

Towards Experimental Analysis of Challenge Scenarios in Robotics

Geoffrey A. Hollinger and Sanjiv Singh

Abstract In this paper, we explore the idea of simulated experimental analysis for challenge scenarios in robotics using the search and secure problem from the Multi Autonomous Ground-robotic International Challenge (MAGIC). The MAGIC problem requires a team of heterogeneous robots to locate, classify, and secure a number of targets in an urban environment with indoor and outdoor areas. We introduce a framework for solving the coordination aspects of the challenge by providing guaranteed clearing strategies (i.e., strategies that ensure coming into contact with any adversarial target). The proposed method allows for repair of the clearing schedule after robot failure, as well as a fall-back strategy if clearing is no longer possible. We analyze scenarios taken directly from the competition, and we utilize repeated simulated trials to validate the hypothesis that strategies designed for locating worst-case targets tend to be more robust to failure than strategies designed for locating average-case targets. Thus, more conservative worst-case methods would tend to perform better if the competition were run many times. However, riskier average-case strategies may win in a single competition. These results demonstrate how insight can be gained from repeated simulated analysis of challenge scenarios in robotics.

1 Introduction

If recent competitions in robotics, such as the Grand Challenge [19] or Urban Challenge [20], were run one-hundred times each, which system would win the majority

Geoffrey A. Hollinger
Computer Science Department, Viterbi School of Engineering, University of Southern California,
Los Angeles, CA 90089, e-mail: gahollin@usc.edu

Sanjiv Singh
Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA
15213, e-mail: ssingh@ri.cmu.edu

of the time? Are some strategies for these competitions risky with high payoff some of the time (i.e., they could come out ahead if the team “got lucky”), while other more conservative strategies are robust to failures and capable of repeated success? Unfortunately, running the actual challenges themselves multiple times to determine the most repeatable strategy is cost prohibitive. However, simulated analysis of scenarios taken from the challenges can provide insight into which strategies are more robust. For such simulated experiments to be meaningful, we must also develop a principled methodology for modeling robotics through computer simulation.

In this paper, we explore the idea of experimental simulated analysis of challenge scenarios by examining the Multi Autonomous Ground-robotic International Challenge (MAGIC) competition [3]. The MAGIC competition requires a team of heterogeneous robots to locate, classify, and secure a number of targets in an indoor/outdoor environment. We extend the G-GSST [6] and FHPE+SA [8] coordinated search algorithms from our prior work to solve the coordination aspects of the challenge. G-GSST provides guaranteed clearing strategies that ensure coming into contact with any adversarial target. FHPE+SA utilizes a non-adversarial model of the target’s motion to optimize the average-case performance of the search.

Our proposed coordination method for the MAGIC competition allows for repair of the clearing schedule after robot failure, as well as a fall-back search based on a non-adversarial motion model of the target if clearing is no longer possible. The key idea behind the proposed coordination strategy is for the robot team to start by generating a guaranteed schedule that ensures mission completion. If a failure occurs during execution, the robots replan to determine if they can still execute a guaranteed schedule by repairing the plan. If no such guaranteed schedule is found, the robots resort to model-based search to complete the mission. If all robots are disabled, then the mission is considered a failure, an outcome that should be avoided as much as possible. Through simulated experiments, we test the hypothesis that a conservative worst-case strategy for locating adversarial targets is more robust to failure than a riskier strategy that assumes the target operates with a known non-adversarial model.

We make a number of improvements to prior techniques in coordinated search for application to the competition scenario. These novel improvements include: the coordination of heterogeneous “disruptor” robots to secure targets, the extension of guaranteed search methods to allow for dynamic replanning, and the integration of a lattice planner for collision avoidance during planning and execution. We repeatedly simulate a scenario taken directly from the competition, and the results demonstrate that conservative clearing strategies are more robust than riskier model-based search strategies. Thus, a model-based search strategy is likely to win some of the time, but a worst-case search strategy would have a higher success rate over many repeated trials. These results demonstrate how insight can be gained through simulated analysis of challenge problems from the robotics community.

The remainder of this paper is organized as follows. We first review related work in multi-robot coordination for search tasks relevant to the MAGIC competition (Section 2). We then discuss some of the design requirements for simulated experiments with challenge scenarios (Section 3). We move on to describe the rules of the

MAGIC competition (Section 4) and implement a scenario from the competition announcement (Section 5). Next, we propose a method for solving the coordination aspect of the challenge (Section 6) and present simulated analysis of variations of this method (Section 7). Finally, we draw conclusions and discuss avenues for future work (Section 8).

2 Related Work

The MAGIC competition [3] draws from problems in the areas of pursuit-evasion, multi-robot coordination, and coordinated search. Search and pursuit-evasion methods typically fall into one of two categories: 1) model-based techniques that maximize the average-case reward of finding a non-adversarial target, and 2) worst-case techniques that guarantee clearing an environment of a worst-case target. Both of these research areas have seen significant treatment in mathematics, theoretical computer science, and robotics.

