

Robotic Swarm Connectivity with Human Operation and Bandwidth Limitations*

Steven Nunnally¹, Phillip Walker¹, Michael Lewis¹, Andreas Kolling², Nilanjan Chakraborty², and Katia Sycara²

¹ University of Pittsburgh, Pittsburgh, PA 15213, USA.

² Carnegie Mellon University, Pittsburgh, PA 15213, USA.

Abstract. Human interaction with robot swarms (HSI) is a young field with very few user studies that explore operator behavior. All these studies assume perfect communication between the operator and the swarm. A key challenge in the use of swarm robotic systems in human supervised tasks is to understand human swarm interaction in the presence of limited communication bandwidth, which is a constraint arising in many practical scenarios. In this paper, we present results of human-subject experiments designed to study the effect of bandwidth limitations in human swarm interaction. We consider three levels of bandwidth availability in a swarm foraging task. The lowest bandwidth condition performs poorly, but the medium and high bandwidth condition both perform well. In the medium bandwidth condition, we display useful aggregated swarm information (like swarm centroid and spread) to compress the swarm state information. We also observe interesting operator behavior and adaptation of operators' swarm reaction.

1 Introduction

Swarm robotic systems consisting of many simple individual units with limited communication capabilities (e.g., limited radio power) may operate in a wide range of environments from indoor to outdoor underwater environments. Thus, swarm robots may operate under conditions where communication bandwidth is limited. Moreover, depending on environmental conditions, there could be differential capacity of inter-robot communication, or human to robot communication. Furthermore, as swarm systems are usually made of simple units, their localization capability may be poor. These limitations lead to two key challenges in human swarm interaction, namely, (a) the state information of the robots available to the human may not be accurate and (b) there may be a mismatch between the intent of the operator and the robots understanding of the human intent. Due to the localization error, any point in the reference frame of the operator will be erroneously interpreted by a robot as some other point. Thus, any effort by the

* This research has been sponsored in part by AFOSR FA955008-10356 and ONR Grant N0001409-10680. An earlier version of this paper has been submitted to SEMCCO.

operator to move the swarm towards a desired goal will be misinterpreted by a robot, thus creating an intent mismatch between the human and the robot. Current human-swarm interaction (HSI) literature [1–9] does not consider the above aspects of HSI and assumes perfect information transfer between the human and the robots. To close this gap, in this paper, we conduct controlled experiments to study the effect on performance of human-swarm *intent mismatch* and error in swarm state displayed to the operator in supervisory control of swarm robotic systems.

In our experimental scenario, a human operator has to guide a robotic swarm to find unknown targets in a given area. The area is divided into a finite number of regions (whose boundaries are unknown to the operator) and the operator has to match the targets found to the corresponding regions. The robots have a single behavior, namely achieving consensus on direction of motion [10, 11]. The humans can guide the swarm by giving them a point in the environment towards which the robots have to travel. The robots are assumed to have a localization error and the robot position and orientation is assumed to be a Gaussian distribution. In our experiment, each subject performs the mission under three conditions (that are presented to them in a random order), namely, (a) low swarm-to-human bandwidth and low intra-swarm bandwidth (low bandwidth condition), (b) low swarm-to-human bandwidth and high intra-swarm bandwidth (medium bandwidth condition) and (c) high bandwidth between swarm and operator (high bandwidth condition). For the low bandwidth condition, we assume that only one robot can send its state information at each time instant. For the medium bandwidth condition, the swarm members communicate among themselves to estimate their mean orientation and standard deviation of orientation, which is displayed on the screen. In the high bandwidth condition, all the robots send their position and orientation information to the operator. Our experimental results indicate that, as expected, there is a degradation of performance in the low bandwidth condition compared to the high bandwidth condition. However, in the medium bandwidth condition, in which operators had additional information about the standard deviation of orientation (and thereby whether the robots were moving in the direction the human desired), they performed as well as the high bandwidth condition.

