

Experimental Implementation of Robotic Sequential Nuclear Search

R.A. Cortez*, X. Papageorgiou†, H.G. Tanner*, A.V. Klimenko‡, K.N. Borozdin‡ and W.C. Priedhorsky‡

*University of New Mexico, Mechanical Engineering Department, Albuquerque, NM, U.S.A.

†National Technical University of Athens, Mechanical Engineering Department, Athens, Greece

‡Los Alamos National Laboratory, Threat Reduction Directorate, Los Alamos, NM, U.S.A.

Abstract—In this paper we build on, and extend our previous work on automated nuclear search, by taking the principle of a bidirectional interaction between perception and the sensor management to the implementation level. In the algorithm that we implement in hardware, a model of the underlying physical process is used to guide measurement, which then closes the loop by updating the model. Realization of our “model-driven-measurement” concept leads to a working robotic prototype, capable of confirming the presence of weak radiation sources at the specified statistical confidence level, within a bounded, two-dimensional area. The contribution of this paper is in demonstrating that automated nuclear search is feasible, can be performed with commercially available robotic hardware and measuring devices, and that experimental data are consistent with existing theoretical and computational results.

I. INTRODUCTION

The world’s geopolitical situation has changed over the past decade to the degree that requires a new approach to the types of threats which we now face. Modern threats are subtle and ephemeral, can be hidden across large areas. Classical information extraction methods, in which data are randomly collected, and then subsequently filtered and analyzed by human operators in search of particular signatures, are no longer effective. Data collection must be *guided* by querying world models that, unlike the human mind, can have the span and resolution needed for multi-scale problems. Physical models of real-world threats will need continuous updating to follow unpredictable human choices and chaotic physical outcomes.

Our response to this challenge is a new methodology of dynamic interplay between model update and data collection, facilitated by autonomous robots (Figure 1). Our sensing and control architecture (illustrated in Figure 3) introduces the concept of “Model-Driven-Measurement,” which goes beyond the classical data collection and assimilation.

Currently, searching for radiation sources is done manually, usually by operators waving radiation counters in front of them. When the target is a weak radiation source like a speck of uranium, this process is highly unlikely to yield any results at all. The strength of the signal in nuclear search relative to noise falls as R^{-2} as distance R to the source increases. For this reason, existing techniques for autonomous mapping and searching that are based on gradient following [3] will fail: there will be no statistically significant gradient measurement to follow. A new approach that combines random and guided search

is needed, to bring the sensor as close to the source as possible [4]. Using mobile robots to carry the sensors close to the source, and position it accurately for required measurement collection, is a natural choice (Fig. 2).




Fig. 1. The Khepera II mobile robot with the radiation sensor attached. Coper tape is used to shield the detector from ambient noise.




Fig. 2. The miniature radiation sensor that is interfaced with the Khepera II mobile robot, and together navigates the area collecting radiation data.




Fig. 3. Model-driven measurement: we must start asking our questions of models, rather than measurements.

In our recent work [9], we automated nuclear search using a strategy based on the classical sequential testing theory which allows to quickly locate microscopic specks of radioactive material scattered over large area. To speed up the search task our motion controller maintains a maximum scanning speed while the observed count rate is consistent with our model of natural background radiation. When the increase in the number of detected counts is observed, the robot decelerates to a level where the exposure time is sufficient to arrive at a definitive conclusion, with very high confidence, as to whether a source is present there. In a sequential search for a weak radiation source, the space is typically divided in cells and the sensor collects measurements at each cell for different time periods. Once a statistically definitive decision is made regarding the presence or absence of a source, the sensor “jumps” to the next cell.

