Probabilistic Map Building for Aircraft-Tracking Radars[†]

Technical Report

João P. Hespanha[‡] hespanha@usc.edu

Yusuf S. Ateşkan ateskan@usc.edu

Hüseyin H. Kızılocak kiziloca@usc.edu

Dept. of Electrical Engineering—Systems, Univ. of Southern California
3740 McClintock Ave., Room 318, MC 2563, Los Angeles, CA 90089-2563
phone: (213) 740-9137, fax: (213) 821-1109

December 2000[§]

Abstract

In this report we describe a formalism to assess the risk of missions that consist of flying an aircraft over a region infested by enemy radars and co-located Surface-to-Air Missile (SAM) launchers. This is accomplished by building several probabilistic maps that will allow us to compute the probability of success for the missions.

1 Introduction

In the control of military air-operations, it is common to find scenarios where it is known that a certain number of enemy radars (often called illuminators), with co-located SAM launchers, are positioned in a particular region but their precise positions are not exactly known. This makes air-missions over such a region particularly dangerous. When high-value aircraft take place in these missions, current military practice employs specially equipped aircraft to create a safe path for the main mission. These aircraft are able to detect if they are being tracked by a radar and destroy the radar by firing at the source of the electromagnetic radiation. Creating this type of safe paths is known as Suppression of Enemy Air Defenses (SEAD).

The future of air-operations is evolving towards a scenario where the battlefield is populated with many unmanned air-vehicles gifted with a large number of cheap sensors whose primary mission is to gather information about enemy assets and possibly execute simple missions in an autonomous fashion. In this scenario, the fusion of information coming from a multitude of sensors into forms suitable for mission planning is a major challenge because the data gathered is likely to be noisy. This report addresses precisely this problem. In particular, we develop a probabilistic formalism to process the information coming from aircraft sensors capable of detecting electromagnetic radiation originating from radars, making it suitable for the planning of air-operations. The need for a probabilistic framework stems from the fact that the sensors producing the measurements are not perfect and are susceptible to provide erroneous information. Because of this, any inferences made from this data can never be assumed 100% correct.

[†] This work was supported by the Defense Advanced Research Projects Agency, JFACC Program.

[‡] Corresponding author.

[§] This report supersedes a previous report with the same title, dated September 2000.

A key concept used in this report is that of a probabilistic map. In usual maps of a region, one finds markings that represent locations of objects (e.g., radars) within the region. By placing one of these markings in a location on the map, one is expressing the fact that the object represented can be found at that particular location with probability one. However, when the information used to build a map (or more precisely to locate interesting objects in it) is not perfect, it is not possible to mark specific positions on the map as the location of the object. In this case, the best one can do is to superimpose the probability density of the location of the object in question on the whole map. It is this probability density that we call the probabilistic map for the object. From a formal point of view, probabilistic maps are actually probability densities of the position of objects, conditioned to the sensor measurements available. Note that these probabilistic maps do not represent the topology of the geographic region, but are often related to it, e.g., the probabilistic map of a radar will most likely exhibit zero probability over areas covered by a river because it is unlikely to find an illuminator radar over water.

The main difficulty in computing probabilistic maps is their size. When one is mapping multiple objects whose positions are not independent of each other—and therefore one needs to keep track of the joint probability density of all the objects—the probabilistic map (i.e., the joint probability density) easily becomes very large. This is because, in principle, the memory requirements to represent a joint probability density grow exponentially with the number of random variables. It should be noted that (conditional) independence of the positions of several radars is not expected in our problem because, e.g., the probability of all the radars being very close to each other is often much smaller (or at least different) than the probability of the radars being spread to cover a certain perimeter.

This report contains two main contributions: First we show that, for the sensor models used here, the probabilistic map of a large number of radars can be represented by a function—called the aggregate measurement function—defined in a low-dimensional space (two-dimensional for the specific setup in this report). We then show that it is possible to assess the risk of given missions in air-operations and even do minimum-risk path planning without ever explicitly computing the high-dimensional probabilistic map and working simply with the aggregate measurement function. The algorithms presented here were implemented in a simulator of military air-operations developed by Honeywell Technology Center and were used to guide aircraft in strike missions.

Probabilistic map building is becoming an important topic in the robotics literature. Especially for indoors robotics, map building is often combined with the problem of localizing the robot within a global coordinate system [1, 2, 3, 4, 5, 6]. This is required since the measurements used to build the map are a function of the position of the robot, which is hard to determine accurately due to accumulated odometry errors. In outdoors robotics, the localization problem can be mitigated through the use of GPS. Occupancy grid mapping is currently the most widely used class of mapping algorithms for mobile robots (cf. , e.g., [3, 7]). Occupancy maps simply provide the probability of each individual cell in the grid being occupied by an object. This information is not sufficient to compute the probability of a particular configuration of objects, unless it is assumed that the existence of objects in different cells correspond to independent events. As mentioned above, it is difficult to justify this type of independence assumption for the problem at hand.

This report is organized as follows. In Section 2, we formally define the radar probabilistic map in terms of a joint conditional distribution and then proceed to show how this map can be computed from data acquired by a noisy sensor. The key result in this section is that, although a probabilistic map is a function defined in a high-dimensional space, it can be represented compactly by another function that takes values in a space with the same dimension as the region where the radars are located. In Section 3, we define the radar occupancy map and show how this map is related to the radar probabilistic map. The main result of this section is the derivation of an algorithm to efficiently compute the occupancy map directly from the compact representation of the radar probabilistic map. In Section 4, we show how to compute a danger or kill-map that measures the danger of flying over a certain point in the mapped region. This map is then used to compute the probability of success of a particular flight mission. This probability (or more precisely the logarithm of it) is linear on the values of a simple transformation of the kill-map and is therefore computationally amenable to do minimum-risk path planning. The algorithms presented here were implemented in Honeywell's simulator of military air-operations and the results are shown in Section 5. Finally, Section 6 contains some concluding remarks and directions for future work.

