Robotic surveillance and Markov chains with minimal first passage time

Pushkarini Agharkar[1] Rushabh Patel [1] Francesco Bullo [2]

Abstract— **We propose stochastic surveillance strategies for quickest detection of anomalies in discrete network environments. Our surveillance strategy is determined by optimizing the mean first passage time also known as the Kemeny constant of a Markov chain. We generalize the notion of mean first passage time to environments with heterogeneous travel and service times and develop analogous theories and relations to previous results in the literature. For reversible Markov chains, we show that both the Kemeny constant and its heterogeneous counterpart can be formulated as convex optimization problems and, moreover, can be expressed as semidefinite programs (SDPs). We illustrate the performance of the proposed surveillance strategies with numerical simulations.**

I. INTRODUCTION

A. Problem description and motivation

The subject of this paper is the minimization of the mean first passage times between any two nodes over a network. This problem is akin to minimizing the travel time over a network, a field of research which has applications in a broad range of areas. In the area of robotics, key applications include surveillance tasks such as reducing detection time of anomalies, and quickest detection of intruders [14], [18]. Other applications that arise in single and multi-agent systems include tasks such as minimizing emergency vehicle response times [2] as well as various servicing tasks which can arise in robotic warehouse management [20]. In areas of research outside of robotics, minimizing travel time has potential applications in analyzing how quickly information can propagate in social networks [1] or determining how an epidemic spreads through the population [19].

B. Literature review

In this paper we consider the design of a surveillance strategy over an environment in order to detect an intruder while simultaneously providing a desired coverage frequency. The surveillance problem has appeared in the literature in various manifestations. Authors in [18] look at minimizing time of detection of noisy anomalies via persistent surveillance strategies. Robot networks are utilized for intruder detection in [13] and a game theoretic approach is presented in [16], [5].

This work was supported in part by NSF grant CNS 1035917 and ARO grant W911NF-11-1-0092.

[1] Pushkarini Agharkar and Rushabh Patel are with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, CA 93106-5070, USA agharkar@umail.ucsb.edu, r_patel@engineering.ucsb.edu

[2] Francesco Bullo is with Faculty of Mechanical Engineering, University of California Santa Barbara, CA 93106-5070, USA bullo@enigneering.ucsb.edu

In our setup we utilize a discrete set of locations and thus formulate our problem as a random walk on a graph. Other works with a similar setup optimize surveillance strategies using Markov chain Monte Carlo methods [7], [17]. In [7] they also utilize convexity results for symmetric matrices to further optimize their strategy. We take an alternate approach and look at minimizing the property of a Markov chain known as the mean first passage time. The mean first passage time, also known as the Kemeny constant of a Markov chain, first appeared in [10] and has since been studied by several other groups [9], [11]. It has also been linked to the mixing rate of a Markov chain [11], a quantity which has received wide interest [3]. Results analyzing bounds on the mean first passage time for an arbitrary Markov chain for various network topologies appear in [9], [12], however, there are no results known that focus on the design of the mean first passage time.

C. Contribution

Before stating our contributions, it should be noted that all work done to date on the *mean first passage time* and mixing rate of a Markov chain has been for graphs which assume homogeneous travel time along edges in the graph. The contributions of this paper are then six fold. First, we provide a convex optimization framework to minimize the Kemeny constant of a reversible Markov chain given the underlying graph topology of the random walk and the desired stationary distribution. Second, using doubly-weighted graphs we extend the formulation of the *mean first passage time* to the case where there are non-homogenous travel times along the edges of a graph, a property not yet looked at in the literature. Third, we derive a closed form solution for the mean first passage time for the doubly-weighted graph scenario and show its relation to the Kemeny constant. Fourth, we provide a convex optimization framework to minimize the mean first passage time of a reversible Markov chain with desired stationary distribution for the doubly-weighted graph setup. Finally, we provide a SDP formulation for the optimization of the Kemeny constant and the mean first passage time for doubly-weighted graphs. Due to space constraints, all proofs are omitted and will be made available in a forthcoming full-length paper.

D. Organization

The paper is organized as follows. In section II we provide relevant mathematical background and notation which we use throughout the paper. In section III we present relevant known results as well as provide a convex optimization formulation to find the Kemeny constant and provide and show it can be written as an SDP. In section IV we provide a closed form expression for the mean first passage time of a doubly-weighted graph. We then provide a convex optimization framework for the mean first passage time of a doubly-weighted graph and show it can be written as an SDP. In section V we provide examples which tie the mean first passage time for doubly-weighted graphs to practical applications. In section VI we present simulation results comparing our surveillance strategies against other established Markov chain algorithms, and in the final section we present our conclusions and future research directions.

E. Notation

We use the notation $A = [a_{ij}]$ to denote a matrix A with the element a_{ij} in its i-th row and j-th column and, unless otherwise indicated, use bold-faced letters to denote vectors. Letting δ_{ij} denote the Kronecker delta, $A_d = [\delta_{ij} a_{ij}]$ represents the diagonal matrix whose diagonal elements are the diagonal elements of the matrix A. The column vector of all ones and length n is denoted by $\mathbb{1}_n \in \mathbb{R}^{n\times 1}$ and I represents the identity matrix of appropriate dimension.

