

Hedonic coalition formation for optimal deployment

M. Ouimet^a J. Cortés^a

^a*Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA, 92093, USA*

Abstract

This paper presents a distributed algorithmic solution, termed COALITION FORMATION AND DEPLOYMENT ALGORITHM, to achieve network configurations where agents cluster into coincident groups that are distributed optimally over the environment. The motivation for this problem comes from spatial estimation tasks executed with unreliable sensors. We propose a probabilistic strategy that combines a repeated game governing the formation of coalitions with a spatial motion component governing their location. For a class of probabilistic coalition switching laws, we establish the convergence of the agents to coincident groups of a desired size in finite time and the asymptotic convergence of the overall network to the optimal deployment, both with probability 1. We also investigate the algorithm's time and communication complexity. Specifically, we upper bound the expected completion time of executions that use the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS coalition switching law under arbitrary and complete communication topologies. We also upper bound the number of messages required per timestep to execute our strategy. The proposed algorithm is robust to agent addition and subtraction. From a coalitional game perspective, the algorithm is novel in that the players' information is limited to neighboring clusters. From a motion coordination perspective, the algorithm is novel because it brings together the basic tasks of rendezvous (individual agents into clusters) and deployment (clusters in the environment). Simulations illustrate the correctness, robustness, and complexity results.

Key words: sensor networks, hedonic games, coalition formation, optimal deployment, spatial estimation

1 Introduction

This paper is motivated by optimal spatial sampling problems under possibly failing communications. Consider a group of mobile robotic sensors that take point measurements of a random field over an environment and relay them back to a data fusion center. Assume that because of the features of the medium and the limited agent communication capabilities, it is known that only a fraction of these packets will arrive at the center, but it is not a priori known which ones will. Given that some sensors are not working and their identity is unknown, a reasonable strategy consists of grouping sensors together into clusters so that the likelihood of obtaining a measurement from the position of each cluster is higher. In this paper, our aim is to design a distributed algorithm that makes the network autonomously create groups of a desired size such that (i) members of each individual group become coincident, and (ii) the groups deploy optimally with regards to the spatial estimation objective.

Literature review: There is an increasing body of re-

search that deals with spatial estimation problems with possibly failing communications where packets are either received without corruption or not received at all, see e.g., [Smith and Seiler, 2003, Schenato et al., 2007, Gupta et al., 2009, Cortés, 2012]. In particular, Cortés [2012] shows that, for the problem motivating our algorithm design, the clustering strategy outlined above is optimal in some cases: the configurations that maximize the expected information content of the measurements retrieved at the center correspond to agents grouping into clusters, and the resulting clusters being deployed optimally. Achieving such desirable configurations is challenging because of the spatially distributed nature of the problem and the agent mobility. In this regard, the notions of spatial coverage and agent clustering (the latter understood as physical co-location), as well as our proposed algorithmic solution, are different from those in typical hierarchical clustering problems, see e.g. [Younis and Fahmy, 2004, Bandyopadhyay and Coyle, 2003], where sensors are static and the objective is to minimize the cost incurred when relaying messages to a data fusion center. Closer to our setup, Heo and Varshney [2003] define clusters as groups of mobile sensors in locations such that their density is above the expected average density. Using a control law based on whether sensors are in a cluster or not, the network minimizes the distance traveled by the sensors to deploy. Our tech-

* Parts of this paper were presented in preliminary form at the 2011 IEEE Conference on Decision and Control.

Email addresses: miouimet@ucsd.edu (M. Ouimet), cortes@ucsd.edu (J. Cortés).

nical approach combines elements of spatial facility location [Okabe et al., 2000], rendezvous and deployment of multi-agent systems [Bullo et al., 2009], and coalition formation games [Bogomolnaia and Jackson, 2002, Banerjee et al., 2001]. From a game-theoretic perspective, our analysis of the coalition formation dynamics is novel because of the consideration of evolving and partial interaction topologies. From a motion coordination perspective, the novelty relies on the coupled dynamics between the coalition formation, the clustering, and the network deployment. Other works in cooperative control employ game-theoretic ideas to solve tasks such as formation control, target assignment, self-organization for efficient communication, consensus, and sensor coverage, see e.g. [Gu, 2008, Marden et al., 2009a, Arslan et al., 2007, Saad et al., 2011]. Given the algorithmic design choice of the agents’ utility function, our work has connections to weakly acyclic games [Marden et al., 2009b, 2007]. Specifically, under a fixed, complete communication graph where all agents can join any coalition they wish, our game can be cast as a weakly acyclic game. However, in general, the limited information available to agents, the dynamic interaction topology, and the dependence of the individual action sets on this topology makes the framework of weakly acyclic games not directly applicable. We build on the well-established notions of time and communication complexity in distributed algorithms [Peleg, 2000, Lynch, 1997, Bullo et al., 2009] to characterize the performance of our. Since in the repeated coalition formation game agents take probabilistic actions, we consider the expected time complexity. In principle, our algorithm can be described as a Markov chain, where the coalition formation time can be exactly defined as the first hitting time for the set of goal states [Meyn and Tweedie, 2009, Lawler, 2006]. However, defining the probabilistic transition function becomes difficult as the number of total agents grows. Thus, we adopt a drift analysis approach [Hajek, 1982] to provide an upper bound on the time complexity.

Statement of contributions: The main contribution of the paper is the design and analysis of the COALITION FORMATION AND DEPLOYMENT ALGORITHM. The aim of this synchronous and distributed strategy is to allow robotic agents to autonomously form groups of a given desired size while clustering together and deploying optimally in the environment. The deployment objective is encoded through a locational optimization function whose optimizers correspond to circumcenter Voronoi configurations. The algorithm design combines a repeated game component that governs the dynamics of coalition formation with a spatial motion component that determines how agents’ positions evolve. In the coalitional game, agents take probabilistic actions and seek to join a neighboring coalition that most closely resembles one with the desired size. According to the motion coordination law, agents not yet in a well-formed coalition cluster together, while agents in a coalition of the desired size also move towards the circumcenter of their Voronoi cell. Our

main result, cf. Theorem 5.1, establishes that, for a large class of probabilistic coalition switching laws, the executions of the COALITION FORMATION AND DEPLOYMENT ALGORITHM converge in finite time to a configuration where agents are coincident with their own coalition and all coalitions are the desired size, and asymptotically converge to an optimal deployment configuration, each with probability 1. For a specific probabilistic coalition switching law, termed PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS, we provide upper bounds on the expected coalition formation time under arbitrary and complete communication topologies. For any switching law, we also upper bound the total number of messages sent per timestep during an execution on an arbitrary communication topology. The algorithm does not require the agents to have a common reference frame, and is robust to agent addition and deletion. Finally, we illustrate the correctness, robustness, and time and communication complexity results in simulation.

Organization: Section 2 presents basic notions from computational geometry, probability, and hedonic coalition games. Section 3 states the problem setup, and Section 4 contains the description of our algorithm. Section 5 analyzes its correctness and Section 6 characterizes its complexity. Section 7 illustrates our results. Finally, Section 8 contains conclusions and ideas for future work.

2 Preliminaries

We present facts from computational geometry, probability, and coalition games that are key in the discussion.

2.1 Basic geometric notions

We denote by \mathbb{R} , $\mathbb{R}_{>0}$, $\mathbb{R}_{\geq 0}$, \mathbb{Z} , and $\mathbb{Z}_{\geq 1}$ the sets of real, positive real, nonnegative real, integer, and positive integer numbers, respectively. Let $\|\cdot\|$ be the Euclidean distance. Given a set $S \subset X$, let $\mathbb{F}(S)$ denote the collection of finite subsets of S , $S^c = X \setminus S$ its complement, and $|S|$ its cardinality. Let $\text{vr} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be defined by $\text{vr}(u) = u/\|u\|$ for $u \in \mathbb{R}^d \setminus \{0\}$, and $\text{vr}(0) = 0$. We let $B(x, r) = \{p \in \mathbb{R}^d \mid \|x - p\| \leq r\}$. The circumcenter of a set of points P , denoted $\text{CC}(P)$, is the center of the ball of minimum radius, denoted $\text{CR}(P)$, which encloses all points in P . Next, we define the **get-together-toward-goal** function $\text{gttg} : S \times \mathbb{F}(S) \times S \rightarrow S$ that will help us later to get a set of points P closer to each other while moving towards a goal q . Define $\text{gttg}(p, P, q) = p + w_1 + w_2$, where we use the shorthand notation $P_0 = P \cup \{p\}$,

$$\begin{aligned} w_1 &= \min\{\|\text{CC}(P_0) - p\|, d_1(r)\} \text{vr}(\text{CC}(P_0) - p), \\ w_2 &= \min\{\|q - (p + w_1)\|, d_2(r)\} \text{vr}(q - (p + w_1)), \end{aligned}$$

and $r = \text{CR}(P_0)/\|q - \text{CC}(P_0)\|$. Here, $d_1 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is an increasing function, continuous on $\mathbb{R}_{>0}$, satisfying

$$d_1(0) = 0, \quad \lim_{s \rightarrow \infty} d_1(s) = d_{\max}, \quad \lim_{s \rightarrow 0^+} d_1(s) = d_{\min},$$

for $d_{\max} > d_{\min} > 0$, and $d_2 : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{\geq 0}$ is defined by $d_2(s) = d_{\max} - d_1(s)$. Figure 1 illustrates the definition of gttg . Appendix A gathers some its relevant properties.

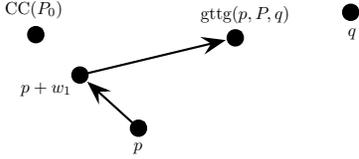


Fig. 1. Illustration of the action of the function gttg .

