Automatica 46 (2010) 2068–2076

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Decentralized coordination of autonomous swarms using parallel Gibbs sampling[☆]

Xiaobo Tan[a,*], Wei Xi[b], John S. Baras[c]

[a] 2120 Engineering Building, Department of Electrical & Computer Engineering, Michigan State University, East Lansing, MI 48824, USA

[b] Western Digital Corporation, 20511 Lake Forest Dr, Lake Forest, CA 92630, USA

[c] Institute for Systems Research, and Department of Electrical & Computer Engineering, University of Maryland, 2247 AV Williams Building, College Park, MD 20742, USA

ARTICLE INFO

Article history:
Received 26 June 2009
Received in revised form
9 May 2010
Accepted 27 July 2010
Available online 20 October 2010

Keywords:
Agents and autonomous systems
Gibbs sampling
Markov random fields
Multi-vehicle coordination
Simulated annealing

ABSTRACT

In this paper we present analysis of a discrete-time, decentralized, stochastic coordination algorithm for a group of mobile nodes, called an autonomous swarm, on a finite spatial lattice. All nodes take their moves by sampling in parallel their locally perceived Gibbs distributions corresponding to a pairwise, nearest-neighbor potential. The algorithm has no explicit requirements on the connectedness of the underlying information graph, which varies with the swarm configuration. It is established that, with an appropriate annealing schedule, the algorithm results in swarm configurations converging to the (global) minimizers of a modified potential energy function. The extent of discrepancy between the modified and original potential energy functions is determined by the maximum node travel between time steps, and when such distance is small, the ultimate swarm configurations are close to the global minimizers of the original potential energy. Simulation results are further presented to illustrate the capability of the sampling algorithm in approximate global optimization for swarms.

1. Introduction

Rapid technological advances have made it possible to build and deploy a large number of mobile robots or unmanned vehicles at an affordable cost. Networks of such autonomous vehicles, called *autonomous swarms* in this paper, can have a multitude of applications, ranging from surveillance and reconnaissance, to search and rescue, to weather forecast, and to oceanography. In recent years significant advances have been made in collaborative control of autonomous swarms, where tools in optimization, control, dynamical systems, and algebraic graph theory are applied to formally analyze or synthesize interaction rules for mobile nodes. Various problems have been studied, including, e.g., flocking (Jadbabaie, Lin, & Morse, 2003; Olfati-Saber, 2006), rendezvous or aggregation (Cortes, Martinez, & Bullo, 2006; Gazi & Passino, 2003), consensus (Moreau, 2005; Olfati-Saber & Murray, 2004; Ren & Beard, 2005), formation control (Egerstedt & Hu, 2001; Justh & Krishnaprasad, 2004; Xi, Tan, & Baras, 2006), and deployment (Cortes & Bullo, 2005; Martinez & Bullo, 2006). Many of the studies have considered the requirement of distributed, local interactions among nodes, which is dictated by the otherwise prohibitive cost for centralized coordination of large-scale networks, and by the need to ensure robustness against single-node failures.

It is notable that a number of convergence results in multi-agent control accommodate a time-varying information-sharing topology, which is an important concern in practice. This has been achieved by adopting and developing different tools, including the theory of stochastic matrices (Jadbabaie et al., 2003; Ren & Beard, 2005), common Lyapunov functions for switched systems (Olfati-Saber & Murray, 2004), set-valued Lyapunov theory (Moreau, 2005), generalized Lyapunov analysis for nonsmooth and set-valued dynamics (Cortes & Bullo, 2005; Cortes et al., 2006), and passivity theory (Arcak, 2007). Despite the progress made, most results provide only convergence to local minima of potential or objective functions (Baras, Tan, & Hovareshti, 2003; Cortes & Bullo, 2005; Ogren, Fiorelli, & Leonard, 2004; Olfati-Saber, 2006), and global objectives are achievable only if initial configurations are sufficiently close to the desired ones.

In an attempt to overcome the aforementioned problem of nodes being trapped at local minima of potential/objective functions, Baras and Tan (2004) explored by simulation a stochastic, decentralized approach to coordination of multiple mobile nodes. The approach exploits the concept of Markov Random Fields (MRFs) to

[☆] The work of X. Tan was supported in part by NSF (IIS 0916720). The material in this paper was partially presented at the 2008 American Control Conference, June 11–13, 2008, Seattle, Washington, USA. This paper was recommended for publication in revised form by Associate Editor Hideaki Ishii under the direction of Editor Ian R. Petersen.

[*] Corresponding author. Tel.: +1 517 432 5671; fax: +1 517 353 1980.
E-mail addresses: xbtan@msu.edu, xbtan@egr.msu.edu (X. Tan), Wei.Xi@wdc.com (W. Xi), baras@isr.umd.edu (J.S. Baras).

0005-1098/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2010.09.004

capture the local interactions among mobile nodes. MRFs can be thought of as a generalization of Markov chains with the temporal index replaced by a spatial index (Kindermann & Snell, 1980). An MRF is thus a collection of random variables, each located at a spatial *site*, where the conditional probability of any random variable given the values of all the other random variables is equal to its conditional probability given only the values of its *neighbors*. Clearly, a neighborhood system needs to be specified in defining an MRF. Another related notion is the Gibbs Random Field (GRF). A GRF is a collection of random variables at sites where their joint probability for each configuration x is proportional to $e^{-H(x)/T}$, where T represents the temperature in the context of statistical physics and H is the Gibbs potential energy as a function of the configuration x. Such a probability measure is called a Gibbs distribution. It turns out that an MRF is equivalent to a GRF (Kindermann & Snell, 1980), and hence one can capture the local interactions among neighbors, as typical of an MRF, through an appropriate definition of the potential for a GRF. *Gibbs sampling* is a Markov Chain Monte Carlo (MCMC) method, where, by sampling the local characteristics of the Gibbs distribution, one produces a Markov chain on the configuration space of the random field. Starting from any initial distribution, the Gibbs sampling-induced Markov chain converges to the Gibbs distribution. Furthermore, performing Gibbs sampling while reducing T with a proper cooling schedule can result in configurations with globally minimal energy. The latter was applied with great success in image processing and computer vision (Chellappa & Jain, 1993; Geman & Geman, 1984; Winkler, 1995), and it largely motivated the work in Baras and Tan (2004).

The work of Baras and Tan (2004) extends the concept of MRFs to the context of autonomous swarms. A group of mobile nodes is assumed to move in discrete time on a finite spatial lattice, where each node is considered as a (mobile) site of the random field. A nearest-neighbor potential energy, as a function of the swarm configuration, is defined to encode desired configurations as its minimizers, and consists of the sum of terms involving only each node and its proximity-based neighbors. At each time instant, each node decides its next move by sampling a locally perceived conditional Gibbs distribution given the current locations of its neighbors. With Gibbs sampling under a suitable cooling schedule, simulation results have indicated that the approach is promising in achieving global objectives (without being trapped at local minima) through purely local interactions (Baras & Tan, 2004). However, no analytical results were presented to explain why the approach worked in Baras and Tan (2004). A primary difficulty in the analysis is that the neighborhood system, or the information graph, varies with the swarm configuration. This is fundamentally different from classical MRFs (Kindermann & Snell, 1980; Winkler, 1995), where the neighborhood systems are assumed to be fixed. Another difficulty is due to the parallel sampling in the autonomous swarm setting, where all nodes update their locations at the same time. Even for classical MRFs, convergence results are typically available for the case of sequential sampling (one site updating at a time) only.

In prior work of the authors (Xi et al., 2006), analytical results were obtained for sequential Gibbs sampling of autonomous swarms with an additional assumption that global communication is available for forwarding relevant information to a newly selected node at each time step. Sequential sampling, however, has two limitations in practice: (1) it takes too long to complete one round of updating for large networks, and (2) it requires explicit indexing of nodes, which is often impossible due to dynamic addition/removal of nodes. The global communication requirement, despite the limited amount of information transmitted, defeats the goal of full decentralization.

