Automated Sequential Search for Weak Radiation Sources

A. Kumar, H.G. Tanner,
A.V. Klimenko, K. Borozdin, and W.C. Priedhorsky

***Abstract*— We demonstrate the principle of a bidirectional interaction between the perception model that describes the environment, and the sensor data collection that shapes the model, through the first automated scheme for sequential nuclear search. The objective is to detect and localize a weak radioactive point source, with consideration to the time spent for the search. Numerical simulations verify that desired false positive and false negative rates can be achieved using the control algorithm developed for deployment on prototype system. The simulation results are also consistent with previous theoretical, discrete simulation of the sequential search strategy that our continuous search control algorithm is based on.**

I. Introduction

Changes in the geopolitical atmosphere around the world over the past decade has lead to qualitative changes in the types of threats with which we are now faced. Modern threats are subtle and ephemeral, and can be hidden across large areas. These threats can no longer be easily characterized using classical information extraction methods where the data are usually randomly collected, filtered and then analyzed by human operators in search for relevant signatures. We must guide data collection by querying 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. The new approach of model-driven measurement offers the dynamic interplay between model update and data collection (Fig. 3) that goes beyond the classical data assimilation.

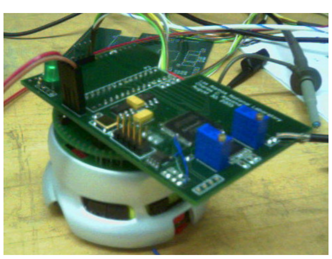

Fig. 1. The Khepera II mobile robot with a custom built turret to interface the radiation sensor.

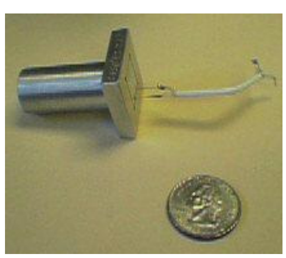

Fig. 2. The miniature radiation sensor that is to be interfaced with the mobile robot.

Amit Kumar and Herbert G. Tanner are with the Department of Mechanical Engineering, MSC01 1150, the University of New Mexico, Albuquerque NM 87131, USA. Their work was partially supported by the Los Alamos National Laboratory, under Subcontract No. 17547–001–05.

Alexei Klimenko, Konstantin Borozdin, and William Priedhorsky are with the Los Alamos National Laboratory, Nonproliferation and International Security Division, Los Alamos NM 87545, USA.

In robotic search and exploration, 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], [10], [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.

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).

We automate the nuclear search using the strategy based on the classical sequential testing theory which allows to

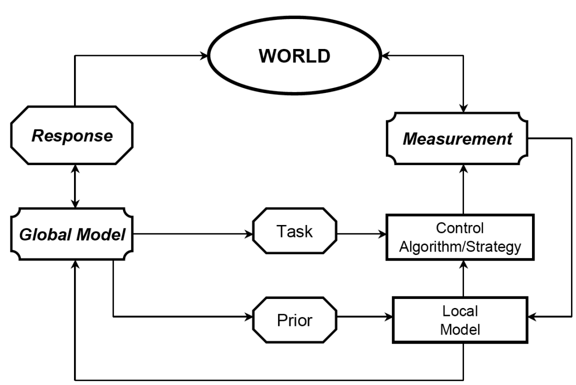

Fig. 3. Model-driven measurement: we must start asking our questions of models, rather than measurements.

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 emissions is observed the robot decelerates to a level where the exposure time is sufficient to produce a definitive answer at a very high confidence level (10^{-7} or better) as to whether a source is present there.

In a sequential search for a weak radiation source, the space is divided in cells and the "sensor" collects measurements at each for different time periods. These time periods are determined by the need to reach a statistically definitive conclusion on whether a radioactive source is present in the particular cell. Once a decision is made, the sensor "jumps" to the next cell.

One of the first problems that has to be resolved for realizing an automated sequential search by means of a mobile robot, is how to modify the method to make it applicable in a continuous-space/time framework, without decomposing into cells; a robot could never instantenously "jump" from one cell to the next, and measurements are collected continuously. We approached 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 linked the robot motion controller to the statistics of the radiation measurements, and let the latter determine if the robot should accelerate or slow down. This paper presents our motion control strategy and a set of extensive simulations that verify consistency between the search task parameters and the search outcome. We are moving in the direction of experimentally testing our automated search strategy using a retrofited Khepera II mobile robot shown in Figure 1.