2 Probabilistic maps for aircraft-tracking radars

Suppose that there exist n radars in a region $\mathcal{R} \subset \mathbb{R}^2$, each radar capable of tracking any aircraft that enters its range. For now we assume that the positions of these radars are fixed but not accurately known. For simplicity we take $\mathcal{R}$ to be a finite collection of points, each representing a cells in $\mathbb{R}^2$. We model the position of each radar as a random variable[1] taking values in $\mathcal{R}$. We denote by $\mathbf{r} := \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\} \in \mathcal{R}^n$ the positions of the n radars, where $\mathbf{r}_i \in \mathcal{R}$ denotes the position of the ith radar. We assume that an *a priori* joint probability distribution for the radars

$$p_0(r) := \mathrm{P}(\mathbf{r}_1 = r_1, \mathbf{r}_2 = r_2, \dots, \mathbf{r}_n = r_n), \qquad r := \{r_1, r_2, \dots, r_n\} \in \mathcal{R}^n,$$

is known. We also assume the radars to be *a priori indistinguishable*, meaning that for any permutation $j := \{j(1), j(2), \dots, j(n)\}$ of $\{1, 2, \dots, n\}$ we have

$$p_0(\{r_1, r_2, \dots, r_n\}) = p_0(\{r_{j(1)}, r_{j(2)}, \dots, r_{j(n)}\}). \tag{1}$$

Our knowledge about the position of the radars improves with time through observations collected by the aircraft flying over $\mathcal{R}$. These observations are denoted by $\mathbf{y}_k$, $k \in \mathcal{K} := \{1, 2, \dots, \bar{k}\}$, and each $\mathbf{y}_k$ consists of a triple of the form

$$\{\mathbf{t}_k, \mathbf{a}_k, \mathbf{s}_k\},$$

where $\mathbf{t}_k \geq 0$ denotes the time at which the observation was made and $\mathbf{a}_k \in \mathcal{R}$ the position of the aircraft. The variable $\mathbf{s}_k$ can either have the value $\mathbf{s}_k = \emptyset$ if the aircraft was not being tracked at that time or some value $\mathbf{s}_k \in \mathcal{R}$ if it was being tracked from a radar positioned at $\mathbf{s}_k$. However, we assume that the sensors used to detect tracking radars are not perfect and therefore $\mathbf{s}_k = \emptyset$ simply means that the sensors were not able to detect a radar and one should consider that this may be a "false negative." Similarly, $\mathbf{s}_k \in \mathcal{R}$ may be caused by a "false positive," or there may be a radar close to $\mathbf{s}_k$ but not exactly at that location. The set of all possible observation triples is denoted by $\mathcal{Y}$.

Sensors of this type can be modeled by a stationary likelihood function of the form

$$\ell_n(y, r) := \mathrm{P}(\mathbf{y}_k = y \mid \mathbf{r} = r), \qquad k \in \mathcal{K},\ r \in \mathcal{R}^n,\ y \in \mathcal{Y}.$$

We refer to ℓ_n as the *n-radar sensor likelihood function.* Here, we assume that $\mathbf{y}_k$ is conditionally independent of any other measurements, given a particular configuration $\mathbf{r} = r$ for the radars. This implicitly assumes that sensor errors occur independently of each other. The stationarity assumption could easily be lifted but we keep it here for simplicity of notation.

The goal of this section is to compute the *a posteriori probabilistic map for the radars* defined by

$$m_r(r, Y) := \mathrm{P}(\mathbf{r} = r \mid \mathbf{Y}_K = Y), \qquad r \in \mathcal{R}^n,$$

where $\mathbf{Y}_K$ denotes the list $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_K\}$ of all sensor observations available. Since we will compute m_r recursively, it will be convenient to build probabilistic maps with just a subset of the observations available. With a slight abuse of notation we will use the same symbol m_r for the *k-observations probabilistic map for the radars*, $k \in \mathcal{K}$, defined by

$$m_r(r, Y_k) := \mathrm{P}(\mathbf{r} = r \mid \mathbf{Y}_k = Y_k), \qquad r \in \mathcal{R}^n,$$

where $\mathbf{Y}_k$ denotes a list $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k\}$ with the first k sensor observations. In the sequel we denote by $\mathcal{Y}^*$ the set of all possible lists of measurements with any length.

[1] In this report we denote random variables by **bold** symbols and their realizations by non-bold ones.

2.1 Single-radar sensor likelihood functions

In practice, the precise shape of the sensor likelihood function must be estimated experimentally. We derive next a simple parametric model for a single-radar (i.e., $n = 1$) sensor likelihood functions that could be "best-fitted" to experimental data. Consider the following parameters:

$\rho_{\text{radar}} > 0$	range over which the radar is capable of tracking aircraft.
$p_{\text{track}} \in [0, 1]$	probability that the radar will track an aircraft within its range (in general, could be a function of the radar position and the enemy's intention).
$\rho_{\text{sensor}} > 0$	range over which the sensor can detect a radar tracking the aircraft.
p_{true}	probability that the sensor will detect the radar, given that it is being tracked. We assume that within the range defined by ρ_{sensor} this probability is approximately constant[2].
e_{loc}	maximum localization error for the position of the radar when the sensor correctly detects that the aircraft is being tracked. The actual localization error is assumed uniform within the range specified by e_{loc}. We also assume here that e_{loc} increases linearly with the distance between the aircraft and the radar, i.e., $e_{\text{loc}}(\\|r - a\\|) = e_{\min} \frac{\sqrt{\\|r-a\\|^2+h^2}}{h}$, where h denotes the aircraft altitude. This would be consistent with a fixed uniform distribution for the error in units of angle.
p_{false}	probability that the sensor will detect a radar tracking the aircraft when this is not happening. We assume here that the distribution for the position of the radar "detected" is uniform over the region $\{\bar{s} \in \mathcal{R} : \\|a - \bar{s}\\| < \rho_{\text{sensor}}\}$. For any useful sensor we should always have $p_{\text{false}} < p_{\text{true}}$.

We consider first the case of a measurement $\mathbf{y}_k = \{t, a, s\}$ for which a radar was detected, i.e., $s \in \mathcal{R}$. By definition,

$$\ell_1\big(\{t, a, s\}, r\big) = \mathrm{P}(\mathbf{s}_k = s \mid \mathbf{r}_1 = r, \mathbf{t}_k = t, \mathbf{a}_k = a)\,\mathrm{P}(t_k = t, \mathbf{a}_k = a \mid \mathbf{r}_1 = r). \tag{2}$$

