



NRL/MR/6440--11-9326

The How and Why of Nomografs™ for CT-Analyst®

J. BORIS

Laboratory for Computational Physics and Fluid Dynamics

G. PATNAIK

K. OBENSCHAIN

Center for Computational Physics Developments

Laboratory for Computational Physics and Fluid Dynamics

May 3, 2011

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 03-05-2011		2. REPORT TYPE Memorandum Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE The How and Why of Nomografs™ for CT-Analyst®				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER 64-9087-00	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) J. Boris, G. Patnaik, and K. Obenschain				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320				8. PERFORMING ORGANIZATION REPORT NUMBER NRL/MR/6440--11-9326	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR / MONITOR'S ACRONYM(S)	
				11. SPONSOR / MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT An urban-oriented emergency assessment system, called CT-Analyst®, was developed to evaluate airborne contaminant transport (CT) threats and aid rapid decisions for regions such as cities where other methods are slow and inaccurate. CT-Analyst gives both greater accuracy and much greater speed than alternate prediction tools because it embodies entirely new principles to function in the information-starved situations that characterize the first few minutes of a terrorist or accident scenario. CT-Analyst was designed for the military prior to 9/11 to use verbal and sensor reports, to use mobile sensors, and to function in realistic situations, such as the first few minutes of a harbor spill, where information about the airborne contaminant or chemical, biological, radiological (CBR) agent is highly uncertain. These improvements are made possible by pre-computing very accurate three-dimensional flow solutions that include solar heating, buoyancy, complete building geometry, trees, and impressed wind fluctuations. Detailed 3D simulations for 18 wind directions are pre-computed for coverage regions where CT-Analyst is to be installed. CT-Analyst extends these results to all wind directions, speeds, sources, and source locations through a new data structure called Dispersion Nomografs™. We generate these “nomografs” for cities, ports, and industrial complexes well in advance so a manager in an emergency need not wait for supporting analyses. CT-Analyst also provides new real-time functions such as sensor data fusion, “backtracking” reports and observations to an unknown source location, and even evacuation route planning. The resulting capability is faster, more accurate, more flexible, and easier to use than Gaussian and particle-based dispersion models.					
15. SUBJECT TERMS Contaminant transport Dispersion modeling Emergency response					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Gopal Patnaik
Unclassified	Unclassified	Unclassified	UL	17	19b. TELEPHONE NUMBER (include area code) (202) 767-3531

The How and Why of Nomografs™ for CT-Analyst®

I. Introduction

Dispersion Nomografs™ (or just nomografs) are two compact, pre-computed contour maps (data bases) that capture the aerodynamic and turbulent effects of terrain, buildings, vegetation and surface types on the transport and dispersion of on contaminant plumes in cities. Standard, well-validated fluid dynamic principles are being used to solve the plume prediction problem in an entirely new way that directly yields the entire “hazard area” at risk of contamination. Most other methods require costly, repetitive computing of many separate contaminant parcels to build up a predicted hazard area. Further, these older methods are only as good as the accuracy of the urban airflow models that they use. With nomografs, improved accuracy and much greater speed are achieved for urban-oriented emergency assessment. By interpolating into these patented nomograf contour maps, we can perform plume predictions and related assessments in milliseconds for wide areas with complex terrain such as cities, ports, and important facilities.

The NRL high-fidelity 3D, Computational Fluid Dynamics (CFD) model called FAST3D-CT underpins our current implementation of dispersion nomografs within the CT-Analyst® software system. FAST3D-CT computes multi-gigabyte 3D, contaminant flow-path databases from which the high-resolution dispersion nomografs are extracted for a city well before CT-Analyst is ever deployed. Other models that can provide the same fluid dynamic information could, in principle, also be the source of data to build nomografs. For that matter, if enough data could be taken in field trials or experiments, equivalent to three-dimensional fields of the important flow variables over the region, nomografs could then be made from field data.

This short paper describes the principles behind the use and computation of nomografs for predicting the transport and dispersion of airborne contaminants in cities such as Hamburg. It also compares this new approach to older approximate approaches that use Gaussian puffs, Gaussian plumes, or moving particles to predict the plumes associated with contaminant releases. This comparison is offered to show that the same fundamental airflow information is absolutely necessary in all cases but that this information can be used much more efficiently and accurately in the nomograf representation under the conditions that apply in the first hour or so of a contaminant release within a city.

II. Common Aspects of Using Transport and Dispersion Models

In an accident or other emergency, the starting conditions for most airborne Contaminant Transport (CT) scenario are very poorly defined. The agent, its amount, the source location, the method of release and the release time may not be known. Nevertheless, a rapid, informed response can be crucial because roughly three-fourths of the casualties will be caused in the first 15 minutes in case of a major rupture or other rapid release. Current CT models take a few minutes to set up and run and thus are too slow to support most crises directly. They are also unnecessarily inaccurate because the building geometry and other local details are largely ignored to reduce the necessary computation. Even if one accurate plume can be computed in 5 minutes, we will need 30 minutes to process reports and sensor readings to find (“backtrack” to) an unknown source location and then to predict the plume’s path and consequences. In this time most casualties will have occurred. Furthermore, other, better data will have come in that we can’t use without starting over.

In some applications, finding the unknown location of a source from sensor readings and reports may be a crucial issue. In others, such as with ships or rail cars, moving and even multiple sources may be involved. When a source location and release time are uncertain and the winds are shifting, such as with expected biological sources, or when a number of sensors have to be placed optimally, many thousands of plume evaluations are needed. These considerations all suggest strongly that faster methods than Gaussian puff, Gaussian plume, or Lagrangian particle models are needed.

Model accuracy is just as important as speed. All methods for computing contaminant transport must know the average air (wind) velocity at each place where the contaminant goes and must have a useable model of the wind fluctuations (i.e. the turbulence and local gusts) in that region. The building geometry determines the velocities and fluctuations in cities and the detailed terrain geometry, including trees will be just as crucial away from cities. Accurate plume predictions can come only from accurate winds, which require accurate geometry. In the interests of speed, other plume prediction schemes generally use large area winds from external sources, such as weather predictions, or else simple formulae to approximate the airflow with random-walk approximations to mock-up the dispersion. Computing the detailed urban Computational Fluid Dynamics (CFD) while you wait is simply far too slow – even when you do know the source and wind initial conditions.

