Loss of Distributed Coverage Using Lazy Agents Operating Under Discrete, Local, Event-Triggered Communication

Extended Abstract

Edward Vickery
SHL
London, United Kingdom
ted.vickery@gmail.com

Aditya A. Paranjape
Tata Consultancy Services Ltd.
Pune, India
aditya.paranjape@gmail.com.

ABSTRACT

In the context of continuous surveillance of a spatial region, this paper investigates a practically-relevant scenario where robotic sensors are introduced asynchronously and inter-robot communication is discrete, event-driven, local and asynchronous. The robots are assumed to be lazy; i.e., they seek to minimize their area of responsibility by equipartitioning the domain to be covered. We construct a non-trivial example which shows that coverage guarantees for a given algorithm might be sensitive to the number of robots and, therefore, may not scale in obvious ways. It also suggests that when such algorithms are to be verified and validated prior to field deployment, the number of robots or sensors used in test scenarios should match that deployed on the field.

KEYWORDS

Distributed coverage, gossip-based communication, loss of coverage

ACM Reference Format:
Edward Vickery and Aditya A. Paranjape. 2023. Loss of Distributed Coverage Using Lazy Agents Operating Under Discrete, Local, Event-Triggered Communication: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

We address the problem of dynamically partitioning an environment for the purpose of continuous surveillance using a distributed scheme which relies on local, event-triggered communication between mobile robots. The robots are introduced asynchronously and seek to minimize their individual areas of coverage while ensuring that the environment as a whole is covered. The distributed nature of the communication and task allocation between the robots leads to the following question with practical ramifications: *if a reasonably designed coverage algorithm is proven to work for some non-trivial range of numbers of robots, can it fail to work when the number of robots is changed to outside the proven range?* We answer this question in the affirmative by constructing an example.

Optimal sensor placement problems map the number of sensors and their placement to an objective function which needs to be either maximized or minimized subject to constraints related to the sensors and the environment [4, 5, 9–11]. When the sensors are mounted on mobile robotic platforms, the sensor placement problem needs to be solved dynamically together with the associated path planning problems for the mobile robots [6, 15]. These *coverage* problems have been extended to accommodate non-convex domains [2, 3, 13], robots with finite communication radii [7], environments with unknown sensory functions [14], and cases where the necessity for covering the area has to be balanced against the spatial distribution of events in that area [1]. A continuous flow of information may not be required for coverage; instead, communication triggered by individual agents [12] or gossip between random pairs of agents [8] may suffice.

1.1 Contribution

This paper considers lazy robots operating under a coverage algorithm based on [8]. By lazy, we mean that a robot's share of the total area at any given time is proportional to the reciprocal of the total number of agents that it is aware of, which can be viewed as a proxy for the robot minimizing its energy consumption. We examine the conditions under which this lazy behavior leads to a loss of coverage. We construct a simplified example and a sequence of events which leads to an instantaneous loss of coverage when the number of robots exceeds a non-trivial threshold. The same sequence of events actually leads to an equipartition of the domain (i.e., the optimum solution) for a smaller number of robots. This demonstration suggests that the success of multi-agent algorithms operating in the presence of restricted communication might be sensitive to the number of agents involved, above and beyond the known complexities that arise due to the "scale" of the problem or the geometry of the environment. We refer the reader to [16] for the proofs of the mathematical results presented here.

2 MAIN RESULTS

Let $Q \subset \mathbb{R}^2$ be a closed, bounded domain containing $N \geq 1$ agents or sensors. Let $p_i \in Q$ denote the position of the i^{th} agent, and let $V_i \subseteq Q$ denote the area assigned to the i^{th} agent, where $p_i \in V_i$. A sufficient condition for coverage by N agents is that $\cup_{i=1}^{N} V_i = Q$. This assignment is carried out in a distributed manner by the agents to ensure an equipartition; i.e., $|V_i| = |V_j|$ for all i, j, where $|V_i|$ denotes the area of V_i. We assume that the agents are *lazy*; i.e., $V_i = |Q|/n_i$ for all i at any given time, where $n_i \geq 1$ denotes the number of agents that agent i is aware of, including itself. Unlike [12], we assume that an interaction event is a random occurrence, which is a proxy for two agents entering their mutual communication radius in the course of, for instance, responding to an environmental event.

