Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Stochastic Strategies for Autonomous Robotic Surveillance

Jeremy Grace*
gracej@bu.edu

John Baillieul*
johnb@bu.edu

Dept. of Aerospace and Mechanical Engineering,
Boston University,
Boston, MA 02215

Abstract— **In this paper, we consider a team of autonomous mobile robotic agents engaged in a surveillance mission. It is desirable not to have the agents move in a predictable fashion so no intruder or invader can plan their movements to avoid the surveillance agents. This paper investigates the use of stochastic rules to guide the motions of the agents throughout their surveillance missions. The research emphasizes methods which minimize centralized computation and communication requirements by focusing on local rules for each agent. We define a formal mathematical approach to analyzing the surveillance problem. We take a general environment and impose some mathematical structure on it. We then define a measure of non-uniformity by which we can compare the surveillance coverage of different systems as well as employ a technique to measure the parametric randomness of each strategy. Drawing on the existing literature on stochastic techniques for searching a graph by a single agent, we study the problem of finding optimal rules for searching a graph using a team of agents. In certain cases, search strategies may be evaluated in terms of explicit closed-form expressions for steady-state distributions. In general, however, strategies must be evaluated by Monte-Carlo methods. We present a decomposition technique by means of which any of the graphs types under discussion may be uniquely decomposed into a collection of complete subgraphs. We developed a surveillance strategy for handling general graphs through a hybrid solution to our less complex decompositions.**

I. INTRODUCTION

Research into multiple agent autonomous robotic systems has reached a point where the systems can be expected to perform increasing complex missions. These missions have the potential to range from nuclear reactor cleaning to planetary exploration; from clearing minefields to flying bombing expeditions. Many researchers have studied the application of groups of mobile agents, acting cooperatively, to perform search and surveillance [8-14]. In this paper, we consider the case where the motion of the surveillance agents is unpredictable and investigate the balance between the predictability of the motion and the efficiency to which the surveillance is accomplished.

We consider a group of surveillance agents that must provide adequate coverage of an environment in order to monitor the condition of the space and detect any intruders. However, the movement of the sentries should not be so well orchestrated that an intruder can plan a path that avoids being seen by the sentries. In addition, we look to reduce the overhead associated with centralized computation and heavy communication requirements. This leads us to consider a team of agents with local, stochastic rules. We investigate the issue of surveillance coverage as well as the ability to capture an intruder for a few stochastic strategies. In particular, can we define local probabilistic rules of motion that yield a specified expected coverage of an area under surveillance? How quickly do we approach the specified surveillance coverage?

There is an extensive background of published work which is related to the stochastic surveillance problem. Chvatal [1] introduced the so-called art gallery problem by considering the number of guards required to cover a n-sided polygon. O'Rourke published an excellent monograph on all the associated art gallery problems and algorithms. Chin and Ntafos [3,4] expanded the realm of possibilities by considering mobile guards. They introduced the concepts of optimal watchman routes. Carlsson et al [5] expanded those ideas into guard covers and multiple watchman routes for a certain set of polygons.

Hespanha [13] et al investigate the use of multiple agents with probabilistic behavior in pursuit evasion games. Kim et al [14] continue the idea in a 3D environment with unmanned ground vehicles and aerial vehicles. The space is divided into cells and the probabilistic behavior comes from the uncertainty of knowledge regarding certain cells. The vehicles behavior is guided by optimizing the increase in knowledge.

We formulate the surveillance problem abstractly as a random walk on a hypergraph, where each node on a hypergraph corresponds to a section of the environment (e.g. a corridor segment of the building) and where each edge of the graph is labelled with the probability of transition between the nodes. We shall be interested in the rate at which the Markov chain converges to its steady state distribution, as discussed by Rosenthal [17]. In [19], Boyd, Diaconis and Xiao, consider the problem of designing a symmetric Markov chain on a graph so that the associated probability distribution on the nodes approaches the uniform steady state as rapidly as possible. They approach the optimal solution to the fastest mixing problem using semi-definite programming techniques.

