An Operational Event Reconstruction Tool:
Source Inversion for Biological Agent Detectors

Michael J. Brown, Michael D. Williams, Gerald E. Streit, Matt Nelson, and Steve Linger
Los Alamos National Laboratory, Los Alamos, New Mexico

1. INTRODUCTION

A terrorist attack in a U.S. city utilizing biological weapons could have severe consequences. A biological agent could be aerosolized and emitted into the air in the middle of a city, invisibly traveling with the winds, and dosing an unknowing populace. The magnitude of the problem would only be revealed as sick people started arriving several days later at hospitals with symptoms, many already too ill to be saved. A national program has deployed a network of biological agent sensors in US cities to provide early detection of a bio-weapon attack, thereby hastening medical intervention and potentially saving many thousands of lives.

If a biological attack were to occur in a city, one or more detectors may register hits with specific dosages and the city would be alerted that an attack had taken place. This information alone, however, would not be enough to determine how serious the attack was, i.e., how much biological agent was released into the air and where the bio plume traveled. The first responder and public health communities will want to know what regions were impacted, how many persons might get sick, which persons most need medical supplies, and where to clean up. The law enforcement community will want to look for evidence at the release location.

The Bio-agent Event Reconstruction Tool (BERT) has been developed in order to recreate what might have happened during an airborne biological agent attack based on biological agent detector measurements and wind sensors mounted around a city. The tool can be used to estimate possible release areas while eliminating other areas, and can estimate bounds on the amount of material released. The tool can then be used to project forward from the possible source areas to estimate potential hazard zones. Due to a unique source inversion technique, the tool runs fast. Source regions can be determined in a few minutes.

In this paper, we provide an over of the BERT, followed by a description of the source inversion technique. We then show some example calculations and conclude with a real environmental background case.

Corresponding author address: Michael Brown, LANL, MS K551, Los Alamos, NM 87545, mbrown@lanl.gov

2. BERT OVERVIEW

BERT consists of a diagnostic wind solver, a source inversion model, and a graphical user interface. The tool is built inside of the ESRI ArcGIS mapping environment, so that the end-user can rapidly display detector measurements, winds from meteorological stations, and results on top of topography, street maps, orthophotos and other geospatial maps and data. A button inside of ArcGIS starts the BERT program and a tabbed window interface is used to input data and run the models.

As part of the BERT system, a server downloads hourly wind measurements for many cities from AIRNOW, MESOWEST, and city-specific servers. The wind measurements from the met stations are read into BERT and the diagnostic wind solver produces a mass consistent interpolated field. This wind field along with the detector dosages are then used by the source inversion code to determine potential release areas and to estimate bounds on the mass of biological agent released into the air.

The source inversion code is based on the concept of using detector footprints – a region upwind of each detector in which a release of a certain size would be detected by the detector – in order to quickly determine regions where the release might have occurred. Due to the uncertainties in both the wind and biological agent detector measurements and the underlying models, the BERT has been designed to provide answers within ranges associated with these uncertainties. The ideas behind the source inversion code will be described in Section 3. Limitations of the approach follow in Section 4.

Once potential source regions are identified, the BERT can be run in "forward" mode showing where the plume might have traveled. Since there is uncertainty in where the release took place, an automated routine can be used to look at the envelope of all potential plumes from all the point source releases inside the potential source region. Using the LANL National Day-Night Indoor-Outdoor population database derived by McPherson and Brown (2004), the total population that might have been dosed can be computed along with their spatial distribution. In a real event, this information may prove useful to the public health and medical communities, who will need to plan and prepare for the ensuing event.

Testing of the BERT has been accomplished using synthetic "data" computed by a variety of plume dispersion models. These idealized tests show that for most cases the tool can determine areas which encompass the actual release location. Examples of a few of these tests will be shown in Section 5.

The BERT has been successfully used in several real events in which naturally-occurring biological materials found in the environment were detected. For one case, BERT was used to eliminate regions where biological material could not have come from and to highlight areas where it might have originated. A team searching for the source of the naturally-occurring material utilized this information and found positive samples within one of the predicted regions. More discussion on this case will be given in Section 6.

BERT also contains a "forward" mode in which a segmented Gaussian plume model is used to create a concentration field and dosages at the biological agent detectors. The forward mode is useful when performing table top exercises, training, and testing of the BERT source inversion code.

