Available online at www.sciencedirect.com
ScienceDirect
Journal of Hazardous Materials 158 (2008) 417–429
Journal of Hazardous Materials
www.elsevier.com/locate/jhazmat

A risk-based sensor placement methodology☆

Ronald W. Lee*, James J. Kulesz

Computational Sciences and Engineering Division, Oak Ridge National Laboratory, United States

Received 15 October 2007; received in revised form 26 January 2008; accepted 28 January 2008
Available online 12 February 2008

Abstract

A risk-based sensor placement methodology is proposed to solve the problem of optimal location of sensors to protect population against the exposure to, and effects of, known and/or postulated chemical, biological, and/or radiological threats. Risk is calculated as a quantitative value representing population at risk from exposure at standard exposure levels.

Historical meteorological data are used to characterize weather conditions as the frequency of wind speed and direction pairs. The meteorological data drive atmospheric transport and dispersion modeling of the threats, the results of which are used to calculate risk values. Sensor locations are determined via an iterative dynamic programming algorithm whereby threats detected by sensors placed in prior iterations are removed from consideration in subsequent iterations.

In addition to the risk-based placement algorithm, the proposed methodology provides a quantification of the marginal utility of each additional sensor. This is the fraction of the total risk accounted for by placement of the sensor. Thus, the criteria for halting the iterative process can be the number of sensors available, a threshold marginal utility value, and/or a minimum cumulative utility achieved with all sensors.

Keywords: Sensor placement; Risk

1. Introduction

The placement of sensors at storage facilities and other locations with known threats is a critical aspect of any strategy to protect population potentially exposed to those threats. However, the prohibitive costs associated with purchasing, deploying, and maintaining a large number of sensors in a network make it impractical to place sensors everywhere. We propose a sensor placement methodology that optimizes sensor locations to protect population against the exposure to, and effects of, known and/or postulated chemical, biological, and/or radiological threats. Our motivation is two-fold. First, sensor placement techniques for detection of chemical and/or biological agents have yet to be standardized or universally adopted. Second, in order to determine population effects, a sensor placement strategy must account for the transport and dispersion of hazardous materials, the population distribution, and the toxicity of materials. Existing sensor placement algorithms and methods do not account for all of these factors.

Current optimization approaches to sensor placement include heuristics, genetic algorithms, dynamic programming, and diffusion boundary methods [1–3]. Placement optimization objective functions vary for different types of sensors. For example, visual or geometric coverage is typically the objective function to be maximized for placement of surveillance cameras [4–7]. For point or area detections the objective is to maximize coverage of a geographic or geometric area [8,9]. Other approaches to plume detection focus on criteria such as time to detection or total sensor area coverage and are concerned only with detecting any part of a plume [1,3].

However, mere detection of a chemical or biological agent is insufficient to fully characterize the threat posed by the corresponding release. We suggest that *effects on population* of undetected releases of hazardous materials represent the *risk* posed by the hazards, and this risk should be the primary consideration when placing sensors.

In order to account for population effects, a sensor placement method must consider the variety of meteorological conditions under which a threat might occur in a manner that is independent

☆ This work is funded under U.S. Department of Energy proposal number ERD-04-2465 for the Memphis and Shelby County Port Commission at Oak Ridge National Laboratory which is managed by UT-Battelle, LLC for the U.S. Department of Energy under Contract No. DE-AC0500OR22725.
* Corresponding author. Tel.: +1 865 241 5096; fax: +1 865 576 0003.
E-mail address: leerw@ornl.gov (R.W. Lee).

0304-3894/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jhazmat.2008.01.111

of the time of occurrence. Other methods assume a particular wind condition [1] and/or solve the sensor placement problem for a particular instant of time. This is useful if sensors can be relocated quickly with each meteorological change. It is more common that sensors cannot be moved easily, and thus a solution accounting for all meteorological conditions is needed.

The proposed algorithm and methodology places sensors for situations in which threats, defined as releases of hazardous material, are known or can be postulated. Sensors are placed to detect the releases whenever they might occur, thereby protecting potentially affected population. The key elements of the proposed approach are:

- risk, defined as population exposure and effects, is the basis of the objective function to be minimized;
- a *wind rose* specifying probabilities for wind (direction, speed) pairs is derived from historical meteorological data and used to represent wind conditions at a site;
- the transport and dispersion, as well as population exposure effects, are computed for each threat at each wind (direction, speed) pair; and
- sensor locations are determined via an iterative dynamic programming algorithm, where each iteration optimizes the location of the next sensor.

The proposed methodology offers: a defensible and reasonable foundation for a systematic, risk-based placement of sensors for known threats; and a quantifiable mechanism for determining the marginal utility of each additional sensor. Moreover, the proposed methodology is generally applicable and is not dependent on locations, types of threats, sources of data, or specific models.

Sensors are assumed to be *perfect detectors*, so that any amount of an agent or material passing through a sensor will be detected. Thus, the methodology does not account for a sensor's response to the concentration or mass of the material required for detection. There are two justifications for this assumption. First, the nature of the problem to be solved, namely toxic agents in sufficient quantities to pose a risk to population, implies material masses above any reasonable detection threshold for a sensor. The sensors used in the reference application of the methodology respond to quantities in the nanogram range, respond within a few seconds, and are quickly saturated. Second, the problem to be solved is the location for sensors in general, irrespective of the particular sensing technology deployed.

The remainder of the paper is organized as follows. We start with a description of the background and concepts upon which the methodology is built in Section 2. Next, Section 3 describes the inputs, parameters, and calculations required to start the placement process. Section 4 presents the risk-based optimization function, and the placement algorithm is explained in Section 5. Steps in the proposed methodology are summarized in Section 6, and results from applying the methodology at the Port of Memphis are presented in Section 7. Comparison with other approaches and concluding comments are given in Section 8.

2. Background and concepts

Managers at the Port of Memphis sought to place a cost-limited number of sensors in locations that would provide the greatest protection from airborne releases to the population in the area. Warning of a release would allow them to take appropriate measures for potentially affected people. Hazardous chemicals stored at three separate locations were the threats to be addressed. Based on these minimal criteria, a methodology was developed to select sensor locations in deployment order, such that the airborne releases that are intercepted represent the highest reduction in exposure risk to the population. Air dispersion analyses were based on site-specific release quantities and locations and 30 years of hourly wind direction and speed observations at Memphis International Airport.

The first step was to identify the hazardous materials, quantities, and locations. A threat weighting factor was assigned to each release to represent its likelihood. (A value of one for all releases means they are equally likely.) A wind rose was calculated using historical data. Then, the transport and dispersion of each release was computed for each wind (direction, speed) pair in the rose. The dosage field resulting from each transport computation was contoured by dosage values corresponding to standard (qualitative) exposure levels, and the number of people contained within each exposure level was calculated. Each exposure level was assigned a factor representing the severity of the exposure. For each exposure level contour, a risk value was calculated as the product of the threat factor, the exposure level factor, the probability of occurrence of the associated wind (direction, speed) pair, and the count of the population for the contour.

The final step before choosing the location of the first sensor was to computationally overlay all contours over a suitable two-dimensional grid to represent the plane at the location of interest. The geographic center of each grid cell was assigned as its reference point. Each cell was checked against each exposure level contour to determine if the cell reference point lies within the contour geographically. If so, the contour's risk value was added to the risk value associated with the cell. The first sensor was assigned to the cell with the highest risk value. If restrictions at the site prevented placing a sensor in the chosen grid cell, the cell with the next highest risk value was selected until a cell was found where a sensor could be placed.

The location of the second sensor was determined in a similar fashion, except that all contours that were used in calculating the risk value for the cell where the first sensor was assigned were removed from the calculations. Subsequent sensors were placed by following the same process described above, excluding contours used in previous sensor placement cells. The process was terminated when all available sensors were placed.

