

A Cooperative Distributed Approach to Target Motion Control in Multirobot Observation of Multiple Targets

Stefan Markov and Stefano Carpin

Abstract—This paper addresses a largely ignored perspective in multiple robot observation of multiple targets, i.e. that of the evaders. We present a robust distributed approach to target motion control that utilizes cooperation to minimize the average observation time of the evaders. Targets actively communicate with each other to gain a much richer view of the environment than their immediate sensors provide, and use such information to generate better motion decisions from the group viewpoint. Targets relay information coming from other targets to ensure maximal data spread, and cooperate to escape the pursuers in certain favorable situations. Extensive testing has demonstrated the superiority of the approach over other simpler strategies. Furthermore, robustness with respect to key environmental conditions has been evaluated, and the results suggest the algorithm is a good candidate for real-world applications.

I. INTRODUCTION

Advances in mobile robot technologies have made the use of robots for the execution of jobs performed by humans very appealing, especially for tasks posing danger to the executor. Security, surveillance and data-gathering tasks are examples of such application areas. A growing in popularity scenario is the Cooperative Multi-Robot Observation of Multiple Moving Targets (CMOMMT) problem formalized by Parker in [1], which belongs to the family of pursuit-evasion games. In CMOMMT a group of mobile robots has the goal to observe any intruder within an area of interest. The goal is to keep as many targets as possible under observation by at least one of the robots. Research on the topic has mainly focused on the pursuers' side, while targets are usually assumed to move randomly or according to very simple algorithms. In this paper we explore the other side of the game, i.e. we are interested in developing strategies for the *targets* that minimize the time they are being observed while being in the region of interest.

Solutions to this problem find applications in several areas that require mobile robots to be covert. An example are spying missions that require to collect data about a region of interest while remaining as unobserved as possible. In addition, these algorithms can be implemented on hand-held devices, thus assisting humans in such missions. Furthermore applications are in the gaming industry and computer-based simulations. Last but not least, developing smart target strategies is very important since they provide the second half of the CMOMMT task, thus allowing for a more challenging evaluation for pursuers algorithms.

Our contribution is the development of a *distributed cooperative* control system where targets share information

over a low-bandwidth channel. Evaders exchange and relay messages about the position of intercepting robots and other targets in order to get a much richer view of the world than their immediate sensors provide. They are capable of recognizing certain favorable situations they share and then decide on cooperative evading motion strategies by communicating among each other. For instance, two or more targets can disperse away from each other to escape a single predator following them. The algorithm is realized via a behavioral control system that activates one of six mutually exclusive behaviors. Furthermore, the proposed approach achieves high robustness as it is distributed and minimizes the amount and frequency of data exchange necessary to realize successful avoidance and thus is a good choice for real-world applications on mobile robots.

The rest of the paper is organized as follows: Section II briefly reviews the related work on the topic. Section III gives a formal definition of the problem. In Section IV we present the motion control system. The experimental results and discussion of the findings are shown in Section V. Finally, Section VI concludes the paper and outlines future work on the topic.

II. RELATED WORK

Work related to the investigated research question comes from two separate fields: *covert robotics* and *autonomous surveillance* of mobile targets. The recently introduced field of covert robotics addresses a largely ignored issue in previous research on autonomous navigation systems. The aim is to develop algorithms for robots that need to achieve different tasks covertly, rather than optimally in the sense of minimizing distance, time or power [2]. The authors that established the field have concentrated in covert path planning, where a robot needs to plan a path between two points which has the least exposure to observing sentries or the rest of the free space [3], [4]. Related work on the same topic has been carried in [5]. Research has also been done on dark-spot hiding, in which case the robot needs to plan a path to a place where it can safely hide from being observed for as long as possible [6]. A specific example of an application of covert robotics is demonstrated in [7]. The task proposed in the current paper has not been addressed, even though it is identified as one of the core problems in covert robotics [2], hence the topic we investigate directly relates to the field.