3. DESCRIPTION OF THE DETECTOR FOOTPRINT METHODOLOGY

Given a set of time-averaged dosages and winds at locations distributed over a metropolitan-scale landscape, one could run a plume dispersion model at tens of thousands of locations (x,y) with different release amounts Q and attempt to match the dosage measurements as best as possible in order to determine the source location and source strength. If the biological detectors require a long averaging time to produce reliable results (long relative to the time scale of the mesoscale wind patterns), then more plume dispersion simulations will need to be conducted using wind fields at different times of the day over the bio detector sampling duration. Given that time is of the essence when responding to a real event, a method to find the source characteristics that requires fewer computations is desirable.

We have developed a source inversion technique deemed the detector footprint methodology. The number of plume dispersion calculations that need to be solved in this approach is equal to the number of biological detectors. The concept of the detector footprint is straightforward. For example, the annulus shown in Fig. 1 represents the region in which a release of 5-10 grams of a biological agent would result in the measured dosage (within a factor of 2) at the detector. A 5-10 g release outside this footprint would result in a dosage at the detector that was less than that measured. A 5-10 gram release inside the detector footprint annulus would result in a dosage that was too large in comparison to what was measured.

The closed contour defining the inside of the 5-10 gram annulus denotes the locations (x,y) where a 5 gram release would result in the measured dosage D. Conversely, a 5 gram release at the detector location (x_d,y_d) with the wind direction reversed results in a dosage D anywhere along the 5 gram contour. This can be easily illustrated through use of the straight-line Gaussian plume model:

$$C(x, y, 0) = \frac{q}{\pi U \sigma_y \sigma_z} \exp\left[\frac{-(y - y_d)^2}{2\sigma_y^2} \right], \quad (1)$$

where for simplicity the detector and release height have been set to ground level, x is in the direction of the wind, y is the lateral distance from the plume centerline, U is the wind speed, and for a given stability, the σ_y and σ_z plume spread parameters are only a function of the distance between source location and the detector location. Eqn. (1) indicates that for the same source strength q the concentration $C(x,y)$ for a release at position (x_d, y_d) is the same as the concentration $C(x_d,y_d)$ for a release at position (x,y) if the wind direction is reversed 180 degrees.

Based on this result, it is then straightforward to compute the upwind footprint by reversing the wind direction, rearranging the time-integrated form of eqn. (1) and solving for Q given the measured dosage D. Upwind contours of Q are created and these become the detector footprints. One could also just run the plume dispersion model with the reversed wind specifying a release of mass Q at the detector location, solve for the measured dosage, and then using the symmetry results from above, rename the dosage contour as the Q contour and create footprints this way.

In the BERT, the underlying transport and dispersion code is a segmented Gaussian plume model. The model accounts for spatially varying winds, that is, the plume centerline trajectory is calculated in discrete line segments based on the wind field and the straight-line Gaussian plume model equation is used for each line segment along with a virtual source to account for the spread of the plume at prior time steps. A beta version of the model has been developed to account for time-varying winds as well. The segmented Gaussian plume model is similar to a puff model, but not as robust. The segmented Gaussian plume model was chosen because it was amenable to being used to create detector footprints that contained information on model wind direction uncertainty.

If only one sensor detects a biological agent, then the upwind detector footprints indicate the zones where a potential release could have occurred of a given strength Q. The other detectors in the network that were not hit contain information that can be used to reduce the areas in which a release could have occurred. That is, detectors that did not register a

positive reading provide "null" footprints, or areas, from which an atmospheric release could not have occurred. The null footprint is created by putting in the threshold, or minimum, dosage that the sensor can detect and solving for Q. For example, a release of 5 grams or more within a 5 g null footprint would result in a hit at the detector and hence this area should be eliminated as a potential release region. As shown in the example in Fig. 2, the overall region in which a release could have occurred can be reduced by the null footprints overlapping the hit detector footprints.

If more than one detector registers a hit, then the potential release area can be reduced even further. As shown in Fig. 3, for a single terrorist to have released between 5 and 10 grams at a point in space and hit both detectors, the release would had to have occurred in the overlap region. Note that the 2.5 to 5 gram annuli do not overlap and that implies that a release of this quantity could not have hit both collectors at the reported dosages. This is how the tool finds bounds on the release amount.

