Contrasting the Use of Single-Realization versus Ensemble-Average Atmospheric Dispersion Solutions for Chemical and Biological Defense Analyses

PAUL E. BIERINGER AND ANDREW J. ANNUNZIO

Research Applications Laboratory, National Center for Atmospheric Research, Boulder, Colorado*

NATHAN PLATT

Institute for Defense Analysis, Alexandria, Virginia

GEORGE BIEBERBACH

Research Applications Laboratory, National Center for Atmospheric Research, Boulder, Colorado

JOHN HANNAN

Defense Threat Reduction Agency, Fort Belvoir, Virginia

(Manuscript received 12 June 2013, in final form 2 January 2014)

ABSTRACT

Chemical and biological (CB) defense systems require significant testing and evaluation before they are deployed for real-time use. Because it is not feasible to evaluate these systems with open-air testing alone, researchers rely on numerical models to supplement the defense-system analysis process. These numerical models traditionally describe the statistical properties of CB-agent atmospheric transport and dispersion (AT&D). While the statistical representation of AT&D is appropriate to use in some CB defense analyses, it is not appropriate to use this class of dispersion model for all such analyses. Many of these defense-system analyses require AT&D models that are capable of simulating dispersion properties with very short time-averaging periods that more closely emulate a ''single realization'' of a contaminant or CB agent dispersing in a turbulent atmosphere. The latter class of AT&D models is superior to the former for performing CB-system analyses when one or more of the following factors are important in the analysis: high-frequency sampling of the contaminant, spatial and temporal correlations within the contaminant concentration field, and nonlinear operations performed on the contaminant concentration. This paper describes and contrasts these AT&D modeling tools and provides specific examples in which utilizing ensembles of single realizations of CB-agent AT&D is advantageous over using the statistical, ''ensemble-average'' representation of the agent AT&D. These examples demonstrate the importance of using an AT&D modeling tool that is appropriate for the analysis.

1. Introduction

In recent decades the materials of concern and delivery methods associated with the use of chemical and biological (CB) agents have continued to evolve. To combat these threats, the U.S. Department of Defense makes significant investments in technologies designed for CB-agent detection and defeat. CB defense-system analysis is a critical element in the defense-system acquisition process that includes identifying technology gaps, determining technology investment direction, and providing information that ultimately directs system acquisition and deployment decisions. While the use of live agents is the most advantageous approach for conducting CB defense-system analyses, their use is clearly difficult, and therefore agent simulants are frequently used in outdoor field-data-collection efforts (Przybylowicz et al. 2003). The use of agent simulants

* The National Center for Atmospheric Research is sponsored by the National Science Foundation.

Corresponding author address: Paul E. Bieringer, National Center for Atmospheric Research, 3450 Mitchell Ln., Boulder, CO, 80301.
E-mail: paulb@ucar.edu

DOI: 10.1175/JAMC-D-13-0201.1

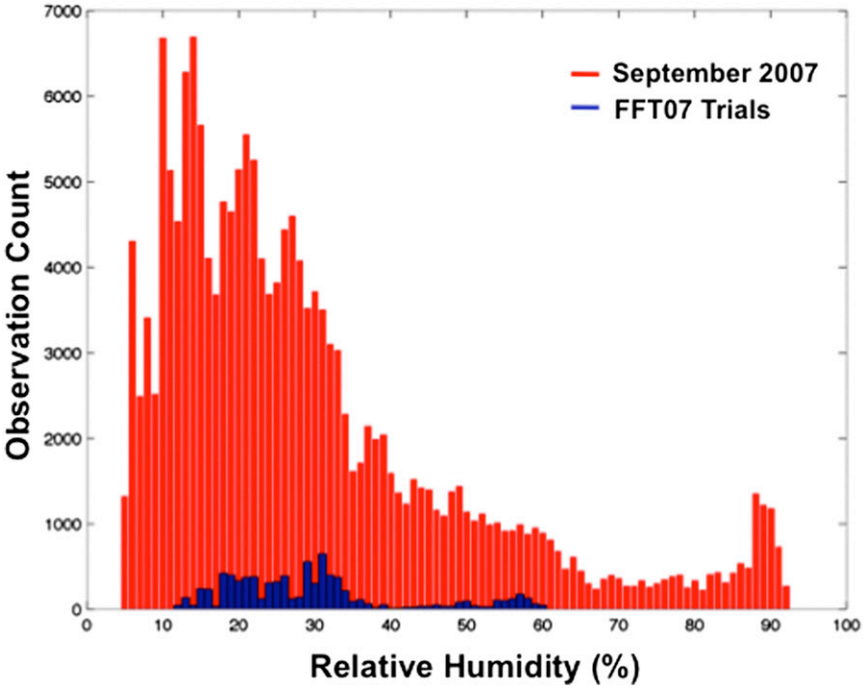

FIG. 1. Relative humidity observations at Dugway Proving Ground during the month of the FFT07 experiment and during the experiment itself.

makes it possible to collect data that can be used to assess system performance, but the permutation space influencing the agent atmospheric transport and dispersion (AT&D) is large. This permutation space includes factors such as variability in meteorological variables (temperature, wind, and moisture), local topography, threat delivery mechanisms, and background interferences such as dust, pollen, and exhaust. These field experiments are unfortunately limited by monetary, logistical, and time constraints and hence cannot capture all of these permutations. As a result, it is not possible to test these systems across the full range of conditions in which they are likely to be deployed.

As an example, biological agents and the systems used to detect them can be sensitive to changes in relative humidity in the atmosphere. Relative humidity influences the hygroscopic growth of background aerosols that are made up of inorganic salts (Malm et al. 2003). The growth and shrinkage of these particles change the scattering properties of these background aerosols and in turn influence the ability of a lidar-based system to detect agents in these environments. Understanding the performance of these systems requires that field data cover the range of atmospheric conditions that can influence system performance. The challenge, however, is that this is difficult to accomplish in ''open air'' tests for which the environmental conditions cannot be controlled. Figure 1 provides an illustration of this limitation using the relative humidity observations recorded at Dugway Proving Ground in Utah during the Fusing Sensor Information from Observing Networks (FUSION) Field Trial 2007 (FFT07) field experiment. FFT07 is one of the more comprehensive field experiments for short-range AT&D and included extensive meteorological measurements and agent-simulant AT&D measurements (Storwold 2007). Figure 1 contrasts the distribution of relative humidity observations during the periods in which CB-agent simulants were released with the distribution of relative humidity observations during the month of the field experiment. Although the data collection in this field trial was extensive, it was not logistically or economically feasible to maintain an experiment of this magnitude for a long enough period of time to collect data when the relative humidity was higher. Furthermore, this trial only included weather conditions experienced in the autumn and may not be representative of conditions that occur at other times of the year or at other locations. Thus, even for a comprehensive field experiment such as FFT07, data collected from a field program do not result in a dataset that captures a significant portion of the environmental variability needed for the defense analysis. This example illustrates a common issue encountered when analysts attempt to use field-trial datasets in

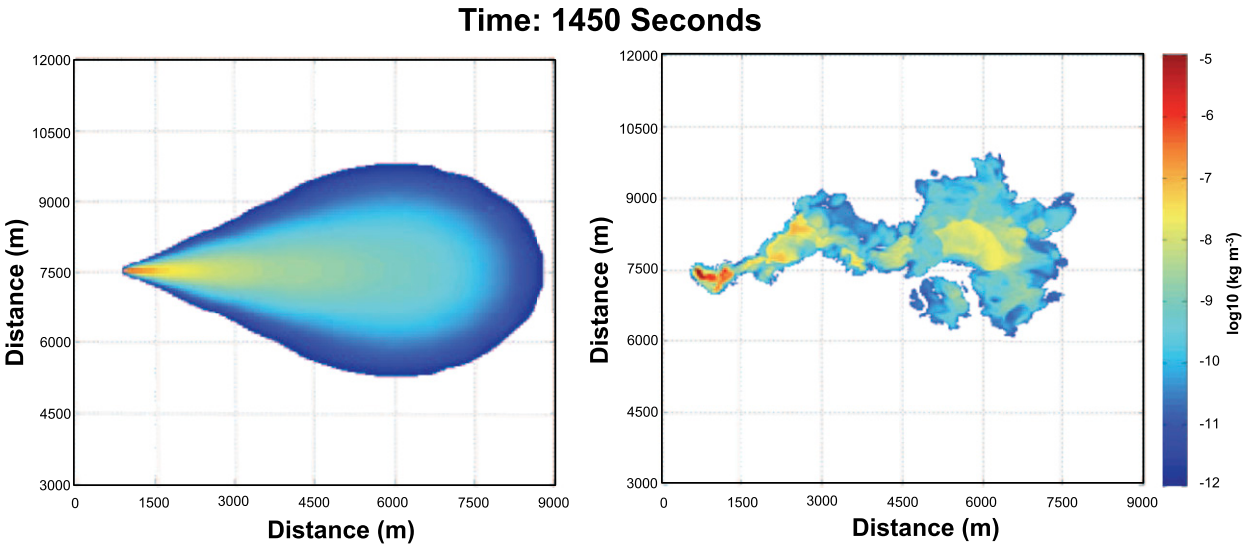

FIG. 2. Two simulations of continuous-release AT&D for the same meteorological conditions: (left) generated using the SCIPUFF Gaussian puff AT&D model and (right) generated using the WRF model with an embedded Eulerian AT&D model.

their analyses and motivates why they often rely on numerical simulations of CB-agent AT&D to extend the observational datasets.

