the propellers on and off. In this case, the thrust values were set to the vehicle’s nominal hover thrust of 55%.

The resulting fitted speed, u , and heading, θ , calibrations are given in equations 3.1 and 3.2, respectively.

$$u = 2.15(P_x^2 + P_y^2)^{0.235} \quad (3.1)$$

$$\theta = \arctan2(P_y, P_x) \quad (3.2)$$

The resulting wind speeds and headings for both the propellers on and propellers off are shown in figure 3-5. As can be seen, the induced flow from the propellers has almost no effect on the heading and minimal effect on the speed measurement below 10 m/s. The discrepancy at lower wind speeds is likely due to the smaller relative difference between the induced flow through the rotors and over them. In this case, the induced flow through the rotors at hover has a velocity of 12.5 m/s. As the momentum of the horizontal flow decreases, it is more easily diverted by this downward flow and the measurement is distorted. As of now, it is not clear how to improve the sensor's response at lower wind speeds. Finally, the underlying variation in the heading is likely due to setup error in the experiment and the inherent inaccuracy of the tripod orientation mechanism.

The second test was done over a narrower range of wind speeds that only covered the vehicle's maximum airspeed and was meant to approximate forward flight conditions. Each wind speed trial was performed with a yaw of 0 degrees and the appropriate pitch angle, thrust command, and pitch command were derived from past, unconstrained flight tests in the Wright Brothers wind tunnel. The average pitch angle during these tests can be seen for each wind speed in figure 3-6. The full test configuration for each speed trial can be found in table 3.1.

The final set of estimated average wind speeds and headings for the forward flight test are shown in figure 3-7. In this case, the propellers apply a nearly constant offset to the measured wind speed, outside of the the highest wind speed. Again, the innaccuracy of the tripod mechanism imparted some underlying variation to the heading measurements, but the propellers still have little effect on heading.

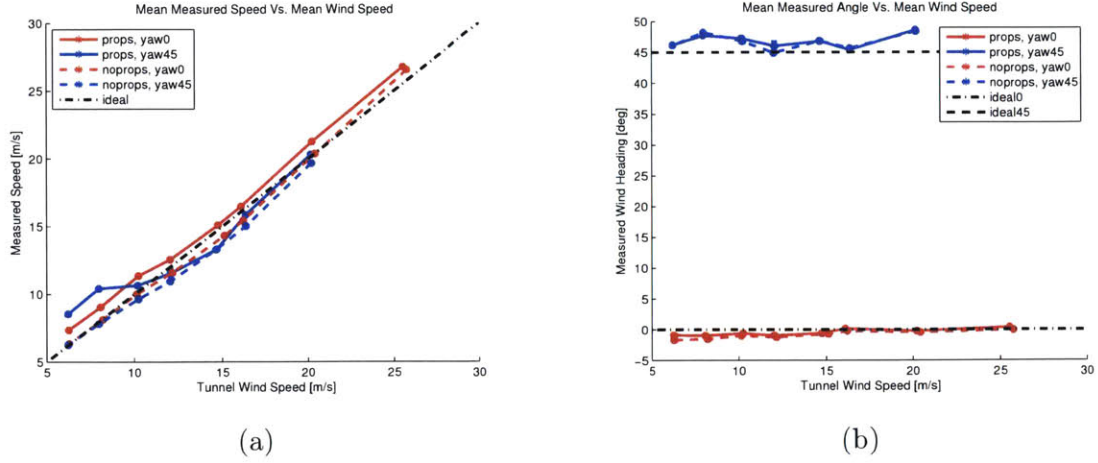


Figure 3-5: The estimated speed and heading, along with their standard error, are shown in the left and right panels, respectively. Each of these values were measured for both 0 and 45 degree vehicle orientations with respect to the wind and over a range of wind speeds between 6 and 25 m/s. Note that the induced flow through the propellers has a much larger effect at lower wind speeds.

Speed [m/s]	Pitch [deg]	Pitch [0,1]	Thrust [0,1]
2.00	3.04	0.15	0.55
4.00	6.88	0.29	0.54
6.00	11.46	0.43	0.51
8.00	16.79	0.56	0.50
10.00	22.87	0.70	0.56
12.00	29.69	0.84	0.72
14.50	39.26	1.00	1.00

Table 3.1: This table lists the testing configuration for the forward airspeed tests in the Wright Brothers wind tunnel. These values were computed for each target speed from prior, free-flight experiments.

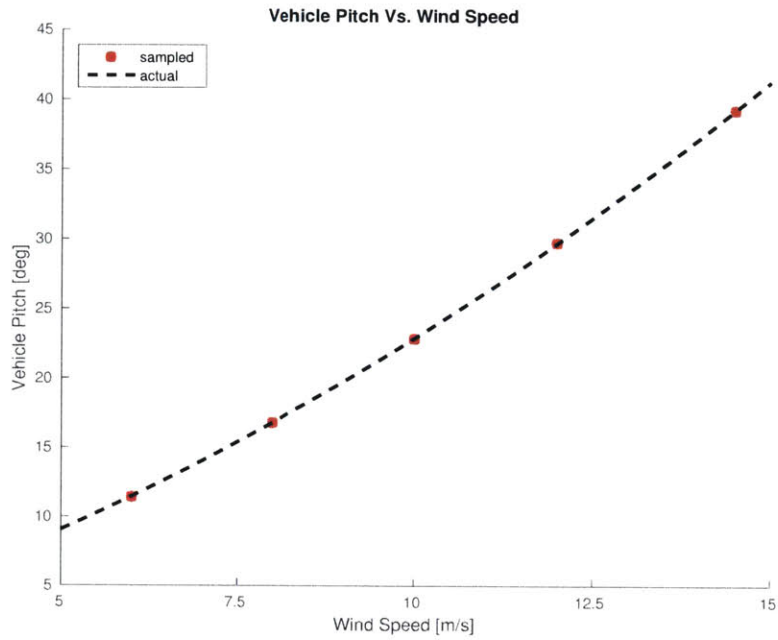


Figure 3-6: This curve was generated during a separate experiment where the quadro-rotor was in free flight in the Wright Brothers wind tunnel over a range of wind speeds. The resulting data generated the fit shown here that provides the natural pitch angle for the vehicle for a given wind speed. The conditions for the pitch tests for the wind sensor were found by sampling this curve at the target wind speeds and referencing logged data for the appropriate motor commands.

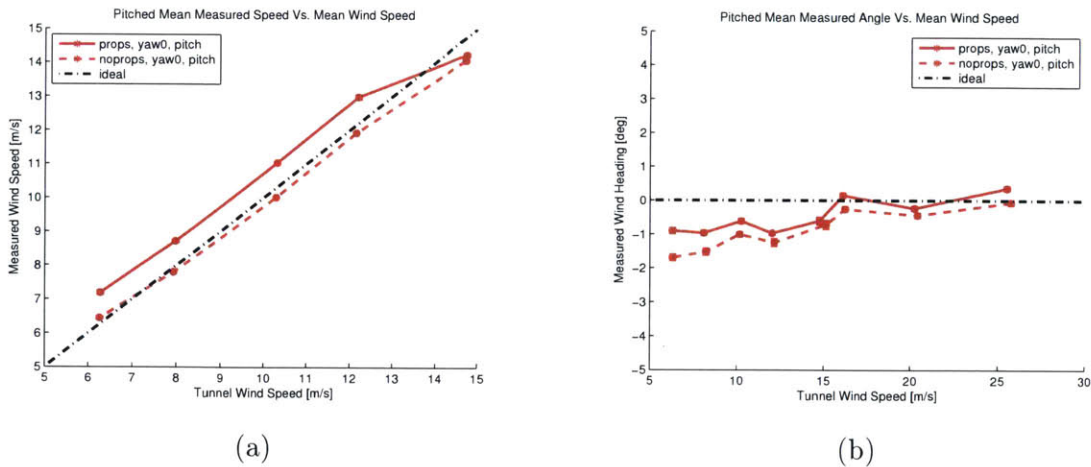


Figure 3-7: These plots show the estimated wind speed and heading in the left and right panels, respectively. Note the constant offset created by the induced flow through the spinning propellers. The error in heading is likely due to experimental error in the alignment of the vehicle and fixture during the trials.

Chapter 4

Wind Field Simulation

As previously discussed, this work relies on CFD to compute the wind field using a 3D model of the environment and a prevailing wind condition at some altitude. Although CFD is a well-established technique for urban wind field estimation, there are many possible wind solvers available and a number of important modelling constraints and conditions that must be met in order to ensure an accurate wind field estimate. This chapter presents a brief summary of the simulation environment, the prevailing wind conditions in the selected environment, and the CFD solver used in the work presented in this thesis. For reference, a thorough discussion of best practices in wind field simulation can be found in Appendix A.

4.1 Simulation Environment

All wind field simulation and trajectory planning was done on the portion of the MIT campus shown in figure 4-1. This region offers a combination of green spaces and dense urban canyons. The average structure height is approximately 30 meters, with a blend of regular, rectilinear structures and more irregular structure geometries. Although they are outliers, the shortest structure is 7 meters and the tallest, the Green building marked in figure 4-1, is 100 meters. As discussed below, the absence of a regular rectilinear grid layout and existence of irregular structure geometries makes this environment a challenge for less sophisticated CFD techniques that might

otherwise be sufficient.




Figure 4-1: A satellite image of the portion of MIT campus used in the work presented in this thesis. The red circle indicates the location of the MIT weather station mounted on the roof of the Green building. Note the mix of dense urban canyons, more open green spaces, and the completely open region above the river. Also note the variety of structure geometries.

4.2 Prevailing Wind Conditions

The prevailing wind conditions define the inflow boundary that drives the flow within the simulated environment. Therefore, the prevailing wind speed and heading distributions must be estimated in order to capture a set of wind fields spanning the probable conditions. These distributions can be found using a history of measurements, and the simplicity, reliability, and ubiquity of weather stations frequently make them the best source for this historical data. Despite their strengths, urban weather stations are challenging to site properly and often suffer from both noisy and biased measurements. Although a typical urban environment usually contains several sta-

tions with publicly available data, it is often the case that many of these stations are unusable. Aside from stations that have either gone out of calibration or failed entirely, the most common cause of these issues is poor siting relative to the nearby structures and obstacles which add bias and noise to the station's measurements. To avoid these issues, weather station's must be selected after carefully assessing past performance given a history of measurements.

The primary goal for selecting a weather station is to obtain wind measurements that are free from the influence of specific surface anomalies and are representative of the spatially averaged wind conditions upstream of the target area. This means wind measurements should be made above the roughness sublayer, but within the internal boundary layer of the selected urban climate zone (UCZ). The following section presents a brief summary of Oke's [78] recommendations on urban weather station selection and siting and their relevance to the weather station on top of the MIT Green building shown in figure 4-1.

4.2.1 Mixing and Reference Heights

The height of the weather station relative to its surrounding environment is critical to achieving wind measurements that are independent of the effects of any given structure on the flow and represent the local, spatially averaged conditions. This requires that the appropriate height be selected on a case by case basis. For relatively dense urban environments, Oke [78] states that a meaningful spatial average is only obtained when sensor height is well above the greater of 10 meters or 1.5 times the mean structure height with special care being taken to avoid the wakes of particularly tall structures. Fortunately, the MIT weather station's 100 meter altitude is significantly higher than Oke's [78] recommendation of 1.5 times the 30 meter mean structure height.

4.2.2 Proximity to Obstacles

In open terrain, Oke [78] suggests that obstacles be no closer than 10 obstacle heights to the weather station, but qualifies this by stating that it is usually not possible in urban environments, and would actually ensure that the measurements were not representative of the UCZ of interest. As such, the guideline reverts to those discussed above with measurements taken at the larger of 10 meters or 1.5 times the mean structure height, which the MIT weather station comfortably achieves.

Another consideration for sensors mounted on buildings is their proximity to the perturbed zone above the rooftop. Wieringa [104] suggests that rooftop sensors should be mounted at a height equal to the roof's maximum horizontal dimension in order to avoid this turbulent zone. As Oke [78] notes, this would require an expensive mast system with support guys, but is technically the correct approach. Further qualifications admit the impracticality of this with additional encouragement to mount and select sensors that are not overly compromised by their proximity to their supporting structure. Unfortunately, the MIT weather station is not only mounted in relatively close proximity to the roof, but the roof also has several large obstacles interfering with the flow to the anemometer.

