Fault tolerant area coverage control for multi-agent systems

Sotiris Papatheodorou^{1,1}, and Anthony Tzes¹

¹New York University Abu Dhabi, Electrical and Computer Engineering, P.O.129188, Abu Dhabi, United Arab Emirates

Abstract. The fault tolerance characteristics of a distributed multi-agent coverage algorithm are examined. A team of sensor-equipped mobile agents is tasked with covering a planar region of interest. A distributed, gradient-based control scheme is utilized for this purpose. The agents are assumed to consist of three subsystems, each one of which may fail. The subsystems under examination are the actuation, sensing and the communication subsystem. Partial and catastrophic faults are examined. Several simulation studies are conducted highlighting the robustness of the distributed nature of the control scheme to these classes of faults, even when several of them happen at the same time.

1 Introduction

The problem of multi-agent area coverage can be categorized as blanket [1], sweep [2] or barrier [3] coverage. In all categories the common theme is the use of a team of sensor-equipped agents for a sensor coverage task. The agents are able to communicate with their neighbors in order to coordinate their movements and attain their common objective in a collaborative manner.

There have been varied approaches to the problem of blanket area coverage, which is the one examined in this article. These include game-theoretic methods [4], event-triggered control [5] or more classical methods such as distributed geometric optimization [6] and optimal control theory [7]. Several special cases of the blanket coverage problem have been examined, including coverage for non-convex domains [8,9], taking into account heterogeneous or anisotropic sensing patterns [10,11], accounting for more realistic agent dynamics [12] and ensuring communication among the agents [13].

In this work the agents are assumed to be dimensionless and are governed by a single integrator kinematic model. Each agent is equipped with an omnidirectional, limited-range sensor, allowing it to sense a circular region around it, while the sensor effective range is assumed to be common among agents. Agents are also equipped with transceivers in order to communicate with nearby agents and exchange the necessary information for coordinating their movement. The goal of the mobile agent team is to cover as much of a given convex, planar region of interest as possible using their on-board sensors. To that extent the power diagram [14] is used in order to assign disjoint regions of responsibility to

1 Corresponding author: sotiris.papatheodorou@nyu.edu

the agents. The computation of the coverage objective can be then distributed over the power diagram cells allowing the use of geometric optimization techniques in order to derive a distributed control law [15]. In order to evaluate the robustness of the resulting control law, each agent is assumed to be comprised of a set of subsystems, each of which can potentially fail. The actuation subsystem of an agent is responsible for its movement inside the region of interest. The sensing subsystem is tasked with sensing the region of interest, which is the goal of the team deployment. Finally, the communication subsystem is tasked with transferring information required for implementing the control law among the agents. Simulation studies are used to show that the resulting control law is robust to faults in the four subsystems described previously.

The first class of faults considered in this article are actuation faults. These can be categorized into partial or complete actuation failures. In a partial actuation failure, one or more of the mobile agent's actuators operate at lower than full capacity which results in the agent moving with reduced speed in the respective direction. In a complete actuation failure on the other hand, one or more of the agent's actuators stop functioning completely in which case the agent may be unable to move at all. Both cases are examined and it is shown through simulations that the agents are able to reconfigure themselves and achieve coverage despite some of them losing part of or all of their actuation capabilities.

Another type of fault considered in this article is the loss of sensing performance. In this case a faulty sensor is modeled as having a reduced effective range. Although the Voronoi diagram could be used in the case of identical sensing radii, it could not take into account the loss of sensing performance due to a faulty sensor, since it is dependent only on the agents' positions. The power diagram on the other hand takes into account the agent's sensing performance in addition to their positions. Thus agents whose sensors have degraded performance are assigned smaller regions of responsibility, resulting in the other agents covering the region the faulty ones no longer can. This behavior is also verified through simulation studies.

Another fault that may arise during the operation of a multi-agent system is a failure of the communication subsystem. In such a case the faulty agent will be unable to receive information from the other agents, as well as send them the information they require in order to implement the control law, essentially becoming completely isolated from the team. Since the neighbors of the faulty agents cannot distinguish a failure of the communication subsystem from a total agent failure, there are two approaches to handling this fault. The safe approach would be to always assume that the faulty agent is unable to perform its assigned sensing task, thus the other agents should implement the control scheme as if it no longer exists. The alternative approach is to always assume that the faulty agent can still complete its sensing task. In this case, the other agents should compute their control law assuming the faulty agent is located at its last reported position. In both cases the faulty agent should remain immobile since the others have no means of learning its new position. The performance of the agent team for both cases is shown through simulation studies.

