ARTICLE

Comparing Two Implementations of a Micromixing Model. Part I: Wall Shear-Layer Flow

John V. Postma · John D. Wilson · Eugene Yee

Abstract A Lagrangian stochastic (LS) micromixing model is used for estimating concentration fluctuations in plumes of a passive, non-reactive tracer dispersing from elevated and ground-level compact sources into a neutral wall shear-layer flow. SPMMM (for sequential particle micromixing model) implements the familiar IECM (interaction by exchange with the conditional mean) micromixing scheme. The parametrization of the scalar micromixing time scale is identical to that proposed in a previously reported LS–IECM model (Cassiani et al., Atmos Environ 39:1457–1469, 2005a). However, while SPMMM is mathematically equivalent to the previously reported model, it differs in its numerical implementation: SPMMM releases N independent particles sequentially, whereas the previously reported model releases N independent particles simultaneously. In both implementations, the trajectories of the N particles are governed by single-point velocity statistics. The sequential particle implementation is computationally efficient, but cannot be applied to the case of reacting species. Results from both implementations are compared to experimental wind-tunnel dispersion data and to each other.

1 Introduction

We build on earlier efforts by others to model the dispersion of a passive scalar using the 'interaction by exchange with the conditional mean' (IECM) micromixing model (Fox 1996; Pope 1998) coupled to a single-particle Lagrangian stochastic (LS; see Wilson and Sawford 1996 for a review) trajectory model. Each particle has a position, a velocity, and a concentration variable assigned to it. The LS trajectory model governs the evolution of the position and

J. V. Postma (✉) · J. D. Wilson
Department of Earth and Atmospheric Sciences, University of Alberta, Edmonton, AB T6G 2E3, Canada
e-mail: jpostma@ualberta.ca

E. Yee
Defence R&D Canada – Suffield, P.O. Box 4000, Medicine Hat, AB T1A 8K6, Canada

the velocity, while the IECM micromixing model governs the evolution of the concentration. Combined, the LS and IECM models allow for the prediction of the higher-order moments of the scalar concentration field.

Heuristically, conditional scalar mixing can be thought of as occurring between all fluid elements that occupy the same eddy. Given their similar velocities and their proximity to one another, these fluid elements are more likely to remain together for times comparable to the Lagrangian integral time and, therefore, are more likely to mix (Fox 1996; Pope 1998). It is important for a mixing model to properly represent the bulk motions (such as meandering) and the in-plume, variance-dissipating motions (such as entrainment and viscous dissipation). By conditioning on velocity, the IECM model segregates out the relevant length and velocity scales that contribute to the scalar dissipation (Sawford 2004a). Bulk motions of the plume are captured by the LS model.

Sawford (2004b) showed that the IECM model can be related to a meandering plume model (Gifford 1959) and applied the technique successfully to simulate concentration statistics due to a continuous line source in grid turbulence. Recent applications of IECM models to atmospheric flows include simulations of dispersion of passive (i.e. non-buoyant), non-reactive (i.e. no chemistry) scalars within the neutral boundary layer (Cassiani et al. 2005a), within the convective boundary layer (Cassiani et al. 2005b; Luhar and Sawford 2005a,b) and within a canopy layer (Cassiani et al. 2005c, 2007).

There are two approaches to LS–IECM modelling to be considered in this article, both of which involve the release of N particles to compute the required conditional mean concentrations, but differ in how the particle trajectories are treated. One approach computes the trajectories of the N particles *simultaneously* (meaning that the particles move together, albeit on independent trajectories), calculating the conditional mean concentrations at each timestep (Cassiani et al. 2005a for example). The other approach computes the trajectories of the N particles *sequentially* and pre-calculates the conditional mean concentrations (Luhar and Sawford 2005a for example). In both approaches, the trajectories of the particles are calculated using single-point velocity statistics. The model presented herein utilizes the latter approach. Regardless of the technique, the IECM model must not alter first-order statistics such as the mean concentration. In other words, the mean concentration field produced by the LS–IECM model must be consistent with the mean concentration field produced by the underlying LS model.

In theory, the IECM model can predict all moments of the scalar concentration field. In practice however, this is limited by the availability of the computational resources. The use of sequential particle trajectories results in computational simplicity, but comes at the cost of prohibiting the calculation of the concentration field of reactive species, which could be computed if simultaneous particle trajectories were used. Since we only consider non-reactive species in this article, this is not a cause for concern.

While the use of simultaneous trajectories allows for the incorporation of chemistry, the computer code of such a model is much more difficult to parallelize since the particles are interactive in that at each timestep the conditional mean concentration must be calculated based on the particles that occupy a particular region of space. If the particles within a region are being processed by different computer processors, then there will be a large computational communication overhead since at each timestep, the processors will have to pause, share particle data, then continue. That said, modern parallel algorithms (e.g. Rembold et al. 2008) frequently utilize domain decomposition which keeps the computer processor communication to a minimum, and can result in approximately 80% parallel efficiency.

However, the sequential particle trajectory framework allows for trivial parallelization and an almost linear increase in performance—the time required to run the sequential particle

trajectory LS–IECM on N_p computer processors is a fraction ($\sim 1/N_p$) of the time required to run it on one computer processor. The only time the processors must communicate is when sharing the plume extent data at the beginning of the simulation. Decreased simulation times would be invaluable in situations where simulation results are needed as soon as possible, such as in the case of an emergency where evacuations may be required: an accidental or intentional release of chemical, biological, nuclear or radiological agents into a densely populated area, for example.

2 Model Equations

Under the assumption that the velocity and position of a fluid element are jointly a continuous Markov process, and assuming validity of the Kolmogorov similarity relationship for the Lagrangian second-order structure function, it follows that the motion of N independent tracer particles is governed by the following stochastic equations:

$$dU_i' = a_i(\boldsymbol{X}, \boldsymbol{U}'t)dt + b_{ij}(\boldsymbol{X}, \boldsymbol{U}', t)d\xi_j(t), \tag{1}$$

$$dX_i = \left(\langle u_i \rangle + U_i'\right) dt, \tag{2}$$

where X_i is the particle position, U_i' is the Lagrangian velocity fluctuation relative to the Eulerian mean (viz., $U_i' = u_i - \langle u_i \rangle$), dt is a small timestep and $d\xi_j(t)$ represents an incremental Wiener process with zero mean and variance dt. On the right-hand side of Eq. 1, we have the deterministic term $a_i dt$, and the stochastic diffusion term $b_{ij} d\xi_j$.