For simplicity we will assume here that the measurements $\mathbf{t}_k = t$ for the time and $\mathbf{a}_k = a$ for the aircraft position are error free and therefore $\mathrm{P}(t_k = t, \mathbf{a}_k = a \mid \mathbf{r}_1 = r) = 1$. Using this and the stationarity assumption, we can expand the right-hand side of (2) to conclude that

$$\begin{aligned}\ell_1\big(\{t, a, s\}, r\big) = {} & \mathrm{P}(\mathbf{s}_k = s \mid \mathbf{r}_1 = r, \mathbf{a}_k = a, \mathfrak{T})\,\mathrm{P}(\mathfrak{T} \mid \mathbf{r}_1 = r, \mathbf{a}_k = a) \\ & + \mathrm{P}(\mathbf{s}_k = s \mid \mathbf{r}_1 = r, \mathbf{a}_k = a, \neg\mathfrak{T})\,\mathrm{P}(\neg\mathfrak{T} \mid \mathbf{r}_1 = r, \mathbf{a}_k = a)\end{aligned}$$

where $\mathfrak{T}$ denotes the event that the radar is tracking the aircraft and $\neg\mathfrak{T}$ its complement. Suppose first that $\|r - a\| > \rho_{\text{radar}}$. In this case the radar will never track the aircraft and we have

$$\ell_1\big(\{t, a, s\}, r\big) = \mathrm{P}(\mathbf{s}_k = s \mid \mathbf{r}_1 = r, \mathbf{a}_k = a, \neg\mathfrak{T}) = \begin{cases} \frac{p_{\text{false}}}{A(a)} & \|a - s\| \le \rho_{\text{sensor}} \\ 0 & \text{otherwise} \end{cases}$$

where $A(a)$ denotes the area of the region $\{\bar{s} \in \mathcal{R} : \|a - \bar{s}\| \le \rho_{\text{sensor}}\}$. Suppose now that $\|r - a\| \le \rho_{\text{radar}}$. According to our sensor model,

$$\mathrm{P}(\mathbf{s}_k = s \mid \mathbf{r}_1 = r, \mathbf{a}_k = a, \mathfrak{T}) = \begin{cases} \frac{p_{\text{true}}}{B(\|r-a\|)} & \|s - r\| \le e_{\text{loc}}(\|r - a\|) \\ 0 & \text{otherwise} \end{cases}$$

[2] A more realistic model would be to assume that this probability is proportional to the ratio

$$\frac{\text{energy that hits aircraft from radar}}{\text{energy that hits aircraft from radar} + \text{energy of background noise}}.$$

Since the energy that hits the aircraft decreases with the square of the distance between the radar and aircraft, this would lead to

$$p_{\text{true}}(\|r - a\|) := \frac{e_r}{\|r - a\|^2 + h^2 + e_r},$$

where h denotes the aircraft altitude and e_r a measure of the energy emitted by the radar, normalized by the energy of the background noise. However, since e_r should be much larger than $\|r - a\|^2 + h^2$ over the range of distances for which the radar would track the aircraft, the previous formula yields an approximately constant value for p_{true}.

where $B(b)$ is the area of the region $\{\bar{s} \in \mathcal{R} : \|\bar{s} - r\| \leq e_{\text{loc}}(b)\}$. We then conclude that, for $\|r - a\| \leq \rho_{\text{radar}}$,

$$\ell_1\big(\{t, a, s\}, r\big) = \begin{cases} \frac{p_{\text{true}} p_{\text{track}}}{B(\|r-a\|)} & \|s - r\| \leq e_{\text{loc}}(\|r - a\|) \\ 0 & \text{otherwise} \end{cases} + \begin{cases} \frac{p_{\text{false}}(1-p_{\text{track}})}{A(a)} & \|a - s\| \leq \rho_{\text{sensor}} \\ 0 & \text{otherwise} \end{cases}$$

Once the likelihood function is known for each $s \in \mathcal{R}$, its values for $s = \emptyset$ can be obtained using

$$\ell_1\big(\{t, a, \emptyset\}, r\big) = 1 - \sum_{\bar{s} \in \mathcal{R}} \ell_1\big(\{t, a, \bar{s}\}, r\big) \tag{3}$$

$$= \begin{cases} 1 - p_{\text{true}} p_{\text{track}} - p_{\text{false}}(1 - p_{\text{track}}) & \|r - a\| \leq \rho_{\text{sensor}} \\ 1 - p_{\text{false}} & \text{otherwise} \end{cases} \tag{4}$$

A typical single-radar sensor likelihood function is plotted in Figure 1. The parameters used for these plots were: $\rho_{\text{radar}} = 3$, $p_{\text{track}} = .9$, $\rho_{\text{sensor}} = 4$, $p_{\text{true}} = .95$, $h = 2$, $e_{\min} = 1$, and $p_{\text{false}} = .05$.

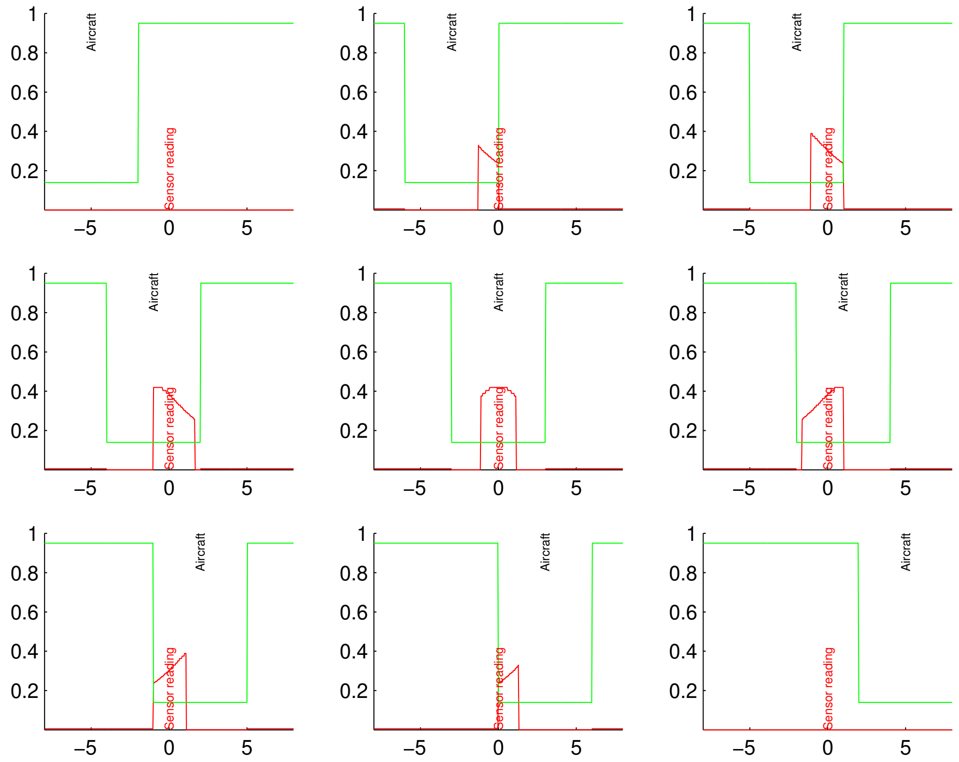

Figure 1: Typical single-radar sensor likelihood function. The plots represent $\ell_1(\{t, a, \emptyset\}, r)$ (light line) and $\ell_1(\{t, a, 0\}, r)$ (dark line) for several positions of the aircraft.