To address this deficiency, scientists and engineers conducting this type of work have traditionally utilized two classes of AT&D models to supplement the data collected in field experiments like FFT07. The first class includes models that describe the statistical representation of a contaminant release over many possible AT&D realizations. Models of this type acknowledge the apparent stochasticity of atmospheric turbulence and, via Reynolds averaging, calculate ensemble-average AT&D. An example of a simulation from this class of AT&D models can be seen in the left panel of Fig. 2.

The second class of models used for this type of application explicitly computes turbulent motions that cause the fluctuations in contaminant concentration instead of modeling their average effect. This approach relies on the development of detailed atmospheric flows in which large-eddy simulations (LES) are used to solve a filtered set of the Navier–Stokes equations (Wyngaard 2010). LES can successfully model many types of atmospheric flows in the boundary layer where many CB defense-system applications are relevant (Bieberbach et al. 2010; Platt et al. 2012). Flow solutions from this modeling approach, which can be produced at spatial grid increments approaching 1 m, can be used to drive dispersion models (in both Lagrangian and Eulerian reference frames) to produce CB-agent plume AT&D. These models are capable of producing solutions that more closely resemble a ''single realization'' of the AT&D that qualitatively resembles the visual depictions of smoke dispersion. The right panel in Fig. 2 shows an example of a simulation from this class of model. The same input meteorological conditions (e.g., mean wind speed, mean wind direction, and atmospheric stability) were used for the simulations shown in the right and left panels of Fig. 2.

The choice of which modeling tool to use in a CB analysis or testing application is not always obvious. Because of their wider availability and less-demanding computational and data-storage requirements, most prior CB defense analyses have relied on the former class of AT&D models. The former class of models has been extensively used and extensively cited in the literature and requires fewer computational resources, but one should be cautious of its use for some types of defense analyses. The ensemble-average representation of contaminant dispersion can deviate significantly from a realization, and it is not possible to resurrect all of the details of a realization from ensemble-average statistics (Haupt et al. 2013). As a result, the modeling tool may provide a misleading or incorrect outcome for some CB defense analyses. These scenarios include applications in which the following occur or are important: 1) high-frequency sampling of a contaminant, 2) spatial and temporal correlations of the contaminant, and 3) nonlinear operations performed on the contaminant concentration field. These scenarios are common in many of the defense analyses that inform technology gaps, set technology investment direction, inform system acquisition decisions, and support the testing and evaluation of the systems being acquired. The CB defense analyses

involving high-frequency sampling of a contaminant refer to situations in which the problem requires information on a time scale that is much less than the eddy-turnover time scale of the boundary layer. Spatial and temporal correlations in the contaminant concentration field are important in applications such as multisensor CB detection enhancement and false-alarm mitigation and remote sensing of CB contaminants for the protection of a location downwind of the initial detection. Nonlinear operations performed on the contaminant concentration values refer to situations such as agent toxicity in which there can be a nonlinear relationship between the agent concentration/dosage and the human response to the exposure to that material. Because recent advances in computational capabilities now make it possible for an analyst to possess the necessary computational power and storage to complete an analysis using ensembles of single realizations of CB-agent AT&D, we believe that the CB defense community should be aware of the analysis limitations and the options for addressing them.

In this paper, we discuss a set of scenarios in which using the single-realization representation of contaminant AT&D affects the outcome of the defense analysis. We complement each scenario with an example. The two models used to generate the AT&D solutions are the Second-Order Closure Integrated Puff (SCIPUFF) model and an inline-tracer LES, both of which are discussed in section 2. In section 3, we demonstrate the utility of using the single-realization AT&D representation for high-frequency sampling with a simple point detection system. Then, in section 4, we demonstrate the importance of capturing spatial and temporal correlations of the contaminant in the context of false-alarm mitigation and CB detection to inform a response decision at a downwind location. In section 5, we use the CB-agent toxic-load application to demonstrate the importance of utilizing contaminant single realizations when nonlinear operations are performed on the concentration field. Last, we discuss the implications for future CB defense analyses.

2. Atmospheric transport and dispersion models

We use two AT&D models to create the contaminant dispersion data for the examples in the following sections. The first utilizes an inline-tracer model within the Weather Research and Forecasting (WRF) model that directly solves the advection–diffusion equation concurrently with flow-field predictions. Within WRF is an LES capability that we use to provide high-resolution, turbulent flow-field predictions (Skamarock et al. 2008). LES explicitly resolves the flux-carrying eddies that evolve the contaminant concentration and parameterizes the smaller eddies that carry little flux and are responsible for energy dissipation (Wyngaard 2010). In its nascent stages, LES provided a means to further explore atmospheric boundary layer (ABL) flows (Deardorff 1972). Similar to the issues discussed in this paper, many field experiments on ABL flows did not capture all of the features of the flow environment, and LES helped to close the gap in our understanding of such flows (Deardorff 1972).

In WRF, there is an option to perform LES calculations under idealized atmospheric settings. We chose this option since the goal for this work is to investigate how the different model solutions affect the defense-analysis results. It is easiest to perform this comparison when minimizing the number of free parameters; therefore, the LES is run with a horizontally homogeneous lower boundary (i.e., topography, roughness, and surface heat flux), a specified inflow condition, and a specified potential temperature profile. Given these inputs, we ''spin up'' the model until turbulent statistics become steady state before releasing the tracers. Within WRF is an embedded Eulerian dispersion model that numerically solves the tracer-gas evolution equation and provides a high-resolution and realistic AT&D solution. All LES AT&D solutions are generated using 30-m horizontal resolution and 10-m vertical resolution. Since the spatial filter is equal to the grid resolution, these solutions represent short-time-average AT&D solutions. Further, the time step used to compute these solutions is 5 Hz, although the output frequency of the LES AT&D solutions is 15 s. One single-realization AT&D solution is not sufficient to draw conclusions on CB system performance. For this reason, this work leverages concepts employed in Warner et al. (2002), wherein an ensemble of mesoscale weather simulations is used to develop synthetic AT&D datasets to compute statistics on contaminant dosage. In that work, each member of the ensemble provides a different AT&D solution, and the average of these solutions was used to map out a hazard area defined in terms of dosage. Here, we utilize the same concept. An ensemble of single-realization AT&D solutions is generated using the identical input parameters (e.g., meteorological and agent source) with statistically stationary flow field conditions to compute statistics for an analysis application. This work, however, differs in several ways from Warner et al. (2002). First, the spatial domain of the numerical model used here is designed to capture the atmosphere primarily within the ABL, and the spatial resolution is set such that the largest, energy-containing eddies in the ABL can be modeled explicitly. Second, the variability present in each of the AT&D simulations

in Warner et al. (2002) was mesoscale model variability that resulted from perturbed boundary conditions and changing model physics. For the simulations used in the examples presented here, the variability in the AT&D simulations results from shifting releases in space and/or time such that contaminant releases occur in uncorrelated eddies. The model physics and inflow conditions remain the same, and the turbulence provides the variability among solutions. Last, the CB defense applications considered here require the explicit modeling of smaller-scale concentration fluctuations that are needed to derive statistics describing the dispersion patterns, whereas Warner et al. (2002) investigated longer-time-integrated concentration statistics. For the purpose of this paper, it is implied that all defense analyses that utilize this class of AT&D modeling solutions must generate ensembles. In the following sections, if an ensemble of single-realization AT&D solutions is not used, it is because the example is meant to be illustrative.

The second type of AT&D model used in this work is an ensemble-average model that represents a concentration average over many CB dispersion realizations. When meteorological conditions and the contaminant threat parameters are stationary such that ergodic theory applies, then the time average of the concentration values approaches the ensemble-average values (Wyngaard 2010). Hence, we use the terms longer-time-average solution and ensemble-average solution interchangeably in this paper. This averaging tends to smooth the gradients in the concentration field and lower the peak concentration values (Slade 1968). The model used to produce the ensemble-average solutions in this paper is the SCIPUFF model of Sykes et al. (2008). SCIPUFF is a Lagrangian puff model that uses a collection of three-dimensional Gaussian puffs to represent an arbitrary time-varying concentration field. Turbulent diffusion is parameterized using second-order turbulence closure techniques devised by Donaldson (1973) and Lewellen (1977) that utilize available velocity statistics to predict the associated dispersion rates. SCIPUFF supports the modeling of various material types, including gases, particles, and liquid droplets, and their associated size distributions. Particle and droplet dynamic effects can also be simulated with this model, particularly dry/wet deposition and material decay. Urban building effects are modeled using an urban canopy model that estimates the bulk impacts of the buildings on the wind and turbulent dispersion profiles.

For the sake of consistency and simplicity in directly comparing and contrasting how the ensemble-average and the single-realization contaminant representations affect defense-analysis results, SCIPUFF calculations are also performed under idealized settings. Similar to the LES model, we use a horizontally homogeneous lower boundary and specify the mean wind, surface heat flux, and boundary layer depth. The output frequency of the horizontal output grid is identical to that of the LES AT&D solutions. Although not used in the illustrative examples below, SCIPUFF can also utilize time-varying meteorological conditions.