The contribution of the current paper is the rigorous analysis of the parallel Gibbs sampling-based swarm coordination algorithm, for the special but important case of pairwise potentials. A pairwise potential energy function consists of contributions from singletons or pairs of nodes, and is widely adopted in the literature of multi-agent control. It is established that, under a constant temperature, the parallel sampling algorithm results in a unique stationary distribution for the swarm configuration. Furthermore, if the temperature follows an appropriate annealing schedule, the configuration converges to the (global) minimizers of a modified potential energy, where the extent of discrepancy between the modified and original potential energy functions is determined by the maximum node travel distance per time step. In particular, when the maximum travel range per update is sufficiently small, or equivalently, when the information about neighbors' locations is updated frequently enough, the ultimate swarm configuration is close to the global minimizers of the original potential energy function. We also note that the algorithm and the proof do not require explicitly the connectedness of the information graph.

Simulation results on examples of rendezvous and line formation are further presented to support the analysis, where the Gibbs sampling algorithm is compared to a deterministic gradient descent-type algorithm. With the gradient descent-type algorithm, the swarm configuration is often stuck at local minima of the potential energy function, while in contrast, the Gibbs sampling algorithm always results in configurations close to the desired ones.

The remainder of the paper is organized as follows. In Section 2, the background on MRFs is briefly reviewed and the application of MRFs to modeling of autonomous swarms is described. Analysis of the parallel sampling algorithm is carried out in Section 3. Simulation results are presented in Section 4. Finally, Section 5 provides concluding remarks.

2. MRFs and application to swarming control

2.1. Review of classical MRFs and Gibbs sampling

2.1.1. MRFs

Intuitively, a random field can be thought of as a (spatial) field of random variables. Let α denote the set of spatial sites, with a cardinality of N. Without loss of generality, we will denote each site with its index s, with $s = 1, 2, \ldots, N$. Consequently, we have $\alpha = \{1, 2, \ldots, N\}$. At each site s, there is an associated random variable X_s, which takes values in a set Λ_s. Λ_s is called the *phase space* for site s. The collection of the random variables, $X = \{X_s\}_{s=1}^N = \{X_s\}_{s\in\alpha}$, is called a random field. A realization $x = \{x_s\}_{s\in\alpha}$ of X is called a *configuration* or a *state* of the field. The configuration space of the random field is defined as $\chi \triangleq \{(x_1, x_2, \ldots, x_N) : x_s \in \Lambda_s, s \in \alpha\}$.

A *neighborhood system* on α is defined to be a family Γ of sets, $\Gamma = \{\Gamma_s\}_{s\in\alpha}$. Here $\Gamma_s \subset \alpha$ is the set of neighbors for site s. The neighborhood system satisfies the following conditions: for $s, r \in \alpha$, (1) $s \notin \Gamma_s$, and (2) $r \in \Gamma_s \Leftrightarrow s \in \Gamma_r$. The neighborhood system induces an undirected graph with vertices $s \in \alpha$, where an edge exists between vertices s and r if and only if $r \in \Gamma_s$. A set $C \subset \alpha$ is called a *clique* if all elements of C are neighbors of each other. A random field X is called a *Markov random field* (MRF) with respect to the neighborhood system Γ if, $\forall s \in \alpha$,

$$P(X_s = x_s | X_r = x_r : r \in \alpha, r \neq s) = P(X_s = x_s | X_r = x_r : r \in \Gamma_s). \tag{1}$$

Here $P(E_1|E_2)$ denotes the conditional probability of E_1 given E_2. The right-hand side of (1) is often referred to as the *local characteristics* of the MRF.

A *potential* U is a family $\{U_A : A \subset \alpha\}$ of functions on the configuration space χ, where $U_A : \chi \to \mathbb{R}$, and $U_A(x)$ depends only on $x_A \triangleq \{x_s : s \in A\}$. In other words, U_A is only a function of the values at the sites contained in the set A. For convenience, we will denote $U_A(x)$ as $U_A(x_A)$ from here on. If $U_A \equiv 0$ whenever A is not a clique or a singleton, U is called a *nearest-neighbor* potential. If $U_A \equiv 0$ whenever A is not a pair or a singleton, U is called a *pairwise* potential. U is called a *pairwise*, *nearest-neighbor* potential if it is both a pairwise potential and a nearest-neighbor potential.

Given a potential U, the *potential energy* $H(x)$ for configuration x is defined as

$$H(x) = \sum_{A \subset \alpha} U_A(x_A). \tag{2}$$

In particular, for a pairwise, nearest-neighbor potential U, we can write H as

$$H(x) = \sum_{s \in \alpha} U_{\{s\}}(x_s) + \sum_{(s,t) \in \alpha \times \alpha,\ t \in \Gamma_s} U_{\{s,t\}}(x_s, x_t). \tag{3}$$

A random field X is called a *Gibbs random field* (GRF) if

$$P(X = x) = \frac{\mathrm{e}^{-H(x)/T}}{Z}, \quad \text{with } Z = \sum_z \mathrm{e}^{-H(z)/T}. \tag{4}$$

The probability measure shown in (4) is called a Gibbs distribution, and the underlying potential U is called a Gibbs potential. T has the interpretation of temperature in the context of statistical physics. Eq. (4) implies that a higher-energy state has a lower probability, and that the influence of energy on probability increases as T decreases. The Hammersley–Clifford theorem establishes the equivalence between an MRF and a GRF (Kindermann & Snell, 1980): any MRF can be shown to have a Gibbs distribution for an appropriately defined potential energy function H, and conversely, any GRF can be shown to have the Markovian properties (1) with the neighborhood system determined by the potential U. This equivalence provides a tangible, convenient characterization of local interactions among neighboring sites in MRFs through Gibbs potentials.

2.1.2. Gibbs sampling

In statistical physics, a GRF is often used to describe the distribution of system configurations at the thermodynamic equilibrium. However, direct evaluation of (4) and related ensemble averages is often impossible due to the high cardinality of the configuration space (the latter rendering the computation of Z intractable). Markov Chain Monte Carlo (MCMC) methods, such as the Metropolis algorithm (Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953) and the Gibbs sampler (Geman & Geman, 1984), can generate Markov chains on the configuration space, with (4) as the limiting probability measure. Next we illustrate such a process with the example of sequential Gibbs sampling.

Given a configuration $x = (x_1, \ldots, x_N)$ at time instant n, one can update it to a different configuration y by updating the values at each site sequentially. For example, for site s, one can update x_s to some value y_s at time $n + 1$ based on the following probability (this is what is meant by *sampling*): for $y_s \in \Lambda_s$,

$$P(X_s(n+1) = y_s \mid X_{\alpha\backslash s}(n) = x_{\alpha\backslash s}) = \frac{\mathrm{e}^{-H\left(y_s, x_{\alpha\backslash s}\right)/T}}{\sum_{z_s \in \Lambda_s} \mathrm{e}^{-H\left(z_s, x_{\alpha\backslash s}\right)/T}}. \tag{5}$$

In (5), $\alpha\backslash s$ denotes the set of all sites other than s:

$$\alpha\backslash s \triangleq \{r \in \alpha : r \neq s\},$$

and $x_{\alpha\backslash s}$ represents the components of x corresponding to the sites in $\alpha\backslash s$:$x_{\alpha\backslash s} \triangleq \{x_r : r \in \alpha\backslash s\}$. The notation $\left(y_s, x_{\alpha\backslash s}\right)$ (likewise for $(z_s, x_{\alpha\backslash s})$) represents a configuration where site s takes the value y_s while other sites take the values $x_{\alpha\backslash s}$. Note that the right-hand side of (5) is precisely the conditional probability of X_s given the values at other sites for a Gibbs distribution (4). It can be verified easily that the evaluation of (5) involves only $\{x_r : r \in \Gamma_s\}$ for a Gibbs field with a nearest-neighbor potential, and thus can be performed efficiently. Following the above procedure, one can update all sites in a prescribed order within N time steps. This process generates a (homogeneous) Markov chain in the configuration space χ, with positive N-step transition probabilities $\mathscr{P}(x, y), \forall x, y \in \chi$. Here $\mathscr{P}(x, y)$ denotes the probability of reaching y after N steps of sequential sampling, starting from the configuration x. It can be shown that the following *detailed balance* equation holds for $\mathscr{P}$ and the Gibbs distribution P:

$$P(x)\mathscr{P}(x, y) = P(y)\mathscr{P}(y, x), \tag{6}$$

which implies that the Gibbs measure $P(x)$ is the (unique) stationary distribution for the Gibbs sampling-induced Markov chain (Winkler, 1995). In other words, starting from any initial distribution, the sampling process will ultimately result in a Gibbs distribution for the configurations.

