

GIBBS SAMPLER-BASED PATH PLANNING FOR AUTONOMOUS VEHICLES: CONVERGENCE ANALYSIS

Wei Xi * Xiaobo Tan ** John S. Baras *

* *Department of Electrical & Computer Engineering, and
Institute for Systems Research
University of Maryland, College Park, MD 20742, USA
Email: {wxi,baras}@isr.umd.edu*

** *Department of Electrical & Computer Engineering
Michigan State University, East Lansing, MI 48824, USA
Email: xbtan@msu.edu*

Abstract: Simulation has indicated that distributed self-organization of autonomous swarms might be achieved through Gibbs sampler-based simulated annealing. However, the dynamic nature of the underlying graph presents significant challenges in convergence analysis. As a first step toward such analysis, convergence of the algorithm is established in this paper for two special cases: single vehicle with limited sensing/moving range, and multiple vehicles with full sensing/moving range. The impact of Gibbs potential functions on the convergence speed is also investigated, which provides insight into the design of these functions.
Copyright© 2005 IFAC.

Keywords: Markov random fields; Gibbs sampler; Decentralized control; Autonomous vehicles; Convergence

1. INTRODUCTION

With the rapid development of sensing, communication, computation, and actuation capabilities, autonomous unmanned vehicles (AUVs) are expected to cooperatively perform dangerous or explorative tasks in various hazardous, unknown or remote environments (Schoenwald, 2000). Distributed approaches to control of individual vehicles and to coordination among vehicles are especially appealing considering the large scale of the vehicle networks and the bandwidth constraint on communication (Olfati-Saber and Murray, 2002; Jadbabaie *et al.*, 2003; Tanner *et al.*, 2003; Baras *et al.*, 2003). Such approaches often involve artificial potentials (Shahidi *et al.*, 1991; Rimón and Koditschek, 1992; Leonard and Fiorelli, 2001), and vehicles follow the negative gradients of po-

tentials mimicking the emergent behaviors (e.g. foraging) demonstrated by swarms of bacteria, insects, and animals (Passino, 2002). In these schemes, all interactions among vehicles, target attraction, and collision avoidance are encoded in potential functions. Despite the simple nature, an essential problem in such approaches is that the system dynamics may get trapped at local minima of the potential functions (Koren and Borenstein, 1991).

To deal with the above problem, a stochastic approach was proposed recently based on the theory of Markov Random Fields (MRFs) (Baras and Tan, 2004). In this approach a swarm of vehicles is treated as an MRF on a graph where each vehicle is a node and their communication/detection links constitute the edges. Similar to the artificial po-

tential approach, global objectives (e.g., reaching the target or making formations) and constraints (e.g., obstacles or threats) are reflected through the design of Gibbs potentials. The movement of each vehicle is then decided through simulated annealing based on the Gibbs sampler. With this scheme, successful self-organization and mission execution with only local sensing and communication have been observed (Baras and Tan, 2004).

It is of interest to analyze the convergence behavior of the stochastic algorithm since this may not only provide theoretical explanation to the simulation results, but also offer insight into the proper design of the Gibbs potentials. However, the dynamic nature of the underlying graph presents significant challenges in analysis and classical results on MRF theory do not apply directly. As a first step toward analytical understanding of the algorithm behavior, two special cases are investigated in this paper: the case of single-vehicle with limited sensing/moving range and the case of multi-vehicle with full sensing/moving range. For both cases it is established that the vehicle(s) will achieve the global objective defined through the Gibbs potential function. Furthermore, for the single-vehicle case, the relationship between the convergence speed and the potential function is derived analytically and then verified through simulation. This can have a direct impact on the optimal design of potential functions in practice.

The remainder of the paper is organized as follows. The Gibbs sampler-based path planning algorithm is first reviewed in Section 2. Convergence analysis for two special cases is presented in Section 3, followed by discussion on the convergence speed in Section 4. Section 5 concludes the paper.