As shown in figure 4-2, the flow across the MIT weather station is obstructed by two radar domes on the N.E. and S.W. corners of the building. Because of their size and proximity to the weather station, these obstacles might introduce significant noise or bias the wind measurements. The station's relatively low elevation above the rooftop is another concern and probably places it within the perturbed zone, or recirculation zone, of the rooftop. The sensor is only mounted approximately 5 meters above the rooftop and the maximum horizontal extent is roughly 35 meters. Because these concerns raise questions about the validity of the station's measurements, an analysis of the station's historical data was performed to better assess the data quality.



(a) Green Building N.N.W. Perspective (b) Green Building W.S.W. Perspective

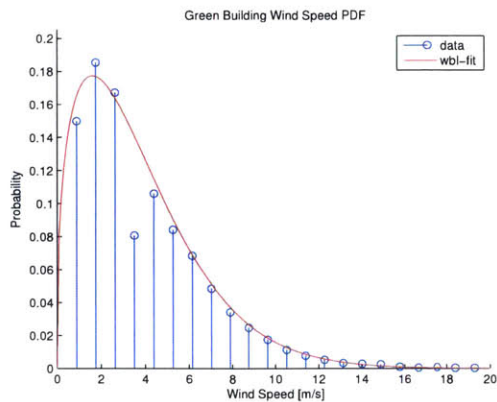
Figure 4-2: Two shots from on top of the Green building showing the positioning of the weather station relative to the roof, radar domes, and other obstacles.

4.2.3 Speed and Heading Distributions

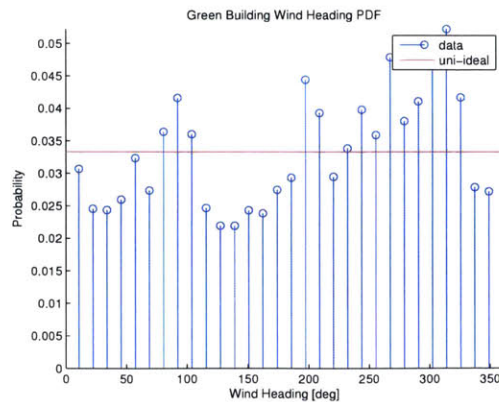
A useful way of checking the validity of the prevailing wind data provided by a weather station is to analyze the distributions of the wind speed, heading, and individual velocity components. Synoptic wind speed probability distributions can be modeled accurately as a Weibull distribution, as seen in equation 4.1, with both the κ and λ parameters constrained to be greater than zero. Justus et al. [52] and many others have demonstrated that, once zero velocity measurements are rejected, this same analysis can be used at lower altitudes in relatively open areas. It is also useful to note that when there are no obstacles and few terrain features, such as on the plains or in the open ocean, the wind heading distribution is uniform.

$$f_X(x) = \left(\frac{\kappa}{\lambda}\right) \left(\frac{x}{\lambda}\right)^{\kappa-1} e^{-\left(\frac{x}{\lambda}\right)^\kappa} \quad (4.1)$$

Despite real world data often resulting in an imperfect fit, the manner and degree to which the data violates the ideal case strongly informs the quality of the weather stations measurements. Given the data, the best fit model can be achieved using a maximum likelihood approach outlined by Dorvlo [36] and Seguro and Lambert [97]. In the absence of sufficient information, the shape parameter is often assumed to be 2 which simplifies the expression to a Rayleigh distribution. In general, the shape parameter is usually between 1.5 and 3, depending on the variability of the wind. Smaller values correspond to gustier conditions.



(a) PDF for Prevailing Wind Speed



(b) PDF for Prevailing Wind Heading

Figure 4-3: Prevailing wind velocity distributions in polar coordinates.

Recognizing the improper siting of the MIT weather station, a brief analysis was done on data taken at 5 minute intervals over the course of 2014. Given the significant reduction in wind found at night due to the cooling of the earth’s surface, only the hours between 8 A.M. and 6 P.M. were considered. Following the previously discussed guidelines, the best fit probability density functions for both the speed and heading can be seen in figure 4-3. Unfortunately, the station’s poor siting and nearby obstructions introduce significant error at approximately 3.5 m/s and at 135 and 45 degrees. These errors are most likely caused by the turbulence generated from the radar domes and cause the data to fail a goodness-of-fit test, but it is clear that the distributions bear a strong likeness to their ideal counterparts. As could be expected, there are significant reductions in the frequency at wind headings corresponding to the location of both of the radar domes. The best fit Weibull for the wind speed data has a shape parameter, κ , of 1.37 and a scale parameter, λ , of 4.15. This discrepancy likely reflects the turbulence generated by the radar domes and the sensors proximity to the turbulent region above the rooftop.

Despite its imperfections, the MIT weather station data offers the best available estimate of the prevailing wind conditions at the target site and the resulting speed PDF was used to characterize the expected conditions. Given the speed and heading distributions, wind speeds of 5, 7.5, 10, 12.5, and 15 m/s and headings of 90 and 270 degrees were selected. The speeds were chosen as a uniform sample of the non-trivial

wind speeds with non-zero rates of occurrence. Specifically, each of these wind speeds occur in approximately 27%, 11%, 3.5%, 1.3%, and 0.35% of all the measurements, respectively. Although the Easterly and Westerly headings are only marginally more likely, they were selected because it is common for wind in coastal cities to switch between two opposite headings as the land mass heats and cools over the course of the day [4].

4.3 QUIC

The Quick Urban and Industrial Complex Dispersion Modeling System (QUIC) [18] is a suite of fast response urban dispersion modelling tools that were originally developed at Los Alamos national labs to quickly calculate the exposure from the release of a chemical or biological agent in an urban environment. Although the QUIC software package contains a range of tools, the work presented in this thesis draws primarily on the QUIC-URB and QUIC-CFD wind solvers. The following section covers the details of these solvers, the results of some initial investigation into their applicability to the work presented here, and a validation study.

4.3.1 QUIC-URB

QUIC's goal of quickly modelling dispersion in urban environments fundamentally requires an abnormally fast CFD solver. In its initial release, QUIC addressed this requirement with QUIC-URB (Updated Rockle-style Building-aware), a parametric solver, based on the work of Rockle [92], that uses empirical algorithms and mass conservation to estimate the wind velocities around buildings. Unlike more traditional CFD approaches, QUIC-URB uses empirically derived parameterizations of the flow around structures and small structure groups to define an initial, time-averaged flow field without solving any components of the Navier-Stokes equations. The empirical relationships are primarily based on structure height, width, and length and the inter-structure spacing. Once the initial flow field has been defined based on the prevailing wind conditions and the structure geometry, it is forced to satisfy mass conservation.

This final step imparts physicality to the flow field in the form of smoothness and eddies.

In many ways, QUIC-URB delivers in its ability to compute urban wind fields in a fraction of the time of more traditional CFD solvers. Neophytou et al. [76] performed a comparison of three wind models using the data from the Joint Urban 2003 Field Campaign in Oklahoma City. This trial benchmarked both QUIC-URB and QUIC-CFD, a more precise CFD solver within QUIC that is discussed in greater depth below, against a high fidelity and computationally expensive LES solver. It was concluded that QUIC-URB was able to reproduce the *in situ* measurements nearly as well as the LES solver with solution times 2-3 orders of magnitude faster. Given these results, a brief investigation was done into QUIC-CFD's applicability to the work being done here. A $300 \times 300 \times 60$ meter cluster of buildings around the North Court on MIT campus, shown in figure 4-4, was modeled in QUIC as a test environment. Unlike the portion of Oklahoma City used in Neophytou et al. [76], these structures were irregularly shaped and not aligned on a rectilinear Manhattan grid.




Figure 4-4: A 3D model of a small region of MIT campus used by QUIC-CFD to compute a wind field given a prevailing wind speed and heading.

A horizontal slice of the QUIC-URB wind field at an altitude of 2 meters can be seen in figure 4-5. The prevailing wind is 15 m/s at 100 meters from a heading of

100 degrees North. As shown in the red ellipses, the solver generates two strongly non-physical wind features. The first feature is the recirculation zone in the upper right of the map that seems to ignore the existence of the building to the right of the figure. This is likely due to the solver imperfectly combining the parametric flow regions around each individual structure due to the irregular geometries and relative positions of the structures. The second feature is the sharp discontinuity in the middle of the figure. This is likely caused by the very irregular geometry of the structure on the left side of the figure. As previously stated, Neophytou et al. [76] was done in an environment that lacked any of these irregularities and therefore did not expose these issues. In the end, although QUIC-URB was able to find a wind field solution in approximately 24.7 seconds compared to QUIC-CFD's 293.9 second solution time for this environment, these clearly erroneous portions of the wind field prevented further consideration of QUIC-URB for the work presented in this thesis. Fortunately, QUIC-CFD provided a more robust solution for a reasonable increase in solution complexity.

4.3.2 QUIC-CFD

Developed by Gowardhan et al. [45], QUIC-CFD is based on the work of Chorin [27] and uses the steady Reynolds-Averaged Navier-Stokes turbulence model to generate a time-averaged solution. Unlike the standard $k - \epsilon$ turbulence model which uses two transport equations, QUIC-CFD uses a zero equation (algebraic) model based on Prandtl's mixing length theory [85]. In this case, the RANS equations are solved explicitly in time until steady state is reached using Patankar's [83] projection method. Unlike QUIC-URB, this approach is mostly agnostic to the structure shapes and configurations.

An identical test to the one performed with QUIC-URB was done with QUIC-CFD in the Northcourt environment within the MIT campus using identical prevailing wind conditions. The results of this test can be seen in figure 4-6. Unlike the solution provided by QUIC-URB, there are no clearly non-physical portions of the wind field and the entire flow pattern appears smoother and more intuitive. Although this is



Figure 4-5: QUIC-URB sample wind fields over a small region of MIT campus. The approximations taken in the solver introduce clearly non-physical regions within the wind field that are circled in red.

not a rigorous validation, it clearly demonstrates the need for a solver that addresses the physics of the flow and its interaction with the complex urban environment by solving some portion of the Navier-Stokes equations. A more rigorous validation of the QUIC-CFD solver is presented below.

Validation

Although QUIC-CFD emphasizes solution time over absolute accuracy, it matched the ability of more complex models in its ability to recreate wind conditions at *in situ* measurement locations during a large field trial in Oklahoma City [76]. Further validation was also done by Gowardhan et al. [45] which concluded that QUIC-CFD achieved its initial development goals of having 90% of in situ measurement samples be within a factor of two of the simulated result. To complement these validation efforts, we performed a small scale validation experiment around the region of MIT

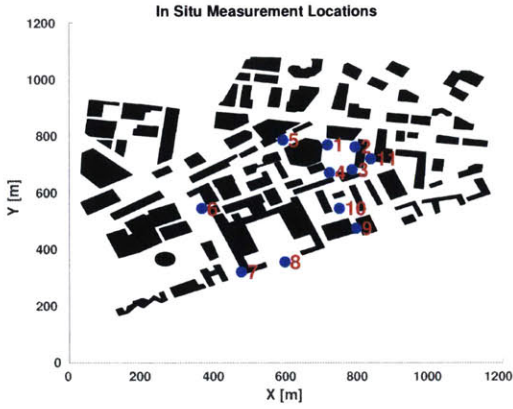


Figure 4-6: QUIC-CFD sample wind fields over a small region of MIT campus. The more traditional approach of QUIC-CFD eliminates the non-physical regions found in the QUIC-URB wind field for this environment.

campus under consideration.

To validate our wind model, a series of *in situ* wind speed and heading measurements were taken with a LCJ Capteurs CV7 ultrasonic anemometer at an altitude of 2 meters and a rate of 4 Hz. An image of the mobile measurement setup can be seen in figure 4-7b. Each measurement period was approximately 20 minutes, and all measurements were done within a 2 hour period. The prevailing wind was measured to have an average speed and heading of 2.9 m/s at 100 m and 102 degrees, respectively, and was compared against a steady wind field generated using QUIC-CFD with a prevailing wind speed of 5 m/s at 100 m and a heading of 90 degrees. The measurement locations can be seen in figure 4-7a.