In the current paper the search is one-dimensional, but we envision it to be facilitated by a topographic map of the known workspace in the form of a navigation function [7]. We link the "Measurement" with the "Physical Model" in Figure 3 through an automated search strategy that aims at formally refining our local interpretation of the world to a certain level of statistical confidence, by driving the measurement. This will be our starting point in developing new, fully 2-D and 3-D nuclear search algorithms within our *model-driven measurement* paradigm, in which robots navigate in partially known cluttered environments using navigation functions.

The paper is organized as follows: in section II we formulate the nuclear search problem and review the sequential search strategy. In section III we describe our motion control algorithm that will enable a mobile robot to perform nuclear search as suggested by sequential nuclear search strategy. Section IV presents our simulation results where we estimate the probabilities of false alarms and false negatives that our method achieves. Finally, section V gives an overview of the results of this work and highlights our current and future research directions.

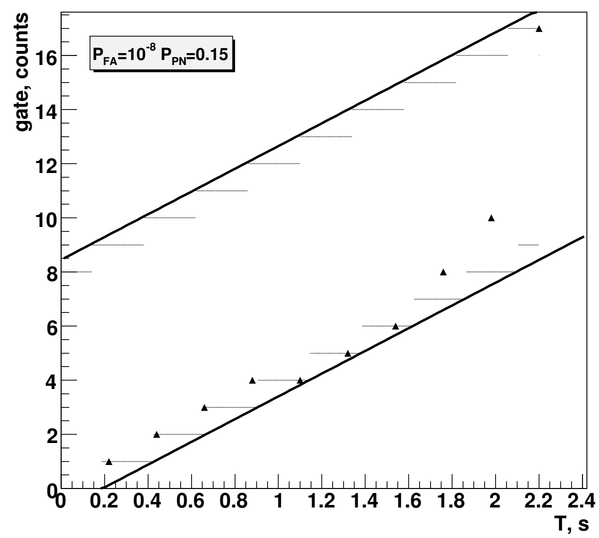

Fig. 4. By applying sequential testing theory to our search problem, we calculate thresholds for a positive confirmation or rejection of thesource hypothesis (from eqn. 2). The set of gates for positive (top set of thin horizontal lines) and negative (bottom set of thin horizontal lines) identification 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.

II. Nuclear Search Strategy

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:

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

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.

Classical sequential testing theory [9] suggests the "stopping rules" that allow for rejection of certain sequences of observations at early stages. 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:

$$C = \frac{P_{FN}}{1 - P_{FA}} \qquad A = \frac{1 - P_{FN}}{P_{FA}} \tag{2}$$

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.

III. Robotic Implementation of the Sequential Search

Presently, the nuclear search and radiological mapping is performed through direct human control and relies strongly on human intuition. The use of mobile robotic agents instead of human operators, allows for faster, more reliable search with statistically definitive outcome. We construct a motion controller for the *sequential search strategy* that is compatible with modern swarm navigation and cooperative control techniques [8]. The robot controls the sensor location in the continuous mode, and regulates the exposure time of each point of the search area by changing the cruising speed. The nuclear search problem we consider in this paper is one dimensional for simplicity, but higher dimensional generalizations are straightforward.

The strategy is implemented on the robot as follows: as long as the counts received at the end of the sampling period remain below a certain threshold, a constant speed is maintained. A sudden increase in the rate of change could bring the system in the region of uncertainty, between the two lines in Figure 4. The presense or absense of a source cannot be verified there, and the robot decelerates, increasing the exposure time of the region that gave the last counts T_t, which is approximately equal to 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 to nominal speed.
- 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:

$$x[k+1] = x[k] + v[k]\Delta T + \frac{1}{2}a[k]\Delta T^2, \tag{3}$$

$$v[k+1] = v[k] + a[k]\Delta T. \tag{4}$$

In the above, $x[k]$ 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 the k sampling period, calculated from observations made in the $(k-1)$ time step.