3. Inputs and parameters

The proposed sensor placement methodology is built around a dynamic programming algorithm where each iteration is an optimization solution for the next sensor location. Data gathering and computations preparing the inputs for the placement

algorithm are essential to understanding and implementing the methodology. The inputs and parameters of the risk-based sensor placement methodology are threat characteristics, exposure levels, a site-specific wind rose, population data, transport and dispersion computations, and a suitable sensor placement grid.

3.1. Threats

We begin with a set of known or postulated threats to detect. Each identified threat is modeled as an agent release for input to an atmospheric transport and dispersion model. The nature of the threat dictates the release definition. For example, a puncture in a storage tank might be modeled as a series of instantaneous releases at decreasing masses. Complete failure of a container holding liquid under pressure might be represented as an instantaneous gaseous release and a pooled liquid release. For the Port of Memphis, threats were derived from an Environment Protection Agency (EPA) study and modeled as continuous releases of specific materials at specific release rates using the Hazard Prediction and Assessment Capability (HPAC) [10].

Regardless of the choice of dispersion model, the threat must be defined as accurately as the chosen model allows and should account for the material mass as well. Realistic release rates must be established for releases from containers, where the dynamics of phase changes that occur as the material is expelled from a container are applied.

3.1.1. Threat factor

In many cases, each threat may be considered equally likely and/or of equal concern. If not, a weighting factor may be assigned to each threat to represent its likelihood of occurrence or importance relative to other threats. This is a means of quantifying threat reduction measures, such as administrative and/or engineering controls. The factor is applied in the risk value calculation described in Section 4.1.

3.2. Exposure levels

Dispersion models typically produce concentration, deposition, and/or dosage fields. However, these results do not relate directly to effects on population, and correlation of the fields to effects is necessary.

For this methodology, a model must produce a dosage field. The dosage field is contoured by material-specific dosage values associated with standard exposure limits and levels. These standards are defined by organizations such as the American Industrial Hygiene Association (AIHA), the National Research Council's Committee on Toxicology (NRC COT), and the Department of Energy (DOE). AIHA publishes the Emergency Response Planning Guidelines (ERPG). NRC COT publishes Acute Exposure Guideline Levels (AEGL), Emergency Exposure Guideline Levels (EEGLs) for the Department of Defense, and Short-Term Public Exposure Guidance Levels (SPEGL) [11]. Temporary Emergency Exposure Limits (TEELs) are defined by DOE and used when AEGLs and ERPGs are unavailable [11].

Dosages corresponding to the standard exposure levels vary by material or agent. HPAC material descriptions include exposure level dosage values used for contouring the dosage field.

3.2.1. Exposure level factor

These exposure levels are qualitative in nature and therefore must be quantified to contribute a term to an objective function. Thus, we associate a quantitative factor with each qualitative exposure level, where baseline levels such as TEEL-0 and ERPG-0 are given a factor of 1.0. Higher exposure levels, LCt90 (90% lethality) being the highest, are given higher factors applied to the risk value term. For example, in application of the methodology at the Port of Memphis, LCt90 was assigned a value of 5.0, meaning one person exposed at LCt90 is equivalent to five persons exposed at TEEL-0. The exposure level factor is applied in the risk value calculation described in Section 4.1.

3.3. Wind rose

To fully account for the range of meteorological conditions occurring over time at a location is impractical. However, wind conditions are a useful subset for the purposes of the proposed sensor placement methodology. A common representation of the range of wind conditions at a location is a *wind rose*. "A wind rose gives a very succinct but information-laden view of how wind speed and direction are typically distributed at a particular location" [12]. Specifically, it specifies wind direction and speed pairs and their percentage of occurrence. The Natural Resources Conservation Service (NRCS) uses data from the Solar and Meteorological Surface Observation Network (SAMSON) to produce wind roses. Refer to Fig. 1 for an example from the NRCS Web site. SAMSON consists of hourly observations from 1961 through 1990 at 237 National Weather Service (NWS) stations in the United States, Guam, and Puerto Rico. Thus, SAMSON is a good data source for computing wind roses at the stations covered, although any source of data can be used.

3.3.1. Direction and speed bins

Direction and speed bins in the rose must be chosen to be sufficiently small to account for variations in the observation values. The NRCS chooses 16 directions and six speeds, the 16 directions yielding a bin size of 22.5°. For the Port of Memphis, the six speed bins were sized at two m/s with centroid values one, three, five, seven, nine, and 11. Speeds above the high bin boundary are assigned to the highest bin.

3.3.2. Station locations

A single rose derived from 30 years of data provides a needed succinct representation of wind speed and direction in terms of percentage of occurrence. However, the NWS station(s) chosen to produce the rose must be close enough to (and representative of) locations of interest to be applicable. SAMSON data for Memphis International Airport were used for the Port of Memphis. Note the airport is approximately eight miles away, and there is no intervening terrain of consequence.

