https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Implementation of a resolved building capability in the Energy Research and Forecasting model for accelerated high-fidelity urban-scale simulation

Adam S. Wise[1], Harish Gopalan[2], Aaron M. Lattanzi[3], Ann S. Almgren[3], David Wiersema[1], Akshay Gowardhan[1], Soonpil Kang[1], Jeffrey D. Mirocha[1], David J. Gardner[1], and Katherine A. Lundquist[1]

[1]Lawrence Livermore National Laboratory, Livermore, CA, USA
[2]National Laboratory of the Rockies, Golden, CO, USA
[3]Lawrence Berkeley National Laboratory, Berkeley, CA, USA

Correspondence: Adam S. Wise (wise14@llnl.gov)

Abstract. Accurate predictions for wind and particle dispersion must be obtained quickly to inform emergency response to atmospheric releases of hazardous material in urban areas. In this study, we demonstrate and validate a newly implemented resolved building capability added to the exascale Energy Research and Forecasting (ERF) model. To represent buildings, an immersed forcing method is employed, which enforces near-zero velocities within buildings. We modify the immersed forcing method to implement a wall-model where we enforce a calculated velocity for building walls determined using the log law. The immersed forcing method is evaluated by simulating cases from the Joint Urban 2003 Oklahoma City tracer release experiment, analyzing cases with both stable and unstable thermodynamic conditions. For the stable case, we use adaptive mesh refinement for regions above a certain scalar concentration. Additionally, we conduct two sensitivity studies: one on building wall roughness length, finding that 0.1 m results in the most accurate scalar concentration predictions, and another on model configurations, where computational cost decreases by an order of magnitude when either not refining the grid or when using anelastic mode compared to fully compressible mode, with model skill only decreasing slightly. For the convective case, model skill is not as high as the stable case largely due to underprediction of scalar concentration at stations to the west and upwind of the release site. Ultimately, model skill for both cases are still among the highest in the published literature, demonstrating the efficacy of our method for accelerated high-fidelity urban atmospheric modeling applications.

1 Introduction

Large populations living in urban areas can be exposed to high risk or airborne contaminants. Accurate simulation is required to provide guidance to officials to ensure public safety; however, existing high-fidelity urban modeling tools are too slow to be used for time-sensitive applications. This is because buildings have highly complex geometries and properly resolving these structures requires very small grid spacing, typically on the scale of a few meters, which is challenging using large-eddy simulation (LES) codes developed to run on traditional high-performance computing (HPC) architectures. Therefore, given recent technological advancements in high-performance computing (exascale architectures), we are developing accelerated,

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

high-fidelity simulation tools for time-sensitive applications. Moreover, these tools enable running ensembles of simulations and performing sensitivity studies to provide more robust guidance and to inform lower-order faster running models.

In recent years, there has been momentum from the atmospheric science community to develop LES codes that take advantage of exascale systems (Sauer and Muñoz-Esparza, 2020; Norman et al., 2025). Codes that leverage graphics processing units (GPUs), accelerated processing units (APUs), or other exascale architectures can often achieve much shorter computational runtimes for fixed problem sizes than those that only use central processing units (CPUs). The Energy Research and Forecasting (ERF) model (Almgren et al., 2023; Lattanzi et al., 2025) is one such code and is built on the AMReX framework (Zhang et al., 2019, 2021) for massively parallel block-structured adaptive mesh refinement (AMR) applications. The AMR capabilities of ERF enable further performance gains by simulating only certain, necessary regions with higher resolution while simultaneously simulating other regions at coarser grid spacing. In the present work, we extend ERF to represent built infrastructure for urban simulations.

There exist three main approaches for representing buildings in numerical codes. The first is to represent buildings using solid boundaries. When using solid boundaries, grid refinement and/or unstructured meshes are typically employed near building boundaries to resolve flow features. The grids are generally conformal to the building geometry, allowing no-slip boundary conditions to be applied (Gorlé et al., 2010; García-Sánchez et al., 2018). Within the AMReX framework, there is also support for a cut-cell method, also known as an embedded boundary (EB), where cells that are intersected by the boundary are cut to conform to the boundary (Zhang et al., 2021; Natarajan et al., 2022; Kang et al., 2025). A cut-cell method computes the flux through the cut surface to impose the desired boundary conditions (no-slip, slip-wall, etc.). A solid boundary approach is more common in the computational fluid dynamics community as opposed to the atmospheric science community as atmospheric models typically have more complex surface boundary conditions or use pressure-based vertical coordinates making conformal grids with steep slopes (e.g., buildings) unachievable. Representing urban environments using an exact method such as solid boundaries is generally superfluous as there are a number of subgrid-scale factors that need to be parameterized, and for which limited data exist.

The second approach, more common in the atmospheric science community, is to use an immersed boundary method (IBM) (Peskin, 1972) that reconstructs a boundary condition at the location where the building intersects the cell (Senocak et al., 2004; Letzel et al., 2008; Lundquist et al., 2010, 2012; Anderson, 2013; Bao et al., 2016; Jähn et al., 2016; Wiersema et al., 2020). These methods generally work by computing source terms near the building boundary to enforce an exact flux or dirichlet condition. A comprehensive review on immersed boundary methods can be found in Iaccarino and Verzicco (2003) and Mittal and Iaccarino (2005). Immersed boundary methods are advantageous over solid boundaries as immersed boundaries do not depend on the specific geometry of interest; however, their implementation is ultimately complex with an interpolation scheme that can add some computational cost.

A class of methods that have gained traction lately for accelerated urban modeling are immersed forcing methods (sometimes called immersed body force methods) that follow, in principle, many of the approaches of an immersed boundary method but with a more straightforward and less computationally expensive implementation. These methods act like a canopy model, where a large drag force is applied to the momentum equations for cells immersed as buildings. By using a large enough drag

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

force, the forcing term acts similarly to a solid obstacle with near-zero velocities for immersed cells. This method was first investigated by Chan and Leach (2007) who demonstrated that for larger drag terms, the flow around a cube acted similarly to when modeling the cube as a solid boundary. Each cell in the domain is masked based on how immersed the cell is, which is then used in the immersed forcing formulation. As a result, the method is applied to all grid cells, allowing for efficient parallelization on accelerated hardware. The downside of immersed forcing methods is that the boundary condition is smeared over a computational cell and thus they are less accurate than the previously described methods. However, this loss in accuracy can be offset, or even outweighed, in an accelerated modeling framework by larger gains in performance that enable improvements to other aspects of the modeling setup (such as forcing, grid spacing, etc.), as explored in this study. Muñoz-Esparza et al. (2020) were the first to use this method in a GPU-accelerated framework, recognizing the efficacy of the method. This work built on the method of Chan and Leach (2007) by evolving the method to include a scale-independent drag coefficient and also introducing temperature and density source terms to force building temperatures to a specified value. Recently, Norman et al. (2025) introduced a GPU-accelerated large-eddy simulation (LES) tool with an immersed forcing capability. They modified the LES closure and surface friction schemes so that non-immersed cells adjacent to immersed cells experience friction and thermal forcing, effectively acting as a wall-model.

In the present work, we use the method of Chan and Leach (2007) and Muñoz-Esparza et al. (2020) for fully immersed cells, but extend the method to include a new wall-model formulation for building wall surface roughness based on the log law, which is accomplished by adding a forcing term to the momentum equation in computational cells that represent walls. This wall-modeling method is described in further detail in Section 2. We validate our wall-modeling implementation with different building wall roughness lengths based on data from two intensive observational periods during the Joint Urban 2003 experiment in Oklahoma City (Allwine and Flaherty, 2006) as discussed in sections 3, 4, and 5). Additionally, in this validation, we leverage AMR-based plume refinement to demonstrate dynamic refinement over a region of interest. Lastly, also in section 5, we compare different model configurations (e.g., compressible or anelastic mode) to explore accuracy and performance tradeoffs.

2 Methods

2.1 The Energy Research and Forecasting model

The Energy Research and Forecasting (ERF) code is a limited-area atmospheric modeling code that leverages accelerated exascale architectures (i.e., GPUs, APUs, etc.). ERF contains many of the standard discretizations and basic physics needed to model general atmospheric dynamics for applications such as extreme weather events, renewable energy modeling, and urban climates. The atmospheric physics parameterizations coupled to the ERF dynamical core include cloud microphysics, land surface models, radiation, and turbulence models for meso- and microscale resolutions. The code-base is open-source and new features are continually being developed, expanding its use to additional applications.

ERF is built on the modern and well-supported software framework AMReX. This framework uses modular and flexible design, enabling use of a variety of physics parameterizations and computational algorithms. Additionally, AMReX allows for both static and dynamic grid refinement. Static grid refinement refers to a time-invariant refinement capability, similar to

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

grid nesting in the Weather Research and Forecasting (WRF) model (Skamarock et al., 2021), whereas dynamic refinement allows the refinement regions to change over the course of a simulation. As in WRF, ERF uses an Arakawa-C grid for spatial discretization, with scalar fields at cell-centers and velocity fields at cell faces. ERF also allows for use of terrain-following coordinates and/or grid stretching.

ERF can be run in one of two modes, solving either the fully compressible equations or the anelastic approximation, which removes sound waves and assumes that density perturbations are small in comparison to the background state. For simulations using the compressible solver, the temporal discretization is the same third-order Runge-Kutta (RK3) scheme with acoustic substepping as in the WRF model. For anelastic simulations, a second-order Runge-Kutta (RK2) scheme is used with no substepping. Fully compressible simulations are often thought of as necessary for deep-atmosphere simulations; however, there are a number of applications targeted by ERF where anelastic mode may be appropriate, such as idealized atmospheric boundary layer simulations, which are the subject of the present study. The tradeoff from a performance perspective is that fully compressible mode is nearly fully explicit (the perturbational equations in the vertical are implicit), which is computationally cheaper but requires smaller timesteps to account for sound waves. For anelastic mode, larger timesteps can be taken as the speed of sound does not need to be accounted for, but each timestep is more expensive because of the global linear (Poisson) solve for the perturbational pressure. The anelastic method used here employs a fast Fourier transform based solver for the perturbational pressure (but a multigrid solver is also available). Ultimately, using anelastic mode is less expensive and up to an order of magnitude faster than fully compressible mode for the type of simulations in the present study, which is discussed in the following paragraph and later on this manuscript. In this study, we will explore both the performance and accuracy tradeoffs when using the fully compressible and anelastic modes for urban simulations.

To demonstrate the performance of ERF, we run a simple scaling test similar to that conducted by Sauer and Muñoz-Esparza (2020). For this test, we use ERF v26.05, and run a doubly periodic convective boundary layer with isotropic grid resolution of 5 m (as opposed to the 20 m spacing used in Sauer and Muñoz-Esparza (2020)) and a geostrophic wind speed of 8 m s^{-1}. The number of vertical points is held constant at 122 but the number of horizontal grid points is increased. Buildings are also included in the simulations (we include a subset of buildings from Torino, Italy) to ensure that the performance demonstrated is what a user should expect when running building-resolving simulations in ERF. The simulations are run for 300 s and for the fully compressible simulations, a main timestep of approximately 0.25 s and a substep timestep of approximately 0.006 s (these timesteps were confirmed to be stable after running 3600 s simulations). For the anelastic simulations, the timestep is approximately 0.33 s. Note that the timesteps used are adaptive and scaled based on the acoustic Courant–Friedrichs–Lewy (CFL) condition, which is why the aforementioned values are approximate, but these values are relatively steady and do not vary significantly other than during the startup of the simulation. This timestepping configuration for the fully compressible simulations results in approximately 30 substeps per main step and was found to be optimal in terms of performance for the fully compressible simulations. A relatively large number of substeps per main step is feasible for this problem because advection and diffusion, which are evaluated on the main step, do not change much over the timescale of the substeps.

Figure 1 shows the performance benchmarking results for the scaling test using a varying number of GPUs. Results with a simulation time to wall-clock time ratio greater than one indicate real time performance or better. For fully compressible

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

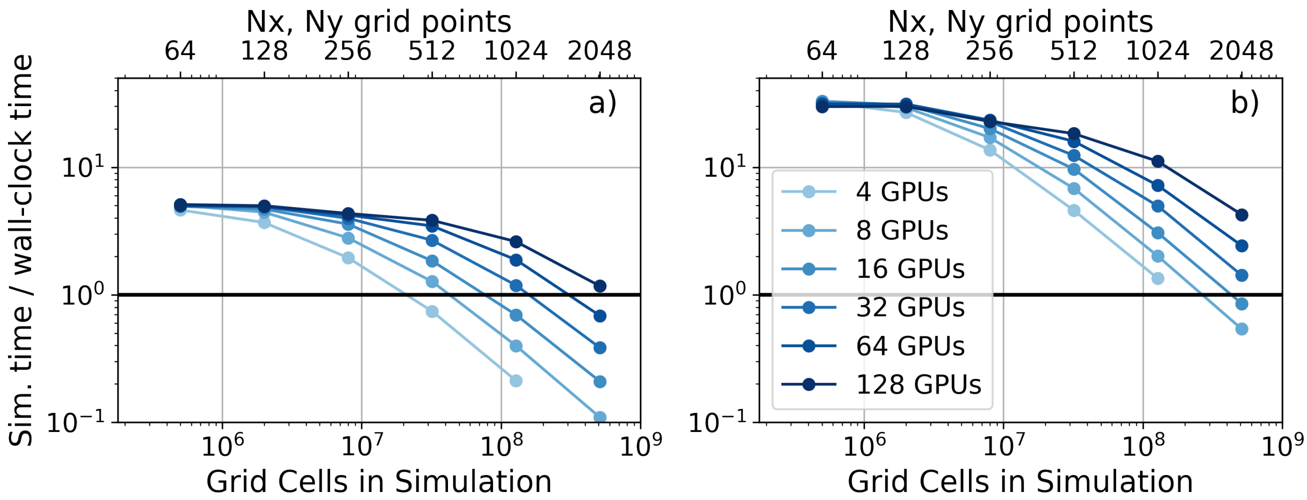

Figure 1. ERF simulation time over the elapsed real time as a function of the number of grid cells in the simulations and number of GPUs used for (a) fully compressible and (b) anelastic simulations. Values greater than 1 represent simulations that are faster than real-time. Simulations were conducted on AMD AMD MI300A accelerated processing units on LLNL's Tuolumne HPC machine.

simulations, this threshold for a 5120 x 5120 x 610 m (Nx, Ny, Nz = 1024 x 1024 x 122) domain size is achieved when using 32 GPUs. For anelastic simulations, real-time simulations for the same domain size are achieved with just 4 GPUs. Even for a domain four times larger in area, nearly 10 km x 10 km, real-time simulations can still be achieved with just 16 GPUs. The increase in performance for using anelastic mode compared to fully compressible mode will vary and need to be optimized from problem to problem; however, significant performance gains should be expected. In later sections, we explore the performance versus accuracy tradeoffs of the different modes.

2.2 Immersed Forcing

We denote the method used to represent resolved buildings in the present study as the immersed forcing method. In this method, a large body force or sink is applied at velocity points to the momentum equations to force the velocity to near zero or to a desired value. This method is used over directly prescribing velocities to be exactly zero within buildings as that introduces sharp, unphysical gradients that can cause numerical instability. This method was first implemented by Chan and Leach (2007) and was recently modified and improved by Muñoz-Esparza et al. (2020). Muñoz-Esparza et al. (2020) made two significant improvements: 1) the forcing term is a function of the grid spacing to be scale independent and 2) the forcing term accounts for the volume fraction of the cell V_f that is immersed as a building. The volume fraction takes a value of 1 for a cell fully immersed in a building (solid cells), 0 for a fluid cell, and $0.0 < V_f < 1.0$ for cells partially immersed in a building and representing a wall (boundary cells). Additionally, Muñoz-Esparza et al. (2020) also implemented thermal forcing for the potential temperature where they force immersed cells to a reference value. In the present work, no thermal forcing is applied; however, thermal boundary conditions for building walls are a subject of ongoing work.

The advantage of using a forcing method (as opposed to a method that sets an explicit boundary) is that its implementation is straightforward and it parallelizes extremely efficiently. By not setting an explicit boundary, the building boundary when using this forcing method is essentially smeared over a grid cell, which is less accurate than an explicit boundary and is a limitation of the method. However, this approach is especially well-suited for codes running on exascale architectures where very small grid spacing is achievable with relatively large domains and compensates for accuracy loss from smearing. In Shin et al. (2021), the authors tested this method in a variety of configurations, including different atmospheric stability conditions, various grid spacings, and advection schemes, and found that the method was robust across a range of configurations. They also tested the number of grid cells required to resolve a building, and found that buildings were well represented with just 6 grid points per building wall with their building being shaped as a cube. The method has been applied more recently by Muñoz-Esparza et al. (2021), Welch et al. (2024), and Norman et al. (2025), demonstrating its versatility in a range of urban modeling applications. Additionally, because of the scale independent formulation of immersed forcing, the method can be applied at different grid spacings making it well suited for both static and dynamic mesh refinement with no special treatment required at grid boundaries.

In this work, we extend the forcing method of Muñoz-Esparza et al. (2020) to include a wall-model appropriate for use with buildings. Generally, this modified immersed forcing method introduces source terms to drive the fluid velocity towards a target velocity,

$$F_{u_i} = -C_d\,\rho\,|U_s|\,(u_i - u_{i,\text{target}}), \tag{1}$$

where u_i is the velocity component in $i,j,k = 1,2,3 = x,y,z$, ρ is the air density, and C_d is a drag coefficient formulated as $\alpha_m\left(\sqrt[3]{\Delta x_1 \Delta x_2 \Delta x_3}\right)^{-1}$ with α_m being a tuning parameter and set to a value of 500 and Δx_i the grid spacing in each dimension. The goal of the forcing term is ultimately to force the velocity at a given velocity points u_i to a target velocity u_{target}. Note that the value α_m is smaller than the value of 1000 determined by Muñoz-Esparza et al. (2020), as during testing using larger values resulted in numerical instability and using a value of 500 similarly achieved near-zero velocities within buildings. The unit velocity scale U_s keeps the units consistent and follows the form of the scalar forcing used in Muñoz-Esparza et al. (2020).

In applications for urban environments, the target velocity is not aware of the near-wall shear layer. Thus, we extend the method to determine u_{target} from one of two ways. For solid cells (V_f = 1.0), the target velocity u_{target} is 0.0 m s^{-1}. For boundary cells ($0.0 < V_f < 1.0$), a building wall is represented at the velocity point by using the log law to determine the target velocity as

$$u_{i,\text{log law}} = \frac{u_{*,i+1}}{\kappa}\left[\log\left(\frac{0.5\Delta x_i}{z_{0,\text{w}}}\right)\right], \tag{2}$$

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

where κ is the von Karman constant (taken to be 0.41) and $z_{0,\mathrm{w}}$ is the building wall surface roughness. The friction velocity u_* is calculated using fully fluid cells adjacent to boundary cells since we assume that the log law is being applied in the constant flux layer,

$$u_{*,i+1} = U_{i+1,\mathrm{tangent}} \frac{\kappa}{\log\left[\frac{1.5\Delta x_i}{z_{0,\mathrm{w}}}\right]}, \tag{3}$$

where U_{tangent} is the tangential wind speed. All building walls are assumed to be coordinate-aligned; therefore, the tangential velocity for a roof is $\sqrt{u_1^2+u_2^2}$, the tangential velocity for an east- or west-facing wall is $\sqrt{u_2^2+u_3^2}$, and, lastly, the tangential velocity for a north- or south-facing wall is $\sqrt{u_1^2+u_3^2}$.