Suppose that the robot initially moves with a predetermined speed v_o. We assume that the length of the sensor is L_S [mm] (Figure 2). If the robot is moving with speed $v[k]$ [mm/s], each point on the robot's path will be given a uniform exposure time of

$$T_{exp} = \frac{L_S}{v[k]}.$$

The radiation detector collects photon hits (counts) during the full sampling period. These counts are added up for the entire period of ΔT and are read by the host computer at the end the sampling period, with negligible delay. After reading the counts the detector buffer is cleared and the sensor starts accumulating the new counts during the next sampling period.

Denote the sum of the counts in the i sampling period, cts_i. The rate of change of counts in the i sampling period, is estimated by

$$\left[\frac{\Delta c}{\Delta t}\right]_i = \frac{cts_i - cts_{i-1}}{\Delta T}. \tag{5}$$

To distinguish between the background noise and the source, $[\frac{\Delta c}{\Delta t}]_i$ is calculated at the end of each sampling period. The expression for the line used for negative identification is

$$cts = \lambda t + \nu,$$

where λ and ν are positive constants. If the following condition is satisfied:

$$\left[\frac{\Delta c}{\Delta t}\right]_i \Delta T > \lambda \Delta T + \nu, \tag{6}$$

then the count sample collected may have been drawn from a (Poisson) distribution with a mean significantly different than that of the expected background. This would indicate that the counts emitted are coming from a source. To confirm the latter, the exposure time is increased to at least $T_t = 2.4$ s, a level that the negative detection gate (6) in the graph of Figure 4} indicates as appropriate in order for the count increase to be classified as noise.

Let μ_s denote the average number of counts emitted by the source each second. Then the estimated time T_p for which the sensor was already collecting counts from the source, before the sudden increase was detected, is

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

By looking at the Figure 4, we know that within the first 0.2 seconds, the source might have emitted nothing. So the time T_p is bounded between $0.2 < T_p < 0.2 + \Delta T$. The additional exposure time needed according to (6) will then be,

$$T = T_t - T_p.$$

$$n = \text{ceil}\left[\frac{T}{\Delta T}\right]$$

We redefine T as, $T = n\Delta T, \quad \forall n \in \mathcal{N}$. If the robot was travelling with speed $v[k]$ for the time period T_p, during which the sensor was exposed to the suspected 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 been exposed to the source yet is then $S = L_S - S_p$. In our motion control strategy we assume that the robot travels S_1 distance in ΔT time decelerating, and then covers S_2 distance in $(T - \Delta T)$ time moving at constant speed. Then,

$$\begin{aligned} S_1 &= v[k]\Delta T + \frac{1}{2}a[k]\Delta T^2 \\ v[k+1] &= v[k] + a[k]\Delta T \\ S_2 &= v[k+1](T - \Delta T). \end{aligned}$$

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

$$\begin{aligned} S &= S_1 + S_2 \\ &= v[k]\Delta T + \frac{1}{2}a[k]\Delta T^2 + v[k+1](T - \Delta T) \\ &= v[k]T + \frac{1}{2}a[k](2T\Delta T - \Delta T^2). \end{aligned}$$

The above equation gives us the required acceleration input at step k:

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

The above acceleration expression can be further simplified and written in terms of L_S, $v(k)$ and T_t. Substituting, the deceleration for the k sampling period is found to be:

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

In order for the robot to be able to decelerate, we need $v[k] > \frac{L_S}{T_t}$. Thus, we must set the initial speed so that $v_o > \frac{L_S}{T_t}$. If needed, the robot will decelerate to a speed $v[k+1]$, given by (4). We also impose an additional condition, in order to exclude the case where the robot has to move backwards: this suggests that $v[k+1]$ is always positive, namely,

$$v[k] + a[k]\Delta T > 0.$$

Substituting $a[k]$ using (7),

$$\begin{aligned} v[k] + \frac{2(L_S - v[k]T_t)}{\Delta T(2T - \Delta T)}\Delta T &> 0 \\ \Rightarrow -v(k)\Delta T - 2v(k)(T_t - T) + 2L_S &> 0, \end{aligned} \tag{8}$$

Since, $T_p = T_t - T$, $-v(k)\Delta T - 2v(k)T_p + 2L_S > 0$. This above inequality gives us another condition on $v(k)$,

$$v(k) < \frac{L_S}{\Delta T/2 + T_p} = \frac{L_S}{\Delta T/2 + \Delta T + 0.2}$$