From another perspective, a significant amount of research has been devoted to autonomous surveillance of mobile targets, which also calls for evasive target algorithms. Most of the approaches though take the view point of the observing robots and assume either simple motion strategies for the

Stefan Markov is with the School of Engineering and Science, Jacobs University Bremen, Germany s.markov@jacobs-university.de

Stefano Carpin is with the School of Engineering, University of California, Merced, USA scarpin@ucmerced.edu

targets or a single evader. For instance, in [1] and [8] the A- and B-CMOMMT pursuers' algorithms are presented and there targets are assumed to move either randomly or through the local force method. The idea of the latter is that a target is experiencing a repulsion force from each robot in its field of view inversely proportional to the distance to this robot. Then, the target will move in the direction of the vector sum of all the forces. In [9] a *region-based* robot strategy is described. It assumes that targets have four simple behaviors, *Random-Wandering*, *Wall-Following*, *Random-Turning*, and *No-Move*, and randomly switch between them. Work done in [10] also focuses on pursuer strategies and assumes that the single evader moves either randomly or according to a probabilistic strategy, in which case it goes to the neighboring free cell with the lowest probability of containing a pursuer. Another example where strategies for the pursuer are presented is [11], but there a single target which is either fully or partially predictable is used. In [12] one of the few formal analysis of pursuers' behaviors is presented.

It is clear that work on smart and cooperative target architectures has been lacking and this paper presents the first approach in this direction within CMOMMT.

III. PROBLEM FORMALIZATION

We start from the definition of CMOMMT presented by Parker in [1]:

- 1) S : two dimensional, bounded, enclosed region of interest
- 2) V : team of m robot¹ vehicles, v_i , $i = 1, \dots, m$, with sensors (also called pursuers or predators in the following)
- 3) $robot_sensor_coverage(v_i)$: subset of S observable by robot v_i . This region varies as the robot v_i moves inside S
- 4) $O(t)$: a team of n targets, $o_j(t)$, $j = 1, 2, \dots, n$ (also called preys or evaders in the following)
- 5) $B(t) = b_{ij}(t)$ such that $b_{ij} = 1$ if robot v_i is observing target $o_j(t)$ in S at time t , 0 otherwise. Robot v_i observes target $o_j(t)$ if $o_j \in sensor_coverage(v_i)$.

We extend the definitions above to accommodate that targets also have sensors and a coverage area in which they can detect robots:

- 6) $target_sensor_coverage(o_i)$: subset of S observable by target o_i . This region varies as the target o_i moves within S . If robot $v_i \in target_sensor_coverage(o_i)$, then v_i is being seen by o_i .

To measure the performance of the control strategies Parker also defined the following surveillance metric:

$$A = \sum_1^T \sum_{j=1}^n \frac{g(B(t), j)}{T} \quad (1)$$

where $g(B(t), j)$ is 1 if there exists an i such that $b_{ij}(t) = 1$ and 0 otherwise. More informally stated, this metric measures the average time each target is observed by at least one

¹throughout the paper the term *robot* indicates a pursuer. However targets are also assumed to be autonomous robots. We nevertheless restrict the term *robot* for pursuers to avoid ambiguities.

robot. The pursuers and the evaders have opposite goals with respect to the surveillance metric: pursuers want to *maximize* A , while the targets aim at *minimizing* A .

Furthermore, it is assumed that the sensor coverage of a pursuer is much smaller than S . Otherwise, the task becomes trivial for predators, since several of them can cover the whole environment by standing still at fixed positions. Also, $target_sensor_coverage$ is smaller than $robot_sensor_coverage$ in order to allow the pursuers to approach a target before they reveal themselves to it. The opposite case would mean that predators might not be able to find targets at all. Another assumption is that the pursuers' highest speed is greater than the targets' speed, since otherwise a target can easily escape by just moving forward at maximum speed. Both predators and preys are able to communicate with other robots and targets respectively as long as they are within a certain *communication range* from each other. Lastly, all targets share a common global reference system and use GPS localization, which allows for the unique identification of pursuers based on position estimates.