Finally, we determine the target velocity $u_{i,\mathrm{target}}$ as a function of the volume fraction,

$$u_{i,\mathrm{target}} = (1-V_f)\, u_{i,\mathrm{log\ law}}. \tag{4}$$

This form has the effect of forcing velocities towards the value calculated from the log law for lower volume fractions (less immersed boundary cells) and towards 0.0 m s^{-1} for higher volume fractions (more immersed boundary cells).

2.3 Verification of the Immersed Forcing Method

To verify the immersed forcing method, we create a similar setup to the $\Delta = 5$ m modeling setup from Shin et al. (2021). In this set up, an isolated cube of height $H = 120$ m is located $20H$ from the intlet and $10H$ from the outlet. The width of the domain is $12H$ with slip wall boundary conditions used at the spanwise boundaries. The inflow profile is from a fully developed doubly periodic neutral boundary layer with a geostrophic wind speed of 10 m s^{-1} and surface roughness $z_{0,\mathrm{s}} = 0.1$ m (see Fig. 2 in Sauer and Muñoz-Esparza (2020)); however, we fully align the wind direction with the x-axis to remove asymmetries introduced by wind veer. There is a capping inversion at 500 m to contain turbulence within the boundary layer. The time step used in ERF is 0.125 s with a substepping CFL number of 0.5 used to determine the substepping timestep (which tended to be approximately 0.007 s). We use the Smagorinksy LES closure (Smagorinsky, 1963) with a coefficient of 0.17 and use the cell perturbation method (CPM) (Muñoz-Esparza et al., 2014, 2015) at the inlet to accelerate the development of realistic turbulence. The simulations are run for 900 s to spin up the flow field and then run for an additional 3600 s for analyzing time-averaged outputs.

To verify the wall-model component of the immersed forcing method, we intentionally increase the size of the cube by 0.25 m to ensure that the edge of the cube is located in a partially immersed cell. We vary the wall boundary condition by modifying the building wall roughness length $z_{0,\mathrm{w}}$. We also run a reference simulation where the wall cell is forced to 0 m s^{-1} representing a no slip boundary condition. Figure 2 shows time-averaged vertical profiles along the centerline of the domain and horizontal profiles at 7.5 m above the ground of the x velocity u. The largest $z_{0,\mathrm{w}}$ value of 2.49 m is chosen because, mathematically, as the building wall roughness length approaches $\Delta z/2$, the target velocity approaches 0 m s^{-1} (see Equation 2). Accordingly,

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

in Fig. 2, the profiles for the no slip and $z_{0,\mathrm{w}} = 2.49$ m cases are nearly identical. Additionally, as the building wall roughness length decreases, the top recirculation zone disappears as the amount of drag introduced by the boundary is reduced. The rear recirculation zone also has differing characteristics with the negative velocities persisting further distances for larger building wall roughness lengths. Overall, the wake from the cube extends to higher heights and further spanwise distances the larger the building wall roughness length is.

The results presented in Fig. 2 demonstrate the effect of different building surfaces and how these characteristics can modify the flow surrounding a building. These characteristics will vary from building to building and the method does not aim to exactly reproduce flow from the field. The immersed forcing method is not an exact boundary; however, the extension of the method to include a log law wall-model provides a physically grounded parameterization for representing rough building surfaces while remaining computationally efficient with the method also working seamless across scales and grids when using adaptive mesh refinement (i.e., no special treatment is required when using refinement). For a validation of the method at wind tunnel scales, see Appendix A.

3 Joint Urban 2003 Atmospheric Dispersion Study

The Joint Urban 2003 atmospheric dispersion study was a field campaign jointly conducted by the U.S. Department of Defense-Defense Threat Reduction Agency (DTRA) and the U.S. Department of Homeland Security (DHS). Planning and execution of the study included scientists from universities, government laboratories, and private industry. The field campaign focused on collecting high-resolution data to better understand dispersion and flows around tall buildings and through street canyons. Observations were collected in Oklahoma City in late June through July 2003. More information on the experiment can be found in Allwine and Flaherty (2006).

There were ten intensive operational periods (IOPs), each 8 hours in duration during the 34 day span of the study. During these IOPs, sulfur hexaflouride (SF_6) was released as a tracer gas in a puff or continuous release. The release location varied depending on the IOP, but the locations were all within OKC's central business district (CBD) (see Figures 3 and 5 for the tracer release sites for IOPs analyzed in this study). To measure the tracer gas, there were 19 Lawrence Livermore National Laboratory (LLNL) "bluebox" integrating gas samplers and over 120 National Oceanographic and Atmospheric Administration (NOAA) Air Resources Laboratory Field Research Division (FRD) programmable integrating gas samples (PIGS). The majority of the PIGS devices were located approximately 3 m above the ground but some devices were located on top of buildings; however, in this study, only stations at street level are used for analysis. Blue box samplers were also located at various heights, but only samplers at street level (1 m above the ground) are used in this study. The bluebox devices were located mainly in the central district while the PIGS devices were located both in the central business district and in 1, 2, and 4 km arcs from the CBD.

In addition to all of the instrumentation for the tracer gas, there were a number of other meteorological instruments present during the study to measure atmospheric data. Relevant to this study are 20 Dugway Proving Ground (DPG) super Portable Weather Information Display Systems (super PWIDS) with three-dimensional sonic anemometers and temperature sensors. The two other instrument types used to characterize inflow and initial conditions were mainly instruments that provided information

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

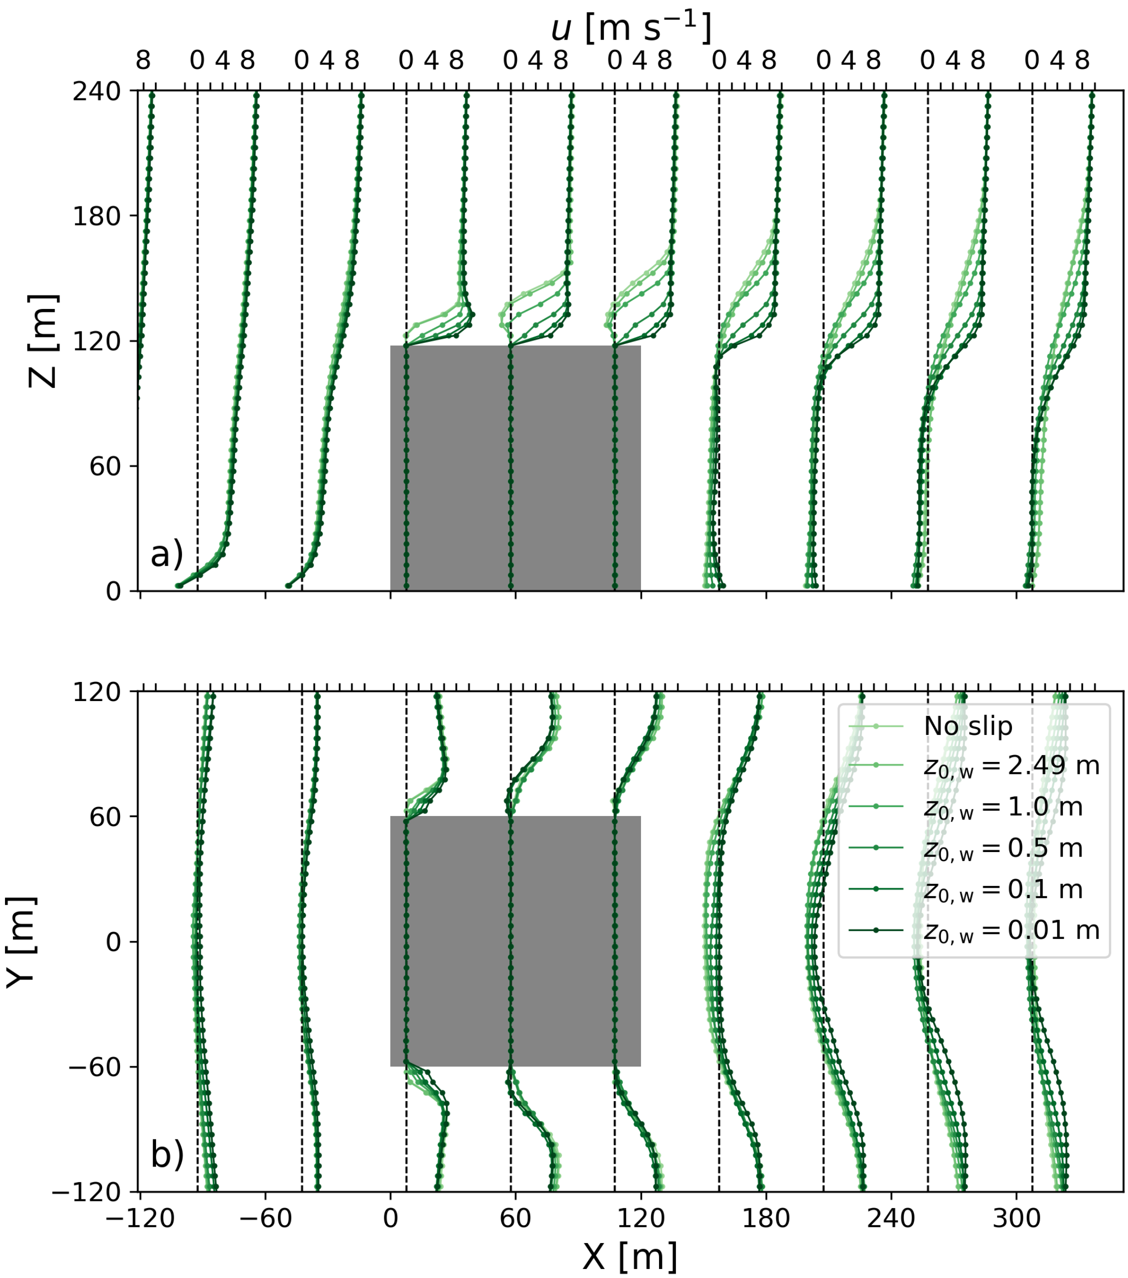

Figure 2. 1 hr time-averaged (a) vertical profiles along the centerline and (b) horizontal profiles at 7.5 m above the ground of the x velocity with different boundary conditions and building wall roughness lengths.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

about the vertical structure of the atmosphere. The first was a crane instrumented by LLNL with 3D sonic anemometers and temperature sensors located roughly every 7-10 m in height up to 83.2 m. The other instruments used to characterize the vertical structure were soundings launched by Argonne National Laboratory (ANL) and the Pacific Northwest National Laboratory (PNNL). These soundings did not measure wind speed but did measure the potential temperature, which was important for understanding the atmospheric stability during each IOP. Both of the radiosonde launch locations are outside of our simulation domains but are still useful for understanding the general behavior of the atmosphere (the ANL radiosondes were launched roughly 4 km north of the CBD and the PNNL radiosondes were launched roughly 1.5 km southwest of the CBD). The location of all of the aforementioned measurement devices (other than the sounding launch sites) can be seen in Figure 3.

4 Modeling Setup

The IOPs investigated as part of this study include IOP9, which was a nighttime IOP with typical stably stratified conditions and IOP3, which is a daytime IOP with typical convective or unstably stratified conditions. Both of these IOPs have been analyzed rigorously in the literature (Muñoz-Esparza et al., 2020; Chan and Leach, 2007; Neophytou et al., 2011; Lundquist et al., 2012; García-Sánchez et al., 2018; Wiersema et al., 2020, 2022). We use two distinct modeling setups depending on the atmospheric conditions, as shown in Figures 3 and 5. The different setups are used due to the different scales of turbulence in stable and convective conditions with smaller scale $O(50\text{m})$ structures dominant in stable conditions while much larger $O(1\text{km})$ structures are important during convective conditions. Additionally, other factors influencing the design of the domains are the different wind directions for the two cases (southerly for the stable case and more southwesterly for the convective case) and the different tracer release locations (Park Avenue for the stable case and the Botanical Garden for the convective case).

4.1 Stable case set up

For the stable case study, we use an inflow/outflow setup with periodic spanwise boundary conditions. As the flow is southerly, the southern boundary is treated as an inflow plane with specified profiles for u, v, and θ based on observations (shown in Figure 4); the northern boundary is treated as an outflow plane. The specified wind direction is a constant 171 degrees with height. García-Sánchez et al. (2017) demonstrated that in dispersion studies, the largest source of uncertainty tends to be with the wind direction. Both García-Sánchez et al. (2018) and Muñoz-Esparza et al. (2020) similarly used a constant wind direction with height but chose a value of 172 degrees. Our value of 171 degrees differs slightly and is based on running a number of preliminary simulations and comparing the tracer gas plume behavior in simulations and observations. For the bottom boundary condition (just at street level), we use Monin-Obukhov similarity theory (MOST) (Monin and Obukhov, 1954) with a surface roughness $z_{0,\text{s}} = 0.29$ m, following Muñoz-Esparza et al. (2020) and García-Sánchez et al. (2018). The initial surface temperature is 306 K and the surface is cooled at a rate of 1 K h^{-1}, consistent with measurements from the super PWIDS stations and measurements on the crane. To decrease the fetch distance required to generate turbulence from the smooth inflow profile, we use CPM, which has been shown to be effective at spurring realistic turbulence consistent with atmospheric forcing conditions (e.g., Muñoz-Esparza and Kosović (2018), Arthur et al. (2020), Wiersema et al. (2022), and

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Wise et al. (2025)). We also use a short 50 m sponge layer to dampen mechanical turbulence at the outflow. The domain extent is $L_x \times L_y \times L_z = 1750\ \mathrm{m} \times 4000\ \mathrm{m} \times 300\ \mathrm{m}$ with uniform grid spacing $\Delta = 5$ m, resulting in $N_x, N_y, N_z = 350, 800, 60$ cells, respectively. The spinup time is 900 s (similar to the 720 s used by Muñoz-Esparza et al. (2020)) with the simulation then run for an additional 1800 s to provide a 30 minute analysis period corresponding to 0800-0830 UTC on 27 July 2003. During the 1800 s period, the SF_6 passive tracer is continuously released at a rate of 5 g s^{-1}. The domain along with a more detailed view of the central business district is shown in Figure 3.

4.2 Convective case set up

For the convective case, we leverage a two-domain nested setup with one-way coupling in an attempt to capture the large-scale turbulent structures characteristic of convective boundary layers. An outer domain with doubly periodic boundary lateral boundary conditions is used to spin up convectively generated turbulence. After this initial spin up period, an inner domain with buildings is then turned on. The one-way nature of the nesting traps any mechanically generated turbulence from the buildings in the inner domain, thus it does not recirculate as inflow. The extent of the outer domain is $L_x \times L_y \times L_z = 8000\ \mathrm{m} \times 8000\ \mathrm{m} \times 1600$ m and the extent of the inner domain is $L_x \times L_y \times L_z = 2750\ \mathrm{m} \times 4000\ \mathrm{m} \times 300$ m as seen in Figure 5. The extent of the inner domain encompasses the vertical reach of the tallest buildings as well as a number of tracer sensors in the 2 km arc. The horizontal grid spacing is 10 m for the outer domain and 5 m for the inner domain. The vertical grid spacing is stretched with a near-surface grid spacing of 10 m for the outer domain and then stretched at a rate of 1.025. The vertical grid spacing for the inner domain is based on the outer domain and refined by a factor of 2 resulting in a near-surface grid spacing of 5 m. The bottom boundary condition also uses MOST with $z_{0,\mathrm{s}} = 0.29$ m but with a specified heat flux of 0.34 K m s^{-1}. Although some studies have considered this time period to be neutral (Lundquist and Chan, 2007; Lundquist et al., 2012; Wiersema et al., 2020), Wiersema et al. (2022) found better agreement between their model and the observations when a heat flux of 407 W m^{-2} (0.34 K m s^{-1}) was prescribed at the surface compared to having no heat flux. Further investigation of the observations demonstrated convective conditions with the potential temperature profile from the sounding launched at 1607 UTC showing a super-adiabatic layer closer to the surface (Figure 6). The geostrophic wind speed is 8 m s^{-1} with a uniform wind direction of 209 degrees specified. Turbulence is spun up on the outer domain for 6900 s (allowing the resolved heat flux 100 m above the ground to become pseudo-steady) and with the inner domain then turned on and also allowed to spin up for an additional 300 s, which is long enough for the flow to traverse the majority of the domain. The simulation is then run for an additional 1800 s (7200-9000 s) as the 30 minute analysis period corresponding to 1600-1630 UTC on 7 July 2003. During the 1800 s period, the SF_6 passive tracer is continuously released at a rate of 2.1 g s^{-1}. An average wind speed and potential temperature profile during the analysis period at the tracer release site is shown in Figure 6.

