EVALUATION OF MOISTURE AND HEAT TRANSPORT IN THE FAST-RESPONSE BUILDING-RESOLVING URBAN TRANSPORT CODE QUIC ENVSIM

by

Kevin A. Briggs

A thesis submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Master of Science

Department of Mechanical Engineering

The University of Utah

May 2015

The University of Utah Graduate School

STATEMENT OF THESIS APPROVAL

The thesis of **Kevin A. Briggs**

has been approved by the following supervisory committee members:

Eric R. Pardyjak	, Chair	**12/22/2014** Date Approved
James R. Stoll	, Member	**12/22/2014** Date Approved
Marc Calaf	, Member	**12/22/2014** Date Approved

and by **Tim A. Ameel** , Chair of

the Department of **Mechanical Engineering**

and by David B. Kieda, Dean of The Graduate School.

ABSTRACT

QUIC EnvSim (QES) is a complete building-resolving urban microclimate modeling system developed to rapidly compute mass, momentum, and heat transport for the design of sustainable cities. One of the more computationally intensive components of this type of modeling system is the transport and dispersion of scalars. In this paper, we describe and evaluate QESTransport, a Reynolds-averaged Navier-Stokes (RANS) scalar transport model. QESTransport makes use of light-weight methods and modeling techniques. It is parallelized for Graphics Processing Units (GPUs), utilizing NVIDIA's OptiX application programming interfaces (APIs). QESTransport is coupled with the well-validated QUIC Dispersion Modeling system. To couple the models, a new methodology was implemented to efficiently prescribe surface flux boundary conditions on both vertical walls and flat surfaces. In addition, a new internal boundary layer parameterization was introduced into QUIC to enable the representation of momentum advection across changing surface conditions. QESTransport is validated against the following three experimental test cases designed to evaluate the model's performance under idealized conditions: (i) flow over a step change in moisture, roughness, and temperature, (ii) flow over an isolated heated building, and (iii) flow through an array of heated buildings. For all three cases, the model is compared against published simulation results. QESTransport produces velocity, temperature, and moisture fields that are comparable to much more complex numerical models for each case. The code execution time performance is evaluated and demonstrates linear scaling on a single GPU for problem sizes up to 4.5 x 4.5 km at 5 m grid resolution, and is found to produce results at much better than real time for a 1.2 x 1.2 km section of downtown Salt Lake City, Utah.

For Amy, Ryder and Easton.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would first and foremost like to thank my research advisor Eric Pardyjak for the interest he has continually taken in my research efforts and personal development. He has given me the freedom and opportunity to pursue personal research interests, while providing direction all along the way. Additionally, I wish to thank Rob Stoll for inspiring me to pursue research in the area of CFD, and his insights and support. I acknowledge the work of Chad Allen which laid the foundation for fast response urban transport modeling using the QUIC Dispersion modeling system. None of the QES framework would exist without the dedicated efforts of Scot Halverson and Matt Overby, as overseen by Pete Willemsen.

I would like to thank Prof. Dr. Bernd Leitl from the Meteorologisches Institut at the Universität Hamburg for sharing a supplementary wind tunnel dataset for the Richards et al. (2006) test case. I also thank Dr. J.J. Kim from the Department of Environmental Atmospheric Sciences at Pukyong National University for his expert advice and dataset sharing. This material is based upon work supported by the National Science Foundation under Grant Nos. 0828214, 1134580, and 1133590.

CHAPTER 1

INTRODUCTION

1.1 Background

A large body of work has been produced in recent decades in an effort to better understand the physics of the urban boundary layer (UBL) and to model its various complexities (For an excellent review see Barlow (2014)). As the field's knowledge base and understanding has grown, it has become possible to develop tools that are capable of explaining complex urban processes as well as providing predictive capabilities. The predictive capabilities are particularly exciting, as they have the potential to help provide engineering and planning solutions to many of the most pressing environmental challenges that humans are currently facing (Council, 2001). These include: land-use dynamics, water and air quality, water scarcity, climate change, and energy use. We wish to be able to make place-based predictions for engineering type optimizations, and to potentially perform operational prediction. The simulation tool must be capable of simulating the most critical place-based physics, at relevant scales, in a reasonable amount of time (much faster than real time), so that scenarios can be explored for novel solutions as well as hidden unintended consequences.

Our approach toward addressing these challenges is to develop a robust, fast-running, complete numerical modeling system. Its aim is to produce physically realistic results much faster than real time. The issue of high computational expense is partially addressed by utilizing Graphics Processing Units (GPUs), which provide an inexpensive, highly-parallel computing architecture. Our modeling system, called QUIC EnvSim (QES), builds on and couples to the QUIC (Quick Urban and Industrial Complex) Dispersion modeling system (Brown et al., 2013). QES accounts for the transfer of mass, momentum, and heat in urban areas, and represents a complex system of modules, where each is devoted to modeling a specific component of the urban system. These include an explicit radiative energy exchange model that explicitly includes vegetation (Bailey, 2014; Bailey et al., 2014), a land surface model (LSM) (Shingleton, 2010), a mean wind and turbulence model (Brown,

2004; Pardyjak and Brown, 2001; Singh, 2012; Singh et al., 2008), and a Reynolds-averaged Navier-Stokes (RANS)-based turbulent scalar transport model, QESTransport. Unlike traditional RANS models, however, the velocity field is assumed to be quasi-steady while the scalar field evolves, essentially decoupling momentum and buoyancy effects. This is an approximation that may be reasonable for well mixed urban canopy flow (Roth, 2000), and a major component of this paper involves quantification of this assumption.

This paper focuses on the development and evaluation of QESTransport. We hypothesize that by using the QES techniques mentioned above, we can simulate physically realistic results much faster than real time. Section 1.2 describes current similar efforts in the literature. Section 2.1 describes the details of the scalar transport methodology. A description of the relevant components of the QUIC Dispersion System and improvements needed to couple the system with QESTransport are presented in Section 2.2. Section 2.3 describes the validation of QESTransport using three experimental studies, and more complex computational fluid dynamics (CFD) methods. Finally, the parallel implementation of QESTransport is discussed, and the code's execution time and scalability are investigated in Section 2.4.

1.2 Microscale Modeling Efforts

While many mesoscale urban models exist in the literature (Barlow, 2014; Grimmond et al., 2010, 2011; Martilli et al., 2002; Masson, 2000), far fewer models exist that include full urban microclimate physics and resolve urban features such as buildings and vegetation. Most CFD models (Coirier et al., 2006; Kim and Baik, 2010; Park et al., 2012) explicitly resolve buildings, making them well equipped to capture the effects of spatial heterogeneity within the roughness sublayer (RSL). CFD models are classified, among other features, by their turbulence modeling approach. The most common CFD codes used for urban studies consist of RANS and large eddy simulation (LES) turbulence models. Urban LES approaches generally have higher numerical complexity and have been shown to replicate turbulent and mean flow characteristics with better accuracy than RANS models (Ikegaya et al., 2010; Letzel et al., 2008; Park et al., 2012; Raasch and Schröter, 2001). Lundquist et al. (2012) is one of the few examples that directly couple an LES model to a mesoscale code. Due to their high computational demands, CFD studies have primarily been limited to wind engineering and dispersion applications at single building ($\sim$100 m) to neighborhood ($\sim$1 km) scales with grid spacing $\sim$5 m (Barlow, 2014).

There are currently very few fine-scale CFD urban modeling systems that include full

microclimate physics, likely due to their complexity and expense. The only one to the authors' knowledge is ENVI-met (Bruse, 1995; Bruse and Fleer, 1998), which uses a RANS model with standard $k - \epsilon$ turbulence closure to simulate the velocity and scalar transport (Conry et al., 2014). Although this approach provides results much faster than a comparable LES model, it is still considerably computationally demanding.

For wind fields, fast-response empirical-diagnostic models (Kaplan and Dinard, 1996; Tinarelli et al., 2007) which produce realistic three-dimensional wind fields at a fraction of typical CFD computational expense, have received considerable attention. Neophytou et al. (2011) conducted a study comparing the mean wind fields generated by one such model, QUIC-URB (Brown, 2004; Pardyjak and Brown, 2001) with both RANS (Q-CFD(RANS)) and LES (Q-LES) models, in replicating measurements from the Joint Urban 2003 field campaign held in downtown Oklahoma City (Allwine and Flaherty, 2006). All three models are part of the QUIC modeling system (Nelson and Brown, 2006). It was shown that Q-CFD(RANS) and Q-LES outperformed QUIC-URB, but only slightly. This is an especially noteworthy finding when computation times are compared– QUIC-URB took $\sim$1 minute, Q-CFD(RANS) $\sim$30 minutes, and Q-LES $\sim$30 hours, where QUIC-URB and Q-CFD(RANS) were both run on the same standard PC, and Q-LES on a 8 node parallel cluster (Neophytou et al., 2011).

CHAPTER 2

QESTRANSPORT DESCRIPTION AND EVALUATION

2.1 Turbulent Transport Model

2.1.1 Numerical Method

The primary objective of QESTransport is to model the mean and turbulent transport of heat and water vapor through a complex urban canopy. The Reynolds-averaged form of the advection-diffusion equation is solved using a finite volume method on a Cartesian grid

$$\frac{\partial \overline{\xi}}{\partial t} + \overline{u}_i \frac{\partial \overline{\xi}}{\partial x_i} = \alpha \frac{\partial^2 \overline{\xi}}{\partial x_i \partial x_i} - \frac{\partial}{\partial x_i}(\overline{u_i' \xi'}) + S_\xi. \tag{2.1}$$

Here, t represent the time, x_i the ith Cartesian coordinate, $\overline{u}_i$ the mean velocity component in the ith direction, and $\overline{\xi}$ the mean scalar of interest. $\overline{u_i' \xi'}$ is the turbulent scalar flux, α is the molecular diffusivity of $\overline{\xi}$, and S_ξ represents a volumetric source term.

The turbulent flux term in (2.1) is parameterized by $\overline{u_i' \xi'} = K_\xi(\partial \overline{\xi}/\partial x_i)$, where K_ξ is the turbulent scalar diffusivity. This is computed in the 3D GLE code (Singh, 2012) using the first-order K-theory approach to get the eddy viscosity, K_m, as outlined in (Singh, 2012; Stull, 1988). Substitution into (2.1) results in

$$\frac{\partial \overline{\xi}}{\partial t} + \overline{u}_i \frac{\partial \overline{\xi}}{\partial x_i} = \frac{\partial}{\partial x_i}\left[(\alpha + K_\xi)\frac{\partial \overline{\xi}}{\partial x_i}\right] + S_\xi. \tag{2.2}$$

K_ξ is obtained using the turbulent Prandtl number (Pr_t) relation given as $K_\xi = K_m/Pr_t$ (Kundu and Cohen, 2008). Here, Pr_t is prescribed as 0.9 everywhere, following Kim and Baik (Jae-Jin Kim, personal communication, July 9, 2014).

In line with the objective of building a light-weight, rapid-response solver, a hybrid differencing scheme (HDS) was selected for spatial discretization, which switches between a first-order upwind differencing scheme (UDS) and second-order central differencing scheme (CDS) depending on the local Peclet number (Pe) (see Versteeg and Malalasekera (2007) and Section B.1.1 for details).

The unsteady scalar term on the left-hand side of (2.2) was discretized using a three time level Leapfrog method with a DuFort Frankel approximation, as outlined in Ferziger and Peric (2002) and Section B.1.2. Despite the scheme's unconditional stability, the time step is restricted by the minimum of that allowed by both the CFL condition and the diffusion condition to ensure physically correct behavior is maintained.

Wall boundary conditions enter into (2.2) through the source term S_ξ as a sum of the scalar fluxes through each face of a control volume adjacent to a wall. While it is common practice to apply Monin-Obukhov Similarity Theory (MOST) to urban surface flux models, we decided against it due to its questionable applicability in urban regions (Grimmond et al., 2011; Martilli et al., 2002; Roth, 2000). Instead, we follow the same technique used by Kim and Baik (2010) by applying the thermal wall function outlined in Versteeg and Malalasekera (2007) as

$$\overline{u_i'\xi'} = \frac{-C_\mu^{1/4} k_P^{1/2} (\overline{\xi}_P - \xi_{wall})}{T^+}, \tag{2.3}$$

where C_μ is an empirical constant ($= 0.0845$) and k_P is the turbulent kinetic energy at the near-wall node. The near-wall scalar distribution is modeled by universal function T^+ for high Reynolds number flow , defined as

$$T^+ = Pr_t \left(u^+ + P \left[\frac{Pr}{Pr_t} \right] \right). \tag{2.4}$$

Here Pr is the molecular Prandtl number and u^+ is the nondimensional velocity

$$u^+ = \frac{1}{\kappa} \ln(9.8 z^+), \tag{2.5}$$

where κ is the von Karman constant ($= 0.4$) and z^+ is the nondimensional distance from the wall

$$z^+ = \frac{y_P C_\mu^{1/4} k_P^{1/2}}{\nu}. \tag{2.6}$$

Here, ν is the kinematic viscosity of air. P in (2.4) is given as

$$P = 9.24 \left[\left(\frac{Pr}{Pr_t} \right)^{0.75} - 1 \right] \left\{ 1 + 0.28 \exp \left[-0.007 \left(\frac{Pr}{Pr_t} \right) \right] \right\}. \tag{2.7}$$

The scalar flux model of (2.3) is applicable in the log-law region of the boundary layer where both u^+ and T^+ assume a log-linear shape in z^+. Since k_P is not computed directly, it is approximated by Prandtl's mixing length hypothesis (Pope, 2000; Singh, 2012)

$$k_P = \left(\frac{K_m}{c l_m} \right)^2, \tag{2.8}$$

where c is a constant (≈ 0.55) (Pope, 2000), and l_m is the mixing length, based on the distance to the nearest surface L_{min}. It is taken here as $l_m = \kappa L_{min}$.