2. GIBBS SAMPLER-BASED PATH PLANNING ALGORITHM

2.1 MRFs and Gibbs Sampler

Let S be a finite set of cardinality σ , with elements indexed by s and called *sites*. For $s \in S$, let Λ_s be a finite set called the *phase space* for site s . A *random field* on S is a collection $X = \{X_s\}_{s \in S}$ of random variables X_s taking values in Λ_s . A *configuration* of the system is $x = \{x_s, s \in S\}$ where $x_s \in \Lambda_s, \forall s$. The product space $\Lambda_1 \times \dots \times \Lambda_\sigma$ is called the *configuration space*. A *neighborhood system* on S is a family $\mathcal{N} = \{\mathcal{N}_s\}_{s \in S}$, where $\forall s, r \in S, \mathcal{N}_s \subset S, s \notin \mathcal{N}_s$, and $r \in \mathcal{N}_s$ if and only if $s \in \mathcal{N}_r$. \mathcal{N}_s is called the *neighborhood* of site s . The random field X is called a *Markov random field* (MRF) with respect to the neighborhood system \mathcal{N} if, $\forall s \in S, P(X_s = x_s | X_r = x_r, r \neq s) = P(X_s = x_s | X_r = x_r, r \in \mathcal{N}_s)$.

A random field X is a *Gibbs random field* if and only if it has the Gibbs distribution:

$$P(X = x) = \frac{e^{-\frac{U(x)}{T}}}{Z}, \forall x,$$

where T is the temperature variable (widely used in simulated annealing algorithms), $U(x)$ is the potential (or energy) of the configuration x , and Z is the normalizing constant, called the *partition function*: $Z = \sum_x e^{-\frac{U(x)}{T}}$. One then considers the following useful class of potential functions $U(x) = \sum_{s \in \Lambda} \Phi_s(x)$, which is a sum of individual contributions Φ_s evaluated at each site. The Hammersley-Clifford theorem (Bremaud, 1999) establishes the equivalence of a Gibbs random field and an MRF.

The *Gibbs sampler* belongs to the class of *Markov Chain Monte Carlo* (MCMC) methods, which sample Markov chains leading to stationary distributions. The algorithm updates the configuration by visiting each site sequentially (or randomly visiting sites with certain *proposal* distribution (Winkler, 1995)) and sampling from the local specifications of a Gibbs field. One round of sequential visits to all sites is called a *sweep*. The convergence of the Gibbs sampler was studied by D. Geman and S. Geman in the context of image processing (Geman and Geman, 1984). There it was shown that as the number of sweeps goes to infinity, the distribution of $X(n)$ converges to the Gibbs distribution Π . Furthermore, with an appropriate cooling schedule, simulated annealing using the Gibbs sampler yields a uniform distribution on the set of minimizers of $U(x)$. Thus the global objectives could be achieved through appropriate design of the Gibbs potential function.

2.2 Path Planning: Problem Setup

Consider a 2D mission space (the extension to 3D space is straightforward), which is discretized into a lattice of square cells. Without loss of generality, each cell is assumed to have unit dimensions. Label each cell with its coordinates (i, j) , where $1 \leq i \leq N_x, 1 \leq j \leq N_y$, for $N_x, N_y > 0$. There is a set of vehicles (or *mobile nodes*) S indexed by $s = 1, \dots, \sigma$ on the mission space. To be precise, each vehicle (node) s is assumed to be a point mass located at the center of some cell (i_s, j_s) , and the position of vehicle s is taken to be $p_s = (i_s, j_s)$. At most one vehicle is allowed to stay in each cell at any time instant.

The distance between two cells, (i_a, j_a) and (i_b, j_b) , is defined as $R \triangleq \sqrt{(i_a - i_b)^2 + (j_a - j_b)^2}$. Each node s can communicate with other nodes that lie within a communication/sensing range R_s . These nodes form the set \mathcal{N}_s of *neighbors* of node s . A vehicle can go at most R_m ($R_m \leq R_s$) within

one move. For ease of discussion, $R_m = R_s$ is assumed in this paper. There might be multiple *obstacles* in the space, where an obstacle is defined to be a set of adjacent cells that are inaccessible to vehicles. For instance, a “circular” obstacle centered at (i^o, j^o) with radius R_o can be defined as $O \triangleq \{(i, j) : \sqrt{(i - i^o)^2 + (j - j^o)^2} \leq R_o\}$. A node can detect obstacle cells within R_s . There can be at most one *target area* in the space. A target area is a set of adjacent cells that represent desirable destinations of mobile nodes. A “circular” target area can be defined similarly as a “circular” obstacle. Fig. 1 shows an example scenario, containing an obstacle (formed by two overlapping circular obstacles) and a circular target.