The Gibbs distribution (4) depends on the temperature T. The lower T is, the higher probabilities for the lowest-energy configurations. In the limit of $T \to 0$, (4) produces probabilities concentrating solely on configurations of minimum energy. Taking the idea of simulated annealing (Kirkpatrick, Gebatt, & Vecchi, 1983), Geman and Geman proposed decreasing T gradually during Gibbs sampling and established the convergence to the lowest-energy configurations (Geman & Geman, 1984).

2.2. Extension to autonomous swarms

Motivated by the promise of Gibbs sampling in achieving globally minimizing configurations through local interactions, Baras and Tan extended the concepts of MRFs and Gibbs sampling to the context of autonomous swarms (Baras & Tan, 2004). Consider a group of mobile nodes moving in a bounded region within the two-dimensional (2D) or three-dimensional (3D) space. The region is discretized into a lattice, and for ease of presentation, each cell is assumed to be square with unit dimensions. A mobile node is assumed to be a point that moves from the center of one cell to that of another. Each node has a *sensing range* R_s: it can sense the locations of obstacles and other nodes within distance R_s. It also has an *interaction range* $R_i \leq R_s$: the moving decision of a node is only influenced by nodes within the distance R_i, which form its set of *neighbors*. In addition, each node can travel by at most $R_m \leq R_s$ within each time step. Fig. 1 illustrates the definitions of the three ranges in a 2D grid. The distances on the lattice are defined using the Euclidean norm based on the center locations of the cells.

The R_i-neighborhood relations induce a graph structure, where the nodes form the vertices of the graph and an edge exists between two nodes if and only if they are neighbors of each other. A random field is then defined on this graph, where each node is considered as a site and the random variable associated with this site represents the location of the node. In particular, we will denote the set of nodes as $\alpha = \{1, 2, \ldots, N\}$, where N is the total number of nodes. X_s (or x_s) will denote the center location of the cell in which node s resides. Similar notation, such as y_s, z_s, etc., will also be used in the later discussion. We will use $x = (x_1, \ldots, x_N)$ to denote the configuration of the swarm. Given x, the set of neighbors $\Gamma_s(x)$ is defined as

$$\Gamma_s(x) \triangleq \{r \in \alpha : r \neq s, \|x_r - x_s\| \leq R_i\}.$$

The set of lattice cells within R_m from node s form the phase space Λ_s. Unlike in the classical MRF case, the phase space Λ_s here will vary with x_s, the location of node s. A suitable potential U can be defined to reflect the swarm coordination objectives, from which the potential energy $H(x)$ can be evaluated, as explained in Section 2.1.1.

One can then perform Gibbs sampling with simulated annealing to update the locations of the nodes. For a nearest-neighbor

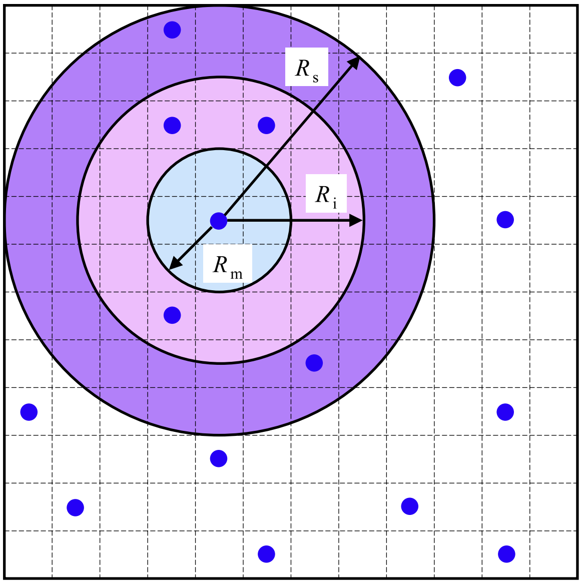

Fig. 1. Illustration of the sensing range R_s, the interaction range R_i, and the moving range R_m on a 2D grid.

potential, the sampling can be done by each node locally. While simulation results were promising (Baras & Tan, 2004), the analysis is challenging because the neighborhood system $\Gamma(x) \triangleq \{\Gamma_s(x), s \in \alpha\}$ varies with the swarm configuration. Analytical results for classical MRFs cannot be applied directly, since the neighborhood system is assumed to be fixed there (Kindermann & Snell, 1980; Winkler, 1995). In particular, even with sequential sampling, the detailed balance (6) no longer holds and the Gibbs distribution is no longer the stationary distribution. In our prior work (Xi et al., 2006), we analyzed a special sequential sampling scheme with an assumption on limited global communication. However, due to feasibility considerations, one will be mostly interested in parallel sampling schemes. Analysis of parallel sampling is involved even for classical MRFs (Winkler, 1995). The goal of this paper is to provide rigorous analysis of a parallel Gibbs sampling scheme for swarm coordination, for the special but popular case of pairwise, nearest-neighbor potentials.

3. Analysis of the parallel Gibbs sampling scheme

3.1. The parallel sampling algorithm

Let n denote the index of time steps. Let $X(n) = x = (x_1, \ldots, x_N)$ be the swarm configuration at time n. Let $F_s(x) \triangleq \{z_s : \|z_s - x_s\| \le R_m\}$ be the set of accessible cell locations for node s given the configuration x, determined by the mobility constraint. Let $F(x)$ be the set of configurations that are accessible from x within one time step:

$$F(x) \triangleq \{z = (z_1, \ldots, z_N) : \|z_s - x_s\| \le R_m,\ s \in \alpha\}.$$

Under parallel Gibbs sampling, all nodes will simultaneously update their locations based on the configuration x at time n; in particular, the node s will move from x_s to y_s at time $n+1$ with probability

$$P_{T,s}(x_s, y_s|x) = \begin{cases} \dfrac{e^{-H(y_s, x_{\alpha\backslash s})/T}}{\sum\limits_{z_s \in F_s(x)} e^{-H(z_s, x_{\alpha\backslash s})/T}} & \text{if } y_s \in F_s(x) \\ 0 & \text{if } y_s \notin F_s(x). \end{cases} \tag{7}$$

For simulated annealing, the temperature variable T will be a function of the time step n. The following assumptions are made:

- ($\mathbf{A_1}$) The total number of lattice cells is bounded;
- ($\mathbf{A_2}$) $R_i + R_m \le R_s$;
- ($\mathbf{A_3}$) U is a pairwise, nearest-neighbor potential.

Remark 3.1. ($\mathbf{A_1}$) requires that the nodes move in a bounded region, which is a reasonable assumption. It will allow us to establish the ergodicity of the Markov chain induced by Gibbs sampling under a constant temperature, and consequently the convergence of the chain to a unique stationary distribution.

Remark 3.2. ($\mathbf{A_2}$) implies that a node s at x_s is able to evaluate the set of new neighbors should it move to $y_s \in F_s(x)$ while other nodes stay put.