It is shown that the examined control scheme based on the power diagram and distributed optimization techniques is robust to multiple classes of faults. In all cases examined the agent team is able to reconfigure and adapt despite some of its members having severe faults.

2 Problem statement

We assume a convex region of interest $\Omega \subset \mathbb{R}^2$ and a space density function $\phi: \Omega \rightarrow \mathbb{R}_+$ which encodes all a priori knowledge about the importance of points inside

Ω . A team of n mobile ground agents which is deployed inside the region of interest. Each agent $i \in I_n = \{1, 2, \dots, n\}$ is approximated by a point mass located at $q_i \in \Omega$ and is governed by the following kinematic model

$$\dot{q}_i = a_i u_i, q \in \Omega, u_i \in \mathbb{R}^2, \quad (1)$$

where u_i is the velocity control input of agent i and $a_i \in [0, 1]$ is the actuator efficiency parameter. Actuator efficiency values of 1, (0,1) and 0 correspond to normal operation, partial actuator fault and catastrophic actuator fault.

All agents are equipped with identical omnidirectional sensors with limited sensing radius R , resulting in a circular sensing pattern

$$S_i(q_i, R) = \{q \in \Omega : \|q - q_i\| \leq s_i R\}, \quad (2)$$

where $s_i \in [0, 1]$ is the sensor efficiency parameter of agent i . Sensor efficiency values of 1, (0,1) and 0 correspond to normal operation, partial sensor failure and complete sensor failure.

Agents are also equipped with identical wireless transceivers in order to exchange information with their neighbors. Given the fact that the sensing radii of all agents are not greater than R , agents with a distance greater than $2R$ will not have any meaningful interaction with each other. Thus the communication radius of all agents is

$$R^c = c_i 2R, \quad (3)$$

where $c_i \in \{0, 1\}$ is used to represent normal operation and communication failure for $c_i = 1$ and $c_i = 0$ respectively.

The goal of the mobile agent team is maximizing the sensor coverage of the region of interest while simultaneously accounting for the space density function. Thus the coverage objective is defined as

$$H = \int_{\Omega} \max_{i \in I_n} \mathbf{1}_{S_i}(q) \phi(q) dq, \quad (4)$$

where $\mathbf{1}_{S_i}(q)$ is the indicator function for the set S_i . The control objective is the design of a distributed control law for the mobile agents so that monotonic increase over time of the coverage objective H is guaranteed.

3 Distributed control law

The first step in designing a distributed control law is the distribution of the computation of the coverage objective among the agents. To that end, a power partitioning of the region Ω using the agent positions q_i as generators and their sensing radii $s_i R$ as weights will be used. The power cell of agent i is defined as

$$V_i = \left\{ q \in \Omega : \|q - q_i\|^2 - (s_i R)^2 \leq \|q - q_j\|^2 - (s_j R)^2, j \in N_i \right\} \quad (5)$$

where N_i is the set of neighbors of agent i . It should be noted that under the absence of sensor faults, the power diagram converges to the Voronoi diagram, thus we can define the

neighbor set N_i as the Delaunay neighbors of agent i . In the case of sensor faults for the new set of neighbors $\hat{N}_i$ it will hold that $\hat{N}_i \subseteq N_i$, thus for this particular use case of the power diagram, it is sufficient to consider only the Delaunay neighbors of the generators. It is also useful to define the R -limited power cell as

$$V_i^R = V_i \cap S_i,$$

which is the intersection of each agent's power cell with its sensed region. Using the power partitioning and the R -limited cell definition, the coverage objective (4) can be written as

$$H = \sum_{i \in I_n} \int_{V_i^R} \phi(q) dq \quad (6)$$

It is now possible to design a gradient-based, distributed control law so that the coverage objective function (6) increases monotonically over time. The complete control law design procedure can be found in [15] and results in

$$u_i = g_i \int_{\partial V_i^R \cap \partial S_i} n_i \phi(q) dq \quad (7)$$

where ∂S denotes the boundary of the set S , n_i the outwards unit normal vector on V_i^R and g_i is a positive constant.

4 Operation under faults

In this section three kinds of faults are examined through simulation studies and for each one the agent team's performance is compared to its performance under normal operation.