From the robotics perspective, the problem of locat-

ing such a target is at the intersection of search and exploration. Contrary to common approaches in robotic mapping and exploration, here it is not the topology of the environment that is of interest, but rather the distribution of a physical quantity over an area. The conceptual similarities between the two problems motivate a very brief (and therefore incomplete) review of robotic approaches to robotic search and exploration, where existing approaches differ depending on the *a priori* information available about the environment. If the boundary of the environment is known, a robot can follow a variety of pre-specified paths to cover the entire space [1], [11], [2]. But when the environment boundaries are not known, exploring the area in minimum time is known to be an NP-complete problem, even for the simplest, discretized environments with graph structure. Since an efficient, time-optimal exploration algorithm is unlikely to exist, locally optimal “greedy” approaches and heuristics are being used. One of the most sophisticated approaches is that of [6], where a single robot decides the new search directions by weighting the information gain against the cost of moving along each particular direction.

Among the first problems to be resolved for realizing an automated sequential search using a mobile robot, is how to modify the method to make it applicable in a continuous-space/time framework; a robot can not instantaneously “jump” from one cell to the next, and measurements are collected continuously. We approach this issue by regulating the velocity and acceleration of the moving sensor, to approximate the execution of the discrete algorithm as close as possible. We link the robot motion controller to the statistics of the radiation measurements, and let the latter determine if the robot should accelerate or slow down. In this paper, we develop an experimental platform with which we test our methodology, and present an application where radiation levels are mapped using the Khepera II mobile robot (shown in Figure 1) along one, and two dimensions.

II. PROBLEM STATEMENT

Low-rate counting of radiation from nuclear decay is described by the Poisson statistics, where the probability to register n counts in the detector in t seconds from the source that is known to emit an average of μ counts per second (cts/s) is

$$P(n, t) = \frac{(\mu \cdot t)^n}{n!} e^{-(\mu \cdot t)}. \quad (1)$$

The simplest way to find the radiation source is to search the area uniformly, exposing each location for a fixed duration of time. When no time constraints are present, uniform search is the reasonable strategy to employ. The width of Poisson distribution is defined as $\sigma = \sqrt{\mu \cdot t}$. At known average expected background μ_b , signal μ_s and exposure time (t), the threshold on the number of observed counts can be set that satisfies the required confidence level of the search outcome. However in real life scenarios the time allowed for the search is limited. We need methods that are optimized with respect to search time.

The objective is to use a mobile robot to scan a two dimensional area and confirm the presence of tiny, weak radiation sources, at a very high level of confidence. The robot is supposed to carry a radiation detector, capable of registering gamma rays as they pass through it. Based on the measured radiation, the robot should decide if a source exists at its location.

III. SEQUENTIAL NUCLEAR SEARCH

Classical sequential testing theory [10] suggests the “stopping rules” that allow for rejection of certain sequences of observations at early stages, providing for optimized solution that we seek. Either positive or negative identification can be made based on the likelihood ratio $\kappa_k = P(N_k|S)/P(N_k|B)$, where $P(N_k|S)$ is the probability to observe N_k counts within time period t_k , given that the location contains source with average number of counts per unit time μ_s , $P(N_k|S) = \frac{(t_k \cdot \mu_s)^{N_k}}{N_k!} e^{-t_k \cdot \mu_s}$, and the probability to observe N_k counts assuming location k has only background is $P(N_k|B) = \frac{(t_k \cdot \mu_b)^{N_k}}{N_k!} e^{-t_k \cdot \mu_b}$. The stopping rule is determined from the desired false negative and false alarm rates

$$C = \frac{P_{FN}}{1 - P_{FA}} \quad A = \frac{1 - P_{FN}}{P_{FA}}. \quad (2)$$

For the probability ratio κ_k in location k , the condition $\kappa_k \leq C$ rejects the hypothesis that the source is present, while the condition $\kappa_k \geq A$ confirms the presence of the source. When $C < \kappa_k < A$, longer exposure is required to make a decision. An example of the stopping rules is illustrated in Figure 4, where the straight lines represent the limits of confidence intervals for certain statistical hypotheses: when the radiation counts collected within a certain time period are in the upper region, the presence of a source is verified with a given confidence; if the counts fall in the lower region, then they are most likely due to background radiation; in between, no conclusion can be confidently drawn until more measurements are collected.