II. Problem Setup

Consider an environment under surveillance where the environment is divided into n regions defined by the set of vertices $\mathcal{V} := \{v_1, \ldots, v_n\}$. These points can represent small areas of interest, and are assumed to be connected by weighted edges. Let $\mathcal{G} = (\mathcal{V}, E, P)$ be an undirected weighted connected graph with edge set $E \subset \mathcal{V} \times \mathcal{V}$ and weight matrix $P = [p_{ij}]$ with the property that $p_{ij} \geq 0$ if $(i, j) \in E$ and $p_{ij} = 0$ otherwise. The weight of edge (i, j) can be thought of as the distance that a vehicle has to travel in order to traverse the edge or the probability of moving along that edge.

Consider an autonomous vehicle that surveils these regions. The trajectory of the vehicle is viewed as a random walk which is assumed to be a Markov chain. The vehicle visits a new region, or possibly stays in its current region, at each time step (*hop*) in its trajectory. The surveillance objective is to achieve a long term (infinite horizon) visit frequency $\boldsymbol{\pi}_i$ at region v_i. Here, the frequency $\boldsymbol{\pi}_i$ is the ratio of visits to region v_i divided by the total number of visits to all the regions. We refer to the vector of visit frequencies $\boldsymbol{\pi} \in \mathbb{R}^n$ as the surveillance criterion.

Given an environment of interest, the routing policy of a vehicle depends on the level of sophistication of the intruder it is trying to detect as well as the desired performance objective for surveillance of the environment. We study a specific surveillance objective: if there exists a stationary intruder continuously performing some malicious activity in some part of the environment, we wish to design a stochastic policy for the vehicle in order to minimize the detection time of that intruder while simultaneously achieving a desired pre-determined visit frequency to all regions in the environment.

Throughout the paper, we describe random walks of the vehicle using realizations of a Markov chain with transition matrix $P = [p_{ij}]$. The element p_{ij} in the matrix represents the probability with which the vehicle visits region v_j from region v_i.

A. Background

A finite *Markov chain* is a sequence of random variables where X_t take values in $\{1, \ldots, n\}$ with the Markov property, namely that, the future state depends only on the present state. In this paper, we consider only finite row stochastic Markov chains.

We summarize some terminology regarding finite Markov chains for later reference. For more details refer to [8].

i. A *time-homogenous* Markov chain has the following property: $\mathbb{P}(X_{n+1} = j | X_n = i) = \mathbb{P}(X_n = j | X_{n-1} = i) = p_{ij}$, where $P \in \mathbb{R}^{n\times n}$ is the transition matrix of the Markov chain.
ii. The vector $\boldsymbol{\pi} \in \mathbb{R}^{n\times 1}$ is a *stationary distribution* of P if $\sum_{i=1}^n \pi_i = 1$, $0 \leq \pi_i \leq 1$ $\forall i \in \{1, \ldots, n\}$ and $\boldsymbol{\pi}^T P = \boldsymbol{\pi}^T$.
iii. A time-homogenous Markov chain is said to be *reversible* if $\pi_i p_{ij} = \pi_j p_{ji},\ \forall\ i, j \in \{1, \ldots, n\}$. For reversible Markov chains, $\boldsymbol{\pi}$ is always a steady state distribution.
iv. A Markov chain is *irreducible* if it has a finite state space and there exists a t such that for all $i, j \in \{1, \ldots, n\}, (P^t)_{ij} > 0$.
v. If the Markov chain is *finite irreducible* of dimension n, then there is a unique stationary distribution $\boldsymbol{\pi}$, and the corresponding eigenvalues of the transition matrix λ_i for $i \in \{1, \ldots, n\}$ are such that $\lambda_1 = 1$, $|\lambda_i| \leq 1$ and $\lambda_i \neq 1$ for $i \in \{2, \ldots, n\}$.

In this paper we consider finite irreducible time-homogenous Markov chains in the analysis and design of stochastic routing policies. We are now ready to state some standard results on generalized inverses. We begin with the following definition.

III. Minimizing the Kemeny constant of a Markov chain

Consider a scenario where an intruder wishes to perform malicious activities somewhere in an environment but if detected, it is permanently deactivated. Assume also, that if a surveillance vehicle is in the same region as the intruder, then it detects the intruder. Then, the optimal strategy for a surveillance vehicle should minimize the detection time of the intruder. Let $X_k = i$ denote that the vehicle is in region v_i at time k. Suppose the vehicle starts at an arbitrary region v_i in the environment so that $X_0 = i$. Let T_{ij} be the random variable tracking the first passage time from state i to state j, that is, $T_{ij} = \min\{k \geq 1 \text{ such that } X_k = j \text{ given that } X_0 = i\}$ and let $m_{ij} = \mathbb{E}[T_{ij} | X_o = i]$. Then, regardless of which region v_i the surveillance vehicle starts in, we wish to minimize the mean first passage time, $\boldsymbol{k}_i = \sum_{j=1}^n m_{ij} \boldsymbol{\pi}_j$, from that region to any other region. Remarkably, for a Markov chain associated with a graph, the mean first passage time $\boldsymbol{k}_i = K$ for all $i \in \{1, \ldots, n\}$, where K is known as the Kemeny constant [9].