Before determining the deterministic coefficient in Eq. 1, the form of the probability density function (PDF) of the Eulerian velocity fluctuations must be specified. For many cases of practical importance and interest, a stationary Gaussian form is sufficient:

$$g_a(\boldsymbol{x}, \boldsymbol{u}') = \frac{[\det(R^{-1})]^{1/2}}{(2\pi)^{3/2}} \exp\left(-\frac{1}{2} u_i' R_{ij}^{-1} u_j'\right), \tag{3}$$

where $R_{ij} = \langle u_i' u_j' \rangle$ is the Reynolds stress tensor, and R_{ij}^{-1} is its inverse. Using the well-mixed condition (Thomson 1987), the 'simplest' (but non-unique) three-dimensional solution for a_i is obtained:

$$a_i = T_i^{(0)} + T_{ij}^{(1)} U_j' + T_{ijk}^{(2)} U_j' U_k', \tag{4}$$

with

$$T_i^{(0)} \equiv \frac{1}{2} \frac{\partial R_{i\ell}}{\partial x_\ell}, \tag{5}$$

$$T_{ij}^{(1)} \equiv -\frac{1}{2}(C_0\varepsilon) R_{ij}^{-1} + \frac{1}{2} R_{j\ell}^{-1} \frac{\partial R_{i\ell}}{\partial x_k} \langle u_k \rangle,$$

$$= -\frac{1}{2}(C_0\varepsilon) R_{ij}^{-1} + T_{ijk}^{(2)} \langle u_k \rangle, \tag{6}$$

$$T_{ijk}^{(2)} \equiv \frac{1}{2} R_{j\ell}^{-1} \frac{\partial R_{i\ell}}{\partial x_k}, \tag{7}$$

where C_0 is the universal[1] Kolmogorov constant, ε is the turbulent kinetic energy (TKE) dissipation rate, and $\langle u_k \rangle$ is the mean Eulerian velocity. Consistency of the LS model and Kolmogorov's theory of local isotropy (Monin and Yaglom 1975) is assured by the following specification for the coefficient of the stochastic term in Eq. 1,

$$b_{ij} = \delta_{ij}(C_0\varepsilon)^{1/2}. \tag{8}$$

Together, Eqs. 1 and 2 can only provide information about first-order statistics such as the mean concentration. A micromixing model may be used to calculate the higher-order moments of the concentration field, in which case the compound Markovian state variable is enlarged to (U_i', X_i, ϕ), where ϕ represents the scalar concentration.

The rate of change in concentration as calculated by the IECM micromixing model is

$$\frac{d\phi}{dt} = -\frac{1}{t_m}(\phi - \langle \phi | \boldsymbol{u} \rangle), \tag{9}$$

where t_m is the scalar micromixing time scale, and $\langle \phi | \boldsymbol{u} \rangle$ is the mean scalar concentration conditioned on the local velocity (also called the conditional mean concentration).

The parametrization of the micromixing time scale used here was originally proposed by Cassiani et al. (2005a; hereafter referred to as CASS). For completeness, we summarize their parametrization for non-homogeneous, non-isotropic turbulence (local equilibrium and local isotropy are assumed). For short and medium times we assume that

$$t_m = \mu \left(\frac{\sigma_r^2}{\sigma_{U_r}^2} \right)^{1/2}. \tag{10}$$

The micromixing constant, μ, is empirically determined and depends upon the type of turbulence, the source configuration, and the stage of development of the plume. It is treated as a 'tuning' parameter. The instantaneous plume width is denoted by σ_r. The variance of the Lagrangian relative velocity fluctuations $\sigma_{U_r}^2$ is in general very difficult to calculate. However, it represents the fraction of the energy responsible for expansion of the plume about its instantaneous centreline. By constructing a ratio of the instantaneous plume width to the length scale of the largest eddies L, Franzese (2003) and Cassiani et al. (2005a) modelled $\sigma_{U_r}^2$ as

$$\sigma_{U_r}^2 = \sigma^2 \left(\frac{\sigma_r}{L} \right)^{2/3}, \tag{11}$$

with

$$\sigma^2 = \frac{\sigma_u^2 + \sigma_v^2 + \sigma_w^2}{3} = \frac{2k}{3}, \tag{12}$$

where the streamwise, spanwise and vertical velocity variances are represented by σ_u^2, σ_v^2 and σ_w^2 respectively, and k is the TKE. The length scale of the most energetic eddies is calculated as

$$L = \frac{(3\sigma^2/2)^{3/2}}{\varepsilon}, \tag{13}$$

---

[1] While considered as being a universal constant, its value has yet to be determined with great certainty and values from 2 to 7 are considered acceptable, depending on the particulars of the Lagrangian model (e.g. first-order or second-order), the prevailing context (e.g. degree of complexity of the turbulent flow), and the focus (e.g. mean concentration, or higher-order statistics).

and wherever $\sigma_r > L$, the constraint $\sigma_{U_r}^2 = \sigma^2$ is imposed. The instantaneous plume spread is modelled as

$$\sigma_r^2 = \frac{d_r^2}{1 + (d_r^2 - \sigma_0^2)/(\sigma_0^2 + 2\sigma^2 T_L t)}, \tag{14}$$

where $T_L \equiv 2\sigma^2/(C_0\epsilon)$ is the Lagrangian time scale and d_r is the root-mean-square separation between particle pairs in the instantaneous plume. This separation is calculated using Richardson's law as

$$d_r^2 = C_r\epsilon(t + t_0)^3, \tag{15}$$

where C_r is the Richardson constant. This equation is discretized by invoking linearization, viz.

$$d_r^2(t + \Delta t) = d_r^2(t) + 3C_r\varepsilon(t + t_0)^2\Delta t, \tag{16}$$

and enforcing the constraint $\sigma_r^2(t + \Delta t) \geq \sigma_r^2(t)$. The constant $t_0 = t_s/C_r^{1/3}$ (where $t_s = (\sigma_0^2/\varepsilon)^{1/3}$ is the characteristic time scale of the source and σ_0 is the initial source distribution) ensures that $t_m \to t_s$ as $t \to 0$.