Remark 3.3. Similar to (3), ($\mathbf{A_3}$) implies that the corresponding potential energy $H(x)$ for configuration x can be written as

$$H(x) = \sum_{s \in \alpha} U_{\{s\}}(x_s) + \sum_{(s,t) \in \alpha \times \alpha,\ t \in \Gamma_s(x)} U_{\{s,t\}}(x_s, x_t). \tag{8}$$

We can show that ($\mathbf{A_2}$) and ($\mathbf{A_3}$) together ensure the local computability of (7) by node s. In particular, for $y_s, z_s \in F_s(x)$,

$$\begin{aligned} H(y_s, x_{\alpha\backslash s}) &= U_{\{s\}}(y_s) + \sum_{t \in \Gamma_s(y_s, x_{\alpha\backslash s})} U_{\{s,t\}}(y_s, x_t) \\ &\quad + \text{terms not involving node } s, \\ H\left(z_s, x_{\alpha\backslash s}\right) &= U_{\{s\}}(z_s) + \sum_{t \in \Gamma_s(z_s, x_{\alpha\backslash s})} U_{\{s,t\}}(z_s, x_t) \\ &\quad + \text{terms not involving node } s. \end{aligned}$$

Since $H(y_s, x_{\alpha\backslash s})$ and $H(y_s, x_{\alpha\backslash s})$ share the terms not involving s,

$$\begin{aligned} &\frac{e^{-H(y_s, x_{\alpha\backslash s})/T}}{\sum\limits_{z_s \in F_s(x)} e^{-H\left(z_s, x_{\alpha\backslash s}\right)/T}} \\ &= \frac{e^{-\left(U_{\{s\}}(y_s) + \sum\limits_{t \in \Gamma_s(y_s, x_{\alpha\backslash s})} U_{\{s,t\}}(y_s, x_t)\right)\Big/ T}}{\sum\limits_{z_s \in F_s(x)} e^{-\left(U_{\{s\}}(z_s) + \sum\limits_{t \in \Gamma_s(z_s, x_{\alpha\backslash s})} U_{\{s,t\}}(z_s, x_t)\right)\Big/ T}}. \end{aligned} \tag{9}$$

Evaluation of (9) thus only requires the knowledge of

$$\bar{U}_s(z_s, x_{\alpha\backslash s}) \triangleq U_{\{s\}}(z_s) + \sum_{t \in \Gamma_s(z_s, x_{\alpha\backslash s})} U_{\{s,t\}}(z_s, x_t) \tag{10}$$

for all $z_s \in F_s(x)$ (since $y_s \in F_s(x)$). From ($\mathbf{A_2}$), node s at the current location x_s is able to evaluate the set of its neighbors should it move to $z_s \in F_s(x)$, implying the local computability. Note that $\bar{U}_s(z_s, x_{\alpha\backslash s})$ is well defined even if $\Gamma_s(z_s, x_{\alpha\backslash s}) = \emptyset$, in which case $\bar{U}_s(z_s, x_{\alpha\backslash s}) = U_{\{s\}}(z_s)$.

Using (9), we can write (7) as

$$\begin{aligned} &P_{T,s}(x_s, y_s|x) \\ &= \frac{\mathbf{1}(y_s \in F_s(x)) \cdot e^{-\left(U_{\{s\}}(y_s) + \sum\limits_{t \in \Gamma_s(y_s, x_{\alpha\backslash s})} U_{\{s,t\}}(y_s, x_t)\right)\Big/ T}}{\sum\limits_{z_s \in F_s(x)} e^{-\left(U_{\{s\}}(z_s) + \sum\limits_{t \in \Gamma_s(z_s, x_{\alpha\backslash s})} U_{\{s,t\}}(z_s, x_t)\right)\Big/ T}}, \end{aligned} \tag{11}$$

where $\mathbf{1}(\cdot)$ denotes the indicator function. Since the nodes make independent moving decisions at time n for given x, the kernel $\mathscr{P}_T(x, y) \triangleq \text{Prob}(X(n+1) = y|X(n) = x)$ for the parallel Gibbs sampling-induced Markov chain can be obtained from (11):

$$\begin{aligned}\mathscr{P}_T(x,y) &= \prod_{s\in\alpha} P_{T,s}(x_s, y_s|x)\\ &= \mathbf{1}(y\in F(x))\\ &\quad\cdot \frac{\mathrm{e}^{-\sum_{s\in\alpha}\left(U_{\{s\}}(y_s)+\sum_{t\in\Gamma_s(y_s,x_{\alpha\setminus s})}U_{\{s,t\}}(y_s,x_t)\right)\Big/T}}{\sum_{z\in F(x)}\mathrm{e}^{-\sum_{s\in\alpha}\left(U_{\{s\}}(z_s)+\sum_{t\in\Gamma_s(z_s,x_{\alpha\setminus s})}U_{\{s,t\}}(z_s,x_t)\right)\Big/T}}\end{aligned} \tag{12}$$

$$\begin{aligned}&= \mathbf{1}(y\in F(x))\\ &\quad\cdot \frac{\mathrm{e}^{-\sum_{s\in\alpha}\left(U_{\{s\}}(x_s)+U_{\{s\}}(y_s)+\sum_{t\in\Gamma_s(y_s,x_{\alpha\setminus s})}U_{\{s,t\}}(y_s,x_t)\right)\Big/T}}{\sum_{z\in F(x)}\mathrm{e}^{-\sum_{s\in\alpha}\left(U_{\{s\}}(x_s)+U_{\{s\}}(z_s)+\sum_{t\in\Gamma_s(z_s,x_{\alpha\setminus s})}U_{\{s,t\}}(z_s,x_t)\right)\Big/T}}\end{aligned} \tag{13}$$

$$= \mathbf{1}(y\in F(x))\cdot\frac{\mathrm{e}^{-G(x,y)/T}}{\sum_{z\in F(x)}\mathrm{e}^{-G(x,z)/T}}, \tag{14}$$

where

$$G(x,y) \triangleq \sum_{s\in\alpha}\left(U_{\{s\}}(x_s)+U_{\{s\}}(y_s)+\sum_{t\in\Gamma_s(y_s,x_{\alpha\setminus s})}U_{\{s,t\}}(y_s,x_t)\right). \tag{15}$$

The denominator of (12) is derived from that, for $y \in F(x)$, $\mathscr{P}_T(x,y)$ is proportional to the expression in the numerator, and that $\sum_{z\in F(x)} \mathscr{P}_T(x,z) = 1$. Eq. (13) is obtained from (12) by multiplying both its denominator and numerator by $\mathrm{e}^{-\sum_{s\in\alpha} U_{\{s\}}(x_s)}$. Note that the transition matrix $\mathscr{P}_T$ has dimensions of $q \times q$, where q denotes the cardinality of the configuration space χ.

Lemma 3.1. *For $y \in F(x)$, the function G is symmetric, i.e., $G(x,y) = G(y,x)$.*

Proof. A key observation is that $t \in \Gamma_s(y_s, x_{\alpha\setminus s}) \Rightarrow s \in \Gamma_t(x_t, y_{\alpha\setminus t})$. In other words, node t being a neighbor of node s for the configuration $(y_s, x_{\alpha\setminus s})$ implies that node s will be a neighbor of node t for the configuration $(x_t, y_{\alpha\setminus t})$ (or vice versa). This is because both conditions mean $\|y_s - x_t\| \le R_\mathrm{i}$. One can then write

$$\begin{aligned}G(x,y) &= \sum_{s\in\alpha}\left(U_{\{s\}}(x_s)+U_{\{s\}}(y_s)+\sum_{t\in\Gamma_s(y_s,x_{\alpha\setminus s})}U_{\{s,t\}}(y_s,x_t)\right)\\ &= \sum_{s\in\alpha}\left(U_{\{s\}}(y_s)+U_{\{s\}}(x_s)\right)+\sum_{t\in\alpha}\sum_{s\in\Gamma_t(x_t,y_{\alpha\setminus t})}U_{\{t,s\}}(x_t,y_s)\\ &= \sum_{t\in\alpha}\left(U_{\{t\}}(y_t)+U_{\{t\}}(x_t)+\sum_{s\in\Gamma_t(x_t,y_{\alpha\setminus t})}U_{\{t,s\}}(x_t,y_s)\right)\\ &= G(y,x),\end{aligned}$$

where the aforementioned observation is used in arriving at the second equality. □

3.2. Stationary distribution under constant-T sampling

Parallel Gibbs sampling produces a Markov chain $X(n)$ for the swarm configuration. We first characterize the stationary distribution of $X(n)$ for a fixed temperature T. This can then be used to analyze the limiting behavior as $T \to 0$ during simulated annealing.