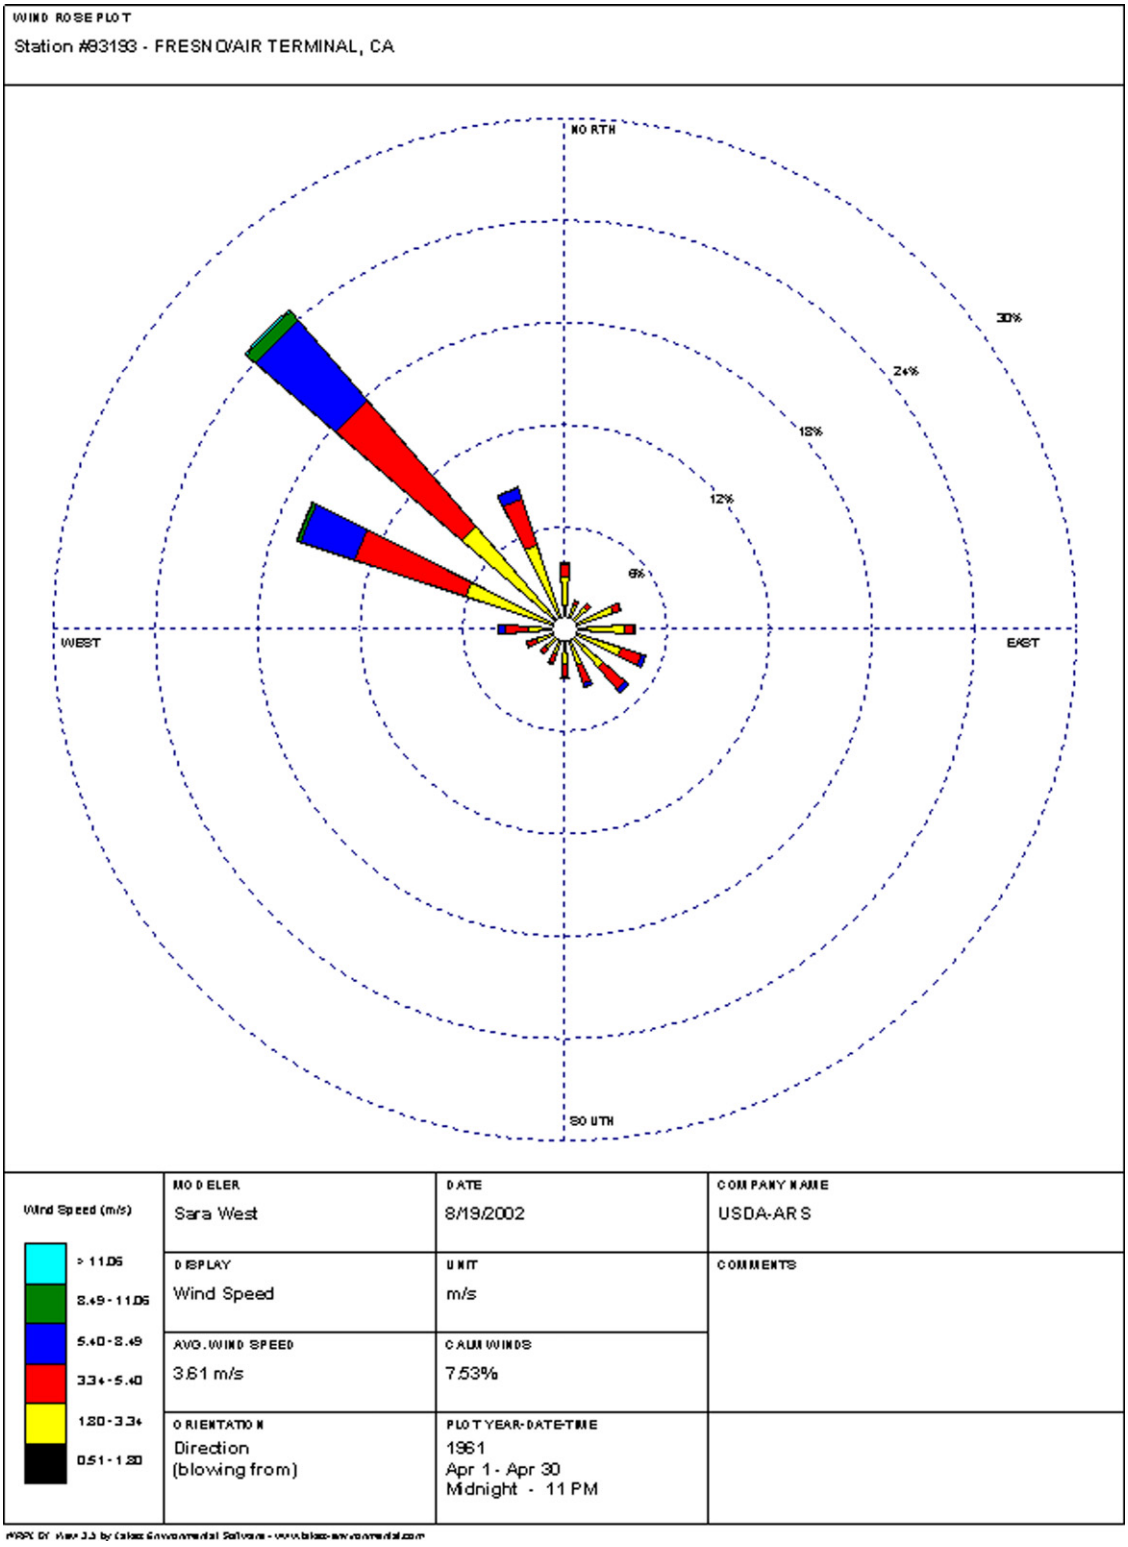

Fig. 1. Sample wind rose.

3.4. Population data

A geographic population distribution is necessary to count the population affected by a hazardous release. As described above, dispersion model dosage fields for each threat are contoured by standard exposure levels, and the population within the contours is counted. LandScan 2003 data at 30 arc-second resolution were used for sensor placement at the Port of Memphis [13].

It should be noted that a population distribution that varies by time of day, most simply day versus night, is particularly useful. The population count within a contour is applied in the risk value calculation described in Section 4.1.

Fig. 2. Example contours resulting from dispersion computation.

3.5. Transport and dispersion computation

As mentioned above, each identified threat is modeled as an agent or material release to be fed to a dispersion model. The dispersion model should account for material characteristics such as evaporation rates and buoyancy as well as weather and environmental conditions during the dispersion process.

Many dispersion models are available, such as the Hybrid Single-Particle Lagrangian Integrated Trajectory Model (HYSPLIT) [14], the Areal Locations of Hazardous Atmosphere (ALOHA) [15], Vapor Liquid Solid Tracking (VLSTRACK), and many others [16]. Any system able to competently model the identified threats can be used. Our implementation of the proposed methodology uses the Second-order Closure Integrated Puff Model (SCIPUFF) [17,18], the transport and dispersion engine in the Hazard Prediction and Assessment Capability (HPAC) [10]. In addition to the Defense Threat Reduction Agency's (DTRA) standardization of HPAC for hazard analysis, reasons for choosing HPAC include:

- built-in modules for contouring calculated dosage fields;
- an extensive material library with dosage values corresponding to the standard exposure levels identified above;
- built-in calculation of population using LandScan 2003 data [13]; and
- a validated transport and dispersion model accounting for meteorology, terrain, land cover, and other environmental conditions.

3.5.1. Computing threat effects

For each threat, the dispersion must be computed for each (direction, speed) bin in the wind rose. Thus, a rose with 16 directions and six speeds requires 96 dispersion computations for each threat. Each dispersion computation results in a dosage field, from which exposure level contours are derived as described in Section 3.2.

With HPAC, contour calculations are provided in the dispersion engine library, and the population contained within each contour is counted against the 30 arc-second night time LandScan 2003 distribution [19]. For each exposure level, one or more contour polygons result. Once the computations are complete there is a set of exposure level contours for each threat and for each wind rose bin.

The duration for the model computation is specified in the threat definition and must be long enough to account for the full effects of a release on population. HPAC defaults the duration to 4 h, but 2 h was sufficient for the threats at the Port of Memphis.

3.5.2. Exposure level contours

Exposure standards specify multiple levels representing a range of toxicity, and a dosage value for each level is associated with the material released. Thus, contouring the dosage field (resulting from dispersion computation for a particular threat and wind condition) results in a contour for each exposure level.

For example, sensor placement at the Port of Memphis was based on nine threats and a wind rose of 96 bins. Materials for the nine threats are such that six have three defined exposure levels and three have seven defined exposure levels. Thus, there are potentially 39 contours for each of the 96 meteorological conditions for a total of 3744, although dosages for higher exposure levels are not always reached. For the Port of Memphis placement, well over 3100 individual contours were generated.[1]

Each contour consists of one or more polygons, and each polygon is defined as an outer ring and zero or more interior or hole rings. Fig. 2 shows three representative contours.

[1] All dosage contour values are taken from HPAC material files which reference DOE document DKC-04-003.

Contours themselves provide primary input to the placement optimization's objective function described in Section 4.

3.6. Placement grid

The proposed methodology assumes point-detector sensors will be deployed on a two-dimensional grid covering the spatial domain of potential locations. The height at which sensors are deployed is chosen to: ensure their detection of any release of sufficient magnitude to threaten nearby population; and to account for other placement constraints, such as protection from tampering and accessibility for scheduled maintenance. Since the sensor height is fixed, the methodology selects placement in two dimensions.

Although a rectilinear grid is a natural fit for two dimensions, *cells* in the grid can be of any shape or size. However, all space within a cell is treated equally, meaning a sensor may be placed anywhere within the cell and achieve the same result. One location within each cell is chosen as the *reference point*, representative of all the space covered by that cell. Note the grid must be granular enough such that cell reference points will account for threat effects, as described in Section 4.2.

A uniform rectangular grid composed of roughly 27 m × 30 m cells was used at the Port of Memphis. This adequately represents the range of space in which discrete sensor locations are possible. Each cell's center is chosen as its reference point. Environmental constraints such as power and network availability might eliminate some grid cells as possible sensor locations, as was the case in Memphis.

4. Risk-based objective function

The key concept behind the sensor placement optimization is *risk*. Each threat's dispersion is computed for each (direction, speed) pair in the wind rose, and each dispersion computation results in a set of contours representing exposure levels appropriate for the agent/material released. For each contour we calculate a *risk value* in order to quantify the effects on population of the exposure represented by the contour. The cumulative risk across all contours for all threats and wind conditions is the total risk, and the objective is to detect or *account for* as much of the total risk as possible.