Henceforth, we will use the term *first passage time matrix* to denote the matrix $M = [m_{ij}]$ where m_{ij} is the mean time to go from v_i to v_j. Mathematically,

$$m_{ij} = p_{ij} + \sum_{k \neq j} p_{ik}(m_{ik} + 1) = 1 + \sum_{k \neq j} p_{ik} m_{kj},$$

or in matrix notation,

$$(I - P)M = \mathbb{1}_n \mathbb{1}_n^T - P M_d, \tag{1}$$

where P is the transition matrix of the Markov chain associated with the graph. Moreover, if the Markov chain is irreducible with stationary distribution $\boldsymbol{\pi}$, then $M_d = \mathbf{diag}[\{1/\boldsymbol{\pi}_1, \ldots, 1/\boldsymbol{\pi}_n\}]$ Further $\boldsymbol{\pi}^T M \boldsymbol{\pi} = \sum_{i=1}^n \pi_i \sum_{j=1}^n \boldsymbol{\pi}_j m_{ij} = \sum_{i=1}^n \boldsymbol{\pi}_i \boldsymbol{k}_i = K$. Clearly the Kemeny constant can be written as the function $P \mapsto K(P)$, however, to ease notation we simply write K and use $K(P)$ when we wish to emphasize the constant's dependence on P.

Finding $K = \boldsymbol{\pi}^T M \boldsymbol{\pi}$ from (1) can be simplified to finding the eigenvalues of the transition matrix P as is stated in the following theorem.

Theorem 1 (Kemeny constant of a Markov chain [9]): Let $\mathcal{G}$ be a graph with associated irreducible transition matrix P with eigenvalues λ_i for $i \in \{1, \ldots, n\}$ where $\lambda_1 = 1$, then the Kemeny constant of the Markov chain is given by

$$K = 1 + \sum_{i=2}^{n} \frac{1}{1 - \lambda_i}.$$

Using Theorem 1, we derive the following equivalent expression for *reversible* Markov chains in terms of the trace of a symmetric positive definite matrix. Before stating our result, we first introduce some notation. Given a stationary distribution vector $\boldsymbol{\pi} \in \mathbb{R}^n$ for a Markov chain with transition matrix $P \in \mathbb{R}^{n \times n}$, we define the matrix $\Pi \in \mathbb{R}^{n \times n}$ as $\Pi = \mathbf{diag}[\boldsymbol{\pi}]$ and the vector $\boldsymbol{q} \in \mathbb{R}^n$ as $\boldsymbol{q}^T = (\sqrt{\boldsymbol{\pi}_1}, \ldots, \sqrt{\boldsymbol{\pi}_n})$. We are now ready to state our first result.

Theorem 2: (Kemeny constant of a reversible Markov chain): The Kemeny constant of a reversible irreducible Markov chain with transition matrix P and stationary distribution $\boldsymbol{\pi}$ is given by

$$K = \mathbf{Tr}\left[(I - \Pi^{1/2} P \Pi^{-1/2} + \boldsymbol{q}\boldsymbol{q}^T)^{-1}\right]. \tag{2}$$

Given the above result, we are now ready to state our first problem of interest.

Problem 3: (Optimizing the Kemeny constant of a reversible Markov chain): Given the stationary distribution $\boldsymbol{\pi}$ and graph $\mathcal{G}$ with vertex set $\mathcal{V}$ and edge set E, determine the transition probabilities $P = [p_{ij}]$ solving:

$$\begin{aligned}
\text{minimize} \quad & \mathbf{Tr}\left[(I - \Pi^{1/2} P \Pi^{-1/2} + \boldsymbol{q}\boldsymbol{q}^T)^{-1}\right] \\
\text{subject to} \quad & \sum_{j=1}^{n} p_{ij} = 1, \text{ for each } i \in \{1, \ldots, n\} \\
& \boldsymbol{\pi}_i p_{ij} = \boldsymbol{\pi}_j p_{ji}, \text{ for each } (i,j) \in E \\
& 0 \leq p_{ij} \leq 1, \text{ for each } (i,j) \in E \\
& p_{ij} = 0, \text{ for each } (i,j) \notin E.
\end{aligned} \tag{3}$$

The following theorem establishes the convexity of the Kemeny constant for transition matrices with fixed stationary distribution.

Theorem 4 (Convexity of Problem 3): Let $\mathcal{P}_{\boldsymbol{\pi}}$ denote the set of matrices associated to irreducible reversible Markov chains with stationary distribution $\boldsymbol{\pi}$. Then, $\mathcal{P}_{\boldsymbol{\pi}}$ is a convex set and $P \mapsto K(P)$ is a convex function over $\mathcal{P}_{\boldsymbol{\pi}}$.

Problem 3 includes constraints on the stationary distribution of the transition matrix, a notion which has not been looked at in the literature before. [11] provides bounds to determine the set of transition matrices such that K is minimized and [9] gives special matrices for which the optimal Kemeny constant can be found, but these are all approached for the general setting with no constraint on the actual stationary distribution. In the following we introduce the notion of distance between regions when calculating the mean first passage time.

A. SDP framework for optimizing the Kemeny constant

We can re-write Problem 3 for solving the Kemeny constant as an SDP by introducing a symmetric slack matrix $X \in \mathbb{R}^{n \times n}$ and writing the problem as in the following.