Theorem 3.1. *Let the assumptions $(\mathbf{A_1})$–$(\mathbf{A_3})$ hold. Under parallel Gibbs sampling with a fixed T, the swarm configuration $X(n)$ has a unique stationary distribution Π_T. Starting from any distribution ν for the swarm configuration,*

$$\lim_{n\to\infty}\nu\mathscr{P}_T^n \to \Pi_T, \tag{16}$$

where $\mathscr{P}_T$ represents the transition matrix of the Markov chain, as determined by (14). Furthermore, the explicit form of Π_T is

$$\Pi_T(x) = \frac{\sum_{z\in F(x)}\mathrm{e}^{-G(x,z)/T}}{\sum_{x'\in\chi}\sum_{z\in F(x')}\mathrm{e}^{-G(x',z)/T}}, \tag{17}$$

where χ denotes the space of swarm configurations.

Proof. For a constant T, $X(n)$ generated under the parallel Gibbs sampling is a homogeneous Markov chain with the transition matrix $\mathscr{P}_T$. Given any current configuration x, the probability of reaching any $y \in F(x)$ within one time step is strictly positive. From $(\mathbf{A_1})$, there exists a finite integer $\tau > 0$, such that, given any x and y in the configuration space χ, the probability of reaching y from x within τ sampling steps is positive. In other words, $\mathscr{P}_T$ has a strictly positive power $\mathscr{P}_T^\tau$. Hence the Markov chain $X(n)$ is ergodic and has a unique, stationary distribution Π_T (Horn & Johnson, 1985); furthermore, (16) follows.

Next we verify that Π_T has an explicit expression as in (17). Denote the right-hand side of (17) as $\hat{\Pi}_T$, and its denominator as Z_T. For $y \in F(x)$,

$$\begin{aligned}\hat{\Pi}_T(x)\mathscr{P}_T(x,y) &= \frac{\sum_{z\in F(x)}\mathrm{e}^{-G(x,z)/T}}{Z_T}\cdot\frac{\mathrm{e}^{-G(x,y)/T}}{\sum_{z\in F(x)}\mathrm{e}^{-G(x,z)/T}}\\ &= \frac{\mathrm{e}^{-G(x,y)/T}}{Z_T}\\ &= \frac{\sum_{z\in F(y)}\mathrm{e}^{-G(y,z)/T}}{Z_T}\cdot\frac{\mathrm{e}^{-G(y,x)/T}}{\sum_{z\in F(y)}\mathrm{e}^{-G(y,z)/T}}\end{aligned} \tag{18}$$

$$= \hat{\Pi}_T(y)\mathscr{P}_T(y,x), \tag{19}$$

where Lemma 3.1 is used in (18). If $y \notin F(x)$, $\mathscr{P}_T(x,y) = \mathscr{P}_T(y,x) = 0$ and (19) still holds. With (19), we can write, $\forall x \in \chi$,

$$\begin{aligned}\sum_{y\in\chi}\hat{\Pi}_T(y)\mathscr{P}_T(y,x) &= \sum_{y\in\chi}\hat{\Pi}_T(x)\mathscr{P}_T(x,y)\\ &= \hat{\Pi}_T(x)\sum_{y\in\chi}\mathscr{P}_T(x,y)\\ &= \hat{\Pi}_T(x),\end{aligned} \tag{20}$$

since $\sum_{y\in\chi}\mathscr{P}_T(x,y) = 1$. Eq. (20) confirms that $\hat{\Pi}_T$ is a stationary distribution for the Markov kernel $\mathscr{P}_T$. $\Pi_T = \hat{\Pi}_T$ then follows from the uniqueness of the stationary distribution. □

Remark 3.4. From (16), the swarm configuration under parallel, constant-T, Gibbs sampling converges to the distribution Π_T. Note that while it takes a form similar to a Gibbs distribution, Π_T is not a Gibbs distribution. This illustrates the difference of the parallel sampling algorithm from the sequential Gibbs sampling of a classical MRF, which would result in a Gibbs distribution.

3.3. Convergence under annealing

Let τ be the minimum integer such that all entries of $\mathscr{P}_T^\tau$ are strictly positive. Note that the definition of τ is independent of T. In annealing, the temperature $T(n)$ will drop as a function of time n.

Theorem 3.2. *Let the assumptions* ($\mathbf{A_1}$)–($\mathbf{A_3}$) *hold. Define*

$$\Delta \triangleq \max_x \max_{y,z\in F(x)} |G(x,y) - G(x,z)|.$$

Let $T(n)$ *be a cooling schedule such that*

$$T(n) = T_k, \quad \tau k \le n < \tau(k+1), \tag{21}$$

where $\{T_k\}$ *is a sequence decreasing to* 0 *and satisfying*

$$T_k \ge \frac{\Delta}{\ln k}. \tag{22}$$

Then for any initial distribution ν *for the swarm configuration,*

(1)
$$\lim_{k\to\infty} \nu Q_1 \cdots Q_k \to \Pi_0, \tag{23}$$
where $Q_i \triangleq \mathscr{P}_{T_i}^{\tau}$, *and* Π_0 *represents the limit of* Π_T, (17), *as* $T \to 0$;

(2) *Define* $\hat{H}(x) \triangleq \min_{z\in F(x)} G(x,z)$ *and* $m_0 \triangleq \min_x \hat{H}(x)$. *The support* $\mathscr{M}$ *of the limiting distribution* Π_0 *is*
$$\mathscr{M} = \{x : \hat{H}(x) = m_0\}. \tag{24}$$

Proof. Claim (1) concerns the characterization of the limiting behavior of $\|\nu Q_1 \cdots Q_k - \Pi_0\|_1$, where $\|\cdot\|_1$ denotes the 1-norm of a vector. The proof uses the contraction property of the Markov kernel Q_k, which is where the annealing schedule (22) comes in. The full proof follows closely the steps in proving Theorem 3.2 in Xi et al. (2006), and is omitted here in the interest of brevity.

To establish the support of Π_0, one can rewrite Π_T as

$$\Pi_T(x) = \frac{\sum_{z\in F(x)} \mathrm{e}^{-(G(x,z)-m_0)/T}}{\sum_{x'\in\chi}\sum_{z\in F(x')} \mathrm{e}^{-\left(G(x',z)-m_0\right)/T}}. \tag{25}$$

As $T \to 0$, $\mathrm{e}^{-(G(x,z)-m_0)/T}$ approaches 1 if $G(x,z) = m_0$, and approaches 0 otherwise. As a result, the numerator of $\Pi_0(x)$, expressed as in (25), will be nonzero if and only if $x \in \mathscr{M}$. Claim (2) follows by noting that the denominator of $\Pi_0(x)$ is always positive and finite. □

Remark 3.5. From Theorem 3.2, the swarm configuration under the parallel Gibbs sampling algorithm with annealing converges to the (global) minimizer of $\hat{H}(x) = \min_{z\in F(x)} G(x,z)$. Recall the definition of G in (15) and note the expression for the original potential energy

$$H(x) = \sum_{s\in\alpha}\left(U_{\{s\}}(x_s) + \frac{1}{2}\sum_{t\in\Gamma_s(x)} U_{\{s,t\}}(x_s,x_t)\right), \tag{26}$$

where $\frac{1}{2}$ accounts for the fact that each pairwise interaction is counted twice in (26). It can be seen that $G(x,z)$ represents an energy term obtained by combining the components from configurations x and z. In particular, $G(x,x) = 2H(x)$. Consequently, $\hat{H}(x)$ can be regarded as a potential energy function that is modified from $2H(x)$. Since for each $s \in \alpha$, $\|z_s - x_s\| \le R_{\mathrm{m}}$, the difference $|\hat{H}(x) - 2H(x)|$ depends on the moving range R_{m} per time step.

Remark 3.6. Given the physical constraint on the speed of a mobile node, R_{m} is inversely proportional to the actual time between steps n and $n+1$. The latter indicates how frequent the nodes get feedback information about their neighbors for making moving decisions. More frequent update would imply a smaller R_{m}, and consequently, the difference between the modified energy and the original energy, $|\hat{H}(x) - 2H(x)|$ becomes smaller. This can be interpreted as a tradeoff between the cost for information gathering and the achievable performance in minimizing H.