2.2 Voronoi partitions and deployment objective

Here, we introduce some computational geometric notions that play an important role in the formalization of the deployment problem. Given $Q \subset \mathbb{R}^d$ and a finite set of points $P = \{p_1, \dots, p_N\} \subset Q$, the Voronoi partition $V(P) = \{V_1(P), \dots, V_N(P)\}$ of Q is defined by

$$V_i(P) = \{q \in Q \mid \|q - p_i\| \leq \|q - p_j\|, \forall p_j \in P\}.$$

Note that $V_i(P)$, the Voronoi cell of p_i , is the set of points in Q closer to p_i than to any of the other points in P . The points p_i and p_j are (Voronoi) neighbors if the boundaries of their cells intersect. To compute the Voronoi cell of p_i , all that is required is the location of its neighbors in P . Cortés et al. [2004] introduces a procedure, that we term the ADJUST RADIUS strategy, which does the following: starting from $r = 0$, it repeatedly grows r until all Voronoi neighbors of p_i are guaranteed to be contained in $B(p_i, r)$.

Given a partition $\{W_1, \dots, W_N\}$ of Q , the disk-covering function $\mathcal{H}_{\text{DC},N}$ is defined by

$$\mathcal{H}_{\text{DC},N}(p_1, \dots, p_N, W_1, \dots, W_N) = \max_{i \in \{1, \dots, N\}} \max_{q \in W_i} \|q - p_i\|_2.$$

The value of $\mathcal{H}_{\text{DC},N}$ solves the following problem: cover the whole environment with balls centered at the points in $P = \{p_1, \dots, p_N\}$ with minimum common radius such that $W_i \subset B(p_i, r)$, for $i \in \{1, \dots, N\}$. For convenience, we set $\mathcal{H}_{\text{DC},N}(p_1, \dots, p_N) = \mathcal{H}_{\text{DC},N}(p_1, \dots, p_N, V_1, \dots, V_N)$. Two properties are worth noting [Bullo et al., 2009]: for a fixed configuration, the Voronoi partition is optimal among all partitions,

$$\mathcal{H}_{\text{DC},N}(p_1, \dots, p_N, V_1(P), \dots, V_N(P)) \leq \mathcal{H}_{\text{DC},N}(p_1, \dots, p_N, W_1, \dots, W_N).$$

For a fixed partition the cells' circumcenters are optimal:

$$\mathcal{H}_{\text{DC},N}(\text{CC}(W_1), \dots, \text{CC}(W_N), W_1, \dots, W_N) \leq \mathcal{H}_{\text{DC},N}(p_1, \dots, p_N, W_1, \dots, W_N).$$

Under certain technical conditions, see [Graham and Cortés, 2009], optimizing $\mathcal{H}_{\text{DC},N}$ corresponds to minimizing the maximum error variance in the estimation

of a random spatial field. Similarly, in the situation of discrete events happening with equal likelihood everywhere in an environment, placing sensors at an optimizer of $\mathcal{H}_{\text{DC},N}$ yields a configuration which minimizes the worst-case distance (error) between an event and the nearest sensor. The deployment objective function that motivates our algorithm is given by

$$\mathcal{H}_{N,g}(p_1, \dots, p_N) = \frac{1}{\binom{N}{g}} \sum_{\{s_1, \dots, s_g\} \in C(N,g)} \mathcal{H}_{\text{DC},g}(p_{s_1}, \dots, p_{s_g}), \quad (1)$$

where $C(N, g)$ denotes the set of unique g -sized combinations of elements in $\{1, \dots, N\}$. This function corresponds to the expected disk-covering performance of a network of N agents where only g of them are working and their identity is unknown. Optimizers of $\mathcal{H}_{N,g}$ correspond to grouping agents into coincident clusters of a specific size, say κ , that themselves are optimally deployed according to $\mathcal{H}_{\text{DC}, \lceil \frac{N}{\kappa} \rceil}$, see [Cortés, 2012]. The cluster size κ is a function of N , g , and Q . For instance, for the case where Q is an interval, if only 1 agent is expected to be working correctly, all agents should form one coalition of size N . If 2 agents are expected to function, the optimal coalition size is $N/2$. Finally, if $N - 1$ are expected to function, the optimal coalition size is 2. In this paper, we assume that the optimal cluster size κ is known, and so forming coincident clusters of size κ and deploying these groups appropriately optimizes (1).

2.3 Probability notions

Here we gather some probability notions from Rosenthal [2000], Billingsley [1995]. Let X be a random variable that has outcomes $\{x_1, x_2, \dots\}$ with probabilities $\{p_1, p_2, \dots\} \subset \mathbb{R}_{\geq 0}$. An event E is a set of outcomes of X . For brevity, we use $\text{P}(E) = \text{P}(X \in E)$. Given a sequence of events $\{E_n\}_{n=1}^{\infty}$, let

$$\begin{aligned} \limsup_n E_n = \{E_n \text{ i.o.}\} &= \bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} E_k, \\ \liminf_n E_n = \{E_n \text{ a.a.}\} &= \bigcup_{n=1}^{\infty} \bigcap_{k=n}^{\infty} E_k. \end{aligned}$$

Here 'i.o.' stands for infinitely often, and 'a.a.' stands for almost always. Note that $\{E_n \text{ i.o.}\}^c = \{E_n^c \text{ a.a.}\}$.

Lemma 2.1 (Borel-Cantelli Lemma) *Given a sequence of events $\{E_n\}_{n=1}^{\infty}$ satisfying $\sum_{n=1}^{\infty} \text{P}(E_n) < \infty$. Then $\text{P}(\limsup_n E_n) = 0$.*

2.4 Hedonic coalition games

Hedonic coalition formation games [Bogomolnaia and Jackson, 2002] are N -player noncooperative games [Fudenberg and Tirole, 1991, Başar and Oldser, 1982] where

players attempt to join/stay in preferable coalitions. Each player is hedonic because the utility it assigns to a given network coalition partitioning is only a function of its own coalition. Each player’s action set is finite: it can stay in the current coalition or join another coalition. For a finite set of players $\mathcal{A} = \{1, \dots, N\}$, a finite coalition partition is a set $\Pi = \{S_k\}_{k=1}^K$, $K \in \mathbb{Z}_{\geq 1}$, that partitions \mathcal{A} . The subsets S_k are called coalitions. For player i and partition Π , let $S_{\Pi}(i)$ be the set $S_k \in \Pi$ such that $i \in S_k$. Agent i ’s preference is defined by an ordering \succeq_i over the set $\mathfrak{S}_i = \{S \in \mathbb{F}(\mathcal{A}) \mid i \in S\}$. A coalition partition Π is called Nash stable if, for each $i \in \mathcal{A}$,

$$S_{\Pi}(i) \succeq_i S_k \cup \{i\}, \quad \forall S_k \in \Pi \cup \{\emptyset\}. \quad (2)$$

In coalition formation games, a player has full information about which coalitions all other players are in and may join any of them. This is in contrast to our scenario, where coalition information is only partial due to the limited capabilities of individual agents. Let us introduce definitions which help capture the spatially-limited nature of coalition information. We say that (S_1, \dots, S_N) is a *consistent* coalition state if $i \in S_i$ and $S_j = S_i$, for each $j \in S_i$, for each $i \in \mathcal{A}$. Note that for a consistent coalition state, $\{S_1, \dots, S_N\}$ reduces to a finite coalition partition of \mathcal{A} . Let $\tau_i \subseteq \mathcal{A}$ denote the set of agents whose coalition information i has access. Letting $S_0 = \emptyset$, the function best-set defines the players whose coalitions i most prefers to be a member of,

$$\text{best-set}(\succeq_i, \{(k, S_k)\}_{k \in \tau_i}) = \{j \in \tau_i \cup \{0\} \mid S_j \cup \{i\} \succeq_i S_k \cup \{i\}, \forall k \in \tau_i \cup \{0\}\}.$$

3 Problem statement

A group of robotic sensors with unique identifiers $\mathcal{A} = \{1, \dots, N\}$ moves in a convex polygon $Q \subset \mathbb{R}^2$. Let p_i denote the location of agent i and $P = (p_1, \dots, p_N)$ denote the overall network configuration. We consider arbitrary agent dynamics, assuming each agent can move up to a distance $d_{\max} \in \mathbb{R}_{>0}$ within one timestep,

$$p_i(\ell + 1) \in B(p_i(\ell), d_{\max}), \quad \ell \in \mathbb{Z}.$$

Through either sensing or communication, we assume each agent i can get the relative position and identity of agents within distance $r_i \in \mathbb{R}_{>0}$. During the coalition formation process, agents can interact with other agents within this radius. Agent i can adjust r_i but the cost of acquiring information is an increasing function of it. Inter-agent communication occurs instantaneously.

Given the problem scenario described in Section 1, the network’s objective is dual. On the one hand, agents want to cluster into groups of a predefined size κ . Equivalently, the network wants to self-assemble into $\lfloor \frac{N}{\kappa} \rfloor$ clusters of size κ , with possibly one additional cluster of size z , $0 \leq z < \kappa$, with $N = \lfloor \frac{N}{\kappa} \rfloor \kappa + z$. On the other hand,

the resulting clusters should be positioned in the environment so as to minimize $\mathcal{H}_{\text{DC}, \lfloor \frac{N}{\kappa} \rfloor}$. As discussed in Section 2.2, such deployments correspond to optimizers of (1) for a class of spatial estimation problems with unreliable sensors. For convenience, we define a partition to be a *goal coalition partition* if the cardinality of m of its coalitions is κ , with the cardinality of the remaining one equal to z , if it exists.

A trivial solution to this problem would be to first elect $\lfloor \frac{N}{\kappa} \rfloor$ leaders and have each leader recruit $\kappa - 1$ followers. Then each group could rendezvous, and afterwards, the overall network would deploy. However, this method is neither distributed nor robust to agent failures. Our aim is to create a distributed algorithm that accomplishes the dual network objective in a robust and efficient way.

4 Coalition formation and deployment

Here, we solve the problem posed in Section 3 with the COALITION FORMATION AND DEPLOYMENT ALGORITHM. This distributed, synchronous strategy specifies for each agent the dynamics of coalition formation and spatial motion. Section 4.1 outlines the logic used by agents to determine which coalition to join as well as the supporting inter-agent communication and Section 4.2 discusses how agents decide to move depending on their coalition size and the deployment objective.