Problem 5: (Optimizing the Kemeny constant of a reversible Markov chain (SDP)): Given the stationary distribution $\boldsymbol{\pi}$ and graph $\mathcal{G}$ with vertex set $\mathcal{V}$ and edge set E, determine $X = [x_{ij}]$ and the transition probabilities $P = [p_{ij}]$ solving:

$$\begin{aligned}
\text{minimize} \quad & \mathbf{Tr}[X] \\
\text{subject to} \quad & \\
& \begin{bmatrix} I - \Pi^{1/2} P \Pi^{-1/2} + \boldsymbol{q}\boldsymbol{q}^T & I \\ I & X \end{bmatrix} > 0 \\
& \sum_{j=1}^{n} p_{ij} = 1, \text{ for each } i \in \{1, \ldots, n\} \\
& \boldsymbol{\pi}_i p_{ij} = \boldsymbol{\pi}_j p_{ji}, \text{ for each } (i,j) \in E \\
& 0 \leq p_{ij} \leq 1, \text{ for each } (i,j) \in E \\
& p_{ij} = 0, \text{ for each } (i,j) \notin E.
\end{aligned}$$

The first inequality constraint in Problem 5 represents a linear matrix inequality (LMI) and denotes that the matrix is positive definite and. The LMI holds true if and only if $X - (I - \Pi^{1/2} P \Pi^{-1/2} + \boldsymbol{q}\boldsymbol{q}^T)^{-1}$ is positive definite therefore the SDP given by Problem 5 minimizes the Kemeny constant.

IV. STOCHASTIC SURVEILLANCE USING DOUBLY-WEIGHTED GRAPHS

In most practical robotic applications, distance/time traveled and service times are important factors in designing

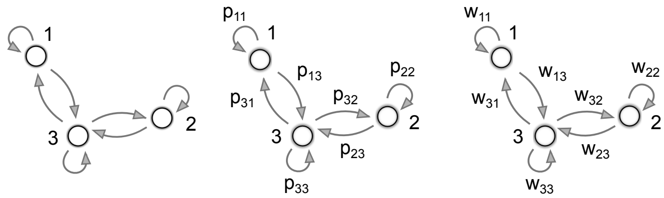

Fig. 1. Example of a doubly-weighted graph $\mathcal{G} = (\mathcal{V}, E, P, W)$ with three nodes: (a) shows the edge set, E, allowed for the graph with three nodes, (b) shows the probabilities, p_{ij} to move along each edge, and (c) shows the time (i.e. distance traveled), ω_{ij} to move along each edge.

routing policies. We incorporate these concepts by allowing for an additional set of weighted edges in our graph in addition to the edge weights which describe the transition probabilities. Such an environment can be represented by the doubly-weighted graph $\mathcal{G} = (\mathcal{V}, E, P, W)$, where $W = [\omega_{ij}]$ is a weight matrix with the property that $\omega_{ii} \geq 0$, $\omega_{ij} > 0$ if $(i, j) \in E$ and $\omega_{ij} = 0$ otherwise. The weighted adjacency matrix $P = [p_{ij}]$ has the same interpretation as before as a irreducible row-stochastic transition matrix P which governs the random walk on the graph. An example of a doubly-weighted graph is shown in Figure 1. In the following, we will interpret ω_{ij}, $i \neq j$ as time to travel between two regions in the graph and ω_{ii} as service time.

If X_k is the state of the vehicle at the $k-$th step in its random walk, then a sample trajectory of the vehicle is $X_0 = v_i$, $X_1 = v_j$, $X_2 = v_k$, and the time instant at which the vehicle is in state X_2 is $\omega_{ij} + \omega_{jk}$. Thus the time interval between two consecutive steps of this random walk depends on the weighted adjacency matrix, W, of the graph and is not constant.

In the following analysis, we consider two objectives for surveillance: The first involves minimizing the mean first passage time (a quantity that is analogous to the Kemeny constant for the described in section III) and the second looks at minimizing the mean time for a single *hop* of the vehicle. The first objective is motivated by the need to minimize detection time of malicious activities in the environment and the second objective can realize aims such as minimizing fuel consumption of the vehicle, minimizing refresh times (time interval between consecutive visits) for all the regions under surveillance. We seek transition matrices P with stationary distribution $\boldsymbol{\pi}$ for realizing the vehicle's trajectory. We start with developing some foundation for achieving the first objective.

A. *Minimizing the mean first passage time for a doubly-weighted graph*

The mean first passage time for the Markov chain on a weighted graph $\mathcal{G} = (\mathcal{V}, E, P)$ by definition, is simply its Kemeny constant. Recall that the mean first passage time for node i, defined by $\boldsymbol{k}_i$, is determined by taking the expectation over the first passage times m_{ij}, from node i to all other nodes j. We present an analogous notion of the first passage time between two nodes on a doubly-weighted graph. We start with defining the first passage time random variable for a random walk on a doubly-weighted graph.

Definition 6 (First passage time random variable): The first passage time T_{ij} to go from v_i to v_j is

$$T_{ij} = \begin{cases} \omega_{ij} & \text{with probability } p_{ij}, \\ T_{kj} + \omega_{ik} & \text{with probability } p_{ik}, k \neq j. \end{cases}$$

Then taking expectations over T_{ij} gives the following result.

Lemma 7: (First passage time for a doubly-weighted graph): If $n_{ij} = \mathbb{E}[T_{ij} | X_0 = i]$ is the mean first passage time to go from v_i to v_j for a graph with weight matrix W and transition matrix P, then

$$n_{ij} = p_{ij}(\omega_{ij}) + \sum_{k \neq j} p_{ik}(n_{kj} + \omega_{ik}), \tag{4}$$

or in matrix notation,

$$(I - P)N = (P \circ W)\mathbb{1}_n \mathbb{1}_n^T - PN_d, \tag{5}$$

where $(P \circ W)$ denotes an element-wise product and $N_d = [\delta_{ij} n_{ij}]$.