We consider the gossip-based coverage algorithm described in Algorithm 1. We are interested in the domain Q which is a unit disc with a hole in the centre, and the partitions are assumed to

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

be angular slices. This is equivalent to partitioning S^1 (the unit circle) into sectors. Algorithm 1 includes a specific sequence of interaction events aimed at generating an equipartition. Once all N agents are introduced sequentially (clockwise, without loss of generality), they interact in the opposite order (i.e., agent j with its immediate anticlockwise neighbor) until either coverage is lost, or an equipartition is achieved, or the number of interactions crosses a prescribed threshold. A related sequence, which proceeds *clockwise*, is also presented in [16].

Algorithm 1 Agent interaction and area allocation

Require: Domain $Q = S^1$ and $N > 1$ agents
  Initialize: add agent 1 to Q; $p_1 = \pi/2$ and $V_1 = 2\pi$
  while last agent in $Q < N$ **do**
    Add new agent j; $n_j = n_{j-1} = j$
    $p_{j-1} \leftarrow p_{j-1} \ominus \pi/n_j$, $p_j = p_{j-1} \oplus 2\pi/n_j$
    $|V_j| = |V_{j-1}| = 2\pi/n_j$
  end while
  Initialize $t = 0$, $l = N - 1$
  while Stopping condition not reached **do**
    l_n = anticlockwise neighbor of l
    Update: $n_{l_n} = N$; $|V_{l_n}| = 2\pi/N$; $p_{l_n} = p_l \ominus 2\pi/N$;
    Update $l \leftarrow l_n$; $t \leftarrow t + 1$
    Stopping condition reached if $t = t_{\max}$ or Q is equipartitioned or coverage is lost
  end while

In Algorithm 1, the operators $\oplus$ and $\ominus$ represent clockwise and anticlockwise rotations. Thus, when we write $\theta_3 = \theta_2 \oplus \theta_1$ for $\theta_{\{1,2,3\}} \in S^1$, it is understood that $\theta_{\{1,2,3\}} \in [0, 2\pi)$ (and likewise for $\ominus$).

We proceed to state the main results. The reader is referred to [16] for the complete proof.

Theorem 1. *Suppose $N > 1$ agents are introduced in $Q = S^1$ and that they interact as per Algorithm 1. Then,*

(1) *the addition of a new agent (i.e., the first* `while` *loop of Algorithm 1) does not lead to a loss of coverage.*
(2) *if $N \leq 7$, then Algorithm 1 terminates with the loss of instantaneous coverage if and only if $N = 7$, and with equipartition of Q for $N \leq 6$.*

While Algorithm 1 ensures that coverage is not lost for $N \leq 6$, we can find an alternate sequence of events which leads to loss of coverage for $N = 5$. This occurs when the interaction described in the second `while` loop of Algorithm 1 involves each agent, starting with agent N, interacting with its immediate *clockwise* neighbor.

The theoretical machinery leading to Thm 1 is difficult to extend to cases when N increases beyond 7, and we use a numerical parametric study instead. The application of Algorithm 1 for $N \in [8, 19]$ (the case $N \leq 7$ is covered using Thm 1) shows that it terminates prematurely with loss of continuous coverage as follows: between agents 1 and 2 for $N \in [7, 11]$; between agents 2 and 3 for $N \in [12, 16]$; and between agents 3 and 4 for $N \in [17, 19]$. We plot the uncovered area after the termination of Algorithm 1, as a function of N, in Fig. 1. Although the size of the uncovered area reduces rapidly (albeit not monotonically) with increasing N, the partition at the end of Algorithm 1 is seen to not be an equipartition.

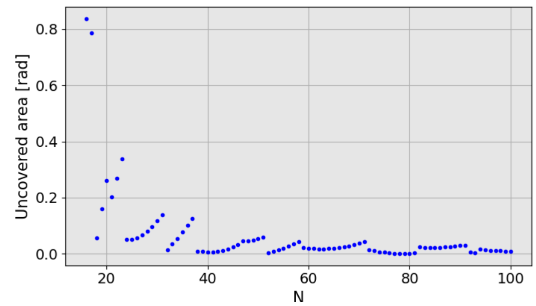

Figure 1: The uncovered area after Algorithm 1 terminates.

3 CONCLUDING DISCUSSION

Although carried out in a simplified setting, our work illustrates how the performance guarantees of coverage algorithms may be sensitive to the number of agents. It is difficult to obtain coverage guarantees rigorously for an arbitrary number of lazy agents when the domain is non-convex and inter-agent communication is limited and local. In a practical setting, our results suggest that the verification and validation of an algorithm prior to field deployment needs to be carried out with the same number of robots as that deployed on the field.