4. LIMITATIONS OF THE APPROACH

The BERT and the collector footprint approach have lots of deficiencies. First, it can't always find an answer and if it does, it doesn't find just one answer, but rather it produces a range of possibilities. Another big limitation is that it doesn't work for line sources or multiple point sources. However, the collector footprint approach has been adapted to work with area sources by looking at the union of hit footprints rather than the intersection. Furthermore, the approach is only as good as the measurements, and biological detectors can have considerable uncertainty in the measured dosage. To account for some of this uncertainty we only specify our answers for release amount to within a factor of two. The wind measurements can also contain errors, but more importantly the wind model is diagnostic and does not account for complex flow effects, unless there are wind measurements at high density to account for them. In order to account for the probable wind direction error in the wind field, we have developed a simple scheme which "fattens" the detector annuli as a function of uncertainty in the modeled wind direction. Finally, the segmented Gaussian plume model that is used to create the detector footprints has many limitations, including not performing well for light and variable winds, for cases where the wind direction changes rapidly over short distances, and when the wind direction changes rapidly with height.

5. IDEALIZED TEST CASES

The BERT has been tested by running the code using "synthetic" detector data created by transport and dispersion models. As expected, when we create the synthetic data by running the BERT segmented Gaussian plume model using wind fields computed by the BERT diagnostic wind model, the BERT source inversion model produces potential source regions that enclose the actual source region.

We have performed other tests with different types of dispersion models, including the Urban Dispersion Model (Hall et al., 2000), the Quick Urban & Industrial Complex dispersion modeling system (Pardyjak and Brown, 2003), and the ADAPT/LODI wind and dispersion modeling system (Nasstrom et al., 2000). The first model is a puff dispersion model, while the latter two have Lagrangian random-walk codes. In all cases, the actual source location was within the computed potential source regions or just slightly outside. In a few cases, answers could only be obtained when the wind direction uncertainty option was activated, indicating that the wind fields produced by the different models were different enough to significantly affect plume transport. Figure 4 shows the results from a synthetic case produced by the ADAPT/LODI dispersion modeling system for a fictitious layout of bio detectors, but using real winds. For this case, we were able to find solutions that enveloped the actual release location only if a 10 degree wind direction uncertainty was added to the footprints. Note that for this particular case, the winds changed substantially over the bio detector sampling duration, so that other solutions in different regions were also found.

6. REAL-WORLD EXAMPLE

Biological agents occur naturally in the environment and are often endemic to different wild or domesticated animal populations. In one case that occurred in 2003, three detectors registered hits of a specific biological agent that was later discovered in soil and water samples by Barns et al. (2005). Due to the knowledge that this particular biological agent was endemic in the area, the BERT was utilized at that time to make maps of potential areas to take environmental samples. The graphic in Fig. 5 shows one of the computed potential source regions along with the locations where the environmental samples tested positive for the agent. In addition, large swaths of land were deemed unlikely natural repositories of the agent due to the null footprints. This information was utilized by the sampling team to focus their collection efforts. Although not a validation of the modeling system, it does illustrate how such a tool can be used to help direct limited resources.

7. CONCLUSION

Biological agent detectors are being placed in large cities throughout the US and at military bases both in the US and overseas. In the event of a detection, one of the first needs by the emergency response community for consequence management is

3

determining where the plume went and how much was released in order to evacuate regions, to devote medical resources, and/or to undertake clean up. But, to determine the location and aerial extent of the hazard zone, the size of the release and where it occurred first needs to be computed. Our team has created the Bio-agent Event Reconstruction Tool (BERT) for helping to reconstruct what happened (i.e., where the release may have occurred, how much was released, where the bio-agent cloud may have traveled) based on wind and bio detector measurements.

The BERT is embedded into a GIS mapping environment and utilizes a detector footprint source inversion methodology that is fast. The tool has been tested using synthetic detector data created by several different transport and dispersion models. The tool is being used for reachback applications as part of a large national program and has demonstrated utility in several cases. Work is ongoing to improve the BERT, including utilizing prognostic mesoscale model wind fields, accounting for line and area sources in the source inversion code, and improving the handling of uncertainty in the measurements and the models.

8. REFERENCES

Barns, S., C. Grow, R. Okinaka, P. Keim, C. Kuske, 2005: Detection of diverse new Francisella-like bacteria in environmental samples, J. Appl. & Env. Microbiology, **71**, 9, 5494-5500.

Hall, D., A. Spanton, I. Griffiths, M. Hargrave, S. Walker, 2000: The UDM: A model for estimating dispersion in urban areas, Tech. Report No. 03/00 (DERA-PTN-DOWN).

McPherson, T. and M. Brown, 2003: U.S. day and night population database (Revision 2.0) – Description of methodology, LA-CP-03-0722 & LA-UR-03-8389, 30 pp.