Remark 3.7. The parallel Gibbs sampling algorithm achieves global minimizers of the modified energy $\hat{H}(x)$, which could be sufficiently close to the global minimizers of the original energy $H(x)$ when R_{m} is small. This result does not require explicitly the connectedness of the information graph during sampling. However, the assumption of a bounded lattice implies a positive transition probability between any two configurations over τ steps, which consequently implies that there is a positive probability for any configuration (including configurations with connected information graph) to exist every τ steps. In other words, the information graph associated with the swarm configuration is connected with a positive probability every τ steps. In some sense, this is analogous to the joint connectedness condition in the deterministic setting (Jadbabaie et al., 2003).

Corollary 3.1. *Let the assumptions* ($\mathbf{A_1}$)–($\mathbf{A_3}$) *hold, and let* Δ *be defined as in* Theorem 3.2. *Let* $T(n)$ *be a cooling schedule such that* $T(n) \ge \frac{\Delta'}{\ln n}$, *with* $\Delta' > \Delta$. *Then for any initial distribution* ν, *the swarm configuration under parallel Gibbs sampling converges to* $\mathscr{M}$, *the set of global minimizers of* $\hat{H}$.

Sketch of proof. First consider a schedule $T_1(n)$ that meets the conditions in Theorem 3.2: $T_1(n) = T_k = \frac{\Delta}{\ln k}$, $\tau k \le n < \tau(k+1)$. It can be shown that there exists a finite n_0, such that when $n \ge n_0$, $T(n) \ge T_1(n)$. Define $Q_k \triangleq \mathscr{P}_{T(\tau k)}\mathscr{P}_{T(\tau k+1)}\cdots\mathscr{P}_{T(\tau(k+1)-1)}$, and $\hat{Q}_k \triangleq \mathscr{P}_{T_k}^{\tau}$. From $T(n) \ge T_1(n)$, one can establish that the contraction property of Q_k is non-weaker than that of $\hat{Q}_k$, $\forall k > n_0/\tau$. The rest of the proof follows similarly as for Theorem 3.2. □

4. Simulation results

Simulation has been further performed to corroborate the analysis and verify the effectiveness of the parallel sampling algorithm. For comparison purposes, a deterministic, gradient descent-type algorithm has also been implemented. Under the gradient-type algorithm, the nodes are subject to the same constraints on R_{i}, R_{s}, and R_{m}. The only difference from the Gibbs sampling algorithm is that, when updating its location, each node moves to the location that would minimize the potential energy if other nodes stay put. Two examples are presented next: (1) rendezvous, and (2) line formation, both on a 50 × 50 square lattice. For the sampling algorithm, in view of Corollary 3.1, we have adopted schedules of the form: $T(n) = T_0/\ln(n)$. T_0 is chosen empirically since the analytically determined values are found to be too conservative in simulation.

4.1. Rendezvous

In the rendezvous problem, the potential is designed as, $U_{\{s\}}(x_s) = 0$, $\forall s \in \alpha$, and for $t \in \Gamma_s(x)$,

$$U_{\{s,t\}}(x_s,x_t) = \begin{cases} 10 & \text{if } \|x_s - x_t\| = 0 \\ -\dfrac{1}{\|x_s - x_t\|} & \text{otherwise.} \end{cases} \tag{27}$$

The equation $U_{\{s\}}(x_s) = 0$ implies that there is no pre-specified gathering point. By setting the potential of an overlapping pair to be high in (27), we discourage multiple nodes from occupying the same cell and thus avoid over-crowding. Figs. 2 and 3 show the snapshots of swarm configurations at different times for the gradient-type algorithm and the sampling algorithm, respectively. The initial configurations for both algorithms were the same. The number of nodes was $N = 40$, and the parameters used in simulation were: $R_{\mathrm{s}} = 13\sqrt{2} + 2$, $R_{\mathrm{i}} = 13\sqrt{2}$, $R_{\mathrm{m}} = 2$, and $T_0 = 5$. From Fig. 2, the nodes formed multiple clusters locally and were stuck at local minima of the potential energy under the gradient-type algorithm. In comparison, under the sampling algorithm, while the

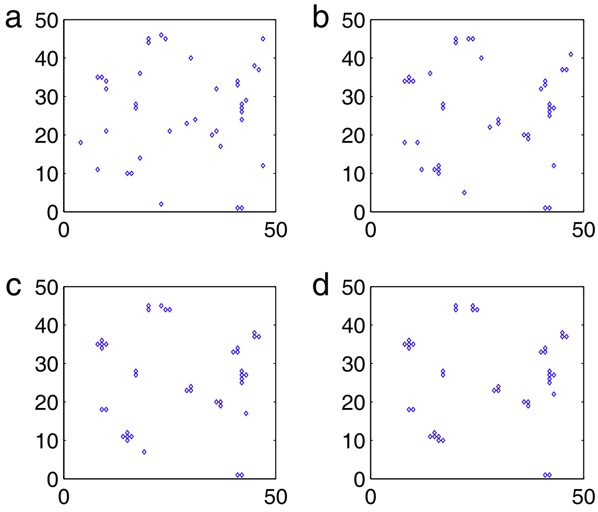

Fig. 2. Snapshots of a swarm of 40 nodes during rendezvous under the gradient-type algorithm: (a) $n = 2$; (b) $n = 5$; (c) $n = 10$; (d) $n = 15$.

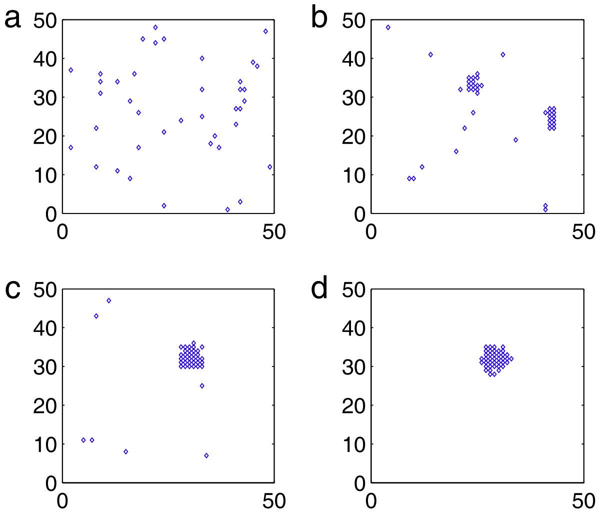

Fig. 3. Snapshots of a swarm of 40 nodes during rendezvous under the parallel Gibbs sampling algorithm: (a) Initial configuration; (b) $n = 200$; (c) $n = 300$; (d) $n = 650$.

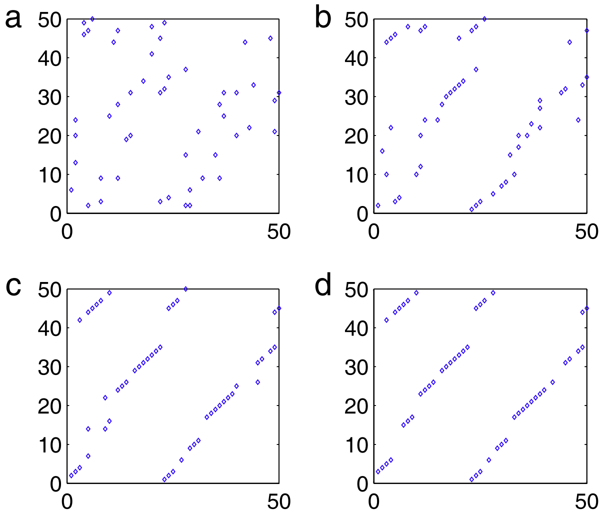

Fig. 4. Snapshots of a swarm of 50 nodes during line formation under the gradient-type algorithm: (a) $n = 2$; (b) $n = 5$; (b) $n = 10$; (c) $n = 15$.