Finally, this is the first study in the direction of target motion control and we assume an obstacle free environment. This choice is reinforced by the fact that work until now in the CMOMMT framework on predator strategies, such as [1], [8], [12], [13], has also been exclusively concerned with such environments, and extending pursuers' algorithms in CMOMMT to handle obstacles is not an easy task.

IV. ESCAPING THE PURSUERS

Two assumptions about the differences between predators and preys make the task of minimizing the surveillance metric a very challenging one. First, robots are faster than targets. In that way if a target is being followed by a robot it is impossible for it to escape the robot alone, given that the sensor field of view is 360° , has meaningful range, and both are equally mobile. The second assumption is that the robots have a higher sensing range than the targets. As a result, a target can sense a robot only after it has already been seen, after which it cannot really escape by just acting in a selfish mode. Having identified these limitations, the real boost for target performance can only come from efficient *cooperation* among them by compromising between what seems locally best for a target and the benefit for the whole group.

A. Cooperation Gains and Goals

We identify several scenarios where cooperation can increase the overall performance of targets if they cooperate instead of acting greedily. First, an unobserved evader should remain such as long as possible, and targets aware of robots approaching the location of the unobserved prey should warn the evader to prevent predators from revealing it. Second, unobserved targets should not form tight clusters since a single predator approaching such a cluster can gain coverage over many targets simultaneously. Instead, they should disperse while remaining within communication range in order to warn the other evaders about approaching robots. Third, if a group of targets is observed by a single predator, the targets

should move away from each other instead of each trying to escape the predator on its own. Two goals are achieved in this way: first, the time for the group to escape the sensor coverage of the robot is minimized; second, the robot is experiencing decision pressure since it needs to decide which target to follow or it risks to lose all the targets at the same time. The latter scenario is extremely beneficial and also very plausible, especially for A-CMOMMT pursuers. We reiterate that with this behavior each target is compromising between what is best for itself and what is best for the group. The robot may decide to follow a specific target and thus be able to shorten the distance to it much faster. On the other hand, all other targets will escape the pursuer's coverage quicker and the whole group will reduce the average observation time per target. Lastly, while targets are moving away from observing predators they should attempt to disperse them as much as possible and thus decrease the overlap in their sensor coverages. This behavior increases the chance for the favorable scenario of a single predator observing multiple targets, while again compromising between selfish and globally better motion commands.

The implementation should be robust and achieve the above tasks with as little communication as possible. Since we target implementation on real mobile robots, feasibility of computation and stability of the system are very important issues. A distributed approach is the most appropriate in this case, as central planning algorithms have a single point of control, and thus introduce a single point of system failure. Furthermore, communication to the planning point is necessary, which is very difficult to fulfill for real-world systems as the communication range is limited and environmental conditions can disrupt the connectivity to the planning station.

B. Cooperative Distributed Algorithm

The target motion control algorithm described here implements the above-discussed concepts. It consists of two modules: *sensing* and *decision*.

The sensing module has two goals: collect as much information about the world as possible and propagate this information to other targets. It maintains a data structure, called *World State*, which has two arrays containing the known positions within S of robots and other targets respectively. Each position entry is augmented with a timestamp indicating its freshness. Every target broadcasts all the information that it currently has about the environment in a single message called *World State Broadcast (WSB)*. When a target receives a WSB it examines the message looking for additional or more recent information. If needed, it updates its local world state and broadcasts it. In addition, the target reads its own sensors to detect nearby objects and puts these data in the world state. As a result, targets form a communication network in which information is hopping from target to target reaching as many targets as possible.

After the update of the local world state, the *primary follower* of the target is identified. This is the robot closest to the target that can observe it. Any other robot that can see the target is stored as *secondary follower*. Once these data

is collected it is broadcasted on the communication channel as a *Target State Broadcast (TSB)*. In addition, the sensing module reads all TSBs coming from other targets and stores them locally. The TSBs are used in the decision module to make control decisions.