IV. AUTOMATION OF SEQUENTIAL METHOD

We develop a (velocity) motion controller for a mobile robot, which regulates sensor exposure time, and implements the sequential search strategy described in the previous section within a two-dimensional environment. The search is performed by scanning an $N \times N$ planar grid, row by row (Figure 5), regulating the time the sensor collects data from each cell by controlling the speed of the robot.

A. Overview

To each region in Figure 4 we associate a robot mode of operation (behavior). As soon as the robot can reach a decision regarding the presence or absence of a source, it keeps moving at a constant, nominal speed. This speed is maintained as long as the counts received at the end of each sampling period remain below a certain threshold. If this threshold is exceeded, and counts begin to accumulate at an increased rate, the robot decelerates to a slower speed, allowing it to collect more measurements at that




Fig. 4. By applying sequential testing theory to our search problem, we calculate thresholds for a positive confirmation or rejection of the source hypothesis (from (2)). The set of gates for positive (top set of thin horizontal lines) and negative (bottom set of thin horizontal lines) identification of 10cts/s source within 1cts/s background compared with the set of gates obtained from numerical calculation (triangles). The bold solid lines are the linear fit to the outer limit of the gates.




Fig. 5. The mobile robot moves the radiation sensor along the rows of a planar grid. The collected measurements facilitate a decision on the presence or absence of a radioactive source, at every cell scanned, and at a given confidence level.

particular region, thus reducing the measurement uncertainty, and enabling a more confident decision regarding the possible existence of a source there.

Once in this mode, the robot slows down to provide a total maximum exposure time of the sensor over the given cell, approximately equal to $T_t = 2.4$ seconds. If during this interval the number of counts collected from this region falls below the negative detection (lower) line, the hypothesis that a source is present is rejected and the robot accelerates back to its nominal speed; if the number of counts registers above the positive detection (upper) line, the presence of the source is verified, the position is marked and the robot accelerates to nominal speed; and if the number of counts is still in the region of uncertainty at the end of the T_t exposure interval, the presence of a source is rejected by default, and the robots accelerates to nominal speed.

The robot dynamics is modeled in discrete time as follows:

$$x[k+1] = x[k] + v[k]\Delta T + \frac{1}{2}a[k]\Delta T^2 \quad (3)$$

$$v[k+1] = v[k] + a[k]\Delta T, \quad (4)$$

where $x[k]$ and $v[k]$ are the position and speed of the robot at the end of the k sampling period, respectively. The sampling time is ΔT and $a[k]$ is the acceleration input at step k computed based on observations made during step $(k-1)$.

Let v_o be the nominal, initial speed of the robot, and let L_S be the length of the sensor (Figure 2). If the robot is moving with constant speed $v[k]$, each point on the

robot's path will be given a uniform exposure time of $T_{exp} = \frac{L_S}{v[k]}$. During each sampling period ΔT , the counts recorded by the radiation detector are added up. At the end of each step, the sum is stored and the detector's buffer is cleared to begin recording the new sum. Let c_i be the sum of the detector counts during step i . The rate of change (increase) of counts in the i^{th} step is estimated at the end of each time step as

$$\left. \frac{\Delta c}{\Delta t} \right|_i = \frac{c_i - c_{i-1}}{\Delta T}. \quad (5)$$

The line expressing the negative identification gate in Figure 4 can be expressed in the form $c = \lambda t + \nu$, where λ and ν are constants that can be calculated from the graph of Figure 4, as the slope and intercept of each straight line, respectively.

To determine whether the sample collected by the detector during step i is actually drawn from a Poisson distribution with a mean significantly larger than the assumed background, we test the condition

$$\left. \frac{\Delta c}{\Delta t} \right|_i \Delta T > \lambda \Delta T + \nu. \quad (6)$$

If (6) is satisfied, (some of) the counts recorded could be coming from a nearby source. To confirm whether this is the case, we increase the exposure time to T_t . By then, the sample size would either have increased enough to cross the positive identification gate (upper line in Figure 4) and confirm the presence of the source with the given confidence, or remain below that line, in which case it is classified as noise by default.