Motivated by this work, we consider the problem of parallel Markov chains and fastest mixing when a team of agents moves among states on a graph. Whereas [19] considers the problem of designing a fixed set of transition probabilities so

Both authors gratefully acknowledge support from the Army Research Office under the ODDR&E MURI01 Program Grant No. DAAD19-01-1-046 to Boston University, the Center for Networked Communicating Control Systems, NSF Grant No. DGE-0221680 through the Center for Computational Science and NSF ITR Program Grant No. DMI-0330171 through the Center for Information and Systems Engineering.

0-7803-9568-9/05/$20.00 ©2005 IEEE

that a Markov process on a graph tends to uniform steady state as fast as possible, we shall be interested in parametrically varying the local transition probabilities over time so as to provide accelerated convergence to steady state.

The paper is organized as follows. In section 2, we will describe the details of the problem and the formal mathematical structure. In section 3, we discuss some measures for analysis. In section 4, we investigate simple, 1D problems. In section 5, we analyze complete graphs and their convergence rates. In section 6, we describe our hybrid strategy for general environments. In section 7, we conclude by highlighting the important aspects of the research.

II. MODEL

The motivation for this research is the problem of deploying a team of robotic agents to provide surveillance of an area or environment of interest. In its most general form, this problem is about parallel random walks on a graph. In many surveillance scenarios, it is desirable to have the agents move in a non-deterministic way. This can prevent an intruder from planning a path that successfully will avoid the sentries for all time.

We first need to formalize the problem and provide some mathematical background. The methods we propose will apply to surveillance domains in one, two and three dimensions. Consider a typical art gallery floor (See Figure 1(a)). To any interior structure such as this, we shall associate a hypergraph to support a mathematical analysis of the problem. This is a special case of a general class of search environments which are modelled as sets of line segments (some or all of which intersect) in a bounded domain in the plane. The lines may be thought of as corridors, with their intersection points corresponding to corridor intersections. There is an intermediate association that can be made with a simple graph in which the vertices correspond to corridor endpoints and intersections, and the edges correspond to corridor segments between pairs of vertices (See Figure 1(b)).

A more general structure is shown in Figure 1(c). To pursue our analysis, it is useful to take the dual of this graph, which is a hypergraph in which the vertices correspond to corridor segments, and whose edges are sets of corridor segments corresponding to possible transitions. (I.e. if an agent is in a corridor segment x_i at time k, it can move, at time $k+1$ into any other corridor segment in an edge containing x_i.) Recall from [20]

Definition: Let $X = x_1, x_2, \ldots, x_n$ be a finite set, and let $E = (E_i | i \in I)$ be a family of subsets of X. The family E is said to be a hypergraph on X if

1) $E_i \neq \phi$

2) $\bigcup_{i \in I} E_i = X$.

The couple $H = (X, E)$ is called a hypergraph [20].

Since the states are directional corridor segments, the transition between states can be viewed as a set of turning probabilities (i.e. the probability of turning from one corridor segment to another). Given this representation, we can consider a Markov chain with the following transition probability matrix:

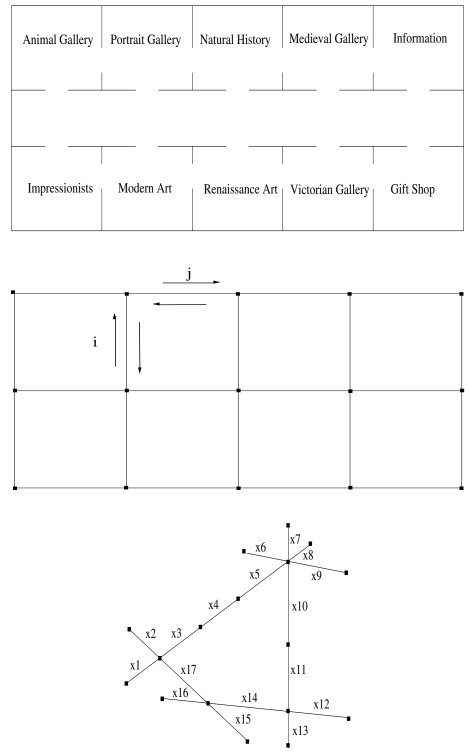

Fig. 1. (a) Typical art gallery floor with (b) the graph representation of the environment (c) a general structure.

$$P = [P_{ij}] \tag{1}$$

$$\sum_i P_{ij} = 1 \tag{2}$$

$$0 \leq P_{ij} \leq 1, \tag{3}$$

where P_{ij}, a stochastic matrix, represents the probability of the agent going to state i from state j. Note that the states correspond to corridor segments. Constraints (2) and (3) must hold since the sum of the probabilities must be one and all the probabilities must be nonnegative.

We parameterize the problem by defining turning probabilities (surveillance strategies) for each agent. The following questions arise: (1) What types of surveillance coverage can the stochastic agents provide (i.e. what is the steady state distribution for the system of agents with specified turning probabilities)? (2) At what rate does the system of agents converge to this invariant distribution? (3) What are the appropriate measures for comparing different surveillance strategies? (4) How can we capture randomness in the motions of the agents? (5) Can we get an understanding of the trade-off between randomness and speed of convergence?

III. Methods of Analysis

A. Steady State Distributions

Our first interest is to gain an understanding of the type of surveillance coverage (invariant distribution) that the stochastic agents can provide. In addition, we seek to take advantage of using multiple agents with different stochastic rules to provide the specified surveillance coverage more quickly.

The probability distribution at each time *k+1* is determined according to

$$\vec{p}_i^{(k+1)} = P_i \vec{p}_i^{(k)},$$

where $\vec{p}_i^{(k)}$ is the probability distribution for agent i at time k and P_i is the transition probability matrix for agent i.

There is a unique invariant distribution for an irreducible, aperiodic Markov chain, which is the eigenvector associated with the eigenvalue 1 [17]. This invariant distribution represents the steady state probability of the agent being at any state. We can sort the eigenvalues of the Markov chain by magnitude:

$$1 = |\lambda_1|(P) \geq |\lambda_2|(P) \geq \cdots \geq |\lambda_n|(P).$$

The mixing rate for the Markov Chain is given by:

$$\mu(P) = |\lambda_2(P)|.$$

where $|\lambda_2(P)|$ is the eigenvalue which is second largest in magnitude. The smaller the mixing rate, the faster the Markov chain converges to its steady state distribution.

We can explicitly determine the expected composite distribution of a by:

$$\vec{p} = \frac{\Sigma_{i=1}^{a} \vec{p}_i^k}{a}. \tag{4}$$

Here, we assume the agents move randomly and independently. The case of interdependence will be treated elsewhere.

IV. Problems in 1-Dimension

A. Single Agent Investigations

We can gain an understanding of how tuning problem parameters can affect the invariant distribution. Consider the case of a probabilistic agent walking on an n-node, one-dimensional lattice, taking steps to the right with probability ρ and steps to the left with probability $1-\rho$. Assume reflecting barriers. One can write down explicit equations for the steady-state probability, and these may in principle always be solved explicitly. For a given initial probability distribution $(p_1(0), \ldots, p_n(0))$, the probability distribution at each time evolves according to

$$\vec{p}(k+1) = P\vec{p}(k),$$

where

$$P = \begin{pmatrix} 0 & 1-\rho & 0 & \ldots & 0 \\ 1 & 0 & 1-\rho & \ldots & 0 \\ 0 & \rho & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \ldots & \rho & 0 \end{pmatrix}.$$

The steady state (invariant) distribution satisfies

$$\vec{p} = P\vec{p}. \tag{5}$$

The components of this steady state distribution can be found explicitly as solutions of recursive equations.

