SOCIAL CONTROL OF A GROUP OF COLLABORATING MULTI-ROBOT MULTI-TARGET TRACKING AGENTS

K. Madhava Krishna and Henry Hexmoor
CSCE Dept., University of Arkansas
Fayetteville AR 72701

1. Introduction

We are concerned with controlling large groups of fast-paced software agents with loosely regulated interaction protocols. These settings are commonplace in the military such as controlling multiple robot reconnaissance and multi-target detection, and tracking. The objective is to enable a human controller oversee large groups of individuals.

Offline strategies and contingency plans are not adequately flexible in these settings. We are seeking online methods for commanding and advising collectives in order to adjust how individuals relate to one another and to their group. These methods provide social control. In this paper we offer design and initial implementation of a framework for social control in the domain of reconnaissance and multi-target tracking.

In our framework the human controller is given the following capabilities:

1. Selecting the group to address

2. An urgency level for individuals in the selected group to adopt the guideline for social control (items 4-8).

3. The degree to which individuals may deviate from the guideline for social control (items 4-8).

4. Assignment of level of sociality and benevolence

5. Assignment of level of independence and autonomous decision-making

6. Assignment of level of trust required between group members

7. Assignment of roles

8. Assignment of interaction norms

In the remainder of this paper we will briefly outline related work in section 2. We will then outline our approach to social control of a group in section 3. Finally, we will draw some conclusions in section 4.

2. Related Work

While a lot of research activities are reported in the areas of multi-robotics [1, 2, 3] there has not been a significant volume of literature on multi-robot applications and utility for surveillance operations. Significant among the related work is Konolige's CentiBOTS [4] for indoor reconnaissance, Sukhatme's algorithms for multi-target tracking through mobile sensor networks [5], and collaborative schemes for human robot interaction for a multi-robot system as proposed by Thorpe's group [6]. Victor Lesser's group has made advances to the area of distributed sensor network for object tracking [10, 11]. Lesser's distributed sensor network [10] for real-time tracking is closely related here. Our work is different as it considers multiple targets and mobile sensors (moving robots) in the light of [10] which deal with stationary sensors and single object surveillance.

The work reported in [6] addresses enhanced user-interface and improved dialogue between human supervisor and the robots [6]. The control has been uni-modal and does not scale up well to multiple robots. Specifically, it leads to issues in autonomy when a human gets preoccupied with one amongst the several robots. In our proposed scheme the human user is permitted to alter the social reasoning parameters to produce predictable behaviors. The human user will observe the behavior of the group and modifies the characteristics of the group by proposing an alteration in the social features of the group. The individual robots locally decide how best to change their behavior to comply and yet maintain desirable local properties.

In [5] a strategy for distributing robots across a region using density estimates is reported that also maximizes the number of targets tracked within a

region. The decision for a robot to move to another region or to stay in its current region is based on certain heuristics. The method presented does not address any collaborative or shared reasoning strategies for decision-making and action selection such as the decision for moving to a new area. The only communication between the robots is about their respective positions. In our system decisions such as *whether a robot stops tracking an intruder* or *whether a new robot starts tracking an intruder that was detected by another* are based on collaborative techniques which involve reasoning about their commitments, resources and capabilities.

This work extends several aspects of our cooperative architecture with social functions that we have successfully used for coordinated collision avoidance of several moving robots [7, 8] that crisscross each other frequently and often in large numbers.

3. Social Control of a group for a surveillance task

Outline of our social control architecture is shown in Figure 1. The architecture has embedded in it multi-agent based notions such as cooperation and shared reasoning with a role for social functions that aid in the robot's decision making process. The robot's tracking decision as well as navigational decision is affected by the multi-agent approach. Although our primarily concerns are the tracking issues, since the environment contains many bodies in motion collision avoidance issues are also addressed. We have developed a novel negotiation based scheme for multi-robot collision avoidance [7, 8] that involves cooperation between robots for decision-making as well as a provision for propagating conflicts or requesting robots not involved in a conflict to help resolving a collision conflict between two robots. We have also shown the utility for social functions such as *benevolence* in reducing the number of collision conflicts that arise in a multi-robot setting as well as depicting their role in decreasing the frequency of cooperation and negotiation amidst robots.