2.2 Multiple-radar sensor likelihood functions

The computation of the likelihood function for multiple radars presented below is based on the following basic assumptions:

1. From the sensor point of view, the radars are indistinguishable from each other.

2. The distance between any two radars is larger than their range ρ_{radar} [3].

[3] In case some radars are placed very close to each other, we can just regard these radars as a single air-defense unit.

We consider first the case of a measurement $\mathbf{y}_k = y := \{t, a, s\} \in \mathcal{Y}$ for which a radar was detected, i.e., with $s \in \mathcal{R}$. As before, we can expand

$$\begin{aligned}\ell_n\big(y, r\big) = \mathrm{P}(\mathbf{s}_k = s \mid \mathbf{r}_{i^*} = r_{i^*}, \mathbf{a}_k = a, \mathfrak{T})\, \mathrm{P}(\mathfrak{T} \mid \mathbf{r} = r, \mathbf{a}_k = a) \\ + \mathrm{P}(\mathbf{s}_k = s \mid \mathbf{a}_k = a, \neg\mathfrak{T})\, \mathrm{P}(\neg\mathfrak{T} \mid \mathbf{r} = r, \mathbf{a}_k = a), \qquad r \in \mathcal{R}^n,\end{aligned}$$

where $\mathfrak{T}$ denotes the event that one of the radars is tracking the aircraft, $\neg\mathfrak{T}$ its complement, i^* the index of the actual radar that is doing the tracking, and r_{i^*} its positions. Reasoning as before we conclude that

$$\ell_n\big(y, r\big) = \ell_1\big(y, r_{i^*}\big),$$

where ℓ_1 is the single-radar likelihood function computed before. Since the aircraft is within the range of at most one radar, we also have

$$\prod_{i=1}^{n} \ell_1(y, r_i) = \left(\frac{p_{\text{false}}}{A(a)}\right)^{n-1} \ell_1(y, r_{i^*}) = \left(\frac{p_{\text{false}}}{A(a)}\right)^{n-1} \ell_n\big(y, r\big), \qquad \|a - s\| \le \rho_{\text{sensor}}.$$

This leads to

$$\ell_n\big(y, r\big) = c \prod_{i=1}^{n} \ell_1(y, r_i), \tag{5}$$

where c is a normalizing constant[4].

Let us consider now the case $\mathbf{y}_k = y := \{t, a, \emptyset\} \in \mathcal{Y}$. We can expand

$$\begin{aligned}\ell_n\big(\{t, a, \emptyset\}, r\big) = \mathrm{P}(\mathbf{s}_k = \emptyset \mid \mathbf{r}_{i^*} = r_{i^*}, \mathbf{a}_k = a, \mathfrak{T})\, \mathrm{P}(\mathfrak{T} \mid \mathbf{r} = r, \mathbf{a}_k = a) \\ + \mathrm{P}(\mathbf{s}_k = \emptyset \mid \mathbf{a}_k = a, \neg\mathfrak{T})\, \mathrm{P}(\neg\mathfrak{T} \mid \mathbf{r} = r, \mathbf{a}_k = a),\end{aligned}$$

where $\mathfrak{T}$, $\neg\mathfrak{T}$, i^*, and r_{i^*} have the same meaning as before. From the definitions of the parameters introduced at the beginning of this section we conclude that

$$\begin{aligned}&\mathrm{P}(\mathbf{s}_k = \emptyset \mid \mathbf{r}_{i^*} = r_{i^*}, \mathbf{a}_k = a, \mathfrak{T}) = 1 - p_{\text{true}} \\ &\mathrm{P}(\mathbf{s}_k = \emptyset \mid \mathbf{a}_k = a, \neg\mathfrak{T}) = 1 - p_{\text{false}} \\ &\mathrm{P}(\mathfrak{T} \mid \mathbf{r} = r, \mathbf{a}_k = a) = \begin{cases} p_{\text{track}} & \exists i : \|r_i - a\| \le \rho_{\text{radar}} \\ 0 & \text{otherwise} \end{cases}\end{aligned}$$

and therefore

$$\ell_n\big(\{t, a, \emptyset\}, r\big) = (1 - p_{\text{false}}) - (p_{\text{true}} - p_{\text{false}}) \begin{cases} p_{\text{track}} & \exists i : \|r_i - a\| \le \rho_{\text{radar}} \\ 0 & \text{otherwise} \end{cases} = \ell_1\big(\{t, a, \emptyset\}, r_{i^+}\big),$$

where i^+ is the index of the radar closest to the aircraft. Since the aircraft is within the range of at most one radar, we also have

$$\prod_{i=1}^{n} \ell_1(y, r_i) = (1 - p_{\text{false}})^{n-1} \ell_1(y, r_{i^+}) = (1 - p_{\text{false}})^{n-1} \ell_n\big(y, r\big).$$

We therefore conclude that (5) also holds in this case.

[4] With some abuse of notation, in this report we use the same symbol c to denote different normalizing constants when this does not lead to ambiguity.

2.3 Probabilistic map for the radars

Pick now some $k \in \mathcal{K}$ and a list $Y_k := \{y_1, y_2, \dots, y_k\}$ of k measurements. Using Bayes rule we conclude that, for the sensor model described above,

$$m_r(r, Y_k) = c_k \, \mathrm{P}(\mathbf{y}_k = y_k \mid \mathbf{r} = r, \mathbf{Y}_{k-1} = Y_{k-1}) \, m_r(r, Y_{k-1}) = c_k \, \ell_n(y_k, r) \, m_r(r, Y_{k-1}) \tag{6}$$

where $Y_k := \{Y_{k-1}, y_k\}$, $r \in \mathcal{R}^n$, and $c_k > 0$ is a constant chosen so that $\sum_{r\in\mathcal{R}} m_r(r, Y_k) = 1$. Iterating the above from 1 through k we get

$$m_r(r, Y_k) = \bar{c}_k \Big(\prod_{j=1}^{k} \ell_n(y_j, r) \Big) p_0(r), \qquad r \in \mathcal{R}^n \tag{7}$$

where $\bar{c}_k$ is a normalizing constant. From (1) and (7) it is also straightforward to conclude that the radars remain *a posteriori indistinguishable*, meaning that for any permutation $j := \{j(1), j(2), \dots, j(n)\}$ of $\{1, 2, \dots, n\}$ we have

$$m_r(\{r_1, r_2, \dots, r_n\}, Y_k) = m_r(\{r_{j(1)}, r_{j(2)}, \dots, r_{j(n)}\}, Y_k).$$