A table of the average anemometer and model speed, S , and average anemometer and model heading, Θ , can be seen in table 4.1 and figure 4-8. The heading values



(a)



(b)

Figure 4-7: Panel (a) shows the *in situ* measurement locations across MIT campus. The measurements were taken over the course of 2 hours with an ultrasonic anemometer. Panel (b) shows the mobile measurement setup used to record the wind data.

are in the standard meteorological form of degrees North with clockwise as positive. Note also that heading refers to the source of the wind, and not to the direction of flow. To account for uncertainty in the test location, the model average was computed over a $5 \text{ m} \times 5 \text{ m}$ area. So as to assess the model's ability to capture the bulk advection effect over its ability to recreate a particular flow feature at a specific position, the measurement location was adjusted within several meters for a more favorable comparison. Given these allowances, the model performs similarly to the results shown in [45], with approximately 36% of measurement locations showing wind heading errors of less than 15 degrees, and 72% less than 45 degrees. Approximately 36% of measurement locations show wind speed errors of less than 10%, and 81% show wind speed errors of less than 50%. Also, the magnitudes of the speed error's standard deviation suggests that the discretizing our ground velocity, ν_g , by 0.5 m/s was appropriate. Although the magnitudes of the heading error's standard deviation are large, this is to be expected in the turbulent urban canyon and reinforces the approach of planning over the expected wind conditions.

Table 4.1: Measured and Model Wind Speed and Heading

Point	Measured		Model		Mean Error		Std. Dev. Error	
	S	Θ	S	Θ	μ_S	μ_θ	σ_S	σ_θ
1	1.17	207	1.26	220	-0.09	-13	0.62	37
2	1.52	346	1.11	317	0.41	29	0.52	24
3	1.68	96	1.54	101	0.14	-5	0.60	11
4	1.08	65	1.24	37	-0.16	28	0.56	39
5	0.94	293	0.90	289	0.03	4	0.36	18
6	0.87	44	0.81	93	0.06	-50	0.30	26
7	1.01	11	0.77	339	0.24	72	0.27	15
8	1.11	29	0.78	127	0.33	-99	0.43	17
9	0.86	198	0.61	158	0.25	40	0.24	15
10	0.58	232	1.33	211	-0.75	21	0.37	103
11	0.87	196	1.53	203	-0.65	-7	0.31	40

All speed values are in m/s and heading values are in degrees north.

Solver Setup

As with any CFD solver, QUIC-CFD has a range of solver parameters that must be selected and have an important role in determining the quality and validity of the solution. QUIC-CFD, in particular, makes several strong assumptions in order to reduce solution time and it is these that warrant additional discussion. Accordingly, this section will provide an overview of the solver setup and focus on QUIC-CFD’s deviations from the best practices presented in Appendix A.

The portion of the MIT campus used as the simulation domain is $1100 \text{ m} \times 1100 \text{ m} \times 24 \text{ m}$. Although the height of the domain is relatively low compared to the guidelines covered in Appendix A, the low height significantly reduced computation time and did not significantly affect the 2D flow at the 10 m flight simulation height. Given the guidance provided Franke et al. [39] and Oke [78], the environment is spatially discretized into hexahedral elements with a resolution of $1 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$. Structure geometry is also resolved to this grid. Although the selected environment has the usual vegetation, cars, and other urban objects, these are left unmodeled to reduce complexity with the assumption that they do not significantly affect the bulk flow characteristics.




Figure 4-8: *In situ* speed and heading values compared with the simulation results. These results show a similar solver accuracy demonstrated by Gowardhan et al. [45].

The simulation domain includes 100 m on each edge as a buffer region that was excluded from flight simulation. Although this is a relatively small buffer region, the dense urban environment dictates that the dominant flow regime is skimming flow. To capture this, the inflow boundary is set to the appropriate logarithmic mean velocity profile as required by the QUIC-CFD solver. Although a blended exponential and logarithmic profile would be more accurate, the simulation height of 10 m is above the altitude of significant error for the domain's mean structure height.

Chapter 5

Planning Over MIT Campus with Steady Wind Fields

Having established approaches for measuring and estimating wind fields in the previous chapters, the effects of these wind fields on a quadrotor must be investigated. In order to characterize the potential utility of incorporating urban wind fields into planned trajectories, a method of finding minimum-energy trajectories must first be discussed. Each instance of trajectory planning will be performed over a given map of the environment with a known wind field.

5.1 Problem Formulation

Given both the map and wind field, we seek to efficiently compute the minimum-energy, collision-free, 2D trajectory between fixed start and goal locations at a given altitude and subject to constraints imposed by the quadrotor's capabilities and the operator's notion of safety. Because we desire to minimize total energy consumption over the trajectory, the cost function for our optimization, $C(\cdot)$, is defined as the change in platform energy between any two states, χ_i and χ_j . In our case, the vehicle state, χ , is defined as the vehicle's x position, y position, and scalar ground velocity, ν_g . Each element of the state is real valued and is discretized to enable efficient search over the space. In order to reduce computational complexity, the planning is done

over the same regular grid used by the wind field solver such that each xy position also has a corresponding planar wind vector, $\nu_w(\chi_i)$, from the precomputed wind field estimate¹. Also, recall that the vehicle’s air speed, $\nu_a \in \mathbb{R}$, is the norm of the wind speed, $\nu_w \in \mathbb{R}^2$, minus the vehicle’s ground speed, $\nu_g \in \mathbb{R}^2$. For these simulations, we chose to compute the full 3D wind field and then incorporate the resulting wind vector, along with the planar ground velocity, into the planner because the vertical velocity component of urban wind fields has a significant affect on the advection in environment. Despite the fact that the vertical velocity component is frequently near zero for a majority of the locations in a given horizontal plane, it allows for a more realistic flow through the modeled environment. For example, the vertical flow up the face of a building only exists very near to the structure face, but provides an important pathway for air to flow up and over the structure or form a frontal vortex as opposed to being driven around the base of the structure in the horizontal plane. Finally, note that we wish to analyze near-terrain flight in the urban environment and assume the vehicle is flying at a constant altitude close to the ground. A result of this assumption is that the quadrotor is forced to fly around obstacles rather than over them. We chose to plan in 2D primarily due to the additional complexity required by the power consumption model when dealing with large vertical airspeeds that put the vehicle in another flight regime which depends on the recirculation of vortices around the rotors. Modeling the effects of these vortices on the quadrotor’s power consumption was beyond the scope of what could be accomplished in the wind tunnel for this work.

$$C(\chi_i, \chi_{i-1}, \nu_w(\chi_i)) = \Delta E_{tot}(\chi_i, \chi_{i-1}, \nu_w(\chi_i)) \quad (5.1)$$

¹We assume energy is always expended at each time step. Different forms of soaring might allow for energy to be harvested, and would therefore require a different objective, but not fundamentally different strategy to the one we use here.

$$\begin{aligned}
& \underset{\chi_{0:N}}{\operatorname{argmin}} && \sum_{i=0}^N C(\chi_i, \chi_{i-1}, \nu_w(\chi_i)) \\
& \text{subject to:} && \nu_{gmin} < \nu_{g,i} < \nu_{gmax} \\
& && \nu_w(\chi_i) - \nu_{g,i} < \nu_{amax} \\
& && \nu_{g,i} - \nu_{g,i-1} < \nu_{step} \\
& && \chi_N = \chi_G
\end{aligned} \tag{5.2}$$

In order to guarantee progress towards the goal and define a window of safe operating speeds, the ground speed, ν_g , is constrained to be greater than 0.5 m/s and less than 8 m/s. The air speed, ν_a , is also subject to a maximum value of either 10 or 20 m/s, representing the vehicle's thrust limits for each set of simulation results. Additionally, the instantaneous change in velocity between two adjacent grid points, $\nu_{g,t} - \nu_{g,t-1}$, is constrained to 0.5 m/s. With every iteration, the planner can move in the xy grid along 8 possible headings and can increment or decrement its ground velocity, ν_g , by ν_{step} . Finally, the trajectory must terminate in the goal state, χ_G . The QUIC-CFD wind solver will provide the wind field, ν_w , and the following section presents an approach to finding the energy consumption, $\Delta E_{tot}(\cdot)$, for a quadrotor moving in a known wind field.

5.2 Power Consumption Model

In order to find a minimum energy trajectory, it is necessary to compute the quadrotor's energy consumption along a trajectory between two points in the map. Although this clearly requires a power consumption model, the standard approach of developing an accurate system model through system identification and first principles would result in the need to perform a complete simulation of the vehicle dynamics along a trajectory in order to characterize the power consumption. Specifically, a vehicle dynamics model would need to be created after performing extensive system identification and then the planner would need to simulate the vehicle dynamics along each trajectory in order to solve for the vehicle's state and assess the vehicle's resulting

power consumption at each time step. In an effort to avoid this computational complexity, the work presented here develops an empirically derived power consumption model that is only a function of air speed. In doing so, transients in the dynamics are ignored and the vehicle’s heading is assumed to be aligned with the vehicle’s air speed vector. We first present a simple analysis of the underlying physics and then discuss the experimental procedure for characterizing the vehicle’s power consumption as a function of air speed.

In order to assess the change in the vehicle’s energy between states, we begin by expressing the vehicle’s total energy at any point, E_{tot} , as the sum of the potential, kinetic, and stored energy as shown in equation 5.3. The stored energy component, E_s , represents the capacity in J of the on-board battery,

$$E_{tot} = mgh + \frac{1}{2}m|\nu|^2 + E_s. \quad (5.3)$$

Assuming a constant ground velocity, constant altitude, and no acceleration between nodes in the planning graph, the change in total energy, ΔE_{tot} , is equal to the change in stored energy, ΔE_s . The change in stored energy is the difference in the battery’s stored energy between two points in time and will be treated as the total change in the vehicle’s total energy between two states. We can assess the change in stored energy by finding the vehicle’s power consumption between nodes in the planning graph and multiplying it by the time to traverse the edge between those nodes.

For a rotary wing vehicle, and especially a quadrotor, the act of spinning the rotors consumes the vast majority of the stored energy and a power consumption model of the controller, motor, and rotor system is required. Furthermore, simply keeping the vehicle aloft takes a large portion, usually greater than 50% of the vehicle’s total thrust. During hover, when the thrust from the four rotors is equal to the weight of the vehicle, the induced velocity through the rotors can be found using momentum theory. The induced velocity, generally denoted by ν_i , is the rate at which the motion of the rotor moves air perpendicular to the rotor plane. The induced velocity at hover,

denoted ν_h in this case, is defined in equation 5.4 with the hover thrust, T_h , equal to 1/4 of the vehicle weight, S as the swept rotor area, and ρ as the air density. This is a direct result of the thrust being equivalent to the product of the mass flow rate and upstream homogeneous flow velocity, or air speed in this context.

$$\nu_h = \sqrt{\frac{T_h}{2\rho S}} \quad (5.4)$$

Having defined the hover thrust, we can continue to define the induced velocity through the rotors, ν_i , which is a significant factor in determining the power consumption. Although the induced velocity is not directly measured, it can be expressed as the quadratic function, shown in equation 5.5, that depends on the induced velocity at hover, airspeed, and angle of attack, α . Similar to the derivation of the hover thrust, the quadratic expression for the induced velocity is a result of setting the rotor thrust equal to the product of the mass flow rate and upstream flow velocity. However, the velocity through the rotor is now a sum of the induced velocity and airspeed, which results in a quadratic expression.

$$\nu_i = \frac{\nu_h^2}{\sqrt{(\nu_a \cos \alpha)^2 + (\nu_i - \nu_a \sin \alpha)^2}} \quad (5.5)$$

With this expression for the induced velocity, we can solve for the power consumption. As shown in equation 5.6, Huang et al. [51] express the power consumption, P , of a single rotor as a function of induced velocity (ν_i), air speed (ν_a), rotor thrust (T), angle of attack (α), propeller efficiency (η_p), motor efficiency (η_m), and controller efficiency (η_c). The efficiencies, found by Latorre [60] for a very similar configuration, are taken to be 0.6, 0.85, and 0.95, respectively.

$$P = \frac{T(\nu_i - \nu_a \sin \alpha)}{\eta_p \eta_m \eta_c} \quad (5.6)$$

One common approach to assessing a vehicle's flight characteristics is to rigidly mount the vehicle on a sting to directly measure the forces generated during simulated flight in the wind tunnel. Unfortunately, a quadrotor platform generates a significant




Figure 5-1: The quadrotor was flown in the wind tunnel chamber with a Vicon motion capture system in order to measure power consumption as a function of air speed. It executed a position hold command for up to 5 minutes over a range of wind speeds and logged battery voltage and current data, as well as the wind speed within the tunnel.

amount of vibration that can excite vibrational modes at particular motor speeds. The resulting vibration adds both measurement noise and bias. A further issue is that a quadrotor, which must pitch in forward flight, cannot assume its natural pitch angle when attached to a sting. Lastly, the quadrotor does not need to expend the effort to hold position or reject disturbances which lowers its overall power consumption. Given these drawbacks, we chose to setup a motion capture system within the MIT Wright Brothers Wind Tunnel and the fly the quadrotor in place, as shown in figure 5-1. The vehicle maintained position while the air speed was set and held at 5, 10, 15, 20, and 25 mph for durations between 30 seconds and several minutes. A series of 11, constant air speed tests were performed over the course of several days. The vehicle's battery voltage and current were measured at 1 kHz, averaged, and recorded at 100 Hz. The wind speed was closely controlled and recorded at 4 Hz by a LCJ Capteurs CV7 ultrasonic anemometer.