4.1 Simulation 1: Normal Operation

The normal operation of agent team is examined in this section in order to better evaluate the effect of the various faults examined. A team of 8 agents with sensing radius $R = 0.5$ and control law gains $g_i = 1, \forall i \in I_n$ operating inside a convex region of interest is simulated. The region Ω is the same as in [15] for consistency. The agents' initial and final configurations can be seen in Fig. 1. [Left] and [Right] respectively where the agent positions are marked by black dots, the boundaries of their power cells are shown in blue and the boundaries of their sensed regions are shown in red. The agent trajectories are presented in Fig. 1. in red while the initial and final positions are shown as dots and circles respectively. This simulation is used for the evaluation of the agent team's performance under faults. The initial conditions of all following simulation studies are the same as in this one.

Fig. 1. Normal operation. Initial configuration [Left], agent trajectories [Center] and final configuration [Right].

4.2 Simulation 2: Actuation faults

Actuation faults reduce the efficiency of an agent’s actuators, leading to decreased movement speed or complete immobilization. This simulation study is used to examine the effects of actuation faults in the control law (7). Agent 8 suffers from a partial actuation fault $a_8 = 0.5$ at $t = 2$ seconds, essentially reducing its speed by half. Agent 7 suffers from a catastrophic actuation fault $a_7 = 0$ at $t = 3$ seconds, leaving it immobilized for the duration of the simulation. The agent trajectories and final configuration are shown in Fig. 2. [Left] and [Center] respectively, using the same colors and notation as in Fig. 1. The points on the trajectories of agents 7 and 8 at which the actuation faults occur are marked by a black cross and a black triangle respectively in Fig. 2. [Left]. The evolution of the coverage objective (6) over time is shown in Fig. 2. [Right] in solid blue for this simulation and in dashed red for normal operation. The times at which the failures occur are marked by dashed black vertical lines. It is observed that despite the actuation faults, the agents manage to converge to a locally optimal configuration although the final value of the coverage objective is 3.6% smaller than the one achieved under normal operation. It is also observed that the agent team needs more time to converge to its final configuration which can be attributed to the reduced speed of agent 8 and the fact that the immobile agent 7 obstructs the movement of other agents. It should be noted that since the control scheme converged to local optima, even if only partial actuator faults occurred, there would be no guarantee that the final agent configuration will be the same as the one under normal operation. In the static area coverage problem however it is the value of the coverage objective that matters and not the agent positions themselves.

Fig. 2. Actuation faults. Agent trajectories [Left], final configuration [Center] and coverage objective [Right].

4.3 Simulation 3: Sensor faults

Sensor faults limit the coverage performance of an agent by reducing its sensing radius, thus decreasing the area it is able to cover. This simulation study serves to examine the effect of sensor faults on the performance of the agent team. Agent 8 suffers from a partial sensor fault $s_8 = 0.5$ at $t = 5$ reducing its sensed area to 25% of its initial value. Agent 4 suffers from a complete sensor fault $s_4 = 0$ at $t = 10$ reducing its sensing performance to zero. The agent trajectories and final configuration are shown in Fig. 3. [Left] and [Center] respectively, using the same colors and notation as in Fig. 1. The points on the trajectories of agents 4 and 8 at which the actuation faults occur are marked by a black cross and a black triangle respectively in Fig. 3. [Left]. The evolution of the coverage objective (6) over time is shown in Fig. 3. [Right] in solid blue for this simulation and in dashed red for normal operation. The times at which the failures occur are marked by dashed black vertical lines. Sudden drops of the coverage objective are observed at these times due to the effective sensing radius of the faulty agents decreasing instantly. It is observed that despite the sensor faults the agents are able to converge to a locally optimal configuration although the final value of the coverage objective is 12.3% smaller than the one achieved under normal operation. Due to the use of the power diagram, agents are assigned cells according to their sensing capabilities, thus agents with sensor faults are not assigned disproportionately large cells.

Fig. 3. Sensor faults. Agent trajectories [Left], final configuration [Center] and coverage objective [Right].