The matrix N, which we call the *first passage time matrix for a doubly-weighted graph* thus satisfies an equation similar to (1) for the passage time matrix M of a graph which has only one weight matrix, the transition matrix P.

The random variable tracking the time interval between consecutive visits to a region has been referred to as the *refresh time* of the vehicle [15] and n_{ii} is the expected value of the refresh time for a stochastic policy. We now obtain a relation between $\boldsymbol{\pi}$ and the refresh times n_{ii}.

Theorem 8 (Refresh times for doubly-weighted graphs): Given the graph $\mathcal{G} = (\mathcal{V}, E, P, W)$ with stationary distribution $\boldsymbol{\pi}$ and associated first passage time matrix N, the refresh time for region v_i is given by $n_{ii} = (\boldsymbol{\pi}^T (P \circ W)\mathbb{1}_n)/\boldsymbol{\pi}_i$, implying that

$$N_d = \boldsymbol{\pi}^T (P \circ W)\mathbb{1}_n M_d.$$

This means that in order to achieve a desired refresh time n_{ii}, the vehicle should visit region v_i with frequency proportional to $1/\boldsymbol{\pi}_i$. It is interesting to note that the visit frequency is not a function the weight matrix W. We now investigate the properties of the mean first passage times of the trajectory of the vehicle.

The mean first passage time of a vehicle for a doubly-weighted graph $\mathcal{G} = (\mathcal{V}, E, P, W)$ with associated passage times matrix N is given by $K_W = \boldsymbol{\pi}^T \boldsymbol{k}_W$, where $\boldsymbol{k}_W = N\boldsymbol{\pi}$ and the $i-$th entry $\boldsymbol{k}_{W,i}$ in $\boldsymbol{k}_W$ denotes the mean time to go from v_i to any other node. We first provide an analytic expression for the vector $\boldsymbol{k}_W \in \mathbb{R}^{n \times 1}$.

Lemma 9: (First passage times for a doubly-weighted graph): For a graph $\mathcal{G} = (\mathcal{V}, E, P, W)$ with stationary distribution $\boldsymbol{\pi}$ and associated first passage time matrix N, the following equality holds.

$$(I - P)\boldsymbol{k}_W = (P \circ W)\mathbb{1}_n - \mathbb{1}_n \pi^T (P \circ W)\mathbb{1}_n,$$

where $\boldsymbol{k}_W = N\boldsymbol{\pi}$.

This gives the insight that, in general, $\boldsymbol{k}_{W,i} \neq \boldsymbol{k}_{W,j}$ on the doubly-weighted graph, unlike the counterpart for the single-weighted graph (where $\boldsymbol{k}_i = K$ for all $i \in \{1, \ldots, n\}$). Interestingly enough however, there does exist a relation between the mean first passage time K_W and the Kemeny constant K as is stated in the following theorem.

Theorem 10: (Mean first passage time of doubly-weighted graphs): For the doubly-weighted graph $\mathcal{G} = (\mathcal{V}, E, P, W)$, the mean first passage time K_W of a vehicle on the graph is given by

$$K_W = \boldsymbol{\pi}^T (P \circ W) \mathbb{1}_n K, \tag{6}$$

where K is the Kemeny constant associated with the transition matrix P whose stationary distribution is $\boldsymbol{\pi}$.

Remark 11: The expected number of *hops* to go from one region to another in a Markov chain with transition matrix P is its Kemeny constant. The expected distance travelled (and hence time taken) executing one *hop* is $\sum_i \pi_i \sum_j p_{ij}\omega_{ij} = \boldsymbol{\pi}(P \circ W)\mathbb{1}_n$. Hence, it fits with intuition that the expected time taken to go from one region to another should be $K\boldsymbol{\pi}^T(P \circ W)\mathbb{1}_n$ as was formally shown above.

Problem 12: (Optimizing mean first passage time of a reversible Markov chain): Given the stationary distribution $\boldsymbol{\pi}$ and graph $\mathcal{G}$ with vertex set $\mathcal{V}$, edge set E and weight matrix W, determine the transition probabilities $P = [p_{ij}]$ solving:

$$\begin{aligned}
&\text{minimize} \\
&\quad \left(\boldsymbol{\pi}^T (P \circ W)\mathbb{1}_n\right) \left(\mathbf{Tr}\left[(I - \Pi^{1/2} P \Pi^{-1/2} + \boldsymbol{q}\boldsymbol{q}^T)^{-1}\right]\right) \\
&\text{subject to} \quad \sum_{j=1}^{n} p_{ij} = 1, \text{ for each } i \in \{1, \ldots, n\} \\
&\qquad \boldsymbol{\pi}_i p_{ij} = \boldsymbol{\pi}_j p_{ji}, \text{ for each } (i,j) \in E \\
&\qquad 0 \leq p_{ij} \leq 1, \text{ for each } (i,j) \in E \\
&\qquad p_{ij} = 0, \text{ for each } (i,j) \notin E.
\end{aligned}$$

Theorem 13 (Convexity of Problem 12): Given the $\mathcal{G}$ with vertex set $\mathcal{V}$, edge set E and weight matrix W, let $\mathcal{P}_{\mathcal{G},\boldsymbol{\pi}}$ denote the set of matrices associated with $\mathcal{G}$ that are irreducible reversible Markov chains with stationary distribution $\boldsymbol{\pi}$. Then, $\mathcal{P}_{\mathcal{G},\boldsymbol{\pi}}$ is a convex set and $P \mapsto \boldsymbol{\pi}^T(P \circ W)\mathbb{1}_n K(P)$ is a convex function over $\mathcal{P}_{\mathcal{G},\boldsymbol{\pi}}$.