Before specifying the dynamics, we describe the required memory of each agent and appropriate initializations. The memory \mathcal{M}_i of agent i is composed of

- the *coalition set* C_i . Elements of this set are of the form (j, p_j) , i.e., identity and position of the member. For convenience, we set $(i, p_i) \in C_i$ and $C_0 = \emptyset$;
- the *communication radius* r_i at which the agent interacts with other agents not necessarily in its coalition;
- the *neighboring set* \mathcal{N}_i corresponding to agents within distance r_i , i.e., $(j, p_j) \in \mathcal{N}_i$ iff $p_j \in B(p_i, r_i)$;
- the *farthest-away radius* \mathbf{r}_i , corresponding to the maximum distance to members of its coalition set.
- the *flag last*, which indicates if an agent belongs to the single final coalition not of size κ when $\lfloor \frac{N}{\kappa} \rfloor \neq \frac{N}{\kappa}$.

The operators $\text{id}(\cdot)$ and $\text{pos}(\cdot)$ extract identities and positions, respectively, from sets with elements of the form (i, p_i) . Initialization requires a consistent coalition state $(\text{id}(C_1), \dots, \text{id}(C_N))$, $r_i \in \mathbb{R}_{\geq 0}$, and $\text{last} = \text{False}$.

The vocabulary all agents can recognize are

- an agent sends the word *query* to ask for the identities of another agent’s current coalition;
- an agent sends a packet with the word *leave/join* along with an agent identity to indicate that the agent is leaving/joining the recipient’s coalition.

Remark 4.1 (Communication protocol) The communication radius r_i should be thought of as the distance at which agent i must interact with other agents not nec-

essarily in its coalition. We do not enter into the specifics of how this interaction is actually implemented. This might be through direct, one-hop communication or, if the radius is large, through indirect, multi-hop routing involving other network agents. Lemma 4.3 below ensures the radius is kept, at each timestep, at the smallest value that guarantees successful coalition formation. •

4.1 Coalition formation game

The formation of coalitions evolves according to a simultaneous-action hedonic coalition game with partial information. Let us start with an informal description.

[Informal description]: The agents' objective is to be in a κ -sized coalition. There are two rounds of communication per timestep. In the first one, each agent acquires information to determine if any neighboring coalition is more attractive than its current one. In the second one, the agents involved in a coalition change (either because they have decided to switch or because someone else decided to join their coalition) exchange information to update the coalition membership.

Next, we formally describe the hedonic coalition formation game. The agent i 's preference ordering \succ_i over \mathfrak{S}_i is

$$\begin{aligned} \{S \in \mathfrak{S}_i \mid |S| = \kappa\} &\succ \{S \in \mathfrak{S}_i \mid |S| = \kappa - 1\} \succ \dots \\ &\succ \{S \in \mathfrak{S}_i \mid |S| = 1\} \succ \{S \in \mathfrak{S}_i \mid |S| = \kappa + 1\} \succ \dots \\ &\succ \{S \in \mathfrak{S}_i \mid |S| = N\}. \end{aligned} \quad (3)$$

According to (3), agents most prefer to be in κ -sized coalitions. They also prefer to be in a coalition of size 1 over any coalition of size larger than κ .

Algorithm 1 BEST NEIGHBOR COALITION DETECTION

Executed by: agents i if $|C_i| \neq \kappa$

- 1: Acquire \mathcal{N}_i (get location of neighbors)
- 2: Send (query, r_i) at r_i to $\text{id}(\mathcal{N}_i \setminus C_i)$
- 3: Receive $\text{id}(C_j)$ from all $j \in \text{id}(\mathcal{N}_i \setminus C_i)$
(request/receive coalition sizes)
- 4: **if** $i \notin \text{best-set}(\succ_i, \{k, \text{id}(C_k)\}_{k \in \text{id}(\mathcal{N}_i)})$ **then**
- 5: **with probability** p **do**
- 6: Set j^* from $\text{best-set}(\succ_i, \{k, \text{id}(C_k)\}_{k \in \text{id}(\mathcal{N}_i)})$
(identify best coalition to join)
- 7: **if** $j^* \neq 0$ **then**
- 8: $r_i := \|p_{j^*} - p_i\|$
- 9: **end if**
- 10: **end**
- 11: **end if**

Next, we specify the two rounds of communication that take place per timestep. Agents who already are in a coalition of size κ do not actively take part in this process; they only respond to other agents' messages. First, agents execute the BEST NEIGHBOR COALITION DETECTION strategy described as Algorithm 1. According to this strategy (cf. step 5), an agent that finds a neighboring coalition better than its own will decide to join

it with some probability of the form

$$p = \begin{cases} f(|C_1|, \dots, |C_N|, N, \kappa), & \text{if } |C_i| \neq \kappa, \\ 0, & \text{if } |C_i| = \kappa, \end{cases} \quad (4)$$

where the function f takes values in the interval $(0, 1)$ for all finite N . The form of f affects the convergence rate of the algorithm, which we will investigate later.

Remark 4.2 (Justification and tradeoffs for probabilistic actions) The probabilistic model for actions described in (4) helps avoid deadlock situations that may result from the decentralized nature of the game. As an example, in a situation with two groups of size $\kappa - 1$, all agents will desire to join the other group. In such case, a group of size κ would never form. Instead, under (4), there is a positive probability that agents in only one of the groups act, breaking the deadlock. In contrast with a one-agent-acting-per-timestep policy, (4) allows multiple agents to switch coalitions at the same timestep. One tradeoff of probabilistic actions is that the identities of the agents in each coalition cannot be known a priori. Another tradeoff is that, at times, agents do not make the most beneficial action to achieve the ultimate objective given a specific configuration. However, as stated above, it is precisely this occasional sub-optimality that helps eliminate deadlock situations. •

Next, all agents run the COALITION SWITCHING strategy found in Algorithm 2. This strategy builds on the input j^* provided to i by the BEST NEIGHBOR COALITION DETECTION strategy. Agents with $j^* \neq i$ switch coalitions. If $j^* = 0$, i forms its own coalition. Otherwise, i interacts with agent j^* to join its coalition. After switching, agents update coalition memberships and the communication radii required to determine the position of other members so that the coalition state remains consistent.

4.2 Motion control law

Here, we describe how agents move at each timestep, beginning with an informal description:

[Informal description]: At each timestep, agents adjust their communication radius and move. Both actions depend on the size of their coalition. Agents not yet in a coalition of size κ increase their radius to improve the chances of finding a better coalition and move towards their coalition members. Agents in a coalition of size κ adjust their radius to ensure they can calculate their Voronoi cell and move towards both their coalition members and the cell circumcenter.

The RADIUS ADJUSTMENT AND MOTION strategy is formally described in Algorithm 3. Its interaction with the coalition formation dynamics is described in steps 10-16, which governs the set of agents that a robot not yet in a κ -sized coalition interacts with. The next result ensures that agents' communication radius are kept at the smallest values that guarantee successful coalition formation.

Algorithm 2 COALITION SWITCHING

Executed by: all agents i

```
1: if  $j^* \neq i$  then
2:   Send (leave,  $i$ ) at  $\tau_i$  to  $\text{id}(C_i)$ 
                                     (alert old coalition)
3:   if  $j^* \neq 0$  then
4:     Send (join,  $i, r_i$ ) at  $r_i$  to  $j^*$ 
                                     (alert new coalition)
5:   end if
6: end if
7:  $M := \{k \in \mathcal{A} \mid i \text{ received join from } k\}$ 
   (agents relying on  $i$  to aid switching)
8: foreach  $m \in M$ , send (join,  $m, r_m$ ) to  $\text{id}(C_i)$ 
   (alert coalition members of  $m$  via  $\tau_i$ )
9:  $L := \{k \in \mathcal{A} \mid i \text{ received leave from } k\}$ 
   (agents leaving coalition)
10:  $J := \{k \in \mathcal{A} \mid \text{an } m \in \text{id}(C_i) \text{ got join from } k\}$ 
   (agents leaving/joining  $i$ 's coalition)
11:  $\text{id}(C_i) := (\text{id}(C_i) \cup J) \setminus L$  and  $r_i := \tau_i + \max\{r_j\}_{j \in J}$ 
   (update current coalition and radius)
12: foreach  $m \in M$ , send ( $r_i, \text{id}(C_i)$ ) at  $r_m$  to  $m$ 
   (update agents joining  $i$ 's coalition)
13: if  $j^* \neq i$  then
14:   if  $j^* = 0$  then
15:      $C_i = \{(i, p_i)\}$  (form a new coalition)
16:   else
17:      $\text{id}(C_i) := \text{id}(C_{j^*})$  and  $r_i := \|p_{j^*} - p_i\| + r_{j^*}$ 
                                     (update coalition and radius)
18:   end if
19: end if
20: if  $J \neq \emptyset \vee j^* \neq i$  then
21:   Acquire  $\mathcal{N}_i, \text{pos}(C_i)$ 
22:    $j^* := i$  (reset switching variable)
23: end if
```

Lemma 4.3 (Optimality property for communication radius law) For each $i \in \mathcal{A}$ such that $|C_i| \neq \kappa$, let k_i be the closest agent which is in a coalition different from i 's with size different from κ . Let $r_i(P, (C_1, \dots, C_N)) = \|p_i - p_{k_i}\|$. For consistent coalition states not corresponding to a goal coalition partition, such radii guarantee that at least one agent has an incentive to switch coalitions. Moreover, if the radii of these agents were set according to any other function r'_i with $r'_i(P, (C_1, \dots, C_N)) < r_i(P, (C_1, \dots, C_N))$ for some i and P , then this property is no longer guaranteed.