For various 'small' values of n, the results are easily written explicitly. For instance, when $n = 6$

$$p_1 = \frac{(\rho-1)^4}{2\left(\rho^4-2\rho^3+4\rho^2-3\rho+1\right)}$$

$$p_2 = -\frac{(\rho-1)^3}{2\left(\rho^4-2\rho^3+4\rho^2-3\rho+1\right)}$$

$$p_3 = \frac{(\rho-1)^2\rho}{2\left(\rho^4-2\rho^3+4\rho^2-3\rho+1\right)}$$

$$p_4 = -\frac{(\rho-1)\rho^2}{2\left(\rho^4-2\rho^3+4\rho^2-3\rho+1\right)}$$

$$p_5 = \frac{\rho^3}{2\rho^4-4\rho^3+8\rho^2-6\rho+2}$$

$$p_6 = \frac{\rho^4}{2\rho^4-4\rho^3+8\rho^2-6\rho+2}.$$

The mixing rate for the $n = 6$ Markov chain is given as:

$$\mu = \sqrt{\frac{3+\sqrt{5}}{2}\rho(1-\rho)}.$$

While this approach can be taken from simple environments for a single agent, it obviously becomes cumbersome as the environments become even slightly more complex and there are more sentries performing the surveillance.

While $\rho = 1$ or 0 will yield the fastest mixing Markov chain, they will also yield a steady state distribution that is the furthest from uniform. Conversely, $\rho = 1/2$ will yield a uniform distribution but will also be the slowest mixing Markov chain. (Strictly speaking, the assumption that the walk leaves the right-most and left-most states with probability one implies that the steady state distribution with $\rho = 1/2$ is uniform except at the two end states. At the cost of simplicity in computing explicit solutions to (4), one can enforce uniformity in the steady state distribution by allowing a probabilistically determined dwell time in the two end states.) We have a developing - but as yet incomplete - understanding of the trade-off between speed of convergence and a uniform steady state distribution in our strategies for stochastic agents.

We next observe that even a small bias in the turning probabilities leads to a significant skewing of the invariant distribution. Figure 2 shows the invariant distribution for a single agent on a 16-node, 1-D lattice with a probability

$\rho = 0.51$ of moving to the right and a probability of 0.49 of moving to the left. Increasing the bias of motions to the right further increases the deviation of the invariant distribution from uniform.

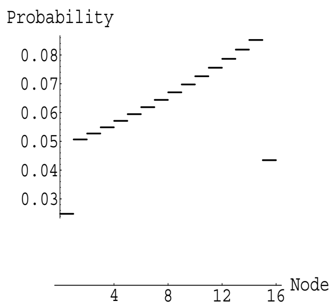

Fig. 2. Invariant distribution for a single agent.

We are also concerned with the transient behavior. Thus, we define an empirical measure of non-uniformity in the following manner:

$$Non-uniformity(k) = \frac{\Sigma_i(\tilde{x}_i - \tilde{\pi}_i)^2}{n^2}, \qquad i = 1, \ldots, \tag{6}$$

$$\tilde{x}_i = \frac{x(i) * n}{k * a} \tag{7}$$

$$\tilde{\pi}_i = \pi_i * n \tag{8}$$

$$\lim_{k \to \infty} \; Non-uniformity(k) \to 0, \tag{9}$$

where $\tilde{x}_i$ is a history of visitation frequency for state i, x_i is the visitation history, $\tilde{\pi}_i$ is a normalized invariant distribution for state i, n is the number of states, a is the number of agents, and k is the number of steps each agent has taken.