CT-Analyst distinguishes itself from other models by using accurate pre-computed wind fields to completely circumvent this speed-vs-accuracy dilemma. In principle, this advantage could also give more accurate predictions cheaply in other plume prediction methods but the accuracy of the wind fields they use is not their only performance drawback. The way they use the wind and turbulence fields is also very inefficient. Figure 1 below has two diagrams contrasting how the velocity and fluctuation information is used in Gaussian and particle models (left) and how the same information is used for Dispersion Nomographs (CT-Analyst – right). The small green rectangle represents a building interacting with the flow and increasing dispersion.

III. Using Wind Velocities for Gaussian Puff and Nomograf Models

In Gaussian puff models, each puff moves with the wind and has a finite size that grows from the local turbulence. Each puff center is moved during a timestep using an average wind velocity at the puff center. In the vertical this average wind velocity is usually computed using knowledge of the atmospheric velocity profile near the ground but the typical horizontal resolution of these winds is one or two kilometers. This is at best a number of city blocks, making accurate treatment of the dispersion caused by urban geometry impossible.

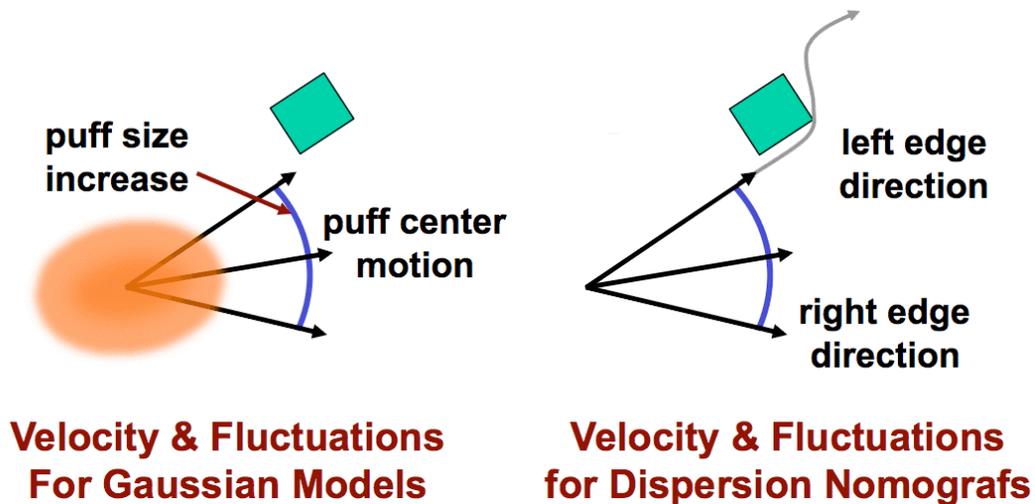


Figure 1. Gaussian puff methods (left) use the wind velocity at the puff center to determine where the center moves and use the fluctuations, represented by the blue arc, to determine how much the puff spreads during a timestep. Using nomographs, the left edge of a contaminant plume moves in the direction of the leftward turbulent deflections of the average velocity and the right edge moves in the direction of the rightward deflections. The grey line shows how the left plume edge is deflected around a building.

The Gaussian puffs each have a different, constantly changing, exponentially decreasing “Gaussian” density profile that they carry around with them. Each profile approaches zero in the outer regions of the puff, as indicated by the shading of the orange ellipse in the figure above. This puff profile is the exact solution of a diffusion equation for the puff material (gas or particles) when the wind velocity and turbulence are constant everywhere but this is only a crude approximation at best. When the puff is small, its motion and spreading results from distortions, vortices, and local shear flows and these are not well represented as diffusion. When the puff is large, diffusion becomes a more correct physical model but the turbulence and velocity are quite variable across the puff. The overall contaminant density at any time is computed as the sum of all these individual density profiles at each point.

By way of contrast, the panel on the right in Figure 1 shows how the local average velocity and turbulent fluctuations (blue arc) are used in CT-Analyst via nomographs. Two composite directions are computed for the contaminant plume, the limiting direction that the right edge would take, and the limiting direction that the left will take. These directions will be the same when the velocity is completely steady; the contaminant then moves with the wind but does not spread. However, when the actual velocity at each point is fluctuating due to turbulence, there is a maximum displacement of the contaminant to the right of the average velocity and there is a maximum displacement to the left. In this approach, the entire plume is defined between the limiting right and left edges, and these edges can be computed as simple one-dimensional integrals. Since our velocity fields, pre-computed using the FAST3D-CT code, are typically resolved to 5 or 6 meters, the geometric accuracy of these edge calculations is far higher than can be approximated by Gaussian puff, particle, or plume approximations.

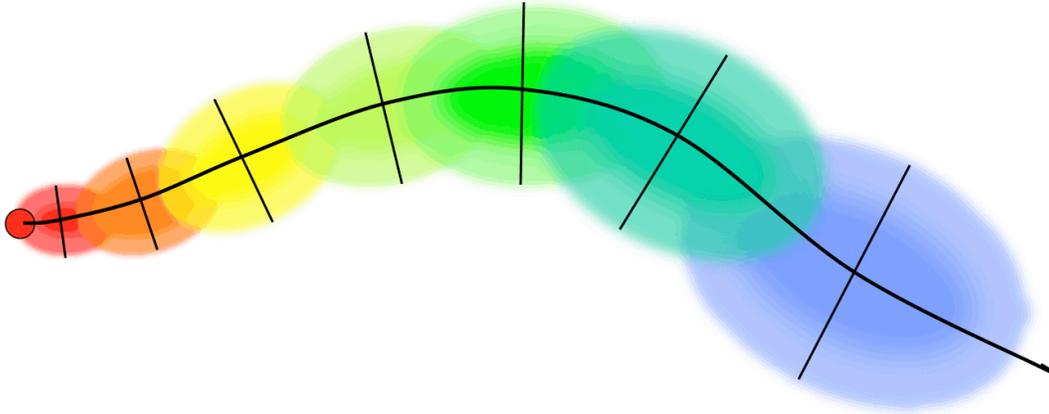


Figure 2. A sequence of Gaussian puffs original at the source circle on the left and move along the average wind direction indicated by the curved black line. These puffs spread according to the strength of the velocity fluctuations along the puff centerline. This spreading is indicated by the color change, as each succeeding puff gets larger.

Figure 2 above schematically indicates the beginning of the detailed solution procedure for Gaussian puff models. Many puffs are generated in the source region, a few each timestep, and are moved away downwind. A number of timesteps are used to track each puff as it moves and expands. These timesteps are extrapolations mathematically and so become less accurate the longer the steps are. To improve accuracy, puffs are split up when they become too large, a process that slows computation. To save unnecessary computer cycles, puffs are also merged when there are a number of them overlapping. This necessarily reduces accuracy and increases program complexity.