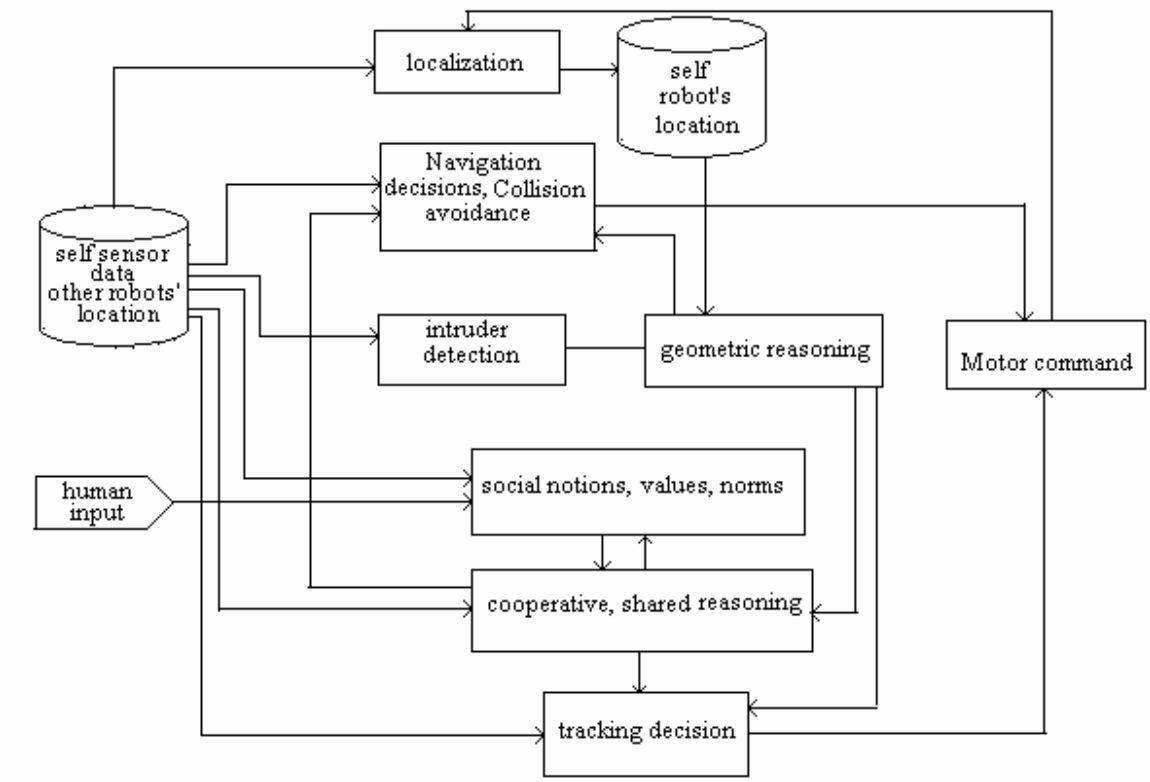

Figure 1. Control architecture for multi-robot surveillance incorporating agent based musings such as cooperation and collaboration with social functions

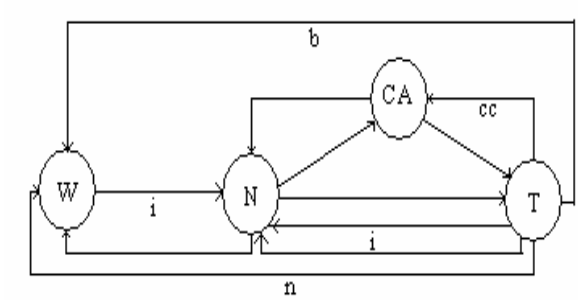

Figure 2. State transition diagram for multi robot surveillance

States	**Transitions**
W: waiting	i: intruder detection
N: negotiating	b: intruder crosses boundary
T: tracking	n: new robot to track the intruder
CA: collision avoidance	cc: collision conflict

Our framework incorporates the existing multi-body collision avoidance strategy with suitable modifications wherever necessary during multi-robot surveillance. Tracking and navigation decisions are also affected by cooperation and negotiation between the robots and the kind of social notions and norms they are implanted with. The state transition model or finite state automata representation for the multi-robot surveillance problem is depicted in Figure 2. We are implementing the processes that occur in those states as well as functions for simulating and detecting the necessary transitions between states. The control architecture and the finite state automata that implements it is distributed across all the robots in the system.