2 Experimental Design

The study described below has three within subject conditions with twenty five participants. The user study explores three levels of bandwidth: low, medium, and high. Participants controlled thirty virtual robots in the robot simulator, Stage, to find targets distributed in an open environment [12]. The open environment, displayed on-screen, is divided into six heterogeneous regions, given to the participants on paper. Each region contains exactly one target and the goal is to match all six different colored targets to each region. The study used the Robot Operating System (ROS) as the controller for the robots in Stage [13].

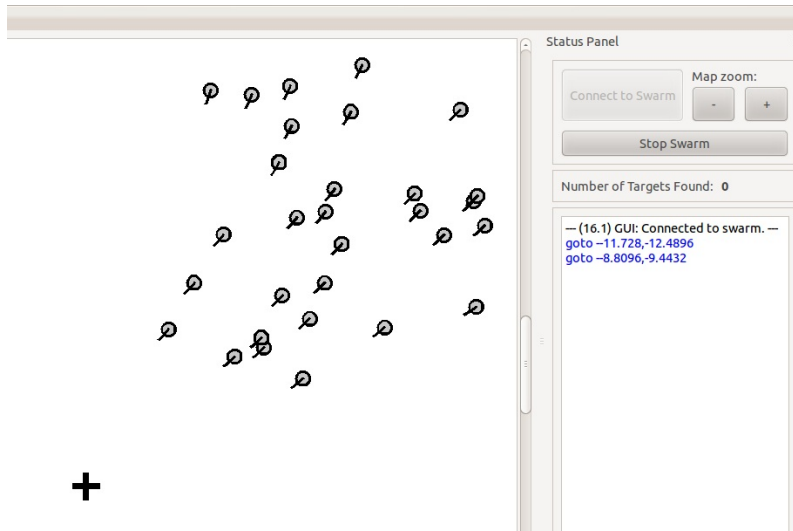


Fig. 1. The GUI used for the study. The left side shows the robots’ estimated position and the right side shows the viewport via which the study participants issued commands. The + shows the endpoint (goal) of the “head-towards” user command.

The operator is given an interface, see figure 1, that displays the states of the swarm from a birds eye, orthographic view. A robots estimated position is displayed as a circle. A line pointing out the front of the robot indicates the current robot heading. The operator can zoom in and out and scroll to manipulate this viewport. The operator can issue two commands: “head-towards” and “stop”. The “head-towards” command is given with a mouse click in the viewport. The “stop” command is issued as a button press.

2.1 Robot Algorithms

The study includes error models for location and orientation, as well as algorithms for the effect of commands on the robots, both as individuals and as a swarm. Location error is simulated with a smoothed, bounded, Gaussian model, with standard deviation of 1.0 meter and mean at the ground truth. The robot shifts in interpolated steps to the sampled error location, then the error is resampled. The location error should make discovery of targets near borders difficult so the participant must use many robots to diminish the error or explore all possible regions for other targets, eliminating regions from the list of possibilities since each regions has one target. When a “head-towards” command is received the robot samples a Gaussian model for its orientation, with standard deviation of $\pi/3$ radians and mean at the orientation vector at the “head-towards” point. The simulation of errors creates a more realistic scenario that considers the constraints of low-cost swarm robots.

The “stop” command is trivial, as all robots halt forward motion. The “head-towards” command starts with the orientation error described above. The robots then move forward at 0.5 m/s and start a standard consensus algorithm, receiving their neighbors’ headings within a communication range of 4.0 meters. Each robot averages their own and their neighbors headings, and adjust their heading to match this average. Robots may lose connectivity before consensus is reached. Due to the nature of Gaussian noise models, the consensed heading will be in general erroneous. It will, however, be closer to the requested “head-towards” point if a greater number of robots are connected. At any point the operator can decide whether the consensus direction is not acceptable and issue a new command, repeating the process.