The important prediction in most crises is the expected hazard area to avoid. Figure 3 below shows schematically how this Gaussian puff information is used to compute a hazard area, indicated by the red line, for the overall plume. This line is generally defined where the summed density of the outermost puffs drops away below some predetermined threshold or where some health consequence is minimal. Each puff has its own density profile that can be expensive to compute; each of these profiles is fuzzy; and these profiles all overlap everywhere. Therefore, to find the hazard area for the plume, conventional models must add the profiles together at each point to get the composite density before identifying, within the fuzzy plume boundary, what defines a particular hazard area.

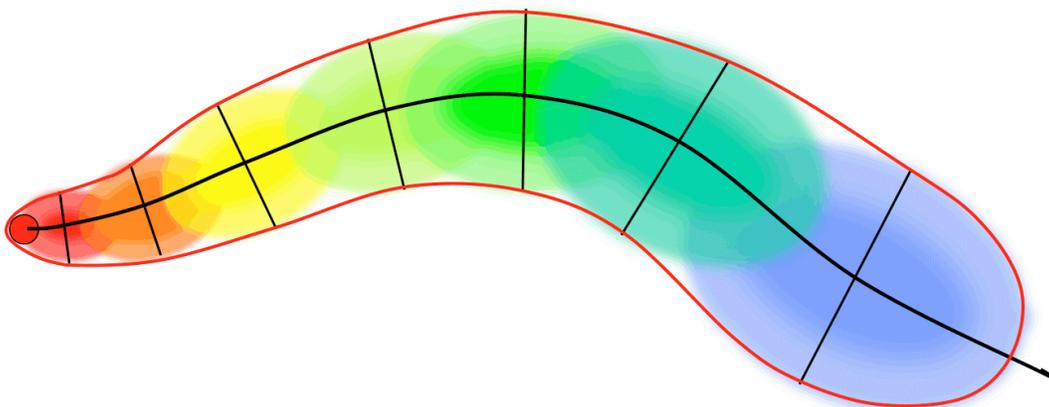


Figure 3. The hazard area for the plume is shown as a red line drawn around the summed density profile of all the puffs at a specified threshold value. This value depends on a definition of the definition of the hazard or health consequences that may require appreciable extra computation.

When different sources of contaminant interact, it is necessary to repeat the above procedure for each source and the overall procedure is very slow, generally taking minutes to compute even on today's computers. The procedure is also relatively inaccurate because the underlying mathematical assumptions are questionable in the urban context and because the wind fields are neither accurate nor well resolved.

III. Fluid Dynamic Principles Used in Dispersion Nomographs

In Figure 1 above we showed schematically how the local average velocity and turbulent fluctuations could determine paths through the urban landscape defining the left and right plume edges rather than computing the center motion and spreading of a large number of puffs. This alone is a major computational saving because the edges can be integrated more accurately at less cost when the hazard areas are computed directly. There is an additional major saving that this alternate approach makes possible. All plume edge paths, since they depend on the underlying wind and turbulence fields that have been pre-computed, can also be pre-computed, tabulated, and stored for instant look-up. Nomographs record these limiting left and right plume edge paths in a very compact data structure that serves for all source locations within the urban "canopy," (i.e. near the ground).

Here we take a couple of minutes to discuss the approximations that make this plume edge approach, as implemented in nomographs, efficient and effective. Please remember that the accuracy of the winds fields used determines the accuracy of the contaminant plume regardless of the computer algorithms used. This was discussed above with regard to the importance of accurate terrain and accurate building geometry. The fluid-dynamic principles that are central to the current dispersion nomograph representation and its subsequent applications are:

- P1. A contaminant transports dynamically via convection with the local air velocity; this is not diffusion. All relevant transport and dispersion arises from resolved (and possibly unresolved) fluid motions. Diffusion, per se, plays only a minimal role.
- P2. Vertical spreading of contaminant is quick on the building scale in an urban environment. Both simulations and field trials support this observation, allowing a two-dimensional computation of the hazard "area" as a footprint for a fully three-dimensional plume representation.
- P3. Maximum lateral dispersion of contaminants occurs near the ground because the buildings cause strong transverse turbulence down where the average wind speed is relatively slow. As a result, the direction of the asymptotic plume edges is determined entirely locally with no need to know where the source is. The further spreading of any small parcel of contaminant that arrives at a location only depends on what happens there.
- P4. Wherever the contaminant goes becomes contaminated and a volume, once contaminated, stays contaminated. This is a conservative approximation favored by first responders to aid in "safe siding" the predictions. It could (perhaps) be relaxed considerably for applications in open regions or over water.

Because fluid dynamic turbulence is not diffusion (P1), plume edges are actually quite sharp, not fuzzy as assumed with puffs, but the actual location of this sharp edge can be rather

uncertain. We may not be able to predict exactly where the edge of the plume is at any instant and we call this variability - but the edge at any instant will be quite sharp. Figure 4 below shows 8 different plumes computed 20 minutes after an SF₆ trial release over downtown Los Angeles using the same input winds and weather. Each of these instantaneous plume snapshots has a quite sharp edge although the overall plume shapes are noticeably different.

Eight different plumes result because the local wind fluctuations were all somewhat different because the contaminant clouds were released at eight different times. These are called different “realizations” and a very important conclusion is that there is no one single “unique” answer. No one of these realizations is more accurate or more “correct” than any other.