B. SDP framework for optimizing the mean first passage time of a doubly-weighted graph

In a similar fashion to Problem 3, we can formulate Problem 12 as a SDP by introducing the symmetric slack matrix $X \in \mathbb{R}^{n \times n}$ and the scalare variable t as shown in the following.

Problem 14: (Optimizing mean first passage time of a reversible Markov chain (SDP)): Given the stationary distribution $\boldsymbol{\pi}$ and graph $\mathcal{G}$ with vertex set $\mathcal{V}$, edge set E and weight matrix W, determine $Q = [q_{ij}]$, X and t solving:

$$\begin{aligned}
&\text{minimize} \quad \mathbf{Tr}[X] \\
&\text{subject to} \\
&\qquad \begin{bmatrix} t(I + \boldsymbol{q}\boldsymbol{q}^T) - \Pi^{1/2} Q \Pi^{-1/2} & I \\ I & X \end{bmatrix} > 0 \\
&\qquad Q\mathbb{1}_n = t\mathbb{1}, \\
&\qquad \boldsymbol{\pi}_i q_{ij} = \boldsymbol{\pi}_j q_{ji}, \text{ for each } (i,j) \in E \\
&\qquad 0 \leq q_{ij} \leq t, \text{ for each } (i,j) \in E \\
&\qquad q_{ij} = 0, \text{ for each } (i,j) \notin E \\
&\qquad \boldsymbol{\pi}^T (Q \circ W)\mathbb{1}_n = 1 \\
&\qquad t \geq 0
\end{aligned}$$

Then, the transition matrix P is given by $P = Q/t$.

Again the first inequality constraint represents an LMI. Before noting when the LMI holds first note that the constraints in Problem 14 imply that $Pt = Q$ and that $t = \frac{1}{\boldsymbol{\pi}^T(P \circ W)\mathbb{1}_n}$. Hence, the LMI constraint holds true if and only if $X - \boldsymbol{\pi}^T(P \circ W)\mathbb{1}_n(I - \Pi^{1/2} P \Pi^{-1/2} + \boldsymbol{q}\boldsymbol{q}^T)^{-1}$ is positive definite therefore the SDP given by Problem 14 minimizes the mean first passage time for doubly-weighted graphs.

V. Interpretations of the Mean First Passage Time for Doubly-Weighted Graphs

The results on mean first passage time for doubly-weighted graphs presented in this work provide a general framework which can potentially be applied to the analysis and design in a myriad of fields. We highlight three in particular in the following, where the first two are inspired by our initial motivation of detecting an intruder in a network. All the scenarios assume we are working on the doubly-weighted graph described by $\mathcal{G} = (\mathcal{V}, E, P, W)$.

a) Environment Service/Surveillance

In practical service/surveillance scenarios, there is often a desire to service some regions more than others while simultaneously minimizing the time any one region has to wait before it is serviced. In this formulation, in every step of the random walk, the vehicle moves to a new region and executes its service task. For region $v_i \in \mathcal{V}$ denote the service time as s_i and denote the time to move from region v_i to region v_j by d_{ij}, where we can assume $d_{ii} = 0$. Then, we can define the weight corresponding to the edge from v_i to v_j as $\omega_{ij} = d_{ij} + s_j$. Optimizing the transition matrix P corresponds to minimizing the average time a region has to wait before getting serviced, regardless of what region the vehicle starts its service task in.

b) Quickest Intruder Detection

Quickest intruder detection is closely related to environment service task and given the appropriate setup, it can be the same. If, for example, there is a persistent intruder in the environment and it takes s_i time to determine if an intruder is in the environment, then minimizing the mean first passage time optimizes the time it takes to capture an intruder.

c) Social Networks

In the social network scenario it is often possible to distinguish different sub-groups in the network (i.e., people who have a certain political view over another). If we let the weighted edges ω_{ij} represent the influence group i has on group j, and let p_{ij} denote the probability that two groups talk to each other, then one could study how changes in influences ω_{ij} affect information propagation in the network. Alternatively, one could use a similar concept to force search results in some optimal way such that it minimizes the time for information to spread through the network regardless of where it is initially planted.

VI. Application to Robotic Surveillance

In this section we demonstrate through simulation how using the mean first passage time is advantageous for the design of surveillance policies over other well known algorithms. We look at both homogeneous and heterogeneous environment cases.

In the first scenario, we work on the discretized environment shown in Figure 2. We assume a set of intruders are distributed throughout the environment. Each intruder performs a malicious activity in its host region for a fixed duration of time, which we call the *intruder life-time*, followed instantaneously by another intruder. The *intruder life-time* is set to 5 time units and 20 intruders appear per simulation run (the sequence in which the intruders appear is determined before each simulation run), for a total simulation run of 100 time units. We also assume that a single surveillance vehicles executes a random walk in the environment, spending 1 time unit in each region, and that the vehicle transition between two connected regions instantaneously. The intruder is caught if the surveillance vehicle is in the same region as the intruder for any portion of the *intruder life-time*. To compare performance we generate a random walk for the environment using the Metropolis-Hastings, fastest mixing Markov chain (FMMC), and Kemeny constant algorithms, where we assume a stationary distribution on the transition matrix (each node in the region must be visited with equal likelihood). Figure 3 shows a representative example of how each algorithm performs for one simulation run. The red squares correspond to when a intruder and vehicle are in the same region at the same time (when a intruder is captured). As can be seen, the Kemeny constant algorithm captures intruders more frequently than the other two algorithms. Table I summarizes the statistical performance of each algorithm after 100 runs of the simulation and justifies our use of the Kemeny constant algorithm as an valid surveillance strategy; the Kemeny constant algorithm captures intruders more frequently than the other two algorithms, and its worst case performance is still better than the worst case performance of the other two algorithms. .