2.2 Procedure

25 paid participants from the University of Pittsburgh participated in the study. The participants were familiarized with the task and the robot algorithms, and were shown how to use the GUI to issue commands. Every time a robot member of the swarm was close to a target, the robot icon (a circle) on the display would turn the corresponding color, visible to the participants. At the end of each session, the participants were asked to match the color target to each region. Participants were urged to only record a non-answer if they never saw that color target, but if seen the participants were instructed to guess the region. The importance of maintaining one connected swarm of robots was stressed for maintaining lower consensed orientation errors. Participants were told that new “head-towards” commands issued before consensus was reached could adversely affect the connectivity of the swarm. Participants were then given ten minutes to adjust to the interface and train for the task.

The study had three experimental conditions with trials lasting ten minutes each. Participants were given the following conditions in a random order: low bandwidth, medium bandwidth, and high bandwidth. In the low bandwidth condition the robot owning the unique token could update its information on the interface. At each update step the token was transmitted to a random neighboring robot, causing robots with more neighbors to update with higher probability. The interface stored the previous 21 updates. In those 21 updates, one robot could update many times and some would not update at all. Reaching consensus could be difficult to observe since it takes a few updates to see robots moving in the same direction.

In the medium bandwidth condition the swarm aggregated information on location and heading using its local communication network. The standard deviation of the heading was displayed as a proxy variable for the swarm’s heading consensus. The standard deviation of the locations was used to create an ellipsoid around the average location to show the general shape and density of the swarm. Up to four robots could update every half second, which allowed smaller groups that break communication with the main swarm to update their information. Sensed targets were displayed as a colored percentage beside the aggregate display of the number of robots in that group that could sense that color target.

Finally, the third condition was the high bandwidth condition, where all robots updated their position (and the updated position was shown on the screen) every half second. The participant determined when consensus had been reached by observing the movement of the individual robots in the swarm.

3 Results

The data analysis showed differences in the connectivity of the robot swarm across participants and conditions. Connectivity is measured using the second eigenvalue of the communication graph's Laplacian matrix. Since the number of zero eigenvalues shows the number of connected components in a graph and single robots break away early and often, the second eigenvalue is measured on the largest connected component of the graph as long as that group contains over half of the robots.

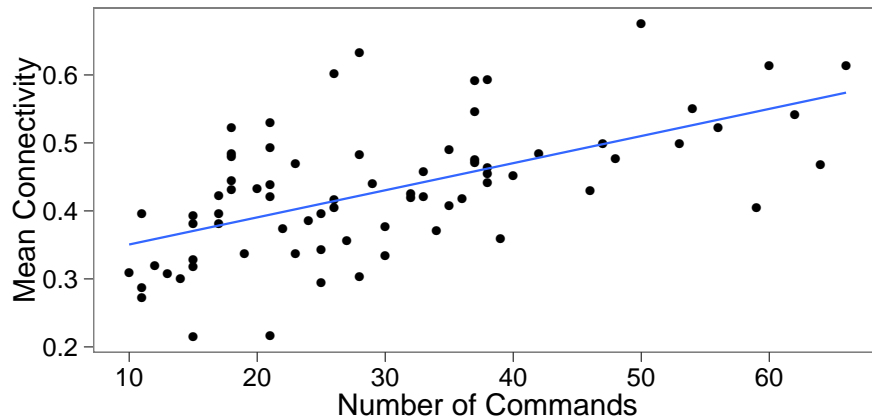


Fig. 2. Each point represents a trial in the experiment. This compares the connectivity throughout each trial to the number of commands sent during the trial. More commands create a more highly connected swarm.

In the instructions participants were warned that every command gave the swarm a chance to break connectivity. Yet, figure 2 shows that connectivity increased with the frequency of commands ($p < 0.001$). This increase in connectivity stems from a difference in participant behavior, see figure 3. Participants who placed reference points closer to swarm enhancing connectivity also issued commands more frequently ($p < 0.001$) keeping the swarm moving. Such a special command, close to the centroid of the swarm, rather than reducing connectivity, can actually improve it. This effective “small and frequent” strategy emerged in a subset of participants despite instructions encouraging infrequent commands.