B. Development of the motion controller

Let μ_s denote the average number of counts emitted by the source at each second. Then the estimated time T_p , for which the sensor has been collecting counts from the source before (6) is satisfied, is

$$T_p = \left. \frac{\Delta c}{\Delta t} \right|_i \cdot \frac{\Delta T}{\mu_s}.$$

Figure 4 suggests that within the first 0.2 seconds, it is probable that the detector recorded no counts at all coming from the source. Therefore, we restrict T_p to satisfy $0.2 < T_p < 0.2 + \Delta T$. At the time when (6) is satisfied, a decision is made to increase the total exposure time to T_t , and therefore the detector has to continue collecting measurements from the same location for $T = T_t - T_p$ additional seconds. Define

$$n = \left\lceil \frac{T}{\Delta T} \right\rceil.$$

We quantize T as $T = n\Delta T$, $n \in \mathcal{N}$. If the robot was travelling with speed $v[k]$ for the time period T_p , during which the sensor was possibly exposed to the source, the latter would have moved relative to the tip of the sensor by a distance $S_p = v[k]T_p$. The part of the sensor which has not yet been exposed to the source is $S = L_S - S_p$. Assume that the robot travels S_1 distance in ΔT time while decelerating,

$$S_1 = v[k]\Delta T + \frac{1}{2}a[k]\Delta T^2,$$

and then covers S_2 distance in $(T - \Delta T)$ time, moving at constant speed:

$$S_2 = v[k+1](T - \Delta T).$$

The sum of the distances ($S_1 + S_2$) should be equal to the length of the unexposed part of the sensor, S , from which

$$\begin{aligned} S &= v[k]\Delta T + \frac{1}{2}a[k]\Delta T^2 + v[k+1](T - \Delta T) \\ &= v[k]T + \frac{1}{2}a[k](2T\Delta T - \Delta T^2). \end{aligned}$$

The required acceleration input at step k is therefore

$$a[k] = \frac{2(S - v[k]T)}{2T\Delta T - \Delta T^2},$$

which can be further simplified and written in terms of L_S , $v(k)$ and T_t :

$$a[k] = \frac{2(L_S - v[k]T_t)}{\Delta T(2T - \Delta T)}. \quad (7)$$

Obviously, if $v[k] \leq \frac{L_S}{T_t}$, the robot does not need to decelerate, because the exposure time is already more than T_t seconds. We therefore assume that $v[k] > \frac{L_S}{T_t}$, and consequently for the nominal speed it should hold $v_o > \frac{L_S}{T_t}$. The speed to which the robot decelerates if more measurements are needed is given by (4). To allow only forward motion, we need $v[k+1]$ to be positive, which based on (7), and given $T_p = T_t - T$ implies

$$v[k] < \frac{L_S}{\Delta T/2 + T_p} = \frac{L_S}{\Delta T/2 + \Delta T + 0.2}.$$

The nominal speed should therefore be set so that

$$\frac{L_S}{T_t} < v_o < \frac{L_S}{3/2\Delta T + 0.2}. \quad (8)$$

After decelerating by $a[k]$ for one sampling period, the robot travels with a constant speed, $v[k+1]$ for the next $(n-1)$ sampling periods. At the end of each of the sampling periods $k+1$ through $k+n$, the total number of counts is $c_j = \sum_{i=k}^{k+j} c_i$, $j \leq n$. By then, the suspected source is exposed for a total time of $t_j = (j+1)\Delta T$. Therefore, if

$$c_j > \lambda t_j + \nu_1,$$

where $c = \lambda t + \nu_1$ describes the positive identification line (Figure 4), then the presence of a source is verified. The robot marks the location and accelerates back to $v[t+1] = v_o$. The acceleration in the $k+j+1$ sampling period is given by $a[k+j+1] = \frac{v_o - v[k+j]}{\Delta T}$.