3. Scenario 1: High-frequency point sampling

The first scenario we address is one in which numerical AT&D simulations are used in an analysis involving point sampling of a contaminant. Point detection systems have a variety of applications in CB defense that range from providing warnings for the presence of a contaminant to determining the spatial extent of its dispersion. The application for which the system is used typically determines the requirements for the point detection system. If the analysis involves point detection systems that provide long time-integrated samples, and environmental conditions are statistically stationary, then it is acceptable to perform the analysis with an ensemble-average AT&D solution. Examples of such analyses include human toxic-load calculations on the basis of Haber's law (see section 5) and the mapping of dosage or hazard areas if the dispersion is assumed to follow a clipped normal distribution. If, however, the analysis involves a point detection system that is designed to sample the environment at a high temporal frequency—for example, the digital photoionization detectors (50-Hz sampling frequency) used for the FFT07 field experiment (Storwold 2007) or the miniature, automatic, continuous air-monitoring system (MINICAM) sensors used at Dugway Proving Ground—then it is no longer appropriate to use the ensemble-average (or long time average) AT&D solutions. In this section we illustrate the impact of employing dispersion-model solutions with varying fidelity on an analysis for the performance of point detection systems like these that sample the environment at a high temporal frequency.

First, to illustrate the impact of dispersion-model fidelity, we utilize three different numerically generated AT&D solutions that are based on an ensemble of single-realization AT&D solutions generated using the WRF LES. The first is a single-realization AT&D solution wherein the inherent averaging time is short (e.g., seconds) and is determined by the scale of the spatial filtering in the WRF LES model. The second is a 25-min time average of one single-realization plume. The averaging time is set to 25 min because this is the time required for the plume to first reach the locations where the simulated measurements occur. For this example, a 25-min average with 15-s model output frequency

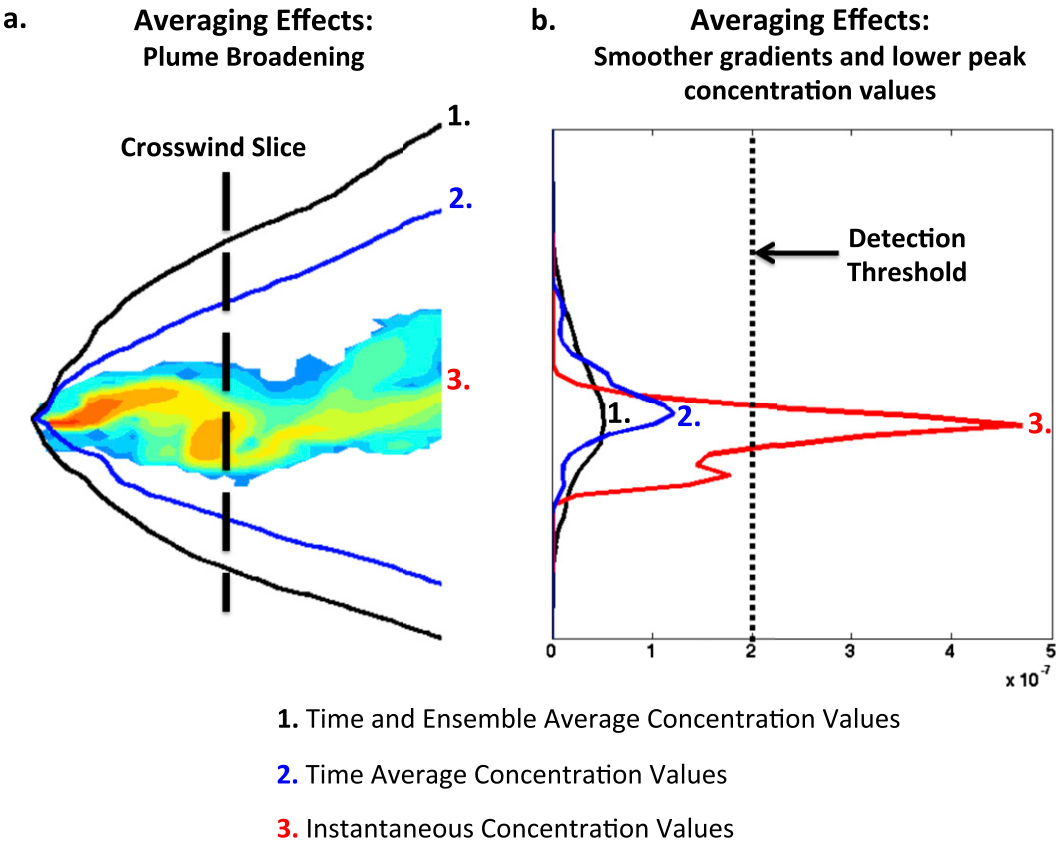

FIG. 3. Impact of averaging on concentration values as one moves across a plume generated using the WRF model with an embedded Eulerian AT&D model. The red lines represent the instantaneous plume, the blue lines represent the time-averaged plume, and the black lines represent the time- and ensemble-average plume.

equates to approximately 100 plume snapshots. The third solution includes an average of 20 plumes, each with 25-min time averaging. The temporal averaging removes the time dependence for the latter two AT&D solutions and provides a solution similar to what is provided by an ensemble-average model driven with statistically stationary and homogeneous flow-field conditions. In this case the temporal variation in the ensemble-average solution is linked solely to the contaminant source terms, and hence the latter two solutions represent surrogates for the ensemble-average solution. This scenario assumes that the contaminant release is 1000 m upwind of the detection system and that the goal of the analysis is to determine whether a high-frequency sampling system can detect the contaminant. This analysis is designed to characterize the detection probability for a given sensor minimum detection threshold. Figure 3 illustrates the effect that averaging has on the contaminant concentration field. Figure 3a shows that the averaging effectively broadens the crosswind extent of the plume, and Fig. 3b shows that the averaging also reduces the peak concentration values and smooths the concentration gradients. Note that the time-average concentration values shown in Fig. 3 do not match that of the time- and ensemble-average plume solution. This occurs because, for the convective boundary layer modeled here, an averaging time of 25 min is not nearly long enough for convergence of the time average to the ensemble representation (Briggs 1993). Figure 3b also displays the crosswind concentration values at a time instant from the single-realization (red line) plume and the average (blue and black lines) crosswind contaminant concentration values 1000 m downwind from the source. The vertical, dotted line in Fig. 3b represents a notional sensor detection threshold, which for this example is set at $2 \times 10^{-7}\,\mathrm{kg\,m^{-3}}$ (note: the actual value is irrelevant for this discussion). In this example, the single-realization concentration values exceed the minimum detection threshold of the sensor at this downwind distance; hence, if an analyst were to use this class of AT&D solution, it would be reported that the sensor is able to make a detection at this range from the source. If an analyst were to utilize the average (time and/or time and ensemble average) solutions to

determine whether a detection system operating at this minimum detection threshold could detect the release, however, the analyst would conclude that it could not.

The differences in the analysis results described above are due to a discrepancy between the duration of the sample taken by the detection system and the temporal fidelity of the modeling solution. Any given detection system will have a sample duration over which the contaminant cloud is interrogated and will make a detection that is typically based (in one form or another) on identifying returned signals that are above a background noise level. When the analysis uses simulations of a contaminant plume that represent a longer time average than that utilized by the detection system, the peak sample concentration values used in the analysis of the detection system will be lower than the peak concentration that would come from the simulated data that are representative of the sample duration of the detector. Because the averaging can potentially decrease the concentration maxima to values that are less than that of the detection threshold, this effect can provide a potentially misleading result. This discrepancy no longer exists when using the single-realization solution. In this case the simulated measurements will contain concentration peaks that may be above the threshold, and hence there is the possibility for the analysis to indicate that a detection is possible. There will, however, be time periods during which the peak of the single-realization plume is less than that of the detection threshold. This is a property of the natural variability in the dispersion at this temporal scale and is not because the spatial and temporal concentration gradients are not physically realistic relative to the sensor averaging time.

This same issue can be a problem in analyses of sensing technologies that are designed to warn by making a rapid detection while minimizing false alarms. For this application, detection thresholds are often used to capture the large peaks in contaminant concentration values associated with a release that is expected to be well above the detection limits of the instrument. Therefore, using longer time-average dispersion solutions to create synthetic observations to test these technologies will also yield inaccurate conclusions in analyses of false-alarm mitigation performance for fast-response sensors.

To overcome this issue, similar analyses have relied on ensemble-average AT&D solutions that also provide the concentration variance (such as SCIPUFF) for the ensemble-mean concentration. The work of Carrano and Jeys (2010) provides a good example in which this was done. Here, the mean and variance are the only inputs to a clipped normal distribution, which, as shown by Lewellen and Sykes (1983, 1986), can be used to describe the concentration probability density function (PDF) at a single spatial location. By integrating over the area of the PDF that exceeds the concentration threshold, it is possible to formulate the probability that exceedance will occur for a specific realization. There are, however, two drawbacks to this approach. The first drawback is that the mean and variance describe the concentration PDF for instantaneous point measurements. Therefore, the probability of exceedance only applies to concentration measurements with an infinitesimal averaging period, as the variance will decrease as the averaging time increases. The second drawback of utilizing the mean and variance is that it is not possible to link the AT&D solution at one place and time with the AT&D solution at another temporal and spatial location. This is important for analyses that are dependent on contaminant travel time that will be discussed in greater detail in the following section.

4. Scenario 2: Spatial and temporal correlations

In this section, we illustrate the importance of using a single-realization synthetic representation of the contaminant dispersion over the ensemble average when the spatial and temporal correlations in the concentration values are needed in the analysis. This can be important in several types of analyses. Examples include analyses that involve the simultaneous examination of multiple sensors in a detection-enhancement application, analyses that involve false-alarm-mitigation methods, and analyses that involve dispersion realizations that are correlated in time such as a CB detection application in which the detection occurs in a different location than the downwind area of concern that the warning is meant to protect.

a. Spatial correlations: False-alarm mitigation

In many CB defense analyses, spatial correlations among the concentration values used in the analysis are important. Examples include sensor siting, determining the probability of detection for a sensing network, and sensor-network false-alarm mitigation (Norige et al. 2009; Carrano and Jeys 2010). Because false-alarm-mitigation applications exploit spatial correlations in a sensor network, we use this application to demonstrate how utilizing an ensemble-average AT&D solution versus the single-realization AT&D solution can influence the analysis results. This is applicable to studies like the one conducted by Norige et al. (2009).