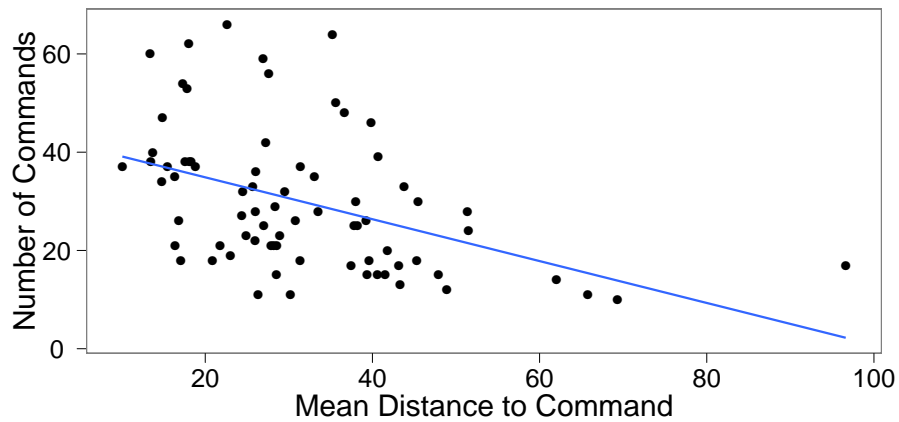


Fig. 3. Each point represents a trial in the experiment. The figure compares the average distance from the swarm’s centroid to each command throughout each trial to the number of commands sent during the trial. More commands have less of a distance to each command explaining the correlation between number of commands and connectivity.

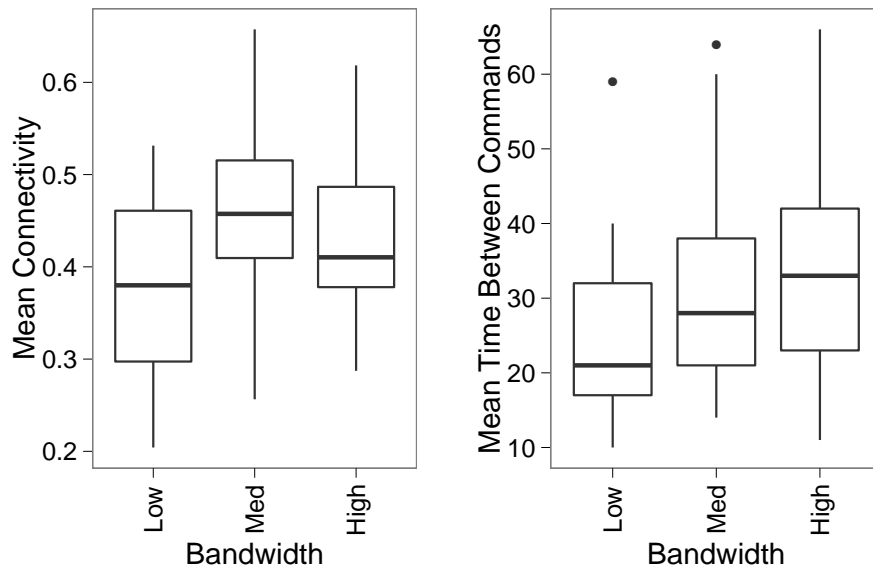


Fig. 4. These graphs display differences of the low condition compared to the other two with regards to connectivity and frequency of commands. Connectivity is lowest in the low bandwidth condition (left). Number of commands over the ten minute trials is lowest in the low bandwidth condition (right). There are 25 participants per condition in each graph.

Between conditions, the low bandwidth condition has a lower average connectivity compared to the other conditions ($p < 0.01$ of the medium condition and $p < 0.05$ of the high condition), see figure 4. The differences between the medium and high bandwidth are not significant. The low bandwidth condition creates an environment where the largest connected component either contains fewer than half of the robots or is sparsely connected. Participants also give fewer commands in the low bandwidth condition ($p < 0.03$ of both medium and high condition), see figure 4. This explains the lower connectivity of the low bandwidth condition as shown in earlier results. The most probable cause of the fewer commands in the low condition case is the time participants must wait for enough robots to update in order to gain knowledge of the state of the swarm before making decisions and giving new commands.