V. EXPERIMENTAL RESULTS

A. Hardware Description

Experimental implementation of the sequential search algorithm was performed using the Khepera II robot with a custom-built turret to interface the CsI radiation sensor (Figure 1). The Khepera II robot features a Motorola 68331, 25MHz processor with 512Kbytes Ram and 512Kbytes Flash memory, running our embedded C code that implements the sequential search motion controller. We use three analog inputs provided by the I/O turret,

and communicate with it through RS232. Data acquisition software is also developed, allowing for real-time 1-D and 2-D radiation mapping during experimentation. The added weight of the sensor, digital board, and power supplies, present a challenge to robot's motors. To reduce friction, a stainless steel ball caster wheel is added at the base of the sensor. Odometry errors tend to cause the robot to deviate from the reference path, if the latter is relatively long (> 75 cm).

The Khepera II is equipped with a CsI radiation sensor. Gamma-rays that pass through the CsI crystal have a probability of depositing some or all of their energy. The deposited energy goes into excitement of the electrons into higher energy levels that decay with characteristic lifetime and emit visible light. The 4 cm long and 1.2 cm in diameter cylindrical CsI crystal was encapsulated into the Al casing with the Hamamatsu S3509 pin photodiode mounted on it to detect light induced in the crystal by passing photons. The sensor was assembled for us by the Alphasp spectra, Inc. The pulse generated by the diode is very weak and needs to be amplified. We used the A250 preamplifier with external FET from Amptek. The pulse is then shaped through a four-stage shaping amplifier based on operational amplifiers. On the output of the amplifier we get almost Gaussian pulse the height of which corresponds to the energy that was deposited by the gamma-ray in the active region of the detector (Fig. 10). The pulse is processed using digital board that is based on low power, high speed, 8bit National Semiconductor ADC08200 and Altera Cyclone 2910 FPGA. The FPGA can be programmed to perform all necessary pulse post-processing. We have developed FPGA code to perform peak finding and pulse counting. We estimated the total power consumption of the electronics to be below 200 mAh at 6V, that allows us to power them for several hours with four rechargeable digital camera batteries. The pin photodiode is in reverse bias and consumes negligible amount of power (nAh at 25V).

B. Tests

The first experimental run is one dimensional, and the robot is driven over a 75 cm straight line. A smooth surface is needed to ensure accurate path trajectory so a 1 m² white board was used. The experiment is performed using a 10 nCi Na-22 source. The source is so weak that we observed an increase in count rate only when the source was directly under our detector. The average count rate from cosmic gamma-rays was approximately 1.8 counts/s, while we measure an average of 6.1 counts/s from the source. The sampling period for the control loop is set at $\Delta T = 0.6$ sec, and the maximum exposure time allowed for each point is set at $T_t = 2.5$ sec. The nominal speed of the robot is set at $v_o = 32$ mm/s. During the deceleration phase, where the presence or absence of a source cannot be verified, the speed of the robot is set to 0 mm/s, since the current experimental setup required the source to be directly under the detector, obscuring the path of the robot. After the maximum exposure time, the radiation source is removed from the path of the robot

and the robot accelerates back to its nominal speed.

Figure 6 shows the counts (background plus source) recorded at each time step. The peak between $t = 14$ s and $t = 16$ s, indicates the presence of a source.

Figure 7 depicts the cumulative counts versus time, in order to highlight the significant jump in the increase rate around $t = 16$ s, which takes the data sample to the region of uncertainty, and triggers the transition to deceleration mode. During the next sample period, the robot decelerates to zero velocity to increase the exposure time, which decreases the uncertainty of the presence of a source. In this sampling period the robot confirms the presence of a source, accelerates back to its nominal speed, and resumes the search.




Fig. 6. The counts (source and background) collected at each time step during the one-dimensional search.




Fig. 7. The total number of counts (source and background) collected as time progresses in the one-dimensional search.




Fig. 8. The distance travelled by the robot as a function of time. Deceleration times correspond to peaks in the number of measured radiation counts.




Fig. 9. The speed of the robot as a function of time.




Fig. 10. Example of the energy spectrum recorded with our CsI sensor.




Fig. 11. A portion of a two dimensional radiation map created by the robot after scanning three rows.