In the second scenario, we work with the discretized environment shown in Figure 4. In this environment the time taken by the vehicle to travel an edge is no longer instantaneous and is equal to the edge's length, where edge lengths are non-uniform. Once in a new region, the vehicle

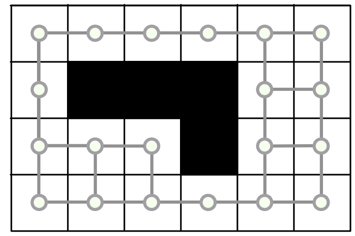

Fig. 2. Environment represented by a graph.

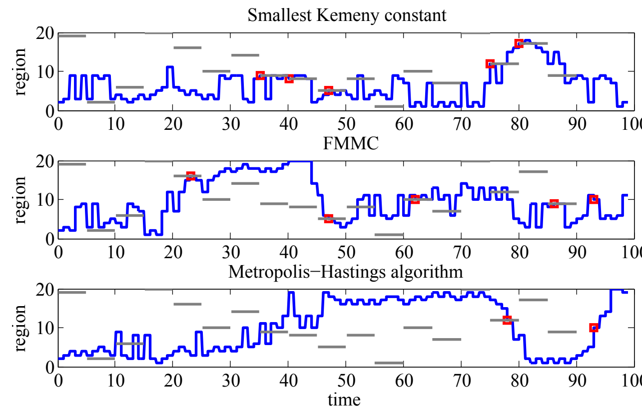

Fig. 3. Performance of three Markov chain design algorithms. The solid lines indicate the vehicle trajectory along the graph, gray dots indicate intruder location with time, and the squares indicate intruder detection.

is required to spend 0.2 time units examining the region for malicious activities. We again assume there are a set of intruders with behavior as described in the first scenario. We again also assume a intruder is caught if the surveillance vehicle is in the same region as a intruder for any portion of the *intruder life-time*, but now set the *intruder life-time* to 10 time units with a intruder appearing 50 times (total of 500 time units per simulation run). Since the design of the FMMC and Metropolis-Hastings algorithms do not inherently account for non-uniform travel weights and service times, we also compare the performance of the random walk generated by the mean first-passage time algorithm against a standard traveling salesman problem (TSP) tour. Figure 5 shows a representative example of how each algorithm performs for one simulation run and Table II summarizes the statistical performance of each algorithm after 100 runs of the simulation. The mean first-passage time algorithm's performance compared to the other two stochastic policies in this scenario is much better than what was seen in the first scenario. This is to be expected since the other two stochastic polices do not account for heterogeneous travel/service times on the graph. However, it is clear that every stochastic algorithm outperforms the TSP tour. This lack of performance in the TSP tour is most likely attributed to the fact that each region is only visited once per tour.

Remark 15: The surveillance task described in the above simulations involves intruders which disappear after a specified *intruder life-time*. However, if an intruder stays fixed in some region until it is detected, then the expected time

Algorithm	Min	Mean	Max	StdDev	K
Kemeny constant	5	11.38	17	2.07	51.99
FMMC	4	10.81	16	2.15	58.03
Metropolis-Hastings	4	11.18	17	2.23	55.80

TABLE I

STATISTICS ON THE TOTAL NUMBER OF INTRUDERS CAUGHT IN 100 SIMULATION RUNS FOR THE ENVIRONMENT IN FIG. 2.

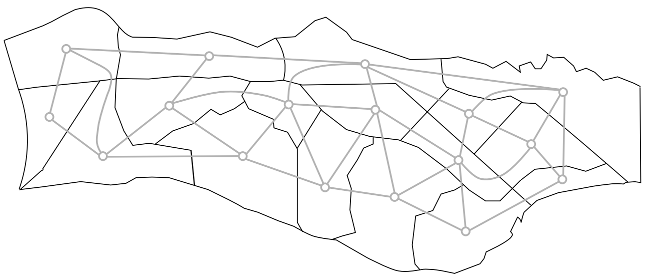

Fig. 4. Environment represented by a doubly-weighted graph. The weight of an edge is its length.

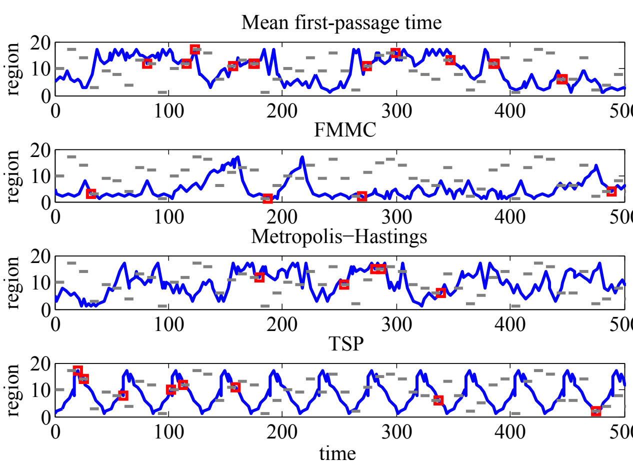

Fig. 5. Performance of mean first-passage time, FMMC, Metropolis-Hastings algorithm and the TSP tour. The solid lines indicate the vehicle trajectory along the graph, gray dots indicate intruder location with time, and the squares indicate intruder detection.

taken by a surveillance vehicle to detect such an intruder is simply the mean first-passage time associated with its random walk. Also, if restricted to reversible Markov chains, then the minimum mean-first passage time is the provably optimal solution for the fixed intruder surveillance problem.