PROOF. If there is at least one coalition of size greater than κ , all agents in this coalition have an incentive to start their own coalition. Consider instead, the case where all coalitions are of size at most κ . An agent i in the smallest coalition has an incentive to join its neighbor k_i and the claimed property follows. Next, we show the minimality property. It is enough to show that there is one non-goal consistent coalition state for which a smaller radius assignment would not work. Consider a consistent coalition state at configuration P where all coalitions but one have been formed, and the re-

Algorithm 3 RADIUS ADJUSTMENT AND MOTION

Executed by: all agents i

```
1: if  $|C_i| = \kappa \vee \text{last} = \text{True}$  then
2:   Update  $r_i$  with ADJUST RADIUS strategy
3:   Acquire  $\mathcal{N}_i$ 
4:    $A_i := (\{\text{CC}(\text{pos}(C_i))\} \cup \text{pos}(\mathcal{N}_i)) \setminus \text{pos}(C_i)$ 
5:    $V_i := V_1(A_i)$  (compute Voronoi cell)
6:   goal =  $\text{CC}(V_i)$ 
7: else
8:   goal =  $\text{CC}(\text{pos}(C_i))$ 
9:    $\mathcal{C}_i := \{j \in \text{id}(\mathcal{N}_i) \mid |\text{id}(C_i)| = \kappa\}$ 
10:  if  $\text{id}(\mathcal{N}_i \setminus C_i) \setminus \mathcal{C}_i \neq \emptyset$  then
11:     $r_i := \min_{k \in \text{id}(\mathcal{N}_i \setminus C_i) \setminus \mathcal{C}_i} \|p_k - p_i\| + 2d_{\max}$ 
                                     (guarantees a neighbor after motion)
12:  else if  $\text{id}(\mathcal{N}_i) = \mathcal{A}$  then
13:    last := True (one non- $\kappa$  coalition)
14:  else
15:     $r_i := r_i + \delta$  (increase radius)
16:  end if
17: end if
18: foreach  $j \in \text{id}(C_i)$ , set  $p_j := \text{gttg}(p_j, \text{pos}(C_i), \text{goal})$ 
   (compute next position)
19:  $\text{pos}(C_i) := \{p_j\}_{j \in \text{id}(C_i)}$  (update positions)
20:  $\tau_i := \max_{p_j \in \text{pos}(C_i)} \|p_j - p_i\|$  (recompute radius)
```

maining agents are in two coalitions, one with the single agent, i , and the other one, C , with the rest. Since $r'_i(P, (C_1, \dots, C_N)) < r_i(P, (C_1, \dots, C_N)) = \|p_i - p_{k_i}\|$, agent i has no agents in \mathcal{N}_i that it has incentive to join. Furthermore, given the coalition state, agent i is the only one who could have an incentive to switch coalitions, concluding the proof. \square

Steps 10-16 in Algorithm 3 implement the result of Lemma 4.3. If agent i is not within r_i of a non-coalition agent that is not in a κ -sized coalition, increase r_i . If agent i is within r_i of such an agent, change r_i to the distance between the two agents plus a constant ensuring they remain within communication range after moving.

Remark 4.4 (Voronoi cell computation) In the Voronoi cell computation of step 5 in Algorithm 3, the coalition's circumcenter replaces all the locations of the individual agents which ensures that all coalition members compute the same cell. However, this also implies that the collection of cells computed by the coalition is not a partition of the environment. This issue gets resolved when the members within each coalition are coincident and is treated in the proof of Theorem 5.1. \bullet

Remark 4.5 (Choice of parameter δ) In step 15 of Algorithm 3, δ describes the amount that r_i increases if i does not have any neighboring candidate agents to join. Several choices are possible. For instance, when agents are roughly uniformly distributed, $\delta \propto \frac{\text{diam}(Q)}{\sqrt{N}}$ makes it likely that the agent discovers at least one new agent. \bullet

The COALITION FORMATION AND DEPLOYMENT ALGORITHM is composed of Algorithms 1-3. This strategy does not require agents to share a common reference frame.

Remark 4.6 (Robustness to addition and subtraction) The COALITION FORMATION AND DEPLOYMENT ALGORITHM is robust to agents joining or leaving the network under the following assumptions: (i) new agents alert the network of their presence by sending a query message, (ii) when an agent fails, the other members of its coalition detect this fact, and (iii) when agents receive a query message they set `last` := `False`. •

5 Correctness analysis

This section analyzes the convergence properties of the strategy designed in Section 4. Our main objective is to establish the following result.

Theorem 5.1 (Algorithm correctness) Consider a network of N agents executing the COALITION FORMATION AND DEPLOYMENT ALGORITHM. Then,

- (i) there exists a finite time after which the agents are in a goal coalition partition and each is coincident with its coalition members, with probability 1;
- (ii) the agents' positions and the induced Voronoi partition asymptotically converges toward the set of minimizers of $\mathcal{H}_{\text{DC}, \lceil \frac{N}{\kappa} \rceil}$, with probability 1.

The proof of this result requires us to establish several intermediate results. Theorem 5.1 states that, with probability 1, the network will not converge to a coalition partition other than the desired one. Agents may stay for some time in a different partition but in finite time they will reach the desired partition with probability 1. This can be traced back to the fact that, in the simplified coalition formation game where agents have both full information and action sets, only the goal coalition partition is Nash stable (see Lemma 5.2 below). Theorem 5.1 implies that, even in the absence of global information, the Nash stable partitions are the desired ones.

Lemma 5.2 (Nash stable partitions of preference ordering) In the N -agent simultaneous-action game where agents have preference orderings satisfying (3), complete knowledge of all other coalition memberships, and their action set is to stay or join any other coalition, only the goal coalition partition is Nash stable.

PROOF. First, let us show that the goal coalition partition is Nash stable. All the agents in coalitions of size κ receive maximal utility, so they satisfy (2). Nash stability follows by noting that the agents in the coalition of size z do not prefer to join either a coalition with size κ or start a new coalition. To show the uniqueness result, we reason with a different arbitrary partition and show it is not Nash stable. This arbitrary partition must have

either at least one coalition with more than κ agents or at least two coalitions with less than κ agents. In the first case, the agents in a coalition with more than κ agents would benefit by joining any coalition with size less than κ (if any exists), and if not, by forming a new coalition. In the second case, the agents in the smallest coalition with less than κ agents would benefit from joining the other. The same argument holds if coalitions are tied for smallest. These two cases show the goal coalition partition is the only Nash stable partition. □

5.1 Analysis of coalition formation dynamics

We define the collection of actions of all agents at a given timestep as a *timestep-event*. Our first result finds a strictly positive lower bound on the probability of any possible timestep-event happening. The result follows by noting that all agents' probabilistic actions are independent and the switching probabilities are given by (4).

Lemma 5.3 (Bound on switching probability) Let E be a timestep-event with $P(E) > 0$. Then $P(E) \geq \min\{f(|C_1|, \dots, |C_N|, N, \kappa), (1 - f(|C_1|, \dots, |C_N|, N, \kappa))^N\}$.

For $\epsilon > 0$, define

$$\Xi_\epsilon(C_1, \dots, C_N) = \sum_{i \in \mathcal{A}_{\leq \kappa}} \frac{(1 + \epsilon)^{|C_i|}}{|C_i|} = \sum_{j=1}^{\kappa} a_j (1 + \epsilon)^j, \quad (5)$$

where a_j are the number of coalitions of size j . Note that coalitions with size strictly larger than κ do not contribute to Ξ_ϵ . Additionally, Ξ_ϵ is upper bounded by $\lceil N/\kappa \rceil (1 + \epsilon)^\kappa$. The next result establishes that one agent joining a coalition of at least its own current coalition's size has positive effect on the overall network reaching a goal coalition partition.

Lemma 5.4 (Ξ_ϵ increase for one switcher) For any $\epsilon > 0$, when exactly one agent joins a new coalition of at least its current coalition's size, this action strictly increases the function Ξ_ϵ by at least $\min\{1 + \epsilon, \epsilon^2\}$.

PROOF. Let i be the size of the coalition being left and j the size of the coalition being joined. We must consider the cases of $i < \kappa$ and $i > \kappa$ separately, starting with the former. Now, by the coalition preference ordering in (3), $i \leq j < \kappa$. After switching, Ξ_ϵ has changed by

$$\begin{aligned} & (1 + \epsilon)^{i-1} - (1 + \epsilon)^i + (1 + \epsilon)^{j+1} - (1 + \epsilon)^j \\ &= \epsilon((1 + \epsilon)^j - (1 + \epsilon)^{i-1}) \\ &\geq \epsilon((1 + \epsilon)^i - (1 + \epsilon)^{i-1}) \geq \epsilon^2. \end{aligned}$$

Now, considering the case where $i > \kappa$, one can see that in the worst case, the agent forms its own coalition of size 1, increasing Ξ_ϵ by $1 + \epsilon$, which completes the result. □

The next result shows that from any consistent coalition state, there is a finite sequence of timestep-events, each with positive probability of occurring, that leads to a goal coalition partition. Also, there exists an upper bound, independent of the coalition state, on this sequence's length.