An issue with the exchange of information as described is that targets need to have a synchronized clock in order to judge if the coming information is newer than what they have. Previous research on clock synchronization in distributed systems has successfully developed algorithms that can maintain a shared clock up to a millisecond accuracy [14], and these approaches can be directly applied to our control system, thus solving the problem.

Once the sensing module has updated the world and target states, the decision module determines the behavior to activate. The six mutually exclusive behaviors are *Avoid Robots*, *Avoid Targets*, *Stay Immobile*, *Maximal Target Repulsion*, *Multiple Robots Repulsion*, *Obstacle Avoidance*. Each behavior, except for *Obstacle Avoidance*, drives the robot according to a variation of the local force method described below.

- **Avoid Robots:** In this mode the target attempts to avoid robots within some range. The magnitude of the force vector is dependent on the distance between the target and the robot and follows the curve shown on Fig. 1. Note that the magnitude becomes zero at distances higher than the sensing range, thus a target will also be repelled by robots that it cannot observe, but has learned about through the communication channel.




Fig. 1. Magnitudes of the force vector acting on a target influenced by a nearby robot

- **Avoid Targets:** When in this mode the target attempts to move away from other targets. The force magnitudes are shown on Fig. 2. Targets stop being repelled by other targets when the distance between them gets close to the communication range, thus aiming to keep connectivity.




Fig. 2. Magnitudes of the force vector acting on a target influenced by another target

- **Stay Immobile:** Targets in this behavior do not move at all.
- **Maximal Target Repulsion:** In this mode the target is repelled by a specific subset of targets with magnitudes of the forces equal to -1. The subset is passed as an argument to the behavior.
- **Multiple Robots Repulsion:** While in this behavior the target is repelled by all robots that are able to observe it with force magnitudes according to figure 1 and by all other robots with halved magnitudes.
- **Obstacle Avoidance:** When in this behavior targets determine their motion based on simple obstacle avoidance algorithm which utilizes a potential field like approach.

The classification of the current state data into one of the above behaviors is depicted in the flow chart in Fig. 3. The

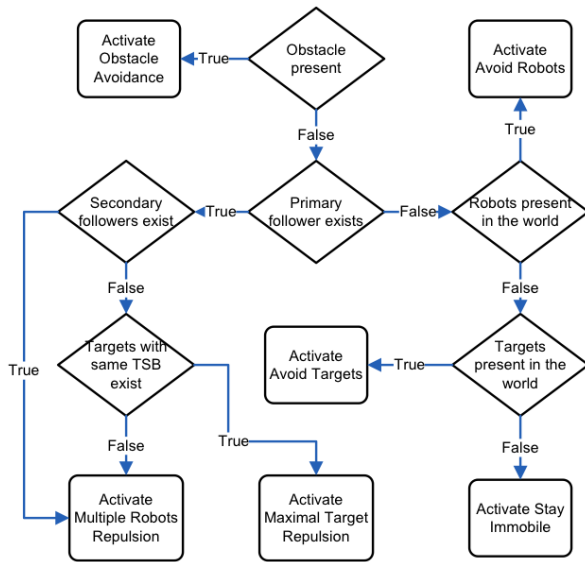


Fig. 3. Decision Module flow chart

decision process follows the concepts described in Section IV-A. Namely, if a target has no primary follower, i.e. it is not observed by any robot, and knows about other predators it will be repelled by them activating Avoid Robots, thus aiming to remain unobserved for as long as possible. If several unobserved targets are not aware of other predators and have formed a cluster they will activate Avoid Targets, thus dispersing the cluster, while keeping within communication range. Lonely targets will stay immobile to avoid stumbling across predators. If two or more targets are followed by the same *single* predator, i.e. they share the same primary follower, but no secondary followers (equivalent TSBs), they will move away from each other with Maximal Target Repulsion. Finally, if a target is observed by multiple robots it will try to avoid them, while still being influenced by the position of other pursuers, thus dispersing them more via Maximal Robot Repulsion.