VII. CONCLUSIONS

We have studied the problem of how to optimally design the surveillance policy which minimizes the mean first passage time to go from one region to any other region in a connected environment. We have presented the first formulation of the mean first passage time for a doubly-weighted graph and have also provided a convex formulation for the design of the mean first passage time for both single-weighted and doubly-weighted graph topologies. Finally, we have shown that both problems can be written as SDPs

Algorithm	Min	Mean	Max	StdDev	K_W
Mean first passage time	3	7.00	11	2.50	82.16
FMMC	2	5.27	10	2.11	110.66
Metropolis-Hastings	2	6.09	10	1.54	86.85
TSP tour	1	3.18	6	1.56	

TABLE II

STATISTICS ON THE TOTAL NUMBER OF INTRUDERS CAUGHT IN 200 SIMULATION RUNS FOR THE ENVIRONMENT IN FIG. 4.

and have demonstrated the effectiveness of our surveillance policy against other well known Markov chain policies.

This work leaves various extensions open for further research. First, this policy looks at the design of one vehicle's trajectory in an environment, it would be of practical interest to consider the case where there are multiple vehicles: [4] is an example of work in this direction. Second, it would be useful to understand bounds on the design of of the mean first passage time for general graph topologies. Finally, we treat only the optimization of the transition matrix of the graph. It would be of interest to study how we can optimize the weight matrix W in conjunction with the transition matrix. This can have the interpretation of optimizing the "capacity" or "resistance" of the graph, a topic in optimization which is of independent interest [6].

REFERENCES

[1] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The role of social networks in information diffusion. In *Int. Conference on World Wide Web*, pages 519–528, Lyon, France, 2012. ACM.

[2] T. H. Blackwell and J. S. Kaufman. Response time effectiveness: Comparison of response time and survival in an urban emergency medical services system. *Academic Emergency Medicine*, 9(4):288–295, 2002.

[3] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004.

[4] G. Cannata and A. Sgorbissa. A minimalist algorithm for multirobot continuous coverage. *IEEE Transactions on Robotics*, 27(2):297–312, 2011.

[5] L. Chen and J. Leneutre. A game theoretical framework on intrusion detection in heterogeneous networks. *IEEE Transactions on Information Forensics and Security*, 4(2):165–178, June 2009.

[6] A. Ghosh, S. Boyd, and A. Saberi. Minimizing effective resistance of a graph. *SIAM Review*, 50(1):37–66, 2008.

[7] J. Grace and J. Baillieul. Stochastic strategies for autonomous robotic surveillance. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 2200–2205, Seville, Spain, December 2005.

[8] J. J. Hunter. *Mathematical Techniques of Applied Probability*, volume 1 of *Discrete Time Models: Basic Theory*. Academic Press, 1983.

[9] J. J. Hunter. The role of Kemeny's constant in properties of Markov chains. *Communications in Statistics - Theory and Methods*, 2013. To appear.

[10] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer, 1976.

[11] S. Kirkland. Fastest expected time to mixing for a Markov chain on a directed graph. *Linear Algebra and its Applications*, 433(11-12):1988–1996, 2010.

[12] M. Levene and G. Loizou. Kemeny's constant and the random surfer. *The American Mathematical Monthly*, 109(8):741–745, 2002.

[13] Y. Y. Li and L. E. Parker. Intruder detection using a wireless sensor network with an intelligent mobile robot response. In *IEEE Southeastcon*, pages 37–42, Huntsville, AL, USA, April 2008.

[14] F. Pasqualetti, A. Franchi, and F. Bullo. On optimal cooperative patrolling. In *IEEE Conf. on Decision and Control*, pages 7153–7158, Atlanta, GA, USA, December 2010.

[15] F. Pasqualetti, A. Franchi, and F. Bullo. On cooperative patrolling: Optimal trajectories, complexity analysis and approximation algorithms. *IEEE Transactions on Robotics*, 28(3):592–606, 2012.

[16] A. Patcha and J. Park. A game theoretic approach to modeling intrusion detection in mobile ad hoc networks. In *Information Assurance Workshop, 2004. Proceedings from the Fifth Annual IEEE SMC*, pages 280–284, June 2004.

[17] K. Srivastava, D. M. Stipanoviȅ, and M. W. Spong. On a stochastic robotic surveillance problem. In *IEEE Conf. on Decision and Control*, pages 8567–8574, Shanghai, China, December 2009.

[18] V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *International Journal of Robotics Research*, 32(12):1438–1458, 2013.

[19] X. Wu and Z. Liu. How community structure influences epidemic spread in social networks. *Physica A: Statistical Mechanics and its Applications*, 387(2):623–630, 2008.

[20] P. R. Wurman, R. D'Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–20, 2008.