Nasstrom, J.S., G. Sugiyama, J.M. Leone, Jr., and D.L. Ermak, 2000: A real-time atmospheric dispersion modeling system, AMS 11[th] Joint Conference on the Applications of Air Pollution Meteorology, Long Beach, CA, 84-89.

Pardyjak, E.R. and M. Brown, 2003: QUIC-URB: Theory and User's Guide, LA-UR-07-3181, 22 pp.

wind
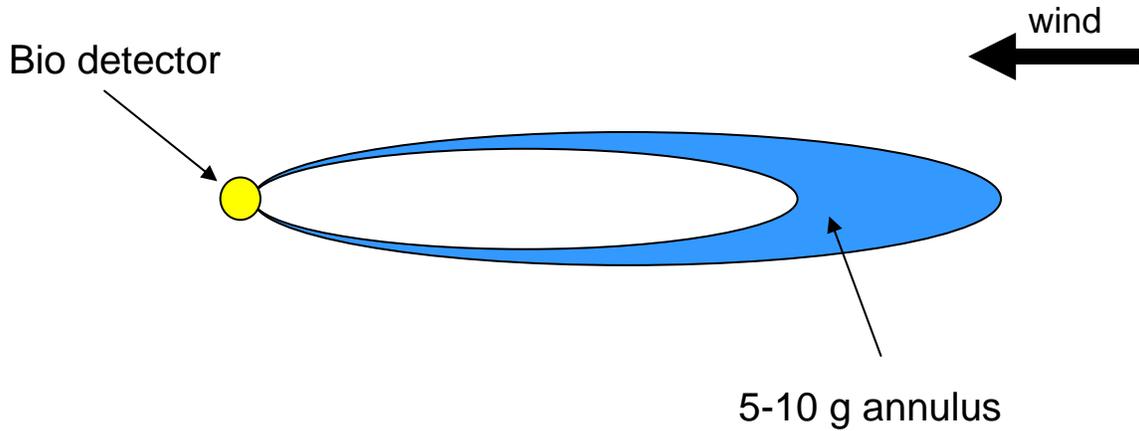
Bio detector

5-10 g annulus

Figure 1. Plan view showing the biological detector and the associated upwind detector footprint annulus. For this example, the annulus defines the region where a 5 to 10 gram release of a specific biological agent would result in the measured dosage at the detector (to within a factor of two). A 5-10 g release inside the annulus will result in a dosage that is higher than that measured, while a release outside the footprint will result in dosage that is smaller than measured. The annulus will change size depending on the magnitude of the dosage (a higher dosage results in a shorter footprint), atmospheric stability (stable stratification results in longer footprints), the wind speed (faster wind speeds result in shorter footprints), and the agent deposition velocity (a higher deposition velocity results in a shorter footprint).
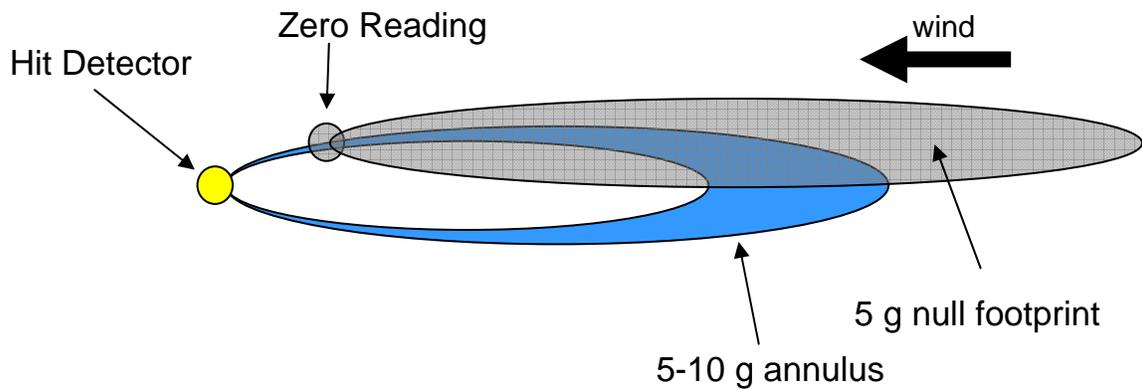