To illustrate these effects qualitatively, our example will use a simulation of a continuous-release scenario that occurs in an unstable boundary layer in which two point sensors are located approximately the same distance downwind of the release location. Time series of

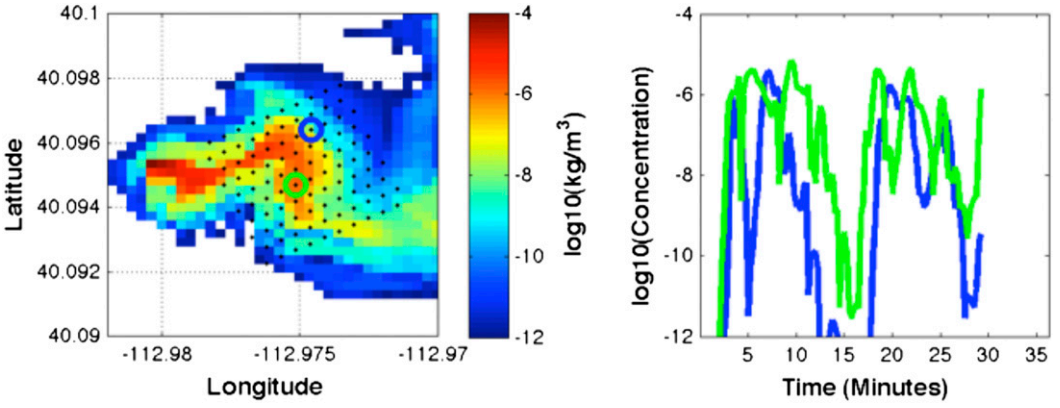

FIG. 4. Experiment setup for testing false-alarm mitigation with single-realization AT&D simulations generated using the WRF model with an embedded Eulerian AT&D model: (left) simulated plume dispersion (snapshot at 30-min mark) and (right) the time series of contaminant concentration values using synthetic sensor observations for the dispersion shown in the left panel.

the contaminant concentration values for the LES-generated single realization and the SCIPUFF ensemble-average plume are depicted graphically in Figs. 4 and 5, respectively. These solutions were generated for convective boundary layer conditions, and both AT&D solutions have a continuous release rate of 1 kg s^{-1} over a 30-min period. In the left panel of Figs. 4 and 5 we illustrate a snapshot of a horizontal slice through the plume at the 30-min mark. This time coincides with the last instant of contaminant emission from the source. Figure 4 shows how the contaminant dispersion exhibits a meandering pattern that is caused by three-dimensional eddies in the boundary layer that are larger than the width of the contaminant plume. The effect of plume meandering is also evident in the right panel of Fig. 4, which depicts a time series of the contaminant concentration values plotted at 15-s temporal resolution at the two sensor locations identified by the green and blue circles in the left panel. The plume meander results in situations in which the concentration at sensors located at approximately the same distance downwind from the release location can simultaneously increase at one sensor and decrease at another. This meandering due to the eddies in the boundary layer is not explicitly resolved in the ensemble-average plume, as shown in Fig. 5. In the ensemble-average case, the absence of mesoscale variability (a product of the simulation design) and a horizontally homogeneous boundary layer (which matches the specifications of the simulation produced by the LES) result in a plume whose temporal variations are linked solely to the contaminant source terms. When this AT&D solution is used, the simulated measurements at similar downwind locations to shown in Fig. 4 now have a highly correlated response. The qualitative example provides insight into how simulated measurements or synthetic observations generated from an ensemble-average AT&D model will lack the realistic gradients in the concentration fields. This in turn leads

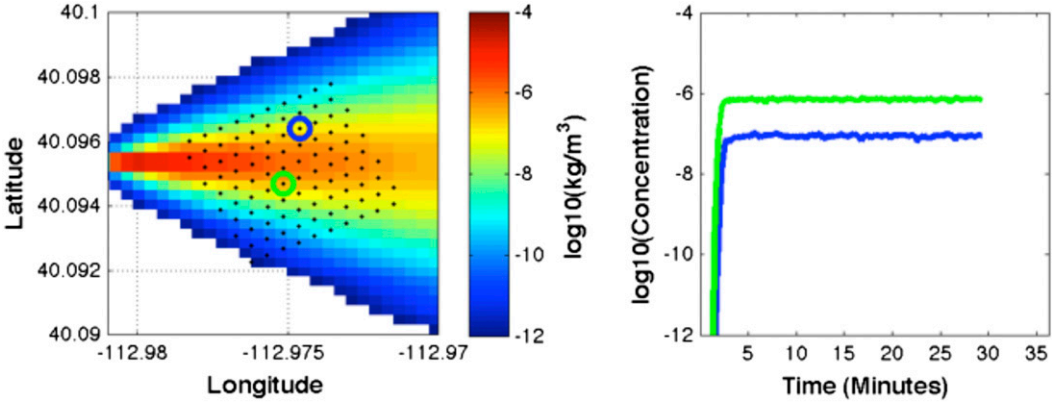

FIG. 5. As in Fig. 4, but for ensemble-average data produced by the SCIPUFF AT&D model.

to spatial correlations in the synthetic sensor measurements that also lack realism. For this reason, the ensemble-average AT&D solution may be of limited use for the networked sensor detection application where such correlations are important.

The impact of using ensembles of single-realization solutions versus ensemble-average AT&D solutions (like those generated by models like SCIPUFF) for sensor-network detection can also be quantitatively shown. We will do this by calculating a variant of Moran's I metric, which gives a measure of spatial autocorrelation (Moran 1950). This Moran I metric provides a weighted sum of the outer product of the measurement deviations from the mean, scaled by the squared sum of the sensor measurement deviations from the mean:

$$I = \frac{N}{\sum_{i=1}^{N}\sum_{j=1}^{N} W_{ij}} \frac{\sum_{i=1}^{N}\sum_{j=1}^{N} W_{ij}(x_i - \bar{x})(x_j - \bar{x})}{\sum_{i=1}^{N}(x_i - \bar{x})^2}, \qquad (1)$$

where N is the number of observations; w_{ij} represent the weights, which are specified by the user; and x_i and x_j are the observed values designated by indexes i and j, respectively, which vary between 1 and N. The value of this metric varies over $[-1, 1]$, wherein a positive value indicates positive autocorrelation, a negative value indicates negative autocorrelation, and a value of 0 indicates that the measurements have no spatial autocorrelation (Moran 1950). Instead of using (1) as specifically defined by Moran (1950), we instead transform the inputs (i.e., x_i, x_j, and $\bar{x}$) by subtracting the time series mean and normalizing by the time series standard deviation at the measurement location. With this modification we can compute the instantaneous spatial autocorrelation of the number of standard deviations that the simulated measurement is from the time series mean. For consistency with the qualitative example above we use the same sensor grid configuration and AT&D solutions as are shown in Figs. 4 and 5. For this calculation, we select three more data points in addition to those used for the qualitative analysis shown in Figs. 4 and 5. These data points were picked such that there is a minimum distance of 150 m between the sensor locations and a maximum distance between the sensor locations of 400 m. The contaminant release location is 850 m upwind from the mean downwind location (e.g., mean location for the Moran I metric) of the sensors. At this downwind distance, the single-realization simulation plume width will be smaller than the spatial scale of the largest eddies, resulting in a meandering of the plume over the chosen data points. The plume meandering creates periods of time during which the simulated measurements are correlated (i.e., when the plume is situated over the whole sensor grid) as well as time periods during which the simulated measurements are uncorrelated (i.e., when the plume is over only a portion of the sensor grid). This plume meandering is not present in the ensemble-average solution generated with the SCIPUFF model. Because the temporal gradients are smoothed by averaging, the sensor measurements provided from the ensemble-average solution will be highly correlated and this correlation will remain constant in time.

To simulate instrument error, we also add white noise to the sensor observations prior to calculating the spatial autocorrelations among the sensors in the network. This is done for both the single-realization and ensemble-average AT&D solutions. The magnitude of the instrument error is such that the average signal-to-noise ratio throughout the simulation is 10. For the synthetic observations derived from both AT&D solutions we expect the sensor autocorrelation metric to vary chaotically in time as a result of the white-noise signal. The single-realization plume, however, will have two dominant modes of variability: one mode associated with the shorter-time-scale variability from the white noise and a second mode associated with the longer-time-scale variability from plume meandering. Because the ensemble-average solution has only the mode associated with the white-noise signal, we expect to see trends in a time series of spatial autocorrelation using the single-realization solution that will not be present in the spatial autocorrelation using the ensemble-average synthetic observations. Further, the upper limit of these trends from the single-realization solution will correspond with the time scale of the largest eddies provided by z_i/w_*, where z_i represents the boundary layer depth and w_* represents the convective velocity scale (Deardorff 1970; Kaimal et al. 1976). In this case, the values of the boundary layer depth and the convective velocity scale for the single-realization simulation are approximately 1200 m and $2\,\mathrm{m\,s^{-1}}$, respectively. This should result in trends in the spatial autocorrelation that persist for time periods of no longer than 10 min.