Figure 8 shows how the robot decelerates in response to the increased rate of count collection. Each sudden increase in radiation count measurements in Figure 6 triggers a transition to the deceleration mode. Increased exposure time allows the robot to make a decision regarding the presense or absense of a source, and proceed with its search at nominal speed. This is evident in Figure 9,

where the speed profile of the robot is plotted versus time.

During the second run, the robot is driven over a $75 \text{ cm} \times 12 \text{ cm}$ area creating a 2-D radiation map. Each cell is considered to be a $2 \text{ cm} \times 4 \text{ cm}$ area. All parameters are kept the same as in the first experiment. The 10 nCi Na-22 source is placed in turns at three different locations: $(32, 3) \text{ cm}$, $(32, 7) \text{ cm}$, and $(45, 11) \text{ cm}$. The robot is scanning the 2-D region following a pattern similar to that shown in Figure 5. The resulting map is shown in Figure 11. The mapping indicates that during the search the robot decelerated 13 times, of which 10 were triggered by background, and 3 by the sources. All three sources were identified at there respective locations and at no time was a false alarm recorded.

C. Comparison to Uniform Search Algorithm

In this section we compare the proposed Sequential search algorithm, with the Uniform search algorithm. We use both algorithms to map arbitrary radiation levels over an area, rather than just confirm the presence of a source.

Uniform search consists of scanning the area cell by cell along each row, spending a constant fraction of time at each cell. At the end of the scan, the maximum value for the variance over all cells is compared to the threshold value, and if found larger, the scan is repeated. The search stops when the variance of all the cells of the workspace, is less than the threshold value.

The sequential search algorithm can be thought of as a variation of the uniform search, in which the time spent in each cell is adjusted to allow sufficient integration time for the sensor, and enough measurements to be collected, so that the cell variance drops below the threshold V_a^0 . Once this is achieved, the robot moves to the next cell in the row and repeats the process. The exploration terminates once the variance in all cells drops below the threshold value; in this case however, each cell is visited only once.

The area to be scanned is a $60 \text{ cm} \times 60 \text{ cm}$ surface, decomposed into a 15×15 grid. The initial robot position is at cell $(10, 10)$. We assume a distribution of radiation levels over this area, with a given mean expected count rate λ_a which is well described by the Poisson distribution, represented in upper left of Figure 12. This distribution is unknown to the system, and the goal of the experiment is to reconstruct it up to a certain confidence level V_a , using measurement data.

There is no initial information about the λ_a distribution available to the system, so a reasonable prior is a uniform distribution, both for λ_a^{mn} and for V_a^{mn} , where (m, n) , $\forall m, n \in \{1, \dots, 15\}$ an arbitrary cell, as shown in the lower left and upper right of Figure 12, respectively. The desired confidence level for the constructed map is set at a value corresponding to a variance of $V_a^0 = 0.5$. The robot therefore has to explore the map until the variance value over the whole area is decrease below this threshold.

The initialization of the experimental test is showed in Figure 13. The picture shows the robot in its initial configurations on the grid.

Figures 14 through 15 present experimental results for the uniform search, and figures 16 through 17 present




Fig. 12. Initial configuration of the area. Upper left: real distribution of λ_a . Upper right: uniform prior information of the distribution of λ_a . Lower left: Uniform prior information of variance V_a .




Fig. 14. Uniform search algorithm. Upper left: real distribution of λ_a . Upper right: updated values of the distribution of λ_a , as the robot is moving and the sensor is collecting radiation data. Lower left: updated values of variance, during the experiment.




Fig. 16. Sequential search algorithm. Upper left: real distribution of λ_a . Upper right: updated values of the distribution of λ_a , as the robot is moving and the sensor is collecting radiation data. Lower left: updated values of variance, during the experiment.

experimental results for the sequential search, for the same mapping scenario. The two search algorithms are compared in terms of completion time. Results indicate that the Sequential search is faster than Uniform. In a typical run, the Sequential search requires approximately 600 time steps (seconds), and the Uniform search roughly 3,000 time steps, to complete the map at the same level of confidence. Another interesting difference between the search methods is the final maximum error (in absolute value terms) between the true (originally assumed, but unknown to the system) distribution of λ_a , and the distribution estimated by means of measurement: typical values for this error are recorded 1.38 for the uniform algorithm,




Fig. 13. Experimental set-up. The robot is in its initial cell on the grid.