2.1 Description and Implementation of the Model

The micromixing model used in this work is called SPMMM, which is an acronym for Sequential Particle MicroMixing Model. In contrast to the simultaneous release of N particles to determine the concentration statistics (as used by CASS), SPMMM releases a single particle at a time. The particle samples a pre-calculated conditional mean concentration field that is provided by a program called MEANS.

The velocity statistics and TKE dissipation rate used for driving MEANS and SPMMM are supplied in discretized form. These velocity statistics can be obtained from discretizations of analytical equations, from interpolations of experimentally measured flow fields, or provided by another model. Perfectly reflective boundary conditions were used on the upstream face and at the top and bottom faces of the simulation domain. Periodic boundary conditions were used on the lateral faces. However, for the compact source simulations described below, care was taken to ensure that the plume did not impinge upon the lateral or top faces of the simulation domain.

2.2 The MEANS Pre-Calculation Program

MEANS is a numerical implementation of Eqs. 1 and 2, and its purpose is to pre-calculate the conditional mean concentration field $\langle\phi|\boldsymbol{u}\rangle$ needed by SPMMM. The first step in MEANS is to determine the extent of the plume in both the spatial (x, y, z) and velocity (u, v, w) domains. To maximize the spatial resolution of the model, the calculation of the conditional mean concentrations is carried out on a dynamic grid that encompasses the plume, as was done in CASS. The spatial extent of the plume is determined by recording the trajectories of a small sub-ensemble of particles, released from the source, and allowed to propagate out of the simulation domain through $x > x_{max}$ (the axes are aligned such that x is the direction of mean motion). The source was chosen to be a two-dimensional isotropic Gaussian distribution in the yz-plane centred on the location $(x_s, y_s, z_s) = (0, y_s, z_s)$ and with a variance of σ_0^2. To account for uncertainties in the exact behaviour of the flow around the source, the initial source distribution was modelled as $\sigma_0 = \mu_s d_s$, where the source constant μ_s is a

tunable parameter and d_s is the diameter of the source. This approach was also used in the CASS model.

The extent of the velocity domain is determined from the driving velocity statistics. For each position in the spatial domain, the mean velocities and the velocity variances are known. From these, the global velocity extrema can be determined. With the extents of the spatial and velocity domains known, the spatial domain is discretized into N_x streamwise bins, N_y spanwise velocity bins and N_z vertical bins. Likewise, the velocity domain is discretized into N_u streamwise velocity bins, N_v spanwise velocity bins and N_w vertical velocity bins. Upper-case subscripts are used to denote this discretization: $(x, y, z) \Rightarrow (x_I, y_J, z_K)$ and $(u, v, w) \Rightarrow (u_L, v_M, w_N)$. Since the motion of the particle is described by six dimensions, the position and velocity are indexed as $(x_I, y_J, z_K, u_L, v_M, w_N)$. Equations 1 and 2 are discretized with a forward difference scheme. The coefficients of Eq. 1 are calculated using the discretized velocity statistics discussed above. Consequently, a_i is constant while the particle occupies a particular position–velocity bin, and b_{ij} is constant while the particle occupies the particular position bin.

To determine the conditional mean concentrations, N_ϕ particles are released from the source and tracked downstream until they leave the spatial domain. The timestep was chosen to be a fraction of the Lagrangian integral time scale

$$\Delta t = \mu_t \min[T_{L_u}, T_{L_v}, T_{L_w}], \tag{17}$$

where $\mu_t \ll 1$ is the timestep constant and the Lagrangian integral time scales associated with the streamwise, spanwise and vertical velocities are calculated as

$$T_{L_u} = \frac{2\sigma_u^2}{C_0\varepsilon}, \tag{18a}$$

$$T_{L_v} = \frac{2\sigma_v^2}{C_0\varepsilon}, \tag{18b}$$

$$T_{L_w} = \frac{2\sigma_w^2}{C_0\varepsilon}. \tag{18c}$$

By accurately accounting for how long each particle spends in a position–velocity bin, MEANS accumulates conditional residence times $t_r^v = t_r^v(x_I, y_J, z_K, u_L, v_M, w_N)$ for each bin in the simulation domain. At the beginning of each timestep, the position and velocity bins of the particle are known. Based on the current velocity of the particle, MEANS calculates how far the particle will step in each spatial direction and checks if the particle will remain in the same position bin, or move into a new one. If the particle is to remain within the same position bin, the full timestep is accumulated in the current position–velocity bin of the particle. If the particle is to change position bins, then the timestep is decomposed into the required number of sub-timesteps (two if the particle is to visit two position bins, three if the particle is to visit three position bins, and so on), and the position of the particle is updated in the same number of substeps, propagating to the nearest position bin boundary at each sub-timestep. At the position bin boundary, the sub-timestep is accumulated into the appropriate position–velocity bin, then the position bin is updated, and then the remainder of the timestep is carried out. For the majority of timesteps, the velocity bin of the particle is constant, since the velocity of the particle is updated at the end of the timestep. The exception is when a particle is reflected off a boundary, in which case the velocity (and velocity bin) of the particle changes during the timestep. MEANS treats reflection boundaries as bin boundaries, and utilizes the sub-timestep process to properly accumulate the conditional residence times in this situation.