This non-uniformity measure essentially quantifies how quickly the surveillance team covers the environment. Mathematically it is the mean 'distance' the ensemble of agents are from the composite invariant distribution. It is important to normalize the visitation history and the invariant distribution because the measure needs to be applicable to both small and large environments. As k increases, the normalized visitation history will approach the normalized steady state distribution. In other words, $\tilde{x}_i$ approaches $\tilde{\pi}_i$. Thus, the non-uniformity approaches 0 as k approaches ∞.

With this non-uniformity measure, we compare different biases in the local transition probabilities and consider adapting those biases as time evolves to minimize the non-uniformity. Figure 3 shows us how the non-uniformity for different biases evolves over time on a 1D, 10 state lattice with a single agent entering from the left. It is apparent that with a significant bias at the start, we can quickly reduce the non-uniformity measure. Then, as time evolves, the distribution will begin to move away from the uniform distribution. At this point, it is better to have equal turning probabilities. In this manner, we can control the behavior of the agents.

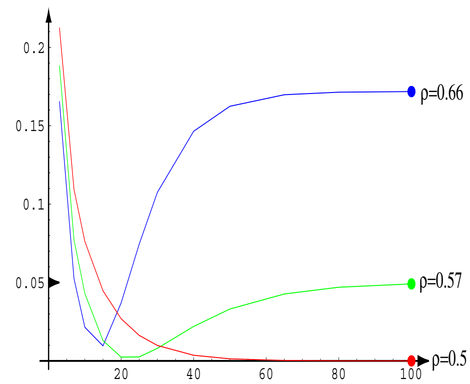

Fig. 3. Plots of the deviations from uniformity of probability distributions after k steps for random walks on a ten-node, one-dimensional lattice. Red denotes the symmetric walk with probabilities of $\rho = 1/2$ that the walker takes a step to the right and $1-\rho$ that she takes a step to the left. Green is a rightward biased walk with probability $\rho = 0.57$ the the walker takes a step to the right at each time instant, and blue represents the rightward biased walk with $\rho = 0.66$. Although only symmetric the random walk has the uniform distribution as its steady state asymptotic limit, both the biased walks (starting with probability 1 at the left-most endpoint) have intermediated distributions that are closer (significantly closer in the 'blue' case) to uniform after a relatively few steps.

A quick look at the entropy for each of the biases shown in Figure 3 and we notice that while a small bias (from uniform) results in a noticeable gain in the initial speed of convergence, there is only a small decrease in the parametric randomness. From the earlier equation for Shannon's entropy, we see that the entropy for the biases of $\rho = 0.66$, $\rho = 0.57$ and $\rho = 0.5$ are 0.278, 0.298, and 0.301, respectively.

Such one dimensional problems are useful in setting the expectations for what will happen in more complex geometries with a more complex model dependence on parameters.

B. Multiple Agent Investigations

Here, we present a strategy for 1D, n-node lattice. While this strategy is ad hoc, it is feasible and sheds some light on aspects of multi-agent strategies. The idea is to investigate 'adiabatically' adapting the turning parameters in order to achieve a more efficient surveillance team. By appropriate choice of parameters, it will be possible to implement a probabilistic strategy in which agents disperse in the lattice domain as fast as possible. That's the idea behind the following strategy. The turning parameters of this strategy are assigned by the following algorithm:

$$\rho_i = \begin{cases} .9 & for \;\; k \le \frac{a+1-i}{a+1} * n \\ .5 & otherwise, \end{cases}$$

where ρ_i is the probability that agent i turns right, k is the number of steps, a is the number of agents and n is the number of nodes in the lattice. This strategy disperses the agents along graph before switching to equal turning probabilities. Clearly, the agents do not have a uniform steady state distribution prior to switching to the equal turning probabilities. After the agents switch to the equal turning probabilities, their steady state distribution will approach the uniform distribution as the initial distributions are suppressed with time.