Fig. 1. An example mission scenario.

The neighborhood system defined earlier naturally leads to a dynamic graph where each vehicle stands for a vertex of the graph and the neighborhood relation prescribes the edges between vehicles. An MRF can then be defined on the graph, where each vehicle s is a site and the associated phase space Λ_s is the set of all cells located within the moving range R_m from location p_s and not occupied by obstacles or other vehicles. An *admissible configuration* is a graph where vehicles stay in different cells. The *accessible area* is the set of cells in the mission space that are not occupied by obstacles. An *accessible-area graph* can then be induced by letting each cell in the accessible area be a vertex and connecting neighboring cells with edges. The accessible area is *connected* if the associated accessible-area graph is connected.

The Gibbs potential $U(x) = \sum_s \Phi_s(x)$, where $\Phi_s(x)$ consists of three terms with each reflecting one goal or one constraint. To be specific,

$$\Phi_s(x) = \lambda_g J_s^g + \lambda_o J_s^o + \lambda_n J_s^n, \quad (1)$$

where J_s^g, J_s^o , and J_s^n account for the attraction from the target area, the repelling from obstacles, and the reaction force from neighbors (could be attractive or repelling depending on the application/distance), respectively, and they can be designed so that $\Phi_s(x)$ depends only on x_s and x_r , $r \in \mathcal{N}_s$. $\lambda_g, \lambda_o, \lambda_n$ are the weighting coefficients for adjusting the potential surface. Note that the

design of these constants is also a challenging and important issue as it may directly impact the nodes behavior and the convergence rate of the algorithm. The movement of each vehicle is then determined by the Gibbs sampler-based simulated annealing (Baras and Tan, 2004).

There are important differences between a classical MRF introduced in Subsection 2.1 and the MRF defined for the vehicle networks. In a classical MRF, both the phase space Λ_s and the neighborhood \mathcal{N}_s are time-invariant; however, for a vehicle network, both Λ_s and \mathcal{N}_s depend on the dynamic graph and therefore vary with time. These differences prevent the classical MRF theory from being adopted directly to analyze the convergence behavior of the path planning algorithm.

3. CONVERGENCE ANALYSIS

3.1 Single Vehicle with Limited Moving Range

For the single-vehicle case, the MRF has only one site and the individual potential Φ_s consists of only first two terms of (1), which makes the analysis simpler. It can be shown that the Markov chain of the MRF defined by the Gibbs sampler converges to a stationary distribution Π_T as time tends to infinity for a fixed temperature T . By choosing an appropriate cooling schedule, the vehicle approaches the configuration(s) of minimum potential. This is made precise by the following theorems. The proofs of the theorems are adapted from some results on classical MRFs (Winkler, 1995).

Theorem 3.1. Assume that the accessible area is connected. For a fixed temperature T , let P_T be the kernel of the homogeneous Markov chain $X(n)$ defined by the Gibbs sampler. Then $X(n)$ has a unique stationary distribution Π_T :

$$\pi_T(x) = \frac{e^{-\frac{U(x)}{T}} \sum_{\|p_z - p_x\| \leq R_m} e^{-\frac{U(z)}{T}}}{Z_T}, \quad (2)$$

where Z_T is the partition function, defined as:

$$Z_T = \sum_{x \in X} \left(e^{-\frac{U(x)}{T}} \sum_{\|p_z - p_x\| \leq R_m} e^{-\frac{U(z)}{T}} \right)$$

Furthermore, starting from any distribution ν

$$\lim_{n \rightarrow \infty} \nu P^n = \Pi_T. \quad (3)$$

Proof. Since the MRF has only one site, the Markov kernel P_T defined by Gibbs sampler is:

$$P_T(x, y) = \begin{cases} \frac{e^{-\frac{U(y)}{T}}}{\sum_{\|p_z - p_x\| \leq R_m} e^{-\frac{U(z)}{T}}} & \text{if } \|p_y - p_x\| \leq R_m \\ 0 & \text{if } \|p_y - p_x\| > R_m \end{cases}. \quad (4)$$

Due to the connectivity of the accessible area, there exists at least one path between any two configurations x and y (i.e., a sequence of moves $\{x, x_1, \dots, y\}$), and the shortest path is bounded by τ moves for some finite τ . This implies that P_T has a *strictly positive power* P_T^τ , i.e., the τ -step Markov chain reaches each state with positive probability from any state. Hence the Markov chain is ergodic and it has a unique invariant distribution Π_T , which implies (3). One can verify that (2) is a stationary distribution for P_T . \square

Due to the limited moving range, τ sweeps will be performed for each temperature $T(n)$ in the simulated annealing to guarantee the convergence.