Therefore, the initial set speed should have a lower bound and a upper bound:

$$\frac{L_S}{T_t} < v_o < \frac{L_S}{3/2\Delta T + 0.2} \tag{9}$$

After decelerating with $a[k]$ 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 from $k+1$ to $k+n$, the total number of counts is added:

$$cts_j = \sum_{i=k}^{k+j} cts_i. \quad j \leq n$$

The suspected source would have been exposed for a total time of $t_j = (j+1)\Delta T$. Therefore, if

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

where $cts = \lambda t + \nu_1$ describes the line used positive identification gate (Figure 4), then the collected counts sample has a mean that is statistically significant from μ_b, and the presence of a source is verified. After marking the location, the robot accelerates to $v[t+1] = v_o$. The acceleration in the $k+j+1$ sampling period is given by

$$a[k+j+1] = \frac{v_o - v[k+j]}{\Delta T}$$

IV. Simulation Results

This section presents simulation results that demonstrate the efficacy of our nuclear search method. The robotic implementation has yielded more conservative probabilities of false alarms (FA) and false negatives (FN) compared to the ones suggested by the sequential search strategy, as long as the cruising speed remained below the maximum allowed: Simulation results indicated that the probability of missing the source (FN) was less than 10%, and was increasing as a function of the robot speed. In a set of approximately 100,000 simulation runs, no false alarms were recorded.

In the simulation example, we drive the robot over a straight line. A radiation source with a mean of $\mu_s = 11$ [counts/s] is positioned at a distance of is placed at 880 [mm] from the starting point. The sampling period of the control loop is set at $\Delta T = 0.5$ s, and the total exposure time allowed for each point (after which we assume the absense of source if no statistical hypothesis regarding the absense or presense of a source can be verified), is set at $T_t = 2.5$ s. The robot begins moving at a speed of $v_o = 33$ [mm/s], and continues for 45 seconds. Background radiation is estimated at $\mu_b = 1$ [counts/s]. Radiation statistics are simulated using MATLAB's `poissrnd` function.

Figure 5 gives the counts (background plus source) recorded at each time step, over the total period of motion. The peaks observed between the 58th and 61st sampling period suggest the presence of a source. Figure 6 depicts the cummulative counts versus time. There is a clear change in the rate of increase around $t = 30$ s. Figure 7 shows the trajectory of the robot, in response to the measurements. After t = 29 s, the increased number of counts recorded,

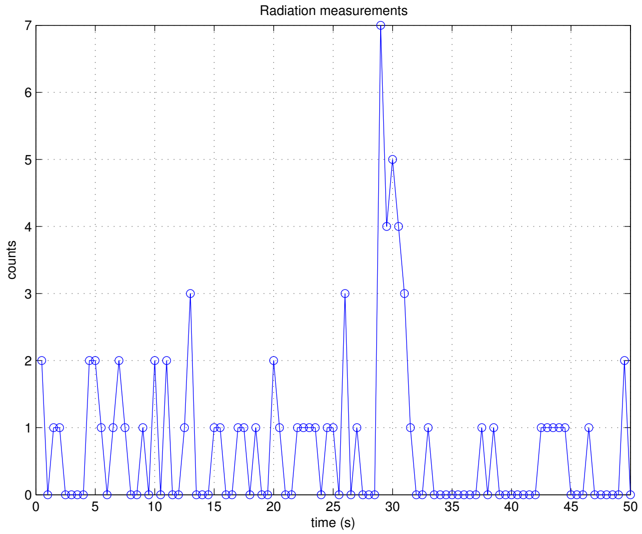

Fig. 5. Time vs Total Counts obtained at the end of each sampling period.

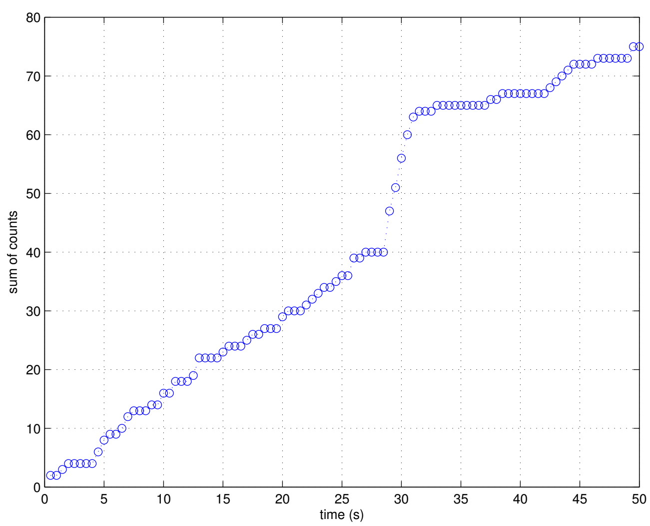

Fig. 6. Time vs Sum of total counts at the end of each sampling period.