Zero Reading

wind

Hit Detector

5 g null footprint

5-10 g annulus

Figure 2. Plan view showing a hit detector with it's upwind hit footprint and a non-hit detector showing it's null footprint. The null footprint shows the region in which a 5 gram release or more could not have occurred or else the bio detector would have recorded a hit. Since the null footprint overlaps part of the 5-10 gram annulus, it means that this part of the annulus is not a potential source region. Only the non-overlapped area in blue is now a potential source region for a 5-10 gram release.
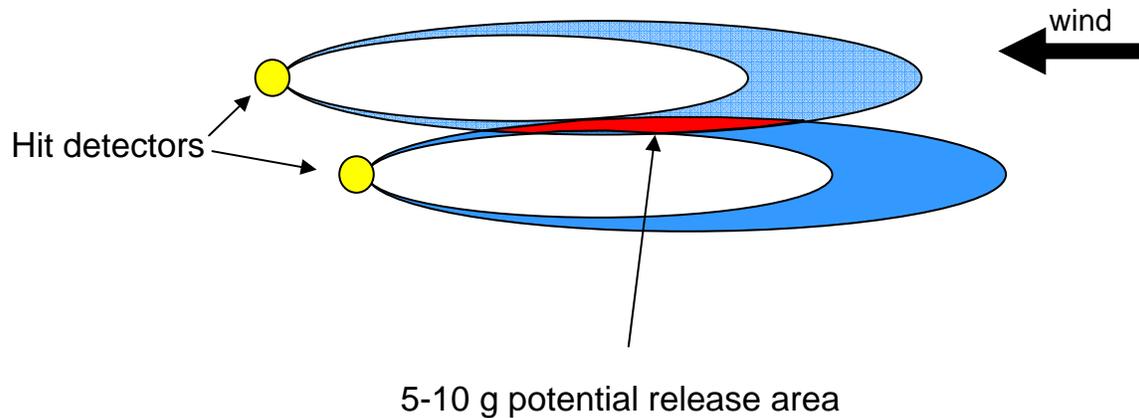
Figure 3. Plan view showing two hit detectors and their associated detector footprint annuli. The overlap region (in red) is the region in which a point source release of 5-10 g could have hit both detectors at the measured dosage. Note that the 2.5-5 g annuli fit within the 5-10 g annuli and they would not overlap. Hence, a lower limit on the release amount can be estimated. A 5 gram or less release from a single point source could not have hit both detectors at the measured dosage.
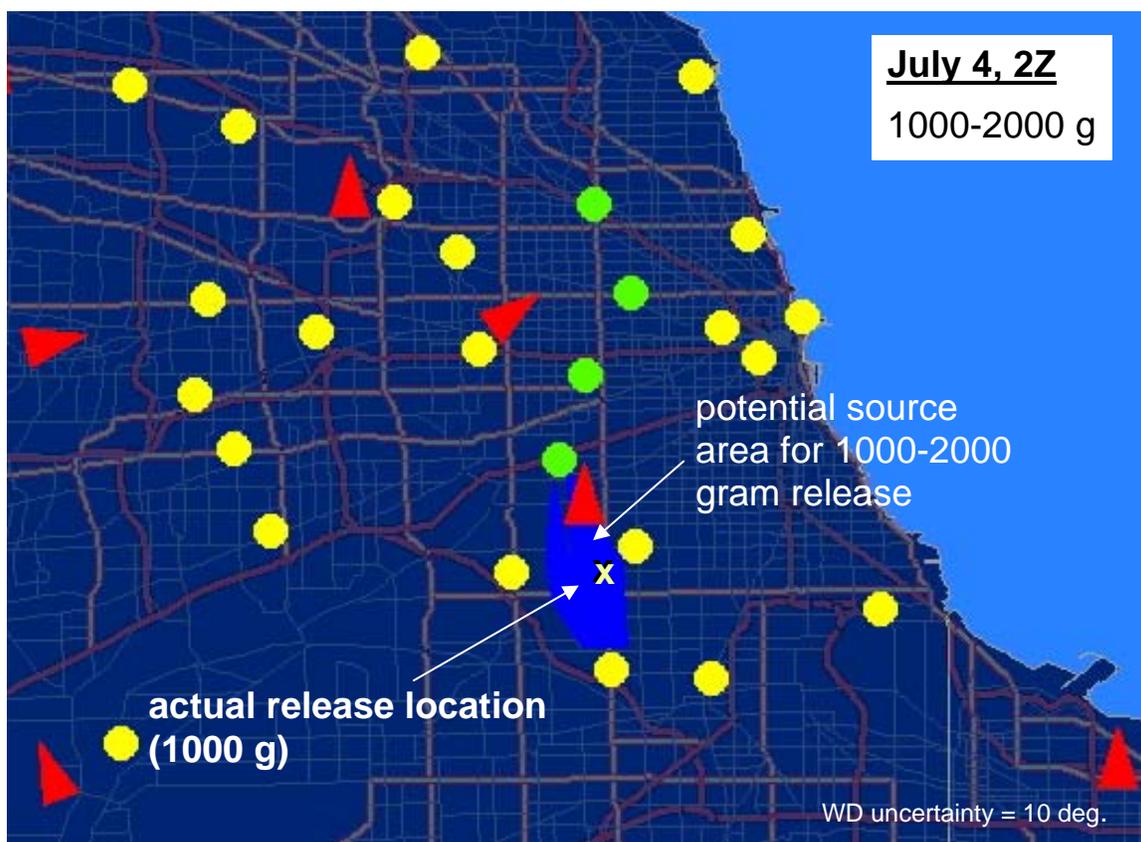