4.3 Adaptive mesh refinement

For the stable case study, we also explore using adaptive mesh refinement (AMR) around the SF_6 plume. For this domain, we use two-way nesting, allowing for a more refined grid based on the concentration of SF_6. We use a refinement ratio of 2 thus the AMR grid has a uniform grid spacing of 2.5 m in all three dimensions. We tag cells for refinement with concentrations greater

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

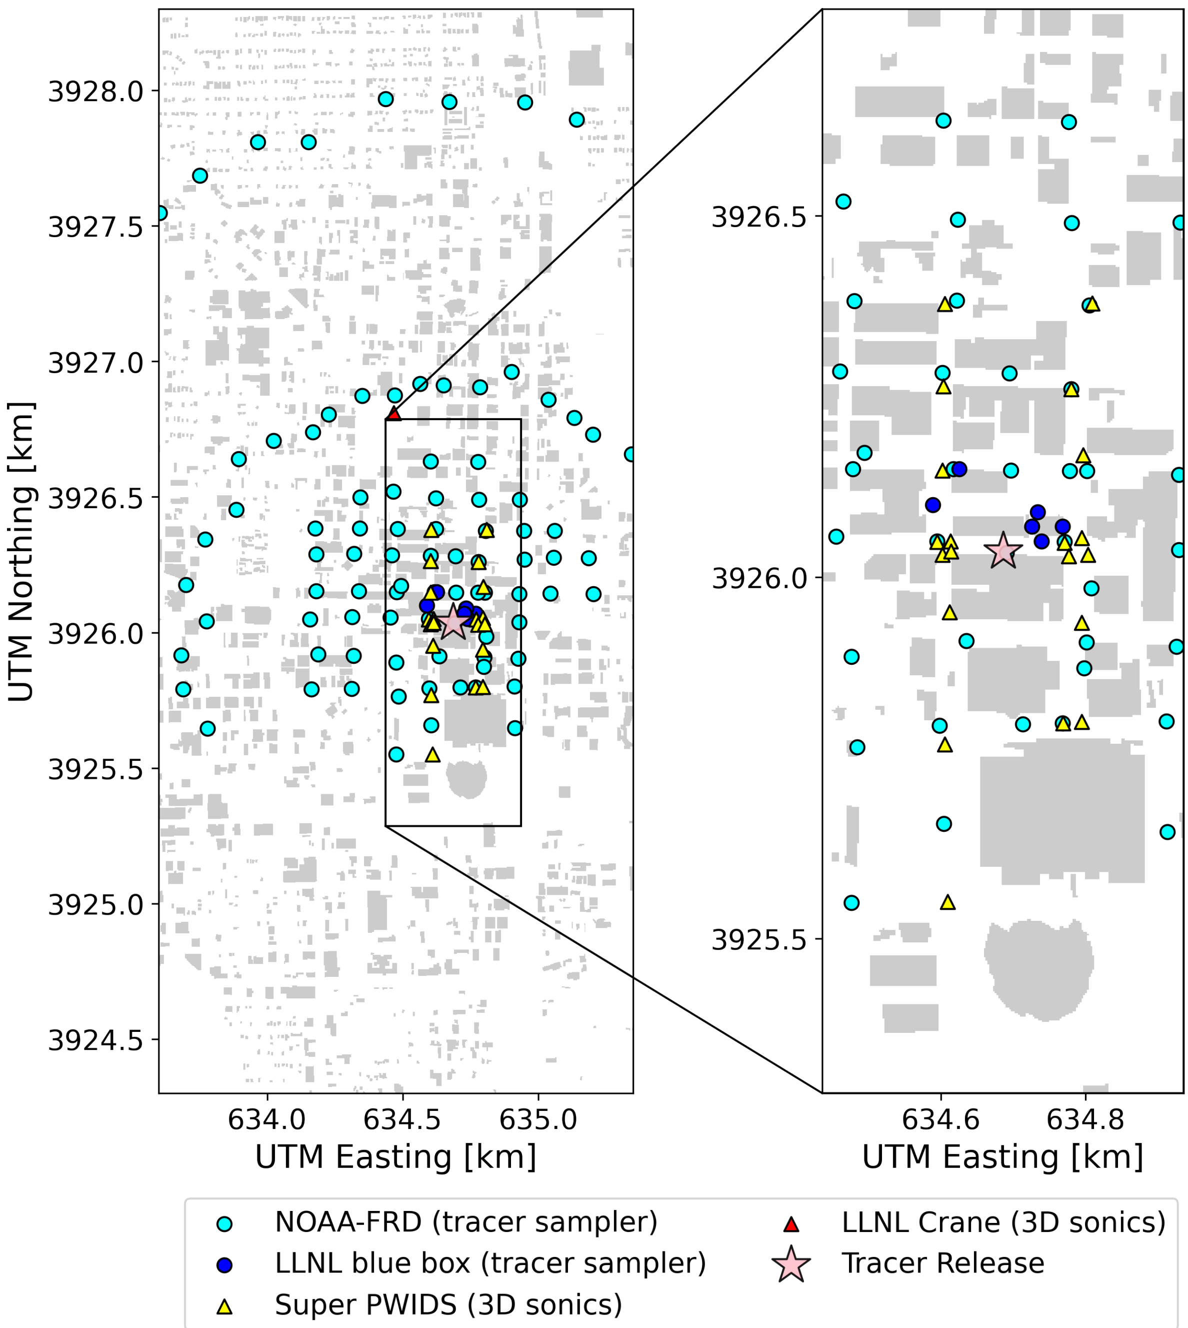

Figure 3. Simulation domain extent for the stable case study with various measurement locations overlaid as well as the tracer release site (Park Avenue). A detailed view of the central business district is provided for clarity, and also represents the initial region refined when using adaptive mesh refinement.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

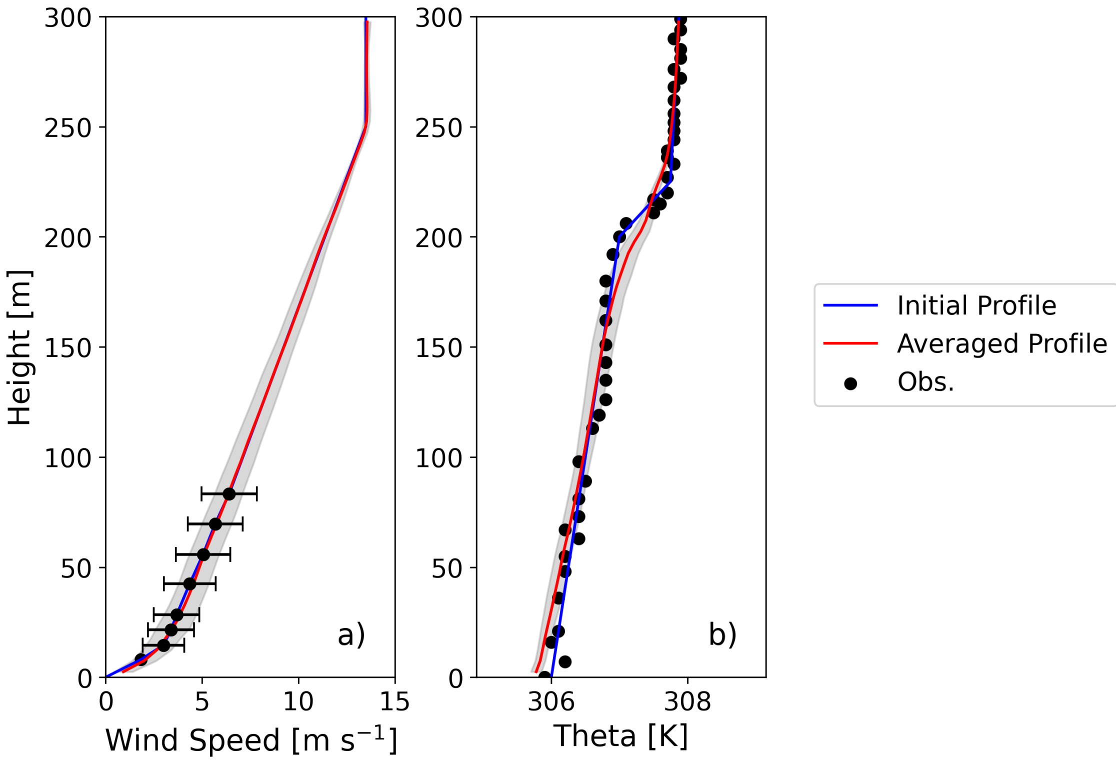

Figure 4. Vertical profiles of (a) wind speed and (b) potential temperature during the stable case study. The profiles include the initial profiles at the beginning of the simulation period and the 30-minute averaged profiles during the analysis period, which are a spanwise average 1 km from the inflow boundary. The wind speed observations are 30-minute averages at the LLNL crane site while potential temperature observations are instantaneous (launched at 04:07 UTC) from the PNNL radiosonde located outside of and upwind of the simulated domain. The error bars represent the standard deviation during the analysis period at each height.

than 5 parts per billion volume (ppbv) and update the grid every 100 time steps (approximately every 1.25 s). We also buffer the tagging regions in the spanwise (x) dimension by 250 m and in the streamwise (y) dimension by 750 m. These buffers serve two purposes. The first is that they tag additional cells for refinement that do not necessarily contain SF_6 but can affect the SF_6 plume behavior. By including the buffer, buildings near the plume are more finely resolved thus impacting the behavior of the flow and passive tracer. This buffer is especially important in the upwind direction. The other purpose is that it makes for more efficient parallelization of the domain, which is important for performance on exascale architectures. We also specify a relatively low grid efficiency for the tagging algorithm and this combined with the relatively large buffer region result in fewer and larger refinement boxes in the AMR algorithm, which more effectively utilize the GPU given their high throughput. The initially refined region using AMR can be seen in Figure 3. How the region evolves over the course of the analysis period is discussed in the following section.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

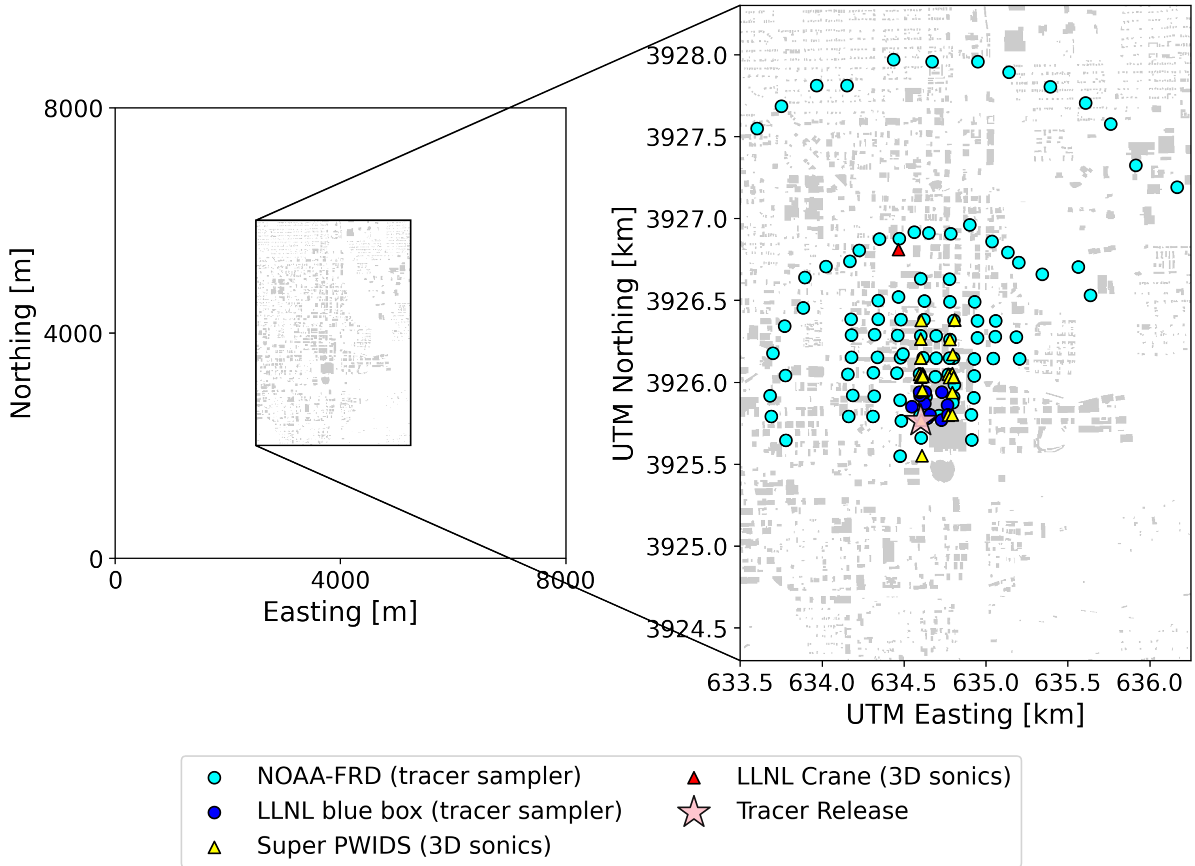

Figure 5. Plan view of the two-domain nested setup used for the convective case study (IOP3) with the tracer release site overlaid (note that the tracer is released from the Botanical Garden site, a different location compared to the stable case study in Figure 3). The vast majority of measurement devices are in the same location as in Figure 3 with the LLNL blue box stations shifting to be closer to the Botanical Garden release site.

4.4 Numerics

For the stable case, the time step is approximately 0.25 s for the outer 5 m grid and 0.125 s for the 2.5 m AMR grid when turned on. For the convective case, the time step is approximately 0.30 s for the outer 10 m grid and 0.15 s for the inner 5 m grid. We use an upwind third-order advection scheme in both the horizontal and vertical directions. For scalars we use a weighted essentially non-oscillatory (WENO) third-order scheme. Lastly, for the LES closure, we use the TKE-1.5 parameterization (Deardorff, 1980).

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

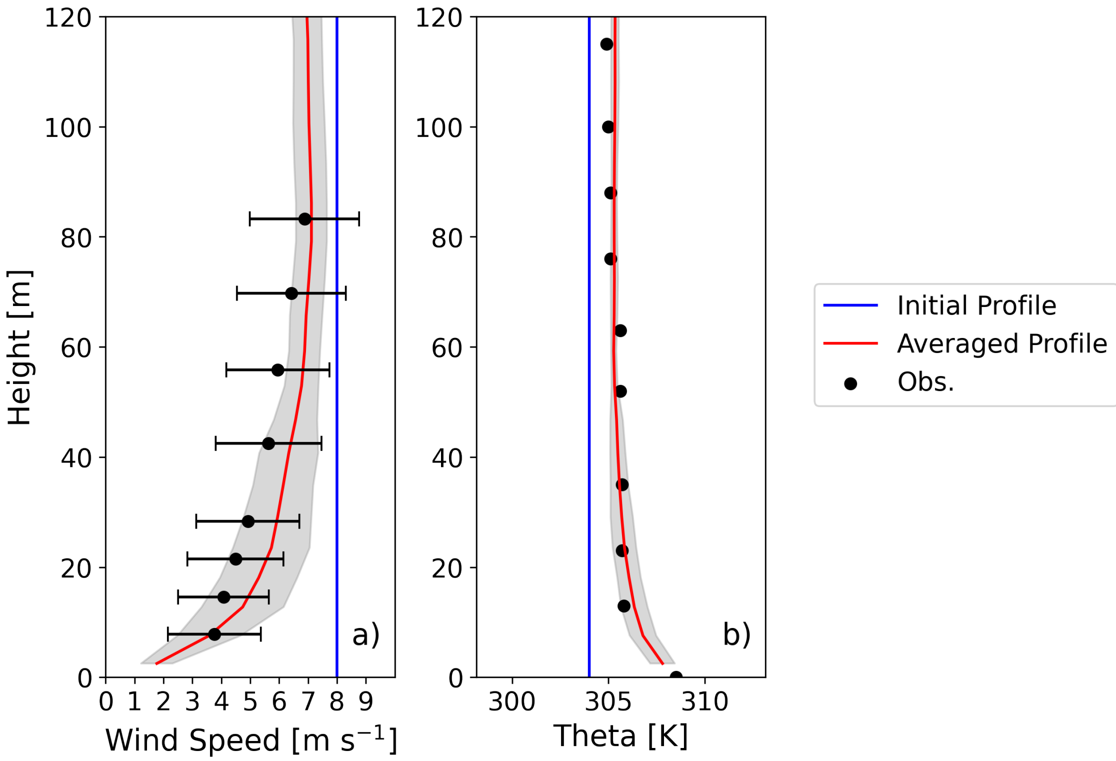

Figure 6. Vertical profiles of (a) wind speed and (b) potential temperature during the convective case study. The profiles include the initial profiles at the beginning of the simulation and the 30-minute averaged profiles during the analysis period. Wind speed data are from the LLNL crane site while potential temperature is from the ANL radiosonde (located outside of the simulated domain); however, data from the simulation are all from the crane site. The error bars represent the standard deviation during the analysis period at each height as in Figure 4.

5 Results and Discussion

5.1 Stable Conditions

The flow in the stable case study (IOP9) is characterized by relatively small-scale, coherent turbulent structures (characteristic of stable boundary layer turbulence) and northward flow. These characteristics, along with the fact that the set up is very efficient from a performance perspective, provide a framework for efficiently conducting sensitivity studies. Because of ERF's performance, we can iteratively improve our modeling parameters for optimal performance.

5.1.1 Sensitivity to building wall roughness lengths

The first sensitivity study is for the building wall roughness length, $z_{0,\mathrm{w}}$. The wall-model implemented with the immersed forcing capability requires a roughness length to represent building surfaces. This parameterization aims to represent the wall

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

material and any unrepresented geometries such as balconies or other subgrid-scale features. While buildings come in many forms and a single roughness value cannot be representative of all buildings, we simply do an order of magnitude analysis to determine the approximate order of magnitude that the roughness value should be. As such, we investigate three different building wall roughness values: 0.01 m, 0.10 m, and 1.00 m.

All simulations conducted for the building wall roughness sensitivity study use AMR. The wind speed and SF_6 concentration horizontal and vertical transects in Figures 7 and 8 demonstrates how the AMR grid changes over the course of the simulation with a building wall roughness of 0.1 m (the AMR grid evolves similarly for the other building wall roughness values and are not shown for brevity). Towards the beginning of the period, the SF_6 concentration is centered around the release location as the SF_6 has not had a chance to advect throughout the domain. As a result, the size of the grid is largely determined by the specified buffer region (250 m in the spanwise dimension and 750 m in the streamwise dimension). As the flow evolves and the SF_6 concentration disperses into a clear plume, the AMR grid continually expands downwind and vertically refining the required region.

Also demonstrated in Figures 7 and 8 are the small-scale stable boundary layer structures interacting with the buildings. Building wakes, lower wind speeds downwind of buildings, are most prominent for the larger and taller buildings in the central business district. The individual building wakes have an aggregate effect with slower wind speeds, in general, downwind of the central business district. Flow speedup or faster wind speeds in the street canyons especially north-south oriented streets of the central business district is also evident. The SF_6 concentration is influenced by the slower wind speeds downwind of the central business district with the concentration largely contained by these slower wind speeds with little lateral mixing.

Figure 9 provides a sense of the three dimensionality of the plume. The plume initially grows in the lateral and vertical directions due to the interactions of the SF_6 concentration and the buildings but its extent does not increase significantly downwind other than due to advection with the mean flow. The plume does not increase significantly laterally or vertically due to the relatively small turbulent scales in the stably stratified conditions.

There are subtle but important differences when analyzing the time-averaged wind speed and concentrations for different building wall roughness lengths in Figure 10. As the building wall roughness length increases the near-surface wind speed decreases, which is most evident in the slower wind speeds downwind of the central business district and near the smaller residential buildings to the north. These slower wind speeds are due to the increased drag imparted on the flow with the larger roughness length. Qualitatively, the general behavior of the wind speeds in the central business and flow speedup in the street canyons is well-predicted by model regardless of building wall roughness length. A more quantitative analysis is discussed later on and available in Table 1. As a result of the slower wind speeds for larger building wall roughness lengths, the SF_6 plume behavior changes. For example, for the largest roughness length ($z_{0,\mathrm{w}} = 1.0$ m), there is less lateral mixing towards the west end of the central business district and this results in underestimates of the SF_6 concentrations compared to observations.