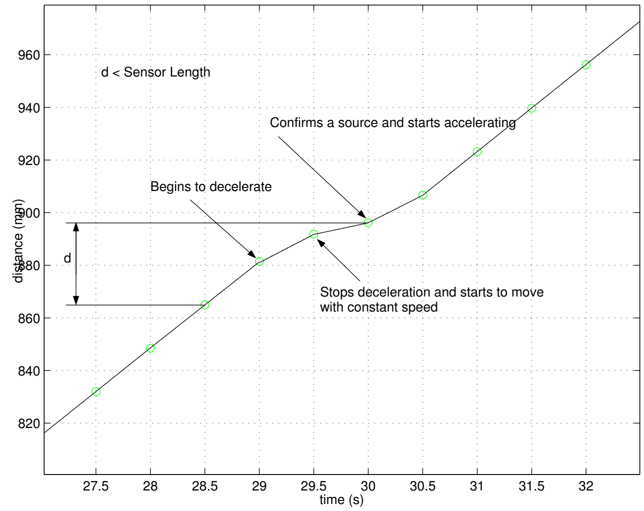

Fig. 7. Trajectory of the robot while passing the source location.

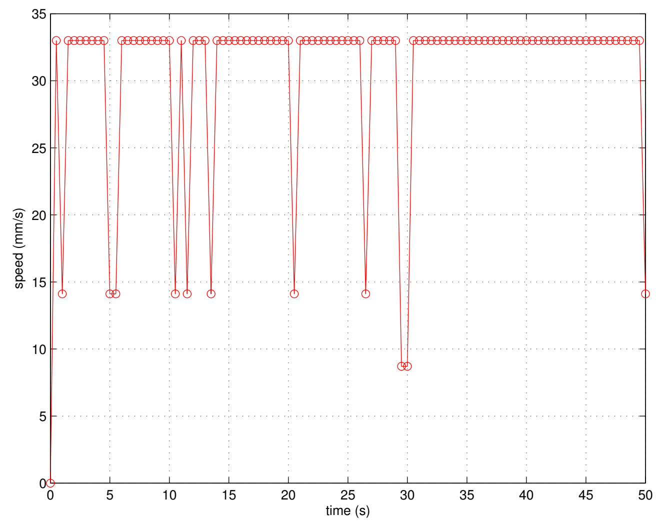

Fig. 8. Robot speed over the simulation time.

registers the data sample within the region of statistical uncertainty. In the next sampling period, the robot decelerates in order to increase exposure time, obtain a sample with a larger time horizon and decrease uncertainty. In the following sampling period, the robot moves with constant speed, and at the end of the 30th second, it confirms the presence of a source. Having decreased the uncertainty of the measurement in this region, and verified one of the two hypotheses, it accelerates again to the initial speed, and resumes the search. Figure 8, indicates that during the whole search, the robot decelerated 9 times, out of which 8 were triggered by background radiation. Each time the robot decelerated, verified one of the two hypotheses and promptly accelerated back to the initial speed to resume the search. The source were correctly identified at its location and no point gave a false alarm.

In a set of 10,000 simulation runs, we estimated the probability of false negatives (FN), correct detection (CD) and false positives (FP). In this set of runs, we varied the initial speed v_o from $v_o = 30$ [mm/s] to $v_o = 80$ [mm/s], we moved the robot for 600 sampling periods, and we measured the percentage of false negatives, using different sampling periods ($\Delta T = 0.25$ s, and $\Delta T = 0.2$ s). In Figure 9 suggests, the percentage of FNs (missing the source) increases with the initial speed v_o. Close to the maximum speed the robot is allowed to travel, we verify a percentage of FNs close to 10%, which is what the theory of sequential search predicts. No false positives (false alarms) were recorded.