All robots are initially in waiting state W, where they are yet to detect an intruder. They could either be in a state of mobile patrol as in the low-density case or as stationary guards in high-density scenario. A detection of an intruder represented by transition 'i' in the figure sets up the negotiation process, (state N in the Figure 2). The negotiation process decides among a set of robots the one that is best suited to track the intruder. The robot can return to its waiting state or move to the tracking state depending on the outcome of the negotiation. The robot continues to be in the tracking state until the intruder crosses the boundary in which case the robot returns to waiting state. The robot can also return to the wait state if another robot accepts its request to track the intruder or comes forward on its own to track. The robot may enter into negotiations with other robots as it continues to track the intruder before any of the two transitions that permit a return to wait (W) state occur. The robot can also detect a new intruder during the tracking process to enter into negotiation. During the process of tracking the robot may detect collision conflicts (denoted by transition 'cc' in Figure 2) with other moving robots, which are moving either on patrol or tracking other intruders and enters into state CA or the collision avoiding state. Collision conflicts could be solved individually or may need further negotiation before they are resolved. Once resolved the robot enters back into state T or the tracking state.

The negotiation[1] process addresses the following:

1. Determining the group of recipient robots who would receive information about an intrusion detection
2. Negotiating mechanism by which a robot is selected from a group for accomplishing a particular task
3. The strategy by which a tracking robot stops tracking the intruder and hands over the responsibility to another robot. For

[1] We note that the negotiation processes that occur for resolving collision conflicts and resolving issues related to tracking vary in details of implementation though they are represented by the same state, N, in the state transition diagram.

example in the simple case where there's only one intruder as shown in Figure 3, the robot 'Ri' first detects and begins to tail the robot and after a while passes the responsibility to either 'Rj' or 'Rk'. However if 'Rj' and 'Rk' had been tied to other intruders a negotiation may ensue between the robots depending on the motion direction of the other intruders.

4. The role for social values and norms such as *reciprocity* and *benevolence* in the negotiation process
5. The utility of shared reasoning in the negotiation process. For example in figure 4 intruder 'a' is detected by 'Rm' and 'Rp' while intruder 'b' is detected by 'Rn' and 'Rp'. Since the information about 'a' and 'b' is available to 'Rp' it can make use of this to advise 'Rm' not to track 'a' but wait for 'b' and track 'b' and similarly advise 'Rn' to wait for 'a' and not to track 'b'.
6. The scalability and performance for a system of large number of robots and many crisscrossing intruders.

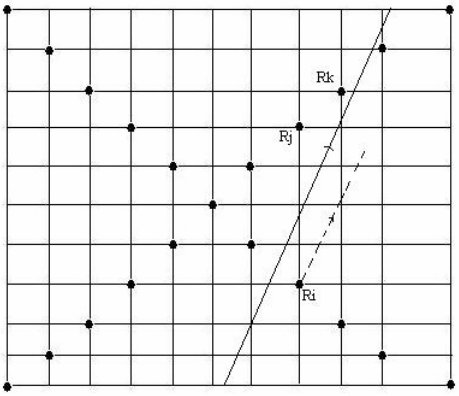

Figure 3. Initial tracking is through robot 'Ri' after which either 'Rj' or 'Rk' take over depending on other commitments

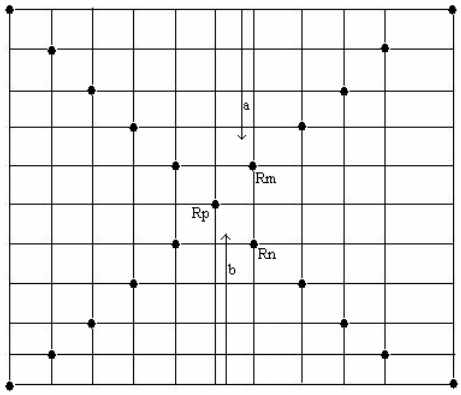

Figure 4. Shared reasoning involving robots 'Rm', 'Rp' and 'Rn' about intruders 'a' and 'b'

Agents in a multi-agent system will benefit by incorporating social cues in their reasoning. Societies and groups often operate and reason on the basis of established conventions, societal norms and values, which enable easier prediction of behavior of the system and the individual agents. In our earlier effort we showed how benevolence [8,9] helps reduce collision conflicts in a multi-robot system. For example, consider a busy cluttered warehouse with mobile robots transporting objects from one place to another. Invariably the robots tend to encounter one another in unexpected ways and may have very little time to avoid collision. Benevolence in such a setting is defined as the propensity of a robot to selectively transfer information to another robot that would be helpful and beneficial to the other robot as well as for the overall group or team objectives. The robot must possess in it a mechanism that enables identification of situations where such selective information transfer would prove beneficial. In a scenario entailing collision avoidance of many robots, we have found that a judicious transfer of information by a robot to robots involved in an impending collision but are unaware of the collision, enhances system performance by reducing collision conflicts and the associated negotiation processes that resolve the conflicts. The robot that transfers the information either advises the robots involved in collision to change their control strategy or itself take a proactive action that resolves collision conflicts amongst others.