QESTransport provides the user with several domain boundary condition options. For the inlet, a uniform, logarithmic, power law, or data defined scalar profile may be prescribed. All other lateral and longitudinal domain faces may either be prescribed as outlet faces (using a standard outflow condition) or as periodic boundaries. Either a zero-flux condition or a floating gradient condition may be set at the domain top, depending on whether the domain extends into the inertial sub-layer or not. The latter condition assigns the top air cell's $\overline{\xi}$ by linear extrapolation from the two cells below (see Section B.1.3.1 for details).

2.2 QUIC-URB Improvements

QUIC-URB is a fast-response 3D mass consistent diagnostic wind model used to generate time-averaged wind fields around buildings, and is based on the methodology developed originally by Röckle (1990). Empirical parameterizations representing various physical features (i.e., building wakes, rooftop recirculations, street canyon vortices, etc.) are applied to generate an initial wind field u_i^o. A mass conservation solver is then run to ensure the field is divergence free (Singh et al., 2008), producing the final mass consistent velocity field u_i.

Expanding the QUIC system beyond traditional contaminant dispersion modeling applications to include moisture and temperature transport processes required new parameterizations to more accurately account for near-wall effects on the flow field. This section presents these improvements, including an internal boundary layer (IBL) model and an explicit wall stress model.

2.2.1 Internal Boundary Layer Model

Urban form is characterized by discontinuities in surface properties such as aerodynamic roughness, land cover/land-use, temperature, humidity, and surface flux of heat and moisture. Horizontal advection across such discontinuities leads to the formation of an IBL, which has a significant effect on the transport of momentum and scalars (Garratt, 1990, 1992; Rao et al., 1974; Rider et al., 1963). Since the QUIC Dispersion Modeling System lacks the functionality to handle variations in aerodynamic roughness length z_0, an IBL model was added to QUIC-URB to include such effects, called QU-IBL.

An in-depth review of IBL growth rate models for near-surface atmospheric boundary conditions is provided by Savelyev and Taylor (2005). Numerous analytical and empirically derived formulae describing widely varying growth rates of the local IBL height (h_b) are presented for both smooth-to-rough and rough-to-smooth transitions (see Table 1 and Fig. 3 of Savelyev and Taylor (2005)). The diffusion analogy class of IBL growth models (Miyake,

1965; Panofsky and Townsend, 1964; Townsend, 1965) define h_b implicitly, requiring computationally costly numerical methods to solve. Instead, an explicit empirically derived model developed by Pendergrass and Aria (1984) was selected for QU-IBL, given as

$$\frac{h_b}{z_{0D}} = 0.32 \left(\frac{x'}{z_{0D}} \right)^{0.8}, \tag{2.9}$$

where x' is the wind-aligned distance from the z_0 discontinuity, and the subscript D refers to the downstream surface (subscript U will hereafter refer to the upstream surface) (Fig. 2.1). IBL 'envelopes' are defined using (2.9), beginning at each z_0 discontinuity and extending along the local mean wind direction to the domain boundary.

The initial velocity field within the IBL, $u_D^o(z)$, is prescribed using the rough-wall log-law model

$$\frac{u_D^o(z)}{u_{*D}} = \frac{\ln(z/z_{0D})}{\kappa}, \tag{2.10}$$

taking $u_D^o\left[h_b(x')\right] = u_U^o\left[h_b(x')\right]$. Following the work of Elliott (1958), $u_{*D}(z)$ is assumed to be constant for $z < h_b$, and u_{*D} is obtained by rearranging and solving (2.10) at $z = h_b$ for all x'

$$u_{*D}(x') = \frac{u_U^o\left[h_b(x')\right]\kappa}{\ln\left[h_b(x')/z_{0D}\right]}. \tag{2.11}$$

Where multiple IBLs exist, $u_D^o(z)$ is assigned beginning with the z_0 transition nearest the inlet, and progressing downstream so that $u_U^o(z)$ is known in (2.11).

Many models include an explicit blending function to transition from the local equilibrium layer to the upwind profile (Chamorro and Porté-Agel, 2008; Panofsky and Townsend, 1964; Townsend, 1965). In QU-IBL, however, blending is achieved by running QUIC-URB's mass conservation solver.

Due to the more dominant recirculation and wake features present in densely built urban areas at neighborhood scales, QU-IBL has limited application in dense urban areas. Regardless, it is used in QUIC-URB's more coarse 'outer' mesh to capture effects of mesoscale surface features in bulk (i.e., coastal regions, rural-urban transitions, etc.), and on the microscale when the length-scale of discontinuous z_0 areas far exceed characteristic building dimensions in close proximity.

This model was validated against the field work of Bradley (1968) for both smooth-to-rough and rough-to-smooth transitions under neutral stability conditions. A roughness transition was created by placing a spiked wire mesh on a smooth tarmac surface, and simultaneous measurements were taken of the mean velocity at several measurement sites downwind (see Bradley's Figs. 2 and 3 for experimental setup).

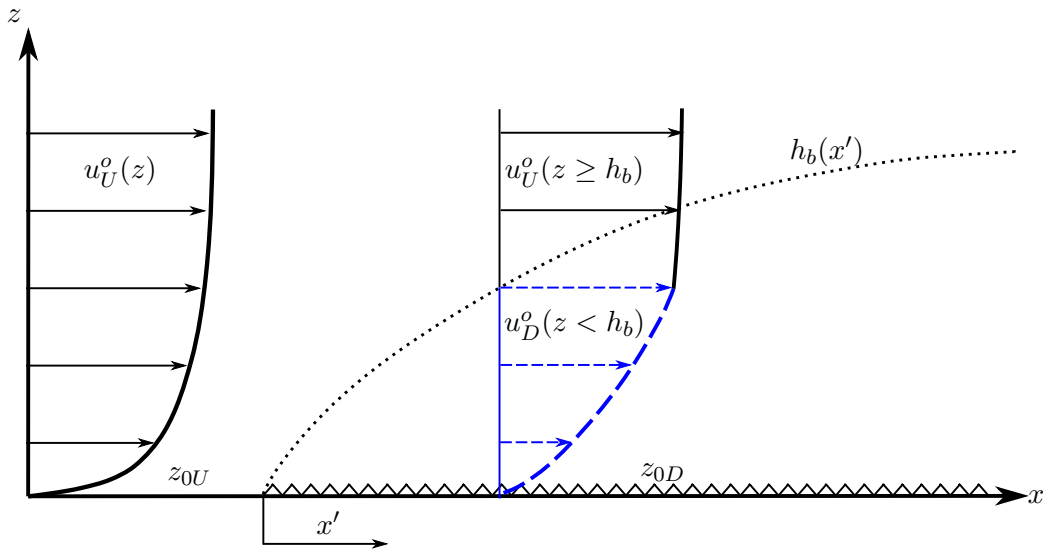

Figure 2.1. Illustration of velocity field assignment in QU-IBL. The dotted line marks $h_b(x')$. The velocity profile at a given x' is prescribed by $u_D(z < h_b)$ (dashed blue curve) and $u_U(z \geq h_b)$ (solid black line). $u_D(z < h_b)$ defines the log-law profile based on the local u_{*D} and $u_U[h_b(x')]$, and $u_U(z \geq h_b)$ defines the profile representing upstream surface conditions.

Profiles of $\overline{u}/u_{ref}$ at four downstream locations give a comparison of the QU-IBL model with Bradley's observations for both the smooth-to-rough (Fig. 2.2) and rough-to-smooth (Fig. 2.3) cases. Key features to note here are 1) the slope of the profile when plotted on a log-linear scale is represented by κ/u_*, and 2) the sudden change in slope seen in Figs. 2.2 and 2.3 is indicative of the IBL height, rising with increased fetch. The lower segment is representative of the local conditions (slope of κ/u_{*D}) and the upper segment of the upwind surface conditions (slope of κ/u_{*U}). QU-IBL captures these features very well in the smooth-to-rough case (Fig. 2.2), indicating good agreement between local and upwind wall stresses, and an appropriate IBL growth model.

For the rough-to-smooth case (Fig. 2.3), the slopes generally agree well, with a slight under prediction of the IBL height. Worth noting is a secondary slope change in the experimental data, which is not replicated by the model. Upon examination of Bradley's Fig. 2 showing the experimental configuration for this case, we see a considerable leading length of smooth tarmac upstream of the wire mesh. This secondary IBL height is likely the result of the initial transition from tarmac to spikes.

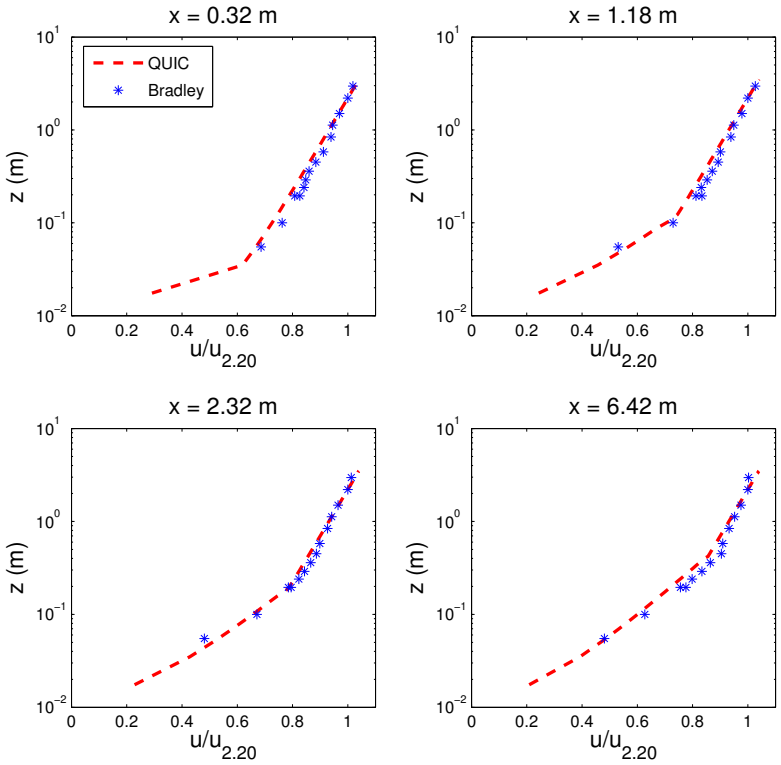

Figure 2.2. Evaluation of the QU-IBL model ($\overline{u}/u_{2.2}$) profiles at four downstream locations over a transition from tarmac to spikes, comparing QU-URB with the experimental data of Bradley (1968). $u_{2.2}$ is the velocity measured at $(x, z) = (0.0, 2.2)$ m.