When the *a priori* probabilistic map precludes configurations where the distance between to radars is larger than ρ_{radar}, the likelihood function ℓ_n derived in Section 2.2 is valid and we can write (7) as

$$m_r(r, Y_k) = \hat{c}_k \Big(\prod_{i=1}^{n} f(r_i, Y_k) \Big) p_0(r), \qquad r := \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\} \in \mathcal{R}^n, \tag{8}$$

where $\hat{c}_k$ is a normalizing constant and

$$f(x, Y_k) := \prod_{j\in\mathcal{K}} \ell_1(y_j, x), \qquad x \in \mathcal{R}.$$

The function $f(\cdot, Y_k)$—called here the *aggregate measurement function*—encapsulates the information contained in all the measurement. This is important since for any region of reasonable size, the posterior probabilistic map $m_r(\cdot, Y_K)$ cannot be represented exhaustively because such representation would require a number of values of the order of $|\mathcal{R}|^n$, where $|\mathcal{R}|$ is the number of points in $\mathcal{R}$. However, using (8) one can represent $m_r(\cdot, Y_K)$ using a single function $f(\cdot, Y_k)$, whose exhaustive representation only requires $|\mathcal{R}|$ values.

2.4 Confidence on the probabilistic map

The maximum *a posteriori* estimate for the position of the radars, given a list $Y_k := \{y_1, y_2, \dots, y_k\}$ of k measurements is given by

$$\hat{r} := \arg\max_{r\in\mathcal{R}^n} m_r(r, Y_k).$$

Such estimate minimizes the *a posteriori* probability of error

$$1 - \mathrm{P}(\mathbf{r} = \hat{r} \mid \mathbf{Y}_k = Y_k).$$

A measure of confidence on the probabilistic map is the ration between this quantity and the probability of error obtained for a uniform distribution of the radar position (complete absence of information):

$$\alpha := \frac{1 - \max \mathrm{P}_{r\in\mathcal{R}^n}(\mathbf{r} = r \mid \mathbf{Y}_k = Y_k)}{1 - \frac{1}{|\mathcal{R}|}},$$

where $|\mathcal{R}|$ denotes the area of $\mathcal{R}$. A value of α close to zero gives high confidence in the *a posteriori* map, whereas a value close to one gives very low confidence.

3 Occupancy maps for aircraft-tracking radars

The goal of this section is to compute the *a posteriori occupancy maps for the radars* defined by

$$m_o(x,Y) := \mathrm{P}(\exists i : r_i = x \mid \mathbf{Y}_k = Y), \qquad\qquad x \in \mathcal{R},$$

where $Y \in \mathcal{Y}^*$ denotes a list of k observations. From the definition of the probabilistic map $m_r(\cdot, Y)$, we conclude that

$$m_o(x,Y) = \sum_{\{r\in\mathcal{R}^n : \exists i : r_i = x\}} m_r(r,Y). \tag{9}$$

When the *a priori* joint probability precludes two radars from being in the same cell, i.e.,

$$r_i = r_j,\ i \neq j \quad \Rightarrow \quad p_0(\{r_1, r_2, \dots, r_n\}) = 0,$$

the same will hold for the *a posteriori* map m_r and (9) can be written as

$$\begin{aligned} m_o(x,Y) &= \sum_{m=1}^{n} \sum_{\{r\in\mathcal{R}^n : r_m = x\}} m_r(r,Y) \\ &= n \sum_{r_j \in \mathcal{R},\ j>1} m_r(\{x, r_2, \dots, r_n\}, Y). \end{aligned} \tag{10}$$

Here, we also used the indistinguishability. Note that, in this case,

$$\sum_{x\in\mathcal{R}} m_o(x,Y) = n. \tag{11}$$

Suppose now that the only *a priori* information about the positions of the radars is that there is a minimum distance $d_{\min}$ between any two of them. The *a priori* joint probability map is then of the form

$$p_0(\{r_1, r_2, \dots, r_n\}) := \begin{cases} 0 & \exists i \neq j : \|r_i - r_j\| \leq d_{\min} \\ c & \text{otherwise} \end{cases} \qquad\qquad \{r_1, r_2, \dots, r_n\} \in \mathcal{R}^n. \tag{12}$$

where c is a positive constant. In this case, using the *a posteriori* map given by (8), we obtain

$$m_o(x,Y) = \bar{c}\, f(x,Y) \sum_{r_2\in\mathcal{R}[x]} f(r_2,Y) \sum_{r_3\in\mathcal{R}[x]\cap\mathcal{R}[r_2]} f(r_3,Y) \cdots \sum_{r_n\in\mathcal{R}[x]\cap\mathcal{R}[r_2]\cap\cdots\cap\mathcal{R}[r_{n-1}]} f(r_n,Y) \tag{13}$$

where $\bar{c}$ is a normalizing constant chosen so that (11) holds and $\mathcal{R}[z] := \{\bar{z} \in \mathcal{R} : \|\bar{z} - z\| > d_{\min}\}$, $z \in \mathcal{R}$. A worst-case upper bound on the *a posteriori* occupancy map for the radars can then be computed by

$$m_o(x,Y) \leq \bar{c}\, f(x,Y) \sum_{r_2\in\mathcal{R}[x]} f(r_2,Y) \sum_{r_3\in\mathcal{R}[x]} f(r_3,Y) \cdots \sum_{r_n\in\mathcal{R}[x]} f(r_n,Y) = \bar{c}\bar{f}^{n-1} f(x,Y),$$

where $\bar{f} := \sum_{z\in\mathcal{R}[x]} f(z,Y)$. However, this bound is difficult to use because it not easy to compute the normalizing constant.

Fix now some $Y \in \mathcal{Y}^*$ and suppose that there exists a partition $\mathcal{R}_0, \mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_m$ of $\mathcal{R}$ such that

1. all points in each $\mathcal{R}_i$, $i \in \{1, 2, \dots, m\}$, are away from each other by no more than $d_{\min}$, i.e.,

   $$\|z - \bar{z}\| \leq d_{\min} \qquad\qquad \forall z, \bar{z} \in \mathcal{R}_i,$$

   therefore there can be no more than one radar in each $\mathcal{R}_i$;

2. all the $\mathcal{R}_i$, $i \in \{1, 2, \ldots, m\}$, are away from each other by more than $d_{\min}$, i.e.,

$$\|z - \bar{z}\| > d_{\min} \qquad\qquad \forall z \in \mathcal{R}_i, \bar{z} \in \mathcal{R}_j,\ i, j \in \{1, 2, \ldots, m\},\ i \neq j,$$

therefore different radars can be in any position in distinct $\mathcal{R}_i$, $i \in \{1, 2, \ldots, m\}$;

