Correction to "A Queue-Stabilizing Framework for Networked Multi-Robot Exploration"

Lillian Clark[1], Joseph Galante[2], Bhaskar Krishnamachari[1], and Konstantinos Psounis[1]

WE call the readers' attention to an error identified in the aforementioned paper which appears in IEEE Robotics and Automation Letters, vol. 6, no. 2, pp. 2091–2098, 2021. The paper introduced a novel distributed controller for multi-robot exploration which uses ideas from queue-based stochastic network optimization to autonomously decide at each time step, based on the current state of the system, how to weight network reliability, connectivity, and new information acquisition. We have demonstrated via extended simulations that this controller achieves better coverage than the state of the art and results in desirable emergent behavior.

Unfortunately, Subsection III.F made formal claims about the limiting behavior of the controller (as time and space grow to infinity) whose proof relied on a stochastic network optimization theorem which may not apply. Specifically, the controller running at a robot uses the current state to decide where to move next. The state includes location of the robot, location of other robots and map that the robot currently knows about, and current size of virtual queues that we introduce to track metrics of interest. To prove that the queues are stable, the stochastic network optimization theorem (drift-plus-penalty method) requires that at each time slot there exist a decision that causes the size of the queues to decrease, or remain constant, in expectation. This requirement is not satisfied, e.g. the robot may reach a state where all possible actions increase the queues sizes. As a result, we cannot make formal statements regarding the limiting behavior of the controller. Note however that the formal results still hold in the context of finite time and/or space.

Let us give a specific example. One of the virtual queues tracks the accumulated delay of acquired data which have not been delivered to the sink yet. It is of interest to investigate whether the time to deliver these data is bounded. While this is the case for a finite space scenario, it is not necessarily the case for infinite space exploration.

We end this note highlighting that regardless of the limited applicability of the formal claims, the performance gains of the proposed controller remain valid in practice. What is more, the use of virtual queues to track metrics of interests, e.g. accumulated delay, provides the controller with the insight required to make the decisions that lead to these performance gains.

This work was supported by NASA Space Technology Research Fellowship Grant No. 80NSSC19K1189.

[1]Clark, Krishnamachari, and Psounis are with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California. {lilliamc, bkrishna, kpsounis}@usc.edu

[2]Galante is with NASA Goddard Space Flight Center. joseph.m.galante@nasa.gov