The spatial autocorrelation is calculated at the temporal resolution of the simulated data, in this case 15 s (note that this is not the model time-integration step, but rather it is the frequency at which data were written from the LES model). Because the value of the temporal derivative over time can show consistencies or inconsistencies in the time series, we display the temporal derivative of this spatial autocorrelation in Fig. 6. When the temporal derivative is close to zero, there is little

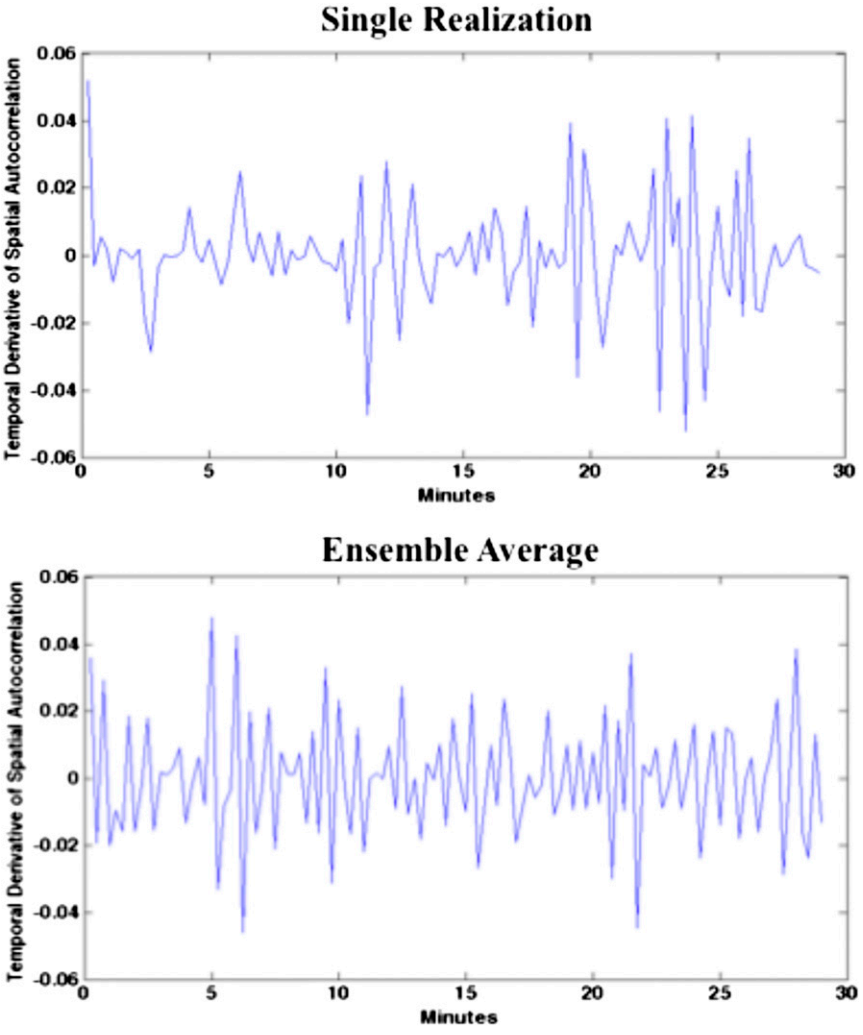

FIG. 6. Calculations of the temporal derivative of spatial autocorrelation for (top) the single-realization plume and (bottom) the ensemble-average plume.

change in the spatial autocorrelation over time, and the opposite is true when the derivative behaves erratically and/or deviates from zero. Figure 6 shows how the temporal derivative of spatial autocorrelation behaves erratically throughout the ensemble-average solution for this example because the signal is dominated only by the sensor (white) noise. As a result, there is no information in the spatial autocorrelation signal to suggest that a contaminant is present. For the single-realization solution, the temporal derivative of spatial autocorrelation has small fluctuations near zero (i.e., minutes 7–11 and 15–20) and other periods during which the derivative exhibits a more random behavior (as seen in the average solution). The former coincides with downdrafts spreading the contaminant over the sensor grid, and the latter coincides with significant meandering of the contaminant plume over the sensor grid.

Many network-based detection algorithms like those described in Norige et al. (2009) exploit spatial correlations from multiple sensors to determine whether a contaminant is present. Therefore, to test and evaluate such algorithms, it is critical that the synthetic observations derived from the simulations have realistic spatial correlations. If sensor-network false-alarm algorithms, like those in Norige et al. (2009), were to be extended and tested using synthetic observations produced by a dispersion model, the example above illustrates that the lack of realistic correlations in the ensemble-average model will lead to incorrect conclusions on sensor-network detection performance. In this case, the single-realization plumes are required to provide the necessary synthetic measurement fidelity for false-alarm mitigation that the ensemble-average plume cannot.

b. Temporal correlations: Time-dependent analyses

In this section, we describe the impact of the simulation's fidelity on defense analyses that involve the use of simulated concentration values at two different spatial locations at different times. Here, the locations of interest are the location at which the detection is made and a separate area that we wish to protect. The distance between these two locations is sufficiently large that the flow field at these two locations is not correlated. In this example, we use a measure of effectiveness (MOE) for a sensing technology that can be represented by

$$\mathrm{MOE} = \frac{\sum_{i=1}^{N} T_i X_i}{\sum_{i=1}^{N} T_i}, \tag{2}$$

where $X_i = 1$ indicates that the threat is detected so that time is available for a response and $X_i = 0$ indicates that there is insufficient time to enact a response. In (2), N represents the number of instances for which the AT&D solutions are run past the sensing technology that necessitate a protective response at the area of interest. These AT&D solutions can represent a set of ensembles for which each ensemble includes AT&D realizations for a given meteorological scenario and/or operational scenario (Lawrence et al. 2013). The variable T_i is assigned a value of 1 when an AT&D solution produces a threat. A valid threat, in this case, is defined as the existence of CB-agent concentration values above a given threshold at the area we wish to protect. The threshold used varies with the material and is often determined by exceedance of a concentration value that will necessitate a protective response (such as putting on a mask, using ''mission-oriented protective posture'' equipment, or moving to a new location). For the examples used in this section, which are meant solely to illustrate how the numerical AT&D solution can affect sensor MOE, we use an arbitrary threat threshold. If an AT&D solution produces a threat at the location we wish to protect, we next determine whether there was a corresponding prior detection of that threat. If the sensing technology detects the contaminant, and from this detection there is sufficient time to enact a protective response, then the value of variable χ_i is 1. If the technology detects the contaminant but not with enough time to provide a warning for a protective response, or if it simply fails to detect the contaminant, then the value of χ_i is set to 0. Using this logic, we then compute the sensor MOE by taking the summation over an ensemble of AT&D solutions. The solution to (2) is the probability that the sensing technology provides sufficient warning time for the approaching CB threat. In the following discussion, we will not provide an explicit calculation of (2); instead, we demonstrate how the numerical AT&D solution used for the analysis can affect the resulting MOE value.

Similar in many respects to the previous examples, this example includes two numerical AT&D solutions of an instantaneous, point contaminant release of an arbitrary agent in an unstable boundary layer. One solution is a single-realization simulation, and the other is an ensemble-average simulation. Figure 7a displays the simulation setup for the single-realization scenario. The red splotch denotes the contaminant release location, the magenta dot represents the location of the sensing technology, and the green dot represents the location of the region that we wish to protect. The contaminant is released upwind of both locations and passes over both locations as it moves downwind. To test the effectiveness of the sensing technology for providing sufficient warning time, we first determine whether the contaminant presents a threat to the area that we wish to protect. This can be determined through an examination of the contaminant concentration time series within this area (Fig. 7c). In this example, the concentration values within the area of interest exceed the threat threshold (denoted by the red line) and hence the simulation is included for the detection analysis. Next, we determined whether there is a corresponding detection for the threat and the amount of time available from that detection to enact a protective measure before the contaminant arrives at the area of interest. Because the single-realization solution is driven by a turbulence-resolving LES, contaminant AT&D is temporally correlated, hence enabling a direct link between the threat and detection. In our example case, the simulated measurements exceed the detection threshold (our criterion for a detection in this sample) and the sensor is able to detect the contaminant. The detection occurs approximately 18.5 min prior to the threat exceedance at the area that we wish to protect, and hence any response that takes less than this time can be enacted. This example illustrates the analysis process by using just one AT&D single-realization solution to compute the sensor MOE. As discussed previously, an analyst will need to repeat this procedure for many more AT&D simulations to provide a representative and stable value for sensor MOE.