For a Markov kernel P , define its *contraction coefficient* $c(P)$ by

$$c(P) = (1/2) \max_{x,y} \|P(x, \cdot) - P(y, \cdot)\|_1,$$

where $P(x, \cdot)$ denotes the vector of conditional distributions $p(\cdot|x)$. The following lemma (Winkler, 1995) will be useful in the proof of Theorem 3.2.

Lemma 3.1. Let μ and ν be probability distributions, and P and Q be Markov kernels. Then

$$\begin{aligned} \|\mu P - \nu P\|_1 &\leq c(P) \|\mu - \nu\|_1, \\ c(PQ) &\leq c(P)c(Q), \end{aligned}$$

$$c(P) \leq 1 - |X| \min\{P(x, y) : x, y \in X\},$$

and for a primitive P ,

$$c(P^n) \rightarrow 0 \text{ as } n \rightarrow \infty,$$

where X is the state space of the Markov chain and $|X|$ denotes its cardinality.

Theorem 3.2. Assume that the accessible area is connected. Let $\{T(n)\}_{n \geq 1}$ be a cooling schedule decreasing to 0 such that eventually

$$T(n) \geq \frac{\tau \Delta}{\ln n} \quad (5)$$

where $\Delta = \max_{x,y} \{ |U(x) - U(y)| : \|p_x - p_y\| \leq R_m \}^1$, and τ is as defined in the proof of Theorem 3.1. Let $Q_n = P_{T(n)}^\tau$. Then from any initial distribution ν ,

$$\lim_{n \rightarrow \infty} \nu Q_1 \cdots Q_n = \Pi_\infty, \quad (6)$$

where Π_∞ is the distribution (2) evaluated at $T = 0$. Let M be the set of configurations achieving the minimum of $U(x)$. Assume that $\|p_x - p_y\| > R_m$, $\forall x, y \in M, x \neq y$. Then

$$\lim_{n \rightarrow \infty} \nu Q_1 \cdots Q_n = \begin{cases} \frac{1}{|M|} & \text{if } x \in M \\ 0 & \text{otherwise} \end{cases}. \quad (7)$$

¹ In this paper, $\|\cdot\|$ denotes the Euclidean norm, and $\|\cdot\|_1$ denotes the 1-norm.

Proof. Define $m_x = \min\{U(z) : \|p_z - p_x\| \leq R_m\}$. Then, if $\|p_y - p_x\| \leq R_m$,

$$P_T(x, y) = \frac{e^{-\frac{U(y)-m_x}{T}}}{\sum_{\|p_z - p_x\| \leq R_m} e^{-\frac{U(z)-m_x}{T}}} \geq |\bar{\Lambda}|^{-1} e^{-\frac{\Delta}{T}} \quad (8)$$

where $|\bar{\Lambda}|$ is the upper-bound on the cardinality of the phase space. For $Q = P_T^\tau$, from Lemma 3.1,

$$\begin{aligned} c(Q) &\leq 1 - |X| \min_{x,y} Q(x, y) \leq 1 - |X| (\min_{x,y} P_T(x, y))^\tau \\ &\leq 1 - \lambda e^{-\frac{\tau \Delta}{T}}, \end{aligned} \quad (9)$$

where $\lambda = |X| \cdot |\bar{\Lambda}|^{-\tau} \leq 1$. Then

$$\begin{aligned} \prod_{k=i}^n c(Q_k) &\leq \prod_{k=i}^n (1 - \lambda e^{-\frac{\tau \Delta}{T(k)}}) \leq \prod_{k=i}^n (1 - \lambda k^{-1}) \\ &\leq e^{-\sum_{k=i}^n \lambda k^{-1}} \leq \left(\frac{i}{n}\right)^\lambda \end{aligned} \quad (10)$$

where the third inequality holds because

$$(1 - z) \leq e^{-z}, \forall z \in (0, 1),$$

and the last one follows by

$$\begin{aligned} \ln(ni^{-1}) &< \ln(n+1) - \ln(i) = \sum_{k=i}^n (\ln(1+k) - \ln k) \\ &= \sum_{k=i}^n \ln(1+k^{-1}) \leq \sum_{k=i}^n k^{-1}. \end{aligned}$$