2.2.2 Wall Stress Model

As discussed in Section 2.1.1, the scalar wall flux model from (2.3) operates under the assumption of a log-law velocity profile (Versteeg and Malalasekera, 2007). Previous versions of QUIC-URB did not explicitly require the computation of wall stresses (Brown et al., 2013; Singh et al., 2008), and while QUIC's computed velocity fields are generally reasonable for dispersion purposes, the near-wall velocity profiles often have unrealistically large gradients. To improve the wall stress estimates, a new approach was implemented which enforces a logarithmic profile at the cells nearest to walls through an iterative procedure. QUIC-URB is run once to rapidly obtain a mass consistent velocity field. A new velocity at the first grid cell above the wall M_1^o is then computed as follows:

First, we define the vector m_i, the projection of the velocity vector u_i onto a plane parallel to the wall of interest.

$$m_{i,2} = R_{ij} u_{j,2}, \tag{2.12}$$

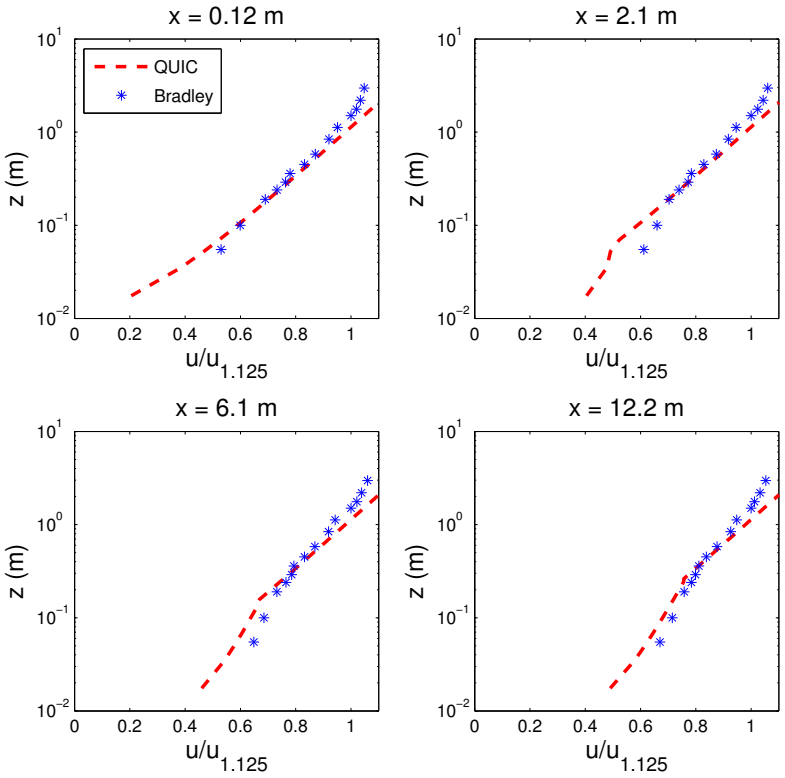

Figure 2.3. Evaluation of the QU-IBL model ($\overline{u}/u_{1.125}$) profiles at four downstream locations over a transition from spikes to tarmac, comparing QU-URB with the experimental data of Bradley (1968). $u_{1.125}$ is the velocity measured at $(x, z) = (0.0, 1.125)$ m.

where $R_{i,j}$ is a matrix of unit vectors defining the planes to be projected onto. Note that subscripts 1 and 2 refer to the cell nearest the wall, and the next cell outward, respectively, and the o superscript refers to the new (non-mass-consistent) velocities. Next, the magnitude of m_i at level 2 is computed by

$$M_2 = (m_{i,2} m_{i,2})^{1/2}.$$

The velocity magnitude nearest the wall, M_1^o, is computed using the log-law approximation and initial guess for u_*:

$$M_1^o = M_2 - \frac{u_*}{\kappa} \ln\left(\frac{n_2}{n_1}\right). \tag{2.13}$$

Finally, M_1^o is partitioned into wall-parallel components

$$u_i^o = \left(\frac{m_{i,2}}{M_2}\right) M_1^o. \tag{2.14}$$

QUIC's mass-consistency solver is run again using u_i^o, and the process is repeated until convergence is reached.

2.3 QESTransport Validation Study

This section presents three validation cases from experimental work, beginning with the simplest case and progressively increasing in geometric complexity. The first considers a step change in surface conditions on flat terrain. The next looks at a single heated cube in a wind tunnel study. The final case involves an array of heated building models in a wind tunnel. The purpose of the validation study was to isolate the QESTransport module of QES as much as possible, and therefore includes only relatively idealized cases with constant surface conditions. More realistic cases with highly spatially and temporally varying conditions will be considered in future works.

2.3.1 Case 1: Advection Across Surface Inhomogeneity

To examine the ability of QESTransport to model the advection of scalars across a discontinuity in surface properties (i.e., aerodynamic roughness and surface temperature and specific humidity), QESTransport was compared to the field data of Rider et al. (1963) and the RANS simulations of Rao et al. (1974). The field experiment of Rider et al. (1963) examined the transition from an extensive tarmac to an adjacent irrigated grassy patch. Measurements of the air temperature ($\overline{T}$) and vapor density ($\overline{\rho}_v$) were taken at the tarmac-grassy patch interface ($x' = 0$ m) and three additional downstream locations ($x' > 0$) at five heights above the ground. The velocity ($\overline{u}$) was measured at $x' = 0$ m and $x' = 16$ m at six heights (Fig. 2 and Table 1 of Rider et al. (1963)). Due to the long fetch over the tarmac, the $\overline{T}$ and $\overline{\rho}_v$ profiles were assumed to be in equilibrium at $x' = 0$ m (Rider et al., 1963).

In the analytical model presented by Rider et al. (1963), the assumption of constant, uniform surface conditions for $x' > 0$ was shown to be appropriate over a relatively short fetch. T_s was measured near $x' = 16$ m and reported for each 10-minute observation period (see Table 2 in Rider et al. (1963)). Due to frequent irrigation and maintenance of the grassy surface, it was treated as freely transpiring vegetation with negligible internal resistance. The humidity boundary condition was then $e_s = e_{sat}(T_s)$, where e_{sat} is the saturation specific humidity at the recorded surface temperature.

The QESTransport simulation domain consisted of a 27 x 5 x 2.25 m volume, with the inlet coincident with the tarmac-grassy patch interface. The grid resolution $(\Delta x, \Delta y, \Delta z) =$ (0.1 m, 0.25 m, 0.02 m), with higher resolution in z to better capture the vertical gradients.

The inlet conditions were fit to the measured profiles for $\overline{T}$, $\overline{\rho}_v$, and $\overline{u}$ at $x' = 0$ m for Rider's observation period 6. Uniform, constant surface conditions were assigned based on T_s and z_0 reported by Rider et al. (1963) (Table 2.1). Due to the step discontinuity in z_0, the QU-IBL model was applied as described in section 2.2.1; however, no $\overline{u}$ measurements were reported at $x' > 0$ for comparison here.

The results produced by QESTransport were compared against a higher-order two-dimensional RANS model developed by Rao et al. (1974). Unlike the analytical approach described by Rider et al. (1963), spatial variation of the surface conditions in the streamwise direction is modeled using equilibrium flux-profile relations.

Modeled vertical profiles of $\overline{T}$ and $\overline{\rho}_v$ are compared with observation in Fig. 2.4. The normalized mean error (NME) and normalized mean bias (NMB) are used to quantify model performance following Dimitrova et al. (2009) (see Appendix for $\overline{T}$ and $\overline{\rho}_v$).

The maximum NME computed for QESTransport of the mean temperature (vapor density) profile comparison was 0.0214 (0.0384) and largest magnitude of the NMB was 0.0194 (-0.0362). Excluding measured values at $x' = 0$, the error for $\overline{T}$ continuously increases and the NMB becomes increasingly positive with fetch. This suggests an under prediction of downward sensible heat flux for the entire fetch after the transition that leads to increasingly over predicted temperatures near the ground with increasing downstream distance (Fig. 2.4, plots *a* - *d*). Rider et al. (1963) noted a sharp decrease in T_s at $x' = 0$ m, with a more gradual decrease to $\sim x' = 8$ m, followed by an increase. Our assumption of a uniform T_s distribution is likely a contributing factor to the increasing error with fetch. Assigning the entire grass surface to the relatively warmer T_s at $x' = 16$ m likely resulted in the downward sensible heat flux underestimation.

QESTransport produces results with similar NME and NMB values as Rao et al. (1974) for both $\overline{T}$ and $\overline{\rho}_v$. Both models capture the advective inversion and growth of a vapor boundary layer of observed by Rider et al. (1963) with good agreement. It should be

Table 2.1. Surface properties/states for tarmac and grass surfaces reported by Rider et al. (1963), where T_s and e_s represent the surface temperature and surface specific humidity, resepectively.

	Tarmac	Grass Surface
z_0 (cm)	0.002	0.14
T_s (°C)	39	28
e_s (g/kg)	6.33*	25.4

*Not specifically reported as a surface value, but extrapolated from measured profile.

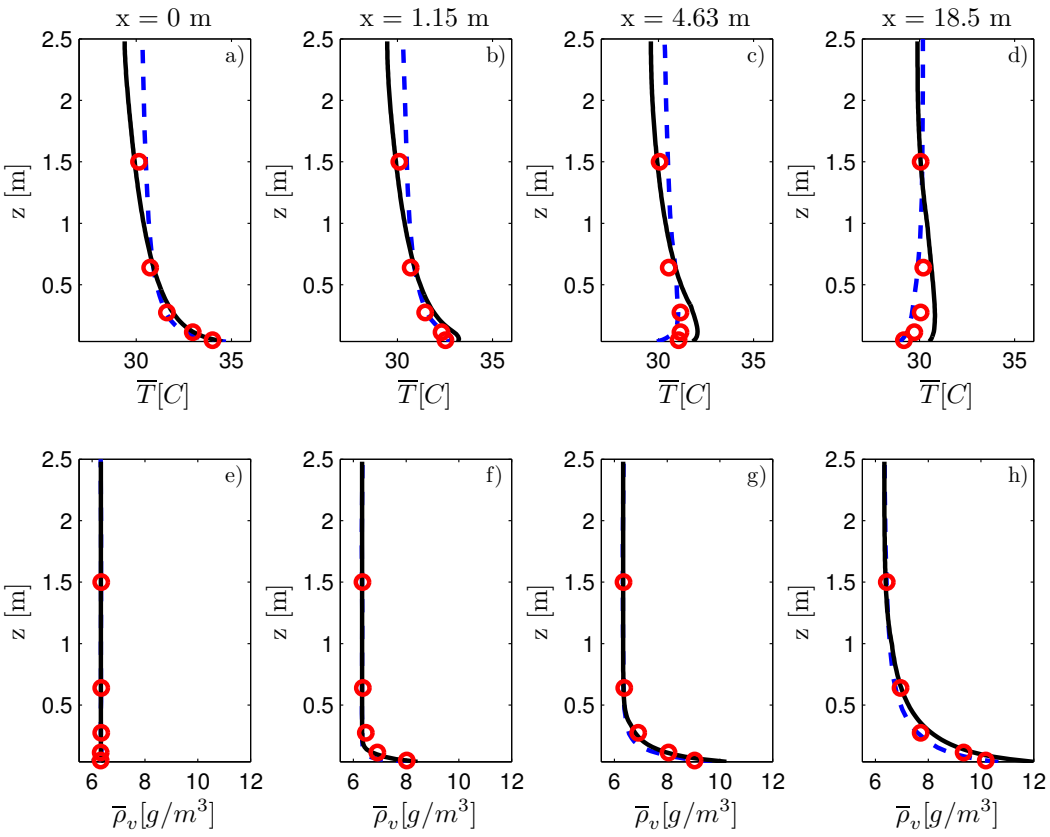

Figure 2.4. Vertical profiles of $\overline{T}$ (plots *a* - *d*) and $\overline{\rho}_v$ (plots *e* - *h*), comparing QESTransport (solid black curves) and Rao et al. (1974) (dashed blue curves) models against Rider et al. (1963) experimental data (open circles).

pointed out that the inlet velocity profile used by Rao et al. (1974) shows poor agreement with experimental measurements, with over predictions near the ground by as much as ∼24 percent (0.65 m/s).

2.3.2 Case 2: Isolated Heated Cube in a Wind Tunnel

The experimental case of Richards et al. (2006) was used to validate QESTransport's ability to simulate turbulent transport around a simple three-dimensional geometry. This involved a wind tunnel study of an isolated cubical building model, scaled to 1:100 of typical building dimensions, with model dimension $H = 0.19$ m. The leeward wall was heated to simulate the effects of solar heating on a wall face. Sharp edged roughness elements and vortex generators were placed upstream to produce a turbulent boundary layer similar to that found in an urban environment. The three-minute averaged inflow wind profile fit a power law with exponent $\alpha = 0.52$ and a roughness length $z_0 = 2.9$ m at full scale (Richards

et al., 2006).

Inertial and buoyant forces were identified as the most dominant effects, represented by the Reynolds number, $Re = \overline{U}_{ref} H/\nu$, and the Grashof number, $Gr = \beta g H^3 \left(\overline{T}_w - \overline{T}_{ref}\right)/\nu^2$, where $Gr \sim 10^9$ signifies transition from laminar to turbulent free convection flow over a vertical flat plate (Bejan and Lage, 1990), $\overline{U}_{ref}$ (≈ 0.5 m/s) is the mean reference velocity at $z = H$, β is the coefficient of thermal expansion for air, and g is the acceleration due to gravity. $\overline{T}_w$ and $\overline{T}_{ref}$ are the mean temperature of the heated wall and reference air temperature, respectively .

Due to the model scale, challenges were faced by Richards et al. (2006) in achieving dynamic similarity for both parameters. Rather than matching full-scale values, Re independence was assumed with $Re = 6291$ (Richards et al., 2006). For a building with $H = 20$ m and a temperature difference ΔT ($= \overline{T}_w - \overline{T}_{ref}$) of 10° C, $Gr \sim 10^{13}$ would be expected. Gr similarity could only be achieved at model scales by using an extreme ΔT, leading to exaggerated buoyancy effects. To avoid this, a limitation of $\overline{T}_w < 200°$ C was applied. For $\overline{T}_{ref} \sim 24°$ C, $Gr \sim 10^8$ was achieved, implying a laminar free-convection boundary layer along the vertical heated wall.