In a surveillance setting we also focus on that aspect of benevolence that involves a robot or an agent coming forward to help another agent and

freeing the other of responsibility as shown in Figure 5. Initially, robot Ra tracks the intruder 'a' and robot Rb tracks intruder 'b'. When 'b' changes its direction course Ra offers to track 'b' in addition to 'a' thereby releasing Rb free to take up other commitments.

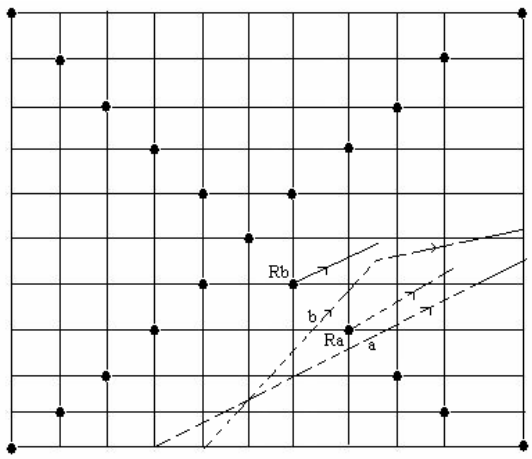

Figure 5. Ra tracks 'a' and Rb tracks 'b' initially. When 'b' changes track Ra offers to track 'a' and 'b' leaving Rb free for other commitments

Reasoning about autonomy, trust, roles, and norms are other social notions we find useful for cohesive group activity of surveillance. We define autonomy as the extent to which a robot's consequent decision or action would be affected by the requests and broadcasts received from other robots. For example, a robot may receive requests to track two objects from two different robots while itself detects a new target to track. Under operation of full autonomy the robot would neglect heeding to other requests and assign paramount importance to track the target it detected. Under variable autonomy the robot reasons with the various requests received and enter a negotiation with the robots concerned. Priorities are assigned to various possible decisions and the decision with highest priority is embarked upon. The robot could also accept the human's choice of decision on the matter. Robots Ra and Rb upon their initial discoveries will announce the targets to other robots in order to intercept targets and to corroborate target trajectory. All robots will report their availability to the target issuing robots. The issuing robot will rank order all incoming offers of assistance and alert the help offering robots. The offering robots must then determine and make known their availability based on their rank ordering they received from the problem issuing robots. The issuing robots consider their availabilities once again and issue another call for help.

Reasoning about trust among agents is useful in determining to delegate and accepting a delegated task. An agent will delegate a surveillance task to an agent in whom it has the most trust. Similarly, an agent is more likely to concede to a task if it trusts the delegating agent. We have developed models of trust and delegation [12]. Trust among agents is developed over a course of interactions and is partly due to agents' performance. Agents performance can be affected due to several factors. For example failure of odometric sensors in a robot can result in poor tracking performance due to localization errors. Other robots would then avoid delegating tracking responsibility to such a robot as far as possible. Trusts can also be affected by autonomy issues. Robots tend to avoid requesting robots for certain tasks that have exhibited a high degree of autonomy in the past.

Patrol robots follow routine patterns that can be considered to compose their roles. For example, horizontal, vertical, perimeter might be three movement patterns as three roles. One consideration of robot role selection is their synergy. If all robots traveled vertically they may not have the best coverage and ability to track targets that cross the region horizontally. Therefore, robots will need to negotiate to discover best patterns of movement.

A norm is a convention that is followed by many agents that govern interaction among agents and may carry a sanction when violated. In surveillance, a norm might be to yield right of way to the robot that is tracking a target by chasing it. A norm's potency is its rate of adherence by the society. For instance, a norm that is universally followed is more potent than one that is followed by half the population.

We conceive of human control at two levels. At a global level the human decides alteration in social functions to evoke better surveillance performance. At local levels he can suggest particular course of action for a robot or a group of robots involved in negotiating over decisions. The

human user also serves as an evaluator of system performance under various degrees of social parameters of benevolence, autonomy, trust and norms. High-level control by a human supervisor allows tuning interactions among agents and control of collective behavior. Since agents are endowed with social abilities and reason about other agents they can be influenced by global parameters of social reasoning. In this section we briefly describe human adjustment of social parameters.