From these conditional residence times t_r^v, the conditional mean concentration field (for a given source configuration and flow) can be determined. The conditional mean concentration in bin $(x_I, y_J, z_K, u_L, v_M, w_N)$ is computed as

$$\langle \phi | \boldsymbol{u} \rangle = \langle \phi | \boldsymbol{u} \rangle (x_I, y_J, z_K, u_L, v_M, w_N) = \frac{Q \, t_r^v}{\mathcal{V} N_\phi^v}, \tag{19}$$

where $N_\phi^v = N_\phi^v(x_I, y_J, z_K, u_L, v_M, w_N)$ is the number of particles during the simulation that visit position–velocity bin $(x_I, y_J, z_K, u_L, v_M, w_N)$, Q is the source strength and $\mathcal{V} = \mathcal{V}(x_I, y_J, z_K)$ is the volume of a position bin. Owing to computer memory limitations, N_ϕ^v is not computed from the stochastic paths. Instead, it is estimated deterministically from the PDF of the driving velocity statistics $f_{\boldsymbol{u}}$ as

$$N_\phi^v = N_\phi f_{\boldsymbol{u}} \Delta u \Delta v \Delta w, \tag{20}$$

giving

$$\langle \phi | \boldsymbol{u} \rangle = \frac{Q \, t_r^v}{\mathcal{V} N_\phi f_{\boldsymbol{u}} \Delta u \Delta v \Delta w}. \tag{21}$$

The result of the MEANS pre-calculation program is a file containing the conditional mean concentrations to be used by SPMMM.

2.3 The SPMMM Micromixing Model

SPMMM is a numerical implementation of Eqs. 1, 2 and 9. Once the conditional mean concentration field has been computed by MEANS, the next step is to use the SPMMM micromixing model to simulate mixing to provide information on the higher-order moments of the concentration field. The first step is to compute the micromixing time scales on the spatial grid. Again a small sub-ensemble of particles is released from the source region sequentially and tracked downstream. As the particles pass through bins in the spatial domain, t_m is computed according to Eqs. 10–16. Since the particles will all follow different trajectories, the computed value of t_m may differ greatly from particle to particle in a given position bin. Once all the particles in the sub-ensemble have exited the spatial domain through $x > x_{\max}$, the mean micromixing time scale in bin (x_I, y_J, z_K) is calculated for use in the next stage of the simulation. If the micromixing time scale is larger than the turbulence time scale $\tau = k/\varepsilon$, then t_m is reset to τ, similar to that done in Cassiani et al. (2007). Furthermore, for regions outside the plume, mixing still occurs and does so at a rate governed by the turbulence time scale, hence for these regions $t_m = \tau$. For the remainder of the SPMMM simulation, the timestep is modified to include the micromixing time scale

$$\Delta t = \mu_t \min[T_{L_u}, T_{L_v}, T_{L_w}, t_m]. \tag{22}$$

These preparations accomplished, to simulate micromixing SPMMM releases N particles (which need not be equal to N_ϕ used for the MEANS pre-calculation simulations), one at a time, uniformly on the upstream face of the spatial domain. If the particle originates outside of the source region, then it is given an initial concentration of $\phi_0 = 0$, otherwise it is given an initial concentration of $\phi_0 = \phi_{src}$, where ϕ_{src} is the source concentration. The SPMMM source region is the area on the upstream face of the simulation domain where particles would have been initialized in the MEANS simulations. Virtually all of the particles initialized within the MEANS source will be within five standard deviations of the centre of the source $(0, y_s, z_s)$. Therefore, a particle is considered to be within the SPMMM source

region if $r \leq 5\sigma_0$. The distance from the particle position to the source centre in the yz-plane is calculated as $r^2 = (y - y_s)^2 + (z - z_s)^2$. The concentration profile for a Gaussian source is

$$\phi_{\text{src}} = \frac{Q}{2\pi\sigma_0^2 U} \exp\left(\frac{-r^2}{2\sigma_0^2}\right). \tag{23}$$

This equation results in $\phi_0 \approx 0$ for particles beyond approximately three standard deviations from the source centre (i.e. $r \gtrsim 3\sigma_0$).

As the particle travels downstream, its concentration is compared to the conditional mean concentration of the position–velocity bin that it presently occupies. Under the assumption that $\langle\phi|\boldsymbol{u}\rangle$ is approximately constant (it is in fact very slowly varying for sufficiently small timestep Δt), Eq. 9 can be solved to give

$$\phi(t + \Delta t) = \phi(t)\exp(-\Delta t/t_{\text{m}}) + \langle\phi|\boldsymbol{u}\rangle(1 - \exp(-\Delta t/t_{\text{m}})), \tag{24}$$

which is used by SPMMM to update the concentration of the particle. Equation 24 always results in the concentration of the particle mixing towards the conditional mean concentration. During downstream propagation, the particles pass through user-specified extraction planes and their positions and concentrations are saved to file for future processing. Given the stochastic nature of the models described above, the data contains statistical noise. A Savitzky–Golay smoothing filter was used to remove this noise while maintaining the signal.

3 Results

3.1 Model Evaluation

We use three performance measures to evaluate SPMMM. In the following definitions, an observed quantity (e.g. from a full-scale, wind-tunnel or water-channel experiment) is denoted by Q_o and a predicted quantity (e.g. a result from a model) is denoted by Q_p. An overbar indicates an arithmetic mean of all (or some subset of) the available observations or predictions.

The fractional bias,

$$FB = \frac{(\overline{Q_o} - \overline{Q_p})}{0.5(\overline{Q_o} + \overline{Q_p})}, \tag{25}$$

is a measure of the systematic bias of the model. A perfect model would have $FB = 0$. However, if the model both underpredicts and overpredicts the results, it is possible for $FB = 0$ due to the cancellation of errors. The normalized mean square error,

$$NMSE = \frac{\overline{(Q_o - Q_p)^2}}{\overline{Q_o}\ \overline{Q_p}}, \tag{26}$$

is a measure of the mean relative scatter of the model results. A perfect model would have no scatter and thus have $NMSE = 0$. The fraction of data within a factor of 2 of the observations,

$$FAC2 = \text{fraction of data that satisfy } 0.5 \leq \frac{Q_p}{Q_o} \leq 2.0, \tag{27}$$

is a robust performance measure as it is not susceptible to outliers in the data. A perfect model would have $FAC2 = 1$, that is 100% of all data would be within a factor of two of the

observations. Chang and Hanna (2004) suggest that for an acceptable model the performance measures would satisfy: $-0.3 < FB < 0.3$, $NMSE < 4$ and $FAC2 > 0.5$.

3.2 Experimental and Computational Set-Up

To evaluate SPMMM, we simulated the concentration fluctuations due to a continuous source emitting into a neutral wall shear-layer flow, which is arguably the simplest regime of atmospheric turbulence. In particular, in the constant stress layer well-known analytic profiles of the key observables (mean wind speed, shear stress, TKE and its dissipation rate) satisfy appropriately simplified governing equations.