The bulk Richardson number, Ri_b, taken here as the ratio of the Grashof number to the square of the Reynolds number, was used in the study to model the thermal influences on the momentum field near the building, and is given as

$$Ri_b = \frac{Gr}{Re^2} = \frac{\beta g H \left(\overline{T}_w - \overline{T}_{ref}\right)}{\overline{U}^2_{ref}}. \tag{2.15}$$

Two cases from Richards et al. (2006), $Ri_b \approx 0.9$ and ≈ 1.6, were selected for model validation to investigate QESTransport's abilities at different thermal stability regimes. The simulation domain included a 1.52 x 0.95 x 0.57 m section of the wind tunnel, with a uniform grid resolution $\Delta = 0.019$ m, maintaining 10 nodes along each building dimension. The building model was placed at $2H$ downstream of the inlet, centered in the lateral direction. A log-law wind profile matching the parameters above was set at the domain inlet, along with an isothermal temperature profile of $\overline{T}_{ref} = 24°$C. The faces adjacent to the leeward wall were heated to values reported in Table 2.2.

QUIC-URB's vertical profiles of the normalized streamwise velocity ($\overline{u}/U_{ref}$) at four locations along the centerline ($y/H = 0$) are compared against isothermal experimental data (Fig. 2.5, plots *a* - *d*), and values of NME and NMB are presented in the Appendix. Similarly, a comparison of NME and NMB for spanwise mean wind profiles for $z/H = 0.5$ at three near-building locations is given in the Appendix.

Table 2.2. Building model surface temperatures for both $Ri_b \approx 0.9$ and ≈ 1.6, where $\overline{T}_{floor}$ is the floor temperature very near the heated wall, $\overline{T}_{roof}$ is the building model rooftop temperature, and $\overline{T}_{LS}$ and $\overline{T}_{RS}$ are the left and right faces, respectively. All temperatures reported in °C

	$Ri_b \approx 0.9$	$Ri_b \approx 1.6$
$\overline{T}_w$	79	176
$\overline{T}_{floor}$	28	38
$\overline{T}_{roof}$	34	50
$\overline{T}_{LS}$	29	39
$\overline{T}_{RS}$	30	40

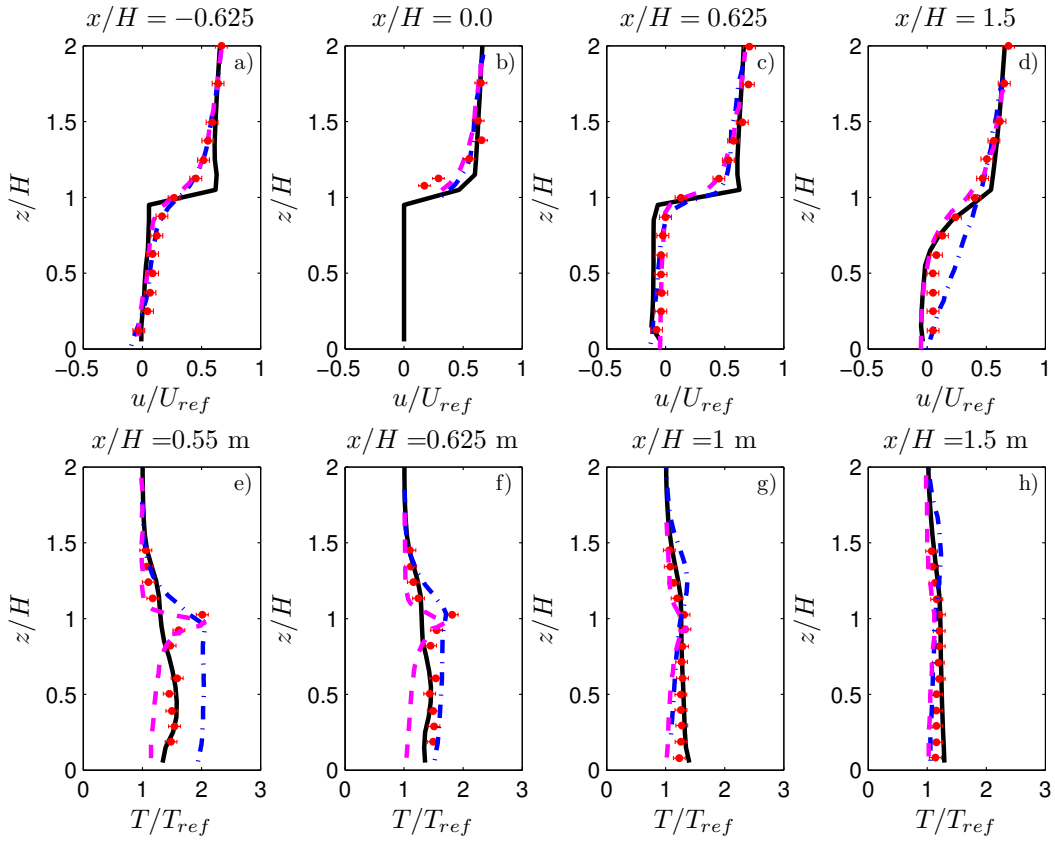

Figure 2.5. Centerline vertical profiles of $\overline{u}/U_{ref}$ at x/H =-0.625, 0.0, 0.625, and 1.5 (plots *a* - *d*), and $\overline{T}/T_{ref}$ at x/H =0.55, 0.625, 1.0, and 1.5 (plots *e* - *h*), with the origin at the building's center. CHENSI (magenta dashed curves), VADIS (blue dash-dotted curves), and QUIC-URB / QESTransport (black solid curves) models are compared against Richards et al. (2006) data (red circles)

The QUIC-URB vertical velocity profiles showed a maximum NME of 0.2878 (at $x/H = 0.625$) with values typically near or below 0.2. These values compared well with NME values from traditional CFD RANS models, CHENSI and VADIS (Dimitrova et al., 2009), that use higher-order turbulence closure models. The largest magnitude of NMB for QUIC-URB was -0.1645 (at $x/H = 0.0$), typically showing less of a bias than these other models (Appendix). In general, we see excessive velocity gradients ($\partial\overline{u}/\partial z$) near $z/H = 1$. As discussed by Singh (2005), this was attributed to the absence of turbulent diffusion in QUIC-URB, and was addressed in Singh et al. (2008).

The vertical profiles of $\overline{T}/T_{ref}$ are compared to the experimental data in Fig. 2.5 (plots e - h) for the $Ri_b \approx 1.6$ case at four locations in the wake of the building. One noteworthy feature of the experimental data is the temperature peak seen near the rooftop height ($z/H = 1$) close to the heated wall ($x/H = 0.55$ and $= 0.625$). See Richards et al. (2006) for a physical explanation. The CHENSI (VADIS) model mimics this behavior, but with severe under prediction (over prediction) of $\overline{T}/T_{ref}$ for $z/H < 1$.

QESTransport misses this near-wall peak altogether due to decoupling of the momentum field from buoyancy effects. It fails to produce a laminar free convection boundary layer along the heated wall, leading to less vertical advection of heat to the rooftop height. Farther downstream, QESTransport slightly over predicts the temperature near $z/H = 0$, potentially due to the over prediction in velocity magnitude near the ground at these locations (Fig. 2.5) causing an over prediction in sensible heat flux. Despite these discrepancies, QESTransport produces results with comparable NME and NMB to the other two models (Appendix), having a maximum NME of 0.1003 (at $x/H = 0.55$) and a maximum NMB of 0.0675 (at $x/H = -0.0675$).

Under moderate thermal forcing ($Ri_b \approx 0.9$), we see much less of a peak observed in the $\overline{T}/T_{ref}$ profiles (Fig. 2.6). The NME and NMB are substantially better under these conditions (Appendix), due to the lower buoyancy effects near the wall.

At this point we return to the discussion on the failure to achieve dynamic similarity with full-scale conditions. With both Re and Gr multiple orders of magnitude larger at full-scale, we would expect to see significantly different transport mechanisms at work. The larger Re is representative of a more turbulent boundary layer throughout the domain. Likewise, Gr at full-scale would suggest a highly turbulent thermal boundary layer near the heated wall. Both of these effects should enhance turbulent mixing and transport away from the wall, thus mitigating the accumulation of heat near the rooftop. Therefore, while QESTransport fails to reproduce the peak observed at wind tunnel scales, it is presumed

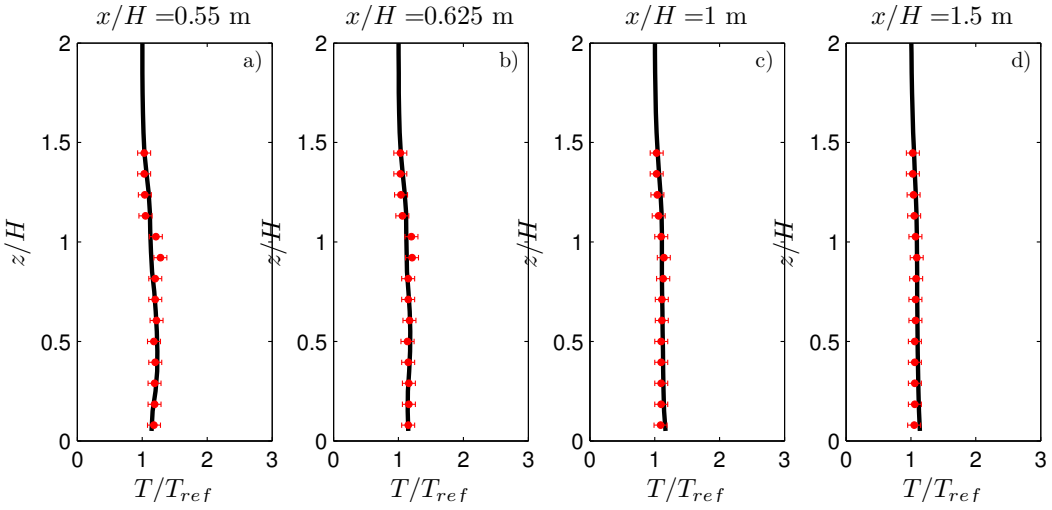

Figure 2.6. Vertical profiles of $\overline{T}/T_{ref}$ comparing QESTransport model (solid black curves) with Richards et al. (2006) data (red circles) for $Ri_b \approx 0.9$

to provide more suitable results at urban scales.

2.3.3 Case 3: Transport in an Array of Heated Buildings

The final case validates QESTransport's ability to represent scalar transport around a more complex morphology. It involves a wind tunnel study conducted by Uehara et al. (2000) examining thermally stratified flow in an urban model, consisting of an array of building models with a heated floor. The experiment was conducted in the atmospheric diffusion wind tunnel at the (Japanese) National Institute for Environmental Studies (Ogawa et al., 1981). Roughness elements were placed along a 10 m fetch upstream of the building array to generate an urban-like turbulent boundary layer at the inlet. The urban model section consisted of Styrofoam cubes with a height H = 100 mm, spaced 100 mm apart in the streamwise (x) direction, and 50 mm in the spanwise (y) direction, forming street canyons perpendicularly aligned to the flow direction (Uehara et al., 2000).

Temperature and wind speed measurements were taken at the center of the street canyon (SC) between the fifth and sixth rows of buildings. For validation, the highly convective case ($Ri_b = -0.21$) was selected where the floor panel temperature, $T_f = 79°$ C, the inlet air temperature, $T_a = 20°$ C, and the building model temperature T_m was set to 39.3° C (Uehara et al., 1997).

For reduced computational expense, the QESTransport domain was limited to a 2 x 0.46 x 0.90 m volume, consisting of a 7x3 array of the building models. A uniform grid with grid spacing Δ = 10 mm was selected, ensuring 10 computational nodes along the

building dimensions. To avoid domain boundary effects, a leading length of $2H$ and trailing length of $5H$ were included, and a symmetric boundary condition was applied on the lateral boundaries. Since QUIC-URB does not allow symmetric boundaries, the velocity domain was expanded to include an extra row of buildings on either side in the spanwise direction (7x5 array of buildings) to reduce edge effects (Singh et al., 2008).

The inlet velocity was specified to match the profile that was provided in Fig. 3 of Uehara et al. (2000), with a reported $z_0 = 3.3$ mm and zero-plane displacement height $d = 35$ mm, using QUIC-URB's urban canopy model (Pardyjak et al., 2008). A reference velocity $U_{ref} = 1.52$ m/s was applied at reference height $z_{ref} = 700$ mm, with an attenuation coefficient $A = 1.75$.

Since only a single $\overline{T}$ profile was reported at the SC test location, the inlet profile was prescribed assuming the profile for $z/H > 1$ was representative of a logarithmic upstream profile, given by Arya (2001)

$$\overline{T}(z) = T_f + \frac{\theta_*}{\kappa}\left[\ln\left(\frac{z}{z_{T0}}\right) - \Psi_H\left(\frac{z}{L}\right)\right], \tag{2.16}$$

where θ_* is the friction temperature, z_{T0} is the thermal roughness length, and Ψ_H is the thermal stability correction function. θ_* at the urban model transition was obtained by rearranging (2.16), and assuming $\overline{T} = T_a$ at the thermal boundary layer height δ_θ. δ_θ was approximated using a self-preserving scalar IBL height model developed by Townsend (1965) which, when integrated, takes the form (Stoll and Porté-Agel, 2006)

$$\delta_\theta\left[\ln\left(\frac{\delta_\theta}{z_0}\right) - 1\right] = 2\kappa^2 x. \tag{2.17}$$

An acceptable fit was obtained using (2.16), with $T_f = 79°$ C, $L = -1.43$ m and $z_{T0} = 0.2$ mm.