4.1. Risk value calculation

The risk associated with an exposure contour depends on many factors. It is convenient to parameterize risk value calculation in the form

$$R = \sum_{c \in C} R(c) \tag{1}$$

$$R = \sum_{c \in C} F_c\, E_c\, P(c)\, N(c), \tag{2}$$

where R is the total risk, a scalar value; C is the set of all contours for all threats and wind conditions; F_c is the threat factor for contour c; E_c is the exposure level factor for contour c; $P(c)$ is the probability of occurrence for the wind condition associated with c; $N(c)$ is the count of population within contour c

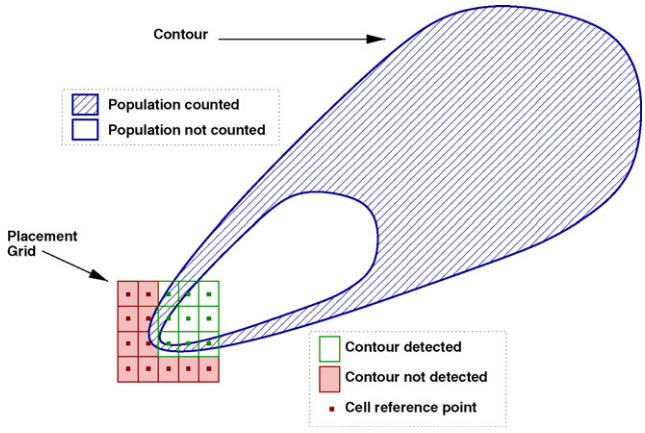

Fig. 3. Contour detection example. Population is not counted in inner polygon rings. Grid cells detect the contour if the reference point lies within the contour's outer polygon ring.

For an example, refer to Fig. 2, a Google Earth display showing contours for 30-min AEGL levels 1–3. Note the contouring shown is from a dosage field for a single dispersion computation, where the wind direction was chosen fortuitously to disperse the material towards a nearby population center. The legend shows the population counts for the respective contours. Note contour population counts do not include areas in inner polygon rings, as illustrated in Fig. 3.

Assume the probability of the wind condition for the associated dispersion computation is 0.1, and the threat factor is 1.0. Since the threat and wind condition are the same for all three contours, $P(c) = 0.1$ and $F_c = 1.0$ for each of them. If we assign exposure level factors (E_c) of 1.0, 2.5, and 4.0 for AEGL-1-30 min, AEGL-2-30 min, and AEGL-3-30 min, respectively, the risk value for each contour calculated according to Eq. (2) is:

$$R(\text{AEGL} - 1\text{–}30\,\text{min}) = 1.0 \times 1.0 \times 0.1 \times 31839 = 3183.9$$
$$R(\text{AEGL} - 2\text{–}30\,\text{min}) = 1.0 \times 2.5 \times 0.1 \times 3575 = 893.75$$
$$R(\text{AEGL} - 3\text{–}30\,\text{min}) = 1.0 \times 4.0 \times 0.1 \times 39 = 15.6$$

4.2. Accounted-for fisk

We wish to detect as much of the total risk as possible by placing sensors in a placement grid. We express the objective function in terms of the placement grid as follows.

A contour is considered to be detected or covered by a sensor in a placement grid cell if the reference point of the cell lies within the outer ring of one of the contour's polygons. We introduce a binary detection function $D(c, g_i)$ where c is a contour, and g_i is a particular grid cell:

$$D(c, g_i) = \begin{cases} 0 & \text{if } r_{g_i} \wedge p = \emptyset \quad \forall p \in P_c \\ 1 & \text{otherwise} \end{cases} \tag{3}$$

where P_c is the set of polygons associated with contour c; r_{g_i} is the reference point for grid cell g_i and; $x \wedge p$ tests for containment of point x in the outer rings of polygon p.

Detection is illustrated in Fig. 3. In the figure, there is a grid of cells in five columns and four rows. Cells with center points not within the contour outer polygon ring are shaded, indicating they do *not* detect or cover the contour.

The risk accounted for by placing one sensor in grid cell g_i can be expressed in terms of contour risk values from Eq. (2):

$$S(g_i) = \sum_{c \in C} D(c, g_i)\, R(c) \tag{4}$$

$$S(g_i) = \sum_{c \in C} D(c, g_i)\, F_c\, E_c\, P(c)\, N(c) \tag{5}$$

Given n sensors, we want to place them in a grid cell configuration $G = \{g_1, \ldots, g_n\}$ that would result in maximum accounted-for risk.

We make no attempt to solve the global optimization problem. Rather, we solve a sequence of weaker optimization problems by finding the optimal grid cell for sequential placement of the sensors in the grid, one at a time. Thus, on the first iteration we find

$$\max_k S(g_k) \tag{6}$$

This fixes the location of the first sensor at $\tilde{g}_1$. On the second iteration $\tilde{g}_1$ is already selected by the previous step, and we solve

$$\max_k S(\tilde{g}_1; g_k),\ g_2 \neq \tilde{g}_1 \tag{7}$$

Generally, for iteration k the cells $g_1, \ldots, g_{k-1}$ have been selected in the previous $k-1$ iterations, and thus we find

$$\max_k S(\tilde{g}_1; \ldots; \tilde{g}_{k-1}; g_k),\ g_k \neq \tilde{g}_1 \ldots \tilde{g}_{k-1} \tag{8}$$

4.3. Marginal and cumulative utility

A byproduct of calculating accounted-for risk in a cell is a simple calculation of the marginal utility of placing a sensor in that cell, the ratio of the accounted-for risk to the total risk. We represent the marginal utility for grid cell g_i as U_{g_i}:

$$U_{g_i} = \frac{S(g_i)}{R} \tag{9}$$

Upon successive iterations in the placement algorithm described in Section 5, the sum of the $S(g_i)$ values for chosen grid cells yields the cumulative accounted-for risk. The quotient of the cumulative accounted-for risk and total risk yields another useful fraction, the cumulative utility. Note this value is monotonically increasing.

5. Placement algorithm

Once the inputs described in Section 3 have been gathered and computed, Eq. (5) can be applied to compute the accounted-for risk for each grid cell and then to determine the optimal cell in which to place a sensor. We execute a dynamic programming algorithm to choose sensor locations in sequence, terminating when one or more criteria are met.

5.1. Sensor placement iterations

Each iteration of the algorithm applies the equations to choose the optimal location (i.e., grid cell) for the next sensor. Locations for any already-placed sensors were determined in prior iterations and are not changed.

The *state* associated with each iteration consists of the set of contours that have yet to be accounted for by a previous grid cell selection. On the first iteration no sensor locations have been chosen yet, and the state consists of all contours for all threats and wind conditions, represented by C in Eq. (5).

An iteration ends with the determination of which contours to include in the contour set for the next iteration. Recalling the detection function of Eq. (3), the contour set for iteration $k+1$ is determined as follows:

$$C_{k+1} = \{c \in C_k :\quad D(c, g_k) = 0,\ g_k \text{ is the selected grid cell for iteration } k\} \tag{10}$$

5.2. Termination criteria

There are three criteria for terminating the placement algorithm. They can be applied independently or in combination.

5.2.1. Fixed number of sensors

There are a fixed number of sensors available to deploy, thus dictating or limiting the number of algorithm iterations.

5.2.2. Minimum cumulative utility

The goal is to achieve a specified cumulative utility described in Section 4.3. Sensors are placed until the minimum is reached. This criterion should be used in combination with the following one, and the specified value must be in the range (0,1).

5.2.3. Marginal utility threshold

Sensors are placed until the marginal utility for the most recent sensor falls below a specified value. The basis for the threshold could be a minimum benefit needed to justify the cost of an additional sensor. This threshold value also must be in the range (0,1).