With some abuse of notation, $\Pi_{T(n)}$ will be written as Π_n to simplify the expressions. For $i \leq n$, one has

$$\begin{aligned} &\|\nu Q_1 \cdots Q_n - \Pi_\infty\|_1 \\ &= \|(\nu Q_1 \cdots Q_{i-1} - \Pi_\infty) Q_i \cdots Q_n + \Pi_\infty Q_i \cdots Q_n - \Pi_\infty\|_1 \\ &\leq \|(\nu Q_1 \cdots Q_{i-1} - \Pi_\infty)\|_1 c(Q_i \cdots Q_n) + \|\Pi_\infty Q_i \cdots Q_n - \Pi_\infty\|_1 \\ &\leq 2c(Q_i \cdots Q_n) + \|\Pi_i - \Pi_\infty\|_1 + \|\Pi_n - \Pi_\infty\|_1 \\ &\quad + \sum_{k=i}^{n-1} \|\Pi_k - \Pi_{k+1}\|_1, \end{aligned} \quad (11)$$

where the last inequality is established by Dobrushin's theorem (Winkler, 1995). From (10), the first term of the last inequality approaches 0 as $n \rightarrow \infty$. By Theorem (3.1), the second and third terms also approach 0 as $i, n \rightarrow \infty$. To show that the fourth term vanishes as $i \rightarrow \infty$, it suffices to prove each component sequence of $(\Pi_n(x))_{n \geq 1}$ decreases or increases eventually. By investigating the derivative of $\Pi_n(x)$ with respect to the temperature T , the ultimate monotonicity can be established (Winkler, 1995). This completes the proof of (6).

Let m denote the (global) minimal value of Gibbs potential $U(x)$. Then

$$\pi_T(x) = \frac{e^{-\frac{U(x)-m}{T}} \sum_{\|p_z - p_x\| \leq R_m} e^{-\frac{U(z)-m}{T}}}{\sum_{y \in X} (e^{-\frac{U(y)-m}{T}} \sum_{\|p_z - p_y\| \leq R_m} e^{-\frac{U(z)-m}{T}})}$$

If $U(x)$ or $U(y)$ is a minimum then the respective exponential equals to 1 no matter what temperature T may be. The other exponentials decrease to 0 as T tends to 0. Eq. (7) then follows from (6) and $\|p_x - p_y\| > R_m, \forall x, y \in M$. \square

Note one can easily characterize Π_∞ when more than one minimizers of $U(x)$ are located within a distance of R_m . Theorem 3.2 establishes that under the algorithm, a vehicle can reach the target ultimately for *arbitrarily shaped* obstacles.

3.2 Multiple Vehicles with Full Moving Range

When each vehicle can sense over the whole mission space and can reach any cell within one move, one can verify that the Gibbs distribution (for a fixed temperature T) and the Markov kernel satisfy the *detailed balance equation* (Winkler, 1995). As the Markov kernel (for a complete sweep) is strictly positive, the Gibbs distribution is indeed the unique stationary distribution. Furthermore, the convergence of the simulated annealing algorithm can be established. These are made precise in the following propositions:

Proposition 3.1. Assume that the accessible area is connected. For a fixed temperature T , the corresponding Markov chain $X(n)$ has a unique stationary distribution $\Pi_T = \frac{e^{-\frac{U(x)}{T}}}{Z}$. From any initial distribution μ , $\lim_{n \rightarrow \infty} \nu P_T^n = \Pi_T$.