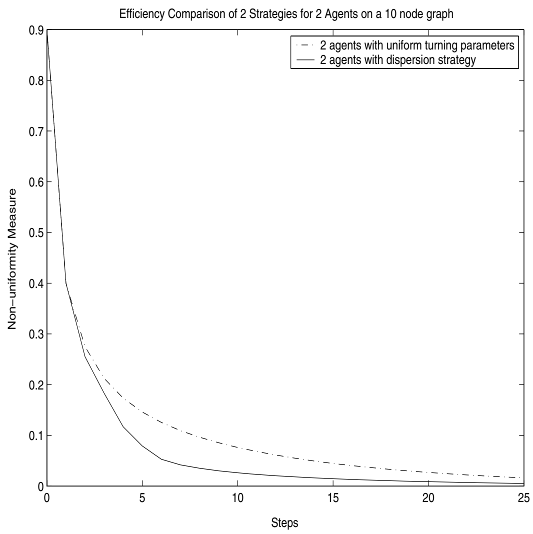

Fig. 4. This figure compares two strategies for controlling the motions of two agents on a 1-D, 10-node graph. Each agent has a parameter which represents the probability of moving to the right. The dotted line shows the non-uniformity measure as a function of the number of transitions taken (a measure of time) for the case where both agents have no bias in there parameters (i.e. $\rho = 0.5$). The solid line shows the non-uniformity measure as a function of transitions taken when the agents employ the dispersion strategy described above. We can see that by inducing an initial bias, the agents approach their steady state distribution more quickly.

V. COMPLETE GRAPHS

As with the problems in one dimension, we wish to investigate complete graphs and the associated turning parameters that yield the fastest mixing and approach the uniform steady state distribution. In the next section, we will combine the strategies from one-dimension and complete graphs to create a more general strategy.

Theorem 5.1: For a complete graph with n vertices, the probabilistic random walk having transition probability matrix

$$p_{ij} = \begin{cases} 1/(n-1) & \text{if} \quad i \neq j \\ 0 & \text{if} \quad i = j, \end{cases}$$

has eigenvalues 1 of multiplicity one and $-1/(n-1)$ of multiplicity $n-1$. The invariant distribution for this Markov chain is uniform, and the eigenvalue $-1/(n-1)$ is smaller in magnitude than the eigenvalue of second largest magnitude corresponding to any other set of transition probabilities.

Proof: Using some basic properties of rank one matrix perturbations, we can show that

$$\left| sI - \begin{pmatrix} 0 & \frac{1}{n-1} & \dots & \frac{1}{n-1} \\ \frac{1}{n-1} & 0 & \dots & \frac{1}{n-1} \\ \vdots & \vdots & \dots & \vdots \\ \frac{1}{n-1} & \frac{1}{n-1} & \dots & 0 \end{pmatrix} \right| = (s + \frac{1}{n-1})^{n-1}(s-1).$$

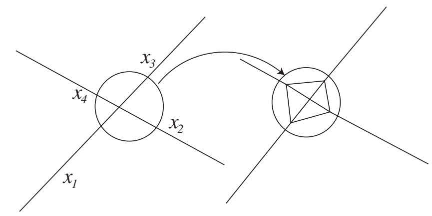

Fig. 5. The hypergraph on the left, with edge set $\{x_1, x_2, x_3, x_4\}$ corresponds to an *augmented* graph which includes edges $\{\{x_1, x_2\}, \{x_1, x_3\}, \{x_1, x_4\}, \{x_2, x_3\}, \{x_2, x_4\}, \{x_3, x_4\}\}$.