5.3. Algorithm steps

The placement algorithm operates on two sets of input data and a parameter. One input is the set of contours, where each contour is a geographic polygon and associated risk value calculated as per Eq. (2). Refer to Section 3.2 for a description of how contours are derived.

The second input is the placement grid, or more succinctly the cells comprising the grid, with a reference point specified for each cell. The parameter is the termination criterion or criteria. Refer to Algorithm 1. The bulk of the algorithm is the while-loop which tests whether or not the termination criteria have been met.

Algorithm 1. Place sensors

Require: R, the total risk
Require: C, the set of all contours and their risk values
Require: G, the set of all placement grid cells
1: $C' \leftarrow C$
2: $PlacementCellList \leftarrow$ `empty`
3: $CumAccountedForRisk \leftarrow 0$
4: $Iteration \leftarrow 0$
5: $TerminateFlag \leftarrow$ `false`
6: **while** $TerminateFlag =$ `false` **do**
7: compute $S(g_i) \quad \forall g_i \in G$ (Eq 5)
8: $g' \leftarrow$ the optimal placement cell (Eq 8)
9: $PlacementCellList \leftarrow PlacementCellList + g'$
10: $CumAccountedForRisk \leftarrow CumAccountedForRisk + S(g')$
11: $G \leftarrow G - g'$
12: $Iteration \leftarrow Iteration + 1$
13: **if** termination criteria are met with $CumAccountedForRisk$, $Iteration$, $S(g')/R$ **then**
14: $TerminateFlag \leftarrow$ `true`
15: **else**
16: $C' \leftarrow$ the set of unaccounted-for contours (Eq 10)
17: **end if**
18: **end while**

The first step in the loop, line 7, is the computation of the accounted-for risk for each unused cell in the placement grid, as per Eq. (5). Applying Eq. (8), the cell with the highest accounted-for risk is selected as optimal for situations where sensors may be located in any grid cell. Alternatively, there may be a finite and discrete set of possible sensor locations due to availability of power, line of sight for wireless networks, and similar constraints. In this case grid cells should be sorted by decreasing accounted-for risk value to identify the next grid cell to be selected for sensor placement.

Lines 9–12 add the optimal cell to the list of cells used to place sensors, mark the cell as used, update the cumulative accounted-for risk, and advance the iteration count. Note the accounted-for risk in the selected cell is used to determine the marginal utility, as per Eq. (9).

Line 16 is the application of Equation 10 to determine the contours which have yet to be accounted-for. Depending on how the exposure contours are stored, it might be easier or more efficient to maintain the list of contours to exclude, rather than include, in subsequent processing. That is, line 16 could compute $\bar{C}$, the set of accounted-for contours.

6. Methodology summary

The methodology is composed of two phases: preparation of inputs for the placement algorithm; and execution of the algorithm. It is applied with the following sequence of steps.

Step 1 Define threats and assign threat factors. As described in Section 3.1, each threat must be defined as one or more releases of hazardous materials. In addition, a weighting factor may be assigned to each defined threat. If all threats are considered equally likely, each threat's factor should be 1.0.

Step 2 Define or compute a wind rose. The next step is to establish the wind rose to use for dispersion calculations and as the $P(c)$ term in Eq. (5). Although the wind rose may be specified manually or derived through any means for use in the placement methodology, the idea is to compute the rose from historical data, such as SAMSON data (refer to Section 3.3).

Step 3 Compute threat dispersions and generate contours. The chosen atmospheric transport and dispersion model must be run for each threat at each (direction, speed) pair in the wind rose. This results in a dosage field for each threat and wind condition (refer to Section 3.5).

Each dosage field is then contoured by dosage values corresponding to the qualitative exposure levels defined for the specific material released in the threat definition. (Note exposure level contour values may be obtained from HPAC's material database.) Contours are represented as geographic polygons, where each polygon consists of an outer ring and zero or more inner rings.

Step 4 Assign exposure level factors. A quantitative weighting factor must be assigned for each unique exposure level associated with a contour resulting from Step 3. Refer to Section 3.2 for a discussion of exposure factors.

Step 5 Compute risk values. Using a suitable population distribution, population contained within each exposure contour is calculated, accounting for any inner polygon rings. Each contour is associated with a threat, wind condition, and exposure level. The threat factor, wind condition probability, exposure level factor, and population count are used to calculate the risk associated with each contour, $R(c)$ in Eq. (2). The sum over all $R(c)$ values yields the total risk, R.

Step 6 Specify the placement grid. Section 3.6 describes the placement grid and how it may be defined. For a simple rectilinear grid, it suffices to specify the corners of the grid's geographic boundary and the number of cell columns and rows.

Step 7 Run iterative placement algorithm. After completion of the first six steps, all the inputs necessary to execute Algorithm 1 are available. Termination criteria must be established as described in Section 5.2, and the algorithm is executed to determine optimal sensor locations.

Table 1
Summary of placement results at the Port of Memphis

Iteration	Highest cell value	Selected sensor cell	Accounted-for risk	Marginal utility (%)	Cumulative utility (%)
1	132300	11	124600	49.43	49.43
2	86420	21	74300	29.48	78.91
3	36090	01	32800	13.01	91.92
4	12340	25	3492	1.39	93.31
5	8854	20	2400	0.95	94.26

7. Results

The approach described here has been applied to the placement of chemical sensors at the Port of Memphis in an attempt to mitigate risks associated with a set of threat scenarios defined in prior analysis of the Port's vulnerabilities. The threats were distilled to nine representative scenarios in three different locations and involving a specific set of materials. Sensor locations were constrained to roughly 30 locations around the Port area due to availability of power and other considerations. Results of applying the methodology are presented in tabular and graphic form.

Table 1 is a summary of the results of applying the methodology at the Port of Memphis. Each successive iteration represents a computation of the optimal cell location against all contours not accounted for by a previously selected cell. The total risk from all threats is 252,060.

A visual representation of the placement process is given in a succession of images showing the grid overlaid on the Port area. Grid cells are colored by accounted-for risk. The colormap is a 224-element spectrum starting with blue for low accounted-for risk values and moving to cyan, green, yellow, orange, and then red for the highest values.

Fig. 4 depicts the grid for the first iteration, when all contours are included in the placement computation. Possible sensor locations, based on site constraints, are indicated with circles. Note Table 1 shows the first-iteration maximum cell accounted-for risk value is 132,300, but that cell does not contain one of the possible sensor locations. Of the possible locations, the cell containing Site *11* has the highest value, 124,600, and thus is the optimal available location for the first sensor.

Fig. 5 depicts the grid for the second iteration. All contours accounted for by the cell containing Site *11* are ignored in this iteration. The figure illustrates the remaining areas of the grid containing high values. Again, the cell with the highest value does not contain a possible location site, but Site *21*'s cell (unlabeled, adjacent to Site *22* in the image) has the highest value among the possible locations.

Fig. 6 depicts the grid for the third iteration. All contours accounted for by the cells containing Sites *11* and *21* are ignored. All that remains is the relatively lower valued area in the western end of the Port area. The cell with the highest value among those containing possible sensor locations contains Site *01*.

Fig. 7 depicts the grid for the fourth iteration. At the resolution of the colormap used for the grid depiction, no significant risk value remains to be accounted for or detected. However, Table 1 shows Site *25* as leading possible locations in accounting for what risk remains. Note the marginal utility has fallen below two percent in this iteration, confirming what one can deduce from the visual representation in the image.

After completion of five placement iterations, 94.26% of the total risk for all nine threats has been accounted for by the five sites. Based on criteria established for the Port, iteration five is the last.

8. Comparison with other approaches

As noted above, there are many approaches and algorithms for optimal placement of sensors. Iterative methods comprise one class of such algorithms and is the class into which the methodology and algorithm proposed here falls. Thus, we wish to compare our risk-based approach to other iterative methods.