The Fackrell and Robins (1982; hereafter referred to as FR82) experiments were designed to investigate the effects of the source size on the concentration fluctuations. They were carried out in the Marchwood Engineering Laboratories' open-circuit wind tunnel, measuring $24\,\mathrm{m} \times 9.1\,\mathrm{m} \times 2.7\,\mathrm{m}$. A neutral wall shear-layer, corresponding to a natural atmospheric boundary layer, was grown within the wind tunnel. The boundary-layer depth was $\delta = 1.2\,\mathrm{m}$ and the mean streamwise velocity at the top of the boundary layer (the free-stream velocity) was $\langle u \rangle_\delta = 4.0\,\mathrm{m\,s^{-1}}$. The friction velocity was reported as $u_* = 0.188\,\mathrm{m\,s^{-1}}$, and the roughness length as $z_0 = 2.88 \times 10^{-4}\,\mathrm{m}$. To ensure Reynolds number independence, a roughness Reynolds number of $Re_\mathrm{r} = u_* z_0 / \nu \geq 1$, where $\nu \approx 1.5 \times 10^{-5}$ is the kinematic viscosity of air, is required (Snyder and Castro 1997, 2002). The roughness Reynolds number for the FR82 experiments was $Re_\mathrm{r} \approx 3.6$.

A neutrally buoyant mixture of propane and helium was released isokinetically from ground-level and elevated compact sources. The sources were horizontally oriented circular pipes with various diameters. The ground-level sources (GLSs) had diameters of 3, 9 and 15 mm, while the elevated sources had diameters of: 3, 8.5, 9, 15, 25 and 35 mm. The height of the elevated sources was $z_\mathrm{s} = 0.19\delta = 0.228\,\mathrm{m}$. Propane concentration measurements were extracted at several downstream locations with a modified flame-ionization detection system described in Fackrell (1980). From these measurements, FR82 calculated experimental concentration statistics.

The measured velocity statistics and the TKE dissipation rate for the FR82 flow are shown as symbols in Fig. 1. The solid lines are the profiles used to drive MEANS and SPMMM. For the mean streamwise velocity and the TKE dissipation rate, the solid lines represent standard analytical formulations for neutral wall shear-layers:

$$\langle u \rangle = \frac{u_*}{k_\mathrm{v}} \ln(z/z_0), \tag{28}$$

and

$$\varepsilon = \frac{u_*^3}{k_\mathrm{v} z}. \tag{29}$$

For the velocity variances and covariance, the solid lines represent cubic spline interpolations to the experimental measurements. For $z/\delta \lesssim 0.05$ (i.e. below the height of the lowest data measurement), the values of the stresses were held constant to represent the constant stress layer. The flow was treated as horizontally homogeneous and stationary.

For the simulations described below, the spatial domain was discretized into 60 bins in the x, y and z directions: $N_x = 60$, $N_y = 60$ and $N_z = 60$. The conditional mean concentrations were calculated in a velocity domain that was discretized into 20 bins in the u, v and w directions: $N_u = 20$, $N_v = 20$ and $N_w = 20$. The timestep constant was $\mu_\mathrm{t} = 0.02$. All MEANS simulations utilized 2×10^7 particles while the number of particles in the SPMMM

Fig. 1 Dimensionless velocity statistics and TKE dissipation rate for the FR82 neutral wall shear-layer flow. The *symbols* are extracted from Fig.1 of Fackrell and Robins (1982). The *lines* are the fitted profiles used to drive the MEANS and SPMMM models. In the case of the mean streamwise velocity and the TKE dissipation rate, the *solid lines* represent Eqs. 28 and 29 respectively. For the stresses, the *solid lines* represent cubic spline interpolations of the experimental data

simulations ranged from 5×10^6 to 5×10^7 depending on the source size, with larger sources having fewer particles.

Altogether there are four parameters that need to be set in SPMMM: the Kolmogorov constant C_0, the initial source distribution σ_0, the Richardson constant C_r and the micromixing constant μ. The best fit of simulated vertical profiles of mean concentration to FR82 experimental data was realized with a Kolmogorov constant of $C_0 = 6.0$, the value used for all simulations herein. The optimal values of $\mu = 0.75$ and $C_r = 0.45$ were determined by comparing FR82 experimental data to a simulated streamwise transect of the concentration fluctuation intensity (defined by FR82 as $\max(\sigma_\phi)/\max(\langle\phi\rangle)$, where $\sigma_\phi = \sigma_\phi(x, y, z)$ is the standard deviation of the concentration) from a 9-mm elevated source. By this same method, $\sigma_0 = \sqrt{(2/3)}d_s$ was determined to provide the best agreement with experimental data for the values of σ_0 examined. Examples of the resulting profiles, and values of the performance measures, can be found below. With the exception of σ_0, the optimized values of the SPMMM free parameters mentioned above differ from those of the CASS simulations, which used $C_0 = 5.0$, $\mu = 0.8(3/2)^{-1/2} \approx 0.65$ and $C_r = 0.30$. There are a few reasons for these discrepancies: CASS optimized the micromixing constant by comparing with two-particle LS data (Thomson 1990) for a different flow and source geometry; they predicted the Richardson constant to agree with the value reported in Borgas and Sawford (1994); and their numerical implementation differs (in the ways described earlier) from that of SPMMM.