The resulting calculated power consumption and measured power consumption data, normalized by the hover power consumption, P_h , of 247 W, can be seen in figure 5-2. To validate the measured power consumption data, a curve was fit to the average pitch values for all trials to serve as an input to equations 5.6 and 5.5. The





Figure 5-2: Panel (a) shows the mean and standard error of the vehicle pitch along with a quadratic fit to the data. These values were used to generate the power consumption model, shown in black, in panel (b). Power consumption data was normalized by the hover power consumption in still air. The blue points denote the results as reported by the current and voltage sensor with the associated standard error. The red points denote the power consumption during the outdoor hover tests and their standard error. Note that the standard error for the wind tunnel trials is negligible due to sampling frequency and test duration.

measured pitch data and quadratic curve fit can be seen in panel figure 5-2(a). The output of this power consumption model, given the fit of pitch as a function of air speed, can be seen, along with the measured power consumption, in figure 5-2(b). Note the initial decrease in power consumption caused by translational lift and the steep increase at higher air speeds caused primarily by the vehicles angle of attack and form drag. Translational lift is present in all rotorcraft and is a natural result of the increased flow over the rotors in forward flight which improves rotor efficiency [64]. The discrepancy between the measured power consumption and the calculated power consumption is probably due to a combination of imperfect modeling of the controller, motor, and rotor efficiencies and the vehicles difficulty holding position at the higher air speeds due to controller constraints.

A series of outdoor hover tests were performed to compare the quadrotor’s energy consumption in the field to the controlled tests from the wind tunnel. In order to comply with FAA regulation, the tests were done on the test frame seen in figure 5-3a. The wind speed was recorded with a Young 81000 ultrasonic anemometer and the



(a)



(b)

Figure 5-3: Panel (a) shows vehicle performing an autonomous hover test on MIT campus. It is attached to a frame to comply with FAA regulations. The blue points represent the vehicle’s power consumption from the test flights in the MIT wind tunnel. The red points denote the power consumption during the outdoor hover tests and their standard error.

vehicle’s battery voltage and current consumption were logged. Because multipath error and obstruction of the horizon prevented accurate localization with GPS, a downward facing optical flow camera, the 3DR PX4Flow, was used instead. The magnetometer was used for the vehicle’s yaw orientation.

The logged voltage and current data was used to compute the vehicle’s power consumption offline. Because both the power data and anemometer data were logged on the same system, each power consumption sample was associated with a corresponding wind speed. Given that the vehicle hovered for several minutes in each test, the vehicle encountered a range of wind speeds up to a max of approximately 4 m/s. Aggregating this data created the power consumption curve shown in figure 5-3b. Although the curve from the hover tests does capture the downward trend generated by translational lift, it is likely that large deviations in the control effort due to turbulence caused a back EMF current to the battery that degraded the measurement accuracy. It is probable that the resulting erroneous measurements artificially lowered the power consumption at higher wind speeds due to the naturally increased turbulence in the wind field from the surrounding obstacles and structures.

5.3 Planning

Given a planar section of the 3D wind field estimate and occupancy map, A* search was used to find the minimum-cost, collision-free trajectory between a start and goal location over xy positions and scalar edge velocity. A velocity discretization of 0.5 m/s allowed for sufficiently fine adjustment of the ground velocity. The vehicle’s altitude was fixed to 10 meters and the planning was performed directly on the xy grid used by the wind model. The vehicle’s air speed was calculated as the ground speed minus the wind speed. To find the stored energy consumption between two states, power consumption as a function of air speed is multiplied by the edge distance and divided by the current ground speed. Recall that our model of urban operations only considers motion at a fixed altitude in the xy -plane, the energy consumed between two states, χ_i and χ_j , can then be expressed as follows with $dist(\chi_i, \chi_j)$ as the Euclidean distance function between two nodes and ν_g as the vehicle ground speed along the connecting edge.

$$\Delta E_{tot}(\chi_i, \chi_{i-1}, \nu_g(\chi_i)) = \frac{P(\nu_{g,i} - \nu_{w,i})dist(\chi_i, \chi_{i-1})}{\nu_{g,i}} \quad (5.7)$$

Two sample sets of 1000 simulated, minimum-energy trajectories using 500 random start and goal locations were planned with both the naive and wind-aware planners. In order to demonstrate the effect of constraining the vehicle’s maximum air speed, the two sets of simulations had max air speeds of 10 and 20 m/s, respectively. The smaller value of 10 m/s was selected to represent a realistic upper bound on the quadrotor’s airspeed, while the large value of 20 m/s is used to investigate the effect that a relatively unconstrained vehicle has on the simulation results. The minimum distance between the start and goal locations was 10 m and the 10 wind conditions were composed of 5 speeds and 2 headings. In order to represent the vehicle’s flight envelope, failure is defined as the case when the planner is unable to lower its ground speed sufficiently to prevent exceeding the maximum air speed, while still making forward progress.

To allow for sufficiently long trajectories, a region of the MIT campus and surrounding city spanning an area of 1100 m \times 1100 m was selected. A three-dimensional

model of this region containing 81 buildings, each constrained to be no taller than 20 m, was created to serve as an input to the wind model and can be seen in figure 5-4. The domain height, the maximum height of the model's air volume, was constrained to 24 m to minimize computation time and a xyz-grid resolution of $1\text{ m} \times 1\text{ m} \times 1\text{ m}$ was used to capture the detailed structure of urban canopy layer (UCL) wind fields. Despite only planning in the plane, the wind model must solve over the full 3D map to account for the complex flow through the urban environment. Although constraining the domain and building height makes for an imperfect model, the focus of this work was on flight within the urban canyon, and *in situ* measurements demonstrate that the resulting wind fields remain sufficiently accurate.




Figure 5-4: The 3D model of the region of MIT campus being planned over that was used as an input to the QUIC-CFD wind model along with the prevailing wind conditions.

As a coastal city, Cambridge typically has either an easterly or westerly wind heading and might only switch once a day due to the diurnal cycle. With this in mind, wind fields with both an Easterly and Westerly wind heading were used here with wind speeds of 5, 7.5, 10, 12.5, and 15 m/s for a total of 10 wind fields. Given the past year of data from the MIT weather station, these wind speeds occur in approximately 27%, 11%, 3.5%, 1.3%, and 0.35% of the measurements, respectively.




Figure 5-5: A visualization of the x-component of the velocity in a planar slice of the three-dimensional wind field for a prevailing wind speed of 10 m/s measured at an altitude of 100 m and a prevailing wind heading of 90 degrees.

A visualization of one of these wind fields can be seen in figure 5-5. The resulting trajectories are used to draw more general conclusions about the benefits of wind-aware planning in the urban environment.

5.4 Results

Because these planners are both seeking to minimize energy consumption, we use the mean percentage difference in energy consumption to compare them. This comparison is made in figure 5-6 for both maximum air speeds over each of the 10 wind fields across all 500 trials. For the simulations with the 20 m/s maximum air speed, the wind planner, as might be expected, outperforms the naive planner over all wind conditions and is able to improve its performance with increasing wind speed. Because a prevailing wind speed of 5 m/s at an altitude of 100 m generates relatively small wind speeds at an altitude of 10 m, the wind planner is still only able to reduce energy




Figure 5-6: The wind planner significantly outperforms the naive planner over all prevailing wind speeds tested. The error bars shown here represent standard error.

consumption by approximately 0.5%. Although there is only a small change at 7.5 m/s, the wind planner’s performance increases significantly with increasing prevailing wind speed and achieves approximately a 4%, and 10% improvement at 10, and 15 m/s, respectively. This is likely because it is able to exploit tail winds, while the naive planner unknowingly flies into headwinds.

The second set of simulations used a maximum air speed of 10 m/s and resulted in a change in the wind planner performance. Again, the wind planner outperforms the naive planner in all cases, but is only able to reduce energy consumption by approximately 0.9% for a prevailing wind speed of 5 m/s. Demonstrating the additional expense incurred by the naive planner’s upwind trajectories, relatively large gains of 7.3% and 9.3% are shown at 10 m/s for the Easterly and Westerly wind headings, respectively. As might be expected, the wind planner shows equivalent or increased gains for 12.5 and 15 m/s in the Easterly wind fields. Surprisingly, the energy consumption decreases for the Westerly heading. This may be due to the Westerly wind being channeled down the aligned urban canyons and causing otherwise costly upwind trajectories to fail. Some evidence of this failure mode can be found in figure 5-8 by noting the cluster of naive planner failures for the Westerly wind heading at the lower left of the map leading into a large diagonal canyon. Also note that the naive planner

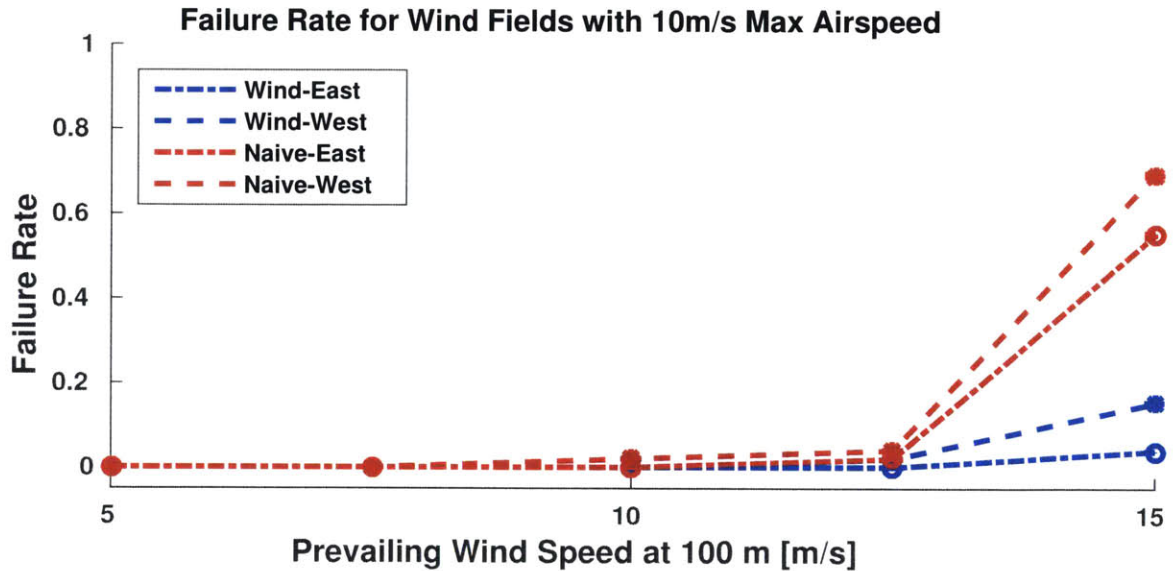


Figure 5-7: Neither planner has significant error rates below 12.5 m/s, but the naive planner strongly underperforms the wind planner for a 15 m/s prevailing wind speed. The error bars shown here represent standard error.

fails to reach the goal location in approximately 55% and 69% of the 15m/s trials, as shown in figure 5-7. These incomplete trials cannot be counted towards the mean energy consumption, but would likely be some of the highest values in the set.