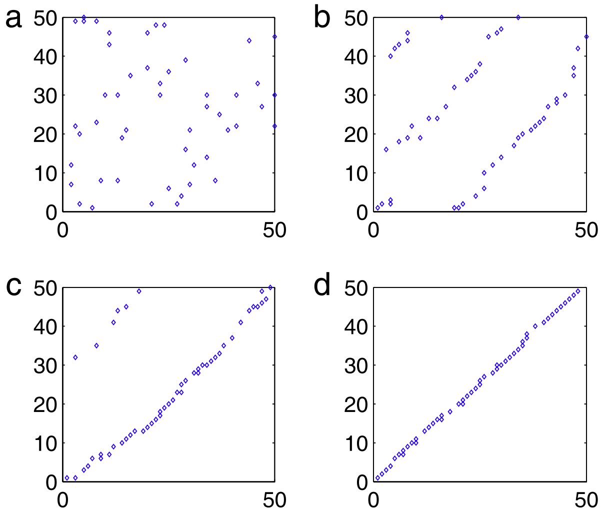

Fig. 5. Snapshots of a swarm of 50 nodes during line formation under the Gibbs sampling algorithm: (a) Initial configuration; (b) $n = 10$; (b) $n = 40$; (c) $n = 140$.

nodes tended to form two clusters at $n = 200$, they successfully managed to escape from the trap and reached rendezvous at $n = 650$. This example illustrates the advantage of the Gibbs sampling algorithm over the traditional gradient-descent-type algorithm in global optimization. Of course, the latter is achieved at a cost—it usually takes the sampling algorithm over 10 times longer to converge than the gradient-type algorithm.

4.2. Line formation

The nodes are required to form a line that makes a 45° angle with respect to the horizontal axis. The potential is designed as, $U_{\{s\}}(x_s) = 0, \forall s \in \alpha$, and for $t \in \Gamma_s(x)$,

$$U_{\{s,t\}}(x_s, x_t) = \begin{cases} 0 & \text{if } \|x_s - x_t\| = 0 \\ -\dfrac{|\langle x_s - x_t, [1,1]^T\rangle|}{\sqrt{2}\|x_s - x_t\|} & \text{otherwise} \end{cases}$$

where $\langle\cdot\rangle$ indicates the inner product. The potential is essentially a measure for the distance between 45° and the angle made by the line connecting a pair of neighboring nodes with respect to the horizontal line. The additive form of the potential energy thus encourages nodes to have more neighbors with desired angles, leading to the formation of a single line; overlapping nodes, however, are discouraged since a connecting line is not well defined in that case.

Figs. 4 and 5 show the snapshots of swarm configurations under the gradient-type algorithm and the Gibbs sampling algorithm, respectively. Here 50 nodes were simulated, with $R_s = 10\sqrt{2} + 3$, $R_i = 10\sqrt{2}$, $R_m = 3$, and $T_0 = 1$. The two algorithms started with the same initial configuration. With the gradient-type algorithm, the swarm evolved into three major stripes, each consisting of line segments of 45°. In contrast, with the Gibbs sampling algorithm, the swarm first self-organized into several parallel line segments, which then merged into a single line of 45°. Fig. 6 compares the evolution of the potential energy with the two algorithms, and it is evident that the configuration was stuck in a local minimum under the gradient-type algorithm.

Multiple simulation runs were carried out, starting from different initial configurations. It is interesting to note that, in all cases, the swarm configuration converged to the diagonal line as in Fig. 5 under Gibbs sampling. This can be explained by the fact that the diagonal line is the only configuration that can accommodate 50 vehicles with minimum inter-vehicle separation larger than zero, thus supporting the (approximate) global optimization capability of the algorithm. On the other hand, the gradient-type

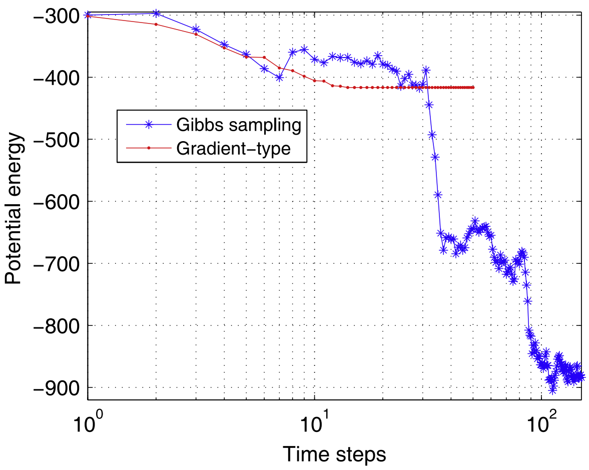

Fig. 6. Evolution of the potential energy under the gradient-type algorithm and the Gibbs sampling algorithm.

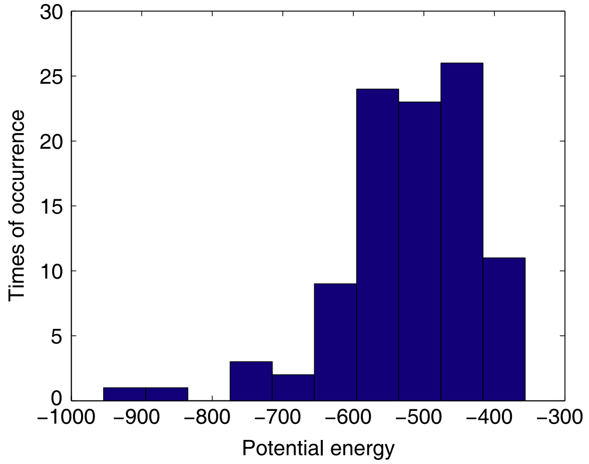

Fig. 7. Histogram of the potential energy for the final swarm configurations achieved with the gradient-type algorithm.

algorithm fails to result in the desired configuration most of the time. Fig. 7 shows the histogram of the potential energy for the final swarm configuration achieved with the gradient-type algorithm. Out of a total of 100 simulation runs, the desired line formation was achieved only once.

5. Conclusion and discussions

In this paper the parallel Gibbs sampling algorithm for swarm coordination was analyzed. The explicit expression for the stationary distribution of swarm configuration was derived for the special but popular case of pairwise potential, and the convergence of the algorithm under appropriate annealing schedule was established. It was found that the algorithm minimizes a modified potential energy $\hat{H}$, where the extent of modification from the original energy H is related to the moving range R_m per time step. When R_m is relatively small, the global minimizers of the modified potential function will be close to those of the original potential function. Simulation results were further presented to compare the sampling algorithm with a deterministic gradient-type algorithm. The Gibbs sampling algorithm showed clear advantage in achieving globally optimal configurations, at the cost of exploration time. While this algorithm can provide high-level path planning for autonomous swarms, it needs to be combined with lower-level path planners and controllers in implementation, where more detailed node dynamics and constraints are incorporated.

While only a pairwise potential is considered in this paper, we note that the class of pairwise potentials encompasses a broad range of interesting problems in swarming, such as rendezvous, dispersion, and formation control. The presented algorithm is decentralized in the sense that there is no centralized decision maker, and that each mobile node makes its moving decisions based only on the locations of its neighbors. In our analysis here, the nodes do need to know the global time and the annealing schedule. However, these assumptions are not considered restrictive. It is possible to relax the assumption on the global time and allow each node to have bounded uncertainties δ on its knowledge about the global time. In that case, we conjecture that the resulting configuration x^* will be close to the set $\mathscr{M}$ of global minimizers of $\hat{H}$, with the bound on the distance between x^* and $\mathscr{M}$ dependent on the bound on δ. The annealing schedule depends only on the potential and is not influenced by the addition or removal of nodes. So it is reasonable for a node to get the annealing schedule when it is informed of the form of the potential (which corresponds to the specific swarm mission).