A sensor placement method representative of iterative approaches is that described by Dhillon et al. [9]. Their method is based on geometry and/or geography without considering human effects. It places sensors using a probability of detection matrix computed as an exponential function of the distance from the grid point to the threat for each proposed placement grid point and each threat location. The Dhillon et al. approach offers many advantages, such as quick calculation and the ability to compute a solution independent of a specific placement grid.

We compare the two approaches by applying the Dhillon et al. method to the problem at the Port of Memphis and using the quantitative accounted-for risk value as defined in Eq. (5) as the basis for comparison. Then, we illustrate the importance of the additional factors considered by our method.

Table 2
Dhillon et al. approach results for the Port of Memphis

Iteration	Selected sensor cell	Miss probability	Accounted-for risk	Marginal utility	Cumulative utility
1	19	0.002137	97390	38.6376	38.6376
2	18	0.002146	18110	7.1848	45.8224
3	21	0.002187	23650	9.3827	55.2051
4	20	0.002191	2016	0.7998	56.0049
5	23	0.002266	181	0.0718	56.0767

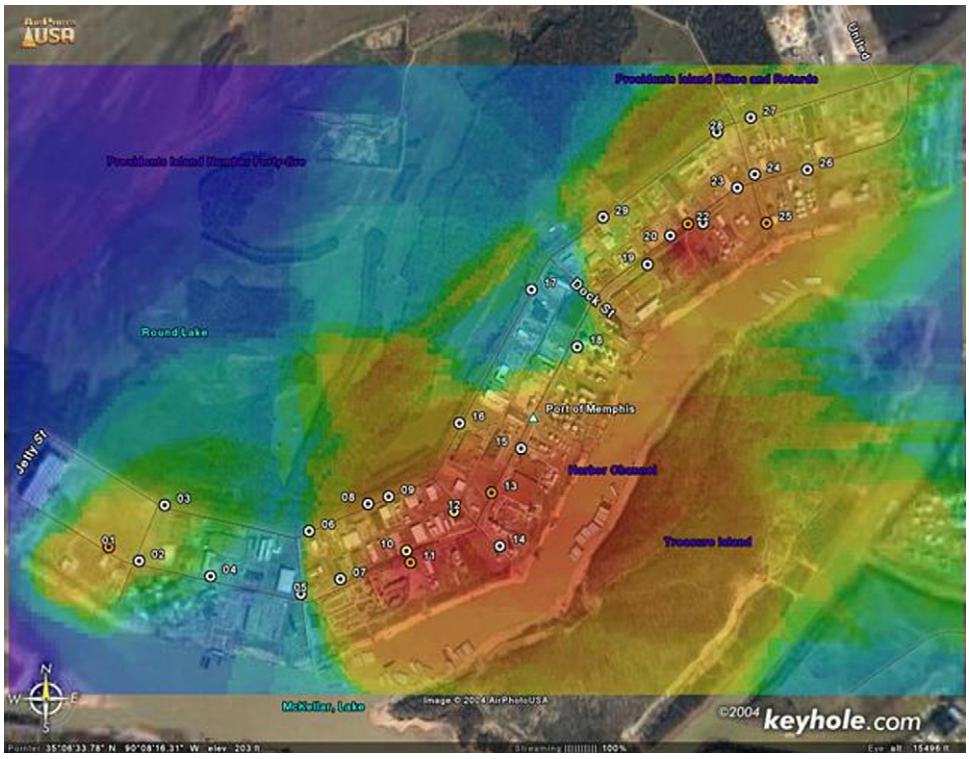

Fig. 4. Placement grid after the first iteration. Grid cells are colored by accounted-for risk in a spectrum from blue (lowest) to cyan, green, yellow, orange, and red (highest). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

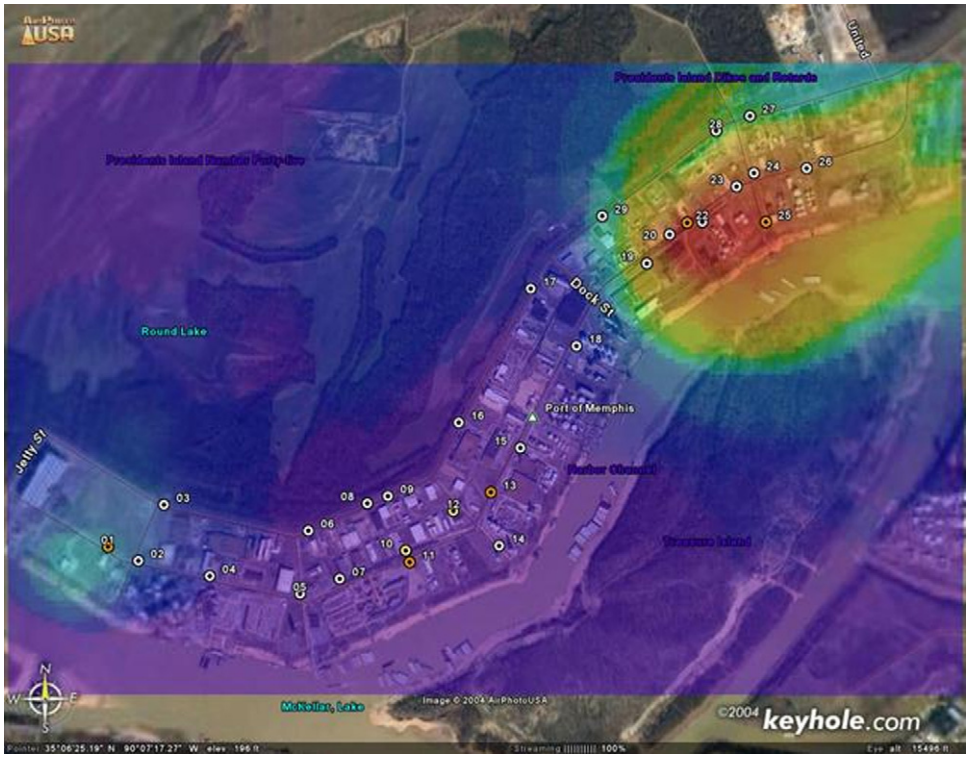

Fig. 5. Placement grid after the second iteration. Grid cells are colored by accounted-for risk in a spectrum from blue (lowest) to cyan, green, yellow, orange, and red (highest). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

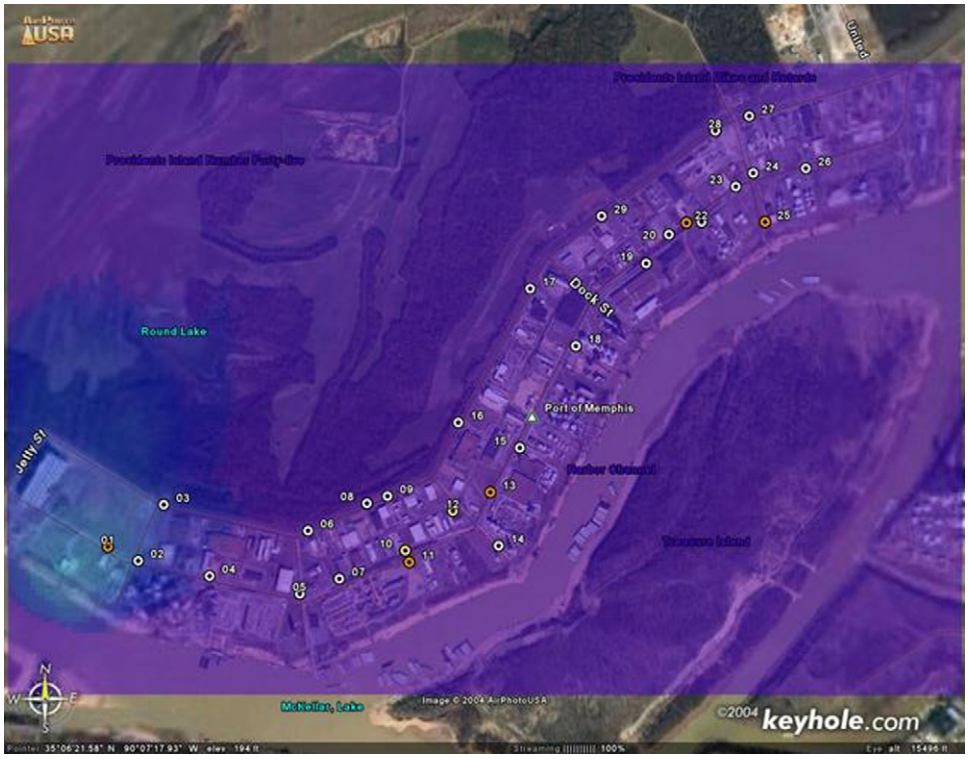

Fig. 6. Placement grid after the third iteration. Grid cells are colored by accounted-for risk in a spectrum from blue (lowest) to cyan, green, yellow, orange, and red (highest). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