Fig. 15. Uniform search algorithm, in which the robot is moving cell by cell along each row, and it is staying there for just one time instant. Variance's values at an intermediate step. More than one scan of the total area is needed to achieve the certainty level.




Fig. 17. Sequential search algorithm, in which the robot is moving cell by cell along each row, and it is staying there until the variance is decreased below the threshold. Variance's values at an intermediate step, with further explored area.

and 2.11 for the sequential search. These numbers could indicate that there could be a trade-off between search time and map accuracy, however proper statistical analysis over a large sample set of test runs is needed for a definitive conclusion on this issue.

VI. CONCLUSIONS

We present an experimental implementation of a sequential nuclear search strategy in two dimensions, using a commercially available desktop mobile robot, fitted with a CsI radiation sensor. Our approach to implementing the sequential nuclear search, is to regulate the speed of the robot (and therefore control the exposure time of the sensor) over each segment of the search area, in real-time, according to the radiation counts collected by the sensor. Further steps include extensions to three-dimensional searches, the use of prior knowledge about the search area, as well as the introduction of Bayesian statistics for the on-line update of the world model. The latter, will enable us to map arbitrary radiation levels over an area, rather than just confirm the presence of a source, using navigation strategies inspired by [8].

ACKNOWLEDGMENT

The work of the three first authors is supported in part by the Los Alamos National Laboratory Award No. STB-UC:06-36, and in part by the DoE URPR grant DE-FG52-04NA25590.

REFERENCES

- [1] H. Choset and P. Pignon, "Coverage path planning: The Boustrophedon Cellular Decomposition," *Proceedings of the International Conference on Field and Service Robotics*, Canberra Australia, 1997.
- [2] H.H. González-Baños and J.-C. Latombe, "Navigation Strategies for Exploring Indoor Environments," *Journal of Robotic Systems*, 21(10-11): 829-848, Oct-Nov 2002.
- [3] P. Ogren, E. Fiorelli, and N.E. Leonard, "Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, 49(8):1292-1302, 2004.
- [4] W.C. Priedhorsky, "Physical and economic tradeoffs between remote sensing and deployed sensor networks", *Journal of Nuclear Materials Management*, submitted (2005).
- [5] A.V. Klimenko et al., "Efficient Strategies for Low-Statistics Nuclear Search", *IEEE Transactions on Nuclear Science*, 53(3): 1435-1442, June 2006.
- [6] S.J. Moorehead, "Autonomous Surface Exploration for Mobile Robots," *Technical Report CMU-RI-TR-01-30*, The Robotics Institute, Carnegie Mellon University, August 2001.
- [7] E. Rimon and D. Koditschek, "Exact robot navigation using artificial potential functions," *IEEE Transactions on Robotics and Automation*, vol 8, pp 501-518, Oct. 1992.
- [8] H. G. Tanner and A. Kumar, "Formation Stabilization of Multiple Agents Using Decentralized Navigation Functions," in S. Thrun, G. Sukhatme, S. Schaal and O. Brock (eds), *Robotics: Science and Systems I*, MIT Press, 2005, pp 49-56.
- [9] A. Kumar, H.G. Tanner, A.V. Klimenko, K. Borozdin, and W.C. Priedhorsky, "Automated Sequential Search for Weak Radiation Sources," 14th IEEE Mediterranean Conference on Control and Automation, Ancona, Italy.
- [10] A. Wald, "Sequential tests of statistical hypotheses," *Annals of Mathematical Statistics*, vol. 16, pp. 117-186, 1945.
- [11] A. Zelinsky, R.A. Jarvis, J.C. Byrne and S. Yuta, "Planning Paths of Complete Coverage of an Unstructured Environment by a Mobile Robot," *Proceedings of the International Conference on Advanced Robotics*, 1993, pp. 533-538.