Looking closer at a specific trial, figure 5-9 (a) and 5-9 (b) show the total energy consumption and speeds over the length of the trajectory for a wind speed of 10 m/s and Westerly heading. By maximizing its ground speed without exceeding the maximum air speed, the wind planner demonstrates a 39.4% reduction in total energy consumption while traveling 10.8% further with a 22% shorter flight time. The resulting trajectories in figure 5-10 show the wind planner sheltering next to and downwind from structures, while the naive planner exposes the vehicle to a strong headwind.

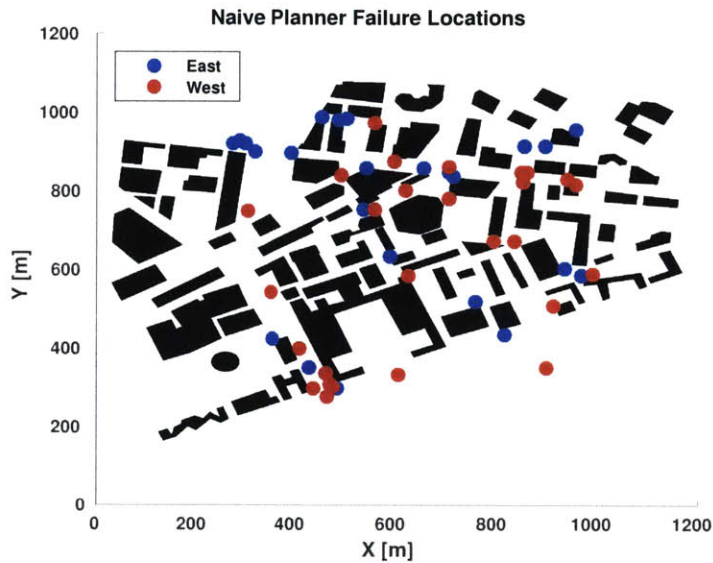


Figure 5-8: The naive planner failure locations over the 2D map. Note that a large fraction of them are at corners of buildings or at the entry points to the structure group near canyons channeling the wind along their axis.

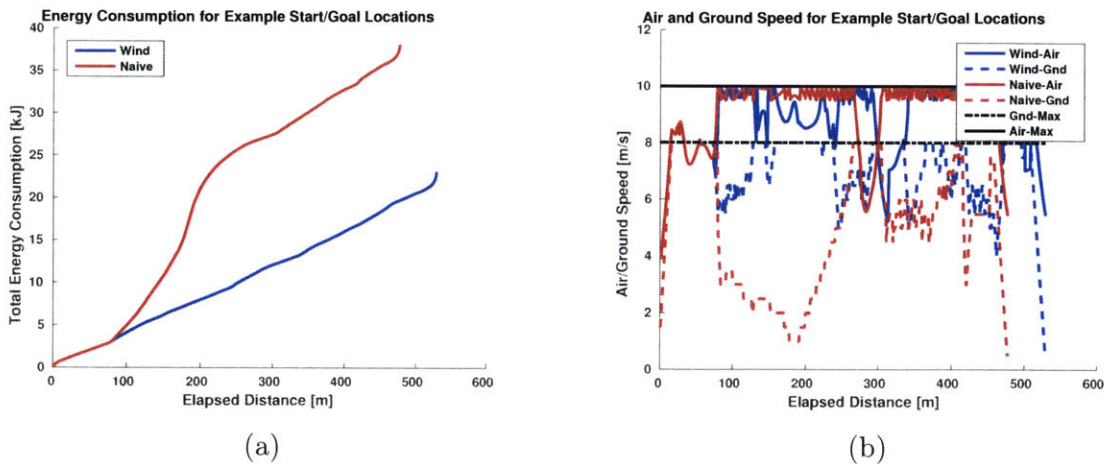


Figure 5-9: Panel (a) demonstrates the wind planner's ability to reduce energy consumption over the course of a flight. Panel (b) shows how the naive planner is unable to increase its ground speed due to its own dynamic constraints and the upwind, exposed trajectory it selected.

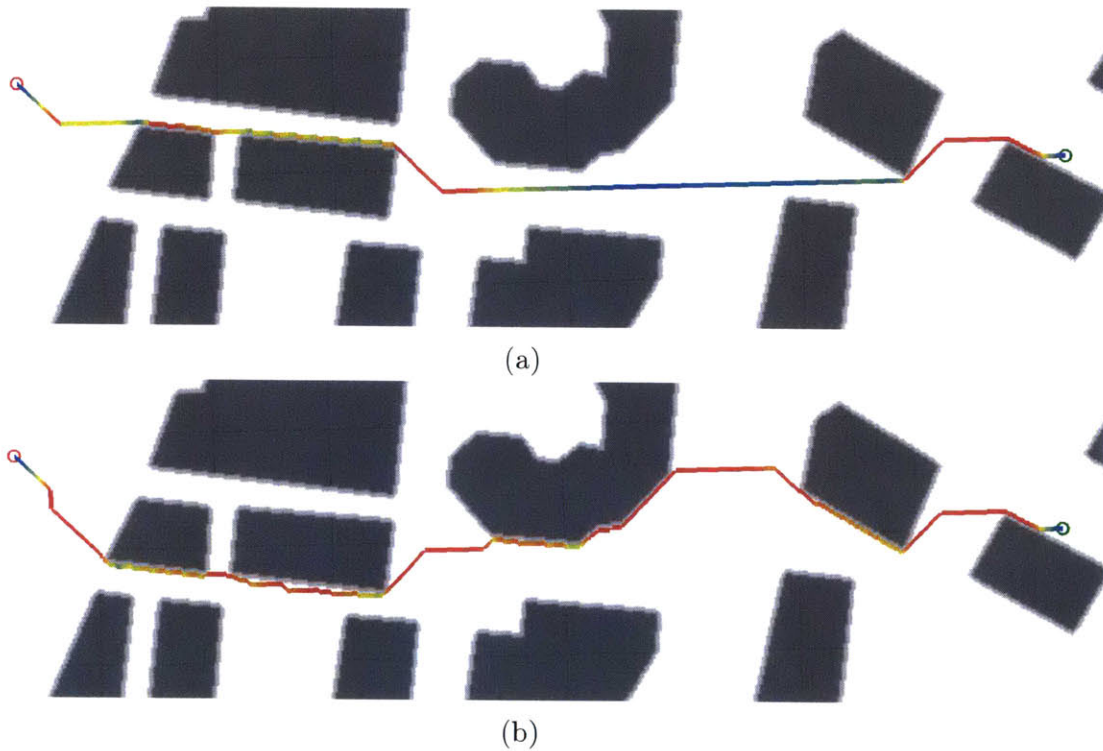


Figure 5-10: Panel (a) and (b) show the naive and wind planner's trajectories, respectively, for an example start and goal location. The trajectories are color coded by ground speed. The green and red circles denote the start and goal locations, respectively.

Chapter 6

Conclusions

The work presented here focused on enabling and assessing the benefits of quadrotor flight in the complex wind fields found in urban environments. In pursuit of these goals, an approach to on-board airspeed measurement was developed and an analysis was performed of the potential benefits and challenges of minimum-energy trajectory planning through known wind fields. Despite the MEMS anemometer's relatively poor performance at lower airspeeds, the resulting wind speed and heading measurements suggest that the propellers only have a significant effect on the estimated speeds below 10 m/s, while heading estimates were practically unaffected in the whole airspeed range. Although it is beyond the scope of the work presented here, it is possible that further development of the sensor housing, additional modelling of the flow around the vehicle, and alternative sensor placements could improve performance at lower airspeeds.

In order to gauge the quadrotor's ability to exploit urban wind fields for improved flight performance, an existing CFD wind solver, QUIC-CFD, was leveraged and verified against *in situ* measurements around campus. QUIC-CFD was used to generate a set of representative wind fields over a portion of MIT campus that served as inputs to the minimum-energy trajectory planner. Furthermore, an empirically derived power consumption model was found through a series of flight tests in the MIT wind tunnel and was also used to inform the minimum-energy trajectory planner. The results of many simulated minimum-energy trajectories through the campus demonstrated

the potential gains of planning over the wind field by comparing a minimum-energy planner naive of the wind field to one aware of it. In this context, the lower prevailing wind speeds offered little opportunity for improvement over the naive planner, but those above 10 m/s allowed the wind planner to significantly outperform the naive planner. Additionally, the wind planner was capable of finding minimum-energy trajectories such that the vehicle was less likely to exceed its flight performance envelope given strong headwinds and more likely to exploit tailwinds for a reduced time to goal and total energy consumption. More generally, our initial investigation suggests that UAVs could benefit from considering wind conditions in complex environments.

Appendix A

Wind Field Simulation Best Practices

Ensuring that a CFD simulation provides an accurate wind field estimate within a complex urban environment is challenging and there is a large body of literature that addresses this issue. Driven by recent work in climatology, particulate dispersion, pedestrian comfort, city planning, and structural engineering in urban environments, several summaries of the best practices for urban wind field simulation from both academia and industry have been created over the past decade. This section attempts to distill these recommendations into a concise set of guidelines relevant for this application.

A.1 Sources of Error

In COST Action 732, Franke et al. [39] present a list of all of the sources of error found in numerical simulations. Although there are many ways to classify these, they selected the approach of Coleman and Stern [29] that uses the two broad categories of errors in the modelling physics and numerical solution. The former derives from erroneous assumptions and approximations concerning the mathematical formulation of the underlying physical processes. The latter is simply a result of incorrect or incomplete solutions to the assumed mathematical model. This list of the individual sources of error, adapted for the purposes of this work, is shown below.

- **Errors and Uncertainties in Modelling Physics:**

- Simplification of physical complexity
- Geometric representation
- Domain size
- Boundary conditions
- Initial conditions

- **Numerical Errors and Uncertainties:**

- Numerical approximation
- Spatial discretization
- Temporal discretization
- Iterative convergence

For the sake of brevity and relevance, several items have been omitted from the version of the list presented here. Programming errors in the CFD code are raised as a potential concern, but we assume here that the CFD models in this work are free of significant coding errors due to their regular use over many years or even decades. Similarly, we ignore round-off issues by always selecting double precision when possible. Lastly, we forgo concerns raised about the use of previous data in parameter tuning because all simulations were done using parameters and supporting data specific to that simulation.

A.2 Simplification of Physical Complexity

Although the Navier-Stokes equations can be used to directly solve for flow in the atmospheric boundary layer, solving them for the turbulent flow found in urban wind fields is computationally infeasible. This rules out approaches such as Direct Numerical Simulation (DNS) and promotes either RANS or LES which perform averaging over time and space, respectively, to lighten the computational burden of solving for flow features on the smallest scales. Although necessary, these approximations

introduce error and care must be taken to ensure that the flow is being resolved appropriately for the application. This is discussed in greater detail in the remainder of this section.

Before selecting a turbulence model, it is useful to make a few basic assumptions about the environment and the flow within it to simplify the model and reduce computational complexity. In all of the work done here, the atmosphere is assumed to have neutral stratification. This means that the flow is not driven by differences in temperature, humidity, or the presence of rain or snow in the atmosphere. Although this is generally the case with relatively strong prevailing winds and some cloud cover [78], this is probably a strong assumption for most urban environments due to their tendency to be a heat island. Regardless, appropriately addressing thermal effects within urban wind fields requires a level of accuracy in terrain modeling and complexity of simulation that are beyond the scope of this work. Similarly, although the Coriolis force has a measurable effect on the flow within an urban environment, Franke et al. claim that it can be ignored for microscale models by properly selecting the incoming flow profile. Lastly, if only the lowest 200 m of the atmosphere are investigated, the assumption of non-divergent flow fields and constant density may be used without losing accuracy in the model results and is commonly applied in the case of urban wind field simulation [21].

Even in the simplest cases, these basic assumptions do not necessarily give rise to a feasible problem without further reduction of the complexity. The inherently turbulent nature of the urban environment prevents using the transport equations to directly solve for the flow at all scales. Accordingly, an appropriate turbulence model must be selected to make the problem feasible. Aside from the basic assumptions discussed above, the turbulence model largely defines what kind of approximations to the model physics will be made, and the errors that will be introduced as a result. This work considers the use of three common turbulence models. These are the Reynolds-averaged Navier-Stokes, Large Eddy Simulation, and Reynolds Stress Model. Defining a continuum of increasing solution accuracy and complexity, RANS comes first followed by RSM and then LES with the highest solution complexity and

potential accuracy. A brief introduction to each of these turbulence models will be discussed below along with their strengths and weaknesses as reported in the literature.