4.4 Simulation 4: Communication faults

The effects of communication faults on the agent team performance are examined in this section. Agents with communication faults are unable to both transmit and receive messages to and from other agents, effectively isolating them from all other agents. It is assumed that each agent is able to detect a fault on its own communication system and that in such a case it will remain stationary indefinitely. The first assumption is needed so that an agent will not interpret a failure of its own communication system as a failure of all of its neighbors while the second assumption is needed so that the other agents know what to expect from an agent with a communication failure. From the perspective of the agents operating normally it is not possible to distinguish a failure in some other agent's communication system from a total failure of that agent since in both cases they will just be unable to communicate with that agent. In case of a communication failure it might be desirable to take the malfunctioning agent into account, given the fact that its sensor is still functioning. In that case all other agents assume it is located at its last known position. Since, control-wise, this is equivalent to the malfunctioning agent having a catastrophic actuation failure, this case is not examined further as the results are the same as those presented in Section 4.1 for agent 7. On the other hand, it might be desirable for the agents operating correctly to assume that every lack of communication from some agent is due to

total agent failure, meaning the malfunctioning agent is unable to move, sense and communicate. In that case the malfunctioning agent is essentially removed from the agent team, which is the case examined in the following simulation study.

This simulation study serves to study the effect of communication failures when assuming that they correspond to total agent failures. Agents 4 and 8 suffer communication failures $c_4 = 0$ and $c_8 = 0$ at times $t = 5$ and $t = 10$ respectively. The agent trajectories and final configuration are shown in Fig. 4. [Left] and [Center] respectively, using the same colors and notation as in Fig. 1. The points on the trajectories of agents 4 and 8 at which the communication faults occur are marked by black crosses in Fig. 4. [Left]. The evolution of the coverage objective (6) over time is shown in Fig. 4. [Right] in solid blue for this simulation and in dashed red for normal operation. The times at which the failures occur are marked by dashed black vertical lines. Sudden drops of the coverage objective are observed at these times due to the malfunctioning agents essentially being removed from the team. It is observed that despite the communication faults the agents are able to converge to a locally optimal configuration although the final value of the coverage objective is 13.6% smaller than the one achieved under normal operation. After each malfunctioning agent is removed from the team, the others are able to adapt and move to cover the region that it was responsible for.

Fig. 4. Communication faults. Agent trajectories [Left], final configuration [Center] and coverage objective [Right].

5 Conclusions

The robustness of a distributed control scheme for area coverage by a mobile agent team was examined in this article. Several classes of faults were considered and their effects examined through simulation studies. It has been shown that the control law is inherently robust and is able to adapt in real time to various types of agent faults.

References

1. F. Abbasi, A. Mesbahi, J.M. Velni, *Aut.* **81**, 342-349 (2017).
2. J. M. Palacios-Gasós, E. Montijano, C. Sagüés, S. Llorente, *IEEE Tr. Rob.* **99**, 1-17 (2016).
3. A. Chen, S. Kumar, T. H. Lai, *ACM International Conference on Mobile Computing and Networking*, 63-74 (2007).
4. V. Ramaswamy, J. R. Marden, *American Control Conference*, 6399-6404 (2016).
5. C. Nowzari, J. Cortés, *Aut.* **48**, 1077-1087 (2012).

6. J. Cortés and F. Bullo, *SIAM J. Con. Opt.* **44**, 1543–1574 (2005).
7. M. T. Nguyen, L. Rodrigues, C. S. Maniu, S. Olaru, *International Symposium on Intelligent Control*, 1-6 (2016).
8. Y. Stergiopoulos, M. Thanou, A. Tzes, *IEEE Tr. Aut. Con.* **60**, 2422-2427 (2015).
9. R. J. Alitappeh, K. Jeddisaravi, F. G. Guimaraes, *Soft Comp.* **21**, 6481-6497 (2017).
10. L. C. Pimenta, V. Kumar, R. C. Mesquita, G. A. Pereira, *IEEE Conference on Decision and Control*, 3947-3952 (2008).
11. Y. Stergiopoulos, A. Tzes, *IEEE International Conference on Robotics and Automation*, 1106-1111 (2014).
12. O. Arslan and D. E. Koditschek, *IEEE International Conference on Robotics and Automation*, 4259-4266 (2016).
13. Y. Kantaros and M. M. Zavlanos, *Aut.* **63**, 209-220 (2016).
14. A. Kwok, S. Martinez, *IEEE Tr. Con. Sys. Tech.* **19**, 792-804 (2011).
15. J. Cortés, S. Martinez, F. Bullo, *ESAIM: Con. Opt. Calc. Var.* **11**, 691-719 (2005).