The QESTransport results are compared against both a RANS CFD model and a LES model. The RANS model, developed by Kim and Baik (2010) (hereafter referred to as KB10), is a three-dimensional solver which uses the renormalization group (RNG) $k-\epsilon$ turbulence model. The LES model, known as the parallelized LES model (PALM) (Letzel et al., 2008; Raasch and Schröter, 2001), uses the 1.5-order Deardorff (1980) SGS eddy viscosity scheme. The LES results are taken from Park et al. (2012). Vertical profiles of the normalized streamwise velocity $\overline{u}/U_{2H}$ and normalized temperature $(\overline{T} - T_{2H})/(T_f - T_{2H})$ at the SC test site are presented in Fig. 2.7, and a quantification of each model's accuracy (NME and NMB) is provided in the Appendix for wind speed and temperature, respectively.

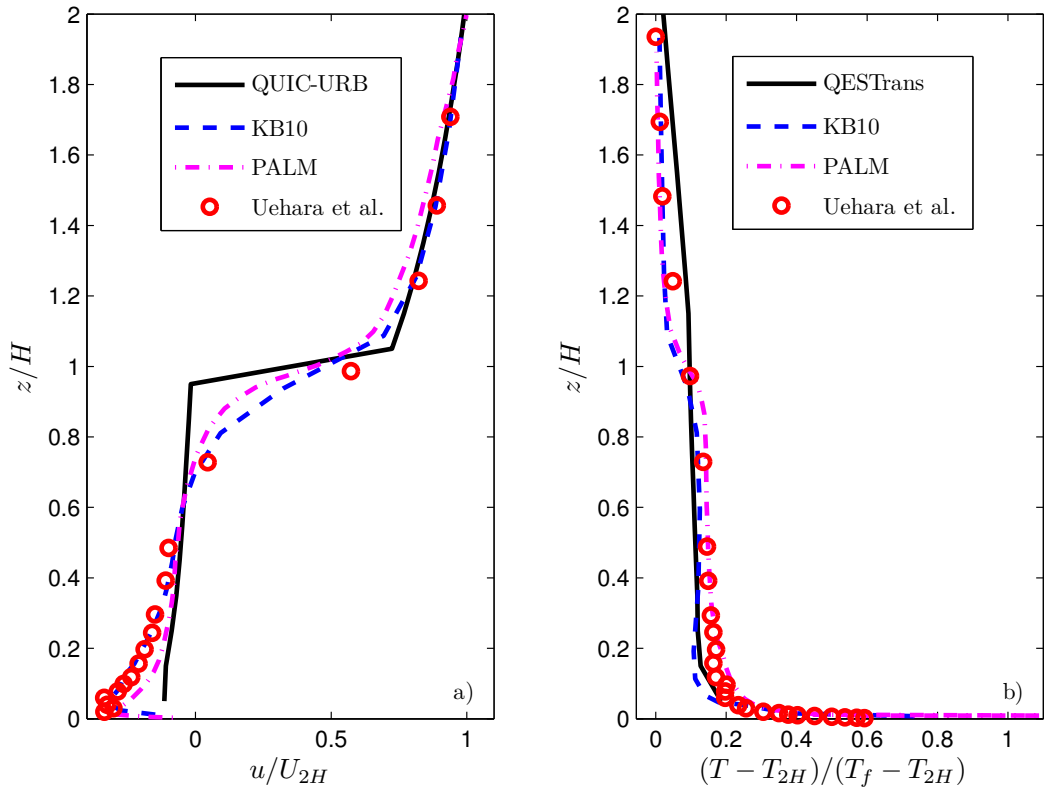

Figure 2.7. Comparison of $\overline{u}/U_{2H}$ (a) and $(\overline{T} - T_{2H})/(T_f - T_{2H})$ (b) for QUIC-URB / QESTransport (solid black curves), KB10 (blue dashed curves), and PALM (magenta dash-dotted curves) against the experimental measurements of Uehara et al. (2000) (red circles), at SC test location.

QUIC-URB appears to under predict the recirculation strength in the SC, especially near $z/H = 0$. As noted in (Uehara et al., 2000), the reverse flow was much stronger in the unstable case versus the neutral case. QUIC-URB's under prediction may be attributed to neglecting buoyancy forces. Additionally, the vertical gradient of $\overline{u}/U_{2H}$ near $z/H = 1$ is rather extreme, consistent with the discussion in Section 2.3.2.

QESTransport slightly under predicts $\overline{T}$ for $z/H < 1$, and slightly over predicts for $1 < z/H < 2$, but otherwise shows good agreement with experimental observations overall (Fig. 2.7). There is an inflection point in the experimental data near $z/H = 1$, which is captured by both the KB10 and PALM models, but not as notably by QESTransport. This is likely the result of enhanced mixing caused by the steep velocity gradients at this height. Quantitatively, QESTransport compares well with the other two models against the experimental measurements, with $NME = 0.0569$ and $NMB = -0.0291$, slightly better than KB10 ($NME = 0.0653$ and $NMB = -0.0646$), and somewhat worse than PALM

($NME = 0.0310$ and $NMB = -0.0221$) (Appendix).

2.4 Parallel Implementation

As discussed above, the two overarching objectives of QESTransort are to provide physically accurate results, and do so very rapidly at full-city scales. The approach taken to achieve the latter is to make use of GPUs. This section describes an initial high-level approach to accelerate the computations utilizing NVIDIA's OptiX and CUDA APIs, including numerical and computational considerations.

OptiX, NVIDIA's GPU accelerated ray-tracing engine, is used heavily by the radiation module of QES (QESRadiant) (Bailey et al., 2014; Overby et al.). OptiX kernels are CUDA kernels with additional macros that define their behavior within the OptiX framework, and provide additional functionality and optimizations for ray-tracing algorithms. In early versions of OptiX, it was difficult to share GPU memory between OptiX and CUDA contexts. Thus, while QESTransport does not use ray-tracing algorithms and therefore makes very little use of OptiX, it was originally developed within the OptiX framework to support interoperation with QESRadiant. Because of this, many of the ray-tracing optimizations provided by OptiX are not applied to the kernels of QESTransport and do not cause errors in execution. However, GPU-based parallelization is still in affect, providing substantial speedup over serial-based computation.

2.4.1 Implementation Details

2.4.1.1 Computational Considerations

Serial CPU-based computing relies on a sophisticated cache hierarchy to accelerate the retrieval of memory. However, as a form of stream processing, GPU computing does not depend on such hardware. A GPU contains multiple types of memory with varying behavior and intended usage. Such types include a thread's individual program stack, shared memory between threads on a multiprocessor, and global memory that is accessible to all threads on the device. Global memory is stored on the GPU's DRAM, and several considerations must be made with respect to its capacity and read/write efficiency. In particular, copying memory from the host (CPU) to the device (GPU) is a notably expensive procedure, and can often be the bottleneck of GPU accelerated software (Yang et al., 2008). In addition, GPUs contain relatively small on-board memory with respect to CPUs. For instance, as of December 2014, a noncustom MacBook Pro might come with 16GB of main memory, but the GPU (NVIDIA 750M) only has 2GB of dedicated memory.

In typical CUDA applications and our implementation, data used in every thread must be resident on the device. In QESTransport, this means physical properties for all air cells and discrete surface "patches" (such as temperature or moisture content) must be copied to the device from the host prior to kernel launch, and the output of each thread must be copied back. To mitigate the cost of this transfer and deal with limited available memory, most simulation input and output is stored as single-precision floating point numbers.

The parallel execution model on NVIDIA's GPUs is similar to the single instruction, multiple data (SIMD) model utilized by early vector supercomputers, in which large numbers of processing elements simultaneously perform the same instructions on many data points. Much higher throughput than serial CPU execution is seen as thousands of execution threads are actively utilized (Kirk and Hwu, 2010). One potential pitfall of this execution model arises with disagreement of thread execution paths, in which each path executes in serial, significantly degrading performance. This creates a condition called branch divergence.

In QESTransport, each thread operates on a specific computation cell. Cells that are on domain boundaries require a separate set of instructions from those on the domain interior. Likewise, cells at the domain inlet require different instructions than those at the outlet or domain top. To avoid branch divergence, the domain is divided such that specialized kernels are launched separately for inlet, outlet, and domain top boundaries, as well as interior cells. We also limit the host-device memory transfer to only the cells needed for each kernel, reducing the communication overhead and overall memory use.

2.4.1.2 Numerical Considerations

The primary numerical consideration for GPU implementation was concerning the time integration scheme. Both implicit and explicit methods were considered. While implicit methods generally have the benefit of remaining computationally stable over large time steps, they typically require an iterative solver and have higher computer memory demands per time step (Ferziger and Peric, 2002). Explicit methods, on the other hand, tend to require less computation time per time step and are much less memory intensive, but have stringent time step requirements to ensure stability (Ferziger and Peric, 2002). Due to the memory restrictions inherent with GPU computing (Kirk and Hwu, 2010), the explicit Leapfrog with DuFort-Frankel point method was determined to be most reasonable for QESTransport, given its relatively low memory demands and unconditional stability.

2.4.2 Computational Performance

2.4.2.1 Execution Benchmark

The Uehara et al. (2000) case (Section 2.3.3) was selected as a benchmark case to evaluate QESTransport's parallel performance against a comparable traditional 3D RANS model. The KB10 model (Kim and Baik, 2010) was chosen for the comparison. As discussed by Lee et al. (2010), comparing GPU vs. CPU times is often misleading. We are cautious to note this study does not serve as a simple comparison of speed increase factor over a serial implementation of the code, as the two models take very different approaches. Rather, the purpose is to evaluate the overall computational cost savings achieved by using the numerical methods described in Section 2.1.1 and the parallel implementation described here.

The QESTransport simulation domain was 200 x 46 x 91 cells (819,000 total) with $\Delta t = 5.2 \times 10^{-4}$ s, as dictated by the CFL condition. The graphics hardware used was a single NVIDIA GeForce GTX Titan on "machine A" (see Table 2.3). The KB10 simulation domain was 152 x 62 x 46 cells (433,504 total), and used a much larger Δt of 0.05 s. This was run in serial on a 2.66 GHz Intel Xeon 5550 CPU. Both models were integrated for $t_f = 180$ s of simulation time. Noting the large differences in domain size and Δt, we define a normalized execution time τ for comparison

$$\tau = \frac{t_{exe}}{(N * P)}, \tag{2.18}$$

where t_{exe} is the total execution time, N is the total number of cells, and P the number of time steps. For QESTransport, t_{exe} encompasses the entire simulation time, including creating the geometry, initialization of data structures, and calculation of boundary conditions at every time step, time integration of the scalar transport on the GPU, and CPU-GPU communication. A cost-saving factor S is then defined as

$$S = \frac{\tau_{Kim}}{\tau_{QES}}, \tag{2.19}$$

Given $t_{exe} \approx 1.44 \times 10^4$ s for the KB10 model and $t_{exe} \approx 5.26 \times 10^3$ s for QESTransport, we achieved $S \approx 500$X.

2.4.2.2 Scaling Study

A scaling study was performed to investigate QESTransport's performance on very large problem sizes. An idealized city was generated consisting of a square array of buildings with dimensions $L = W = 30$ m and $H = 50$ m and uniform spacing in both x and y directions of 30 m. Scaling was evaluated by progressively expanding the dimensions of a square domain

Table 2.3. Hardware specifications for "machine A" used in computations.

OS:	Ubuntu 12.04 LTS
RAM:	64 GB RAM
CPU:	10-core Intel Xeon E5-2690, 3.00GHz
GPU:	NVIDIA GeForce GTX Titan, 2688 CUDA Cores, 6GB GDDR5
Software:	NVIDIA OptiX 3.0.1, CUDA 5.0

ranging from $\sim$ 0.5 to $\sim$ 4.5 km. (Table 2.4), with a constant uniform grid resolution $\Delta = 5$ m. Building faces were heated uniformly to 60 °C above ambient temperature, and the simulation was integrated for $t_f = 12$ s simulation time. The study was performed on "machine A" (2.3).

t_{exe} vs. N is reported in Table 2.4 and Fig. 2.8, where t_{exe} here represents everything but the setup stage (definition of geometry and initialization of data structures). Scaling is nearly linear, with the largest case taking approximately 74 s to run. Assuming linear scaling, the simulation matches real time for a domain dimension of approximately 1.8 km ($\sim$ 6.5 million cells) (red dashed line in Fig. 2.8). With the current implementation and hardware, the problem size is limited to approximately 41 million cells due to GPU memory constraints. This will be addressed in the future through either processing individual "chunks" of the domain on the GPU at a given time, or decomposing the problem across multiple GPUs in parallel.

The uniform, idealized building array was used in the scaling study for convenience and consistency as the problem size was increased. As a more realistic, geometrically complex case, a 1.2 x 1.2 km section of downtown Salt Lake City, Utah, USA was simulated. The geometry was generated with a 5 m grid resolution using Geographic Information System (GIS) data. The tallest building was $\sim$150 m, and the domain height was 250 m. Building faces were uniformly heated to 60 °C above ambient, the ground was heated to 38 °C above

Table 2.4. Setup and results for idealized city scaling study.