Proposition 5.5 (A sequence of switching events leading to the goal coalition partition) *From any consistent coalition state, there exists a finite sequence of timestep-events, each having a positive probability of occurring under the COALITION FORMATION AND DEPLOYMENT ALGORITHM, leading to a goal coalition partition. Furthermore, for any $\epsilon > 0$, the length of this sequence is bounded by*

$$L = \frac{\lceil N/\kappa + 1 \rceil (1 + \epsilon)^\kappa}{\epsilon^2} \left(\frac{\text{diam}(Q)}{\delta} + 1 \right) + \lfloor \frac{N}{\kappa} \rfloor. \quad (6)$$

PROOF. Initially, if any coalitions are larger than size κ , let the first timestep-event E_1 be one where the correct number of agents leave one of these large coalitions and all other agents do not switch, creating a coalition of size κ . From Lemma 5.3, $P(E_1)$ is bounded away from zero. There can be at most $\lfloor \frac{N}{\kappa} \rfloor - 1$ more coalitions larger than size κ , and so $E_2, \dots, E_{\lfloor \frac{N}{\kappa} \rfloor}$ are defined similarly. From step 15 in Algorithm 3, within at most $\frac{\text{diam}(Q)}{\delta}$ timesteps, each agent i will have a radius r_i satisfying Lemma 4.3, so at least one agent has an incentive to change coalitions. In the timesteps in which no agents wish to change coalitions, the corresponding timestep-events, $E_{\lfloor \frac{N}{\kappa} \rfloor + 1}, \dots, E_{\lfloor \frac{N}{\kappa} \rfloor + \alpha}$, $\alpha \leq \frac{\text{diam}(Q)}{\delta}$, occur with probability 1. Define $E_{\lfloor \frac{N}{\kappa} \rfloor + \alpha + 1}$ to be a timestep-event where exactly one agent joins a coalition it has an incentive to and all others do not switch. By Lemma 5.3, the probability of this event is bounded away from zero. Additionally, because all coalitions are at most size κ , the function Ξ_ϵ increases by at least ϵ^2 (c.f. Lemma 5.4). If the configuration is not in a goal coalition partition, within at most $\frac{\text{diam}(Q)}{\delta}$ timesteps, at least one agent will have an incentive to switch coalitions. Because the upper-bounded function Ξ_ϵ monotonically increases each time this sequence of timestep-events occurs, the number of times this can occur is at most $\frac{\lceil N/\kappa + 1 \rceil (1 + \epsilon)^\kappa}{\epsilon^2}$. Therefore, within L timesteps (cf. (6)), the agents will be in a goal coalition partition. \square

The next result uses the sequence constructed in Proposition 5.5 to show that in finite time all agents are in a goal coalition partition, with probability 1.

Theorem 5.6 (Finite-time convergence to goal coalition partition) *There exists a finite time after which N agents using the COALITION FORMATION AND DEPLOYMENT ALGORITHM are in a goal coalition partition with probability 1.*

PROOF. Lemma 5.3 asserts that the probability of a timestep-event occurring is lower bounded by $\rho = \min\{f(|C_1|, \dots, |C_N|, N, \kappa), (1 - f(|C_1|, \dots, |C_N|, N, \kappa))\}^N$. Given an initial consistent coalition state,

Proposition 5.5 guarantees that there exists a finite sequence of timestep-events, whose length is upper bounded by L (cf. (6)), leading to the goal coalition partition. If the length of this sequence is smaller than L , this sequence can be extended to one of exactly length L by considering additional timestep-events where no agents wish to change coalitions. The latter occur with probability 1. Therefore, the sequence of timestep-events leading to a goal coalition partition has a probability of occurring of at least ρ^L , independent of the initial coalition state.

Define a sequence of events $\{A_1, A_2, \dots\}$, where A_n is the event that the coalitions do not exist after nL timesteps. The probability of A_n occurring is at most $(1 - \rho^L)^n$. Now,

$$\sum_{n=1}^{\infty} P(A_n) \leq \sum_{n=1}^{\infty} (1 - \rho^L)^n < \infty,$$

since it corresponds to a convergent geometric series. Thus, by the Borel-Cantelli Lemma, cf. Lemma 2.1, $P(\{A_n \text{ i.o.}\}) = 0$. This means $P(\{A_n \text{ i.o.}\}^c) = 1$ or, equivalently, $P(\{A_n^c \text{ a.a.}\}) = 1$. The result follows by noting that A_n^c is the event that the coalitions occur at some point in nL timesteps and $\{A_n^c \text{ a.a.}\}$ is the event that all but a finite number of events A_n^c occur. \square

5.2 Proof of the main result

We are now ready to prove Theorem 5.1.

PROOF. [Proof of Theorem 5.1] In statement (i), the fact that there exists, with probability 1, a finite time after which all agents are in a goal coalition partition follows from Theorem 5.6. Proposition A.2 allows us to upper bound the number of timesteps it takes for the circumradius of one of these coalitions to vanish by $\lceil \frac{\text{diam}(Q)}{d_{\min}} \rceil$. This implies the fact that in finite time agents become coincident with its coalition members. Once coalitions form and all individual agents are coincident with the members of their respective coalitions, the collection of Voronoi cells that the agents compute correspond to a correct Voronoi partition with $\lceil \frac{N}{\kappa} \rceil$ generators. Statement (ii) then follows from [Bullo et al., 2009, Theorem 5.5]. \square

6 Algorithm complexity analysis

This section investigates the time and communication complexity per timestep of the COALITION FORMATION AND DEPLOYMENT ALGORITHM.

6.1 Time complexity analysis

After having established in Section 5 the correctness of the COALITION FORMATION AND DEPLOYMENT ALGORITHM, here we analyze the expected completion time of the coalition formation dynamics. In general, this time depends on the specific probability law chosen. In this section, we bound the expected completion time for a specific switching probability law that we term PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS. Before specifying this law, let us introduce some useful notation. Given the network state at a certain time, let $N_{\text{left}} \leq N$ denote the number of agents not in a group of size κ at that moment. We assume that each agent i can estimate N_{left}^i within a constant factor of N_{left} , i.e.,

$$N_{\text{left}} \leq N_{\text{left}}^i \leq cN_{\text{left}} \quad (7)$$

uniformly in time, for some $c \in \mathbb{R}_{\geq 1}$. The PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS law is defined as the switching probability given by

$$p = \left(\frac{1}{N_{\text{left}}^i} \right)^{1+\gamma}, \quad \text{if } |C_i| \neq \kappa, \quad (8)$$

where $\gamma > 0$ is a design parameter.

Remark 6.1 (Determination of N_{left}) In the forthcoming analysis, we do not consider a specific way of estimating N_{left} . There are a number of ways to implement this. One possibility is for all agents to initially have an estimate of N . If each time a coalition of size κ is formed, one of the agents pays the one-time broadcast cost to send a message of this to all agents in the environment, all the agents can update N_{left} to $N_{\text{left}} - \kappa$. •

Our strategy to characterize the algorithm's time complexity relies in measuring the effect that agents switching coalitions has on the function Ξ_ϵ , cf (5). When one agent switches coalitions, this action increases Ξ_ϵ . However, when multiple agents switch coalitions at the same time, it is possible that their joint actions decrease Ξ_ϵ . One example of this is when multiple agents join the same coalition, making it larger than size κ . The next result provides an upper bound on how much Ξ_ϵ might decrease when more than one agent switches per timestep. Its proof follows from over-approximating the decrease by removing two coalitions of size $\kappa - 1$ (one for the coalition left and joined) for each agent that is switching.

Lemma 6.2 (Upper bound on decrease in Ξ_ϵ due to multiple switchers) *In a non-goal coalition partition, if exactly $\phi > 1$ switch coalitions, the function Ξ_ϵ does not decrease by more than $2\phi(1 + \epsilon)^{\kappa-1}$.*

Given $N_{\text{left}} \leq N$ agents not yet in a coalition of size κ , the next result shows that the expected number of timesteps until all agents are in coalitions of size κ can be upper bounded by a function of N_{left} .

Lemma 6.3 (Convergence time for $N_{\text{left}} \leq N$ agents) *The expected number of timesteps it takes $N_{\text{left}} \leq N$ agents not yet in a coalition of size κ to all be in κ -sized coalitions is upper bounded by $L_{\text{left}}(cN_{\text{left}})^{N_{\text{left}}(1+\gamma)L_{\text{left}}}$, with*

$$L_{\text{left}} = \frac{\lceil N_{\text{left}}/\kappa + 1 \rceil (1 + \epsilon)^\kappa (\sqrt{cN_{\text{left}}} + 1) + \lfloor \frac{N_{\text{left}}}{\kappa} \rfloor,}{\epsilon^2}$$

under the bound (7), when agents switch using the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS law, and each agent i 's communication parameter is given by $\delta_i = \text{diam}(Q)/\sqrt{N_{\text{left}}^i}$.

PROOF. Following Lemma 5.3, Proposition 5.5, and Theorem 5.6 for the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS switching law, one can define $\rho_{\text{left}} = (\frac{1}{cN_{\text{left}}})^{N_{\text{left}}(1+\gamma)}$ and change (6) to L_{left} to account for the estimate N_{left}^i and δ_i 's dependence on it. Then, the probability that all agents are in the goal coalition after L_{left} timesteps is at least $\rho_{\text{left}}^{L_{\text{left}}}$, independent of the network's state. Thus, the expected value of the first time the network is in the goal coalition is upper bounded by $\frac{L_{\text{left}}}{\rho_{\text{left}}^{L_{\text{left}}}}$, completing the result. □

The upper bound in Lemma 6.3 implies that for $N_{\text{left}} = \mathcal{O}(1)$, the time complexity is also $\mathcal{O}(1)$. We are now ready to upper bound the expected number of timesteps for all coalitions to form under COALITION FORMATION AND DEPLOYMENT ALGORITHM, executed over an arbitrary graph, when the switching probability is defined by (8).