A.2.1 Reynolds-Averaged Navier-Stokes (RANS)

The RANS solution is the most common approach in wind field estimation, and represents the infinite time average of the underlying unsteady flow. The standard $k - \epsilon$ model introduces a transport equation for the turbulent kinetic energy, k , and the specific turbulent kinetic energy dissipation rate, ϵ . These are used to find the turbulent eddy viscosity via the Boussinesq approximation. Although the overall accuracy of this approach in urban wind field estimation is still an open question, wind fields generated this way match well with wind tunnel results that have stable inflow conditions and do so with relatively short solution times [39]. In general, no definite statement can be made about the relative performance of the various turbulence models used within a RANS solver, but some insights are provided in [38]. Specifically, the two equation $k - \epsilon$ model is the industry standard, but according to Castro [21] it only produces good results in wind engineering “fortunately”. Its primary issue being overproduction of turbulent kinetic energy in regions of stagnant flow (stagnation point anomaly). As such, Franke et al. [38] recommends more advanced versions of the $k - \epsilon$ model such as RNG or realizable $k - \epsilon$. Another significant issue is that this model does not provide any estimate of the uncertainty in the wind velocity over the domain. Not only does $k - \epsilon$ RANS not provide an uncertainty estimate, but it also averages out the unsteady fluctuations in the wind velocity found in urban wind fields.

A.2.2 Reynolds Stress Model (RSM)

The Reynolds stress model offers improved fidelity over the simpler one or two equation RANS models, but remains less computationally intensive than LES. Unlike the $k - \epsilon$ model, it does not use the isotropic eddy viscosity approach to approximate

the Reynolds stresses and instead directly solves for them using additional transport equations. Although no work in urban wind field estimation for UAV flight has been done using RSM, it has been frequently used to estimate pedestrian wind comfort in urban environments. Murakami et al. [74] did a study comparing $k - \epsilon$ RANS, RSM, and LES in urban wind field estimation using ANSYS Fluent and, as might be expected, ranked them in the same order of increasing accuracy. Similarly, Reiter et al. [88] compared $k - \epsilon$ RANS, realizable $k - \epsilon$ RANS, and RSM. In this case, RSM outperformed the other turbulence models by not only identifying critical areas but also accurately estimating the wind velocity magnitudes at those locations. An added benefit of using RSM is that the six Reynolds stresses are the unique components of the wind velocity covariance matrix multiplied by the air density. This provides convenient access to an uncertainty estimate of the wind velocity without explicitly computing the unsteady solution.

A.2.3 Large Eddy Simulation (LES)

Unlike $k - \epsilon$ or RSM RANS, LES is an unsteady simulation that performs only spatial averaging as opposed to time averaging. Because small scale eddies are computationally intensive to resolve and tend to be isotropic, LES directly solves for the larger eddies in the flow that are a direct result of the environmental geometry and uses sub-grid scale models to approximate the remaining flow features. Although LES has been actively studied in urban wind field simulation, the general consensus is that it provides improved accuracy over steady solutions but at a large cost in solution time. Franke et al. [38] suggests that LES is capable of more accurate solutions than RANS, but that its computational complexity makes it infeasible for wind engineering problems in the foreseeable future. Seven years later, Franke et al. [39] again supports the same conclusion saying that LES outperforms both RANS and URANS, but only stating that it requires a great deal of computation. It is also noted that the boundary conditions need to be resolved in both time and space and that this is often impossible to obtain. This last conclusion is supported by the work of White et al. [103] in the finding that the LES simulation required the turbulent kinetic energy

to be modeled at the inflow boundary in order to match the *in situ* measurements. Furthermore, Tominaga et al. [101] notes that LES simulations are capable of higher accuracy, but are not practical. Finally, Sutherland [99] attempted to extend the work of Galway [40] by comparing $k - \epsilon$ RANS and LES and, in short, found that although RANS was significantly faster, its time averaging of the wind field reduced the vehicle’s control effort and position deviations by as much as an order of magnitude.

A.3 Geometric Representation

The choice of what environmental geometry to include and its resolution is another important approximation that, if done correctly, acts to reduce the domain size and complexity of the computational mesh without affecting the quality of the solution. Franke et al. [39] notes that a structure will have minimal influence on the target area if its distance upwind of the target area is between 6 and 10 times the structure height. Tominaga et al. [101] states that buildings within a radius of 1 to 2 times the average feature height of the target area should be clearly modeled with an additional street block in each direction. Furthermore, representative roughness elements should be placed from the outer edge of the additional street blocks to the outer edge of the computational domain and the appropriate environmental roughness length should also be specified for the inflow velocity profile.

To properly assess the level of detail required in the environmental geometry, a sensitivity study should be done over a range resolutions. Because this work is focused on navigation in a horizontal plane over a representative urban environment with inflated obstacles to prevent flight near structures, the structures are modeled mostly as extruded 2D sketches with approximately a 1 m resolution. Although the absence of vegetation, cars, and other small scale urban features certainly impacts the realism of the simulation, it should not detract from the primary goal of investigating the effect of representative urban wind fields on quadrotor flight performance over minimum-energy trajectories.

A.4 Domain Size

The choice of domain size is strongly linked with the boundary conditions. For example, a uniform inflow velocity profile could be used if the domain and modeled environment was sufficiently large so as to allow the velocity profile to evolve to match the environmental roughness within the domain. Because of the inherent complexity of wind field simulation over large urban areas, the domain size is limited as much as possible while using the boundary conditions to enforce a representative velocity profile. A primary concern the extents of the domain is the blockage ratio. The blockage ratio is the ratio of the frontal area of the structures over the total area of a vertical plan perpendicular to the wind heading and is used to characterize the effects of the boundaries on the flow through the domain. A large blockage ratio leads to a Venturi effect that increases the flow velocity over the structures. Both Franke et al. [39] and Tominaga et al. [101] give a minimum blockage ratio of 10%, but add that it would ideally be less than 3%. Beyond this basic rule, further guidelines are given for the extent of the domains.

Conservatively, the height of the domain could be the height of the boundary layer determined by terrain roughness category [101]. In the likely case that this creates a prohibitively large domain, the top of the domain should be at least 6 times the maximum structure height [39]. For single buildings and a single flow direction, the lateral extent of the domain should be a between 2.3 and 5 times the maximum structure height and distances up to 10 times the maximum structure height can be used along the longitudinal flow direction for domains with multiple buildings. Although practical, these distances are often insufficient to properly resolve all quantities and others recommend using representative roughness elements out to 10 times the maximum structure height around the target area [39]. This is reinforced by Tominaga et al. [101] which notes that lateral boundaries of 2.3 times the width of the built area in the lateral direction and 15 times the maximum structure height in the longitudinal flow direction.

A.5 Boundary Conditions

The boundary conditions of a simulation define the influence of the unmodeled terrain that surrounds the domain. Although they can be used to encode a great deal of information in order to reduce the size of the explicitly modeled environment, care must be taken to follow the guidelines on domain size to prevent the boundary definitions from affecting the solution in the region of interest. This section presents a summary of best practice guidelines for each type of boundary in the domain.

A.5.1 Inflow

Having obtained a prevailing wind estimate from a suitable, nearby weather station with a known reference height and local roughness height, the inflow velocity profile can then be defined using the characteristics of the target area. The process for obtaining this profile is roughly outlined in Oke [78], with reference to figure 1-1, and covered in greater detail in Macdonald [65]. To start, it must be assumed that for the case of neutral stability, the mean velocity profile above the mean structure height, zh , takes the form given in equation A.1. In this equation, the mean horizontal velocity, u , is a semi-logarithmic function of the altitude, z , the roughness height, z_0 , the von Karman constant, κ , and the displacement height, z_d . The sum of the roughness and displacement height is equivalent to the theoretical zero velocity of the purely logarithmic velocity profile. Oke [78] provides tables to determine the roughness height and displacement distance for several types of urban densities, but further discussion can be found in Grimmond and Oke [46].

$$u(z) = \frac{u_*}{\kappa} \ln\left(\frac{z - z_d}{z_0}\right) \quad (\text{A.1})$$

At lower altitudes below the mean structure height and within the urban canopy layer, the velocity profile takes on the exponential form shown below that was originally developed by Cionco [28] for plant canopies. In this case, the mean horizontal velocity is a function of the mean rooftop velocity, u_H , measured at the mean structure height, z_H , and the Cionco's attenuation coefficient a .

$$u(z) = u_H \exp(a(z/z_H - 1)) \quad (\text{A.2})$$

There are several approaches to finding the attenuation coefficient. Cionco originally proposed the form found in equation A.2. This version has the cube of the attenuation coefficient as a function of the mean structure height, z_H , mixing length scale, l_c , and friction coefficient. Unfortunately, the friction coefficient is itself a function of the sectional drag coefficient, C'_D and the frontal area ratio, λ_f , and the sectional drag coefficient is not easily estimated. Luckily, data from Macdonald et al. [66] shows that the attenuation coefficient follows the linear relationship in equation A.3 with respect to the frontal area ratio up to an approximate value of 0.35. The frontal area ratio, λ_f , is defined as the frontal area of each obstacle exposed to the wind, A_f , divided by the plan area of the obstacles, A_d . Given their experimental data, Macdonald et al. found the constant k to be 9.6. Although the frontal area ratio could be calculated from a 3D structure model, Britter and Hanna [17] give approximate values that can be used in place of exact calculation. They state that typical values of the frontal area ratio range between roughly 0.1 for moderate building density and 0.3 for downtown areas. Further discussion of how to determine the frontal area ratio can be found in Grimmond and Oke [46].

$$u(z) = k * \lambda_f \quad (\text{A.3})$$

Although it is clear that higher altitudes have a logarithmic profile, and lower altitudes have an exponential profile, these two must smoothly blend together between the mean structure height, z_H , and the wake diffusion height, z_w , to form a single, continuous profile. As its name suggests, the wake diffusion height, or blending height, is the altitude at which the structure turbulence resolves into the semi-logarithmic profile that defines the upper bound of the roughness sublayer and lower bound of the inertial sublayer. Addressing this missing component, Macdonald [65] derives the expression in equation A.4 for the blended region and those in equations A.5 and A.6 for its associated constants.




Figure A-1

$$u(z) = \frac{u_*}{B} \ln\left(\frac{A + Bz}{A + Bz_H}\right) + u_H \quad (\text{A.4})$$

$$A = l_c - \left(\frac{z_H}{z_w - z_H}\right)(\kappa(z_w - d) - l_c) \quad (\text{A.5})$$

$$B = \left(\frac{1}{z_w - z_H}\right)(\kappa(z_w - d) - l_c) \quad (\text{A.6})$$

These expressions could be used directly to solve for a component of the velocity profile by approximating the wake diffusion height, z_w . A common approximation proposed by Raupach et al. uses the mean structure height and width and gives $z_w = z_H + 1.5W$, but Oke [78] claims it can range between 1.5 and 4 times the mean structure height. Alternatively, equation A.4 can be set equal to equation A.1 with $z = z_w$ and for the blending height, z_w . This can then be used to obtain a single continuous mean velocity profile. An example of this profile over a range of prevailing wind velocities can be seen in figure A-1.

A.5.2 Wall Boundary Conditions

Both structure surfaces and the ground plane can be treated as no-slip conditions. Because obstacles will be inflated and flow conditions near walls are not important, walls are treated as smooth, low-Reynolds number approaches are not used, and only

standard walls functions are used.

A.5.3 Top Boundary Conditions

Franke et al. [39] suggests that the top boundary condition could be treated as a constant shear stress, fixed velocity corresponding to the inflow profile at that altitude, symmetry condition, or as a zero shear condition. The constant shear stress is used to prevent deviation from the inflow profile and can be derived from the same [48]. It is suggested that using a symmetry or zero shear condition should be treated as an approximation that is only valid if the top boundary is outside the boundary layer.

A.5.4 Lateral Boundary Conditions

Although multi-directional flows often force boundary conditions to be either an inflow or an outflow, symmetry conditions can be used when there is only a single inflow and outflow boundary. Because symmetry conditions enforce parallel velocity, the domain should be sufficiently large in the lateral direction so as not to influence the flow in the target area.