Once a target goes into the state of Maximal Target Repulsion there is a timeout before it can switch to Avoid Targets, Avoid Robots or Stay Immobile. This is to improve

the robustness of the algorithm with respect to the target communication and sensing ranges. For example, when two targets enter Maximal Target Repulsion they will move in opposite directions because there is a robot approaching them, which in turn tries to get into their center of mass. Since robots have higher sensing range, the targets will stop seeing this robot after certain distance traveled, which is not enough to effectively apply the cooperative evading strategy. Furthermore, the targets might lose connectivity between themselves, thus not being able to get each others TSBs. To prevent this from interrupting the escaping plan the targets execute the last sent speed commands for a certain amount of time before allowing the switching to one of the three above-mentioned behaviors.

V. EXPERIMENTAL RESULTS

Extensive testing has been carried out in order to evaluate the proposed target control algorithm. Its performance was compared with the random motion and the local force methods formerly described. Furthermore, we measured how varying the targets' communication range, communication channel quality and frequency of sending WSBs and TSBs affected the performance.

A. Experimental set up

All tests were implemented using the USARSim simulation framework [15], [16], a high fidelity simulator based on a commercial game engine. Previous tests on the topic were performed using the Player/Stage software [17]. We have kept the simulation parameters the same as in previous tests whenever possible. Both the A-CMOMMT [1] and the B-CMOMMT [8], [12] strategies were used as pursuers' control algorithms.

For the experiments, a circular arena with radius of 30 meters was used. The sensing range of the robots was kept at 4 meters, while the targets had 3 meters with 360° field of view. The robots moved at maximum speed of 0.2 m/s and the targets at 0.15 m/s. The communication range for robots was 10 meters, while for targets it varied between 0 and 10 meters. All vehicles were equipped with GPS localization. The radius of deployment of the robots and targets was kept at 10 meters in the center of the environment and vehicle positions were chosen randomly, ensuring a minimum of 0.5 meters clearance between each two starting locations. Furthermore, for the same targets/robots ratios the same spawning configurations were used in order to negate the effects of random initial placements. For the case of randomly moving targets an algorithm with a 5% chance per second to turn between -90° and +90° was used. Each average value reported is the result of 50 simulations.

B. Comparison to other target strategies

We compared the performance of the behavior target control to the local force method and randomly moving targets for target to robot ratios from 1 to 4. Fig. 4 summarizes the results for A-CMOMMT robots. From the graph it is clear that the cooperative strategy outperforms by far both the local force and the random motion algorithms expressed through




Fig. 4. Performance of all three target strategies against A-CMOMMT for target/robot ratios of 10/10, 10/5, 15/5, 20/5. Lower coverage means increased performance for the evaders.

the lower target coverages. This effect becomes stronger for higher targets/robots ratios. This is explained with the usage of the Avoid Robots behavior which increases greatly the chance that initially unobserved targets will remain such for as long as possible. Furthermore, once a target escapes the observation of a robot it is difficult for the robots to re-gain coverage on this target due to the sharing of information about the robots' positions.

For the case of B-CMOMMT robots the same general trend is observed, only with higher coverages due to the better robot algorithm. Fig. 5 displays the experimental data obtained for all three target strategies.




Fig. 5. Performance of all three target strategies against B-CMOMMT for target/robot ratios of 10/10, 10/5, 15/5, 20/5. Lower coverage means increased performance for the evaders.

These results, while very good, can be expected as other methods do not use any communication and cooperation. As mentioned before though, this approach is the first of its kind for developing online cooperative evading strategies and hence no other similar approaches exist to compare with. Furthermore, computing the optimal coverage times for the targets is first, predator-strategy dependent; second, unfeasible taking into account the large number (20 per

scenario) of vehicles involved, which has not been subject to previous theoretical research in pursuit evasion games.

In addition, the results demonstrate that smarter strategies offer much greater challenges to pursuers' algorithms. This is not only expressed through lower coverage times, but also is visible in the decreased gain for B-CMOMMT over A-CMOMMT. The difference between the two algorithms in target coverages for local force driven targets and target-to-robot ratio of 1 is 12.5%, while for cooperating evaders decreases to 4.0%.