Proposition 6.4 (Time complexity on a generic graph) *Under COALITION FORMATION AND DEPLOYMENT ALGORITHM, the expected number of timesteps for the network to enter the goal coalition partition is $\mathcal{O}(N^{\frac{5}{2}+\gamma})$ when agents switch using the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS law, and each agent i 's communication parameter is given by $\delta_i = \text{diam}(Q)/\sqrt{N_{\text{left}}^i}$.*

PROOF. Let S be the number of agents who wish to switch at a given timestep and s be the number of agents who actually do. Note that

$$P(s = \phi | S = \varphi) = \binom{\varphi}{\phi} p^\phi (1-p)^{\varphi-\phi}.$$

Then, using (8), one can bound

$$P(s = 1 | S = \varphi) \geq \frac{j}{(cN_{\text{left}})^{1+\gamma}} \left(1 - \frac{1}{2^{1+\gamma}}\right)^2, \quad (9a)$$

$$P(s = \phi | S = \varphi) \leq \frac{1}{\phi!} \left(\frac{\varphi}{N_{\text{left}}^{1+\gamma}}\right)^\phi, \quad \forall 1 < \phi \leq N. \quad (9b)$$

For $\epsilon > 0$, using Lemmas 5.4 and 6.2, one can bound the expected change in Ξ_ϵ as a function of the number of agents that wish to switch by

$$\begin{aligned} \mathbb{E}[\Xi_\epsilon(\ell + 1) - \Xi_\epsilon(\ell) | S(\ell) = j] \geq & \min\{\epsilon^2, 1 + \epsilon\} \mathbb{P}(s = 1 | S = \varphi) \\ & - \sum_{\phi=2}^{\varphi} \mathbb{P}(s = \phi | S = \varphi) 2\phi(1 + \epsilon)^{\kappa-1}. \end{aligned}$$

Combining this with (9), we get

$$\begin{aligned} \mathbb{E}[\Xi_\epsilon(\ell + 1) - \Xi_\epsilon(\ell) | S(\ell) = \varphi] \geq & \min\{\epsilon^2, 1 + \epsilon\} \left(1 - \frac{1}{2^{1+\gamma}}\right)^2 \frac{\varphi}{(cN_{\text{left}})^{1+\gamma}} \\ & - 2(1 + \epsilon)^{\kappa-1} \sum_{\phi=2}^{\varphi} \frac{1}{(\phi-1)!} \left(\frac{\varphi}{N_{\text{left}}^{1+\gamma}}\right)^\phi. \end{aligned}$$

Using $\sum_{\phi=2}^{\infty} \frac{1}{(\phi-1)!} p^\phi = p(e^p - 1)$, for all $|p| < 1$, one gets

$$\mathbb{E}[\Xi_\epsilon(\ell + 1) - \Xi_\epsilon(\ell) | S(\ell) = \varphi] \geq \frac{\varphi}{N_{\text{left}}^{1+\gamma}} \left(\frac{\min\{\epsilon^2, 1 + \epsilon\} \left(1 - \frac{1}{2^{1+\gamma}}\right)^2}{c^{1+\gamma}} - 2(1 + \epsilon)^{\kappa-1} \left(e^{\frac{\varphi}{N_{\text{left}}^{1+\gamma}}} - 1\right) \right).$$

One can bound the expected change in Ξ_ϵ if $S > 0$ by

$$\mathbb{E}[\Xi_\epsilon(\ell + 1) - \Xi_\epsilon(\ell) | S(\ell) \geq 1] \geq \frac{1}{N_{\text{left}}^{1+\gamma}} \left(\frac{\min\{\epsilon^2, 1 + \epsilon\} \left(1 - \frac{1}{2^{1+\gamma}}\right)^2}{c^{1+\gamma}} - 2(1 + \epsilon)^{\kappa-1} \left(e^{\frac{1}{N_{\text{left}}^{1+\gamma}}} - 1\right) \right).$$

From this expression, one can see that, given A satisfying

$$0 < A < \frac{\min\{\epsilon^2, 1 + \epsilon\} \left(1 - \frac{1}{2^{1+\gamma}}\right)^2}{c^{1+\gamma}},$$

one can find $N_{\text{crit}}(\epsilon, \gamma, c, A)$ such that for all $N_{\text{left}} \geq N_{\text{crit}}(\epsilon, \gamma, c, A)$, the following holds

$$\mathbb{E}[\Xi_\epsilon(\ell + 1) - \Xi_\epsilon(\ell) | S(\ell) \geq 1] \geq \frac{A}{N_{\text{left}}^{1+\gamma}}.$$

From (5), note that the largest that Ξ_ϵ can be with $N_{\text{left}} \geq N_{\text{crit}}$ is

$$\Xi_{\epsilon, \text{crit}} = \frac{N - N_{\text{crit}}}{\kappa} (1 + \epsilon)^\kappa + \frac{N_{\text{crit}}}{\kappa - 1} (1 + \epsilon)^{\kappa-1}.$$

Furthermore, once $N_{\text{left}} \leq N_{\text{crit}}$, it will be for all time after. From (8) and our choice of δ_i , there are at most $\sqrt{cN_{\text{left}}}$ timesteps between each time when at least one agent desires to switch. Therefore, one can say that

$$\mathbb{E}[N_{\text{left}}(\ell)] \leq N_{\text{crit}}, \forall \ell \geq \ell_{\text{crit}} = \lceil \Xi_{\epsilon, \text{crit}} \sqrt{cN_{\text{left}}} \frac{N_{\text{left}}^{1+\gamma}}{A} \rceil.$$

By the definition of the expected value for a non-negative random variable, $\mathbb{P}(N_{\text{left}}(\ell) \leq N_{\text{crit}}) \geq \frac{1}{2}$, for all $\ell \geq \ell_{\text{crit}}$. Defining T_{crit} as the first time that $N_{\text{left}} \leq N_{\text{crit}}$, it is clear that $\mathbb{E}[T_{\text{crit}}] \leq \ell_{\text{crit}} + 1$.

Finally, define T to be the first time that all agents are in a goal coalition partition. By Lemma 6.3, $\mathbb{E}[|T|]$ is finite, for all finite N . With this condition satisfied, one can apply the law of total expectation [Billingsley, 1995] as well as Lemma 6.3 again and see that

$$\begin{aligned} \mathbb{E}[T] &= \mathbb{E}[\mathbb{E}[T | T_{\text{crit}}]] \leq \mathbb{E}[L_{\text{left}}(cN_{\text{crit}})^{N_{\text{crit}}(1+\gamma)L_{\text{left}}} + T_{\text{crit}}] \\ &\leq L_{\text{left}}(cN_{\text{crit}})^{N_{\text{crit}}(1+\gamma)L_{\text{left}}} + \ell_{\text{crit}} + 1. \end{aligned}$$

Finally, given that $N_{\text{left}} \leq N$, noting the order of ℓ_{crit} with respect to N finishes the result. \square

Next, we show that the time complexity bound in Proposition 6.4 can be improved for the complete graph.

Proposition 6.5 (Time complexity on the complete graph) *Under the COALITION FORMATION AND DEPLOYMENT ALGORITHM, the expected number of timesteps for the network to enter the goal coalition partition is $\mathcal{O}(N^{1+\gamma})$ when agents switch using the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS law and each agent can communicate with all other agents.*

PROOF. The proof strategy is the same as for Proposition 6.4, so we only provide a sketch here. There are two differences between the generic case and the complete graph case. The first difference is that in a timestep where at least one agent wishes to switch coalitions, in the complete graph case, we can show that almost all agents wish to switch. More precisely, agents in coalitions larger than size κ have an incentive to at least form their own coalition. Of the coalitions less than size κ , agents desire to join the largest one. If two coalitions have the same cardinality and are both the largest coalition of size less than κ , they mutually want to join each other. This means that at least $N_{\text{left}} - \kappa + 1$ agents have an incentive to switch coalitions. This affects the expected change in Ξ_ϵ as follows,

$$\begin{aligned} \mathbb{E}[\Xi_\epsilon(\ell + 1) - \Xi_\epsilon(\ell) | S(\ell) \geq 1] \geq & \frac{N_{\text{left}} - \kappa + 1}{N_{\text{left}}^{1+\gamma}} \left(\frac{\min\{\epsilon^2, 1 + \epsilon\} \left(1 - \frac{1}{2^{1+\gamma}}\right)^2}{c^{1+\gamma}} \right. \\ & \left. - 2(1 + \epsilon)^{\kappa-1} \left(e^{\frac{1}{N_{\text{left}}^{1+\gamma}}} - 1\right) \right). \end{aligned}$$

The second difference is that agents wish to switch at every timestep (instead of once every $\sqrt{cN_{\text{left}}}$ timesteps, as in Proposition 6.4, given the assumed δ_i and the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS switching law). The result follows from propagating these changes through the proof of Proposition 6.4. \square

6.2 Communication complexity per timestep

Here, we analyze the communication complexity per timestep of the COALITION FORMATION AND DEPLOYMENT ALGORITHM. We begin by stating our conventions regarding how messages are counted along the algorithm execution. First, we make the assumption that an identical message sent at a given moment by an agent to one or more other agents located within some distance of it counts as one. For instance, an omnidirectional communication model fits into this description. Second, in several instances, the algorithm requires the location and identity of the neighbors of an agent. This information may be obtained via either communication or sensing. We assume this service is efficiently carried out by the network and do not count it toward the communication required per timestep.

Before stating our communication complexity characterization, we introduce one slight modification to step 2 of the BEST NEIGHBOR COALITION DETECTION that decreases the number of messages sent without affecting the overall algorithm execution. According to this modification, an agent, instead of querying all neighbors that are not in its coalition, will now query the closest neighbor who is not in its current coalition and also not in a coalition of size κ . This modification requires agents to know what nearby agents have already formed a coalition of the desired size. This can be addressed in at least one of the following two ways. One way is for agents to broadcast to the network that it and its coalition members are in a complete group. The other way is that when agents query the coalition size of other agents (as in step 2 of the BEST NEIGHBOR COALITION DETECTION), if the queried agent is in a complete group, the querying agent notes the identities of all agents already in that complete group, and never needs to ask again.

We refer to the algorithm with this modification as the COALITION FORMATION AND DEPLOYMENT ALGORITHM*. The modification does not affect the algorithm's correctness or time complexity bounds. The basic argument is that, even with the modification, in any network configuration, one can still guarantee that within $\text{diam}(Q)/\delta$ timesteps, at least one agent will have an incentive to join a more desirable coalition. This observation allows to reproduce the technical proofs given above to establish the same correctness and time complexity results. The next result characterizes the communication complexity per timestep of the algorithm.

Proposition 6.6 (Communication complexity per timestep) *Under the COALITION FORMATION AND DEPLOYMENT ALGORITHM*, the network of agents sends at most $6N_{\text{left}}$ messages per timestep and, hence, the communication complexity per timestep is $\mathcal{O}(N)$.*

PROOF. We show the result by simply counting the number of messages sent per timestep. In steps 2 and 3

of the modified version of the BEST NEIGHBOR COALITION DETECTION, an agent not in a coalition of size κ sends one message to one neighbor and receives one back. In steps 2 and 4 of the COALITION SWITCHING, when an agent switches coalitions it sends one broadcast message alerting its former coalition it is leaving, as well as one message to tell one member of its new coalition it is joining. This one member in the new coalition sends one broadcast message alerting the rest of its coalition of the new member, as specified in step 8. Finally, that member sends one message back to the joining member to alert it of any other agents who happened to join/leave in the exact same timestep, as specified in step 12. Thus, if an agent is joining another coalition, 4 messages are required and if the agent is forming its own coalition of size 1, only 1 message is required. The most messages would be sent if all agents switched at the same timestep. Since agents in coalitions of size κ will never switch, executing one timestep of the COALITION FORMATION AND DEPLOYMENT ALGORITHM* generates at most $6N_{\text{left}}$ messages, and the result follows. \square

Note that the upper bound in Proposition 6.6 is a function of the agents not in a completed coalition and, thus, monotonically decreasing as the algorithm evolves and completed coalitions form.