3.3 Dispersion from Elevated Sources

With the free parameters optimized, simulations of dispersion from elevated sources with diameters of 3, 9, 15, 25 and 35 mm, and from a 15-mm ground-level source were carried out. The resulting streamwise transects of concentration fluctuation intensity from these simulations are displayed in Fig. 2. The symbols are FR82 observations, the solid lines are SPMMM simulation results, and the dashed lines are CASS simulation results. As it was shown in FR82

Fig. 2 Streamwise transects of the concentration fluctuation intensity for elevated sources (ES) of various sizes, and for a ground-level source (GLS). The *symbols* are from the FR82 wind-tunnel experiments. The *solid lines* are SPMMM simulation results, and the *dashed lines* are CASS simulation results. Recall that the 9-mm elevated source data was used to optimize the values of μ, C_r and σ_0 in SPMMM

Table 1 Performance measures corresponding to the streamwise transects of the concentration fluctuation intensities from SPMMM simulations of dispersion from compact sources simulations shown in Fig. 2

Source	FB	NMSE	FAC2
3-mm ES	0.046	0.058	1.00
9-mm ES	−0.038	0.0056	1.00
15-mm ES	0.040	0.016	1.00
25-mm ES	0.15	0.033	1.00
35-mm ES	0.17	0.052	1.00
15 GLS	−0.069	0.0093	1.00

that the relative fluctuation intensity for ground-level sources was not sensitive to the source size, we only display the results for a 15-mm ground-level source. The simulation results for the 3-mm and 9-mm ground-level sources agreed well with the observations of FR82. Table 1 lists the performance measures for the six SPMMM simulations shown in Fig. 2. All are within their acceptable ranges.

The profiles for the elevated sources all display the same trend, an initial rise followed by a slow decay. Since the model was tuned to the 9-mm profile it is not surprising that profile displays the best fit to the FR82 observations, showing good agreement with the FR82 experimental data along the full fetch, and by the performance measures in Table 1. When compared to the FR82 data, the initial rise of the concentration fluctuation intensity is too low for the 3-mm elevated source but captured reasonably well for the 15-mm, 25-mm, and 35-mm sources. Farther downstream, it appears that the mixing is too vigorous for $0.80 \lesssim x/\delta \lesssim 2.0$. In this range, the modelled fluctuations from the 15-mm, 25-mm, and 35-mm sources are approximately 10–15% too low. The SPMMM simulations for the elevated

sources consistently produced concentration fluctuation intensities with 10–15% lower maximum intensities than the corresponding CASS results, although the general shapes of the curves are the same.

For the 9-mm results, there is approximately a 9.1% difference between the maximum fluctuation intensities (at $x/\delta \approx 0.6$) of the SPMMM and CASS simulations. An SPMMM simulation with an 8.5-mm source (not shown) had approximately a 4.6% difference with the CASS simulation. Therefore, approximately 50% of the difference between the SPMMM and CASS predictions for the 9-mm source is due to the differing source sizes, the remainder is likely due to differences in model parameters and the numerical implementation of the two models.

The SPMMM model results reproduced the FR82 data with fair accuracy, and are very similar to the results of the CASS model. It is interesting to compare the resulting profiles from the 15-mm elevated source simulations to those from the 15-mm ground-level source simulations. Both profiles start at approximately the same value but the profile for the elevated source increased substantially before undergoing a slow decay whereas the profile for the ground-level source exhibits a very small rise and then stays approximately constant. Physically this is caused by increased stretching, twisting and folding of the material lines due to increased velocity shear near ground level. This figure shows that the parametrization for the micromixing time scale captures reasonably well this effect.

In addition to the simulations discussed above, a simulation with an 8.5-mm elevated source was performed to compare the vertical profiles of mean concentration and concentration variance to the FR82 data. A comparison of the MEANS and SPMMM-simulated vertical profiles of the mean concentration on the plume centreline at five downstream locations is shown in Fig. 3. A characteristic time scale τ_a/T_L is provided for each position, where $\tau_a = x/\langle u(z_s) \rangle$ is the advection time scale. These results demonstrate the required first-order consistency of SPMMM. First-order consistency in SPMMM is dependent upon the complexity of the driving velocity statistics and the resolution of the discretized velocity domain—more velocity bins are needed as the complexity of the PDF of the driving velocity statistics increases. For the FR82 velocity field, $N_u = N_v = N_w \geq 15$ was required to assure first-order consistency. Very slight improvements were seen in the performance measures by increasing the number of velocity bins in u, v and w space to 20, after which no improvements were seen.

The vertical profiles of the normalized mean concentration on the plume centreline at five downstream locations are shown in Fig. 4. The performance measures for the SPMMM simulations are: $FB = -0.0045$, $NMSE = 0.015$ and $FAC2 = 0.92$, all within the acceptable limits. There is a general agreement amongst all three data sets shown in the figure. Small differences between the SPMMM and CASS results are likely due to the different values of C_0 utilized, and perhaps small differences in the velocity statistics used to drive the models.

Figure 5 shows vertical profiles of the normalized concentration variance on the plume centreline at the same locations as in Figs. 3 and 4. The shaded grey areas represent ± 1 standard error. The performance measures for this simulation are: $FB = 0.14$, $NMSE = 0.12$, $FAC2 = 0.74$, poorer than the performance measures for the mean concentration, but still within the acceptable ranges. In the first panel from the left ($x/\delta = 0.96$) the agreement between SPMMM and the FR82 measurements is generally quite good. Farther downstream (last four panels), the modelled maximum variance is always at a greater height than the FR82 measurements and the modelled variance close to the ground is underpredicted when compared with the FR82 measurements. The general shapes of the modelled profiles are very similar to the experimental profiles but they appear to be shifted upwards in height.

Fig. 3 Vertical profiles of the mean concentration on the plume centreline for an 8.5-mm elevated source at five downstream locations from MEANS (*solid line*) and SPMMM (*dashed line*) simulations. These results demonstrate the first-order consistency of SPMMM


Fig. 4 Vertical profiles of the normalized mean concentration on the plume centreline for an 8.5- mm elevated source at five downstream locations. The *open circles* are from the FR82 wind-tunnel experiments. The *solid lines* are SPMMM simulation results and the *dashed lines* are the CASS simulation results

The CASS vertical profiles of concentration variance displayed a better agreement with the FR82 measurements for $z/\delta \lesssim 0.2$. One possible explanation for this lies in the nature of the two models. The CASS model utilized simultaneous particle trajectories, and the conditional mean concentration in a bin was computed 'on the fly' at each timestep by considering the concentrations and velocities of the particles that occupied that bin at a particular timestep. It is therefore conceivable that conditional mean concentration in the bin would vary slightly from step to step. Mixing towards this varying conditional mean concentration may result in increased concentration fluctuations. In contrast, the conditional mean concentrations in SPMMM are pre-calculated and remain the same throughout the simulation. For each step

Fig. 5 Vertical profiles of the normalized concentration variance on the plume centreline for an 8.5- mm elevated source at five downstream locations. The *open circles* are from the FR82 wind-tunnel experiments. The *solid lines* are simulation results from SPMMM and the *dashed lines* are the CASS simulation results. The area *shaded grey* represents ±1 standard error

in a particular bin, a particle mixes with the unchanging conditional mean concentration of that bin. If the concentration of the particle reaches the conditional mean concentration, then no more mixing will occur. In practice, this will not happen due to the exponential nature of the relaxation towards the conditional mean. However, the concentration of the particle will closely approximate the conditional mean concentration. If many particles that travel through this bin reach the conditional mean concentration, then concentration fluctuations will decrease.