Dimension [km]	N [million]	Buildings	t_{exe} [sec]
0.5	0.5	25	2.3
1.6	5.2	529	9.8
3.0	18.4	2116	31.2
3.9	31.0	3721	52.7
4.5	41.3	5041	74.4

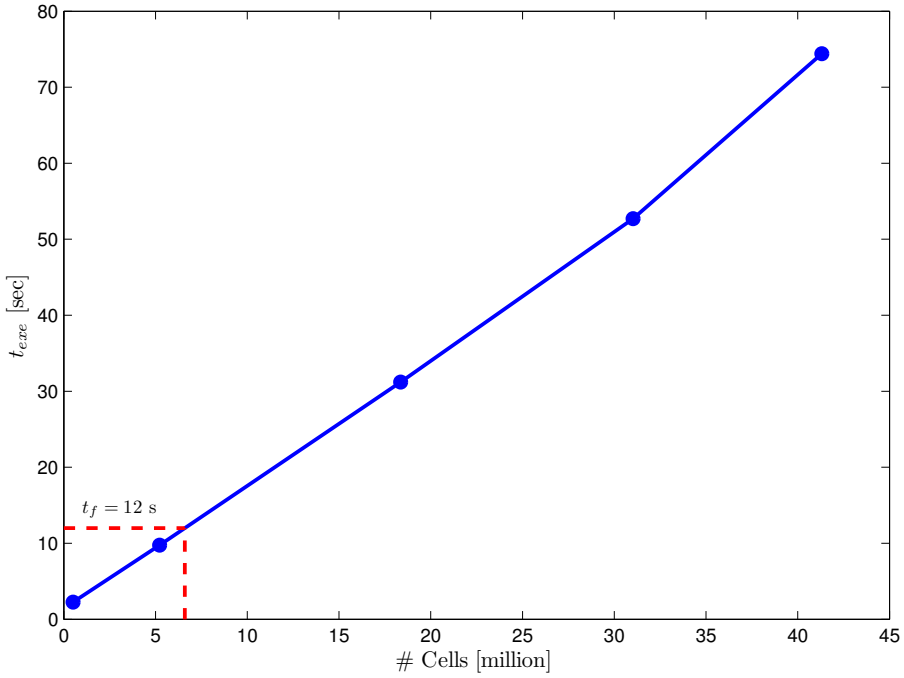

Figure 2.8. Execution time over a range of domain sizes. Red dashed lines correspond to the estimated problem size at real time. Study was run on "machine A" (Table 2.3).

ambient, and the wind direction was set to 135° from north.

The simulation was integrated for $t_f = 600$ s with $\Delta t = 0.205$ s, and required $t_{exe} = 171$ s to run on "computer A" ($\sim$29% of t_f). A contour plot of the temperature field at $z = 5$ m is given in Fig. 2.9.

The parallel optimization approach taken here was relatively naïve in nature, as a more in-depth optimization was beyond the scope of this work. Substantial performance improvements should be attainable as QESTransport is ported to CUDA where more control is afforded over computing resources. Latency will be heavily reduced as memory hierarchy optimizations are leveraged (i.e., utilizing different memory types and caches on the GPU), and larger problems may be considered as computing resources are more judiciously managed. Additionally, host-device communication overhead will be reduced with more control over data transfers between CPU and GPU. Finally, with CUDA it is possible to have multiple kernels (i.e., boundary conditions) simultaneously active on the device allowing for higher throughput, where the current implementation in OptiX does not have this capability.

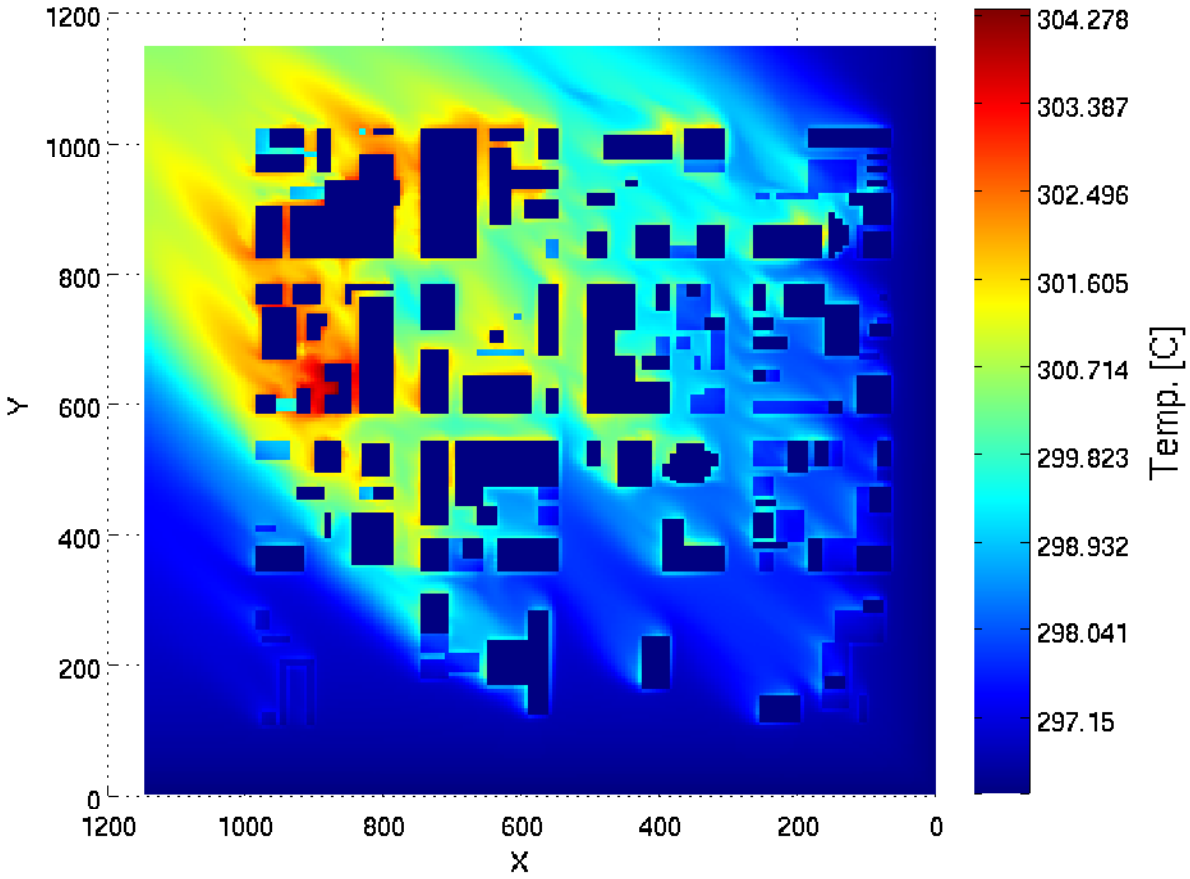

Figure 2.9. Temperature contour plot for 1.2 x 1.2 km section of downtown Salt Lake City, Utah, USA at $z = 5$ m, with a wind angle of 135° from north. All building walls and rooftops were heated to 60 °C above ambient temperature, and the ground surface was heated to 38 °C above ambient

CHAPTER 3

SUMMARY

In order to address place-based urban design questions, a new rapid-response RANS turbulent scalar transport model, QESTransport, has been developed as a component of a complete urban microclimate modeling system, QES. QESTransport is used to compute the mean and turbulent transport of heat and water vapor in a building-resolving complex urban domain, with grid resolutions ∼1-10 m. The key objectives were to produce physically accurate results on an inexpensive computing platform with very low computation times. This was achieved by 1) decoupling momentum and buoyancy effects by assuming a quasi-steady wind field, 2) employing low-cost models and numerical methods, such as a light-weight turbulence closure model and a hybrid differencing spatial discretization scheme, and 3) using GPUs for acceleration of the computations.

A unique feature of QESTransport is that its wind field comes from a fast-response empirical-diagnostic model, QUIC-URB. Previous versions of QUIC-URB do not allow for spatially varying z_0 and therefore do not account for the formation of IBLs. A new model was added to QUIC-URB called QU-IBL which uses an IBL growth rate model at each z_0 transition, and computes the velocity field within the IBL using the rough wall log-law model. QU-IBL was shown to agree well with the experimental measurements of Bradley (1968) for both rough-to-smooth and smooth-to-rough surface transitions. To make QUIC-URB more appropriate for scalar wall flux calculations, a new wall stress model was added to enforce a log-law velocity profile at air cells nearest solid boundaries.

A validation study designed to isolate QESTransport's performance on idealized cases with constant surface conditions was performed. This study consisted of three experimental cases, where QESTransport was shown to replicate physical results with accuracy comparable to much more complex traditional RANS and LES models. A similar validation study involving full-scale observations will be presented in the near future, once all modules of QES are linked together.

While the assumption of negligible buoyancy effects has observational justification, most

of the local discrepancies between the RANS/LES models and QESTransport may be attributed to the lack of buoyancy coupling in the latter. Its applicability is generally limited to regions with relatively high building density. As the ultimate goal of QES is to simulate entire cities, there will certainly be instances where this assumption is more suspect. Future work will include adding parameterizations to QUIC-URB to account for buoyancy effects on the momentum field.

The computations were accelerated on consumer-level GPUs for inexpensive, effective parallelization. Special considerations were taken in the model numerics and implementation to operate within the SIMD-style parallel execution model. A benchmark study was carried out on the Uehara et al. (2000) case, comparing QESTransport's execution speed with the KB10 RANS model (Kim and Baik, 2010). When normalized by number of cells and time step size, QESTransport showed a cost-saving factor of $\sim$500X.

Model execution was tested on problems with up to $\sim$ 41 million cells and shown to scale nearly linearly for domain sizes up to 4.5 x 4.5 km, beyond which the memory is exceeded for a single consumer-level GPU. Compared to the majority of fine-scale CFD urban studies on domains of $<$1 km in extent (Barlow, 2014), this is a noteworthy feat. Under the current implementation, results are produced at better than real time for a domain size of $\sim$ 1.8 km. Work has begun on exploring multi-GPU parallelization for certain components of QES (Overby, 2014) and will continue for QESTransport to handle much larger domains.

Given the performance of QESTransport, it shows promise for use as a fine-scale operational model, especially when nested within a mesoscale model. Work towards this is well under way, as seen in Kochanski et al. (2015) involving a one-way coupling of WRF to the QUIC Dispersion Modeling System.

APPENDIX A

VALIDATION STUDY- TABULATIONS OF ERROR AND BIAS

The metrics used for comparing QESTransport's results against published data were the normalized mean error

$$NME = \frac{\sum_{i=1}^{N} |P_i - O_i|}{\sum_{i=1}^{N} O_i}, \tag{A.1}$$

and the normalized mean bias

$$NMB = \frac{\sum_{i=1}^{N} (P_i - O_i)}{\sum_{i=1}^{N} O_i}, \tag{A.2}$$

where P_i and O_i correspond to the predicted and observed values, respectively, at each measurement location i.

Table A.1. NME and NMB for $\overline{T}$ profiles for Rao et al. (1974) and QESTransport models against Rider et al. (1963) data.

x (m)	NME		NMB	
	Rao	QESTransport	Rao	QESTransport
0.0	0.0063	0.0060	0.0010	-0.0043
1.15	0.0071	0.0103	0.0071	0.0060
4.63	0.0143	0.0158	-0.0041	0.0121
18.5	0.0090	0.0214	-0.0078	0.0194

Table A.2. NME and NMB for $\overline{\rho}$ profiles for Rao et al. (1974) and QESTransport models against Rider et al. (1963) data.

x (m)	NME		NMB	
	Rao	QESTransport	Rao	QESTransport
0.0	0.0026	0.0021	0.0015	0.0002
1.15	0.0368	0.0254	-0.0368	-0.0254
4.63	0.0285	0.0183	-0.0206	-0.0033
18.5	0.0210	0.0384	-0.0198	0.0362

Table A.3. NME and NMB values for vertical $\overline{u}/\overline{U}_{ref}$ profiles from CHENSI, VADIS, and QESTransport models against Richards et al. (2006) data, under isothermal conditions.

x/H	NME			NMB		
	CHENSI	VADIS	QUIC-URB	CHENSI	VADIS	QUIC-URB
-0.625	0.125	0.068	0.1937	-0.119	-0.037	0.009
0.0	0.156	0.171	0.2008	0.021	0.071	-0.1645
0.625	0.125	0.068	0.2878	-0.119	-0.037	-0.0545
1.5	0.213	0.297	0.1626	-0.205	0.252	-0.0595

Table A.4. NME and NMB values for spanwise $\overline{u}/\overline{U}_{ref}$ profiles from CHENSI, VADIS, and QESTransport models against Richards et al. (2006) data, under isothermal conditions.

x/H	NME			NMB		
	CHENSI	VADIS	QUIC-URB	CHENSI	VADIS	QUIC-URB
-0.625	0.165	0.133	0.6802	-0.087	0.116	-0.4019
0	0.105	0.124	0.4318	-0.089	0.116	-0.3619
1.5	0.274	0.227	0.7066	-0.274	0.087	-0.4123

Table A.5. NME and NMB values for spanwise $\overline{T}/\overline{T}_{ref}$ profiles from CHENSI, VADIS, and QESTransport models against Richards et al. (2006) data, with $Ri_b \approx 1.6$

x/H	NME			NMB		
	CHENSI	VADIS	QESTransport	CHENSI	VADIS	QESTransport
0.55	0.115	0.283	0.1003	-0.071	0.264	-0.0372
0.625	0.195	0.082	0.0895	-0.195	0.066	-0.0675
1	0.104	0.11	0.0452	-0.104	-0.007	0.0312
1.5	0.073	0.071	0.0556	-0.073	-0.012	0.0556

Table A.6. Comparison of NME and NMB for $\overline{T}/\overline{T}_{ref}$ profiles produced by QESTransport against Richards et al. (2006) data, at $Ri_b \approx 0.9$ and ≈ 1.6

x/H	NME		NMB	
	0.9	1.6	0.9	1.6
0.55	0.0399	0.1003	-0.0092	-0.0372
0.625	0.0276	0.0895	-0.0012	-0.0675
1	0.0283	0.0452	0.0199	0.0312
1.5	0.0331	0.0556	0.0328	0.0556

Table A.7. Model comparison of NME and NMB for $\overline{u}/U_{2H}$ between KB10, PALM, and QUIC-URB models against Uehara et al. (2000) experimental data at the SC test location.