Figure 8a displays the identical analysis design as shown in Fig. 7a; in this example, however, we use the ensemble-average AT&D solution generated using the SCIPUFF model. The simplest approach to calculating the MOE when using the ensemble-average AT&D solution is to look solely at the ensemble-mean

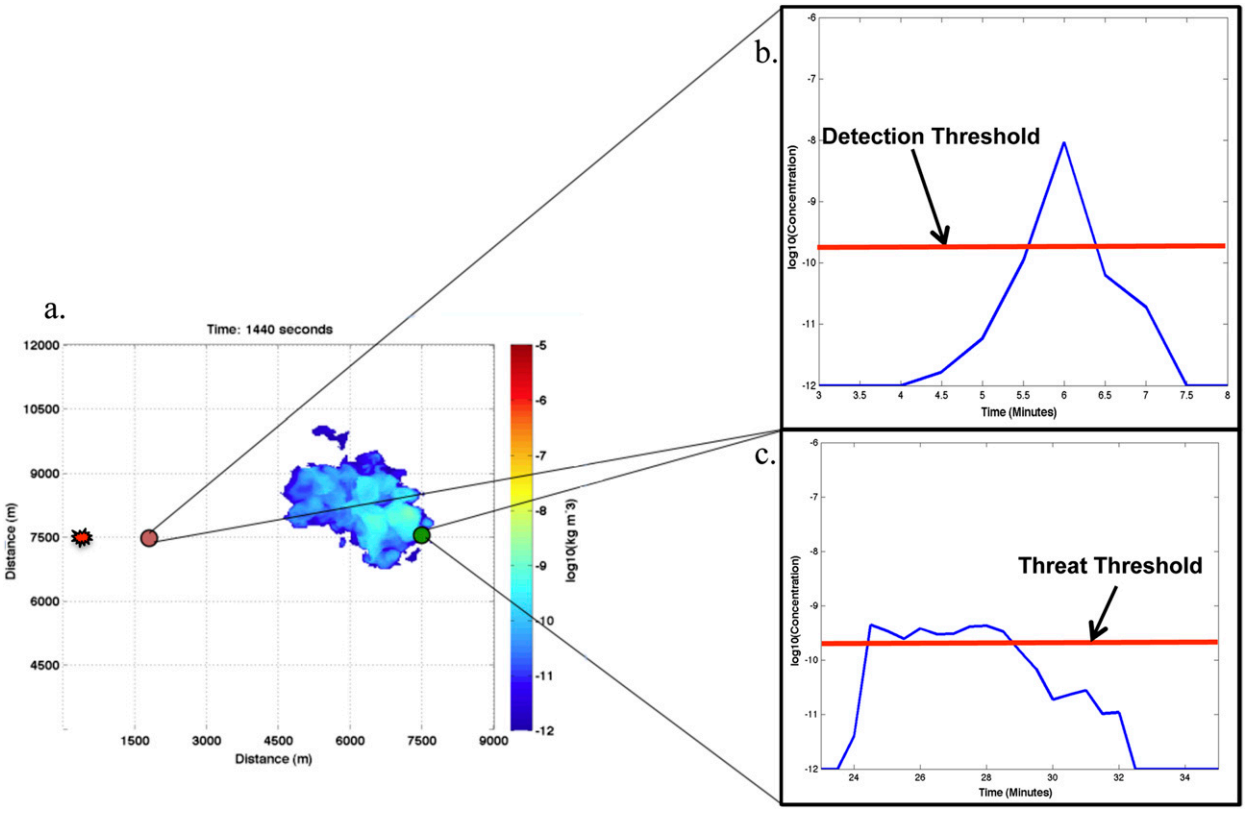

FIG. 7. Experimental setup for analyses involving time using a single-realization representation of CB-agent AT&D: (a) a snapshot of the contaminant (24 min after release from the red splotch) and the location of the sensing technology (magenta dot) and the protective position (green dot), (b) a time series of concentration values at the location of the sensing technology, and (c) a time series of concentration values at the location of the protective position.

concentration values. The ensemble-mean concentration values within the area of interest and for the simulated measurements are represented by the black curves in Figs. 8c and 8b, respectively. Because the mean value of contaminant dispersion is temporally correlated, it is possible to directly link the threat with the detection, as was done in the example that used the single-realization plumes. Although one can utilize the ensemble-mean solution for calculating sensor MOE, the ensemble-mean solution only provides a single outcome to calculate the MOE for a given operational and meteorological scenario, when in reality there exist many more possible outcomes (with higher and lower concentrations at the detection location and the area that we wish to protect). The ensemble-average solution also lacks realistic correlations in the agent concentration and hence, for the same reasons as described in the previous section, may not provide the correct outcome. Further, in this example, the ensemble-mean solution would not even be included in the MOE calculation because the concentration lies below the threat threshold.

The ensemble variance (if available) can be utilized to address these limitations with the ensemble-mean AT&D solution (Carrano and Jeys 2010). SCIPUFF is one such dispersion model that can be used in this way because it uses a clipped normal distribution to describe the concentration PDF (i.e., a combination of the ensemble mean and variance provides a means to describe the PDF of contaminant concentration values at a given point). An example of this is illustrated in Figs. 8b and 8c, where the 5th percentile and 95th percentile of the concentration values are described by the green and blue lines, respectively.

There are two methods to utilize the ensemble mean and variance to better characterize the sensor MOE. The first method is to sample stochastically the concentration PDF at each time step to provide concentration values at the respective locations and times (Carrano and Jeys 2010). Figure 8c illustrates this process within the area of interest for several time steps, as depicted by the black dots on the dashed lines. The concentration values at that point are determined by stochastically sampling the concentration PDF at each time, and the

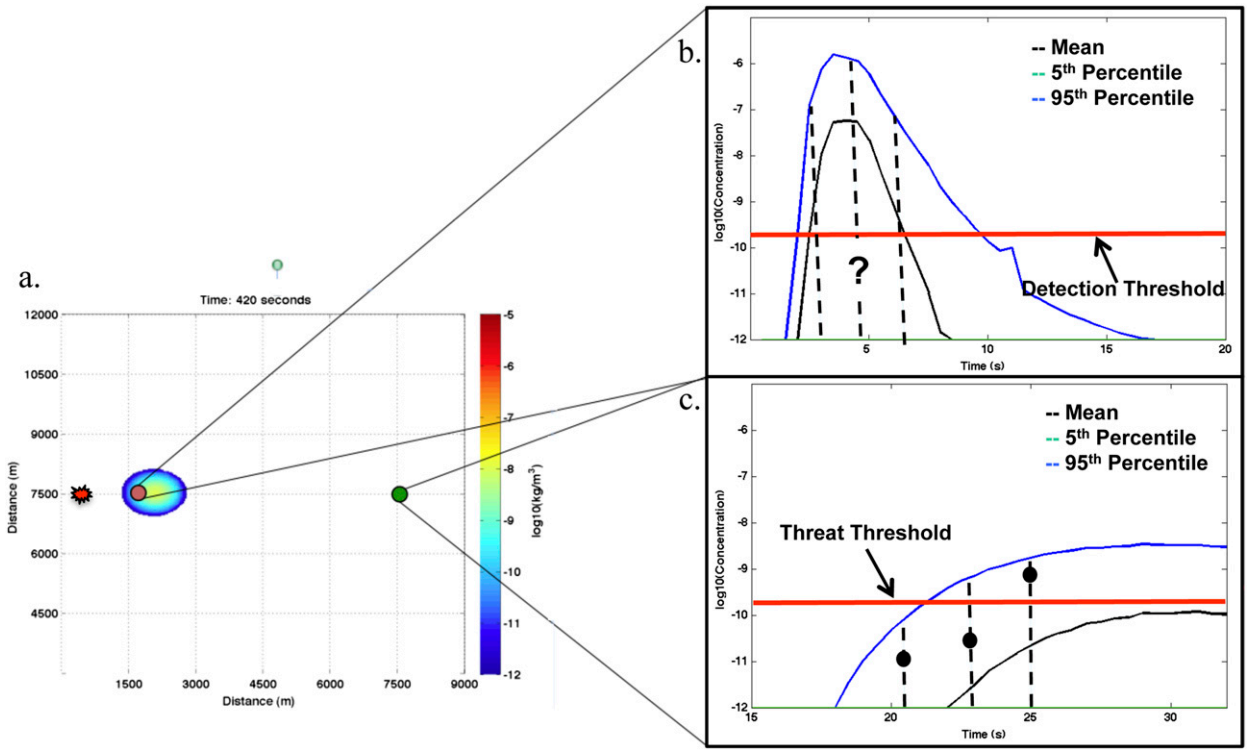

FIG. 8. As in Fig. 7, but using an ensemble-average representation and the contaminant snapshot in (a) is at 7 min after release. The additional content in (b) and (c) is described in the text.

analyst can investigate when (or whether) the concentration values exceed the threat threshold. Upon exceedance of the threat threshold, one must determine whether the sensing technology was able to warn for that threat. As shown in Fig. 8b, we sample the concentration PDF at each time step at the instrument location to produce the simulated measurements. This process, however, results in physical inconsistencies between the simulated measurements and the threat diagnosis because the simulated concentration values at these locations and times lack realistic correlations.

Instead of randomly sampling the concentration PDFs to determine a linkage between the threat and the detection, an alternative method is to calculate probabilities that exceedance of the threat threshold and the detection occur over a specific time interval. In this analysis process, these probabilities again cannot be linked back to a specific AT&D realization and thus do not contain the full distribution of spatial and temporal correlated solutions necessary to provide a precise warning time required to compute the sensing technology's MOE.

The example in this section demonstrates how an analysis of point sensor performance is influenced by the AT&D solution used to generate the simulated measurements. Namely, the ensemble-average solution lacks realistic spatial and temporal correlations that are needed to more accurately determine the sensor MOE. If one were to consider an analysis for standoff sensor detection, in which a remote sensing technology interrogates an environment for anomalies in returned signals, then these effects are compounded since the sensor measurement includes area and volume measurements. As a result, when determining a sensor's ability to warn for downwind response, it is necessary to utilize the single-realization AT&D solutions that are able to capture realistic spatial and temporal concentration gradients. As described earlier, the example used just one AT&D single-realization solution to compute the sensor MOE to illustrate the method. An analyst will need to repeat this procedure for many more AT&D simulations to provide a representative and stable value for sensor MOE. The study by Lawrence et al. (2013) is an example of a CB defense analysis that employed this method.