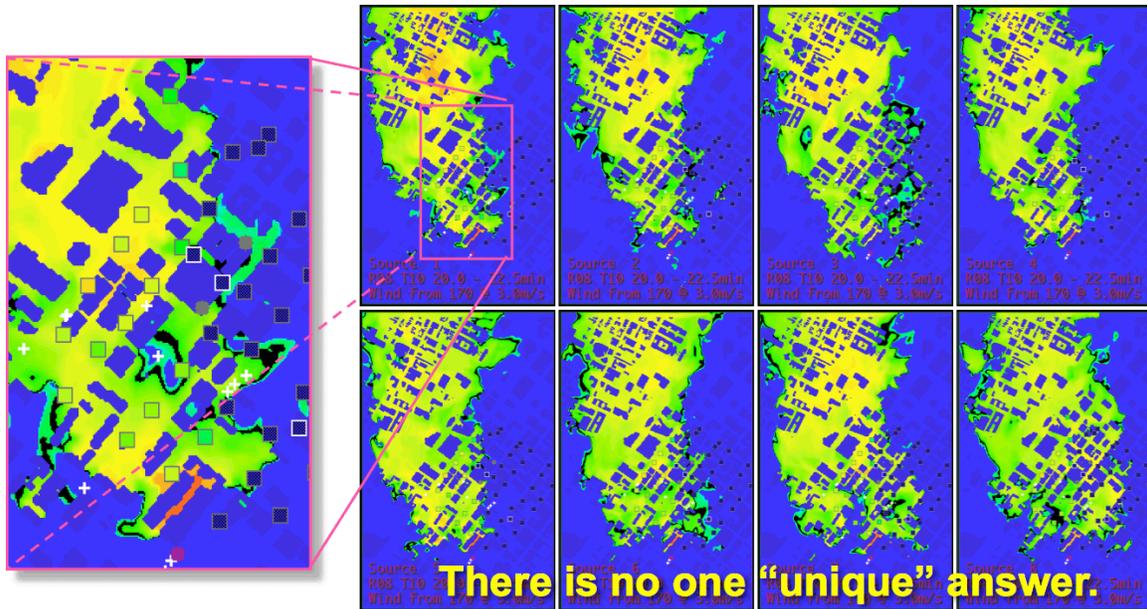


Figure 4. Eight different realizations computed for a single source. These are compared with Tracer ES&T sampler data from downtown Los Angeles (colored squares). The wind is from 170° at 3 m/s with moderate fluctuations. The point releases are 5 min in duration with measurements (2.5 min sampling intervals) at each of 50 locations within about 1 km source at the magenta circle. Estimating variability requires multiple trials. There is no correct or “unique” answer!

Another important fluid dynamic consideration in using plume edges paths (nomographs) to determine a hazard area is to decide that a point stays contaminated once it has been contaminated (P4). In light of the essentially unpredictable variations shown in Figure 4, the safest approach is to take the worst of the realizations as defining the “envelope,” the outer boundary, of the hazard area. Figure 5 below shows this plume envelope construction schematically at a particular time after the contaminant release begins. Four different yellow plumes are shown, labeled (a) to (d). Each is a possible realization of the plume from the source at the red circle. The pink “envelope” drawn in behind each of these realizations is the ideal hazard area we would like to predict if these four realizations were the only possible solutions. Each plume fills out a different fraction (percentage) of the composite hazard area and there is a probability that each point in the hazard will be contaminated by any particular realization.

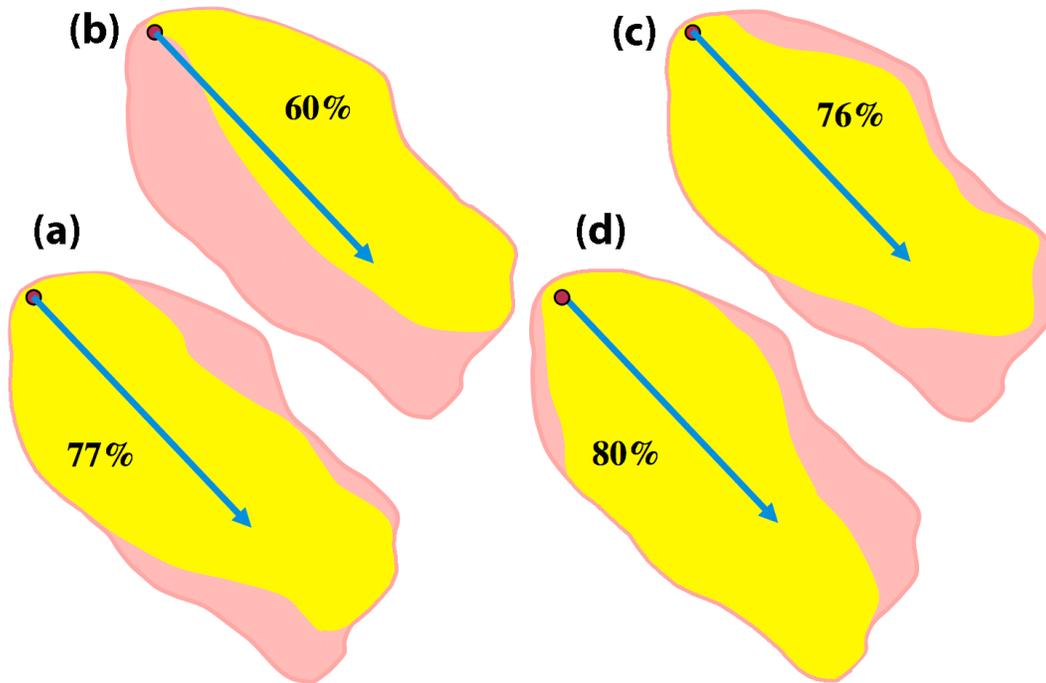


Figure 5. Four different plume realizations (yellow) of a single source are overlaid on the plume envelope, here defined as the union of all these plume realizations. Since there is no one “correct” answer, the plume envelope gives a hazard area that evolves in time and generally provides a safe, conservative estimate of the area that could be contaminated.

The concept of a two-dimensional plume envelope is useful because the maximal lateral spreading of the plume is determined near the ground (P2) and this lateral extent is not very different up to the average height of the buildings (P3). This construct “safe sides” the predicted hazard area because it is not based on an average of the realizations, as with most Gaussian averaging models, but rather is the union of worst cases. The probability that any particular location outside the plume envelope will be contaminated at time t will be small and can be made smaller by making the plume envelope a little bit larger.

IV. Computing Limiting Plume Envelope Edges: Dispersion Nomographs

This new approach, to find the limiting plume envelope edges directly, is both safer and faster. The main limitation is that the prediction is confined to the domain where the CFD computations for the nomographs have been performed and to sources within the urban canopy (i.e. near the ground). Areas as large as 40 km by 30 km have been covered with nomographs and methodology exists to extend the nomograf approach to long-time variations in the winds but consequence the predictions are generally limited to an hour or two so unless the winds are quite steady and to the first 10-20 kilometers downwind of the source.

When an agent density (concentration) is desired, a simple polynomial profile is interpolated into the three-dimensional extension of the plume envelope above the surface. Our algorithms use the Dispersion Nomograf tables to determine the spatial distribution of the agent density within the plume envelope. In this way only one simple polynomial evaluation is computed at each point inside the plume envelope and none are performed outside the envelope. Meanwhile, the nomographs transfer information about the local geometry into the approximate density profiles. This approach implicitly depends on principles P2 and P3 stated above.