3. the values of $f(z, Y)$ are sufficiently uniform over $\mathcal{R}_0$ so that, for any $\bar{\mathcal{R}} \subset \mathcal{R}_0$ with small area compared to $\mathcal{R}_0$,

$$\sum_{z \in \mathcal{R}_0 \setminus \bar{\mathcal{R}}} f(z, Y) \approx \frac{|\mathcal{R}_0| - |\bar{\mathcal{R}}|}{|\mathcal{R}_0|} \sum_{z \in \mathcal{R}_0} f(z, Y) \approx \sum_{z \in \mathcal{R}_0} f(z, Y), \tag{14}$$

where $|\mathcal{R}_0|$ and $|\bar{\mathcal{R}}|$ denote the areas of $\mathcal{R}_0$ and $\bar{\mathcal{R}}$, respectively. The first approximation above comes from the uniformity of $f(z, Y)$ over $\mathcal{R}_0$, whereas the second one comes from the fact that the area of $\bar{\mathcal{R}}$ is small. Note that, even when (14) does not hold, we still have the following upper and lower bounds on $\sum_{z \in \mathcal{R}_0 \setminus \bar{\mathcal{R}}} f(z, Y)$.

$$\Big(\sum_{z \in \mathcal{R}_0} f(z, Y)\Big) - |\bar{\mathcal{R}}| \max_{z \in \mathcal{R}_0} f(z, Y) \leq \sum_{z \in \mathcal{R}_0 \setminus \bar{\mathcal{R}}} f(z, Y) \leq \sum_{z \in \mathcal{R}_0} f(z, Y). \tag{15}$$

Pick now some $x \in \mathcal{R}$. Using the fact that the $\mathcal{R}_i$ form a partition of $\mathcal{R}$ and the indistinguishability, we can write (10) as

$$m_o(x, Y) = n\Bigg(\sum_{\substack{r_2, \ldots, r_n \\ \in \mathcal{R}_0}} m_r(\{x, r_2, \ldots, r_n\}, Y) + \sum_{k=1}^{n-1} \binom{n-1}{k} \sum_{\substack{r_2, \ldots, r_{1+k} \\ \in \cup_{i=1}^m \mathcal{R}_i}} \sum_{\substack{r_{2+k}, \ldots, r_n \\ \in \mathcal{R}_0}} m_r(\{x, r_2, \ldots, r_n\}, Y)\Bigg). \tag{16}$$

Let us denote by $\sigma := \{\sigma_1, \sigma_2, \ldots, \sigma_k\}$ a combination of k elements taken from the set $\{i : x \notin \mathcal{R}_i\} \subset \{1, 2, \ldots, m\}$. We denote by $\mathcal{S}^{(k)}$ the set of all such combinations. Using the fact that there can be no more than one radar in each $\mathcal{R}_i$, $i \in \{1, \ldots, m\}$, we can further expand (16) as

$$\begin{aligned} m_o(x, Y) &= n \sum_{\substack{r_2, \ldots, r_n \\ \in \mathcal{R}_0}} m_r(\{x, r_2, \ldots, r_n\}, Y) + n \sum_{k=1}^{n-1} \binom{n-1}{k} k! \sum_{\sigma \in \mathcal{S}^{(k)}} \sum_{\substack{r_{i+1} \in \mathcal{R}_{\sigma_i} \\ i \in \{1, \ldots, k\}}} \sum_{\substack{r_{k+2}, \ldots, r_n \\ \in \mathcal{R}_0}} m_r(\{x, r_2, \ldots, r_n\}, Y) \\ &= n \sum_{\substack{r_2, \ldots, r_n \\ \in \mathcal{R}_0}} m_r(\{x, r_2, \ldots, r_n\}, Y) + \sum_{k=1}^{n-1} \frac{n!}{(n-1-k)!} \sum_{\sigma \in \mathcal{S}^{(k)}} \sum_{\substack{r_{i+1} \in \mathcal{R}_{\sigma_i} \\ i \in \{1, \ldots, k\}}} \sum_{\substack{r_{k+2}, \ldots, r_n \\ \in \mathcal{R}_0}} m_r(\{x, r_2, \ldots, r_n\}, Y) \end{aligned}$$

In this case, using the approximate *a posteriori* map given by (8), we obtain

$$\begin{aligned} m_o(x, Y) = \bar{c} f(x, Y) \Bigg(& n \sum_{\substack{r_2, \ldots, r_n \\ \in \mathcal{R}_0}} f(r_2, Y) \cdots f(r_n, Y) p_0(\{x, r_2, \ldots, r_n\}) + \sum_{k=1}^{n-1} \frac{n!}{(n-1-k)!} \sum_{\sigma \in \mathcal{S}^{(k)}} \\ & \sum_{\substack{r_{i+1} \in \mathcal{R}_{\sigma_i} \\ i \in \{1, \ldots, k\}}} f(r_2, Y) \cdots f(r_{k+1}, Y) \sum_{\substack{r_{k+2}, \ldots, r_n \\ \in \mathcal{R}_0}} f(r_{k+2}, Y) \cdots f(r_n, Y) p_0(\{x, r_2, \ldots, r_n\}) \Bigg) \end{aligned} \tag{17}$$

where $\bar{c}$ is a normalizing constant chosen so that (11) holds. However, from (12), we conclude that for $x, r_1, r_2, \ldots, r_{k+1}$ away from each other by more than $d_{\min}$,

$$\sum_{\substack{r_{k+2},\dots,r_n\\ \in\mathcal{R}_0}} f(r_{k+2},Y)\cdots f(r_n,Y)p_0(\{x,r_2,\dots,r_n\}) = c \sum_{r_{k+2}\in\bar{\mathcal{R}}[x,r_1,\dots,r_{k+1}]} f(r_{k+2},Y)$$
$$\sum_{r_{k+3}\in\bar{\mathcal{R}}[x,r_1,\dots,r_{k+2}]} f(r_{k+3},Y) \quad \cdots \sum_{r_n\in\bar{\mathcal{R}}[x,r_1,\dots,r_{n-1}]} f(r_n,Y),$$

where $\bar{\mathcal{R}}[z_1,z_2,\dots,z_j] := \{\bar{z}\in\mathcal{R}_0 : \|\bar{z}-z_i\| > d_{\min}, \text{for some } i\}$, with each $z_i \in \mathcal{R}$. When $\pi d_{\min}^2$ is small when compared to the total area of $\mathcal{R}_0$, because of (14), we further conclude that[5]

$$\sum_{\substack{r_{k+2},\dots,r_n\\ \in\mathcal{R}_0}} f(r_{k+2},Y)\cdots f(r_n,Y)p_0(\{x,r_2,\dots,r_n\}) \approx c f_0(Y)^{n-k-1},$$