5. Scenario 3: Nonlinear transformations

In this section, we discuss analyses such as human casualty and injury estimates that are a function of time-integrated exposures, commonly referred to as toxic load. CB defense-analysis results for studies that include

toxic-load calculations can be largely influenced by the fidelity of the AT&D model used for the calculation. Toxic exposure to a chemical through inhalation, upon exceeding a certain value, can cause adverse health effects. Toxic exposure calculations originally were assumed to be a linear function of the total inhaled amount—an assumption that leads to Haber's law, in which contaminant toxic exposure is proportional to the integral of concentration values at a point over a period of time (NRC 2012; Sommerville et al. 2006, 2009; Urban et al. 2014):

$$\mathrm{TL}(x,t) = \int_0^t C(x,t')\,dt'. \tag{3}$$

Equation (3) implies that toxic-exposure calculations are dependent on the total amount inhaled and are independent of the manner in which the contaminant is inhaled. In other words, (3) states that inhaling a small amount of contaminant over a long period of time will have a health effect that is equivalent to inhaling a large amount of the chemical over a small amount of time. This situation is contradictory to observations for most chemicals, and health effects have been shown to be dependent on how the contaminant is inhaled. From experiments, a modification to the toxic exposure was proposed (called the toxic-load model) using a toxic-load exponent (ten Berge and van Heemst 1983; ten Berge et al. 1986; Sommerville et al. 2009):

$$\mathrm{TL}(x,t) = \int_0^t C^n(x,t')\,dt'. \tag{4}$$

The value of the exponent varies with the toxicity of the material, and this formulation makes toxic-load calculations sensitive to contaminant concentration fluctuations that result from atmospheric turbulence. For this reason, toxic-load estimates that adopt the approach specified in (4) are dependent on whether a dispersion model explicitly resolves these contaminant concentration fluctuations or provides an ensemble representation of the contaminant concentration (Urban et al. 2014). We note that when $n = 1$ the toxic-load model reduces to Haber's law.

For toxic-exposure calculations that follow Haber's law, the commutativity of linear operators will apply. When an operator acting on contaminant concentration is linear, the linear operator and the averaging commute, yielding

$$\frac{1}{N}\sum_{i=1}^{N} L(C_{R_i}) = L\left[\frac{1}{N}\sum_{i=1}^{N} L(C_{R_i})\right] = L(C_{\mathrm{EM}}), \tag{5}$$

where L represents a linear operator, C_{EM} is the ensemble-mean contaminant concentration, and C_{R_i} represents the ith ensemble member. The same, however, does not hold true when an operator is nonlinear. A nonlinear operator will not commute with the average, and thus

$$\frac{1}{N}\sum_{i=1}^{N} \mathrm{NL}(C_{R_i}) \neq \mathrm{NL}\left[\frac{1}{N}\sum_{i=1}^{N} (C_{R_i})\right] = \mathrm{NL}(C_{\mathrm{EM}}), \tag{6}$$

where NL represents a nonlinear operation. Equation (6) states that when nonlinear operators are present the ensemble-mean solution will provide a different answer than the mean toxic-load estimate provided from an ensemble of single realizations of contaminant dispersion.

To demonstrate the effects mentioned above, we calculate toxic-load estimates using 20 continuous-release, single-realization plumes, using the mean of these 20 toxic-load estimates, and using one toxic-load calculation with SCIPUFF, our model for producing an ensemble-average model solution. The 20 single-realization solutions result from staggering contaminant releases within an LES of a convective boundary layer in both space and time. We are able to stagger these releases in space because of the assumption of a horizontally homogeneous lower boundary and in time because of statistical stationary in the turbulence statistics. By staggering the releases, each single-realization AT&D solution will be different because these solutions are produced with a turbulence-resolving LES model. The source term for all 20 plumes used in this example was a $1\,\mathrm{kg\,s^{-1}}$ continuous release over a 30-min period. The toxic-load calculations were computed approximately 0.5 km downwind from the source locations. The size of the ensemble was qualitatively determined on the basis of the number of single-realization ensemble members that are required to produce an average dispersion that matches the SCIPUFF-derived average solution.

Figure 9 displays the results when the toxic-load exponent is 1. The black line in Fig. 9 represents the toxic-load estimate calculated using the SCIPUFF ensemble-average AT&D solution, the blue dotted lines represent the toxic-load solutions from each of the single-realization plumes, and the red line represents the average of the single-realization toxic-load estimates (i.e., the average of the toxic-load solutions represented by the blue dots). The toxic-load calculations from the individual single realizations differ considerably from the toxic-load calculation using the ensemble-average AT&D solution. The mean toxic load of the single-realization estimates closely matches the toxic-load calculation computed using the ensemble-average model, however. This

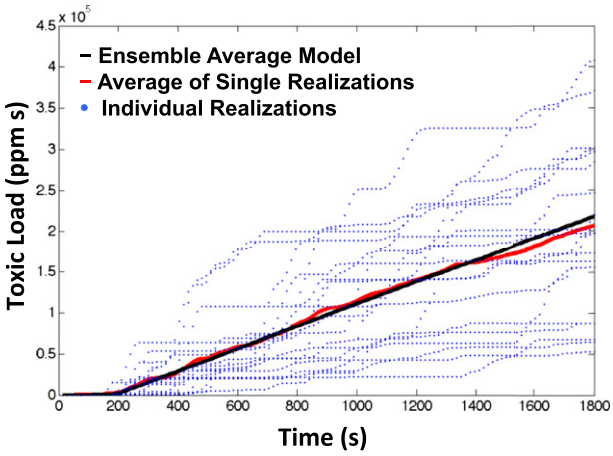

FIG. 9. Toxic-load calculations for a toxic-load exponent equal to 1. The black line represents the ensemble-average solution, the red line represents the average toxic-load calculations among 20 single-realization plumes, and the dotted blue lines represent the toxic-load calculations from each of the single realizations.

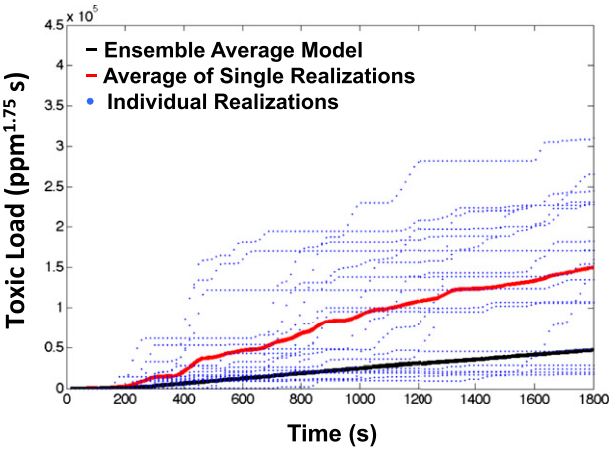

FIG. 10. As in Fig. 9, but for a toxic-load exponent equal to 1.75.

result is consistent with (4) because integration is a linear operator and the toxic-load exponent is unity, and hence all operations commute. The small deviations between the average of the single-realization toxic-load estimates and the toxic-load estimate calculated from the ensemble-average AT&D solution (the red and black lines, respectively) are due to the fact that there are not enough ensemble members to have absolute convergence between the ensemble-average solution from SCIPUFF and the mean toxic-load estimates from the single-realization plumes.

Next, we examine the scenario in which the notional toxic-load exponent of the contaminant is 1.75. From (4), we would expect the toxic-load calculations to be more sensitive to peak concentration values. In section 2 we described how the use of the ensemble-average model has the effect of smoothing the concentration field and reducing the peak concentration values relative to single-realization AT&D solutions generated by the LES model. This change to the toxic-load exponent results in a large impact on the toxic-load calculations (Fig. 10). In Fig. 10, the black line again represents the toxic-load estimates using the ensemble-average model, the dotted blue lines represent the toxic-load estimates from the 20 single-realization AT&D solutions, and the red line describes the mean of the 20 single-realization toxic-load estimates. For this scenario, the averaging and the nonlinear operator do not commute, and the toxic-load estimate calculated using the ensemble-mean AT&D solution differs considerably from the toxic-load estimates calculated using each of the 20 single-realization plumes and the mean of these single-realization toxic-load estimates. In fact, the latter is nearly 2 times the value of the former. This toxic-load example illustrates the effect of not properly incorporating the peak concentration values and has potential implications in any analysis that involves estimating human effects. When making human injury or casualty estimates, the inclusion of the spatial distribution of the dispersion, which is highly nonlinear, compounds the problem described above and leads to the potential for an underprediction of the consequences of the material release (Platt et al. 2013).

Contaminant toxic load provides an example in which it is essential to use a single realization instead of an ensemble-mean solution when an analysis involves nonlinear operations performed on the concentration field. Nonlinear transformations are present in many defense analyses (such as those involving standoff detection). Therefore, it is likely that when these nonlinear transformations are present the use of ensemble-average solutions can provide misleading results.

6. Conclusions

In this paper, we provide the motivation for the use of single-realization CB-agent concentration AT&D representations for defense analyses over the ensemble-average representation of CB-agent AT&D when the following factors are important: high-frequency sampling of a contaminant, spatial and temporal correlations, and nonlinear operations performed on the contaminant concentration field. Sections 3 and 4 use contaminant detection to demonstrate how this is important for high-frequency sampling and spatial and temporal correlations, respectively. This concept was initially motivated by the work of Warner et al. (2002), in which an ensemble

of AT&D solutions driven from numerical weather predictions provided statistics on contaminant dosage in a region of interest. In each section, direct comparisons were made between the single-realization representation and the ensemble-average representation of CB-agent AT&D for each respective example. These examples reveal that conclusions drawn from a CB defense analysis are potentially influenced by the class of dispersion model used for the analysis. Further, these examples show that the use of the single-realization representation of CB-agent AT&D over the ensemble average can provide more accurate results for CB-system performance analyses in some instances.