This proves the statement regarding eigenvalue multiplicities. Suppose

$$P = \begin{pmatrix} 0 & p_{12} & \dots & p_{1n} \\ p_{21} & 0 & \dots & p_{2n} \\ \vdots & \vdots & \dots & \vdots \\ p_{n1} & p_{n2} & \dots & 0 \end{pmatrix}$$

is any other matrix of transition probabilities for the graph. The zeros on the diagonal reflect the fact that there are no self-loops. Let $1, r_1, r_2, \dots, r_{n-1}$ be an enumeration of the eigenvalues of P. Then we know that $r_1 + \dots + r_{n-1} = -1$. If any r_i is less than $1/(n-1)$ in magnitude, there must be some other r_j which is greater than $1/(n-1)$ in magnitude. From this it follows that the eigenvalue of second largest magnitude has magnitude greater than or equal to $1/(n-1)$. This proves the theorem. ■

VI. GENERAL GRAPH STRATEGIES

In the most general case, we can consider the structure of the problem as line segments on a bounded plane. Recall that the *rank* of an edge in a hypergraph is the number of vertices contained in that edge. To each edge of rank greater than or equal to three in $H(X, \mathscr{E})$ we shall associate a complete graph on its vertices. With $H(X, \mathscr{E})$ representing the set of corridor segments in our search environment, the edges of these complete graphs represent the possible choices of transitions from one corridor segment to another. We shall construct a new graph $\hat{H}(X, \hat{\mathscr{E}})$ whose vertex set is X and hose edge set $\hat{\mathscr{E}}$ is the same as $\mathscr{E}$ with each edge of rank three or higher replaced by the collection of two-element subsets of that edge. The construction is illustrated in Figure 5. We shall call $\hat{H}(X, \hat{\mathscr{E}})$ the *augmented graph* associated with our surveillance environment. Any general graph can be decomposed into a system of interconnected complete subgraphs (cliques). We shall distinguish between cliques having two vertices and those having more than two vertices. Those having two vertices will correspond to transitions in which the only choices available to the agent are to move ahead or to move backward. Because we wish to consider models which provide a relatively fine-grained description of long corridors, we shall allow for long strip two-element cliques in our graphs. This allows us to consider agents which collect information from sensors with limited range as part of the surveillance duties.

The intersections of a general graph can be thought of as a complete graph where the number of nodes in the complete graph is equivalent to the number of choices an agent can make. With no restrictions on the movement of an agent, the number of nodes in the complete graph is equal to number of edges incident to the intersection in the graph representation.

We use the ideas developed in the one-dimensional and complete graph investigations to create a strategy for an agent providing surveillance on a general graph. Here, we ask the agent to provide uniform surveillance coverage for all states (i.e. visit all states approximately the same number of times) while maintaining some sense of randomness in the behavior. For the linear segments with intervals of demarkation, we will induce a bias for the agent's transition probability for a specified amount of time then adjust the agent's parameters so the turning probabilities are uniform. This is shown by the following:

$$\rho \;\; = \;\; \begin{cases} 2/3 \;\; for \;\; k \leq \rho * n \\ .5 \;\; otherwise, \end{cases}$$

where ρ is the probability of transitioning toward the center of the linear graph, k is the number of transitions taken on the linear graph, and n is the number of nodes in the linear graph. We have shown in Section IV that, an agent on a linear graph, a small initial parameter bias leads to a significantly faster convergence to the uniform steady state distribution with only a small decrease in agent randomness. Then in order to maintain that steady state distribution, the parameters must be adjusted so that the agent has equal turning probabilities.

For graph sections that decompose into complete graphs, we showed in the previous section that uniform turning parameters not only yield the fastest convergence to the uniform distribution but also has the largest parameter entropy (i.e. the most randomness in agent behavior).

By combining the strategies for both linear graphs and complete graphs, we have a hybrid strategy which will provide uniform coverage of a general graph while achieving this coverage quickly without a large sacrifice in the randomness in the behavior of the agent.