Proposition 3.2. Let $(T(n))_n \geq 1$ be a cooling schedule satisfying $T(n) \geq \frac{\sigma \Delta}{\ln n}$, where σ is the number of vehicles, and Δ is the maximum local oscillation (Winkler, 1995) of $U(\cdot)$. Then from any initial distribution μ ,

$$\lim_{n \rightarrow \infty} \nu P_1 \dots P_n = \begin{cases} \frac{1}{|M|} & \text{if } x \in M \\ 0 & \text{otherwise} \end{cases}.$$

4. POTENTIAL FUNCTIONS AND CONVERGENCE SPEED

In Theorem 3.2 a condition on the cooling schedule $T(n)$ is specified to guarantee convergence. In the condition the maximal energy difference Δ among neighboring configurations plays an important role. It is of interested to study how to improve the convergence speed by appropriately

designing the potential functions. The study here will be focused on the single-vehicle case.

Proposition 4.1. For the single-vehicle case, the convergence speed is characterized by

$$\|\nu Q_1 \dots Q_n - \Pi_\infty\| = O(n^{-\frac{2\lambda \tilde{m}}{2\tilde{m} + \lambda \Delta \tau}}) \quad (13)$$

where $\tilde{m} = \min_{y \notin M} (U(y) - m)$, and m, λ, Δ , and τ are defined as in Theorem 3.2.

Proof. From (12), $\|\nu Q_1 \dots Q_n - \Pi_\infty\|$ is bounded by four terms. From (10), the first term is bounded by $(\frac{i}{n})^\lambda$. The other four terms are each bounded by $\|\Pi_i - \Pi_\infty\|_1$ considering the ultimate monotonicity of the sequence $\{\Pi_n\}$.

Take $T(i) = \frac{\tau \Delta}{\ln i}$. If x is not a minimizer of U , $e^{-\frac{U(x)-m}{T(i)}} \leq e^{-\frac{\tilde{m} \ln i}{\tau \Delta}} = i^{-\frac{\tilde{m}}{\tau \Delta}}$. Then

$$\begin{aligned} & \|\Pi_i(x) - \Pi_\infty(x)\|_1 \\ &= \frac{e^{-\frac{U(x)-m}{T(i)}} \sum_{\|p_z - p_x\| \leq R_m} e^{-\frac{U(z)-m}{T(i)}}}{\sum_{y \in X} (e^{-\frac{U(y)-m}{T(i)}} \sum_{\|p_z - p_y\| \leq R_m} e^{-\frac{U(z)-m}{T(i)}})} \\ &\leq \frac{|\tilde{\Lambda}| i^{-\frac{2\tilde{m}}{\Delta \tau}}}{|M|} = O(i^{-\frac{2\tilde{m}}{\Delta \tau}}). \end{aligned}$$

Similarly, it can be shown that for a minimizer x , the distance converges at the order $O(i^{-\frac{2\tilde{m}}{\Delta \tau}})$ too. And so is $\|\Pi_i - \Pi_\infty\|_1$. Then a bound for $\|\nu Q_1 \dots Q_n - \Pi_\infty\|$ is

$$\left(\frac{i}{n}\right)^\lambda + \text{const} \cdot i^{-\frac{2\tilde{m}}{\Delta \tau}} \quad (14)$$

This becomes optimal for

$$\begin{aligned} i^* &= (\text{const} \cdot \frac{2\tilde{m}}{\Delta \tau})^{\lambda + \frac{1}{2\tilde{m}}} \cdot n^{\frac{\lambda}{\lambda + \frac{1}{2\tilde{m}}}} \\ &= \text{const} \cdot n^{\frac{\lambda}{\lambda + \frac{1}{2\tilde{m}}}}. \end{aligned}$$

Eq. (13) then follows by plugging i^* into (14). \square

Proposition 4.1 shows that the potential surface (in particular: $g \triangleq \frac{2\lambda \tilde{m}}{2\tilde{m} + \lambda \Delta \tau}$) determines the convergence speed of the algorithm. It is thus natural to use g as a design indicator. Simulation has been conducted to verify the above analysis. In the simulation a similar scenario as in Fig. 1 is used but with a 10 by 10 grid. Two overlapping circular obstacles are located in the middle of the field, which forms a non-convex shape. A single vehicle starts from one corner, and wants to reach the target area at the other corner. The potential functions used are: $J_s^g = \|p_s - p^g\|$, $J_s^o = \sum_{k=1}^K \frac{1}{\|p_s - p^{ok}\|}$, where p^g and p^{ok} denote the centers of the target area and of the obstacles. λ_g is varied from 0.05 to 100 while λ_o is fixed to 1. For each pair of coefficients, the algorithm is run




Fig. 2. Convergence *vs.* design parameter - comparison of simulation results with analysis.