4 Conclusions and Future Work

Despite the large errors in localization for our swarm robots, we have shown that operators can successfully interact with the swarm and its consensus algorithms. The observed operator behavior revealed a variety of interaction strategies, with operators adapting their behavior to the swarm, even in the short time the participant used the system. This was especially interesting for operators that issued frequent commands and then discovered to use some of these commands to improve connectivity, mitigating the otherwise negative effect. This suggests that human operators can, in fact, learn and adapt to swarm dynamics and adapt their instructions to improve the swarm's behavior and state.

In addition, we investigated the effects of bandwidth limitations on the interaction. We have shown that a medium bandwidth condition, which only shows aggregated state information from the swarm, is sufficient for a successful interaction in our foraging scenario. Additional information available in the high bandwidth condition about every individual robot did not improve the interaction with the operator. This result encourages an emphasis on aggregate statistics when considering operator interactions with a swarm. In contrast, the low bandwidth condition not only had a low spatial resolution but also a low temporal resolution. This affected the interaction negatively and lead to worse performance. In future work, we plan to further explore how different environmental constraints affect the performance in human supervisory control of swarms.

References

1. Cummings, M.: Human supervisory control of swarming networks. In: 2nd Annual Swarming: Autonomous Intelligent Networked Systems Conference. (2004)
2. Klarer, P.: Flocking small smart machines: An experiment in cooperative, multi-machine control. Technical report, Sandia National Labs., Albuquerque, NM (United States) (1998)
3. Kira, Z., Potter, M.: Exerting human control over decentralized robot swarms. In: Autonomous Robots and Agents, 2009. ICARA 2009. 4th International Conference on, IEEE (2009) 566–571

4. Bashyal, S., Venayagamoorthy, G.: Human swarm interaction for radiation source search and localization. In: Swarm Intelligence Symposium, 2008. SIS 2008. IEEE, IEEE (2008) 1–8
5. Fields, M., Haas, E., Hill, S., Stachowiak, C., Barnes, L.: Effective robot team control methodologies for battlefield applications. In: Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on, IEEE (2009) 5862–5867
6. Naghsh, A., Gancet, J., Tanoto, A., Roast, C.: Analysis and design of human-robot swarm interaction in firefighting. In: Robot and Human Interactive Communication, 2008. RO-MAN 2008. The 17th IEEE International Symposium on, IEEE (2008) 255–260
7. Goodrich, M., Pendleton, B., Sujit, P., Pinto, J.: Toward human interaction with bio-inspired robot teams. In: Systems, Man, and Cybernetics (SMC), 2011 IEEE International Conference on, IEEE (2011) 2859–2864
8. Kolling, A., Nunnally, S., Lewis, M.: Towards human control of robot swarms. In: Proceedings of the 7th international conference on Human-robot interaction, ACM (2012)
9. Coppin, G., Legras, F.: Autonomy spectrum and performance perception issues in swarm supervisory control. Proceedings of the IEEE (99) (2012) 590–603
10. Reynolds, C.: Flocks, herds and schools: A distributed behavioral model. In: ACM SIGGRAPH Computer Graphics. Volume 21., ACM (1987) 25–34
11. Couzin, I., Krause, J., James, R., Ruxton, G., Franks, N.: Collective memory and spatial sorting in animal groups. *Journal of theoretical biology* **218**(1) (2002) 1–11
12. Gerkey, B.P., Vaughan, R.T., Howard, A.: The player/stage project: Tools for multi-robot and distributed sensor systems. In: International Conference on Advanced Robotics, Coimbra, Portugal (2003) 317–323
13. Quigley, M., Gerkey, B., Conley, K., Faust, J., Foote, T., Leibs, J., Berger, E., Wheeler, R., Ng, A.: Ros: an -open source robot operating system. In: International Conference on Robotics and Automation, Kobe, Japan (2009)