To summarize this new approach schematically, look at Figure 6 below. The limiting left plume edge is computed by integrating along a direction given by taking the average wind direction near the ground and deflecting this direction to the left by the strength of the fluctuations. On the upper blue curve, labeled “left edge” in the figure, a black arrow in the local average wind direction is shown. The blue wedge at this location indicates the left and right fluctuations of the wind velocity that go into building up this average wind direction. It is clear that a different right edge can be started at each point along the left edge. The limiting right plume edge that originates at the right (lower) edge of the source is determined just as the left edge is and diverges to the right of the left edge by an amount that depends on the local fluctuations divided by the average wind speed. It is also clear that a new left edge could be started at each point along the chosen right edge.

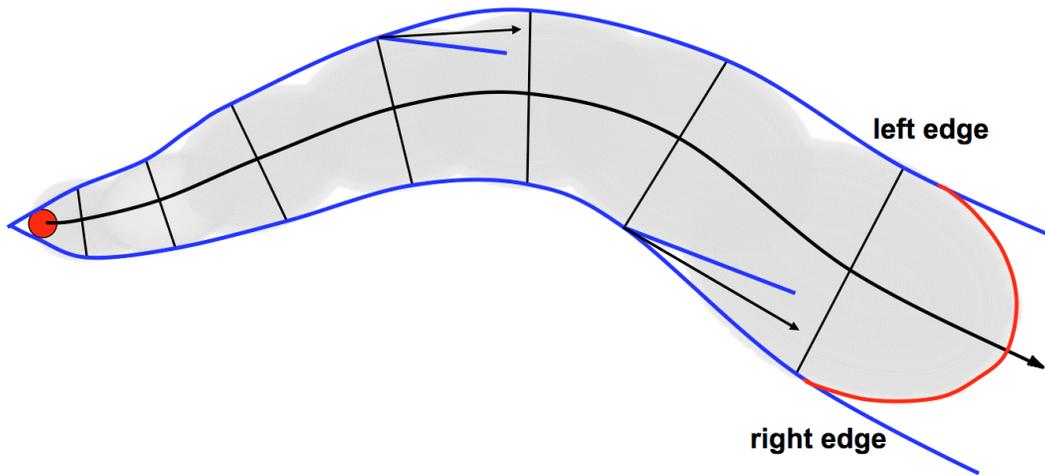


Figure 6. Computing left and right plume edges to bound the transport and dispersion of a contaminant. Individual components (puffs) in the Gaussian puff method are shown in grey to indicate the relationship between the different approaches.

Figure 6 also shows a red line used to cap off the plume envelope at any particular time. The distance of this end cap from the source depends on the time that has elapsed after the release has begun and the distribution of wind speeds over the domain. It should be obvious that these edge integrals could be performed for each source scenario separately and the resulting hazard area prediction would be a lot faster than current methods and more accurate, as long as detailed, pre-computed wind and turbulence fields are used.

As mentioned above, however, P1 and P3 taken together, mean that the limiting left and right edge directions for a small parcel of contaminant near the ground, at any place in the domain, cannot depend on where that parcel came from. As a result, the limiting left and right edges extend all the way across the domain and only need to be computed once and stored. This is what Dispersion Nomographs are designed to do (Figure 7 below). The nomographs are contour maps of the domain, one for the right edges and one for the left, where the edges originating at a location (x^s, y^s) are defined by the contour levels $R(x, y) = R(x^s, y^s)$ and $L(x, y) = L(x^s, y^s)$.

This definition is “Gauge invariant” mathematically, meaning that the particular “height” values used to label the contours are irrelevant, only the geometry of the limiting plume edge curves through the domain matters. In this approach, given a source location, we can look up the plume edges directly from the nomographs so the procedure is very fast! The nomographs can encode very accurate edge shapes and each left and right edge contour can be used for many

sources. Therefore the representation is very compact as well as very fast to apply. Calculation times for hazard areas take only a few tens of milliseconds with accurate edge shapes and computing a density profile inside the plume envelope takes only two or three times longer.

Figure 7 below indicates schematically how the same left and right edge contours can be used, even simultaneously, for multiple sources. For simplicity a single limiting left edge contour (thicker blue line) and a limiting right edge contour (thicker black line) have been duplicated above and below to show what an idealize nomograf would look like. A situation exactly like this would occur if the average wind direction and the corresponding turbulence fluctuation strength were to depend only on the horizontal coordinate (x) but not the vertical coordinate (y).

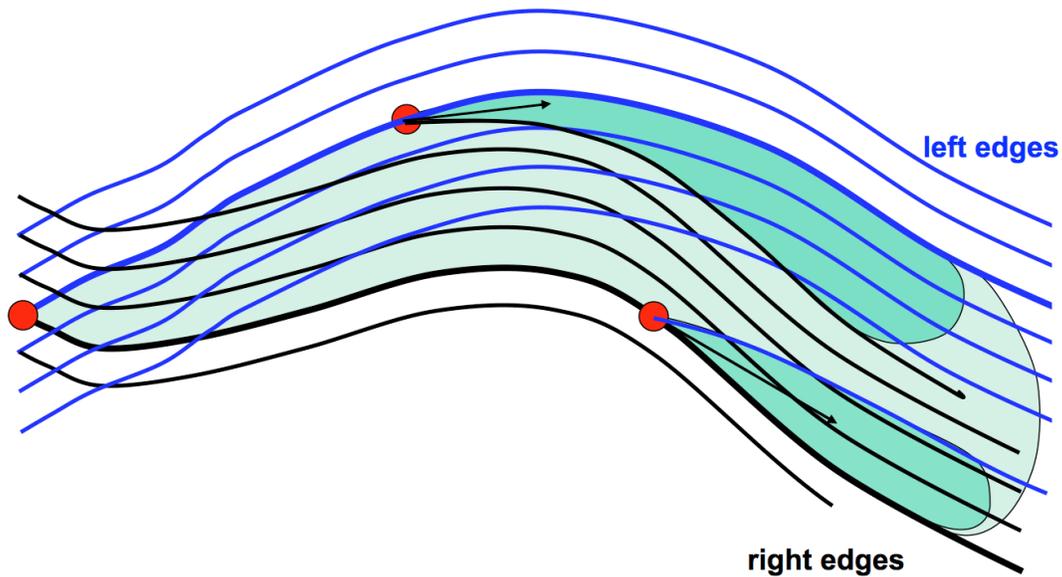


Figure 7. Left (blue) and right (black) limiting plume edge paths can be nested into a data structure containing two contour maps as shown. Because the contaminant plume edge directions are determined near the ground, this direction is the same whether near the source, in the middle of the plume, or near the downwind boundary. Three plume envelopes are shown.