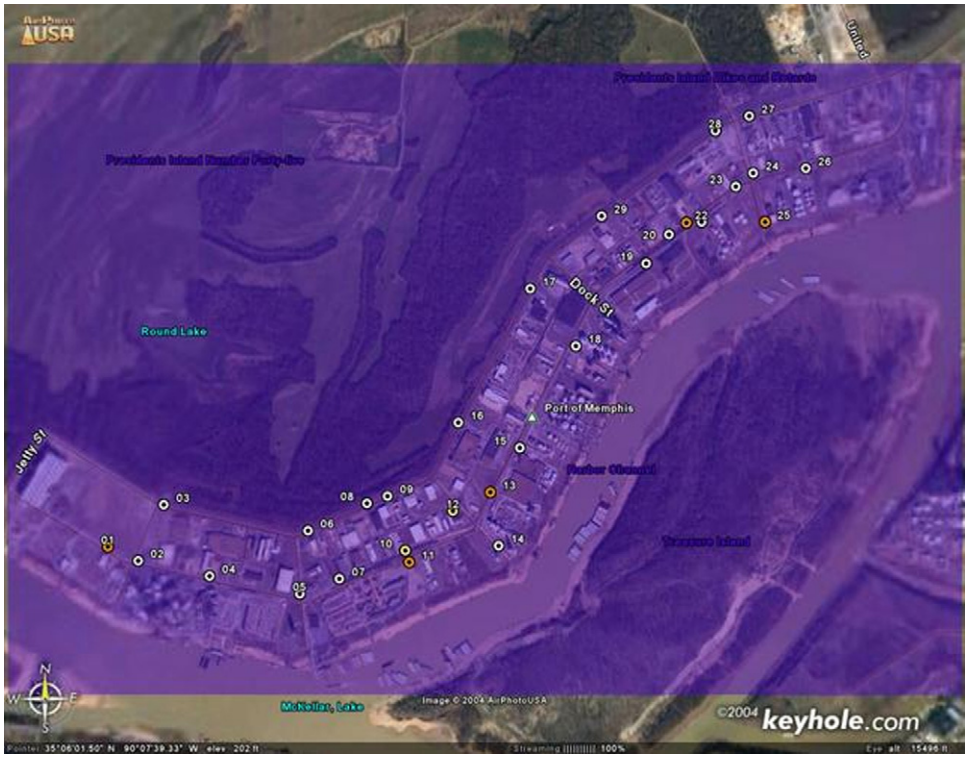

Fig. 7. Placement grid after the fourth iteration. Grid cells are colored by accounted-for risk in a spectrum from blue (lowest) to cyan, green, yellow, orange, and red (highest). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

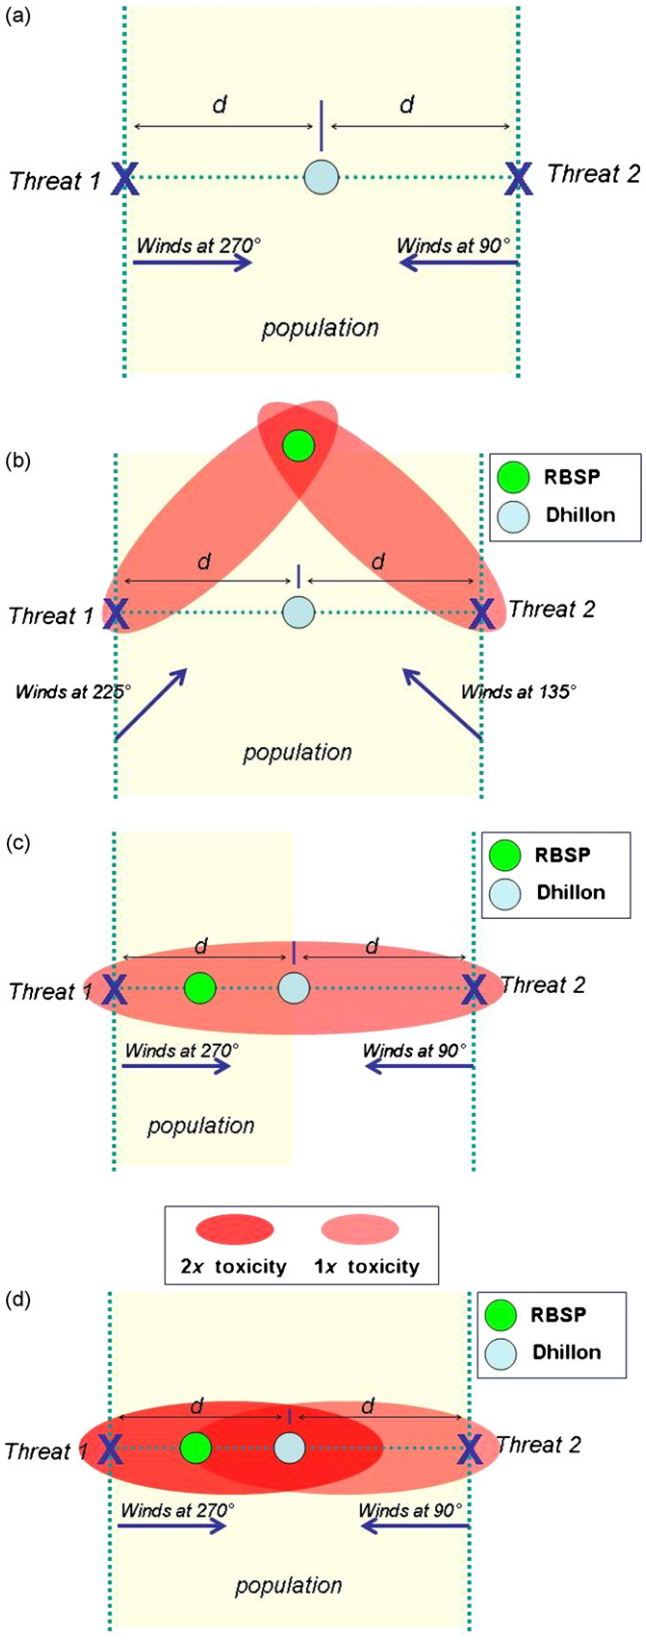

Fig. 8. Scenario in which the algorithms (a) choose the same location, (b) differ due to winds, (c) differ due to population, and (d) differ due to threats.

8.1. Port of Memphis comparison

The Dhillon et al. algorithm applied to the same Port of Memphis problem as described in Section 7, with the available sensor locations illustrated in Fig. 4, yields the results shown in Table 2.