We also reiterate that the single-realization representation of the CB-agent AT&D is just a single snapshot of what the dispersion pattern may look like. The information content from a CB defense-system analysis *based on just* a single-realization representation of CB-agent AT&D is not sufficient to draw comprehensive conclusions. While some of the examples illustrated the differences between just one single-realization AT&D solution and the ensemble-average solution, it is implied that all defense analyses that utilize this single-realization AT&D modeling solution must utilize ensembles. This concept of using an ensemble of AT&D solutions is highlighted in section 5 with toxic-load estimates. Here, toxic-load estimates from a single realization have a very wide range of values, but the mean over many realizations is the most likely answer that is sought by the decision makers.

The case used only 20 realizations for just one meteorological scenario to populate the ensemble. Most CB defense analyses will also require ensembles that include a variety of source terms, operational scenarios, and meteorological conditions. Further, additional sources of uncertainty can be incorporated in each ensemble through model physics diversity to be representative of the full distribution of possible solutions. A similar method in which ensembles of single-realization CB threats were used to evaluate the technology trade space for standoff CB-agent detection systems was conducted by Lawrence et al. (2013). For these types of analyses, only by generating and utilizing a sufficient ensemble is it possible to develop the accurate and consistent conclusions (with potential error bounds) that are required by decision makers who rely on these defense analyses.

Acknowledgments. This work is dedicated to the late Tom Warner, whose early contributions to this work were invaluable. The authors also thank Jeff Weil, Steve Warner, and Sue Haupt for helpful conversations throughout the development of the material used in this paper.

REFERENCES

Bieberbach, G., P. E. Bieringer, R. Cabell, J. Hurst, J. Weil, A. Wyszogrodzki, and J. Hannan, 2010: A framework for developing synthetic chemical and biological agent release data sets. Extended Abstracts, *13th Int. Conf. on Harmonisation within Atmospheric Dispersion Modeling for Regulatory Purposes,* Paris, France, Initiative on Harmonisation within Atmospheric Dispersion Modelling for Regulatory Purposes, 5 pp. [Available online at ftp://ftp.rap.ucar.edu/pub/sdf/JBTDS/HARMO13_H13-248_extended_abs_Bieberbach_etal.pdf.]

Briggs, G. A., 1993: Final results of the CONDORS convective diffusion experiment. *Bound.-Layer Meteor.,* **62,** 315–328, doi:10.1007/BF00705562.

Carrano, J., and T. Jeys, 2010. Chemical and Biological Sensor Standards Study II. Defense Threat Reduction Agency Rep., 76 pp. [Available online at http://www.dtra.mil/docs/system-documents/Chem_Bio_Sensor_Standards_Study_Vol_2_Oct_2010.pdf?sfvrsn=0.]

Deardorff, J. W., 1970: Preliminary results from numerical integrations of the unstable planetary boundary layer. *J. Atmos. Sci.,* **27,** 1211–1213, doi:10.1175/1520-0469(1970)027<1211:CVATSF>2.0.CO;2.

——, 1972: Numerical investigation of neutral and unstable planetary boundary layers. *J. Atmos. Sci.,* **29,** 91–115, doi:10.1175/1520-0469(1972)029<0091:NIONAU>2.0.CO;2.

Donaldson, C. du P., and D. A. Haugen, 1973: Atmospheric turbulence and the dispersal of atmospheric pollutants. *Workshop on Micrometeorology,* D. A. Haugen, Ed., Amer. Meteor. Soc., 313–390.

Haupt, S. E., A. J. Annunzio, and K. J. Schmehl, 2013: Evolving turbulence realizations of atmospheric flow. *Bound.-Layer Meteor.,* **149,** 197–217, doi:10.1007/s10546-013-9845-7.

Kaimal, J. C., J. C. Wyngarrd, D. A. Haugen, O. R. Cote, and Y. Izumi, 1976: Turbulence structure in the convective boundary layer. *J. Atmos. Sci.,* **33,** 2152–2169, doi:10.1175/1520-0469(1976)033<2152:TSITCB>2.0.CO;2.

Lawrence, W. G., and Coauthors, 2013: Scientific evaluation of technology for standoff detection of chemical and biological agents. Massachusetts Institute of Technology Lincoln Laboratory Project Rep. CB-3, 185 pp.

Lewellen, W. S., 1977: Use of invariant modeling. *Handbook of Turbulence,* W. Frost and T. H. Moulden, Eds., Plenum Press, 237–280.

——, and R. I. Sykes, 1983: On the use of concentration variance predictions as a measure of natural uncertainty in observed concentration samples. Preprints, *Sixth Symp. on Turbulence and Diffusion,* Boston, MA, Amer. Meteor. Soc., 47–50.

——, and ——, 1986: Analysis of concentration fluctuations from lidar observations of atmospheric plumes. *J. Climate Appl. Meteor.,* **25,** 1145–1154, doi:10.1175/1520-0450(1986)025<1145:AOCFFL>2.0.CO;2.

Malm, W. C., D. E. Day, S. M. Kreidenweis, J. L. Collett, and T. Lee, 2003: Humidity-dependent optical properties of fine particles during the Big Bend Regional Aerosol and Visibility Observational Study. *J. Geophys. Res.,* **108,** 4279, doi:10.1029/2002JD002998.

Moran, P. A. P., 1950: Notes on continuous stochastic phenomena. *Biometrika,* **37,** 17–23, doi:10.1093/biomet/37.1-2.17.

Norige, A., J. Thorton, C. Schiefelbein, and C. Rudzinski, 2009: High density distributed sensing for chemical and biological defense. *Linc. Lab. J.,* **18,** 25–40. [Available online at https://www.ll.mit.edu/publications/journal/pdf/vol18_no1/18_1_1_Norige.pdf.]

NRC, 2012. *Acute Exposure Guideline Levels for Selected Airborne Chemicals.* Vol. 12. The National Academies Press, 334 pp.

Platt, N., D. DeRiggi, S. Warner, P. Bieringer, G. Bieberbach, A. Wyszogrodzki, and J. Weil, 2012: Method for comparison of large eddy simulation generated wind fluctuations with short-range observations. *Int. J. Environ. Pollut.,* **48,** 22–30, doi:10.1504/IJEP.2012.049648.

——, W. R. Kimball II, and J. T. Urban, 2013: The use of probabilistic plume predictions for the consequence assessment of atmospheric releases of hazardous materials. *15th Int. Conf. on Harmonisation within Atmospheric Dispersion Modeling for Regulatory Purposes,* Madrid, Spain, Initiative on Harmonisation within Atmospheric Dispersion Modelling for Regulatory Purposes, 5 pp. [Available online at http://www.harmo.org/Conferences/Proceedings/_Madrid/publishedSections/H15-75.pdf.]

Przybylowicz, E., and Coauthors, 2003: *Testing and Evaluation of Standoff Chemical Agent Detectors.* The National Academies Press, 66 pp.

Skamarock, W. C., and Coauthors, 2008: A description of the Advanced Research WRF version 3. NCAR Tech. Note NCAR/TN–475+STR, 113 pp.

Slade, D. H., Ed., 1968: *Meteorology and Atomic Energy.* Division of Technical Information, U.S. Atomic Energy Commission, 955 pp.

Sommerville, D. R., K. Y. Park, M. O. Kierzewski, M. D. Dunkel, M. I. Hutton, and N. A. Pinto, 2006: Toxic load modeling. *Inhalation Toxicology,* 2nd ed., H. Salem and S. A. Katz, Eds., CRC Press, 137–158.

——, J. J. Bray, R. E. Jablonski, S. A. Reutter-Christy, and E. E. Shelly, 2009: Review of human lethality estimates for chlorine inhalation. *Proc. Eighth Symp. on the Urban Environment,* Phoenix, AZ, Amer. Meteor. Soc., J14.4. [Available online at https://ams.confex.com/ams/pdfpapers/149892.pdf.]

Storwold, D. P., 2007: Detailed test plan for the Fusing Sensor Information from Observing Networks (FUSION) Field Trial 2007 (FFT 07). U.S. Army Dugway Proving Ground West Desert Test Center Doc. WDTC-TP-07-078, 46 pp.

Sykes, R. I., S. F. Parker, D. S. Henn, and B. Chowdhury, 2008: SCIPUFF version 2.4 technical documentation. Sage Management Enterprise, LLC, Tech Rep., 317 pp.

ten Berge, W. F., and M. V. van Heemst, 1983: Validity and accuracy of a commonly used toxicity-assessment model in risk analysis. *Fourth Int. Symp. on Loss Prevention and Safety Promotion in Process Industries,* Vol. 1, Rugby, United Kingdom, Institute of Chemical Engineers, I1–I12.

——, A. Zwart, and L. M. Appleman, 1986: Concentration-time mortality response relationship of irritant and systematically acting vapors and gases. *J. Hazard. Mater.,* **13,** 301–309, doi:10.1016/0304-3894(86)85003-8.

Urban, J. T., K. Galvin, N. Platt, P. E. Bieringer, G. Bieberbach, and A. J. Annunzio, 2014: Comparison of hazard area and casualty predictions of a small-scale chemical attack using various toxic load toxicity models. *Int. J. Environ. Pollut.,* in press.

Warner, T. P., R. S. Sheu, J. F. Bowers, I. R. Sykes, G. C. Dodd, and D. S. Henn, 2002: Ensemble simulations coupled with atmospheric dynamic and dispersion models: Illustrating uncertainties in dosage simulations. *J. Appl. Meteor.,* **41,** 488–504, doi:10.1175/1520-0450(2002)041<0488:ESWCAD>2.0.CO;2.

Wyngaard, J. C., 2010: *Turbulence in the Atmosphere.* Cambridge University Press, 393 pp.