Even though the wind fields leading to Figure 7 are constrained and artificial, we can still see that plumes originating in different places have different shapes. The red circles indicate three source locations. The source on the left side of the figure shares its right edge with the lower source and shares its left edge with the upper source so the resulting plume envelope (light green in the background) would eventually spread to cover the other two plumes in the upper middle and on the right. In this figure, the source on the left has been active for twice as long as the upper source and about three times as long as the lower (rightmost) source.

IV. Computing Dispersion Nomographs from FAST3D-CT Runs

The data collected for nomograf generation in a FAST3D-CT urban aerodynamics simulation include:

- D1. The three-dimensional means and standard deviations of the three velocity components of the airflow computed from data collected at seven ground-conforming height levels from the ground to above the height of the typical buildings.

D2. Six key quantities are accumulated at two different heights for each of a number of sources, that are initialized after the CFD model “spin-up” period is complete and the statistics of the fluctuation urban aerodynamics have stabilized.

D2a. First arrival time (sec) of contaminant density exceeding a threshold value at each location,

D2b. Time of arrival (sec) of the maximum contaminant value at the location,

D2c. Decay time (sec) after the maximum is reached at the location,

D2d. Local peak contaminant density (gm/m^3) at any time,

D2e. Integrated contaminant density at the location (dose = $\text{sec gm}/\text{m}^3$),

D2f. Local contaminant variability, measured as the integrated total variation (gm/m^3)

A typical run with the FAST3D-CT model for a complex urban area of 30 square km resolved with 6 m cells takes 12 hours on a 16-processor SGI computer system. This is significantly faster per square km than classical CFD models due to the savings achieved by our MILES turbulence model, our efficient treatment of the complex geometry, and other algorithmic improvements including efficient parallel processing. Because of the FAST3D-CT model efficiency, we can compute a set of dispersion nomographs for each eighteen different wind directions in a few days of computing on what is really a rather modest “supercomputer.”

The right and left dispersion nomographs are computed geometrically from a weighted average of the local wind velocity components and standard deviations using the seven height levels recorded in D1 above. The heaviest weighting is at or below the height of the typical buildings, as required by principles P2 and P3 above. While the nomographs themselves depend only on the geometric velocity data, the data items in D2 are used to calibrate the formulae and interpolations used by CT-Analyst to convert the limiting left and right plume envelope edge contours to on-screen predictions. The quantities D2 are virtually the same definitions and are computed exactly as for the short-duration releases in the Hamburg wind tunnel in the recent research projects validating FAST3D-CT for Oklahoma City trials and matching wind tunnel studies. Thus there is a rather direct trail from the experiments and the data taken there directly through the detailed CFD simulations to the CT-Analyst predictions.

Figure 8 below shows an example of the left and right edge nomograph tables for a small island domain. The color coding of the left edge nomograph (on the left) and the right edge nomograph shows that these two arrays are indeed contour maps with blue indicating low (negative) values and red indicating larger (positive) values. The wind is from the north (top) with the broad lines, alternating blue and grey, identifying a few limiting plume envelope edge contours. This visualization helps show the relationship of large lateral dispersion of the plume edges and therefore potential contaminant plumes to the underlying geometry.

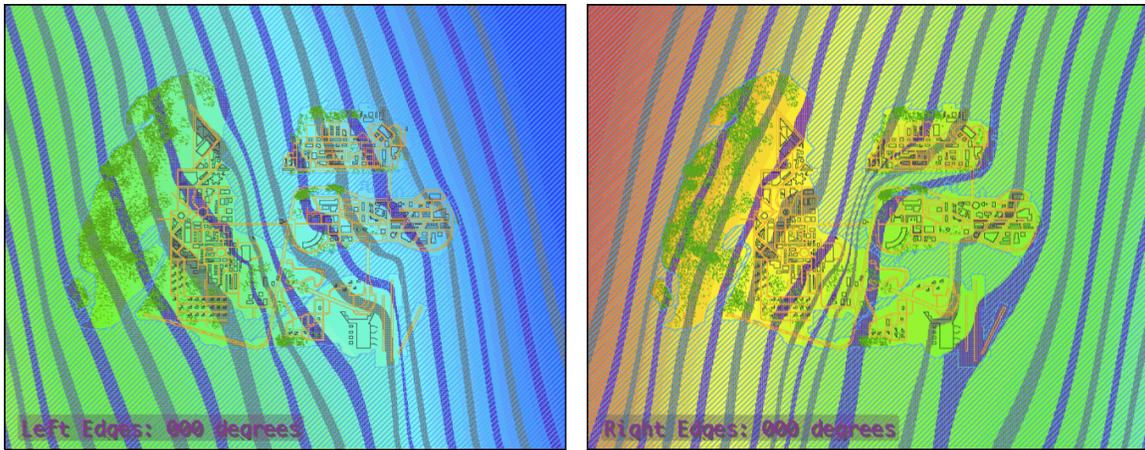


Figure 8. Left edge (left panel) and right edge (right panel) Dispersion Nomograf contours for a fictitious island kingdom called Atlantis. The wind is from the top and the alternating blue and grey contours are the nomograf plume envelope edge paths.

The four steps in generating and using dispersion nomographs are:

1. An accurate geometry database is compiled from LIDAR, stereo imagery, or shape files. The geometry database used by FAST3D-CT is a two-dimensional (typically one meter resolution) array that returns the heights of terrain, buildings, and trees, and surface composition in the computational domain.
2. Detailed 3D computational fluid dynamics calculations (FAST3D-CT) are repeated for 18 wind directions and the results are captured in an extensive database. These simulations include the appropriate urban boundary layer for the region with realistic turbulent fluctuations imposed at the inflow boundaries. Multiple releases are tracked in each case as described above.
3. The salient features from the CFD database are distilled into Dispersion Nomograf plume edge contour maps for rapid interactive access. Time integration is thus replaced by interpolations that capture the aerodynamic effects of the full urban geometry through the nomograf tables.
4. The nomograf tables then are encrypted and input to CT-Analyst, an easy-to-use graphical user interface (GUI) for instantaneous situational analysis. Plume computation, for example, takes less than 50 milliseconds.