C. Robustness of the Approach

Real-world applications hardly allow for perfect communication among targets, hence it is very important and interesting to examine how the developed algorithm performs with deterioration of the communication possibilities among targets.

The most sensitive parameter to the performance is the communication range of the targets. It can be expected that limiting it will degrade the performance, since the hopping of state information will be restricted and the implemented behaviors will not be that effective. We picked the case of targets/robots ratio of 10/5 and A-CMOMMT robots. Table I plots the results of the performed runs.

Comm. Range (m)	2	4	5	6	8	10
Target Coverage (%)	48.5	48.3	44.2	42.6	43.2	42.5

TABLE I

VARYING TARGETS' COMMUNICATION RANGE FOR A-CMOMMT ROBOTS AND TARGETS/ROBOTS RATIO OF 10/5. LOWER COVERAGE MEANS BETTER PERFORMANCE.

The results demonstrate that limiting the communication range up to 5 meters does not really cause performance degradation. One reason is the sufficient minimum time that targets spend in the Maximal Target Repulsion mode, which ensures that once started, the evading motion to escape a single predator will be successful, unless preempted. Furthermore, robots' sensing range is 4 meters, and 5 meters communication range still allows unobserved targets to remain without coverage due to the state broadcasts. When decreasing the sensing range further though we notice a degradation in the performance of the behavior method expressed in higher coverages, although still outperforming the local force strategy.

The second set of parameters that we varied was the quality of the communication channel and the frequency of broadcasting WSBs and TSBs. These parameters are very important for real-world scenarios for two reasons. First, it is often the case that the communication channel deteriorates because of environmental conditions or external disturbances. Second, since we want the targets to stay as covert as possible we should limit the amount and frequency of information exchange among them. In this way the targets have a much lower chance to reveal their positions, due to a triangulation of the source of broadcasts, for instance.

Fig. 6 presents the results of the performed experiments. The quality of the channel was varied by introducing a

probability for a packet drop. As can be seen, for the case of broadcasting with period of $T = 0.5$ seconds, something easily achievable due to the very fast control loop, even packet drops of 90% do not influence the performance significantly. This is explained with the very low bandwidth that is necessary to achieve successful cooperation strategies. A single WSB containing information about 20 objects is not larger than 500 bytes, hence it can be fitted within a single User Datagram Protocol (UDP) datagram. Decreasing the




Fig. 6. Varying the communication channel quality and broadcast frequency for A-CMOMMT robots and targets/robots ratio of 10/5

period of broadcasting to $T = 30$ s degrades the performance, but nevertheless coverage is less by approximately 10% than the local force method. One reason for degraded performance is the decreased effectiveness of the evading behavior when multiple targets escape a single robot. A thirty seconds period is still enough to warn unobserved targets about approaching predators and hence the boost in performance over the local force method. Furthermore, increasing the drop-rate does not cause a significant performance degradation, partly due to the case when multiple targets have information about the same predator and the chance that a broadcast reaches an unobserved target is higher. As a result the system demonstrates high robustness to decreasing broadcast frequency and communication channel quality. For a broadcast period of $T = 60$ s we observe a further increase in the target coverage, and a much faster convergence to the case of no communication at all, which is expected since we almost deprive the targets from communication.

VI. CONCLUSION

In this paper we presented a study towards intelligent evading strategies in the CMOMMT framework. The designed novel target control architecture demonstrated satisfactory results which firmly reinforce the underlying principles for evasion of multiple pursuers. Under the strong assumptions of superiority of the pursuers, cooperation is a very efficient way to gain performance for the targets. Broadcasting the known world state allowed evaders to form a hop-like network in which the information is shared for the benefit of the whole group. Identifying key situations and cooperating to exploit them drastically boosted the performance

of our system over non-cooperative ones. Furthermore, the minimization of the amount of data broadcasted and the demonstrated possibility to broadcast in relatively large intervals makes the control system very robust to environmental deterioration, and good choice for real-world applications. A drawback is its reliance on GPS localization and the unique identification of pursuers. Further testing has to be carried to determine how errors in the localization of pursuers and their uniqueness will affect the overall performance and how to deal with these situations. The developed architecture is a solid starting point for further work on related tasks. Possible directions include complication of the environment by introducing obstacles and merging the current approach with other research from covert robotics, like dark-spot hiding.