First we consider autonomy reasoning. A human user may determine that commitments to follow targets should persist even though new targets are discovered. Therefore, a robot that has committed to track a target will not consider other targets as they are discovered and would not heed to requests from other robots to track. A user can adjust agent commitment levels to fit the surveillance requirements. At higher trust levels, agents readily delegate tasks to one another whereas at low trust levels, agents are reluctant to delegate or to accept delegation from others. Human supervisor is given the power to alter agent reasoning about trust, which in turn affects their attitude toward delegation. He can at local levels influence agents trust with a particular agent whose performance has been tardy due to sensor inaccuracies, odometric errors and other reasons.

The human input also involves rank ordering of roles. For instance, if most targets cross horizontally, the human might introduce rank for vertical patrol higher than horizontal patrol. The robot must not only consider their synergy, they must also consider the human ranking of roles. Norm adherence is controlled by a human supervisor by adjusting its potency. If a norm is determined to be relevant, human controller will increase the potency of such a role enforcing all agents to adhere to it. Norms can also be adjusted for this sanction. An agent that adheres to a norm is less likely to violate it if the sanction for that norm is set high. Sanction rates are adjusted by a human controller.

4. Conclusion

We have presented here an architecture for distributed collaborative and shared reasoning amongst several agents on a reconnaissance or surveillance mission. The architecture facilitates human interaction with the system at two levels. At a higher level the user is allowed to interact with the entire system or a particular group by modifying the social properties of the group such as benevolence and autonomy. At the lower level the user interacts with a particular robot suggesting changes in its tracking strategies such as deciding to track a new target or dropping a target that is already being tracked. The user is also allowed suggest changes in task delegation decisions made by a robot considering the performance measure of the robot to which the task is being currently delegated.

Acknowledgements

This work is supported by AFOSR grant F49620-00-1-0302.

References

[1]. V. Genevose, R. Magni and L. Odetti, “Self-organizing Behavior and Swarm Intelligence in a Pack of Mobile Miniature Robots in Search of Pollutants”, *Proc. 1992, IEEE/RSJ Int. Conf. on Intelligent Robotics and Systems*, Raleigh, NC, 1992, pp. 1575-1582

[2] L. E. Parker, ALLIANCE: An Architecture for Fault Tolerant Multi-Robot Cooperation, *IEEE Transactions on Robotics and Automation*, 14 (2), 1998.

[3] R. Alami, S. Fleury, M. Herbb, F. Ingrand and F. Robert, “Multi Robot Cooperation in the Martha Project”, *IEEE Robotics and Automation Magazine*, 5(1), 1998

[4] K. Konolige, C. Ortiz, et al., “CENTIBOTS: Large Scale Robot Teams”, http://www.ai.sri.com/centibots/papers/MR-workshop-2003.pdf

[5] B. Jung and G. S. Sukhatme, “Multi-Target Tracking using a Mobile Sensor Network”, Proc., *IEEE Intl. Conf. On Robotics and Automation*

[6] T. Fong, C. Thorpe and C. Baur, “Multi-robot Remote Driving with Collaborative Control”, *IEEE Trans. on Industrial Electronics*, pp. 1-6, July 2002

[7] K. M. Krishna, H. Hexmoor and S. R. Pasupuleti, “Avoiding Collision Log-jams through

Cooperation and Conflict Propagation", to be published in the *Proceedings of KIMAS' 03 (Knowledge Integrated Multi Agent Systems)*, Boston, MA, 2003

[8] K. M. Krishna and H. Hexmoor, "Towards Quantification of Cooperation between Robots", accepted for publication in PERMIS'03, *Performance Metrics in Intelligent Systems*, 2003

[9] K. M. Krishna and H. Hexmoor, "Benevolence in a multi robotic scenario", under preparation

[10] B. Horling, R. Vincent, R. Mailler, J. Shen, R. Becker, K. Rawlins and V. Lesser, "Distributed Sensor Network for Real Time Tracking", In *Proceedings of the 5th International Conference on Autonomous Agents*, pp. 417-424. June 2001.

[11] B. Horling, R. Miller, M. Sims and V. Lesser, "Using and Maintaining Organization in a Large-Scale Distributed Sensor Network", In *Proceedings of the Workshop on Autonomy, Delegation, and Control (AAMAS03)*. July 2003.

[12] P. Poli, and H. Hexmoor, 2003. "Effects of Power on Trust and Autonomy", In AAMAS workshop titled: Autonomy, Delegation, and Control, Melbourne, Australia.