As one would expect, the accounted-for risk values for sensor locations chosen solely on the basis of distance from threats are much lower. Miss probabilities calculated according to the Dhillon et al. algorithm are very small, meaning sensors at the selected locations have a very high probability of detecting a threat based on the algorithm's exponential distance objective function. However, sensors at the five locations resulting from the Dhillon et al. approach account for only 56% of the total risk posed by the defined threats, where total risk is defined by Eq. (2) and accounted-for risk by Eq. (5). As shown in Table 1, the risk-based methodology accounts for 94% of the total risk.

Further, the 49% of total risk accounted for by the first risk-based location nearly equals what five sensors achieve in the distance-based algorithm, and placing a second sensor in the risk-based approach exceeds the accounted-for risk achieved with five sensors in the distance-based algorithm.

8.2. Accounting for the placement factors

The three primary factors accounted for in the proposed risk-based methodology and which are missing in other iterative approaches are winds, population distribution, and the relative toxicity of different materials or agents representing threats. We can construct a scenario in which the two methods compared above yield equivalent results and then vary the three factors to demonstrate the advantage of the risk-based approach.

Suppose there are two threats and one sensor to be placed. The threats are identical and have the same latitude coordinate value. There are two wind conditions, one with direction 90° and the other direction 270° and both with the same wind speed. Population is distributed uniformly between the longitudinal boundaries formed by the threat locations. These baseline conditions are depicted in Fig. 8(a). Note the effects for threat 1 and winds at 90° lie west of the population grid and thus have no population effects. Similarly, effects for threat 2 with winds at 270° lie east of the grid and have no effect.

In this contrived situation, both algorithms will place the single sensor on the same latitude as the threat locations and at a longitude exactly between them. Moreover, the Dhillon algorithm will always place the sensor at the mid-point location regardless of changes in the winds, population, or the release(s) associated with the threats.

8.2.1. Changing the winds

Fig. 8(b) depicts the situation where the identical threats in the two locations are subjected to equally likely winds at 135° and 225°. The dose fields for threat 1 at 135° and threat 2 at 225° lie outside the population area and are thus not shown. Effects contours for threat 1 with 225° winds and threat 2 with 135° winds are represented (unrealistically) as ovals. The risk-based approach will place the sensor in the area where the most risk is accounted for, in this case at the intersection of the two effects

contours. Note the center-line location chosen by the Dhillon et al. approach accounts for none of the effects on population.

8.2.2. Changing the population distribution

Fig. 8(c) represents the situation where the winds are at 90° and 270°, but the population lies completely in the western or left half of the grid. Assuming an equal geographic area of coverage for the effects plumes from the respective threat locations, as illustrated by the oval shape, the risk-based approach will choose a location in the center of the western half of the grid.

8.2.3. Changing the threats

Finally, Fig. 8(d) illustrates the situation differing from the baseline only in the toxicity of the threats at the respective locations. Threat 1 is twice as toxic as threat 2. Assuming the geographic reach of the respective effect plumes are the same, represented by the ovals, the risk-based approach will choose a location in the western half of the grid.

Although these comparisons vary one factor at a time, in reality all the factors will vary together in complex combinations. The proposed risk-based sensor placement methodology accounts for these complex combinations by quantifying the risk posed by each threat and determining where to place sensors to account for the most risk.

In most circumstances a placement of sensors that does not account for these factors will result in an inadequate reckoning of exposure risks, as demonstrated in the comparisons above. These factors comprise the objective function for the proposed risk-based sensor placement methodology. Each sensor's location is chosen from a local optimization of the objective function, and the result of this process is a quantified assessment of population protection against risks.

9. Summary

The sensor placement methodology proposed here attempts to solve the problem of locating sensors to protect population against a set of known and/or postulated threats. An objective function based on population exposure and effects is used to solve a series of local optimizations. Historical meteorological data are used to characterize wind speed and direction and thus drive atmospheric transport and dispersion modeling of the threats, the results of which are used to calculate population at risk in various exposure levels. Sensor locations are determined with a dynamic programming algorithm whereby risk accounted for by sensors placed in prior iterations is not considered in subsequent iterations.

Moreover, the proposed methodology provides a quantification of the marginal utility of each additional sensor. This is the fraction of the total risk accounted for by the sensor. Thus, the criteria for halting the iterative process can be the number of sensors available, a threshold marginal utility value for the most recent sensor, and/or a minimum cumulative utility achieved with all sensors.

Acknowledgements

This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

[1] K. Obenschain, J. Boris, G. Patnaik, Using ct-analyst to optimize sensor placement, Proc. SPIE: Chem. Biol. Sens. V 5416 (2004) 14–20.
[2] Y. Chen, K. Moore, Z. Song, Diffusion boundary determination and zone control via mobile actuator-sensor networks (mas-net): challenges and opportunities, in: Proc. SPIE: Intelligent Computing: Theory and Applications, vol. 5421, 2004, pp. 102–113.
[3] H. Ishida, et al., Plume-tracking robots: a new application of chemical sensors, Biol. Bull. 200 (2001) 222–226.
[4] H. Gonzalez-Banos, J. Latombe, A randomized art-gallery algorithm for sensor placement, in: Proceedings of Seventeenth Annual Symposium on Computational Geometry, SCG 2001, ACM, June 2001.
[5] O. Tekdas, V. Isler, Sensor placement algorithms for triangulation based localization, in Proc. IEEE Int. Conf. on Robotics and Automation, April 2007.
[6] F. Prieto, et al., Cad-based range sensor placement for optimum 3d data acquisition, Mach. Vis. Appl. 15 (2) (2003) 76–91.
[7] Y. Zou, K. Chakrabarty, Sensor deployment and target localization for tactical surveillance, ACM Trans. Embedded Comput. Syst. 3 (2004) 61–91.
[8] S.S. Dhillon, K. Chakrabarty, Sensor placement for effective coverage and surveillance in distributed sensor networks, J. Parallel Distributed Comput. 64 (7) (2004) 788–798.
[9] S.S. Dhillon, K. Chakrabarty, S.S. Iyengar. Sensor placement for grid coverage under imprecise detections, in: Proceedings of Fifth International Conference on Information Fusion, vol. 2, July 2002, pp. 1581–1587.
[10] Defense Threat Reduction Agency, Hazard Prediction and Assessment Capability (HPAC), http://www.dtra.mil/Toolbox/Directorates/td/programs/acec/hpac.cfm.
[11] National Oceanic and Atmospheric Administration, Public Exposure Guidelines, http://response.restoration.noaa.gov/cameo/locs/expguide.html.
[12] Natural Resources Conservation Service, U.S. Department of Agriculture, Wind Rose Data, http://www.wcc.nrcs.usda.gov/climate/windrose.html.
[13] Oak Ridge National Laboratory, LandScan Main Page, 2002, http://www.ornl.gov/sci/gist/landscan/.
[14] Air Resources Laboratory, National Oceanic and Atmospheric Administration, Hybrid Single-Particle Lagrangian Integrated Trajectory Model, http://www.arl.noaa.gov/ready/hysp_info.html.
[15] National Oceanic and Atmospheric Administration. ALOHA, http://response.restoration.noaa.gov/cameo/aloha.htm.
[16] Modeling and Simulation Information Analysis Center, Defense Modeling and Simulation Office, Models, simulations, and other resources, 2002, http://www.msiac.dmso.mil/wmd/msres.asp.
[17] R.I. Sykes, C.P. Cerasoli, D.S. Henn, The representation of dynamic flow effects in a lagrangian puff dispersion model, J. Hazard. Mater. 64 (1999) 223–247.
[18] R.I. Sykes, R.S. Gabruk, A second-order closure model for the effect of averaging time on turbulent plume dispersion, J. Appl. Meter. 36 (1997) 165–184.
[19] S.F. Parker. The HPAC Tool Application Programming Interface, HPAC Version 4.0, Titan Corporation, 2000.