NME			NMB		
KB10	PALM	QUIC-URB	KB10	PALM	QUIC-URB
0.1541	0.4577	0.6902	-0.0813	0.1092	0.2863

Table A.8. Model comparison of NME and NMB for $(\overline{T} - T_{2H})/(T_f - T_{2H})$ between KB10, PALM, and QESTransport models against Uehara et al. (2000) experimental data at the SC test location.

NME			NMB		
KB10	PALM	QESTrans.	KB10	PALM	QESTrans.
0.0653	0.0310	0.0569	-0.0646	0.0221	-0.0291

APPENDIX B

NUMERICAL DETAILS

B.1 Numerical Methods

The transport equation (2.2) is solved using a finite volume method on a cartesian grid. It can be rewritten in integral form as

$$\begin{aligned}\int_V \frac{\partial \overline{\xi}}{\partial t} dV &+ [(u\overline{\xi})_e - (u\overline{\xi})_w] A_{ew} \\ &+ [(v\overline{\xi})_n - (v\overline{\xi})_s] A_{ns} + [(w\overline{\xi})_u - (w\overline{\xi})_d] A_{ud} \\ &= \left[\left((\alpha + K_\xi) \frac{\partial \overline{\xi}}{\partial x} \right)_e - \left((\alpha + K_\xi) \frac{\partial \overline{\xi}}{\partial x} \right)_w \right] A_{ew} \\ &+ \left[\left((\alpha + K_\xi) \frac{\partial \overline{\xi}}{\partial y} \right)_n - \left((\alpha + K_\xi) \frac{\partial \overline{\xi}}{\partial y} \right)_s \right] A_{ns} \\ &+ \left[\left((\alpha + K_\xi) \frac{\partial \overline{\xi}}{\partial z} \right)_u - \left((\alpha + K_\xi) \frac{\partial \overline{\xi}}{\partial z} \right)_d \right] A_{ud} \\ &+ \int_V S_\xi dV\end{aligned} \tag{B.1}$$

For readability, we introduce the variables F ($= u_i A$) and D ($= \frac{\alpha + K_\xi}{\Delta x_i} A$). Carrying out the volume integrals and rearranging, (B.1) becomes

$$\begin{aligned}\frac{\partial \overline{\xi}}{\partial t} V = &- [F_e \overline{\xi}_e - F_w \overline{\xi}_w] - [F_n \overline{\xi}_n - F_s \overline{\xi}_s] - [F_u \overline{\xi}_u - F_d \overline{\xi}_d] \\ &+ [D_e(\overline{\xi}_E - \overline{\xi}_P) - D_w(\overline{\xi}_P - \overline{\xi}_W)] \\ &+ [D_n(\overline{\xi}_N - \overline{\xi}_P) - D_s(\overline{\xi}_P - \overline{\xi}_S)] \\ &+ [D_u(\overline{\xi}_U - \overline{\xi}_P) - D_d(\overline{\xi}_P - \overline{\xi}_D)] \\ &+ S_\xi V\end{aligned} \tag{B.2}$$

Note here that the lower-case subscripts correspond to the value at the center of a control volume face (face-centered) while capital subscripts refer to values located at the center of the control volume (cell-centered), with the east and west directions along the x-axis, north and south along the y-axis, up and down along the z-axis, and P at the cell's center. As

each quantity is defined at the center of the control volume, the face-centered values are obtained by taking the algebraic mean of neighboring cell-centered values. For example,

$$\overline{\xi}_e = \frac{\overline{\xi}_P + \overline{\xi}_E}{2}. \tag{B.3}$$

B.1.1 Spatial Discretization

To provide physically correct solutions, a discretization scheme must effectively account for the relative strength between advection and diffusion in determining the influence of the surrounding cells on the local transport of $\overline{\xi}$, a property of the scheme referred to by Versteeg and Malalasekera (2007) as transportiveness. For example, where diffusion dominates, nearly equal influence should be assigned to each of the surrounding cells. When advection dominates, on the other hand, the cells in the upwind direction should have significantly greater influence on the transport calculation. The Peclet number, defined as $Pe = F/D$, provides a ratio of the advective to diffusive strength.

In line with the objective of building a light-weight, rapid-response solver, only lower order truncation error schemes are considered. While a second-order central differencing scheme (CDS) certainly provides a solution with higher order of accuracy than a first-order upwind differencing scheme (UDS), it does not possess the transportive property described above at higher Peclet numbers, giving equal influence to all surrounding cells. An UDS on the other hand is transportive, but has a lower order of accuracy. To take advantage of the desirable properties of both the CDS and UDS, a hybrid differencing scheme (HDS) is used which switches from CDS to UDS at cells where $|Pe|$ exceeds a threshold value of two (Ferziger and Peric, 2002; Versteeg and Malalasekera, 2007)). Equation (B.2) may now be simplified by substituting in the face-centered values illustrated by (B.3), resulting in

$$\begin{aligned}\frac{\partial \xi}{\partial t} V = &- a_P \xi_P + a_E \xi_E + a_W \xi_W + a_N \xi_N \\ &+ a_S \xi_S + a_U \xi_U + a_D \xi_D + S_\xi V.\end{aligned} \tag{B.4}$$

The a coefficients in (B.4) provide the means for implementing the HDS, where

$$
\begin{aligned}
a_E &= max[-F_e, (D_e - F_e/2), 0] \\
a_W &= max[F_w, (D_w + F_w/2), 0] \\
a_N &= max[-F_n, (D_n - F_n/2), 0] \\
a_S &= max[F_s, (D_s + F_s/2), 0] \\
a_U &= max[-F_u, (D_u - F_u/2), 0] \\
a_D &= max[F_d, (D_d + F_d/2), 0] \\
a_P &= a_E + a_W + a_N + a_S + a_U + a_D + \Delta F \\
\Delta F &= F_e - F_w + F_n - F_s + F_u - F_d
\end{aligned}
\tag{B.5}
$$

B.1.2 Time Integration

The unsteady scalar term on the left-hand side of (B.2) was discretized using a three time level Leapfrog method with a DuFort Frankel approximation, as outlined in Ferziger and Peric (2002). Applying this method to (B.4) results in

$$
\begin{aligned}
&\left(1 + \frac{\Delta t}{V} a_P\right) \xi_P^{n+1} \\
&\qquad = \left(1 - \frac{\Delta t}{V} a_P\right) \xi_P^{n-1} \\
&\qquad + (a_E \xi_E^n + a_W \xi_W^n + a_N \xi_N^n + a_S \xi_S^n + a_U \xi_U^n + a_D \xi_D^n \\
&\qquad + \sum_{f=1}^{6} \left(\overline{u_i' \xi'}\right)_f A_f) \frac{2\Delta t}{V}.
\end{aligned}
\tag{B.6}
$$

The n superscript indicates the time level. Using the Von Neumann stability analysis, the method is shown to be unconditionally stable. Despite this, the time step is restricted by the minimum of that allowed by both the CFL condition and the diffusion condition to ensure physically correct behavior is maintained. These are given by

$$
\Delta t_{CFL} = \frac{\mu \left[(\Delta x)^2 + (\Delta y)^2 + (\Delta z)^2\right]^{1/2}}{\left(u_{max}^2 + v_{max}^2 + w_{max}^2\right)^{1/2}}, \tag{B.7}
$$

$$
\Delta t_{diff} = \frac{\left[(\Delta x)^2 + (\Delta y)^2 + (\Delta z)^2\right]^{1/2}}{2K_{\xi,max}}, \tag{B.8}
$$

respectively, where μ is a constant ($= 0.9$) and the grid spacing is assumed to be constant in each direction.

B.1.3 Boundary Conditions

B.1.3.1 Domain Top Boundary Condition

A floating gradient condition is set at the domain top, defined by

$$\frac{\partial \xi}{\partial z} \mid_{z=L_z} = \frac{\partial \xi}{\partial z} \mid_{z=L_z-\Delta z}, \tag{B.9}$$

where L_z is the domain height and Δz is the vertical grid spacing.

BIBLIOGRAPHY

K.J. Allwine and J.E. Flaherty. Joint Urban 2003: Study Overview and Instrument Locations. (Tech Rep PNNL-15967, pp. 192). Pacific Northwest National Laboratory. Technical report, 2006.

S. Pal Arya. *Introduction to Micrometeorology*. Academic Press, second edition, 2001. ISBN 9780120593545. doi: 0120593548.

B.N. Bailey. A new three-dimensional energy balance model for complex plant canopy geometries: Model development and improved validation strategies. *Submitt. Rev. to Agric. For. Meteorol.*, 2014.

B.N. Bailey, M. Overby, P. Willemsen, E.R. Pardyjak, W.F. Mahaffee, and R. Stoll. A scalable plant-resolving radiative transfer model based on optimized GPU ray tracing. *Agric. For. Meteorol.*, 2014.

Janet F. Barlow. Progress in observing and modelling the urban boundary layer. *Urban Clim.*, 44, August 2014. ISSN 22120955. doi: 10.1016/j.uclim.2014.03.011. URL http://linkinghub.elsevier.com/retrieve/pii/S2212095514000558.

A. Bejan and J.L. Lage. The Prandtl number effect on the transition in natural convection along a vertical surface. *J. Heat Transfer*, 112:787–790, 1990.

E.F. Bradley. A micrometeorological study of velocity profiles and surface drag in the region modified by a change in surface roughness. *Q. J. R. Meteorol. Soc.*, 94(401):361–379, 1968. URL http://onlinelibrary.wiley.com/doi/10.1002/qj.49709440111/abstract.

Michael J Brown. Urban Dispersion- Challenges for Fast Response Modeling. In *13th Confernce Appl. Air Pollut. Meteorolgy with Air Waste Manag. Assoc*, pages 1–12, 2004.

Michael J Brown, Akshay A Gowardhan, Mathew A Nelson, Michael D Williams, and Eric R Pardyjak. QUIC transport and dispersion modelling of two releases from the Joint Urban 2003 field experiment. *Int. J. Environ. Pollut.*, 52:263–287, 2013.

M Bruse. *Development of a microscale model for the calculation of surface temperatures in structured terrain, MSc Thesis, Inst. for Geography, Univ. Bochum.* PhD thesis, 1995.

M Bruse and H Fleer. On the Simulation of Surface-Plant-Air Interactions Inside Urban Environments. *Environmental Modelling & Software*, 13:373–384, 1998.

Leonardo P. Chamorro and Fernando Porté-Agel. Velocity and Surface Shear Stress Distributions Behind a Rough-to-Smooth Surface Transition: A Simple New Model. *Boundary-Layer Meteorol.*, 130(1):29–41, December 2008. ISSN 0006-8314. doi: 10.1007/s10546-008-9330-x. URL http://link.springer.com/10.1007/s10546-008-9330-x.

W. J. Coirier, D. M. Fricker, M. Furmanczyk, and S. Kim. A Computational Fluid Dynamics Approach for Urban Area Transport and Dispersion Modeling. *Environ. Fluid Mech.*, 5 (5):443–479, January 2006. ISSN 1567-7419. doi: 10.1007/s10652-005-0299-4. URL http://link.springer.com/10.1007/s10652-005-0299-4.

Patrick Conry, H J S Fernando, L S Leo, and Notre Dame. Multi-scale simulations of climate-change influence on Chicago heat island. In *Proc. 4th Jt. US-European Fluids Eng. Div. Summer Meet. FEDSM2014, August 3-7, 2014, Chicago, Illinois, USA*, pages 1–11, 2014.

National Research Council. *Grand Challeneges in Environmental Sciences.* National Academy of Sciences, 2001.

J.W. Deardorff. Stratocumulus-capped mixed layers derived from a three-dimensional model. *Boundary-Layer Meteorol.*, 18:495–527, 1980.