A.5.5 Outflow Boundary Conditions

Open boundary conditions are usually used in commercial CFD for microscale obstacle-accommodating meteorology models. These are either outflow or constant static pressure [39]. Because of the mechanism of this condition, it should be placed far enough away from the domain to avoid any reverse flow.

A.6 Initial Conditions

Initial conditions in the domain are relatively unimportant for a RANS solver, but can speed up the solution time. As such, the RANS solutions in this work are run without precomputed initial conditions. Unsteady solutions, however, are often simulated for

a period of time proportional to the domain size divided by the inflow velocity at the altitude of primary interest.

A.7 Numerical Approximation

In order to make the basic differential transport equations solvable, they must be discretized and transformed into algebraic equations. There are several approaches to accomplishing this and too much dependence on the application to offer specific guidance that is broadly applicable, but both Franke et al. [39] and Tominaga et al. [101] discourage the use of first-order upwind methods, as they are not appropriate for all transported quantities.

A.8 Spatial Discretization

A significant source of error can be the resolution and distribution of grid points within the domain. The goal of any finite volume discretization is to discretize the domain while introducing as little error as possible. Although Franke et al. [39] and others recommend a primarily hexhedral domain, this work uses a strictly tetrahedral domain for the purposes of making the post processing easier. Outside of the element type, the primary concerns are the growth rate of the cells and the grid size near features. Although the naive solution would be to simply use a very fine mesh, increased solution times encourage the use of grid refinement studies to allow for as coarse a mesh as possible, Furthermore, Bélair et al. [13] notes that RANS solvers do not necessarily converge to the exact solution with increased resolution and it has also been shown that the accuracy of mesoscale models does not necessarily increase with increasing resolution. Franke et al. [39] recommends 10 cells per building separation and Tominaga et al. [101] recommend 10 cells along a buildings wall to properly reproduce flow separation. Franke et al. [39] recommends an expansion ratio of 1.2 in regions of high gradient and Scaperdas and Gilham [96] and Bartzis et al. [12] recommend a max expansion ratio of 1.2.

A.9 Temporal Discretization

Unsteady simulations introduce the need to select a time step with which the flow will evolve. This should be selected through iterative trials over a range of time steps, but there are several guidelines that can be followed in the absence of or to supplement a sensitivity study. In the ideal case that the relevant frequency range can be estimated, then the highest frequency should be resolved with at least 10 to 20 time steps per period [70]. Another approach, shown in equation A.7, relies on the Courant-Friedrichs-Lewy (CFL), or Courant number and relies on the minimum grid dimension, ΔX_{min} , and the maximum velocity, U_{max} .

$$\Delta t = CFL \frac{\Delta X_{min}}{U_{max}} \quad (\text{A.7})$$

A.10 Iterative Convergence

Because CFD simulation is an inherently iterative process, termination conditions are required and are usually based on the change in the transport variables over a single time step. These changes in the variables are called residuals and industrial applications usually impose a termination criteria of 0.001. However, this tends to be too high to allow for a converged solution and a reduction of four orders of magnitude is a better alternative [39].

Bibliography

- [1] Amazon prime air. <https://www.amazon.com/b?node=8037720011>. Accessed: 2016-07-08.
- [2] Global health observatory data. http://www.who.int/gho/urban_health/situation_trends/urban_population_growth_text/en/. Accessed: 2016-03-28.
- [3] Google wing. <http://spectrum.ieee.org/autoton/robotics/drones/google-project-wing-moonshot>. Accessed: 2016-07-08.
- [4] Dr Steve Ackerman. Sea and land breezes. *University of Wisconsin*. Retrieved, pages 10–24, 2006.
- [5] Hubert Airy. The soaring of birds. *Nature*, 28:103, 1883.
- [6] David E Alexander and Steven Vogel. *Nature's flyers: birds, insects, and the biomechanics of flight*. JHU Press, 2004.
- [7] Michael J Allen and Victor Lin. Guidance and control of an autonomous soaring uav. In *NASA Technical Memorandum*, 2007.
- [8] K Jerry Allwine and Julia E Flaherty. *Joint Urban 2003: Study overview and instrument locations*. Pacific Northwest National Laboratory Richland, WA, 2006.
- [9] K Jerry Allwine, MJ Leach, LW Stockham, JS Shinn, RP Hosker, JF Bowers, and JC Pace. Overview of joint urban 2003's atmospheric dispersion study in oklahoma city. In *Preprints, Symp. on Planning, Nowcasting and Forecasting in the Urban Zone, Seattle, WA, Amer. Meteor. Soc., CD-ROM J*, volume 7, 2004.
- [10] K Jerry Allwine, Joseph H Shinn, Gerald E Streit, Kirk L Clawson, and Mike Brown. Overview of urban 2000: A multiscale field study of dispersion through an urban environment. *Bulletin of the American Meteorological Society*, 83(4):521, 2002.
- [11] Jens Bange and Rainer Roth. Helicopter-borne flux measurements in the nocturnal boundary layer over land—a case study. *Boundary-Layer Meteorology*, 92(2):295–325, 1999.

- [12] John G Bartzis, Diamando Vlachogiannis, and A Sfetsos. Thematic area 5: Best practice advice for environmental flows. *The QNET-CFD Network Newsletter*, 2(4):34–39, 2004.
- [13] Stéphane Bélair, Pierre Lacarrère, Joël Noilhan, Valéry Masson, and Joël Stein. High-resolution simulation of surface and turbulent fluxes during hapex-mobilhy. *Monthly weather review*, 126(8):2234–2253, 1998.
- [14] Bert Blocken and Jan Carmeliet. Pedestrian wind environment around buildings: Literature review and practical examples. *Journal of Thermal Envelope and Building Science*, 28(2):107–159, 2004.
- [15] MC Bogstad, WG Habashi, I Akel, D Ait-Ali-Yahia, N Giannias, and V Longo. Computational-fluid-dynamics based advanced ship-airwake database for helicopter flight simulators. *Journal of aircraft*, 39(5):830–838, 2002.
- [16] Jay Boris, Gopal Patnaik, Mi Young Lee, Theodore Young, Bernd Leitl, Frank Harms, and Michael Schatzmann. Validation of an les urban aerodynamics model for homeland security. In *Proceedings of 47th AIAA Aerospace Sciences Meeting*, pages 5–8, 2009.
- [17] RE Britter and SR Hanna. Flow and dispersion in urban areas. In *Annual Review of Fluid Mechanics*, 2003.
- [18] Michael J Brown, Akshay Gowardhan, Matt Nelson, Mike Williams, and Eric R Pardyjak. Evaluation of the quic wind and dispersion models using the joint urban 2003 field experiment dataset. In *AMS 8th Symposium in Urban Environment, Phoenix, AZ*, page 16, 2009.
- [19] Richard E Brown and M Roger Fedde. Airflow sensors in the avian wing. *Journal of experimental biology*, 179(1):13–30, 1993.
- [20] Paolo Bruschi, Michele Dei, and Massimo Piotto. A low-power 2-d wind sensor based on integrated flow meters. *Sensors Journal, IEEE*, 9(12):1688–1696, 2009.
- [21] IP Castro. Cfd for external aerodynamics in the built environment. *The QNET-CFD Network Newsletter*, 2(2):4–7, 2003.
- [22] Nicola Ceccarelli, John J Enright, Emilio Frazzoli, Steven J Rasmussen, and Corey J Schumacher. Micro uav path planning for reconnaissance in wind. In *American Control Conference, 2007. ACC'07*, pages 5310–5315. IEEE, 2007.
- [23] Anjan Chakrabarty and Jack W Langelaan. Flight path planning for uav atmospheric energy harvesting using heuristic search. In *AIAA Guidance, Navigation and Controls Conference, Toronto, Canada*, 2010.
- [24] Ankush Chakrabarty and Jack Langelaan. Uav flight path planning in time varying complex wind-fields. In *American Control Conference*, 2013.

- [25] WL Chan, CS Lee, and FB Hsiao. Real-time approaches to the estimation of local wind velocity for a fixed-wing unmanned air vehicle. *Measurement Science and Technology*, 22(10):105203, 2011.
- [26] Am Cho, Jihoon Kim, Sanghyo Lee, and Changdon Kee. Wind estimation and airspeed calibration using a uav with a single-antenna gps receiver and pitot tube. *Aerospace and Electronic Systems, IEEE Transactions on*, 47(1):109–117, 2011.
- [27] Alexandre Joel Chorin. A numerical method for solving incompressible viscous flow problems. In *Journal of computational physics*, 1967.
- [28] Ronald M Cionco. A mathematical model for air flow in a vegetative canopy. *Journal of Applied Meteorology*, 4(4):517–522, 1965.
- [29] HW Coleman and Fred Stern. Uncertainties and cfd code validation. *Journal of Fluids Engineering*, 119(4):795–803, 1997.
- [30] Clement Crozon, Rene Steijl, and George N Barakos. Numerical study of helicopter rotors in a ship airwake. *Journal of Aircraft*, 51(6):1813–1832, 2014.
- [31] BZ Cybyk, BE McGrath, TM Frey, DG Drewry, JF Keane, and G Patnaik. Unsteady airflows and their impact on small unmanned air systems in urban environments. In *Journal of Aerospace Information Systems*, 2014.
- [32] Dake. Lee wave. <https://commons.wikimedia.org/w/index.php?curid=1343209>. Accessed: 2016-03-28.
- [33] Dake. Pente thermique. <https://commons.wikimedia.org/w/index.php?curid=1343172>. Accessed: 2016-03-28.
- [34] Dake. Thermal. Accessed: 2010-03-28.
- [35] Nicola de Divitiis. Wind estimation on a lightweight vertical-takeoff-and-landing uninhabited vehicle. *Journal of aircraft*, 40(4):759–767, 2003.
- [36] Atsu SS Dorvlo. Estimating wind speed distribution. *Energy Conversion and Management*, 43(17):2311–2318, 2002.
- [37] James S Forrest, Ieuan Owen, Gareth D Padfield, and Steven J Hodge. Ship-helicopter operating limits prediction using piloted flight simulation and time-accurate airwakes. *Journal of Aircraft*, 49(4):1020–1031, 2012.
- [38] J Franke, C Hirsch, AG Jensen, HW Krüs, M Schatzmann, PS Westbury, SD Miles, JA Wisse, and NG Wright. Recommendations on the use of cfd in wind engineering. In *Cost Action C*, volume 14, page C1, 2004.

- [39] Jorg Franke, Antti Hellsten, K Heinke Schlunzen, and Bertrand Carissimo. The cost 732 best practice guideline for cfd simulation of flows in the urban environment: A summary. In *International Journal of Environment and Pollution*, 2011.
- [40] David Galway. *Urban Wind Modeling with Application to Autonomous Flight*. PhD thesis, Carleton University Ottawa, 2009.
- [41] David Galway, J Etele, and Giovanni Fusina. Modeling of urban wind field effects on unmanned rotorcraft flight. In *Journal of Aircraft*, 2011.
- [42] David Galway, J Etele, and Giovanni Fusina. Development and implementation of an urban wind field database for aircraft flight simulation. In *Journal of Wind Engineering and Industrial Aerodynamics*, 2012.
- [43] David Galway, Jason Etele, and Giovanni Fusina. Modeling of the urban gust environment with application to autonomous flight. In *AIAA Paper*, 2008.
- [44] Michael Gewecke and Martin Woike. Breast feathers as an air-current sense organ for the control of flight behaviour in a songbird (*carduelis spinus*). *Zeitschrift für Tierpsychologie*, 47(3):293–298, 1978.
- [45] Akshay A Gowardhan, Eric R Pardyjak, Inanc Senocak, and Michael J Brown. A cfd-based wind solver for an urban fast response transport and dispersion model. In *Environmental fluid mechanics*, 2011.
- [46] CSB Grimmond and Timothy R Oke. Aerodynamic properties of urban areas derived from analysis of surface form. *Journal of applied meteorology*, 38(9):1262–1292, 1999.
- [47] David Gross, Steve Rasmussen, Phil Chandler, and Greg Feitshans. Cooperative operations in urban terrain (counter). In *Defense and Security Symposium*, 2006.
- [48] DM Hargreaves and Nigel G Wright. On the use of the $k-\epsilon$ model in commercial cfd software to model the neutral atmospheric boundary layer. *Journal of Wind Engineering and Industrial Aerodynamics*, 95(5):355–369, 2007.
- [49] F Harms, B Leitl, M Schatzmann, and G Patnaik. Validating les-based flow and dispersion models. *Journal of Wind Engineering and Industrial Aerodynamics*, 99(4):289–295, 2011.
- [50] Walter G Hoydysh and Walter F Dabberdt. Kinematics and dispersion characteristics of flows in asymmetric street canyons. *Atmospheric Environment (1967)*, 22(12):2677–2689, 1988.
- [51] Haomiao Huang, Gabriel M Hoffmann, Steven L Waslander, and Claire J Tomlin. Aerodynamics and control of autonomous quadrotor helicopters in aggressive maneuvering. In *IEEE Internation Conference on Robotics and Automation*, 2009.