We can more quantitatively compare the observations and simulation results for the three building wall roughness length values with a scatter plot (Figure 11). Figure 11 also has lines corresponding to a perfect match (1:1), within a factor of 2, and within a factor of 5. The latter two metrics have been commonly used in JU2003 analyses (initially proposed by Chang and Hanna (2004)) and are defined as follows:

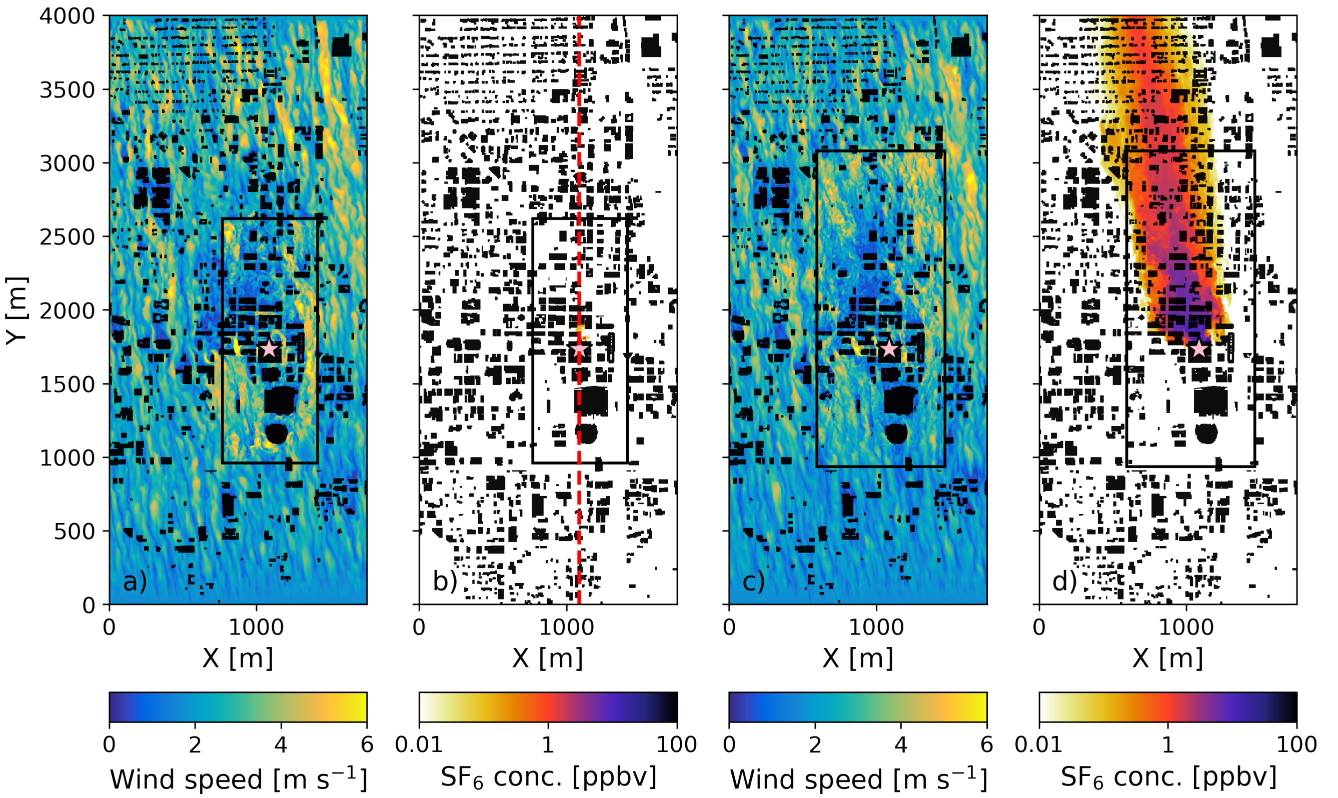

Figure 7. Instantaneous plan view of (a,c) wind speed at z = 7.5 m and (b,d) SF_6 concentration at z = 2.5 m at (a,b) 100 s into the analysis period and (c,d) at the end of the analysis period (1800 s) for IOP9. The black rectangle represents the extent of the refined region based on the AMR configuration. The red dashed line in (b) corresponds to the vertical cross sections shown in Figure 8 are taken from. Note that the visualization is for the $z_{0,\mathrm{w}} = 0.1$ m case.

$$\mathrm{FAC2} = \text{fraction of the data that satisfies } \frac{1}{2} \leq \frac{S_i}{O_i} \leq 2 \quad (5)$$

$$\mathrm{FAC5} = \text{fraction of the data that satisfies } \frac{1}{5} \leq \frac{S_i}{O_i} \leq 5 \quad (6)$$

where S_i and O_i are the simulated and observed values for individual stations, respectively. With the larger building wall roughness length of $z_{0,\mathrm{w}} = 1.0$ m, the simulated concentrations are, in general, underestimated in Figure 11. While the larger drag imparted on the flow does increase mixing in general, the wind speeds at street level and in the central business district are very small, smaller than when using smaller building wall roughness lengths, thus reducing dispersion. Results for the smaller building wall roughness lengths of 0.01 m and 0.1 m are qualitatively similar but there is a slight tendency to overestimate the

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

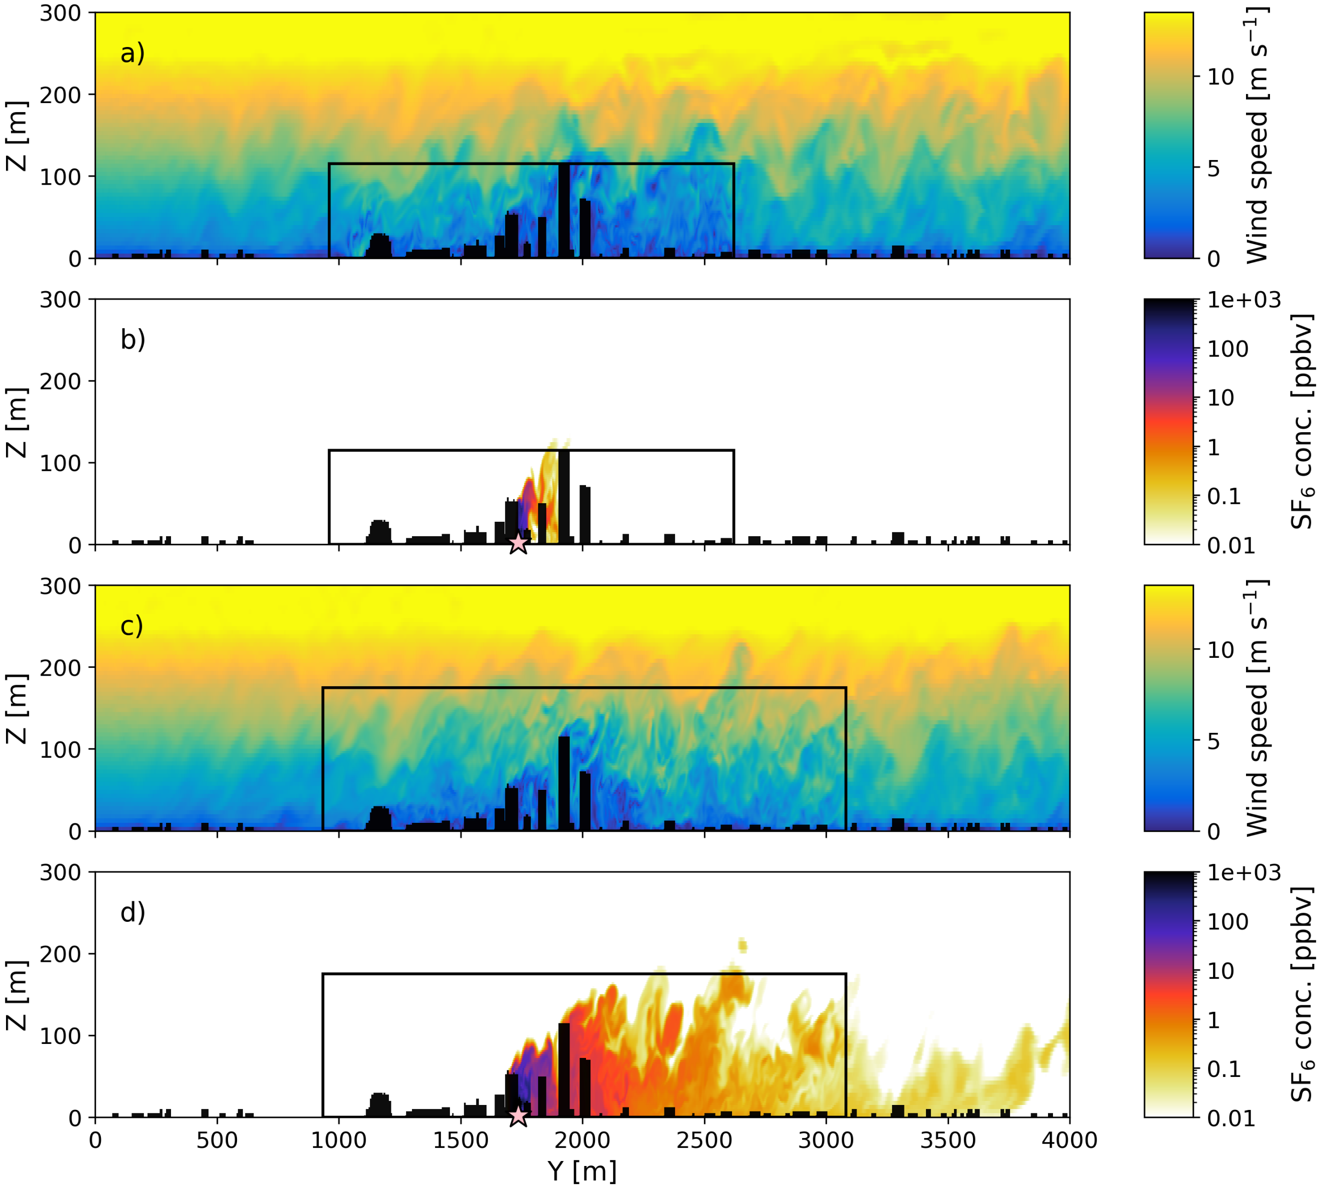

Figure 8. Instantaneous vertical cross section of (a,c) wind speed and (b,d) SF_6 concentration at (a,b) 100 s into the analysis period and (c,d) at the end of the analysis period (1800 s) for IOP9. The black rectangle represents the extent of the refined region based on the AMR configuration. Note that the visualization is for the $z_{0,\mathrm{w}} = 0.1$ m case.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

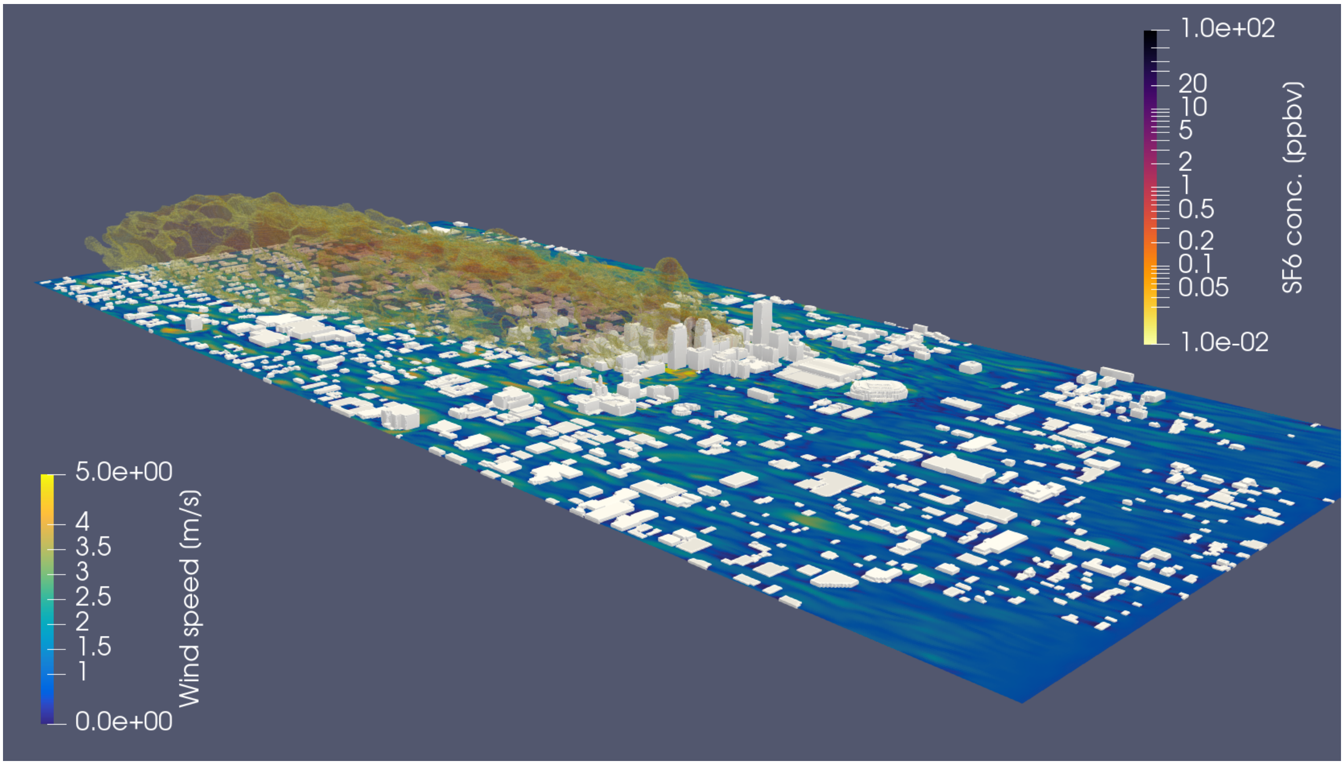

Figure 9. Three-dimensional instantaneous visualization of near-surface wind speed and SF_6 plume dispersion at the end of the analysis period of IOP9. Note that the visualization is for the $z_{0,\mathrm{w}} = 0.1$ m case.

concentrations for the smallest roughness length. This bias is discussed in more detail in the following paragraph. Overall, 81% of stations are within a factor of two for the $z_{0,\mathrm{w}} = 0.1$ m case with 79% for the $z_{0,\mathrm{w}} = 0.01$ m case (there are 68 total stations in the simulated area). For $z_{0,\mathrm{w}} = 0.01$ m case, 97% of stations are within a factor of five while 100% of stations are within a factor of five for the $z_{0,\mathrm{w}} = 0.1$ m case. For the $z_{0,\mathrm{w}} = 1.0$ m case, only 93% of stations are within a factor of five resulting in the lowest model skill of the evaluated configurations (see Table 1).

Other commonly used metrics to assess model skill are the fractional bias (FB), the geometric mean (MG), and the normalized mean squared error (NMSE). These metrics will be discussed in the following paragraph and are defined as follows:

$$\mathrm{FB} = \frac{\overline{O_i} - \overline{S_i}}{0.5(\overline{O_i} + \overline{S_i})} \tag{7}$$

$$\mathrm{MG} = \exp\left[\overline{\ln(O_i)} - \overline{\ln(S_i)}\right] \tag{8}$$

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

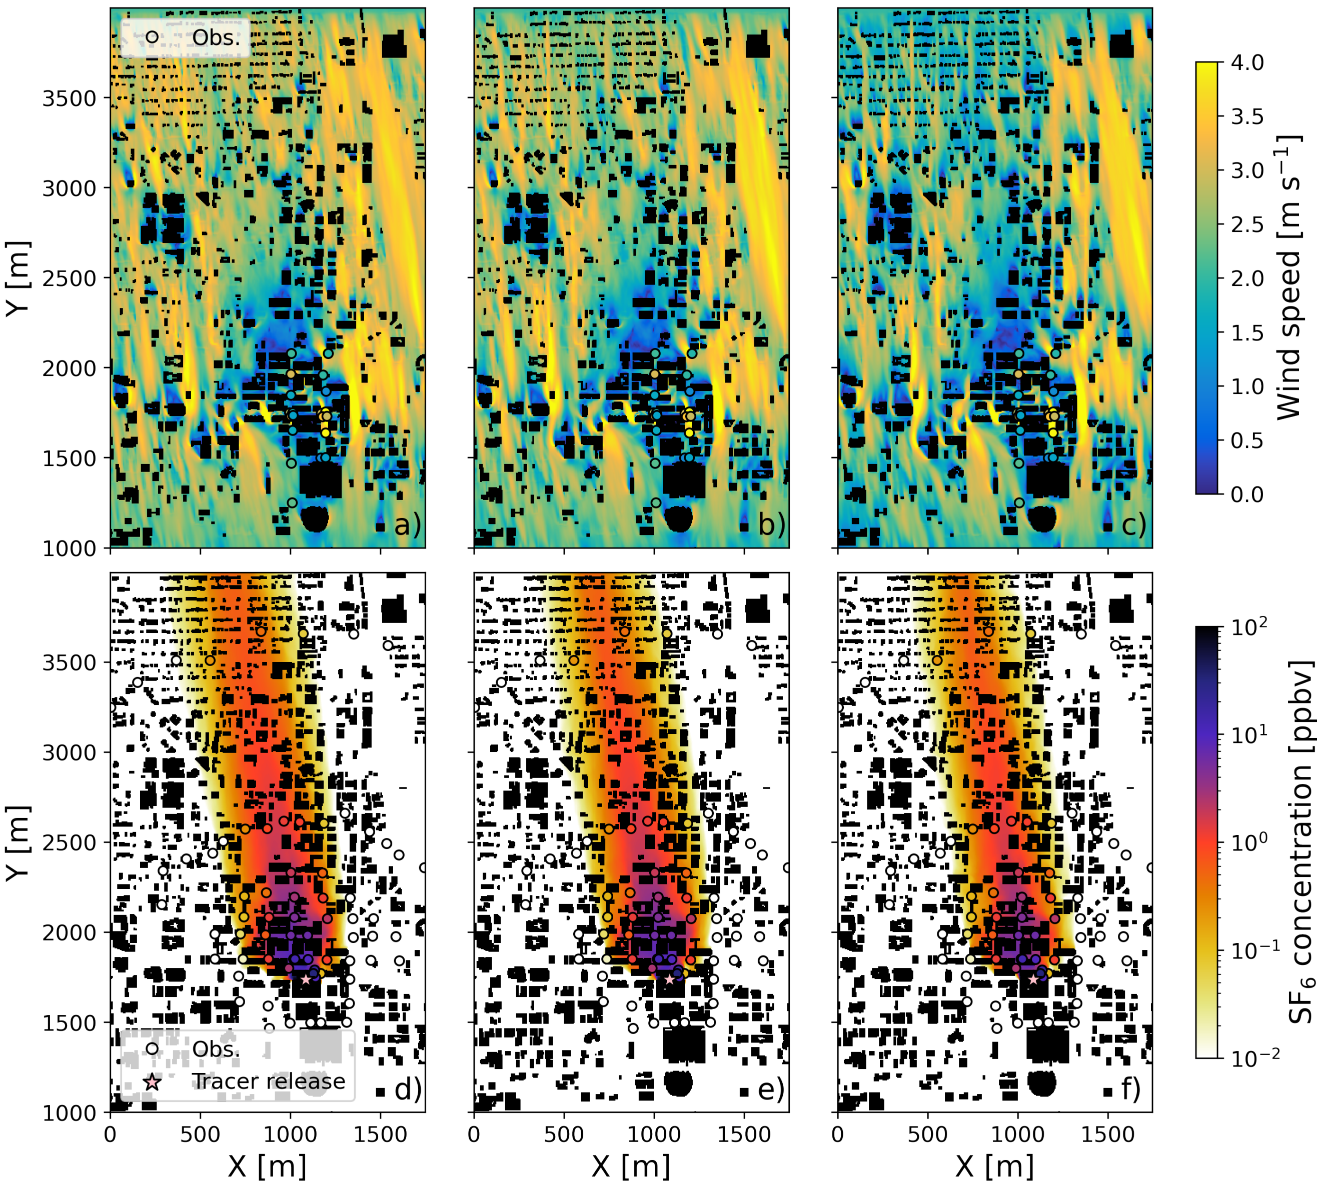

Figure 10. 30 minute averaged plan view of (a,b,c) wind speed at z = 7.5 m and (d,e,f) SF_6 concentration at z = 2.5 m for the (a,d) $z_{0,\mathrm{w}} = 0.01$ m, (b,e) $z_{0,\mathrm{w}} = 0.1$ m, and (c,f) $z_{0,\mathrm{w}} = 1.0$ m cases for IOP9. For wind speed, data from the Super-PWIDS are overlaid. For SF_6 concentration, data from the NOAA-FRD PIGS and LLNL blue box stations are overlaid with only stations measuring greater than 0.01 ppbv included.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

$$\text{NMSE} = \frac{\overline{(O_i - S_i)^2}}{\overline{O_i S_i}} \tag{9}$$

with the overbars denoting averaging over all of the stations. At the top of Table 1 are the values for each error metric that correspond to a perfect model. For FAC2 and FAC5, this value is 1.00 or 100% of the stations are within the specified factor. For the fractional bias and NMSE, a perfect model scores a value of 0.00 while for the geometric mean, a perfect model should score a value of 1.00 with values less than that corresponding to an underestimation on average and value over corresponding to an overestimation.