R. Dimitrova, Jean-François Sini, K. Richards, M. Schatzmann, M. Weeks, E. Perez García, and C. Borrego. Influence of Thermal Effects on the Wind Field Within the Urban Environment. *Boundary-Layer Meteorol.*, 131(2):223–243, March 2009. ISSN 0006-8314. doi: 10.1007/s10546-009-9368-4. URL http://link.springer.com/10.1007/s10546-009-9368-4.

W.P. Elliott. The Growth of the Atmospheric Internal Boundary Layer. *Am. Geophys. Union*, 39:1048–1054, 1958.

Joel H. Ferziger and Milovan Peric. *Computational Methods for Fluid Dynamics.* Springer, 2002. ISBN 978-3-642-56026-2.

J.R. Garratt. The Internal Boundary Layer - A Review. *Boundary-Layer Meteorology*, 50: 171–203, 1990.

J.R. Garratt. *The Atmospheric Boundary Layer*. Cambridge University Press, 1992. ISBN 9780521380522.

C. S. B. Grimmond, M. Blackett, M. J. Best, J. Barlow, J-J. Baik, S. E. Belcher, S. I. Bohnenstengel, I. Calmet, F. Chen, a. Dandou, K. Fortuniak, M. L. Gouvea, R. Hamdi, M. Hendry, T. Kawai, Y. Kawamoto, H. Kondo, E. S. Krayenhoff, S-H. Lee, T. Loridan, a. Martilli, V. Masson, S. Miao, K. Oleson, G. Pigeon, a. Porson, Y-H. Ryu, F. Salamanca, L. Shashua-Bar, G-J. Steeneveld, M. Tombrou, J. Voogt, D. Young, and N. Zhang. The International Urban Energy Balance Models Comparison Project: First Results from Phase 1. *J. Appl. Meteorol. Climatol.*, 49(6):1268–1292, June 2010. ISSN 1558-8424. doi: 10.1175/2010JAMC2354.1. URL http://journals.ametsoc.org/doi/abs/10.1175/2010JAMC2354.1.

C. S. B. Grimmond, M. Blackett, M. J. Best, J.-J. Baik, S. E. Belcher, J. Beringer, S. I. Bohnenstengel, I. Calmet, F. Chen, a. Coutts, a. Dandou, K. Fortuniak, M. L. Gouvea, R. Hamdi, M. Hendry, M. Kanda, T. Kawai, Y. Kawamoto, H. Kondo, E. S. Krayenhoff, S.-H. Lee, T. Loridan, a. Martilli, V. Masson, S. Miao, K. Oleson, R. Ooka, G. Pigeon, a. Porson, Y.-H. Ryu, F. Salamanca, G.J. Steeneveld, M. Tombrou, J. a. Voogt, D. T. Young, and N. Zhang. Initial results from Phase 2 of the international urban energy balance model comparison. *Int. J. Climatol.*, 31(2):244–272, February 2011. ISSN 08998418. doi: 10.1002/joc.2227. URL http://doi.wiley.com/10.1002/joc.2227.

N. Ikegaya, A. Hagishima, J. Tanimoto, S. Raasch, and L. Marcus. Large-eddy simulation of scalar transport phenomena between an urban surface and atmosphere. In *fifth Int. Symp. Comput. Wind Eng. CHapel Hill, North Carolina, USA, May 23-27*, 2010.

H. Kaplan and N. Dinard. A Lagrangian dispersion model for calculating concentration distribution within a built-up domain. *Atmospheric Environment*, 30(24):4197 – 4207, 1996.

Jae-jin Kim and Jong-jin Baik. Effects of Street-Bottom and Building-Roof Heating on Flow in Three-Dimensional Street Canyons. *Adv. Atmos. Sci.*, 27(3):513–527, 2010. doi: 10.1007/s00376-009-9095-2.1.Introduction. URL http://link.springer.com/article/10.1007%2Fs00376-009-9095-2.

David B. Kirk and Wen-mei W. Hwu. *Programming Massively Parallel Processors*. Elsevier, 2010. ISBN 9780123814739. doi: 0123814731.

A.K. Kochanski, E.R. Pardyjak, W.J. Steenburgh, J.R. Stoll, and A. Gowardhan. One-way Coupling of the WRF-QUIC Urban Dispersion Modeling System. *Submitt. Rev. to J. Appl. Meteorol. Climatol.*, 2015.

Pijush K. Kundu and Ira M. Cohen. *Fluid Mechanics*. Academic Press, fourth edition, 2008. ISBN 9780123737359.

Victor W Lee, Changkyu Kim, Jatin Chhugani, Michael Deisher, Daehyun Kim, Anthony D Nguyen, Nadathur Satish, and Mikhail Smelyanskiy. Debunking the 100X GPU vs . CPU Myth : An Evaluation of Throughput Computing on CPU and GPU. In *Proceedings of the 37th annual international symposium on Computer architecture*, pages 451–460, 2010. ISBN 9781450300537.

M.O. Letzel, M. Krane, and S. Raasch. High resolution urban large-eddy simulation studiees from street canyon to neighbourhood scale. *Atmos. Environ.*, 42:8770–8784, 2008.

Katherine a. Lundquist, Fotini Katopodes Chow, and Julie K. Lundquist. An Immersed Boundary Method Enabling Large-Eddy Simulations of Flow over Complex Terrain in the WRF Model. *Mon. Weather Rev.*, 140(12):3936–3955, December 2012. ISSN 0027-0644. doi: 10.1175/MWR-D-11-00311.1. URL http://journals.ametsoc.org/doi/abs/10.1175/MWR-D-11-00311.1.

Alberto Martilli, Alain Clappier, and Mathias W Rotach. An urban surface exchange parameterisation for mesoscale models. *Boundary-Layer Meteorol.*, 104:261–304, 2002.

Valéry Masson. A physically-based scheme for the urban energy budget in atmospheric models. *Boundary-Layer Meteorol.*, 94(3):357–397, 2000.

M. Miyake. Transformation of the Atmospheric Boundary Layer over Inhomogeneous Surfaces. *Sci. Rep.*, 5R-6, 1965.

M.A. Nelson and Michael J. Brown. The QUIC Start Guide (v. 5.6). LA-UR-10-01-01062 (available from ¡http://lanl.gov¿)., 2006.

Marina Neophytou, Akshay Gowardhan, and Michael Brown. An inter-comparison of three urban wind models using Oklahoma City Joint Urban 2003 wind field measurements. *J. Wind Eng. Ind. Aerodyn.*, 99(4):357–368, April 2011. ISSN 01676105. doi:

10.1016/j.jweia.2011.01.010. URL http://linkinghub.elsevier.com/retrieve/pii/S0167610511000122.

Y Ogawa, P.G. Diosey, K. Uehara, and H. Ueda. A wind tunnel for studying the effects of thermal stratification in the atmosphere. *Atmos. Environ.*, 15(5):807–821, 1981.

M. Overby, P. Willemsen, D. Alexander, S. Halverson, and E.R. Pardyjak. Development of a rapid urban energy budget modeling system using a GPU ray sampling computational framework. *Submitt. Rev. to Urban Clim.*

Matthew Overby. *A High Performance Framework for Coupled Urban Microclimate Models*. Ms thesis, University of Minnesota Duluth, 2014.

H.A. Panofsky and A.A. Townsend. Change of Terrain Roughness and the Wind Profile. *Q. J. R. Meteorol. Soc.*, 90:147–155, 1964.

E.R. Pardyjak, S.O. Speckart, F. Yin, and J.M. Veranth. Near source deposition of vehicle generated fugitive dust on vegetation and buildings: Model development and theory. *Atmos. Environ.*, 42(26):6442–6452, August 2008. ISSN 13522310. doi: 10.1016/j.atmosenv.2008.04.024. URL http://linkinghub.elsevier.com/retrieve/pii/S1352231008004779.

Eric R Pardyjak and Michael J Brown. Evaluation of a Fast-Response Urban Wind Model-Comparison to Single-Guilding Wind Tunnel Data. In *Int. Soc. Environ. Hydraul. Conf.*, 2001.

Seung-Bu Park, Jong-Jin Baik, Siegfried Raasch, and Marcus Oliver Letzel. A Large-Eddy Simulation Study of Thermal Effects on Turbulent Flow and Dispersion in and above a Street Canyon. *J. Appl. Meteorol. Climatol.*, 51(5):829–841, May 2012. ISSN 1558-8424. doi: 10.1175/JAMC-D-11-0180.1. URL http://journals.ametsoc.org/doi/abs/10.1175/JAMC-D-11-0180.1.

W. Pendergrass and S.P.S. Aria. Dispersion in Neutral Boundary Layer over a Step Change in Surface Roughness-I. Mean Flow and Turbulence Structure. *Atmos. Environ.*, 18: 1267–1279, 1984.

Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, 2000. ISBN 9780521598866.

S. Raasch and M. Schröter. A large-eddy simulation model performing on massively parallel computers. *Meteor. Z.*, 10:363–372, 2001.

K.S. Rao, J.C. Wyngaard, and O.R. Cote. Local Advection of Momentum, Heat, and Moisture in Micrometeorology. *Boundary-Layer Meteorol.*, 7:331–348, 1974.

K. Richards, M. Schatzmann, and B. Leitl. Wind tunnel experiments modelling the thermal effects within the vicinity of a single block building with leeward wall heating. *J. Wind Eng. Ind. Aerodyn.*, 94(8):621–636, August 2006. ISSN 01676105. doi: 10.1016/j.jweia.2006.02.003. URL http://linkinghub.elsevier.com/retrieve/pii/S016761050600033X.

N.E. Rider, J.R. Philip, and E.F. Bradley. The horizontal transport of heat and moisture - a micrometeorological study. *Q. J. R. Meteorol. Soc.*, 89(382):507–531, 1963.

R. Röckle. *Bestimmung der Stromungsverhaltnisse im Bereich komplexer Bebauungsstrukturen.* Phd dissertation, der Technischen Hochschule Darmstadt, Germany, 1990.

M. Roth. Review of atmospheric turbulence over cities. *Quarterly Journal of the Royal Meteorological Society*, 126:941–990, 2000.

Sergiy a. Savelyev and Peter a. Taylor. Internal Boundary Layers: I. Height Formulae for Neutral and Diabatic Flows. *Boundary-Layer Meteorol.*, 115(1):1–25, April 2005. ISSN 0006-8314. doi: 10.1007/s10546-004-2122-z. URL http://link.springer.com/10.1007/s10546-004-2122-z.

Nick Shingleton. *COUPLING A LAND-SURFACE MODEL TO LARGE-EDDY SIMULATION TO STUDY THE NOCTURNAL BOUNDARY by.* PhD thesis, University of Utah, 2010.

B. Singh. *Testing and development of a fast response Lagrangian dispersion models.* Ms thesis, University of Utah, 2005.

B. Singh. *Development of a fast response dispersion model for virtual urban environments.* Phd dissertation, University of Utah, 2012.

Balwinder Singh, Bradley S. Hansen, Michael J. Brown, and Eric R. Pardyjak. Evaluation of the QUIC-URB fast response urban wind model for a cubical building array and wide building street canyon. *Environ. Fluid Mech.*, 8(4):281–312, August 2008. ISSN 1567-7419. doi: 10.1007/s10652-008-9084-5. URL http://link.springer.com/10.1007/s10652-008-9084-5.

Rob Stoll and Fernando Porté-Agel. Dynamic subgrid-scale models for momentum and scalar fluxes in large-eddy simulations of neutrally stratified atmospheric boundary layers over heterogeneous terrain. *Water Resour. Res.*, 42(1):n/a–n/a, January 2006. ISSN 00431397. doi: 10.1029/2005WR003989. URL http://doi.wiley.com/10.1029/2005WR003989.

Roland B. Stull. *An Introduction to Boundary Layer Meteorology*. Springer, 1988. ISBN 9789027727688.

Gianni Tinarelli, G. Brusasca, O. Oldrini, Domenico Anfossi, Silvia Trini Castelli, and J. Moussafir. Micro-Swift-Spray (MSS): A New Modelling System for the Simulation of Dispersion at Microscale. General Description and Validation. In Carlos Borrego and Ann-Lise Norman, editors, *Air Pollution Modeling and Its Application XVII*, pages pp 449–458. Springer US, 2007.

A. A. Townsend. The Responce of a Turbulent Boundary Layer to Abrupt Changes in Surface Conditions. *J. Fluid Mech.*, 22:799–822, 1965.

K. Uehara, S. Murakami, and S. Oikawa. Wind tunnel test on stratified flow fields around urban street canyons by LVD. *J. Archit. Plan. Environ. Eng.*, 492:39–46, 1997.

Kiyoshi Uehara, Shuzo Murakami, Susumu Oikawa, and Shinji Wakamatsu. Wind tunnel experiments on how thermal stratification affects flow in and above urban street canyons. *Atmos. Environ.*, 34:1553–1562, 2000.

Henek K. Versteeg and Weeratunge Malalasekera. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method*. Limited, Pearson Education, 2007. ISBN 9780131274983.

Z. Yang, Y. Zhu, and Y. Pu. Parallel image processing based on cuda. In *Int. Conf. Comput. Sci. Softw. Eng.*, pages 198–201. IEEE, 2008.