The two panels in Figure 9 below were taken from a CT-Analyst screen shot for a region in downtown Washington DC using two different sets of nomograf tables. Panel (a) shows the three different nomograf displays computed 6 minutes after a release has occurred at the source indicated by the light blue star using the nomographs that encode the actual detailed Washington DC geometry resolved to 1-meter accuracy. Panel (b) shows the corresponding predictions when the detailed geometry is ignored. The source of contamination may be an accident, a leaking sprayer, or a broken container of hazardous chemicals that needs to be analyzed. The closed pink contour in each panel shows the 6-minute plume envelope for a wind of 3 meters per second from the west-northwest. The plume envelopes show the geographical region that the contaminant plume could have reached during its expansion up to the indicated time after release. The corresponding contamination footprints (grey) behind the plume envelopes in

Figures 9a and 9b respectively, represent the full extent of the growing contamination region after the plume envelope has spread to its maximum toxic extent. The actual contaminant plume envelope starts at the source near the upwind corner of the footprint and expands downwind away from the source as time increases.

The plume envelope and contamination footprint in Panel (b) are computed in exactly the same way as in Panel (a) except that the geometry used for the nomographs is flat and featureless. As a result the plume envelope is symmetric, exactly wind aligned, and shows none of the local dispersion caused by the actual building complexes whose effects are illustrated in Panel (a). Such a simple, non-physical result would be expected from Gaussian puff models using average roughness coefficients and even from Gaussian plume models. The results, however, are wrong and even a quick look at such results displayed generally makes most people suspicious.

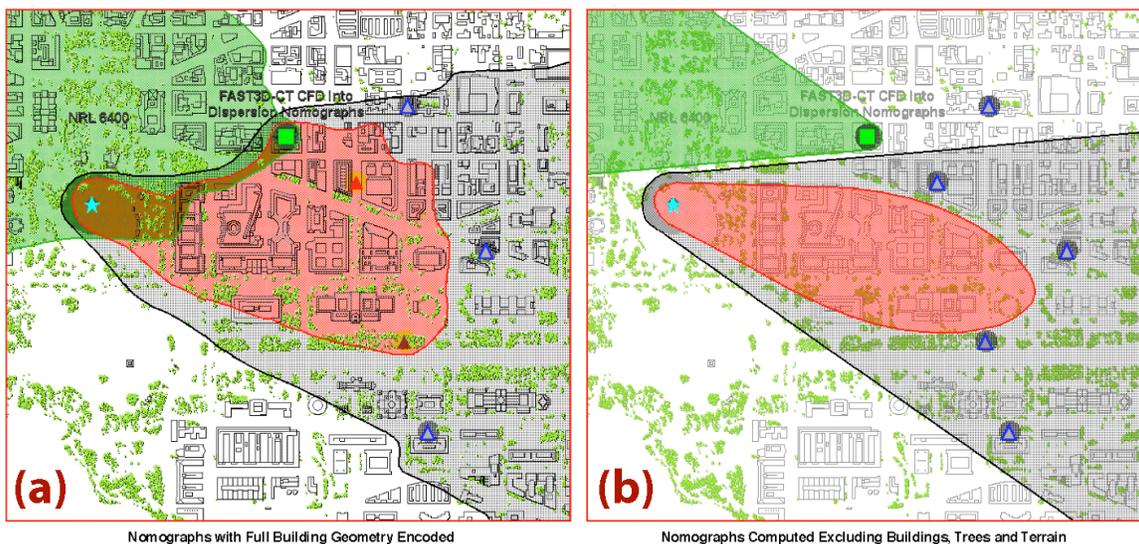


Figure 9. Dispersion nomographs in CT-Analyst capture building aerodynamics. a) Nomograf results with full building geometry. b) Nomograf computed excluding buildings, trees, and terrain. The figure shows the corresponding contamination footprints (gray), plume envelopes after 6 minutes (pink), and the upwind danger zone of a site (green).

The green shaded regions in Figures 9a and 9b estimate the upwind danger zone for the chosen site (bright green square) with and without capturing the effects of the geometry, respectively. The danger zone for a site is the set of all possible positions upwind where a source of contamination could reach the site. Computing the danger zone is an entirely new capability made possible by the nomograf representation and algorithms. Notice the difference caused by the local variations of the city geometry. In (b), without proper representation of the geometry, the source (blue star) is outside the predicted danger zone (green) and the plume envelope (pink) therefore does not reach the site. In (a) however, the extra northward dispersion from the buildings causes the plume envelope to reach the green square (site). In turn, CT-Analyst shows, via the danger zone, that the given site can be threatened from the location of the star. Clearly the predictions using the correct geometry are “safer.”

Sensors are indicated in these figures with triangular icons. Blue sensor icons have not been triggered by contaminant at the time depicted in Figures 9a and 9b as they lie outside of the instantaneous plume envelope. Some sensors and sites will lie outside the footprint with flat-earth geometry but inside the footprint when full geometry is used. In other words, a prediction

that does not account for complex geometry gives a false negative result for many locations, possibly resulting in unnecessarily high losses. An emergency response tool that does not account for complex geometry will be unreliable at best. Some conventional plume software greatly overestimates the size of a plume envelope, particularly near the source, to avoid these geometry-induced false negatives. Severe overestimates can also have dangerous consequences. People may be inclined to stay in a dangerous location, thinking an uncontaminated region is too far away to reach.

VI. Summary

All methods to compute the transport and dispersion of airborne contaminants need to use a representation of the wind field over the region of interest. The more accurate this representation of the wind field, the more accurate the plume prediction can be. Providing an accurate, building-scale approximation of the wind field is a slow, costly computational process, far too slow for emergency applications with the responsible officials (users) wait. Therefore conventional CT models based on conducting forward time-stepping predictions generally use unresolved constant parameter winds (for Gaussian plume models), very poorly-resolved wind field estimates that they can get quickly from other sources (for Gaussian puff models), or possibly very local wind/building interaction models (for Lagrangian particle models). These models still do a lot of computing while the users wait, eventually arriving at hazard-area predictions after several minutes that have obvious inaccuracies and other limitations.

In these earlier rather inaccurate, city-scale predictions, the specification of atmospheric stability classes and many other global properties about the agents and source dispersion mechanisms fosters a kind of intellectual dishonesty that misleads users, giving them a false sense of reliability for the models. For example, specifying the temperature and pressure of a toxic agent like chlorine and then asking a user to state the size hole that the chlorine is coming out of gives the impression that the heavier-than-air aspects of the release are being properly accounted for. Actually the type of surface, such as sand, water, or concrete, its heat transfer properties, and the slope of the terrain are even more important and yet totally neglected.