$N = 10,000$ steps, and the number of times w that the vehicle visited target during the last 100 steps is counted. The empirical distance is then given by $\|\Pi_N - \Pi_\infty\|_1 = 2(1 - w/100)$. Comparison with the numerically calculated design indicator g reveals good agreement with the bound (13) (Fig. 2).

5. CONCLUSIONS

Gibbs sampler-based simulated annealing can potentially achieve distributed control of autonomous vehicles based on local information. The analysis of this algorithm has been complicated by the dynamic graph structure associated with such vehicle networks. As a first step to understand its convergence behavior analytically, this paper has studied two special cases and concluded convergence with specified conditions. Furthermore, the implication of the results in the design of Gibbs potentials has been explored.

Extension of the results to the most general case (multiple-vehicle, limited sensing/moving range, and general tasks) is difficult. However, it is possible that analytical results can be obtained for some special geometries or special missions (i.e., potential functions of particular forms). This can be an interesting direction for future work.

ACKNOWLEDGMENT

This research was supported by the Army Research Office under the ODDR&E MURI01 Program Grant No. DAAD19-01-1-0465 to the Center for Networked Communicating Control Systems (through Boston University), and under ARO Grant No. DAAD190210319.

REFERENCES

- Baras, J. S. and X. Tan (2004). Control of autonomous swarms using Gibbs sampling. In: *Proceedings of the 43rd IEEE Conference on Decision and Control*. Atlantis, Paradise Island, Bahamas. pp. 4752–4757.
- Baras, J. S., X. Tan and P. Hovareshti (2003). Decentralized control of autonomous vehicles. In: *Proceedings of the 42nd IEEE Conference on Decision and Control, Maui*. Vol. 2. Maui, Hawaii. pp. 1532–1537.
- Bremaud, P. (1999). *Markov Chains, Gibbs Fields, Monte Carlo Simulation and Queues*. Springer Verlag. New York.
- Geman, S. and D. Geman (1984). Stochastic relaxation, Gibbs distributions and automation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6**, 721–741.
- Jadbabaie, A., J. Lin and A. S. Morse (2003). Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control* **48**(6), 988–1001.
- Koren, Y. and J. Borenstein (1991). Potential field methods and their inherent limitations for mobile robot navigation. In: *Proceedings of the IEEE International Conference on Robotics and Automation*. Sacramento, CA. pp. 1398–1404.
- Leonard, N. E. and E. Fiorelli (2001). Virtual leaders, artificial potentials and coordinated control of groups. In: *Proceedings of the 40th IEEE Conference on Decision and Control*. Orlando, FL. pp. 2968–2973.
- Olfati-Saber, R. and R. M. Murray (2002). Distributed cooperative control of multiple vehicle formations using structural potential functions. In: *Proceedings of the 15th IFAC World Congress*. Barcelona, Spain.
- Passino, K. M. (2002). Biomimicry of bacterial foraging for distributed optimization and control. *IEEE Control Systems Magazine* **22**(3), 52–67.
- Rimon, E. and D. E. Koditschek (1992). Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation* **8**(5), 501–518.
- Schoenwald, D. A. (2000). AUVs: In space, air, water, and on the ground. *IEEE Control Systems Magazine* **20**(6), 15–18.
- Shahidi, R., M. Shayman and P. S. Krishnaprasad (1991). Mobile robot navigation using potential functions. In: *Proceedings of the IEEE International Conference on Robotics and Automation*. Sacramento, CA. pp. 2047–2053.
- Tanner, H. G., A. Jadbabaie and G. J. Pappas (2003). Stable flocking of mobile agents, Part I: Fixed topology. In: *Proceedings of the 42nd IEEE Conference on Decision and Control*. Maui, Hawaii. pp. 2010–2015.
- Winkler, G. (1995). *Image Analysis, Random Fields, and Dynamic Monte Carlo Methods: A Mathematical Introduction*. Springer-Verlag. New York.