Regardless of building wall roughness length, ERF scores very well in terms of error metrics for wind speed (Table 1). The model wind speeds and observed wind speeds are compared at 20 Super PWIDS that are all largely in or very near the central business district of Oklahoma City. 90% of stations are within a factor of two and 100% are of stations are predicted within a factor of five regardless of building wall roughness value. Most previous modeling studies have reported values of below 75% for FAC2 (see Table 1 in Muñoz-Esparza et al. (2020)) other than certain configurations by Wiersema et al. (2020) and Wiersema et al. (2022) who achieved up to 90-100% FAC2 for the 15 Super PWIDS in their domain. Additionally, the NMSE and FB in this study are among the lowest reported, as well.

For the SF_6 concentration, we compare with 62 NOAA-FRD PIGS gas analyzers, which include 40 in the central business district, 15 in the 1 km arc, and 7 in the 2 km arc. We also compare with 6 LLNL blue box stations. However, note that the PIGS instruments have a method limit of quantification (MLOQ) of 0.005 ppbv (Clawson et al., 2005). Below the MLOQ, any SF_6 cannot be accurately quantified; therefore, 0.005 ppbv acts as a floor for the PIGS instruments. To maintain consistency between observations and model predictions, this floor value of 0.005 ppbv is also applied to simulation results. No floor is needed for the LLNL blue box stations as they measured relatively high concentrations since they are all located relatively close and downwind of the release site. Ultimately, we conclude that the $z_{0,\mathrm{w}} = 0.1$ m case performs the best. While the FB and NMSE for the $z_{0,\mathrm{w}} = 0.1$ m case are slightly less accurate compared to the $z_{0,\mathrm{w}} = 0.01$ m case, a larger number of stations are within both a factor of two and a factor of five, which the authors weigh more heavily. Our study compares simulation results with the largest number of stations as well as the furthest downwind stations for turbulence- and building-resolved modeling studies in the published literature. Considering that other modeling studies have domains that cover different extents of Oklahoma City and different IOPs, it can be difficult or misleading to compare between studies with a vastly different number of measurement stations taken into account when calculating errors. However; the results from the present study compare very favorably. The FB and NMSE for both the $z_{0,\mathrm{w}} = 0.1$ m and $z_{0,\mathrm{w}} = 0.01$ m cases are among the lowest for published modeling studies and the MG for the $z_{0,\mathrm{w}} = 0.1$ m is the closest to 1.0 (see Table 1 in Muñoz-Esparza et al. (2020)).

5.1.2 Sensitivity to model configurations

The simulation configuration in the prior section produces highly accurate results using AMR. In this section we examine the balance between performance and accuracy when considering the parent domain with no refinement (i.e., the grid spacing is 5 m throughout the entire domain). Additionally, we compare the previous setup with a similar configuration (parent domain

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Table 1. Wind speed and SF_6 modeling errors for building wall roughness sensitivity study for IOP9.

	FAC2	FAC5	FB	MG	NMSE
Perfect model	1.00	1.00	0.00	1.00	0.00
		Wind speed			
$z_{0,w}$ = 0.01 m	0.90	1.00	-0.08	1.10	0.08
$z_{0,w}$ = 0.10 m	0.90	1.00	-0.12	1.11	0.11
$z_{0,w}$ = 1.00 m	0.90	1.00	-0.08	1.02	0.13
		SF_6 concentration			
$z_{0,w}$ = 0.01 m	0.79	0.97	0.09	1.07	0.54
$z_{0,w}$ = 0.10 m	0.81	1.00	0.22	0.94	0.73
$z_{0,w}$ = 1.00 m	0.69	0.93	0.93	0.67	3.07

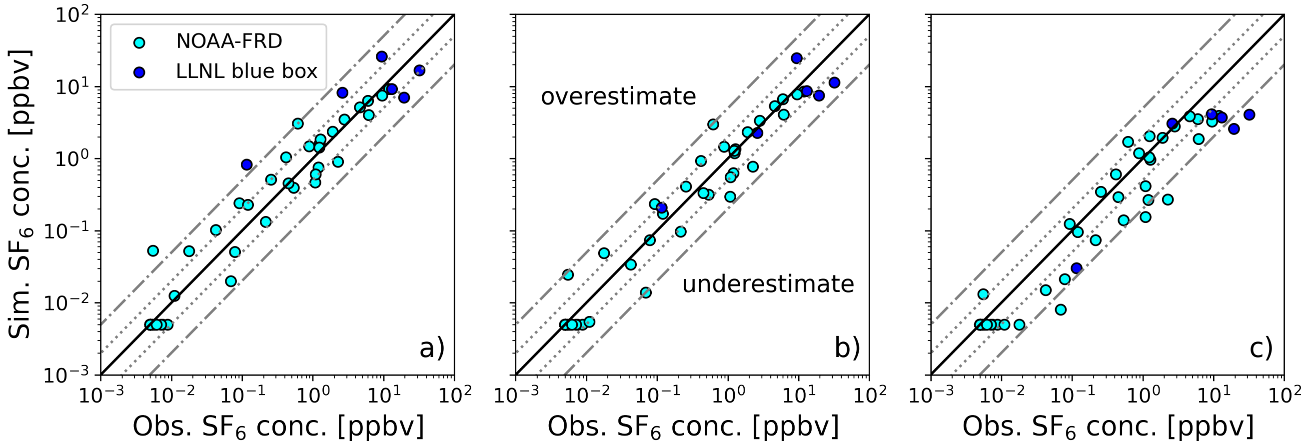

Figure 11. Simulated vs observed SF_6 concentrations for the (a) $z_{0,w}$ = 0.01 m, (b) $z_{0,w}$ = 0.1 m, and (c) $z_{0,w}$ = 1.0 m cases for IOP9. The solid line is a perfect correlation while the dotted lines represent a factor of two (FAC2) and the dashed-dotted lines represent a factor of five (FAC5). Note that the clustering of data near 0.005 ppbv is due to the MLOQ floor imposed on both the observations and simulation results.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

only) but with ERF running in anelastic mode as opposed to fully compressible mode. In this subsection, we will compare three different simulation configurations:

- Config. 1: Fully Compressible Δ = 2.5 m (AMR).
- Config. 2: Fully Compressible Δ = 5.0 m.
- Config. 3: Anelastic Δ = 5.0 m.

For these sets of simulations, they are all run with a building wall roughness length of 0.1 m.

The most consequential impact when not using the adaptive grid refinement is that the turbulence kinetic energy (TKE) is decreased. In Figure 12, the resolved TKE is significantly higher for the Fully Compressible Δ = 2.5 m (AMR) configuration compared to the other configurations, which do not include the 2.5 m grid. Note that only the resolved turbulence is shown; however, we found that at this height (z = 7.5 m), the TKE is largely resolved and at a similar value between the two fully compressible configurations (93% for Δ = 2.5 m (AMR) and 92% for Δ = 5.0 m) and slightly lower for the anelastic configuration (78%). The high percentages of resolved TKE relatively close to the bottom boundary is due to the main contribution of turbulence being building-induced. Ultimately, the higher TKE on the finer grid is due to better-resolved buildings and energetic eddies such as, for example, those from shear-generated turbulence.

Also included in Figure 12 are time-averages of the observed TKE from the Super PWIDS stations. Qualitatively, the agreement between the model and observed values are very good for the Fully Compressible Δ = 2.5 m (AMR) configuration with weaker agreement for the other configurations. While the magnitude for the configurations with 5.0 m grid spacing is underpredicted by the model, the qualitative behavior, such as gradients within street canyons, is still captured well. For example, the intersection of Broadway Ave and Park Ave (the cluster of four measurement locations on the right of the subplots in Figure 12) was outfitted with four Super PWIDS one in each corner and all configurations capture increased TKE in the eastern half of the intersection. The general underestimation of TKE by the configurations with only 5.0 m grid spacing is quantified in terms of the error metrics (Table 2). In terms of magnitude, the configurations with only 5.0 m grid spacing underestimate TKE by approximately 1.2 $m^2\ s^{-2}$ on average across all 20 stations. In general, Table 2 demonstrates that errors are much lower when using the Fully Compressible Δ = 2.5 m (AMR) configuration compared to the other configurations. The 5.0 m grid spacing configurations have a much more positive FB and much lower MG indicating underestimation of TKE. Additionally, 80% of stations are within a factor of 2 when using the Fully Compressible Δ = 2.5 m (AMR) configuration while just 15% are for the 5.0 m grid spacing configurations.

The turbulence can also be visualized by comparing the three cases with point measurements from individual stations. Figures 13(a) and (b) show the vertical velocity at a downwind site near the northeastern edge of the domain shown in Fig. 12 (denoted as a triangle), while Figs. 13(c) and (d) show the vertical velocity at an upwind site near the southern edge of the domain located near the botanical garden (denoted as a square in Fig. 12). Hereinafter, we refer to these as the downwind and upwind sites. The upwind site is more open, with fewer buildings around it than the downwind site. Consequently, the vertical velocity is less energetic at all frequencies compared to the observations (Fig. 12(d)); however, there is a notable increase

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

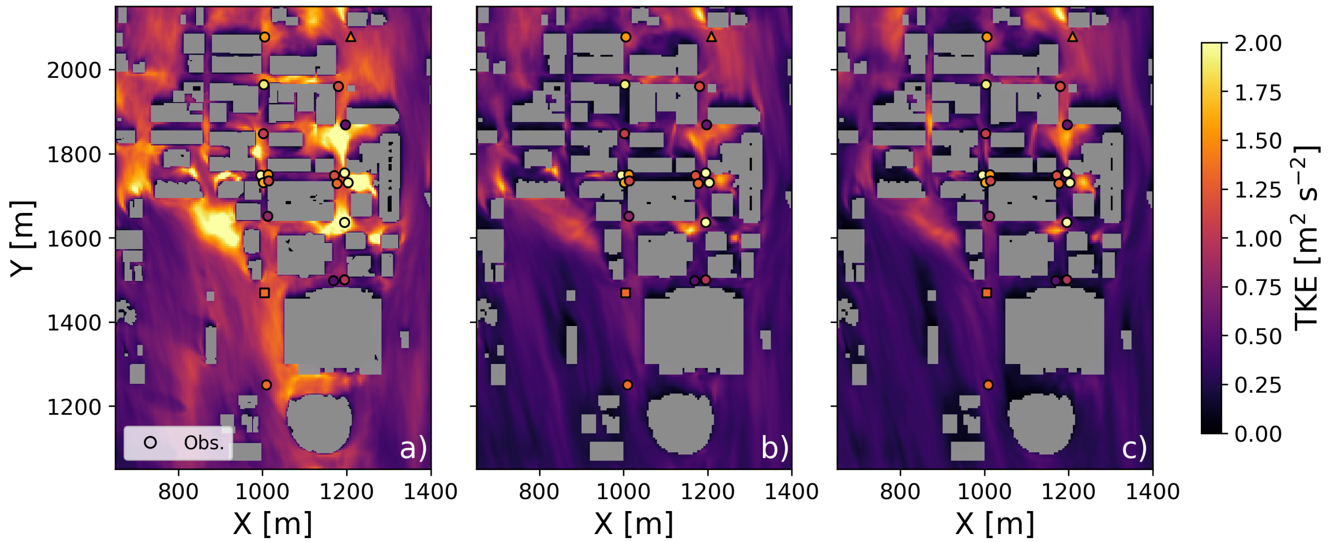

Figure 12. 30 minute averaged TKE for the (a) fully compressible Δ = 2.5 m (AMR), (b) fully compressible Δ = 5.0 m, and (c) anelastic, Δ = 5.0 m configurations with data from the Super-PWIDS overlaid for IOP9 at z = 7.5 m. The TKE from the simulations is the resolved TKE. Note that this visualization is a zoom in on the central business district where the measurements are located.

in energy for the Fully Compressible Δ = 2.5 m (AMR) configuration compared to the Δ = 5.0 m configurations, which is expected given the increase in grid resolution. Additionally, there is a marginal increase in energy for the Fully Compressible Δ = 5.0 m configuration compared to the Anelastic Δ = 5.0 m configuration. In preliminary simulations for a similar setup but without buildings (not shown), we also observed slightly less energy in the inertial subrange when running in anelastic mode compared to fully compressible mode. At the downwind site, the energy in the vertical velocity signal is nearly identical for the Δ = 5.0 m configurations, as the main turbulence contribution comes from the nearby and upwind buildings. For the downwind site, the Fully Compressible Δ = 2.5 m (AMR) configuration is the most energetic and matches well with the observations at lower frequencies. Additionally, the difference in the energy between the Fully Compressible Δ = 5.0 m (AMR) configuration and the Δ = 5.0 m configurations at the downwind site is smaller across all frequencies compared to the upwind site.

The largest difference between the configurations is in the TKE; however, this is not necessarily impactful when it comes to the SF_6 concentration. Figure 14 shows the observed and simulated SF_6 concentrations and we can see that the Fully Compressible Δ = 2.5 m (AMR) and the Fully Compressible Δ = 5.0 m configurations have relatively similar spreads. The Anelastic Δ = 5.0 m configuration also has good performance but with some underestimation of stations with lower concentrations. The Anelastic Δ = 5.0 m configuration also has a slightly lower MG for TKE, indicating a further underestimation relative to the Fully Compressible Δ = 5.0 m configuration. This is consistent with the spectra shown for the upwind site in Fig. 13(d). The reduced turbulence mixing may also contribute to the larger SF_6 errors by decreasing dispersion. Overall, the results are robust across all three configurations, which perform similarly in terms of FAC2 and FAC5 for concentration (Table 2). While the turbulence levels are the main differences between the three configurations, the similar behavior of the SF_6 concentration indi-

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

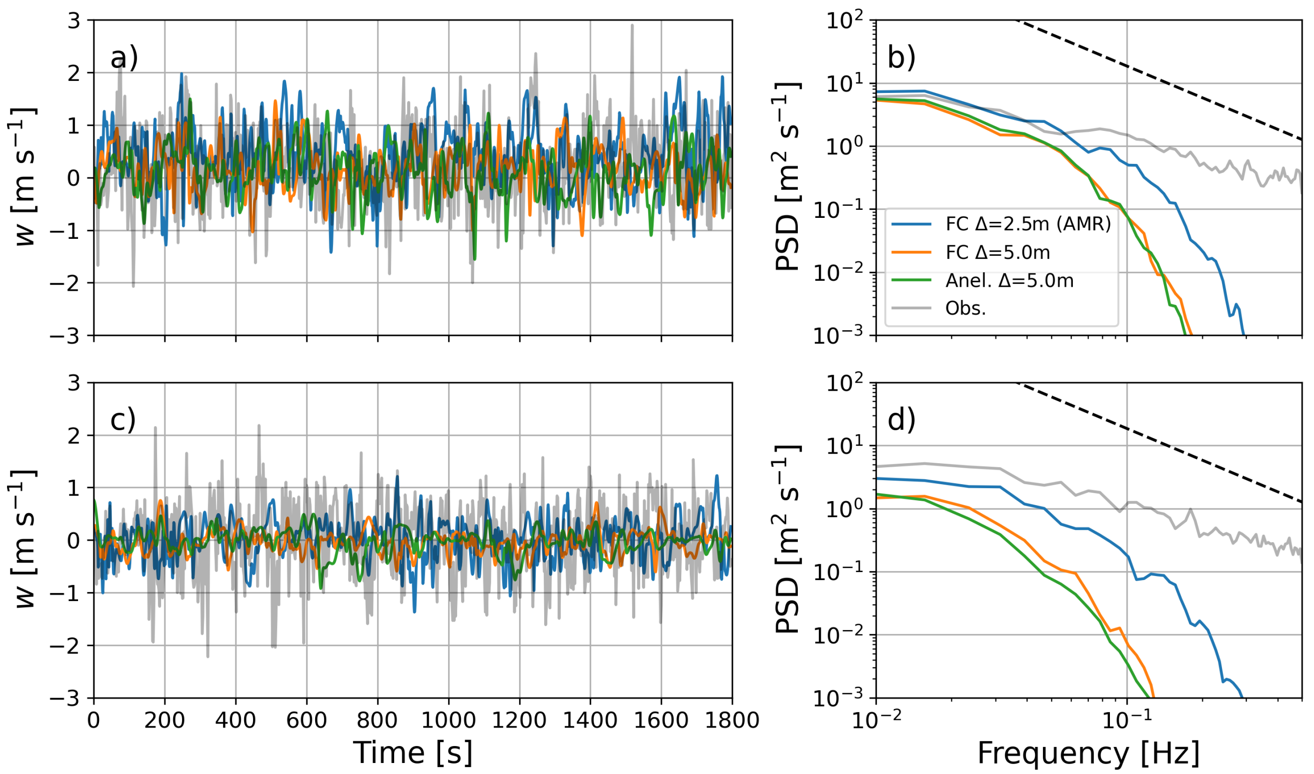

Figure 13. (a,c) Time-series and (b,d) power spectral density (PSD) of the vertical velocity towards the (a,b) northern edge (c,d) and southern edge of the central business district. The -5/3 inertial subrange power law is provided for reference as the black dashed line in (b) and (d).

cates that the plume behavior at the height analyzed (z = 7.5 m) is largely driven by mean flow effects such as the wind speed and wind direction. We also examined the plume's vertical structure and found that the Fully Compressible Δ = 2.5 m (AMR) configuration, with its higher turbulence levels, produces increased vertical SF_6 flux in the central business district and slightly more SF_6 transport to higher elevations (not shown). However, with a lack of measurements available above street-scale, further analysis is left for future work.

While there are minor accuracy losses when not incorporating the refined 2.5 m region, the tradeoff is increased computational efficiency. We use the same CFL scaling as in section 2.1 and the time step coincidentally approaches a similar value of 0.25 s for the Δ = 5.0 m configurations (whereas the Fully Compressible Δ = 2.5 m (AMR) configuration uses a timestep of 0.125 s). The simulations are each run using 8 GPUs and take 7334 s, 2291 s, and 893 s for the Fully Compressible Δ = 2.5 m (AMR), Fully Compressible Δ = 5.0 m, and Anelastic Δ = 5.0 m configurations, respectfully. Using the Fully Compressible Δ = 2.5 m (AMR) configurations as the baseline, the Fully Compressible Δ = 5.0 m configuration offers a 3.2x speedup while the Anelastic Δ = 5.0 m configuration offers a 8.2x speedup. While the accuracy is slightly slower in terms of SF_6 concentration, the decrease in computational cost makes the anelastic configuration a promising option for emergency response where

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

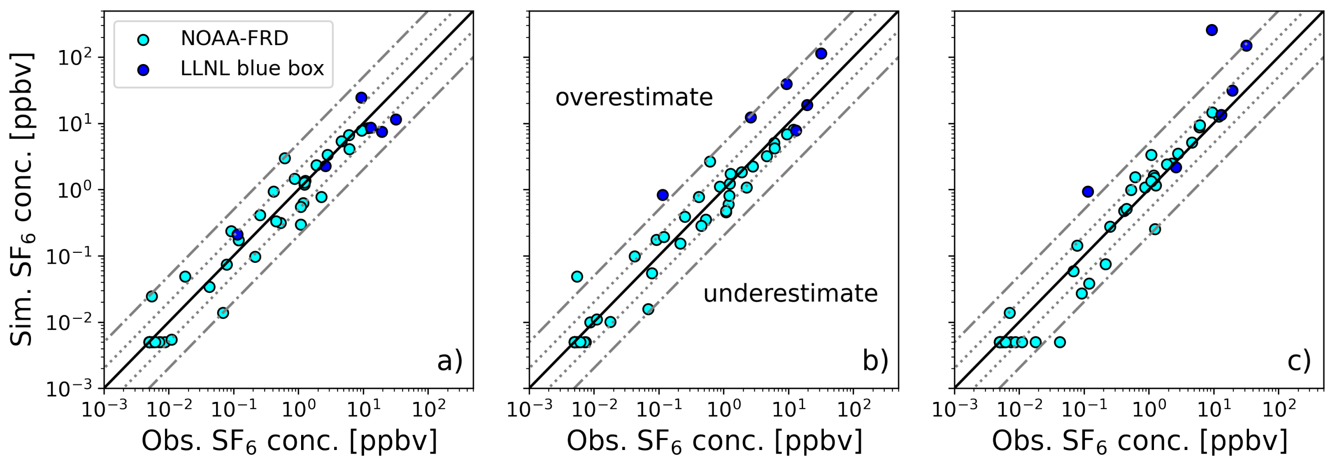

Figure 14. Simulated vs observed SF_6 concentrations for the (a) fully compressible Δ = 2.5 m (AMR), (b) fully compressible Δ = 5.0 m, and (c) Anelastic, Δ = 5.0 m configurations for IOP9. The solid line is a perfect correlation while the dotted lines represent a factor of two (FAC2) and the dashed-dotted lines represent a factor of five (FAC5).