- [52] CG Justus, WR Hargraves, Amir Mikhail, and Denise Graber. Methods for estimating wind speed frequency distributions. *Journal of applied meteorology*, 17(3):350–353, 1978.
- [53] G Kaltsas and AG Nassiopoulou. Novel c-mos compatible monolithic silicon gas flow sensor with porous silicon thermal isolation. *Sensors and Actuators A: Physical*, 76(1):133–138, 1999.
- [54] Seunghyun Kim, Teckjin Nam, and Sekwang Park. Measurement of flow direction and velocity using a micromachined flow sensor. *Sensors and Actuators A: Physical*, 114(2):312–318, 2004.
- [55] Sakbundit Kittiyoungkun. Aerodynamic problems of urban uav operations. 2011.
- [56] Mangal Kothari, Ian Postlethwaite, and Da-Wei Gu. Uav path following in windy urban environments. *Journal of Intelligent & Robotic Systems*, 74(3-4):1013–1028, 2014.
- [57] Jack W Langelaan. Tree-based trajectory planning to exploit atmospheric energy. In *American Control Conference, 2008*, pages 2328–2333. IEEE, 2008.
- [58] Jack W Langelaan, Nicholas Alley, and James Neidhoefer. Wind field estimation for small unmanned aerial vehicles. *Journal of Guidance, Control, and Dynamics*, 34(4):1016–1030, 2011.
- [59] Jack W Langelaan, John Spletzer, Corey Montella, and Joachim Grenestedt. Wind field estimation for autonomous dynamic soaring. In *IEEE International Conference on Robotics and Automation*, 2012.
- [60] Eva Saadé Latorre. Propulsion system optimization for an unmanned lightweight quadrotor. *Univ. Politec. de Catalunya, Master in Aerospace Science and Technology*, 2011.
- [61] NR Lawrance and Salah Sukkarieh. Simultaneous exploration and exploitation of a wind field for a small gliding uav. In *AIAA Guidance, Navigation, and Control Conference*, 2010.
- [62] NR Lawrance and Salah Sukkarieh. Autonomous exploration of a wind field with a gliding aircraft. In *Journal of Guidance, Control, and Dynamics*, 2011.
- [63] NR Lawrance and Salah Sukkarieh. Path planning for autonomous soaring flight in dynamic wind fields. In *IEEE International Conference on Robotics and Automation*, 2011.
- [64] J Gordon Leishman. *Principles of Helicopter Aerodynamics with CD Extra*. Cambridge university press, 2006.

- [65] RW Macdonald. Modelling the mean velocity profile in the urban canopy layer. *Boundary-Layer Meteorology*, 2000.
- [66] RW Macdonald, RF Griffiths, and DJ Hall. An improved method for the estimation of surface roughness of obstacle arrays. *Atmospheric environment*, 32(11):1857–1864, 1998.
- [67] M Mansour, G Kocer, C Lenherr, N Chokani, and RS Abhari. Seven-sensor fast-response probe for full-scale wind turbine flowfield measurements. *Journal of Engineering for Gas Turbines and Power*, 133(8):081601, 2011.
- [68] Matthew Marino, Alex Fisher, Reece Clothier, Simon Watkins, Samuel Prudden, and Chung Sing Leung. An evaluation of multi-rotor unmanned aircraft as flying wind sensors. *International Journal of Micro Air Vehicles*, 7(3):285–299, 2015.
- [69] Stephanie Mayer, Gautier Hattenberger, Pascal Brisset, Marius Jonassen, and Joachim Reuder. A “no-flow-sensor” wind estimation algorithm for unmanned aerial systems. *International Journal of Micro Air Vehicles*, 4(1):15–30, 2012.
- [70] F Menter and B Hemstrom. Cfd best practice guidelines for cfd code validation for reactor safety applications. *Deliverable D01 of the ECORA project*, 2002.
- [71] Richard JD Moore, Saul Thurrowgood, and Mandyam V Srinivasan. Vision-only estimation of wind field strength and direction from an aerial platform. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4544–4549. IEEE, 2012.
- [72] Javier Moyano Cano. Quadrotor uav for wind profile characterization. 2013.
- [73] Jörg Müller, Oliver Paul, and Wolfram Burgard. Probabilistic velocity estimation for autonomous miniature airships using thermal air flow sensors. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 39–44. IEEE, 2012.
- [74] Shuzo Murakami, Ryoza Ooka, Akashi Mochida, Shinji Yoshida, and Sangjin Kim. Cfd analysis of wind climate from human scale to urban scale. *Journal of Wind Engineering and Industrial Aerodynamics*, 81(1):57–81, 1999.
- [75] S Myszchik, M Heller, F Holzapfel, and G Sachs. Low-cost wind measurement system for small aircraft. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, volume 16, page 19, 2004.
- [76] Marina Neophytou, Akshay Gowardhan, and Michael Brown. An inter-comparison of three urban wind models using oklahoma city joint urban 2003 wind field measurements. In *Journal of Wind Engineering and Industrial Aerodynamics*, 2011.

- [77] Patrick P Neumann and Matthias Bartholmai. Real-time wind estimation on a micro unmanned aerial vehicle using its inertial measurement unit. *Sensors and Actuators A: Physical*, 235:300–310, 2015.
- [78] Tim R Oke. *Initial guidance to obtain representative meteorological observations at urban sites*, volume 81. World Meteorological Organization Geneva, 2004.
- [79] TIMOTHY R Oke. *Urban Environments*. McGill–Queen’s University Press, 1997.
- [80] Matthew W Orr, Steven J Rasmussen, Etan D Karni, and William B Blake. Framework for developing and evaluating mav control algorithms in a realistic urban setting. In *American Control Conference*, 2005.
- [81] John Osborne and Rolf Rysdyk. Waypoint guidance for small uavs in wind. *AIAA Infotech@ Aerospace*, 193(1-4):1–12, 2005.
- [82] Meir Pachter, Nicola Ceccarelli, and Philip R Chandler. Estimating mav’s heading and the wind speed and direction using gps, inertial, and air speed measurements. In *Proceedings of the AIAA Guidance, Navigation and Control Conference and Exhibit*, 2008.
- [83] Suhas Patankar. *Numerical heat transfer and fluid flow*. CRC press, 1980.
- [84] M Piotta, G Pennelli, and P Bruschi. Fabrication and characterization of a directional anemometer based on a single chip mems flow sensor. *Microelectronic Engineering*, 88(8):2214–2217, 2011.
- [85] Ludwig Prandtl. Bericht über untersuchungen zur ausgebildeten turbulenz. *Z. Angew. Math. Mech*, 5(2):136–139, 1925.
- [86] Carl Edward Rasmussen. *Gaussian processes for machine learning*. 2006.
- [87] Syed Ali Raza and Jason Etele. Autonomous position control analysis of quadrotor flight in urban wind gust conditions. In *AIAA Guidance, Navigation, and Control Conference*, 2016.
- [88] Sigrid Reiter. Assessing wind comfort in urban planning. *Environment and Planning B: Planning and Design*, 37(5):857–873, 2010.
- [89] Joachim Reuder, Pascal Brisset, Marius Jonassen, Martin Müller, and Stephanie Mayer. The small unmanned meteorological observer sumo: A new tool for atmospheric boundary layer research. *Meteorologische Zeitschrift*, 18(2):141–147, 2009.
- [90] Joachim Reuder and Marius O Jonassen. First results of turbulence measurements in a wind park with the small unmanned meteorological observer sumo. *Energy Procedia*, 24:176–185, 2012.

- [91] J Robadey, O Paul, and H Baltes. Two-dimensional integrated gas flow sensors by cmos ic technology. *Journal of Micromechanics and Microengineering*, 5(3):243, 1995.
- [92] Rainer Röckle. *Bestimmung der Strömungsverhältnisse im Bereich komplexer Bauungsstrukturen*. na, 1990.
- [93] Andrés F Rodriguez, Evan Andersen, Justin M Bradley, and Clark N Taylor. Wind estimation using an optical flow sensor on a miniature air vehicle. 2007.
- [94] Rolf Rysdyk. Unmanned aerial vehicle path following for target observation in wind. *Journal of guidance, control, and dynamics*, 29(5):1092–1100, 2006.
- [95] Gottfried Sachs. Minimum shear wind strength required for dynamic soaring of albatrosses. *Ibis*, 147(1):1–10, 2005.
- [96] Athena Scaperdas and Steve Gilham. Thematic area 4: Best practice advice for civil construction and hvac. *The QNET-CFD Network Newsletter*, 2(4):28–33, 2004.
- [97] JV Seguro and TW Lambert. Modern estimation of the parameters of the weibull wind speed distribution for wind energy analysis. *Journal of Wind Engineering and Industrial Aerodynamics*, 85(1):75–84, 2000.
- [98] Thomas Spiess, Jens Bange, Marco Buschmann, and Peter Vörsmann. First application of the meteorological mini-uav’m2av’. *Meteorologische Zeitschrift*, 16(2):159–169, 2007.
- [99] Mark Sutherland and Jason Etele. Urban wind field generation using les for application to quadrotor flight.
- [100] Nitin Sydney, Brendan Smyth, and Derek A Paley. Dynamic control of autonomous quadrotor flight in an estimated wind field. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 3609–3616. IEEE, 2013.
- [101] Yoshihide Tominaga, Akashi Mochida, Ryuichiro Yoshie, Hiroto Kataoka, Tsuyoshi Nozu, Masaru Yoshikawa, and Taichi Shirasawa. Aij guidelines for practical applications of cfd to pedestrian wind environment around buildings. In *Journal of wind engineering and industrial aerodynamics*, 2008.
- [102] Aline Van den Kroonenberg, Tim Martin, Marco Buschmann, Jens Bange, and Peter Vörsmann. Measuring the wind vector using the autonomous mini aerial vehicle m2av. *Journal of Atmospheric and Oceanic Technology*, 25(11):1969–1982, 2008.
- [103] Caleb White, Simon Watkins, Ee Lim, and Kevin Massey. The soaring potential of a micro air vehicle in an urban environment. *International Journal of Micro Air Vehicles*, 2012.

- [104] J Wieringa. Roughness-dependent geographical interpolation of surface wind speed averages. *Quarterly Journal of the Royal Meteorological Society*, 112(473):867–889, 1986.
- [105] Derrick W Yeo, Nitin Sydney, D Paley, and Donald Sofge. Onboard flow sensing for downwash detection and avoidance with a small quadrotor helicopter. In *Proc. AIAA Guidance Navigation and Control Conference*, 2015.
- [106] Derrick W Yeo, Nitin Sydney, and Derek A Paley. Onboard flow sensing for rotary-wing uav pitch control in wind. In *AIAA Guidance, Navigation, and Control Conference*, page 1386, 2016.
- [107] Derrick W Yeo, Nitin Sydney, Derek A Paley, and Donald Sofge. Downwash detection and avoidance with small quadrotor helicopters.
- [108] George Young, B Gaudet, NL Seaman, and DR Stauffer. Interaction of a mountain lee wave with a basin cold pool. In *Preprints, 13th Conf. on Mesoscale Processes, Salt Lake City, UT, Amer. Meteor. Soc. P*, volume 1, 2009.
- [109] Steven Zan. Technical comment on "computational-fluid-dynamics based advanced ship-airwake database for helicopter flight simulation". *Journal of Aircraft*, 40(5):1007–1007, 2003.
- [110] Ni Zhang. Estimating wind velocities in atmospheric mountain waves using sailplane flight data. 2012.