VII. CONCLUSIONS

We have formalized the mathematical approach to understanding the surveillance problem with stochastic sentries. We showed that we can describe environments of interest as hypergraphs, thereby giving us a basis for analysis. We defined a metric of non-uniformity that provides a measures how close a composite distribution is from its steady state distribution as well as use a measure for parametric randomness. We continue to study the problem of finding optimal rules for searching a graph with a team of agents. We have shown that some 'small' cases, the strategies can be evaluated explicitly but, in general, one must employ the techniques of Monte Carlo simulation. The approach in all cases has been to provide each agent in a team with a time-varying transition bias which will allow the team as a whole to most rapidly provide uniformly random surveillance of an augmented graph. In forthcoming work, we shall discuss transitions with rules of exclusion for multi-agent systems. This work is aimed at developing team strategies based on optimal local transition control of mixing in finite state Markov chains.

REFERENCES

[1] Chvatal, V. A Combinatorial Theorem in Plane Geometry. *Journal of Combinatorial Theory (B)* **18**, 39-41, 1975.

[2] O'Rourke, J. Art Gallery Theorems and Algorithms. Oxford University Press, 1987.

[3] Chin, W. and Ntafos, S. Optimum Watchman Routes. *Symposium on Computational Geometry*, 24-33, 1987.

[4] Chin, W. and Ntafos, S. Shortest Watchman Routes in Simple Polygons. *Discrete and Computational Geometry* **6**: 9-31, 1991.

[5] Carlsson, S., Nilsson, B. and Ntafos, S. Optimum Guard Covers and *m*-Watchmen Routes for Restricted Polygons. *International Journal of Computational Geometry and Applications* **3(1)**:85-105, 1993.

[6] Carlsson, S., Jonsson, H. and Nilsson, B. Finding the Shortest Watchman Route in a Simple Polygon. *ISAAC* 58-67, 1993.

[7] Yamauchi, B. A Frontier-Based Approach for Autonomous Exploration. *IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 1997.

[8] Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M., Thrun, S.M. and Younes, H. Coordination for Multi-Robot Exploration and Mapping. *Proceedings of 17th National Conference on Artificial Intelligence* 852-858, 2000.

[9] Burgard, W. Moors, M., Fox, D., Reid S., and Thrun, S. Collaborative Multi-Robot Exploration. *IEEE International Conference on Artificial Intelligence* 852-858, 2000.

[10] Acar, E., and Choset, H. Robust Sensor-based Coverage of Unstructured Environments. *IEEE International Conference on Robotics and Automation (ICRA)*, 2000.

[11] Yang, Y., Minai, A., and Polycarpou, M. Analysis of Opportunity Method for Cooperative Search by Mobile Agents, *IEEE Conference on Decision and Control*, 2002.

[12] Flint, M., Polycarpou, M. and Fernandez-Gaucherand, E. Cooperative Control for Multiple Autonomous UAV's Searching for Targets. *IEEE Conference on Decision and Control*, 2002.

[13] Hespanha, J., Kim, H., Sastry, S. Multiple-Agent Probabilistic Pursuit-Evasion Games. *IEEE CDC* **3**:2432-2437, 1999.

[14] Kim, H., Vidal,R., Shim, D., Shakeria, O. and Sastry, S. A Hierarchical Approach to Probabilistic Pursuit-Evasion Games with Unmanned Ground and Aerial Vehicles. *IEEE CDC*, 2001.

[15] Kemeny, J. and Snell, J. Finite Markov Chains. Springer-Verlag, 1976.

[16] Feller, W. An Introduction to Probability Theory and Its Applications. Wiley, 1968.

[17] Rosenthal, J. Convergence Rates of Markov Chains. *SIAM Review* **37**:387-405, 1994.

[18] Diaconis, P. and Stroock, D. Geometric Bounds for Eigenvalues of Markov Chains. *Annals of Applied Probability* **1**:36-61, 1991.

[19] Boyd, S., Diaconis, P. and Xiao, L. Fastest Mixing Markov Chain on a Graph. *SIAM Review*, accepted Mar 2004.

[20] Berge, C., Graphs and Hypergraphs. North Holland, 1970.

[21] Shannon, C, The Mathematical Theory of Communication, University of Illinois , 1963.