By way of contrast, this paper shows that the hazard area of a CT scenario can be computed quickly, accurately, and directly using detailed, pre-computed wind fields that properly capture individual building effects. This approach leads to an even more efficient representation in terms of Dispersion Nomographs where the plume envelope edges are recalled from compact, contour-map databases. In our implementation a state-of-the-art CFD model, FAST3D-CT, pre-computes the complete time-dependent 3D urban aerodynamics with full urban geometry at high resolution (2-6 meters). In this way there is no operator delay in using the very best wind velocity fields available. Dispersion Nomographs™ recall the results of these high-resolution 3D simulations with no time delay for transport and dispersion (T&D) calculations of many puffs/particles over many time steps.

The instant response CT-Analyst® format using nomographs also allows quick sensor fusion operations and “backtracks” to unknown source locations with actionable information in a few milliseconds. Dispersion Nomographs are faster because they provide plume edges directly rather than computing them laboriously while you wait. This approach is also more accurate because there is no application-time cost of using the very best wind velocity fields that can be had. The approach is also more robust because there is essentially no computation to go bad, very little

input is needed, and the displays are intuitive. This approach also leads to software that is easier to use because it is simpler.

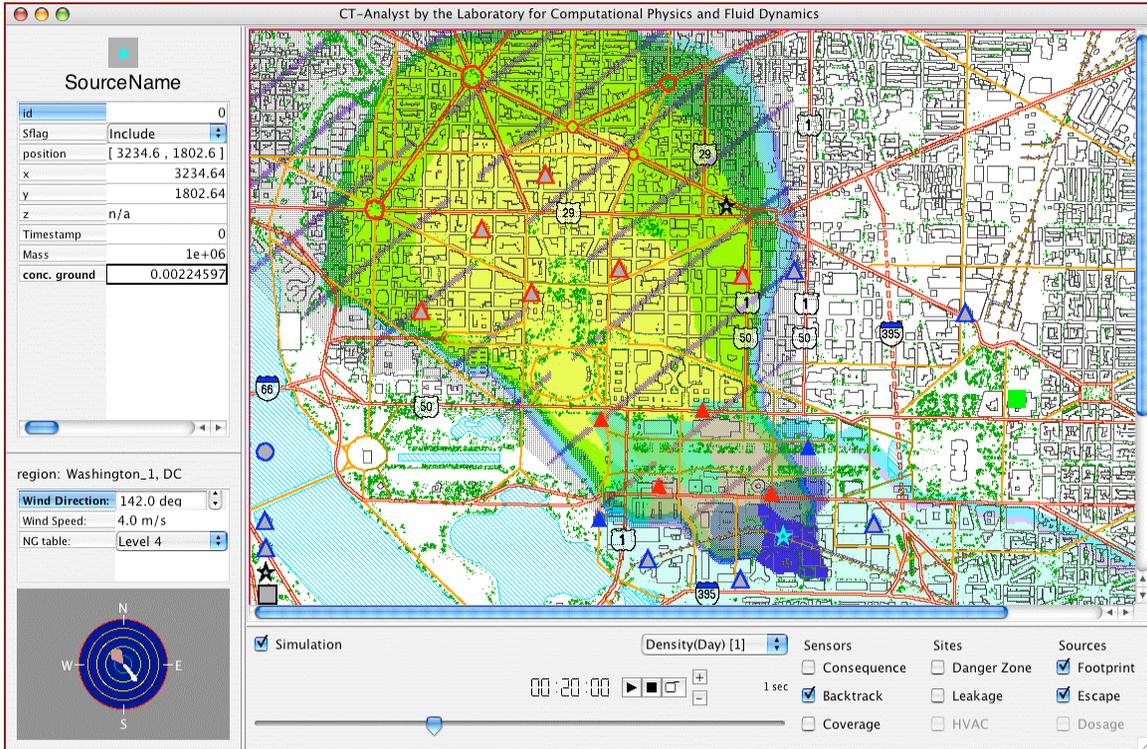


Figure 10. CT-Analyst full-screen display from a moderate area urban nomograf showing contaminant density contours (yellow, green, and blue), the contamination footprint (grey), and evacuation routes (purple lines) overlaid on a city map. The wind is from the southeast. Evacuation routes are optimized to minimize inhaled contaminant doses. The dark blue region in the lower middle of the display is the result of the CT-Analyst backtrace operation. Several “hot” sensors (red) and two “cold” sensors (blue) backtrack to a source location estimate using a union of the upwind “danger zone” capability.

Figure 10 above shows a typical CT-Analyst display for an urban area, in this case a section of downtown Washington DC prepared to support the 2005 and 2009 inaugurations. The contaminant concentration display (yellow-green-blue contours) for the source at the blue star is overlaid on the contamination footprint (gray region). Star-shaped nodes are sources, triangular and circular nodes are sensor reports, and square nodes indicate specific sites. When a source node is active it is colored light blue, as shown above. Footprints, plume envelopes, contaminant concentration plots, and escape routes can be displayed for sources by activating buttons on the lower portion of the CT-Analyst screen. Triangular sensor report nodes inside an active plume envelope are “hot” (red) while those still uncontaminated are “cold” (blue). In the example above all the site displays are inactive as shown by the buttons. Downwind consequence regions (for active “hot” reports) and upwind backtrace estimates (for all active “hot” and “cold” reports) can be displayed for individual sensor nodes, indicated by filled triangles. Contamination zones from down wind leakage and upwind danger zones can be plotted for all square site nodes (bright green when they are active). The diagonal purple lines are the recommended evacuation (escape) routes.

To compute displays such as danger zones, plume envelopes, and backtracks to unknown source locations, knowing the actual concentration of the airborne agent is not necessary. Indeed, until the total amount of the contaminant is known, plotting the actual concentration distribution isn't even possible. Therefore, CT-Analyst provides a relative concentration until the mass of the agent from a specific source can be estimated. Fortunately, this relative concentration and its time history are all that is needed to select civil defense options that minimize the inhaled dose of contaminant.

Acknowledgments

The authors wish to thank Lisa Williams, Mi Young Obenschein, David Fyfe, and Theodore Young, Jr., the NRL contaminant transport team who made this technology practical and easy to use. Further thanks are due Bob Doyle, Julie Pullen, Bo Cybyk, Jack Fulton, Sandy Landsberg, Charles Lind, and Rob Scott for their technical discussions and contributions to various components of this effort. Aspects of the work discussed here were supported by ONR through NRL, Hamburg Germany through the University of Hamburg, and in the past by the DoD High Performance Computing Modernization Office, DHS, NSWC Crane, DARPA, DTRA, MDA, and TSWG.