7 Simulations

This section presents several simulations of the COALITION FORMATION AND DEPLOYMENT ALGORITHM. In all simulations where they are relevant, $\delta = d_{\text{max}} = \frac{.2 \text{ diam}(Q)}{\sqrt{2} \sqrt{N}}$. We use the function

$$\phi(C_1, \dots, C_N) = \frac{1}{N(\kappa - 1)} \sum_{i \in \mathcal{A}} ||C_i| - \kappa|, \quad (10)$$

to illustrate the coalition formation dynamics. This function measures the normalized average absolute difference between the agents' coalition size and the desired size κ .

We begin by illustrating the correctness of the algorithm, i.e., convergence to a desired goal coalition partition and the achievement of the deployment task. Figure 2(a)-(b) show an execution of the COALITION FORMATION AND DEPLOYMENT ALGORITHM on 21 agents forming coalitions of size 2 with PROPORTIONAL-TO-COALITION-SIZE switching law defined by

$$p = \begin{cases} 1 - (1 - b)^{\frac{1}{|C_i|}}, & \text{if } |C_i| \neq \kappa, \\ 0, & \text{if } |C_i| = \kappa, \end{cases} \quad (11)$$

for some $b \in (0, 1)$. Note that this switching law satisfies (4). In this and other simulations where b is constant, we chose $b = 0.5$. The appeal of this switching law is that b is the probability that at least one agent in a coalition will switch, given that all coalition members

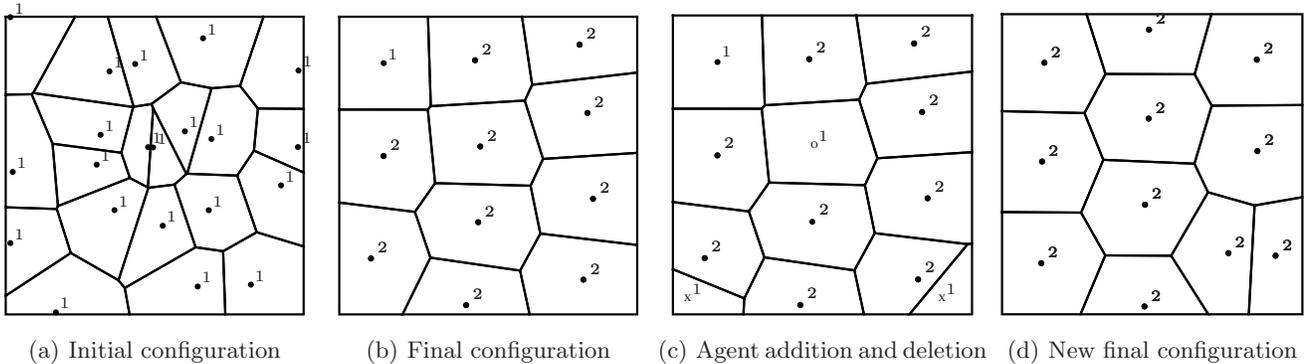


Fig. 2. Execution of the COALITION FORMATION AND DEPLOYMENT ALGORITHM with 21 agents and $\kappa = 2$ using the PROPORTIONAL-TO-COALITION-SIZE switching law (11). The network starts at (a) and converges to the configuration in (b) where all agents are in correctly-sized coalitions and these coalitions are optimally deployed. After this, (c) shows an agent failing in the coalition marked as 'o' and two agents, marked as 'x', joining the network. After the agent additions and subtractions, coalitions adapt and the network re-converges to the optimal deployment configuration in (d).

wish to switch. This switching law makes it likely that several agents (most likely from different coalitions) will get the chance to switch coalitions at each timestep. One can observe in Figure 2(b) that the network converges to both correctly sized groups and coalitions optimally deployed at their Voronoi cell's circumcenters. From Theorem 5.1, the final configuration optimizes $\mathcal{H}_{DC,11}$.

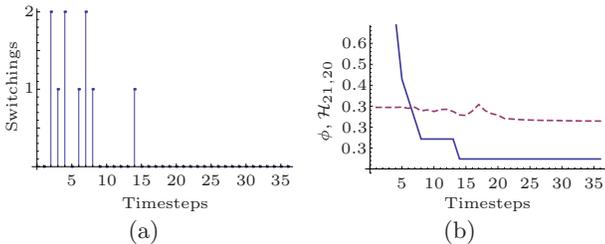


Fig. 3. For the execution in Figure 2(a) and (b), (a) shows the number of agents switching coalitions at each timestep, and (b) shows the evolution of ϕ (solid line) as defined in (10) and $\mathcal{H}_{21,20}$ (dashed line) as defined in (1).

Figure 3(a) shows the number of coalition switches at each timestep for the same run. Many switches happen early, but decrease in frequency as agents form κ -sized coalitions. The evolution of ϕ depicted in Figure 3(b) confirms this by showing how agents join more desirable coalitions over time. It also shows the evolution of the objective function $\mathcal{H}_{N,N-1}$ that, in the language of Section 2.2, corresponds to the situation where $N - 1$ of the sensors are working. This choice of function is motivated by the fact that, in one dimension, it is known that in such a case, forming coalitions of size 2 is optimal [Cortés, 2012]. The bumps in the evolution of $\mathcal{H}_{21,20}$ occur when an agent with no nearby coalitions to join must increase its radius to join a group far away. $\mathcal{H}_{21,20}$ temporarily increases while these agents get together.

Figure 2(c) and (d) illustrate the robustness of the COALITION FORMATION AND DEPLOYMENT ALGORITHM. After agents have achieved the final optimal

configuration seen in Figure 2(b), we let one agent fail and two new agents come into the picture. The agents adapt to the new network composition and optimally deploy according to the available resources.

Next, Figure 4 shows the average number of timesteps required for coalition formation for 4 different probabilistic switching laws under a generic communication topology, cf. (a), and the complete communication topology, cf. (b). For all network sizes, the desired coalition size is 4. Each point is the average of 50 runs, where the agents are initially randomly placed with uniform distribution in a unit square. The time complexity upper bounds in Section 6.1 are corroborated and the bound seems tight for the complete communication case.

Figure 5 illustrates the communication complexity analysis of Section 6.2. (a) shows the number of messages sent per timestep for one execution of 21 agents forming coalitions of size 3 with switching probability $p = .9$. As coalitions form, fewer messages are sent per timestep. (b) depicts the average number of messages sent per timestep as a function of the total network size for switching probability $p = .9$ and desired coalition size of 4. The plot validates the $\mathcal{O}(N)$ -characterization of the communication complexity per timestep stated in Proposition 6.6.

Finally, Figure 6 illustrates the dependency of the average coalition formation time on κ and b for the PROPORTIONAL-TO-COALITION-SIZE switching law. We focus on this law because it is the one that executes the fastest out of the probability laws illustrated in Figure 4. Each point is the average of 200 runs, where agents are initially uniformly randomly placed in a unit square. The error bars correspond to plus and minus one standard deviation. Figure 6(a) shows the average coalition formation convergence time for fixed $N = 20$ and varying κ . This time is roughly equal for all desired coalition sizes, until nearly all agents are joining one coalition, which takes less time on average. Figure 6(b)

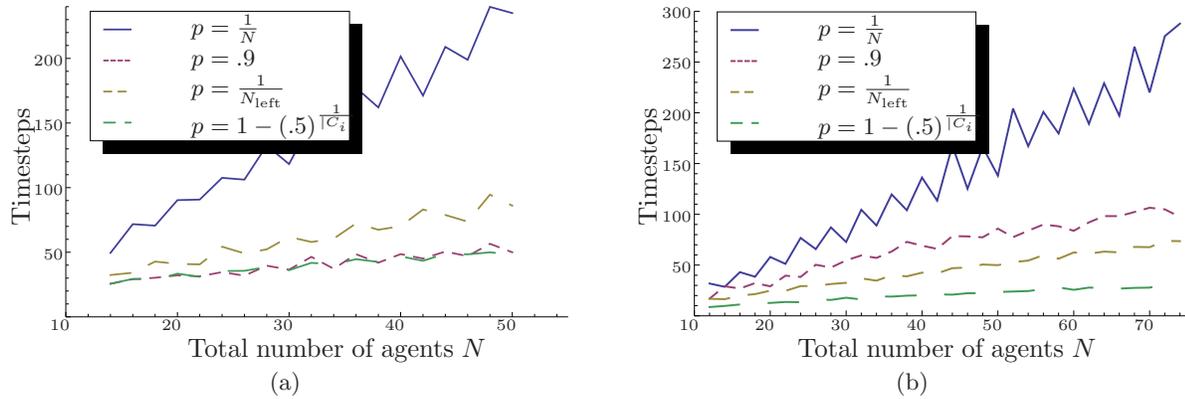


Fig. 4. Average coalition formation time for 4 different probabilistic switching laws under (a) a generic communication topology and (b) the complete communication topology. Each point is the average of 50 runs, where the agents were initially randomly placed with uniform distribution in a unit square. The time complexity upper bounds in Section 6.1 are validated, and that the bound seems tight for the complete communication case.