Table 2. Wind speed, SF_6 concentration, and TKE modeling errors for the three different configurations for IOP9.

	FAC2	FAC5	FB	MG	NMSE
Perfect model	1.00	1.00	0.00	1.00	0.00
Wind speed					
Fully compressible Δ = 2.5 m (AMR)	0.90	1.00	-0.12	1.11	0.11
Fully compressible Δ = 5.0 m	0.95	1.00	-0.04	0.97	0.10
Anelastic Δ = 5.0 m	0.90	1.00	0.03	0.88	0.11
SF_6 concentration					
Fully compressible Δ = 2.5 m (AMR)	0.81	1.00	0.22	0.95	0.73
Fully compressible Δ = 5.0 m	0.84	0.97	-0.57	1.03	1.62
Anelastic Δ = 5.0 m	0.82	0.96	-1.20	1.04	8.94
TKE					
Fully compressible Δ = 2.5 m (AMR)	0.80	1.00	0.24	0.83	0.24
Fully compressible Δ = 5.0 m	0.15	0.65	1.11	0.26	2.92
Anelastic Δ = 5.0 m	0.15	0.60	1.13	0.22	3.26

fast simulations are needed. The anelastic configuration could also be beneficial for those with limited computational resources or for improving ensembles. Overall, the results are robust regardless of configuration demonstration the capabilities of ERF and immersed forcing.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

5.2 Convective conditions

Another commonly simulated IOP for the JU2003 experiment is IOP3, which is characterized by typical daytime convective conditions. The turbulent structures are much larger for this case study compared to the previous stable case study, which is notable as the large-scale structures are important for the lateral mixing and meandering of the SF_6 plume. For IOP3, we only simulate a single configuration (as described in Section 4) using fully compressible mode and with a building wall roughness length of 0.1 m based on findings in the previous sections.

Figures 15, 16, and 17 demonstrate the larger-scale turbulent structures and their effect on the SF_6 dispersion. From the plan view, the lateral meandering of the plume is evident and more prominent when compared to the stable conditions shown in Figure 7. The meandering is illustrated in Figure 15 where the dashed line is aligned with the mean wind direction. At this time instance and further downwind, there is more SF_6 concentration to the west of the dashed line but this varies over the course of the analysis period. This lateral meandering is due wind direction variability from large-scale eddies.

The convective conditions also result in more SF_6 being transported to higher altitudes due to enhanced vertical mixing by buoyancy, relative to the stable case. Updrafts are the main driver of vertical mixing during this IOP and this is observed in the SF_6 field in Figure 16. The SF_6 concentration acting as a passive tracer highlights the buoyant behavior of the flow. Figure 17 also shows the vertical extent of the plume, which is much higher in the atmosphere compared to the stable case study. For the stable case (Figure 9), the plume did not extend much higher than the tallest buildings as the top of the boundary layer was only 200 m above ground level; however, the convective case boundary layer height is over 1 km from the ground allowing for more vertical transport of SF_6. Generally, there is both more horizontal and vertical transport of SF_6 for the convective case, which is expected for convective conditions compared to stable conditions.

The average SF_6 concentration over the 30 minute analysis period near the surface further illustrates the lateral spreading of SF_6. Similar to Figure 10, the simulation results in Figures 18a and 18b are overlaid with time averages from the measurement stations of wind speed and SF_6. For wind speed, the domain includes the same 20 Super PWIDS stations as in the stable case. For SF_6 concentration, the different domain from the stable case includes a slightly different number of stations. There are 63 NOAA ARL FRD gas analyzers, 34 in the central business district, 16 in the 1 km arc, and 13 in the 2 km arc, and 11 LLNL blue box stations. As expected, the lateral extent of the SF_6 plume is much larger for the convective case compared to the stable case. While the lateral spread in the downwind direction of the SF_6 concentration matches well with the observations, there are a number of stations upwind and to the west of the tracer release location that are significantly underestimated. This underestimation is likely because of missing submeso motions as has been similarly discussed by Wiersema et al. (2020) and Wiersema et al. (2022). An idealized setup is limited in capturing the larger-scale turbulent structures that multiscale (mesoscale to microscale) setups are able to capture. Additionally, just a single wind direction drives the flow when, in reality, there is wind direction variability over the 30 minute analysis period. Spectra (not shown) and TKE did show good agreement with observations (see Table 3); however, during such a short analysis period and given that the stations underpredicted typically measured very small concentrations, a single gust or submeso eddie could be responsible for the stations that were missed by the model.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

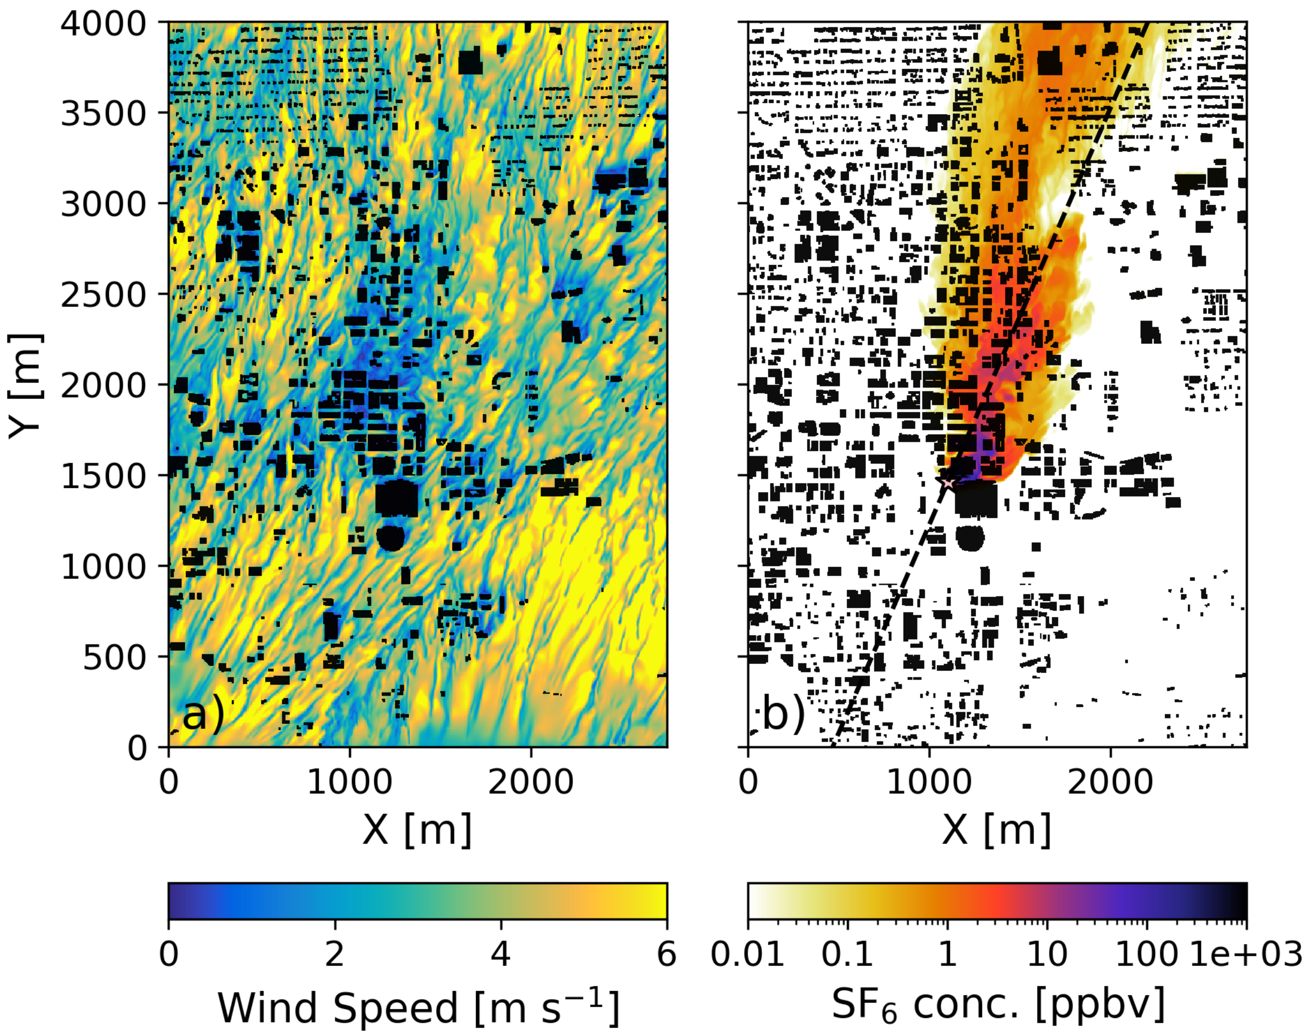

Figure 15. Instantaneous plan view of (a) wind speed at z = 7.5 m and (b) SF_6 concentration at z = 2.5 m at the end of the analysis period for IOP3. The dashed line corresponds to the vertical cross section in Figure 16.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

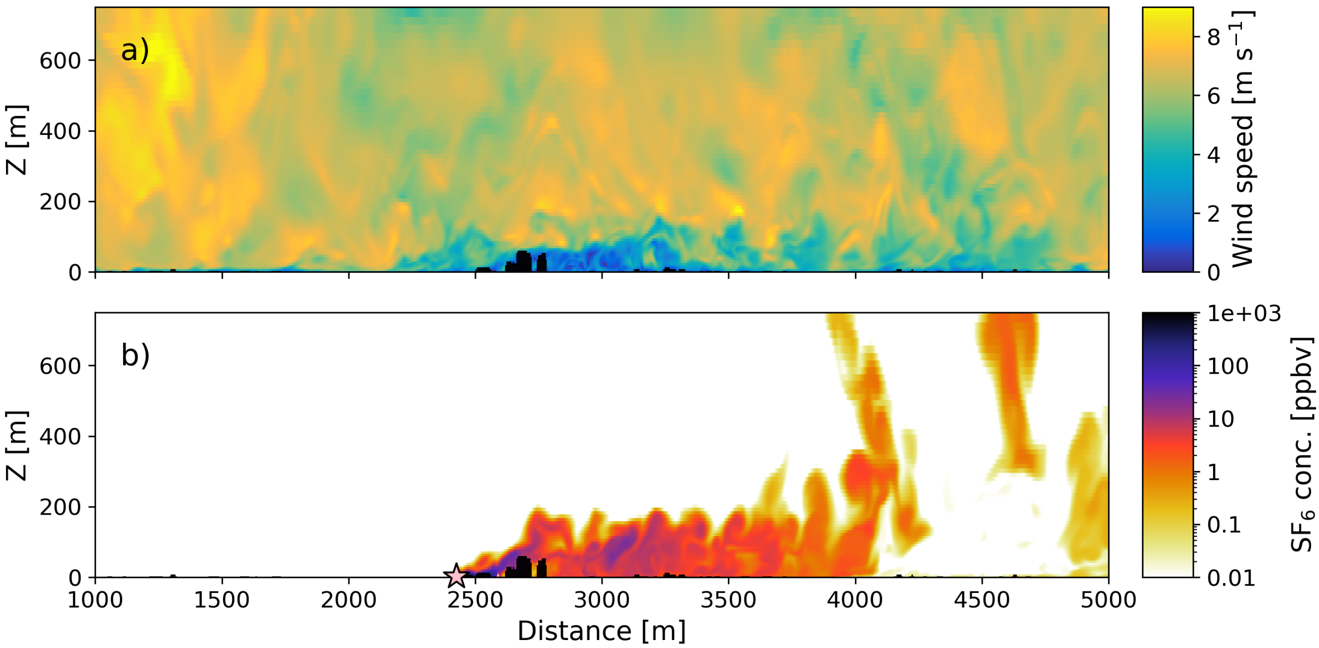

Figure 16. Instantaneous vertical cross section of (a) wind speed and (b) SF_6 concentration at the end of the analysis period for IOP3.

The underestimation of the western and southern stations is illustrated in Figure 18c. There are six stations with observed concentrations between 10^{-2} and 10^{-1} ppbv where the predicted concentrations are effectively 0 (or the floor value of 0.005). While the SF_6 concentration is underpredicted at these stations, it is important to note that the concentration for these stations is very small (< 0.1 ppbv). As a result, the FAC2 and FAC5 values are 0.64 and 0.85, respectively (see Table 3). Due to the large extent of the computational domain, these calculations include a total of 74 stations, which is, again, the largest number of stations used in a model comparison in the published literature (for turbulence- and building-resolved models). While these values are lower compared to the stable case, they are still among best in the published literature for IOP3 (see Table 1 in Muñoz-Esparza et al. (2020)). The wind speed metrics all indicate good model skill with TKE also reasonably well-predicted. However, the positive FB indicates that there is some underestimation of TKE. This underestimation of TKE is due to 1) only using a grid spacing of 5.0 m and 2) the idealized configuration limiting the upper length scale of turbulence as these larger scales of turbulence can significantly influence TKE. Ultimately, with a more realistic setup containing land use variability, terrain, mesoscale physics, etc., the SF_6 measurement stations on the western and southern area of the central business district would likely be better predicted. Additionally, the absence of thermal wall forcing likely leads to decreased lateral dispersion and mixing, in general (such as street canyon updrafts). Running ERF in a more realistic set up and with thermal forcing is a goal of future work.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

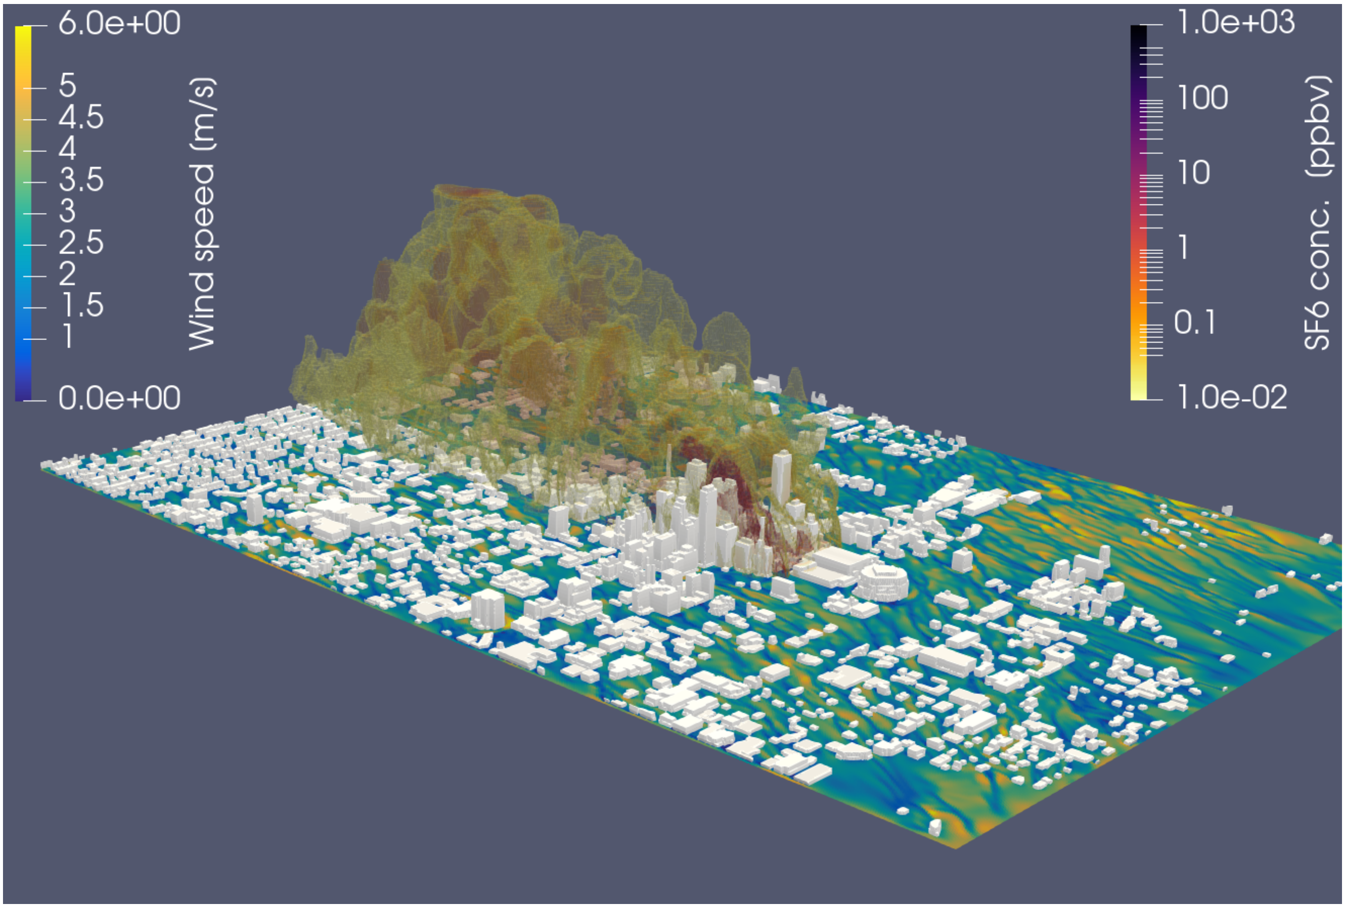

Figure 17. Three-dimensional instantaneous visualization of near-surface wind speed and SF_6 plume dispersion at the end of the analysis period for IOP3.

Table 3. Wind speed, SF_6 concentration, and TKE modeling errors for the convective case study (IOP3).