This hypothesis is supported by two pieces of evidence. First, the agreement between SPMMM and the FR82 measurements in the left-most panel of Fig. 5 is quite good. This panel corresponds to the earliest available travel time, and therefore the concentrations of the particles have not yet had much time to mix towards the conditional mean concentration. Second, increasing the mixing time scale (by increasing μ) resulted in a better fit to the normalized data in the figure. A longer mixing time scale results in the concentrations of the particles mixing more slowly to the conditional mean concentrations. This hypothesis may also explain why Cassiani et al. (2005a) used smaller values of μ and C_r than that were used with SPMMM.

Another possible source of the discrepancy (again related to the implementation of the models) is that by calculating the conditional mean concentration at each timestep, or 'on the fly', the CASS model contains more bias error than SPMMM, which has the pre-calculated conditional mean concentration supplied to it. A bias towards overprediction would improve the results of the SPMMM simulations, although it would be due to error and not model accuracy. Given that both the CASS model and SPMMM utilize the same micromixing model, it is possible that the better results obtained by Cassiani et al. (2005a) for the vertical profiles of normalized concentration variance are due to bias and may be fortuitous. However, differences in numerical implementation, representation of the driving velocity statistics, and model constants are likely responsible for the majority of the discrepancies.

Table 2 Performance measures corresponding to the sensitivity of SPMMM to the number of streamwise position bins

Statistic	N_x	FB	NMSE	FAC2
$\langle \phi \rangle / \max \langle \phi \rangle$	20	−0.0006	0.015	0.92
	60	−0.0045	0.015	0.92
	100	−0.007	0.016	0.92
$\sigma_\phi^2 / \max \sigma_\phi^2$	20	0.18	0.14	0.71
	60	0.14	0.12	0.74
	100	0.120	0.12	0.74

Sensitivity tests to the number of streamwise position bins were performed for $N_x = 20$ and $N_x = 100$. All of the other parameters were kept the same. The results are summarized in Table 2. From these data, it is shown that the ability of SPMMM to predict the normalized mean concentration and the normalized concentration variance in the FR82 flow is not strongly affected by the streamwise resolution of the spatial domain. The performance measures for the normalized mean concentration are very similar for the three simulations, and there is only a slight improvement of the performance measures for the normalized concentration variance with increasing streamwise spatial resolution. Qualitatively, the results from the three simulations were indistinguishable.

3.4 Dispersion from a Ground-Level Source

We now consider dispersion from a 15-mm ground-level source. Figure 6 shows the vertical profile of the normalized mean concentration on the plume centreline at three downstream locations. The filled circles are FR82 data in the downstream region $1.67 \le x/\delta \le 5.92$. The 'plus' signs are from CASS simulation results, and the other symbols correspond to SPMMM simulation results at the specified extraction locations. The vertical coordinate has been scaled by the vertical plume halfwidth δz to display the self-preserving nature of the plume (i.e. the scaled shape of the plume from any extraction location is invariant). Excellent agreement between SPMMM and the FR82 experiments is realized in this figure. Performance measures were not calculated for this simulation.

Shown in Fig. 7 are vertical profiles of the normalized concentration variance from the SPMMM and CASS simulations for the same locations shown in Fig. 6, along with FR82 experimental data. The predicted variance profile for the ground-level source is in very good conformance with the FR82 experimental data, much better than the predicted variance profiles from the elevated source. Performance measures were not calculated for this simulation. Recall that the TKE production and dissipation rates increase towards the ground. This leads to increased mixing and dissipation of the concentration fluctuations near the ground. Evidently, the parametrization of the micromixing time scale used by SPMMM captures this effect.

The results from the SPMMM simulations of dispersion from a ground-level source are slightly more accurate than the CASS results, as seen in Figs. 6 and 7. There is a significant positive bias (overestimation) in the CASS simulation results for both the mean concentration and the concentration variance for $z/\delta z \gtrsim 1.8$, which is absent in the SPMMM simulation results, the latter of which agree very well with the FR82 experimental data. This discrepancy in the CASS simulation results may be due to increased bias and statistical errors in the estimation of the conditional mean concentration obtained from using an 'on the fly' approach

Fig. 6 Vertical profiles of the normalized mean concentration on the plume centreline for a 15-mm ground-level source. The *filled circles* are from the FR82 wind-tunnel experiments in the downstream region $1.67 \leq x/\delta \leq 5.92$. The *plus signs* are from CASS simulation results, and the other *symbols* are SPMMM simulation results. The vertical coordinate has been scaled by the vertical plume halfwidth to display the self-preserving nature of the plume


Fig. 7 Vertical profiles of the normalized concentration variance on the plume centreline for a 15-mm ground-level source. The *filled circles* are from the FR82 wind-tunnel experiments in the downstream region $1.67 \leq x/\delta \leq 5.92$. The *plus signs* are CASS simulation results, and the other *symbols* are SPMMM simulation results. The vertical coordinate has been scaled by the vertical plume halfwidth to display the self-preserving nature of the plume

at each timestep. This seems to provide further supporting evidence that the better results for concentration variance obtained for the elevated source in Fig. 4 for CASS compared to SPMMM are probably rather fortuitous—they may be the result of increased bias and

statistical errors in CASS simulations compensating for model inaccuracies. However, for the ground-level source, this increased bias and statistical error worked against CASS simulations, that are poorer (in the mid to upper plume fringes) than the SPMMM simulations that, (1) predict the proper shape of the self-similar mean concentration and concentration variance profiles, and (2) are completely consistent with the FR82 experimental measurements (both magnitude and shape).