Future work can be carried out in a few directions. First, the analysis in this paper has been focused on the case of pairwise potentials. We plan to extend the work to cases involving nearest-neighbor potentials of other forms. Second, the design of potentials corresponding to given swarm missions is a problem of interest. One needs to understand how to encode configurations that are desired for a task as the global minimizers of some potential function. This could be intuitive for some simple cases (such as the rendezvous), but is nontrivial in general. A more subtle issue is related to the landscape of the potential function, which has an effect on the convergence speed of the sampling algorithm. Finally, we will also investigate the connection between the parallel Gibbs sampling algorithm studied here and the continuous, diffusion-based swarm coordination approach in Tan (2007), when the time step for Gibbs sampling approaches zero.

References

Arcak, M. (2007). Passivity as a design tool for group coordination. *IEEE Transactions on Automatic Control*, *52*(8), 1380–1390.

Baras, J. S., & Tan, X. (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., Tan, X., & Hovareshti, P. (2003). Decentralized control of autonomous vehicles. In: *Proceedings of the 42nd IEEE conference on decision and control. Vol. 2. Maui, Hawaii* (pp. 1532–1537).

Chellappa, R., & Jain, A. (1993). *Markov random fields: theory and applications*. Boston: Academic Press.

Cortes, J., & Bullo, F. (2005). Coordination and geometric optimization via distributed dynamical systems. *SIAM Journal on Control and Optimization*, *44*(5), 1543–1574.

Cortes, J., Martinez, S., & Bullo, F. (2006). Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. *IEEE Transactions on Automatic Control*, *51*(8), 1289–1298.

Egerstedt, M., & Hu, X. (2001). Formation constrained multi-agent control. *IEEE Transactions on Robotics and Automation*, *17*, 947–951.

Gazi, V., & Passino, K. M. (2003). Stability analysis of swarms. *IEEE Transactions on Automatic Control*, *48*(4), 692–697.

Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *6*, 721–741.

Horn, R., & Johnson, C. R. (1985). *Matrix analysis*. New York: Cambridge University Press.

Jadbabaie, A., Lin, J., & Morse, A. S. (2003). Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, *48*(6), 988–1001.

Justh, E., & Krishnaprasad, P. S. (2004). Equilibria and steering laws for planar formations. *Systems and Control Letters*, *52*, 25–38.

Kindermann, R., & Snell, J. L. (1980). *Markov random fields and their applications*. Providence, RI: American Mathematical Society.

Kirkpatrick, S., Gebatt, C. D., & Vecchi, M. (1983). Optimization by simulated annealing. *Science*, *220*(4598), 671–680.

Martinez, S., & Bullo, F. (2006). Optimal sensor placement and motion coordination for target tracking. *Automatica*, *42*(4), 661–668.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, *21*(6), 1087–1092.

Moreau, L. (2005). Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, *50*(2), 169–182.

Ogren, P., Fiorelli, E., & Leonard, N. E. (2004). Cooperative control of mobile sensor networks: adaptive gradient climbing in a distributed environment. *IEEE Transactions on Automatic Control*, *49*(8), 1292–1302.

Olfati-Saber, R. (2006). Flocking for multi-agent dynamic systems: algorithms and theory. *IEEE Transactions on Automatic Control*, *51*(3), 401–420.

Olfati-Saber, R., & Murray, R. M. (2004). Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, *49*(9), 1520–1533.

Ren, W., & Beard, R. W. (2005). Consensus seeking in multiagent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic Control*, *50*(5), 655–661.

Tan, X. (2007). Self-organization of autonomous swarms via Langevin equation. In: *Proceedings of 46th IEEE conference on decision and control. New Orleans, LA* (pp. 1435–1440).

Winkler, G. (1995). *Image analysis, random fields, and dynamic monte carlo methods: a mathematical introduction*. New York: Springer-Verlag.

Xi, W., Tan, X., & Baras, J. S. (2006). Gibbs sampler-based coordination of autonomous swarms. *Automatica*, *42*(7), 1107–1119.

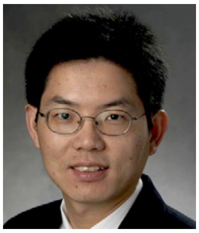

Xiaobo Tan received the Bachelor's and Master's degrees in Automatic Control from Tsinghua University, Beijing, China, in 1995 and 1998, respectively, and the Ph.D. degree in Electrical and Computer Engineering from the University of Maryland, College Park, in 2002.

From September 2002 to July 2004, he was a Research Associate with the Institute for Systems Research at the University of Maryland. In August 2004 he joined the Department of Electrical and Computer Engineering at Michigan State University (MSU) as an Assistant Professor, and founded the Smart Microsystems Laboratory. He has been an Associate Professor at MSU since July 2010. His current research interests include electroactive polymer sensors and actuators, biomimetic robotic fish, mobile sensing in aquatic environments, modeling and control of smart materials, and collaborative control of autonomous systems. Dr. Tan is also keen to integrate his research with educational and outreach activities, and currently leads an NSF-funded Research Experiences for Teachers (RET) Site on Bio-Inspired Technology and Systems (BITS) at MSU. Dr. Tan is an Associate Editor of Automatica, and is serving as the Program Chair for the 15th International Conference on Advanced Robotics (ICAR 2011). He was a guest editor of IEEE Control Systems Magazine for its February 2009 issue special section on modeling and control of hysteresis. Dr. Tan received the NSF CAREER Award in 2006, the 2008 ASME DSCD Best Mechatronics Paper Award (with Yang Fang) in 2009, and the MSU Teacher-Scholar Award in 2010.

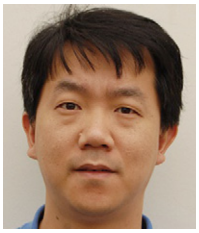

Wei Xi received the B.S. and M.S. degrees from Tsinghua University at Beijing, in 1998 and 2001, respectively, and the Ph.D. degree in Electrical Engineering from the University of Maryland, College Park in 2006. He is currently a Senior Principal Engineer in the Advanced Servo Technology Development Group at Western Digital Corporation, Lake Forest, CA. His research interests include modeling and control of hard disk drives, distributed control of networked systems, and embedded control systems.

John S. Baras received the B.S. in Electrical Engineering with highest distinction from the National Technical University of Athens, Greece, in 1970. He received the M.S. and Ph.D. degrees in Applied Mathematics from Harvard University, Cambridge, MA, in 1971 and 1973 respectively.

Since 1973 he has been with the Department of Electrical and Computer Engineering, University of Maryland at College Park, where he is currently Professor, member of the Applied Mathematics and Scientific Computation Program Faculty, and Affiliate Professor in the Department of Computer Science. From 1985 to 1991 he was the Founding Director of the Institute for Systems Research (ISR) (one of the first six NSF Engineering Research Centers). In February 1990 he was appointed to the Lockheed Martin Chair in Systems Engineering. Since 1991 Dr. Baras has been the Director of the Maryland Center for Hybrid Networks (HYNET), which he co-founded. Dr. Baras has held visiting research scholar positions with Stanford, MIT, Harvard, the Institute National de Reserche en Informatique et en Automatique, the University of California at Berkeley, Linkoping University and the Royal Institute of Technology in Sweden.

Among his awards are: the 1980 George S. Axelby Prize of the IEEE Control Systems Society; the 1978, 1983 and 1993 Alan Berman Research Publication Award from NRL; the 1991, 1994 and 2008 Outstanding Invention of the Year Award from the University of Maryland; the 1996 Engineering Research Center Award of Excellence for Outstanding Contributions in Advancing Maryland Industry; the 1998 Mancur Olson Research Achievement Award from the Univ. of Maryland College Park; the 2002 Best Paper Award for IT/C4ISR at the 23rd Army Science Conference; the 2004 Best Paper Award at the Wireless Security Conference WISE04; the 2007 IEEE Communications Society Leonard G. Abraham Prize in the Field of Communication Systems; the Globecom '08 Best Paper Award for Wireless Networks; the 2008 Best Paper Award for IT/C4ISR at the 26th Army Science Conference.

Dr. Baras' research interests include control, communication and computing systems. He has published more than 500 refereed publications, holds five patents, graduated 60 Ph.D. students and sponsored 40 postdoctoral scholars. He is a Fellow of the IEEE and a Foreign Member of the Royal Swedish Academy of Engineering Sciences (IVA).