	FAC2	FAC5	FB	MG	NMSE
Perfect model	1.00	1.00	0.00	1.00	0.00
Wind speed	0.90	1.00	-0.05	1.02	0.17
SF_6 concentration	0.64	0.84	-0.51	0.77	0.47
TKE	0.60	0.90	0.61	0.55	2.30

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

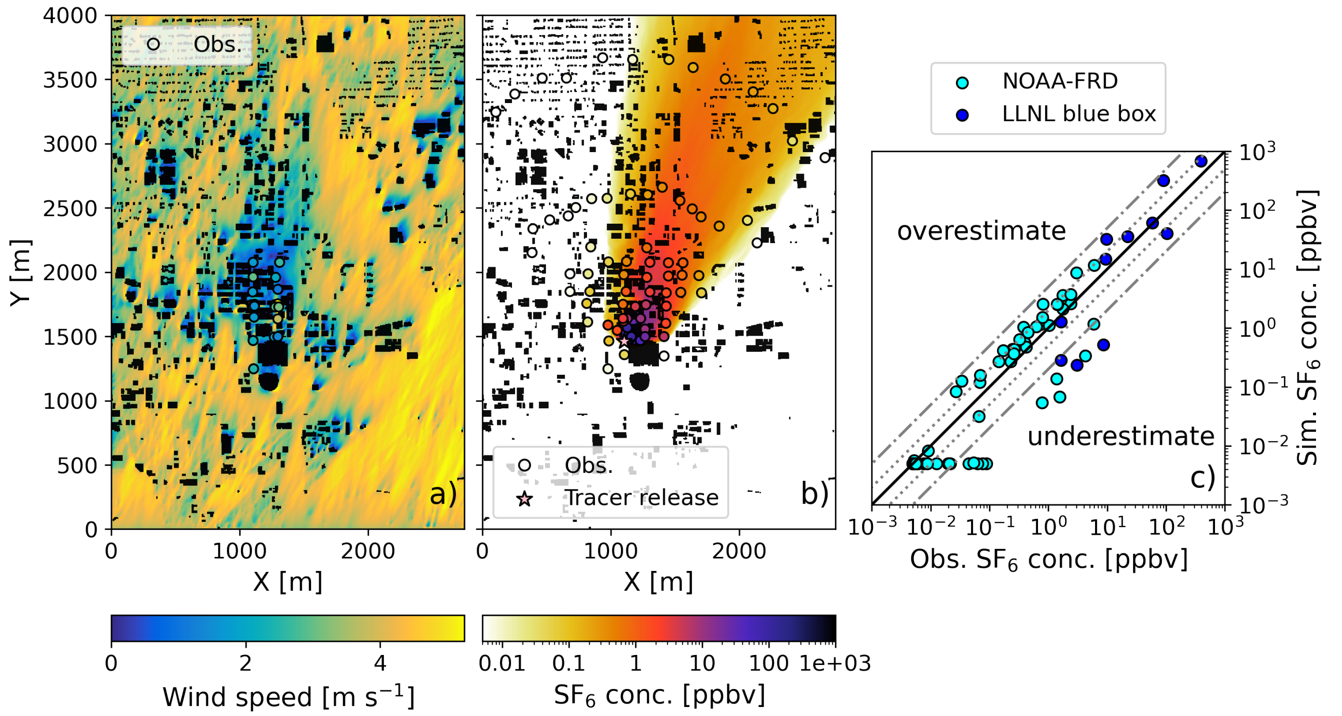

Figure 18. 30 minute averaged plan view of (a) wind speed at z = 7.5 m and (b) SF_6 concentration at z = 2.5 m. (c) Simulated vs observed SF_6 concentrations for IOP3. For the wind speed in (a), data from the Super-PWIDS are overlaid. For SF_6 concentration, data are from the NOAA-FRD PIGS and LLNL blue box stations.

6 Conclusions

In this study, we have demonstrated and validated the immersed forcing capability implemented into the Energy Research and Forecasting (ERF) model for urban applications. The acceleration computations in the ERF model make real-time urban simulations possible given enough computational resources. Moreover, the immersed forcing capability is highly parallelizable and efficient making it adept for models designed to run on GPUs or other exascale hardware. For example, when simulating a 2.5 km × 2.5 km domain at 5 m grid spacing when using ERF in fully compressible mode, real-time speed is achieved with 32 GPUs. Even fewer resources are necessary when using anelastic mode. Additionally, we leverage the adaptive mesh refinement capability of the underlying AMReX framework to further optimize ERF performance.

The immersed forcing capability has been used in one form or another in other atmospheric modeling tools; in the present work, we have extended this approach to include a wall-model. The validation of the immersed forcing with a wall-model capability was carried out using wind measurements and SF_6 tracer concentrations from both stable and convective intensive observational periods (IOP) during the Joint Urban 2003 field campaign in Oklahoma City. Considering the speed and performance of ERF, our simulations were able to include the largest extent of Oklahoma City in the published literature. As a result

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

we were able to include measurements in the central business district (similar to most other studies) as well as stations multiple kilometers downwind.

For IOP9 with stable atmospheric conditions, we used an inflow/outflow setup with a base grid spacing of 5.0 m refined down to 2.5 m based on the SF_6 plume concentration value. This setup proved to be very efficient, providing a framework for two sensitivity studies. The first was to determine an appropriate building wall roughness length ($z_{0,\mathrm{w}}$) for the wall-model aimed at broadly parameterizing building materials and unresolved subgrid scale features such as balconies, for example. We determined that a building wall roughness length of 0.1 m resulted in the most accurate model based on statistical metrics suggested by Chang and Hanna (2004) that are widely used in JU2003 modeling studies. Not only did our simulations exhibit excellent model skill, but the $z_{0,\mathrm{w}} = 0.1$ m case resulted in one of the most accurate models in the published literature. Although, the $z_{0,\mathrm{w}} = 0.01$ m case performed nearly as well in many of the statistical measures demonstrating that there is a range of appropriate building wall roughness values. While the optimum value determined here is specific to Oklahoma City, the results are likely transferrable to other mid-rise cities and future work could map roughnesses to building materials and other geometric properties.

The second sensitivity study compared different configurations of ERF and was aimed at quantifying balance between accuracy and computational speed. As expected, we found that the biggest tradeoff when not including adaptive mesh refinement (i.e., using a fixed grid spacing of 5.0 m) was in the resolved TKE. However, simulation results were robust between configurations with grid spacings of 5.0 m and 2.5 m for wind speed and, more importantly, SF_6 concentration, indicating that underpredicting the resolved TKE field is of secondary importance. When comparing fully compressible and anelastic modes at the same grid spacing, accuracy in terms of SF_6 was slightly lower (96% of stations were predicted within a factor of five and 82% within a factor of two whereas 97% were within a factor of five and 84% within a factor of two for fully compressible mode); however, the model was nearly twice as fast.

The prediction of SF_6 concentration for the convective case study resulted in slightly larger errors or less model skill compared to the stable case. The downwind spreading of the plume was well-predicted, but there were measurement stations on the outer edges of the central business district that were underestimated likely due to missing submeso turbulent structures important for meandering of the plume. However, the stations that were underpredicted all measured average concentrations less than 0.1 ppbv (the minimum threshold for inclusion for comparison was 0.01 ppbv). Using a more realistic, multi-scale setup including land surface variability, terrain, and mesoscale parameterizations should improve model skill and is the subject of ongoing work.

We have demonstrated the immersed forcing capability implemented in ERF is one of the most accurate and fastest urban modeling tools that currently exists. The increased performance has allowed us to simulate large domains not feasible with traditional high resolution, turbulence and building resolving urban models. The speed of ERF has provided a framework to easily perform sensitivity studies to further investigate various model parameters and configurations. Our simulation results were robust to both grid spacing, whether using 5.0 m or 2.5 m, and to whether ERF was run in fully compressible or anelastic mode. The performance and accuracy of ERF highlights its efficacy for emergency response or other time sensitive applications.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Code availability. The ERF model is open-source and is released under a modified BSD license. It is available at https://github.com/erf-model/ERF. Documentation of the ERF code can be found at https://erf.readthedocs.io. The immersed forcing capability is supported in the latest release of ERF prior to submission, v26.06, archived at https://doi.org/10.5281/zenodo.21464614 (Wise, 2026).

Data availability. The necessary input files to reproduce the simulations in the present work along with simulation results and a post-processing script are publicly available via Zenodo https://doi.org/10.5281/zenodo.21385466. The JU 2003 Oklahoma City measurements are available from the Defense Threat Reduction Agency database (https://ftd.cbrncore.com/ju2k3/html/index.php) with the appropriate permission and training.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Appendix A: Validation of the Immersed Forcing Method

The immersed forcing method is validated against a wind tunnel experiment of turbulent flow over an isolated, smooth cube of height $H = 0.2$ m. The wind tunnel produced an atmospheric boundary layer 1.8 m deep with a sheared wind profile characterized by a power-law exponent of 0.16 and a free stream (z = 1.0 m) wind speed of 4 m s^{-1}. The log-law profile in the surface layer corresponded to a roughness length of 1 mm. More details for the experiment are found in Snyder and Lawson (1994).

In ERF, we set up the wind tunnel experiment with inflow/outflow streamwise boundary conditions and slip wall lateral and top boundary conditions. The bottom boundary condition uses Monin-Obukhov similarity theory with a roughness length of 1 mm. The domain extent is $L_x \times L_y \times L_z = 3.2$ m $\times$ 1.2 m $\times$ 1.2 m or $11H \times 6H \times 6H$. The cube is located 1.0 m from the inlet and is centered in the spanwise direction.

To spur turbulence, we use a modified version of the cell perturbation method (CPM). We found that at these wind tunnel scales, unphysically large potential temperature perturbations were required to spin up turbulence. The modified CPM used here instead perturbs the vertical velocity field but keeps the same cell structure ($8 \times 8 \times 3$ grid cells). We tuned the vertical velocity perturbations to match turbulence profiles from Snyder and Lawson (1994) and found that random vertical velocity perturbations within the range [-2.0, 2.0] m s^{-1} were required to match the turbulence levels in the wind tunnel experiment. Simulations are run for a total of 50 s with a spin up period of 5 s and analysis done with outputs from the final 45 s.

We conduct a grid resolution sensitivity to examine how the flow behaves at different grid resolutions and for a different number of cells representing the building. The grid spacings used are $\Delta = 5\ \mathrm{mm}, 10\ \mathrm{mm},$ and $20\ \mathrm{mm}$ in all three directions corresponding to 40, 20, and 10 grid cells spanning the cube, respectively. Note that the domain size is held constant. We use a no-slip boundary condition for the cube considering that the cube is smooth, except for a single configuration as discussed later. It is also worth noting that even with the wall-model, we would not expect exact agreement across grid resolutions as the immersed forcing method is a smeared boundary and not an explicit boundary.

In addition to the flow being sensitive to the grid resolution, preliminary simulations also revealed that the average flow structures were sensitive to a number of other input parameters. Namely, these parameters include the turbulence closure (Smagorinsky vs. TKE-1.5). the solver (anelastic vs fully compressible), and the advection scheme (see Shin et al. (2021) for additional analysis of various sensitivities of flow around a cube using an immersed forcing method). We found that the anelastic solver appeared too diffusive and the expected recirculation zones were slightly larger than desired even for the finest grid spacing (not shown). Ultimately, we found the best agreement using fully compressible mode along with the TKE-1.5 closure and $\Delta = 5\ \mathrm{mm}$, which serves as a reference for our suite of simulations.

Figure A1 show the time-averaged streamwise velocity ($\bar{u}$) for different grid resolutions. As the grid resolution and number of cells spanning the cube decreases, the recirculation zone in the lee of the building increases in size. For the $\Delta = 5\ \mathrm{mm}$ and $\Delta = 10\ \mathrm{mm}$ cases, the flow reattaches at the top of the building at approximately $X = 0.2$ m. For the $\Delta = 20\ \mathrm{mm}$ case, the flow does not reattach. Additionally, for the $\Delta = 20\ \mathrm{mm}$ case, the lateral recirculation or side vortices are not well-resolved. While the flow structure is relatively similar between the $\Delta = 5\ \mathrm{mm}$ and $\Delta = 10\ \mathrm{mm}$ cases, there is a jump qualitatively in

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

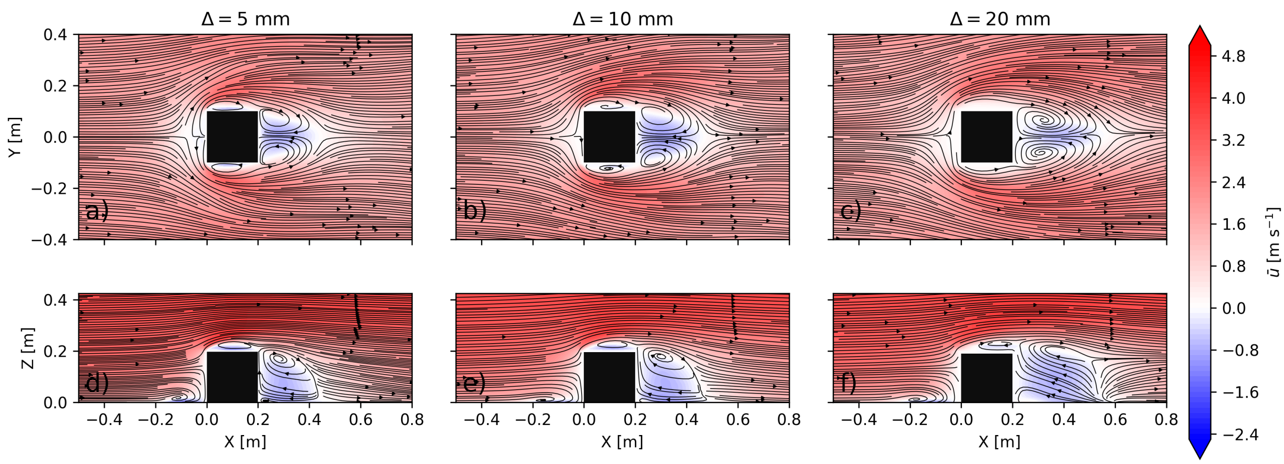

Figure A1. Time-averaged (a-c) horizontal and (d-f) vertical transects of the streamwise velocity for different grid spacings. The horizontal transects are at Z=0.1H or 0.02 m and the vertical transects are taken along the centerline.

the flow behavior between the $\Delta = 10$ mm and $\Delta = 20$ mm cases indicating that it is preferable to have at least 20 grid cells spanning a building. Notably, Shin et al. (2021) also saw a similar shift in flow behavior when having 16 compared to 24 grid cells spanning the cube in their simulations at atmospheric scales (see their Figs. 2 and 3).

To demonstrate convergence, we leverage the mesh refinement capability in ERF. We conduct two additional simulations with base grid spacings of $\Delta = 10$ mm and $\Delta = 20$ mm but include a refined region at $\Delta = 5$ mm covering $(-0.2, -0.22, 0)$ to $(0.5, 0.22, 0.4)$. Figure A2 shows the additional simulations with the refined region and we can qualitatively see that the flow behavior is very similar across all three configurations.

We can quantitatively compare simulation results with measurements taken from the wind tunnel experiment. Figures A3 and A4 show horizontal and vertical profiles of wind speed overlaid with measurements for the cases with and without refinement. As previously mentioned, the best agreement with observations in Fig. A3 (cases without refinement) is for the $\Delta = 5$ mm case. The $\Delta = 10$ mm case behaves similarly as the $\Delta = 5$ mm case with both cases capturing the wake of the cube well in both the horizontal and vertical planes. The $\Delta = 20$ mm case is the outlier with the weakest agreement with observations, most notably for the vertical profiles downwind of the cube. For the flow above the cube, agreement with observations decreases as the grid spacing increases.

When adding refinement for the coarser resolution cases (Fig. A4), the profiles for all cases nearly overlap all with strong agreement with the observations other than the side vortices towards the front of the cube. Even though the turbulence levels are similar across configurations, any differences in the profiles are likely due to differences in the incoming turbulent scales. For example, for the $\Delta = 20$ mm case with refinement, the incoming turbulent structures are larger compared to the $\Delta = 5$ mm case even though the turbulence levels are similar. While all three cases similarly overestimate the size of the side vortices, the

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

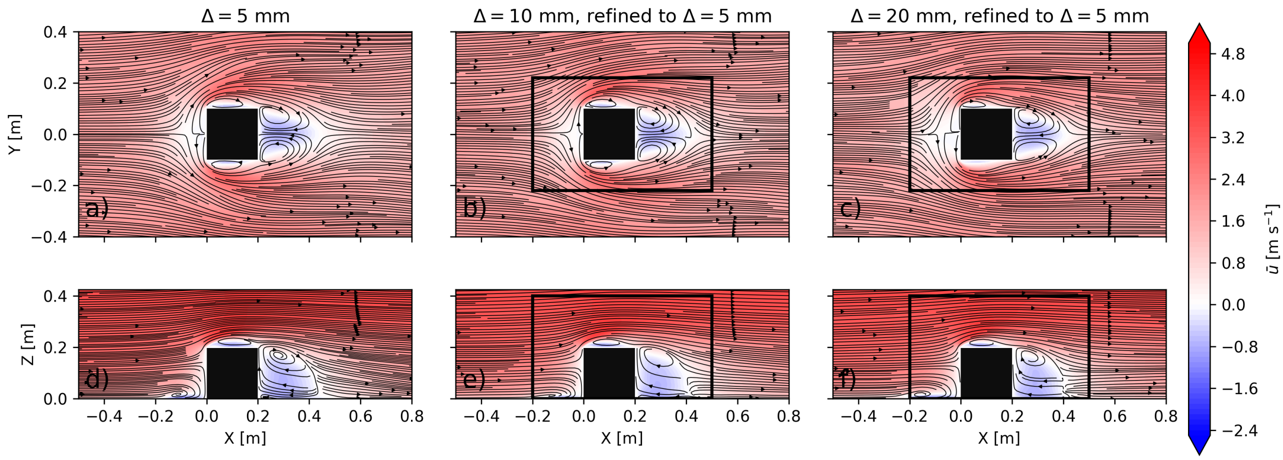

Figure A2. Time-averaged (a-c) horizontal and (d-f) vertical transects of the streamwise velocity (u) for different configurations with mesh refinement. The horizontal transects are at $Z = 0.1H$ or 0.02 m and the vertical transects are taken along the centerline. The black rectangle represents the refined region.

$\Delta = 20$ mm, refined to $\Delta = 5$ mm case shows better agreement with measurements for the $x = 0.04$ m profile indicating that, perhaps, larger turbulent scales are important for accurately capturing the near cube flow for this experiment.

We also conduct an additional simulation at $\Delta = 5$ mm grid spacing where we engage the wall-model and set a the roughness length of the cube walls, $z_{0,\mathrm{w}}$, to 0.1 mm. Since the cube is remarked as smooth in the wind tunnel experiment, there is no quantitative information regarding its roughness. At $\Delta = 5$ mm grid spacing, we swept through a range of $z_{0,\mathrm{w}}$ values and found that 0.1 mm offered the best agreement with observations. While these simulations are high-resolution, there are still not enough grid points to resolve the viscous sublayer around the cube and a wall-model could offer improved flow predictions. Fig. A4 shows that when using the wall-model, there is improved agreement with the measurements closest to the cube.

Author contributions. Writing—original draft preparation and visualization, ASW; writing—review and editing, HG, AML, ASA, DJW, AG, SK, JDM, DJD, and KAL; methodology, software, validation, and formal analysis, ASW, HG, AML, ASA, DW, AG, and SK; conceptualization, ASW, HG, JDM, and KAL. All authors contributed with critical feedback on this research and have read and agreed to the published version of the manuscript.