where $f_i(Y) := \sum_{r\in\mathcal{R}_i} f(r,Y)$, $i\in\{0,1,\dots,m\}$. Replacing this in (17) we obtain

$$\begin{aligned}
m_o(x,Y) &\approx \bar{c}f(x,Y)\Bigg(nf_0(Y)^{n-1} + \sum_{k=1}^{n-1}\frac{n!\,f_0(Y)^{n-k-1}}{(n-1-k)!}\sum_{\sigma\in\mathcal{S}^{(k)}}\sum_{\substack{r_{i+1}\in\mathcal{R}_{\sigma_i}\\ i\in\{1,\dots,k\}}} f(r_2,Y)\cdots f(r_{k+1},Y)\Bigg)\\
&= \bar{c}f(x,Y)\Bigg(nf_0(Y)^{n-1} + \sum_{k=1}^{n-1}\frac{n!\,f_0(Y)^{n-k-1}}{(n-1-k)!}\sum_{\sigma\in\mathcal{S}^{(k)}} f_{\sigma_1}(Y)f_{\sigma_2}(Y)\dots f_{\sigma_k}(Y)\Bigg), \qquad (18)
\end{aligned}$$

where the constant $\bar{c}$ was redefined to absorb c. Equation (18) allows us to compute the *a posteriori* occupancy map $m_o(\cdot,Y)$ directly from the aggregate measurement function $f(\cdot,Y)$, without explicitly computing the *a posteriori* probabilistic map $m_r(\cdot,Y)$.

4 Mission risk assessment

The goal of this section is to compute the probability of success of a particular flight mission using the *a posteriori* occupancy map defined above. To this effect, let us start by computing the *a posteriori kill map* defined by

$$m_{\text{kill}}(a,Y) := \mathrm{P}(\mathfrak{K} \mid \mathbf{Y}_k = Y, \mathbf{a}_k = a)$$

were Y denotes a list of k observations, $\mathfrak{K}$ the event that an aircraft positioned at a will be tracked by a radar, shot at by a SAM co-located with the radar, and destroyed as a result. To compute the kill map we need the parameters ρ_{radar} and p_{track} defined before, together with the following one:

$p_{\text{kill}} \in [0,1]$	probability that a missile will destroy an aircraft that is being tracked (in general, could be a function of the aircraft position and the enemy's intention).

Denoting by $\mathfrak{R}$ the event that the aircraft is within range of at least one radar and by $\mathfrak{T}$ the event that it will be tracked by that radar, we have that

$$\begin{aligned}
m_{\text{kill}}(a,Y) &:= \mathrm{P}(\mathfrak{K}\mid \mathbf{Y}_k = Y, \mathbf{a}_k = a, \mathfrak{T})\,\mathrm{P}(\mathfrak{T}\mid \mathbf{Y}_k = Y, \mathbf{a}_k = a, \mathfrak{R})\,\mathrm{P}(\mathfrak{R}\mid \mathbf{Y}_k = Y, \mathbf{a}_k = a)\\
&= p_{\text{kill}}\, p_{\text{track}} \sum_{\exists i : \|r_i - a\|\le\rho_{\text{radar}}} m_r(r,Y).
\end{aligned}$$

A lower bound on $m_{\text{kill}}(a,Y)$ can be computed by allowing "double-counting" when the aircraft is within range of multiple radars. When this is done we obtain the following "worst-case" bound

$$m_{\text{kill}}(a,Y) \le p_{\text{kill}}\, p_{\text{track}} \sum_{\|x-a\|\le\rho_{\text{radar}}} m_o(x,Y). \qquad (19)$$

[5]When (14) is not valid we can still use (15) to get upper and lower bounds on the *a posteriori* occupancy map.

In the above formula, we actually have an equality when the aircraft is within range of at most one radar. This can be know to occur with probability one, e.g., when the prior probability places the radar far apart.

Suppose now that we are interested in having an aircraft flying a mission whose trajectory is defined by m way points $\{x_1, x_2, \dots, x_m\}$. The conditional *a posteriori* probability p_{survive} that an aircraft will survive this mission, given a list of past measurement $\mathbf{Y}_k = Y \in \mathcal{Y}^*$, can then be computed as

$$p_{\text{survive}} = \mathrm{P}(\mathfrak{S}_k \mid Y_k = Y) = \prod_{j=1}^{k} \mathrm{P}(\mathfrak{S}_j \mid Y_k = Y, \mathfrak{S}_{j-1})$$

where $\mathfrak{S}_j$, $j \in \{0, 1, 2, \dots, m\}$ denotes the event that the aircraft will survive up to the jth way point (by convention $\mathfrak{S}_0$ is the whole sample space). Since $\mathrm{P}(\mathfrak{S}_j \mid Y_k = Y, \mathfrak{S}_{j-1})$ is the conditional probability that an aircraft flying over x_j will not be hit at that point, we have that

$$\mathrm{P}(\mathfrak{S}_j \mid Y_k = Y, \mathfrak{S}_{j-1}) = 1 - m_{\text{kill}}(x_j, Y)$$

and therefore

$$p_{\text{survive}} = \prod_{j=1}^{m} (1 - m_{\text{kill}}(x_j, Y)).$$

It is often convenient to express this as a summation, which can be achieved taking the logarithm of p_{survive}:

$$\log(p_{\text{survive}}) = \sum_{j=1}^{m} \log_{\text{live}}(x_j, Y), \tag{20}$$

where

$$\log_{\text{live}}(x, Y) := \log\big(1 - m_{\text{kill}}(x, Y)\big), \qquad x \in \mathcal{R},$$

is called the *log-live map*. Using (19) we can obtain the following worst case bound for the log-live map

$$\log_{\text{live}}(x, Y) \geq \log\Big(1 - p_{\text{kill}}\, p_{\text{track}} \sum_{\|\bar{x} - x\| \leq \rho_{\text{radar}}} m_o(\bar{x}, Y)\Big), \qquad x \in \mathcal{R}.$$

Note that, because of the monotonicity of the logarithm, optimizing for $\log(p_{\text{survive}})$ is equivalent to optimizing for p_{survive}. We can therefore compute a m-step path of minimum-risk from x_{init} to x_{final} using (20):

$$\min_{\substack{x_1, x_2, \dots, x_m: \\ x_1 = x_{\text{init}},\ x_m = x_{\text{final}}}} \sum_{j=1}^{m} \log_{\text{live}}(x_j, Y),$$

where the minimization is further constrained by the aircraft's kinematic model. It is also straightforward to compute paths that are Pareto-optimal with respect to the two costs: probability of survival p_{survive} and path length m.