Another observation is that the percentage of FNs decreases as the sampling period increases. We believe that this is due to the way we have simulated the radiation measurements and we plan to investigate this effect more

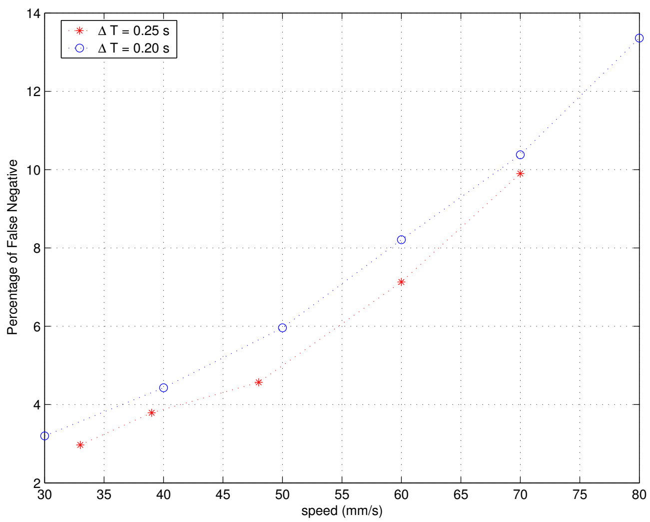

Fig. 9. Speed vs Percentage of False Negative.

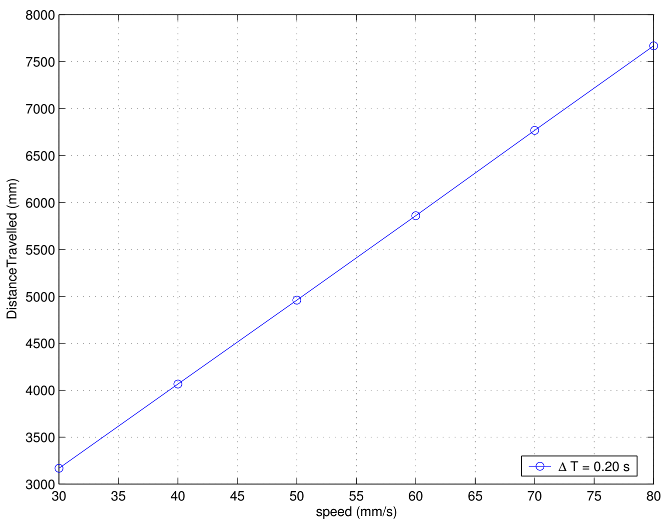

Fig. 10. Average distance travelled (and searched) over a sample of 10^4 runs as a function of the cruising speed of the robot.

thoroughly in experimental implementations, currently in progress.

We use Figure 10 to evaluate the efficiency of our search method. We see a monotonic increase in the area that is searched as the cruising speed increases.

V. Conclusions

We have developed a robot motion control algorithm that implements a sequential nuclear search strategy on a mobile robot carrying a radiation sensor. The algorithm uses stopping rules, provided by the strategy, to decelerate and accelerate the robot based on the "Task" that has to be performed, the "Local Model" of the robot's workspace and the measurements collected. In our 1-D simulation, we consider only the interaction between the "Local Model" and the "Measurement," realized by the motion controller and the search strategy . The "Task," which is to *locate the radioactive source, having a probability of missing it (false negative) equal to* $P_{FN} = 0.1$, *and a probability of a false alarm equal to* $P_{FA} = 10^{-7}$, is relayed from the "Global Model" to our "Control Algorithm/Strategy". The "Local Model" of the environment in which the "Control Algorithm/Strategy" is executed, is based on the "Prior" information that we have. The latter is comprised of the mean emission rates from the source ($\mu_s = 11$ cts/s) and the background ($\mu_b = 1$ cts/s). These emission rates are kept constant throughout the search process. Our simulations indicate that the maximum speed of 80 mm/s, we were able to search 7.7 m, having 13% rate of false negatives, and $< 10^{-2}\%$ rate of false alarms.

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*, 2005. (submitted)

[6] S.J. Moorehead, "Autonomous Surface Exploration for Mobile Robots," *Technical Report CMU-RI-TR-01-30*, The Robotics Institute, Carnegie Melon 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," *Robotics: Science and Systems*, Boston, June 2005, (in print)

[9] A. Wald, "Sequential tests of statistical hypotheses," *Annals of Mathematical Statistics*, vol. 16, pp. 117-186, 1945.

[10] 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.