Competing interests. None.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

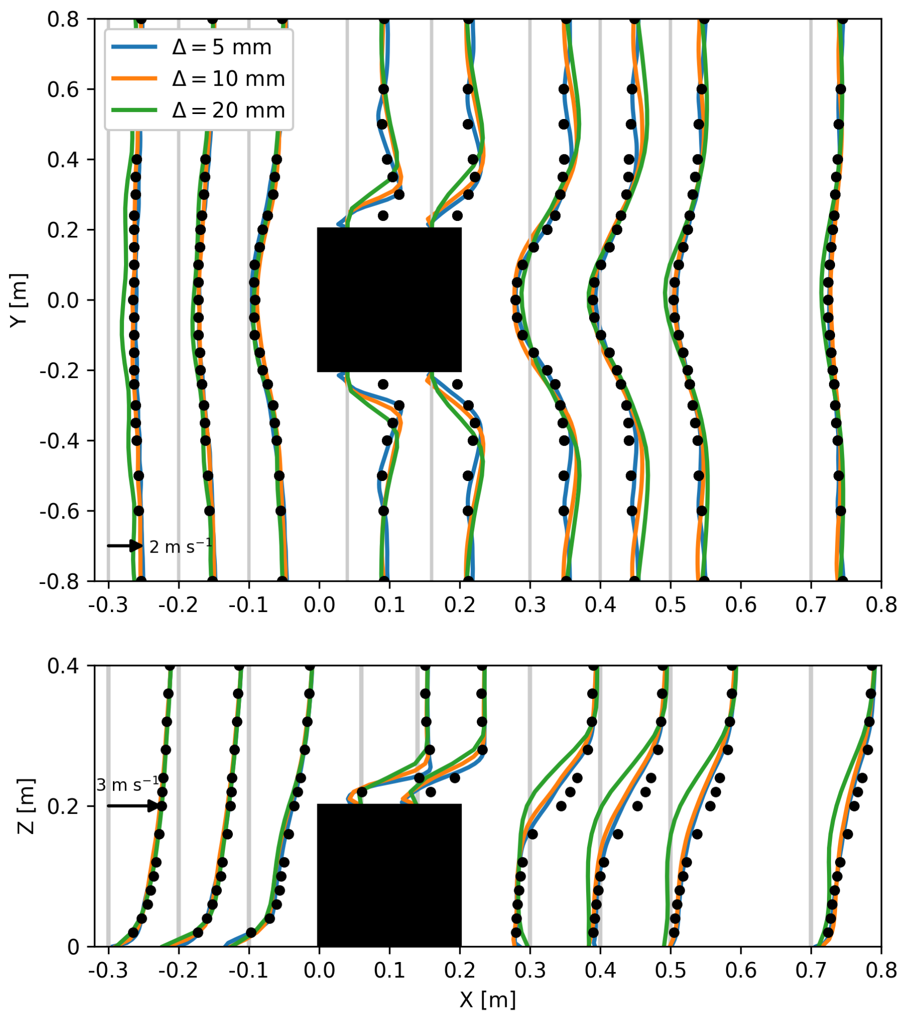

Figure A3. Time-averaged (a) horizontal and (b) vertical profiles of streamwise velocity at various distances for the Δ = 5 mm, 10 mm, and 20 mm cases. Black dots represent measurements from the wind tunnel experiment. The horizontal profiles are at Z=$0.1H$ or 0.02 m and the vertical profiles are along the centerline.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

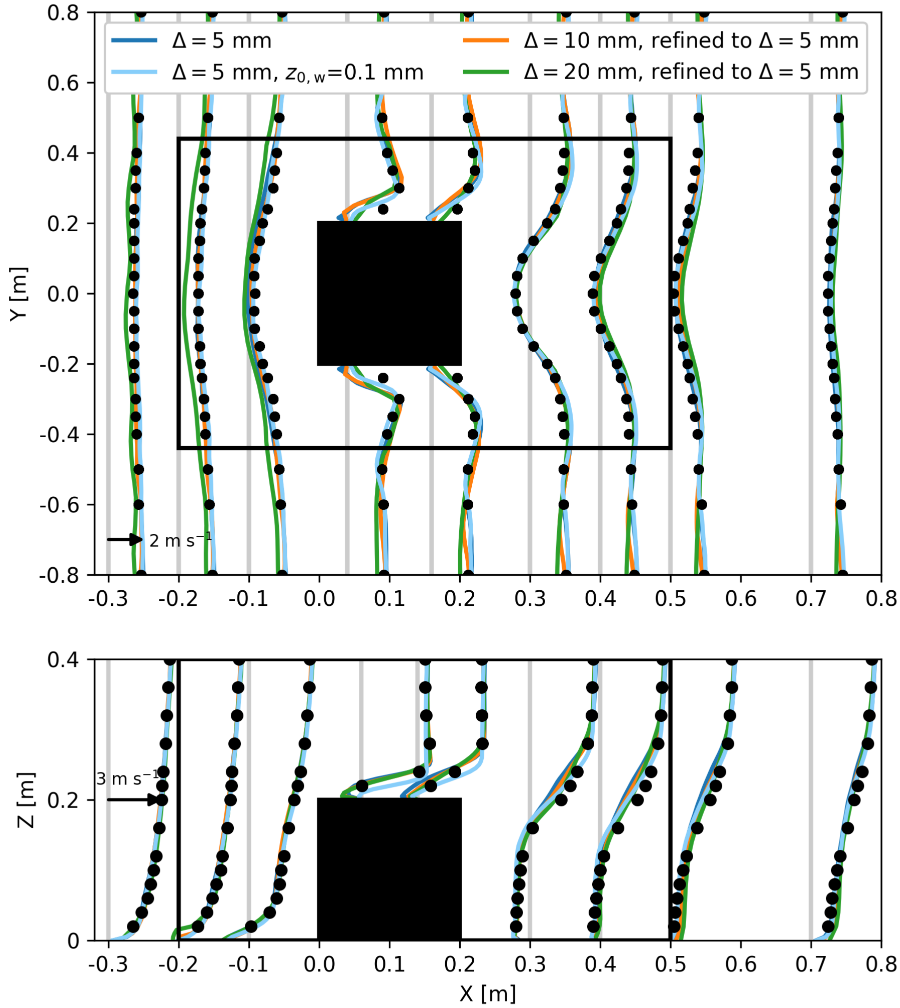

Figure A4. Time-averaged (a) horizontal and (b) vertical profiles of streamwise velocity at various distances for different configurations (with mesh refinement). Black dots represent measurements from the wind tunnel experiment. The horizontal profiles are at Z=0.1H or 0.02 m and the vertical profiles are along the centerline. The black rectangle represents the refined region..

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Acknowledgements. ASW's, DJW's, AG's, SK's, JDM's, DJG's, and KAL's contributions were performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and were supported by the LLNL-LDRD Program under Project No. 24-SI-001. LLNL-JRNL-2012615-DRAFT. The contributions of ASA and AML were performed under the auspices of the U.S. Department of Energy at the Lawrence Berkeley National Laboratory under contract DE-AC02-05CH11231. The contribution of HG was performed under the auspices of the National Laboratory of the Rockies, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308.

References

Allwine, K. J. and Flaherty, J. E.: Joint Urban 2003: Study Overview And Instrument Locations, Tech. rep., Pacific Northwest National Laboratory (PNNL), Richland, WA (US), https://doi.org/10.2172/890732, 2006.

Almgren, A., Lattanzi, A., Haque, R., Jha, P., Kosovic, B., Mirocha, J., Perry, B., Quon, E., Sanders, M., Wiersema, D., Willcox, D., Yuan, X., and Zhang, W.: ERF: Energy Research and Forecasting, Journal of Open Source Software, 8, 5202, https://doi.org/10.21105/joss.05202, 2023.

Anderson, W.: An immersed boundary method wall model for high-Reynolds-number channel flow over complex topography, International Journal for Numerical Methods in Fluids, 71, 1588–1608, https://doi.org/https://doi.org/10.1002/fld.3727, 2013.

Arthur, R. S., Mirocha, J. D., Marjanovic, N., Hirth, B. D., Schroeder, J. L., Wharton, S., and Chow, F. K.: Multi-scale simulation of wind farm performance during a frontal passage, Atmosphere, 11, 245, 2020.

Bao, J., Lundquist, K. A., and Chow, F. K.: Comparison of three different implementations of the immersed boundary method in WRF (WRF-IBM), in: 22nd Symposium on Boundary Layers and Turbulence, Salt Lake City, Utah, USA, session 4B: Results from the Mountain Terrain Atmospheric Modeling and Observations (MATERHORN) Project, 2016.

Chan, S. T. and Leach, M. J.: A Validation of FEM3MP with Joint Urban 2003 Data, Journal of Applied Meteorology and Climatology, 46, 2127 – 2146, https://doi.org/10.1175/2006JAMC1321.1, 2007.

Chang, J. C. and Hanna, S. R.: Air quality model performance evaluation, Meteorology and Atmospheric Physics, 87, 167–196, https://doi.org/10.1007/s00703-003-0070-7, 2004.

Clawson, K. L., Carter, R. G., Lacroix, D. J., Biltoft, C. A., Hukari, N. F., Johnson, R. C., Rich, J. D., Beard, S. A., and Strong, T. W.: Joint urban 2003 (JU03) SF_6 atmospheric tracer field tests, Tech. rep., Air Resources Laboratory (U.S.), https://repository.library.noaa.gov/view/noaa/12573, 2005.

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

García-Sánchez, C., Van Tendeloo, G., and Gorlé, C.: Quantifying inflow uncertainties in RANS simulations of urban pollutant dispersion, Atmospheric Environment, 161, 263–273, https://doi.org/https://doi.org/10.1016/j.atmosenv.2017.04.019, 2017.

García-Sánchez, C., van Beeck, J., and Gorlé, C.: Predictive large eddy simulations for urban flows: Challenges and opportunities, Building and Environment, 139, 146–156, https://doi.org/https://doi.org/10.1016/j.buildenv.2018.05.007, 2018.

Gorlé, C., van Beeck, J., and Rambaud, P.: Dispersion in the Wake of a Rectangular Building: Validation of Two Reynolds-Averaged Navier–Stokes Modelling Approaches, Boundary-Layer Meteorology, 137, 115–133, https://doi.org/10.1007/s10546-010-9521-0, 2010.

Iaccarino, G. and Verzicco, R.: Immersed boundary technique for turbulent flow simulations, Applied Mechanics Reviews, 56, 331–347, https://doi.org/10.1115/1.1563627, 2003.

Jähn, M., Muñoz Esparza, D., Chouza, F., Reitebuch, O., Knoth, O., Haarig, M., and Ansmann, A.: Investigations of boundary layer structure, cloud characteristics and vertical mixing of aerosols at Barbados with large eddy simulations, Atmospheric Chemistry and Physics, 16, 651–674, https://doi.org/10.5194/acp-16-651-2016, 2016.

Kang, S., Almgren, A. S., Natarajan, M., Lattanzi, A. M., Mirocha, J. D., Lundquist, K., Musser, J., and Zhang, W.: An Embedded Boundary Scheme for Terrain in Staggered Meshes, In preparation for the Journal of Computational Physics, 2025.

Lattanzi, A., Almgren, A., Quon, E., Natarajan, M., Kosovic, B., Mirocha, J., Perry, B., Wiersema, D., Willcox, D., Yuan, X., and Zhang, W.: ERF: Energy Research and Forecasting Model, Journal of Advances in Modeling Earth Systems, 17, e2024MS004 884, https://doi.org/https://doi.org/10.1029/2024MS004884, e2024MS004884 2024MS004884, 2025.

Letzel, M., Krane, M., and Raasch, S.: High resolution urban large-eddy simulation studies from street canyon to neighbourhood scale, Atmospheric Environment, 42, 8770–8784, https://doi.org/10.1016/j.atmosenv.2008.08.001, 2008.

Lundquist, J. K. and Chan, S. T.: Consequences of Urban Stability Conditions for Computational Fluid Dynamics Simulations of Urban Dispersion, Journal of Applied Meteorology and Climatology, 46, 1080 – 1097, https://doi.org/10.1175/JAM2514.1, 2007.

Lundquist, K. A., Chow, F. K., and Lundquist, J. K.: An Immersed Boundary Method for the Weather Research and Forecasting Model, Monthly Weather Review, 138, 796 – 817, https://doi.org/10.1175/2009MWR2990.1, 2010.

Lundquist, K. A., Chow, F. K., and Lundquist, J. K.: An Immersed Boundary Method Enabling Large-Eddy Simulations of Flow over Complex Terrain in the WRF Model, Monthly Weather Review, 140, 3936 – 3955, https://doi.org/10.1175/MWR-D-11-00311.1, 2012.

Mittal, R. and Iaccarino, G.: IMMERSED BOUNDARY METHODS, Annual Review of Fluid Mechanics, 37, 239–261, https://doi.org/https://doi.org/10.1146/annurev.fluid.37.061903.175743, 2005.

Monin, A. S. and Obukhov, A. M.: Basic laws of turbulent mixing in the surface layer of the atmosphere, Contrib. Geophys. Inst. Acad. Sci. USSR, 151, e187, 1954.

Muñoz-Esparza, D. and Kosović, B.: Generation of inflow turbulence in large-eddy simulations of nonneutral atmospheric boundary layers with the cell perturbation method, Monthly Weather Review, 146, 1889–1909, 2018.

Muñoz-Esparza, D., Kosović, B., Mirocha, J., and van Beeck, J.: Bridging the transition from mesoscale to microscale turbulence in numerical weather prediction models, Boundary-Layer Meteorology, 153, 409–440, 2014.

Muñoz-Esparza, D., Kosović, B., Van Beeck, J., and Mirocha, J.: A stochastic perturbation method to generate inflow turbulence in large-eddy simulation models: Application to neutrally stratified atmospheric boundary layers, Physics of Fluids, 27, 035 102, 2015.

Muñoz-Esparza, D., Sauer, J. A., Shin, H. H., Sharman, R., Kosović, B., Meech, S., García-Sánchez, C., Steiner, M., Knievel, J., Pinto, J., and Swerdlin, S.: Inclusion of Building-Resolving Capabilities Into the FastEddy® GPU-LES Model Using an Immersed Body Force Method, Journal of Advances in Modeling Earth Systems, 12, e2020MS002 141, https://doi.org/https://doi.org/10.1029/2020MS002141, e2020MS002141 10.1029/2020MS002141, 2020.

Muñoz-Esparza, D., Shin, H. H., Sauer, J. A., Steiner, M., Hawbecker, P., Boehnert, J., Pinto, J. O., Kosović, B., and Sharman, R. D.: Efficient Graphics Processing Unit Modeling of Street-Scale Weather Effects in Support of Aerial Operations in the Urban Environment, AGU Advances, 2, e2021AV000 432, https://doi.org/https://doi.org/10.1029/2021AV000432, e2021AV000432 2021AV000432, 2021.

Natarajan, M., Grout, R., Zhang, W., and Day, M.: A moving embedded boundary approach for the compressible Navier-Stokes equations in a block-structured adaptive refinement framework, Journal of Computational Physics, 465, 111 315, https://doi.org/https://doi.org/10.1016/j.jcp.2022.111315, 2022.

Neophytou, M., Gowardhan, A., and Brown, M.: An inter-comparison of three urban wind models using Oklahoma City Joint Urban 2003 wind field measurements, Journal of Wind Engineering and Industrial Aerodynamics, 99, 357–368, https://doi.org/https://doi.org/10.1016/j.jweia.2011.01.010, the Fifth International Symposium on Computational Wind Engineering, 2011.

Norman, M., Gopalakrishnan Meena, M., Gottiparthi, K., Koukpaizan, N., and Nichols, S.: PortUrb: A Performance Portable, High-Order, Moist Atmospheric Large Eddy Simulation Model with Variable-Friction Immersed Boundaries, EGUsphere, 2025, 1–36, https://doi.org/10.5194/egusphere-2025-1135, 2025.

https://doi.org/10.5194/egusphere-2026-4030
Preprint. Discussion started: 28 July 2026

Peskin, C. S.: Flow patterns around heart valves: A numerical method, Journal of Computational Physics, 10, 252–271, https://doi.org/https://doi.org/10.1016/0021-9991(72)90065-4, 1972.

Sauer, J. A. and Muñoz-Esparza, D.: The FastEddy® Resident-GPU Accelerated Large-Eddy Simulation Framework: Model Formulation, Dynamical-Core Validation and Performance Benchmarks, Journal of Advances in Modeling Earth Systems, 12, e2020MS002 100, https://doi.org/https://doi.org/10.1029/2020MS002100, e2020MS002100 10.1029/2020MS002100, 2020.

Senocak, I., Ackerman, A. S., Stevens, D. E., and Mansour, N. N.: Topography modeling in atmospheric flows using the immersed boundary method, Research brief, Stanford University, Center for Turbulence Research, https://web.stanford.edu/group/ctr/ResBriefs/2004/senocak2.pdf, 2004.

Shin, H. H., Muñoz-Esparza, D., Sauer, J. A., and Steiner, M.: Large-Eddy Simulations of Stability-Varying Atmospheric Boundary Layer Flow over Isolated Buildings, Journal of the Atmospheric Sciences, 78, 1487 – 1501, https://doi.org/10.1175/JAS-D-20-0160.1, 2021.

Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Liu, Z., Berner, J., Wang, W., Powers, J. G., Duda, M. G., Barker, D. M., and Huang, X.-Y.: A Description of the Advanced Research WRF Model Version 4, 2021.

Smagorinsky, J.: GENERAL CIRCULATION EXPERIMENTS WITH THE PRIMITIVE EQUATIONS: I. THE BASIC EXPERIMENT, Monthly Weather Review, 91, 99 – 164, https://doi.org/10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2, 1963.

Snyder, W. H. and Lawson, Jr, R. E.: Wind-tunnel measurements of flow fields in the vicinity of buildings, American Meteorological Society, Boston, MA (United States), https://www.osti.gov/biblio/28752, 1994.

Welch, B. M., Horel, J. D., and Sauer, J. A.: Application of the NCAR FastEddy® Microscale Model to a Lake Breeze Front, Atmosphere, 15, https://doi.org/10.3390/atmos15070809, 2024.

Wiersema, D. J., Lundquist, K. A., and Chow, F. K.: Mesoscale to Microscale Simulations over Complex Terrain with the Immersed Boundary Method in the Weather Research and Forecasting Model, Monthly Weather Review, 148, 577 – 595, https://doi.org/10.1175/MWR-D-19-0071.1, 2020.

Wiersema, D. J., Lundquist, K. A., Mirocha, J. D., and Katopodes Chow, F.: Evaluation of Turbulence and Dispersion in Multiscale Atmospheric Simulations over Complex Urban Terrain during the Joint Urban 2003 Field Campaign, Monthly Weather Review, 150, https://doi.org/10.1175/MWR-D-22-0056.1, 2022.

Wise, A. S.: ERF, https://doi.org/10.5281/zenodo.21464614, 2026.

Wise, A. S., Arthur, R. S., Abraham, A., Wharton, S., Krishnamurthy, R., Newsom, R., Hirth, B., Schroeder, J., Moriarty, P., and Chow, F. K.: Large-eddy simulation of an atmospheric bore and associated gravity wave effects on wind farm performance in the southern Great Plains, Wind Energy Science, 10, 1007–1032, https://doi.org/10.5194/wes-10-1007-2025, 2025.

Zhang, W., Almgren, A., Beckner, V., Bell, J., Blaschke, J., Chan, C., Day, M., Friesen, B., Gott, K., Graves, D., Katz, M. P., Myers, A., Nguyen, T., Nonaka, A., Rosso, M., Williams, S., and Zingale, M.: AMReX: a framework for block-structured adaptive mesh refinement, Journal of Open Source Software, 4, 1370, https://doi.org/10.21105/joss.01370, 2019.

Zhang, W., Myers, A., Gott, K., Almgren, A., and Bell, J.: AMReX: Block-structured adaptive mesh refinement for multiphysics applications, The International Journal of High Performance Computing Applications, 35, 508–526, https://doi.org/10.1177/10943420211022811, 2021.