5 Simulation Results

The algorithms presented above were implemented in a simulator of military air-operations developed by the Honeywell Technology Center. Figure 2 shows a snapshot of the main simulation window displaying the battlefield during a mission that consisted of flying several aircraft from one air-base to another, passing through a region infested by enemy SAMs. In this figure, enemy radars with co-located SAMs are represented by large triangles and the friendly aircraft are represented by small circles or triangles (depending on the

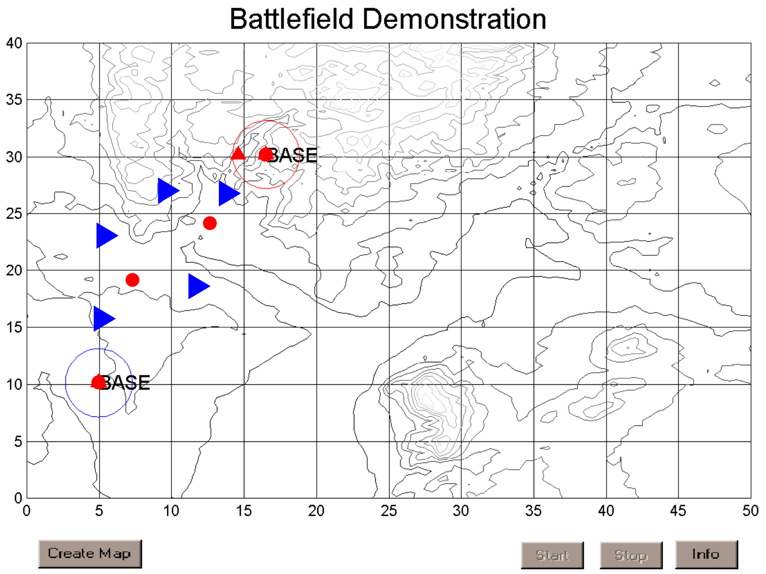

Figure 2: Battlefield

type of weapons they are carrying). In this simulation, all aircraft were carrying radar detection equipment.

Figure 3 shows a snapshot of the occupancy map defined in Section 3 (this snapshot was taken at the same time as the one in Figure 2). The probability of occupancy is encoded in the background color (darker corresponding to a higher probability) and each "x" corresponds to a location where an aircraft "detected" a radar. The occupancy map shown in Figure 3 exhibits large certainty on the positions of the radar since most probabilities in the occupancy map already collapsed to values very close to 0 or to 1. This is true in spite of several false positives that can be seen in the plot.

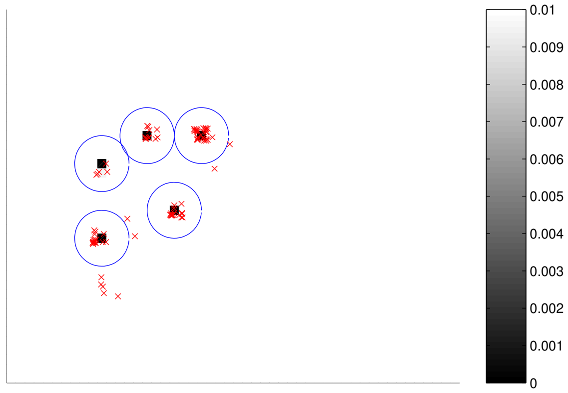

Figure 3: The occupancy map of radars

The log-live map defined in Section 4 is plotted in Figure 4. In this figure, darker regions indicate smaller values for the log-live map and therefore more danger, whereas lighter regions correspond to safer areas. The line superimposed in the log-live map corresponds to a path that is Pareto-optimal for the pair of costs: probability of survival and path length (with distance measured in the Manhattan sense).

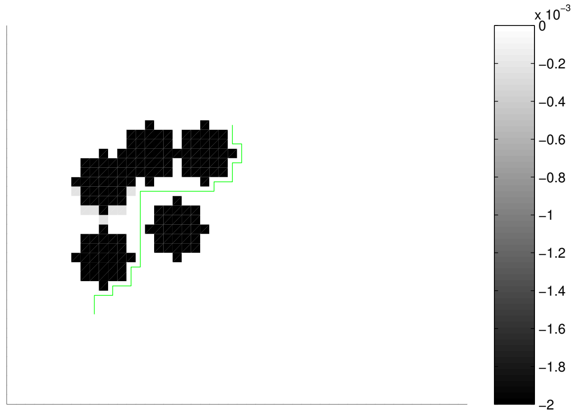

Figure 4: The log-live map of radars

6 Conclusions

In this report, we developed a probabilistic formalism to fuse the information coming from aircraft sensors used to detect electromagnetic radiation originating from illuminator radars. Through the computation of several probabilistic maps, we showed how to assess the risk of missions that consist of flying an aircraft over a region infested by enemy radars and co-located SAM launchers.

We showed that, although a probabilistic map is a function defined in a high-dimensional space, it can be represented compactly by an aggregate measurement function that takes values in a space with the same dimension as the region where the radars are locates. We also showed how to assess the risk of missions in air-operations and even do minimum-risk path planning without ever explicitly computing the high-dimensional probabilistic map and working simply with the aggregate measurement function. The algorithms presented here were implemented in a simulator of military air-operations developed by Honeywell Technology Center and used to guide aircraft in strike missions.

References

[1] S. P. Engelson and D. V. McDermott, “Error correction in mobile robot map learning,” in *Proc. IEEE Int. Conf. on Robot. and Automat.*, pp. 2555–2560, May 1992.

[2] F. Lu and E. Milios, “Globally consistent range scan alignment for environment mapping,” *Autonomous Robots*, vol. 4, pp. 333–349, 1997.

[3] S. Thrun, “Learning metric-topological maps for indoor mobile robot navigation,” *Artificial Intelligence*, vol. 99, pp. 21–71, Feb. 1998.

[4] S. Thrun, W. Burgard, and D. Fox, “A probabilistic approach to concurrent mapping and localization for mobile robots,” *Machine Learning* and *Autonomous Robots* (joint issue), vol. 31, no. 5, pp. 29–53, 1998.

[5] G. Dedeoglu, M. J. Matarić, and G. S. Sukhatme, “Incremental online topological map building with a mobile robot,” in *Proc. of the SPIE*, International Society for Optical Engineering, Sept. 1999.

[6] J. A. Castellanos, J. M. M. Montiel, J. Neira, , and J. D. Tardós, “The SPmap: a probabilistic framework for simultaneous localization and map building.,” *IEEE Trans. Robot. Automat.*, vol. 15, pp. 948–52, Oct. 1999.

[7] B. Yamauchi and P. Langley, "Place recognition in dynamic environments," *J. of Robot. Systems*, vol. 14, pp. 107–120, 1997.