Figure 4. The potential source region (blue) computed by the BERT footprint model using synthetic detector dosages created by Lawrence Livermore National Laboratory's ADAPT/LODI transport and dispersion modeling suite. The green circles show the hit detectors and the yellow circles represent those that did not detect the biological agent. The red triangles denote the wind directions measured at stations throughout the city. The computed 1000-2000 gram potential source region overlaps the actual location of the modeled release.
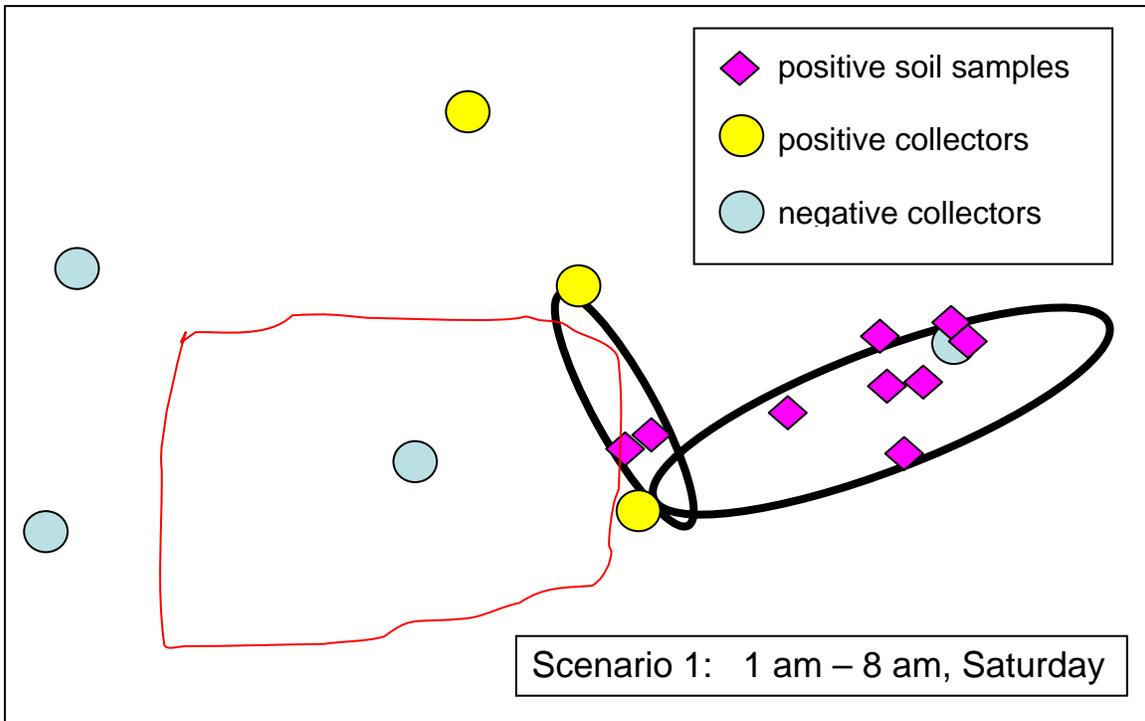
Figure 5. Example from a real event in which the detectors were triggered by naturally-occurring biological agent. The black ellipses represent the potential source regions. Since the winds shifted directions over the 7 hour period, two regions were found. Since we were looking for an area source, instead of overlapping the regions to find a point source location from the intersection of the two polygons, one would take the union of the two polygons. The red curve depicts the approximate outline of the metropolitan region. Note that the underlying map was removed because the locations of the detectors are not to be disclosed.