4 Conclusions

The ability of a sequential particle micromixing model called SPMMM to simulate dispersion from elevated and ground-level sources in the FR82 wall shear-layer flow was investigated. The four free parameters of the model were optimized to: $C_0 = 6.0$, $\mu = 0.75$, $C_r = 0.45$, and $\sigma_0 = \sqrt{(2/3)}d_s$. These optimized values resulted in an acceptable agreement with the measured streamwise transects of the concentration fluctuation intensity for elevated sources of various sizes, as well as good agreement for a ground-level source. The modelled profiles for the elevated sources were 10–15% lower than the measured values in the downstream range $0.80 \lesssim x/\delta \lesssim 2.0$. The initial rise of the concentration fluctuation intensity from the SPMMM simulation results was also about 10–15% lower than the CASS simulation results. After the initial rise, the results from the two models showed a closer agreement.

The SPMMM-simulated vertical profiles of normalized mean concentration from an 8.5-mm elevated source were in good agreement with the FR82 measurements, and with the CASS simulation results. The SPMMM-simulated vertical profiles of normalized concentration variance were in reasonable agreement with the FR82 measurements, although not as accurate as the CASS simulation results. The SPMMM simulations produced variance levels near ground level that were too low compared to the measurements. For dispersion from a 15-mm ground-level source the SPMMM simulated profiles of normalized mean concentration and normalized concentration variance were in excellent agreement with the FR82 experimental results, and slightly more accurate than the CASS results, most notably for $z/\delta z \gtrsim 1.8$.

The majority of the discrepancies between the SPMMM and CASS simulation results are likely due to the representation of the driving velocity statistics, differing model constants, and to the numerical implementations of the models—the CASS model utilized simultaneous particle trajectories and the conditional mean concentration field was computed at each timestep, whereas SPMMM utilized sequential particle trajectories and a pre-calculated conditional mean concentration field. It is conceivable that calculating the conditional mean concentration field at each timestep may result in slight differences in the field and introduce more concentration variance into the model. As shown in Figs. 2, 5 and 7, the CASS model results do display greater concentration variance relative to the SPMMM results.

Overall the sequential particle, pre-calculation approach of the SPMMM model shows promise. All simulations presented in this article produced acceptable results, and in the case of dispersion from ground-level source, good results were obtained. The velocity shear, and TKE production and dissipation rates are maximized near ground level, where SPMMM was shown to produce good results for ground-level sources. This suggests that SPMMM may also be effective at simulating dispersion from ground-level (or near to it) sources in other flows with large shear and TKE production and dissipation rates, such as canopy flow.

References

Borgas MS, Sawford BL (1994) A family of stochastic models for two-particle dispersion in isotropic homogeneous stationary turbulence. J Fluid Mech 279:69–99

Cassiani M, Franzese P, Giostra U (2005a) A PDF micromixing model of dispersion for atmospheric flow. Part I: development of model, application to homogeneous turbulence and to a neutral boundary layer. Atmos Environ 39:1457–1469

Cassiani M, Franzese P, Giostra U (2005b) A PDF micromixing model of dispersion for atmospheric flow. Part II: application to convective boundary layer. Atmos Environ 39:1471–1479

Cassiani M, Radicchi A, Giostra U (2005c) Probability density function modelling of concentration in and above a canopy layer. Agric For Meteorol 133:153–165

Cassiani M, Radicchi A, Albertson JD (2007) Modelling of concentration fluctuations in canopy turbulence. Boundary-Layer Meteorol 122:655–681

Chang CJ, Hanna SR (2004) Air quality performance evaluation. Meteorol Atmos Phys 87:167–196

Fackrell JE (1980) A flame ionisation detector for measuring fluctuating concentration. J Phys E 13:888–893

Fackrell JE, Robins AG (1982) Concentration fluctuations and fluxes in plumes from point sources in a turbulent boundary layer. J Fluid Mech 117:1–26

Fox RO (1996) On velocity-conditioned scalar mixing in homogeneous turbulence. Phys Fluids 8:2678–2691

Franzese P (2003) Lagrangian stochastic modelling of a fluctuating plume in the convective boundary layer. Atmos Environ 37:1691–1701

Gifford FA (1959) Statistical properties of a fluctuating plume dispersion model. Adv Geophys 6:117–137

Luhar AK, Sawford BL (2005a) Micromixing modelling of concentration fluctuations in inhomogeneous turbulence in the convective boundary layer. Boundary-Layer Meteorol 114:1–30

Luhar AK, Sawford BL (2005b) Micromixing modelling of mean and fluctuating scalar fields in the convective boundary layer. Atmos Environ 39:6673–6685

Monin AS, Yaglom AM (1975) Statistical fluid mechanics II. MIT Press, Cambridge, 874 pp

Pope SB (1998) The vanishing effect of molecular diffusivity on turbulent dispersion: implications for turbulent mixing and the scalar flux. J Fluid Mech 359:299–312

Rembold B, Grass M, Jenny P (2008) Parallel hybrid particle/finite volume algorithm for transported PDF methods employing sub-time stepping. Comput Fluids 37(3):181–193

Sawford BL (2004a) Conditional scalar mixing statistics in homogeneous isotropic turbulence. New J Phys 6:1–30

Sawford BL (2004b) Micro-mixing modelling of scalar fluctuations for plumes in homogeneous turbulence. Flow Turbul Combust 72:133–160

Snyder WH, Castro IP (1997) Surface Reynolds number effects in rough-wall boundary layers. Paper presented at Third European fluid mechanics conference, Gottingen, September 1997

Snyder WH, Castro IP (2002) The critical Reynolds number for rough-wall boundary layers. J Wind End Ind Aerodyn 90:41–54

Thomson DJ (1987) Criteria for the selection of stochastic models of particle trajectories in turbulent flows. J Fluid Mech 180:529–556

Thomson DJ (1990) A stochastic model for the motion of particle pairs in isotropic high-Reynolds-number turbulence, and its application to the problem of concentration variance. J Fluid Mech 210:113–153

Wilson JD, Sawford BL (1996) Review of Lagrangian stochastic models for trajectories in the turbulent atmosphere. Boundary-Layer Meteorol 78:191–210