We assumed that the agents repartition and resize their areas of responsibility lazily. Our results show that some degree of altruism, in the sense of partaking a larger area, might be necessary in order to guarantee coverage using a manifestation of Algorithm 1 for an arbitrary number of agents. The extra area, above that recommended by a lazy scheme, could reduce with each interaction so that the eventual area of responsibility matches that of the lazy scheme. It remains an open problem to determine if there exists a sequence of events and an accompanying sequence of reductions which guarantees that Algorithm 1 leads to an equipartitioned domain for an arbitrary number of agents.

REFERENCES

[1] Qier An and Yuan Shen. 2021. Distributed Coverage Control for Mobile Camera Sensor Networks With Anisotropic Perception. *IEEE Sensors* 21, 14 (2021), 16264–16274.
[2] Subhrajit Bhattacharya, Nathan Michael, and Vijay Kumar. 2013. Distributed Coverage and Exploration in Unknown Non-convex Environments. In *Distributed Autonomous Robotic Systems*. Springer Tracts in Advanced Robotics, Vol. 83. Springer, Berlin, 61–75.
[3] Andreas Breitenmoser, Mac Schwager, Jean-Claude Metzger, Roland Siegwart, and Daniela Rus. 2010. Voronoi coverage of non-convex environments with a group of networked robots. In *IEEE International Conference on Robotics and Automation, Anchorage, Alaska*. 4982–4989.
[4] Maria Charitidou and Tamás Keviczky. 2022. An MILP approach for persistent coverage tasks with multiple robots and performance guarantees. *European Journal of Control* 64 (2022), 100610.
[5] Andrew Clark, Basel Alomair, Linda Bushnell, and Radha Poovendran. 2017. Submodularity in Input Node Selection for Networked Linear Systems. *IEEE Control Systems Magazine* 37, 6 (2017), 52–74.
[6] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. 2004. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation* 20, 2 (2004), 243–255.
[7] Jorge Cortes, Sonia Martínez, and Francesco Bullo. 2005. Spatially-distributed coverage optimization and control with limited-range interactions. *ESAIM: Control, Optimisation and Calculus of Variations* 11, 4 (2005), 691–719.
[8] Joseph W. Durham, Ruggero Carli, Paolo Frasca, and Francesco Bullo. 2009. Discrete partitioning and coverage control with gossip communication. In *ASME*

Dynamic Systems and Control Conference, Hollywood, CA, Vol. 2. 225–232.

[9] Andreas Krause and Carlos Guestrin. 2007. Near-optimal observation selection using submodular functions. In *AAAI'07: Proceedings of the* 22nd *National Conference on Artificial Intelligence*, Vol. 2. 1650–1654.

[10] Andreas Krause, Carlos Guestrin, Anupam Gupta, and Jon Kleinberg. 2006. Near-optimal sensor placements: Maximizing information while minimizing communication cost. In *Proceedings of the* 5th *International Conference on Information Processing in Sensor Networks*. 2–10.

[11] Andreas Krause, Amarjeet Singh, and Carlos Guestrin. 2008. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *Journal of Machine Learning Research* 9 (2008), 235–284.

[12] Cameron Nowzari and Jorge Cortes. 2012. Self-triggered coordination of robotic networks for optimal deployment. *Automatica* 48, 6 (2012), 1077–1087.

[13] Luciano C. A. Pimenta, Vijay Kumar, Renato C. Mesquita, and Guilherme A. S. Pereira. 2008. Sensing and coverage for a network of heterogeneous robots. In *Proc. IEEE Conference on Decision and Control, Cancun, Mexico*.

[14] Mac Schwager, Jean-Jacques Slotine, and Daniela Rus. 2009. Decentralized, adaptive coverage control for networked robots. *International Journal of Robotics Research* 28, 3 (2009), 357–375.

[15] Amarjeet Singh, Andreas Krause, Carlos Guestrin, William Kaiser, and Maxim Batalin. 2007. Efficient planning of informative paths for multiple robots. In *International Joint Conference on Artificial Intelligence (IJCAI), Hyderabad, India*. 2204–2211.

[16] Edward Vickery and Aditya A. Paranjape. 2023. Loss of Distributed Coverage Using Lazy Agents Operating Under Discrete, Local, Event-Triggered Communication. arXiv:2302.08729.