REFERENCES

- [1] L. Parker, "Distributed algorithms for multi-robot observation of multiple moving targets," *Autonomous Robots*, vol. 12, no. 3, pp. 231–255, 2002.
- [2] M. Marzouqi and R. Jarvis, "Covert robotics: Developing robots for covert missions," Monash University, Tech. Rep., 2004.
- [3] M. S. Morzouqi and R. Jarvis, "Fast visibility evaluation for covert robotics path planning," in *Proceedings of IEEE SSRR*, June 2005.
- [4] M. Morzouqi and R. Jarvis, "Covert path planning for autonomous robot navigation in known environments," in *Proceedings Australasian Conference on Robotics and Automation*, 2003.
- [5] A. D. Tews, G. S. Sukhatme, and M. J. Matari, "A multi-robot approach to stealthy navigation in the presence of an observer," in *Proceedings of ICRA*, 2004, pp. 2379–2385.
- [6] M. S. Morzouqi and R. Jarvis, "Covert robotics: Hiding in known environments," in *Proceedings of the IEEE Conference on Robotics, Automation and Mechatronics*, 2004.
- [7] E. L. Hall, P. Cao, and M. Ghaffari, "Remote runway survey system to measure soil type, load bearing capacity, slope and grades using a small, stealthy robot," E.L. Hall Associates, Tech. Rep., 2002.
- [8] A. Kolling and S. Carpin, "Multirobot cooperation for surveillance of multiple moving targets - a new behavioral approach," in *Proceedings of ICRA*, pp. 1311–1316, 2006.
- [9] B. Jung and G. S. Sukhatme, "Tracking targets using multiple robots: The effect of environment occlusion," *Autonomous Robots*, vol. 13, no. 3, pp. 191–205, 2002.
- [10] R. Vidal, O. Shakernia, H. J. Kim, D. H. Shim, and S. Sastry, "Probabilistic pursuit-evasion games: theory, implementation, and experimental evaluation," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 662–669, 2002.
- [11] S. M. LaValle, H. H. González-Baños, C. Becker, and J.-C. Latombe, "Motion strategies for maintaining visibility of a moving target," in *Proceedings ICRA*, 1997, pp. 731–736.
- [12] A. Kolling and S. Carpin, "Cooperative observation of multiple moving targets: an algorithm and its formalization," *International Journal of Robotics Research*, accepted for publication.
- [13] Y. Ding, M. Zhu, Y. He, and J. Jiang, "P-CMOMMT algorithm for the cooperative multi-robot observation of multiple moving targets," in *The Sixth World Congress on Intelligent Control and Automation*, vol. 2, 2006, pp. 9267–9271.
- [14] D. L. Mills, "Improved algorithms for synchronizing computer network clocks," *IEEE Transactions on Networking*, vol. 3, no. 3, pp. 245–254, June 1995.
- [15] S. Carpin, J. Wang, M. Lewis, A. Birk, and A. Jacoff, "High fidelity tools for rescue robotics: results and perspectives," in *Robocup 2005: Robot Soccer World Cup IX, LNCS*, 2005, pp. 301–311.
- [16] S. Carpin, M. Lewis, J. Wang, S. Balakirsky, and C. Scrapper, "Bridging the gap between simulation and reality in urban search and rescue," in *Robocup 2006: Robot Soccer World Cup X, LNAI*, 2006.
- [17] B. Gerkey, R. Vaughan, and A. Howard, "The player/stage project: Tools for multi-robot and distributed sensor systems," in *Proceedings of ICAR*, 2003, pp. 317–323.