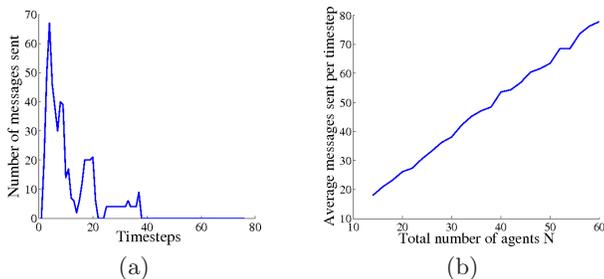


Fig. 5. (a) depicts the number of messages sent per timestep for a run of 21 agents forming coalitions of size 3. (b) illustrates the average number of messages sent per timestep as a function of the network size. Both of these plots validate that the algorithm has an $\mathcal{O}(N)$ message complexity per timestep result, as in Proposition 6.6.

shows the average coalition formation time for 20 agents forming coalitions of size 4 as a function of b . For values of b far from 0 and 1, this time is roughly constant.

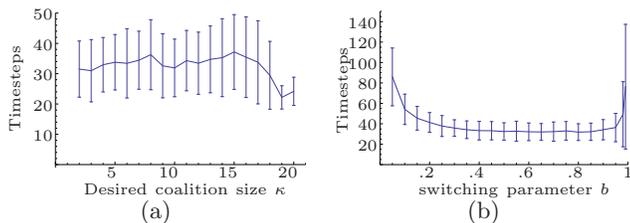


Fig. 6. Average coalition formation time for 20 agents with the PROPORTIONAL-TO-COALITION-SIZE switching law (11) as a function of (a) the desired coalition size κ and (b) the parameter b (with coalitions of size 4). Each point is the average of 200 runs, where the agents were initially randomly placed with uniform distribution in a unit square. The error bars correspond to plus and minus one standard deviation.

8 Conclusions

Motivated by a spatial estimation problem, we have designed a synchronous, distributed algorithm for a net-

work of robotic agents to autonomously deploy in groups over a given region. Our strategy allows agents to autonomously form coalitions of a desired size, cluster together within finite time, and asymptotically reach an optimal deployment, each with probability 1. The algorithm design is a combination of a hedonic coalition formation game where agents only have partial information about other coalition memberships with motion coordination strategies for group clustering and deployment. The proposed algorithmic solution, termed COALITION FORMATION AND DEPLOYMENT ALGORITHM, is provably correct, does not rely on a common reference frame and is robust to agents joining or leaving the environment. We have provided time complexity upper bounds for algorithm executions with the PROPORTIONAL-TO-NUMBER-OF-UNMATCHED-AGENTS switching law under arbitrary and complete communication topologies. We also have upper bounded the communication complexity per timestep for algorithm executions with arbitrary switching laws. Simulations illustrate the correctness, robustness, and complexity results. Future work will be devoted to characterizing the probabilistic switching law that optimizes the speed of the coalition formation process and further explore the use of noncooperative game-theoretic ideas in other motion coordination problems.

References

- G. Arslan, J. R. Marden, and J. S. Shamma. Autonomous vehicle-target assignment: A game theoretic formulation. *ASME Journal on Dynamic Systems, Measurement, and Control*, 129(5): 584–596, 2007.
- S. Bandyopadhyay and E. Coyle. An energy efficient hierarchical clustering algorithm for wireless sensor networks. In *Joint Conference of the IEEE Computer and Communication IEEE Societies*, pages 1713–1723, 2003.
- S. Banerjee, H. Konishi, and T. Sonmez. Core in a simple coalition formation game. *Social Choice and Welfare*, 18:135–153, 2001.
- T. Başar and G.J. Olsder. *Dynamic Noncooperative Game Theory*. Academic Press, 1982.
- P. Billingsley. *Probability and Measure*. Wiley, 1995.
- A. Bogomolnaia and M. Jackson. The stability of hedonic coali-

tion structures. *Games and Economic Behavior*, 38(2):201–230, 2002.

F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009. ISBN 978-0-691-14195-4. Electronically available at <http://coordinationbook.info>.

J. Cortés. Deployment of an unreliable robotic sensor network for spatial estimation. *Systems & Control Letters*, 61(1):41–49, 2012.

J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.

D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

R. Graham and J. Cortés. Asymptotic optimality of multicenter Voronoi configurations for random field estimation. *IEEE Transactions on Automatic Control*, 54(1):153–158, 2009.

D. Gu. A differential game approach to formation control. *IEEE Transactions on Control Systems Technology*, 16(1):85–93, 2008.

V. Gupta, A. F. Dana, J. Hespanha, R. M. Murray, and B. Hassibi. Data transmission over networks for estimation and control. *IEEE Transactions on Automatic Control*, 54(8):1807–1819, 2009.

B. Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, 14(3):502–525, September 1982.

N. Heo and P. Varshney. An intelligent deployment and clustering algorithm for a distributed mobile sensor network. In *IEEE Conference on Cybernetics*, pages 4576–4581, 2003.

G. Lawler. *Introduction to Stochastic Processes*. Chapman and Hall, 2006.

N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1997. ISBN 1558603484.

J. R. Marden, G. Arslan, and J. S. Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Conference on autonomous agents and multi-agent systems*, pages 194–201, Honolulu, Hawaii, May 2007.

J. R. Marden, G. Arslan, and J. S. Shamma. Cooperative control and potential games. *IEEE Transactions on Systems, Man & Cybernetics. Part B: Cybernetics*, 39:1393–1407, 2009a.

J. R. Marden, H. Young, G. Arslan, and J. S. Shamma. Payoff-based dynamics for multi-player weakly acyclic games. *SIAM Journal on Control and Optimization*, 48(1):373–396, February 2009b.

S. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009.

A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley Series in Probability and Statistics. Wiley, 2 edition, 2000. ISBN 0471986356.

D. Peleg. *Distributed Computing. A Locality-Sensitive Approach*. Monographs on Discrete Mathematics and Applications. SIAM, 2000. ISBN 0898714648.

J. Rosenthal. *A First Look at Rigorous Probability Theory*. World Scientific Publishing, Singapore, 2000.

W. Saad, Z. Han, A. Hjørungnes, D. Niyato, and E. Hosain. Coalition formation games for distributed cooperation among roadside units in vehicular networks. *IEEE Journal on Selected Areas in Communications*, 29(1):48–60, 2011.

L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. S. Sastry. Foundations of control and estimation over lossy networks. *Proceedings of the IEEE*, 95(1):163–187, 2007.

S. C. Smith and P. Seiler. Estimation with lossy measurements: jump estimators for jump systems. *IEEE Transactions on Automatic Control*, 48(12):1453–1464, 2003.

O. Younis and S. Fahmy. HEED: a hybrid, energy-efficient, dis-

tributed clustering approach for ad hoc sensor networks. *IEEE Transactions on Mobile Computing*, 3(4):366–379, 2004.

A Properties of gttg

This appendix contains useful properties of gttg.

Lemma A.1 *For $d > 0$ and $p_1, p_2, q \in \mathbb{R}^n$, let $p_i^+ = \min\{\|q - p_i\|, d\} \text{vr}(q - p_i) + p_i$, $i \in \{1, 2\}$. Then $\|p_1^+ - p_2^+\| \leq \|p_1 - p_2\|$.*

Lemma A.1 is used in determining how much the circumradius of a coalition decreases and how much they get closer to the goal point q after moving according to gttg.

Proposition A.2 (Application of gttg decreases circumradius) *Given $P = (p_1, \dots, p_k)$ and $q \in Q$, let $P^+ = (p_1^+, \dots, p_k^+)$ be given by $p_i^+ = \text{gttg}(p_i, P, q)$, $i \in \{1, \dots, k\}$. Then $\text{CR}(P^+) \leq \text{CR}(P) - \delta_1$ and*

$$P^+ \subset B(q, \|CC(P) - q\| + \text{CR}(P) - \delta_1 - \delta_2),$$

with $\delta_1 = \max_{i \in \{1, \dots, k\}} \min\{\|CC(P) - p_i\|, d_1(r)\}$ and $\delta_2 = \min\{\|q - CC(P)\|, d_2(r)\}$.

PROOF. Our strategy is to look independently at the effect of the two halves of the motion defined in gttg. Define the intermediate positions P^* by $p_i^* = p_i + w_{1,i}$, $i \in \{1, \dots, k\}$, where $w_{1,i} = \min\{\|CC(P) - p_i\|, d_1(r)\} \text{vr}(CC(P) - p_i)$. We show that the circumradius decreases a finite amount while moving from P to P^* and does not increase while moving from P^* to P^+ . First, according to the motion prescribed by gttg, we have $P^* \subset B(CC(P), \text{CR}(P) - \delta_1)$, which by definition of circumcenter, implies that $\text{CR}(P^*) \leq \text{CR}(P) - \delta_1$. Second, to show $\text{CR}(P^+) \leq \text{CR}(P^*)$, let us rewrite p_i^+ as $p_i^+ = p_i^* + w_{2,i}$, where $w_{2,i} = \min\{\|q - p_i^*\|, d_2(r)\} \text{vr}(q - p_i^*)$. Let

$$NC^+ = CC(P) + \delta_2 \text{vr}(q - CC(P)). \quad (\text{A.1})$$

By Lemma A.1, for all $i \in \{1, \dots, k\}$,

$$\|p_i^+ - NC^+\| \leq \|p_i^* - CC(P)\| \leq \text{CR}(P^*),$$

where we have used the fact that $CC(P^*) = CC(P)$. Finally, $\text{CR}(P^+) \leq \max_{i \in \{1, \dots, k\}} \|p_i^+ - NC^+\|$ implies that $\text{CR}(P^+) \leq \text{CR}(P^*)$ and the result follows. Next, we study how much closer the points are to q after the application of gttg. Initially, $P \subset B(q, \|q - CC(P)\| + \text{CR}(P))$. After the application of w_1 , the configuration's circumcenter has not moved and the circumradius has decreased by δ_1 , so $P^* \subset B(q, \|q - CC(P)\| + \text{CR}(P) - \delta_1)$. Then, after the application of w_2 ,

$$P^+ \subset B(q, \|q - NC^+\| + \text{CR}(P) - \delta_1).$$

Combined with (A.1), the result follows. \square