One of the first formulations of the pursuit-evasion problem on a graph with an adversarial target was proposed by Parsons [16]. The problem of determining the smallest number of pursuers to guarantee catching a worst-case adversarial evader was later found to be NP-hard [14]. These seminal works led to a subfield of computer science known as graph search (see references [4, 1] for surveys). Graph search has also been applied to robotics problems using a number of different formulations [21, 10]. These works assume that target is acting adversarially, which necessitates optimizing the *worst-case* performance of the search.

In contrast with the adversarial search problem, a large body of literature in robotic search treats the target as non-adversarial. This assumption allows modeling the target’s motion and incorporating uncertainty into the search plans [2]. The non-adversarial search problem can be formulated as a Partially Observable Markov Decision Process (POMDP) with the target’s position as the partially observed state [17, 15]. This formulation provides a fully probabilistic representation of the problem, which can easily reason about uncertainty. An alternative approach is to apply techniques such as dynamic programming [12] or branch and bound [11] to determine optimal search strategies without the overhead of the POMDP formulation. Both types of optimal solvers often suffer from scalability issues as the size of the team and/or the complexity of the environment increases. These algorithms for non-adversarial search optimize the *average-case* performance of the search given a known model of the target’s motion.

To move towards unifying worst-case and average-case search, we introduced the G-GSST and FHPE+SA algorithms [6]. G-GSST is a worst-case algorithm that guarantees clearing a physical environment of any potentially adversarial target, and FHPE+SA optimizes a non-adversarial model of the target’s motion to provide good average-case performance. These algorithms allow for scalable, decentralized, and online operation for robotics applications. In addition, we showed that G-GSST and FHPE+SA can be combined to yield a unified approach for solving search tasks.

Thus, they are natural candidates for solving the coordination aspects of the MAGIC competition challenge.

3 Simulated Experiment Design

We now discuss some of the requirements for conducting experiments using simulated challenge scenarios. Like any experiment, a simulated experiment using a challenge scenario must test a hypothesis. In most cases, this hypothesis will likely be that one algorithm or system outperforms another relative to some metric, or perhaps that a class of algorithms outperforms another class of algorithms. Alternative hypotheses may test the nature of the challenge itself. For instance, a simulated experiment may test whether systems with certain desired characteristics are likely to win the challenge. Thus, simulated experiments can also be a tool for designing challenge scenarios.

Another important design decision is whether to test the average-case performance or the worst-case performance of the methods in question. Simulation environments are well-suited for both types of analysis based on the capability of running them many times and either taking the mean or minimum of the performance metric over all trials. This design decision depends on the failure modes in question. For instance, if failures are particularly costly (e.g., a vehicle is destroyed), the worst-case performance is likely an important metric. However, average-case performance is also important as a representation for the expected long-term performance of the method.

Unlike experiments performed using real-robots, an additional requirement for simulated experiments is that the methods analyzed must cross the “reality gap.” In order to cross the reality gap, a system must qualitatively perform similarly in simulation as it does on physical robots. Jakobi argues that sufficient conditions for crossing the reality gap are that the system be “base set exclusive” and “base set robust” relative to a base set of factors modeled by the simulator [9]. A system is base set exclusive relative to a simulation if its performance is fully determined by factors modeled by the simulation. Similarly, a system is base set robust if it performs similarly when the modeled factors are perturbed from the simulation environment to the real-world.

It is important to note that base set exclusivity and robustness are characteristics of a system in conjunction with a simulation, and not characteristics of the simulation itself. In fact, given a simulation environment, one system may be base set exclusive/robust and another may not. For example, consider a 2D simulation environment. Clearly, any system for 3D robotic manipulation would not be base set exclusive in this simulation. However, a system for ground vehicle path planning may in fact be base set exclusive.

Given these two design criteria, it is important to be aware of the base set of factors for a given simulated experiment. In addition, it is important to ensure that this base set varies between trials of the simulated experiment. In the next section, we

will describe the MAGIC competition and design our simulation environment. We will be sure to model the aspects of the challenge that are required to compare the proposed coordination methods, and we will also make a number of assumptions and simplifications, noting that they do not hinder the base set exclusivity of our methods.

4 MAGIC Competition Rules

The MAGIC 2010 competition requires the coordination of a heterogeneous team of unmanned ground vehicles to locate, classify, and secure a number of targets, some of which are moving. The competition guidelines [3] describe the following rules (paraphrased):

1. To complete each challenge phase a team must: a) explore and map the entire phase area; and b) correctly locate, classify, recognize and secure all targets.
2. The team of unmanned ground vehicles (UGVs) is divided into sensor UGVs (for mapping and detection) and disruptor UGVs (to secure detected targets). The team must have a ratio of at least two sensor UGVs to disruptor UGVs.
3. Mobile targets will move at a maximum velocity of 6 km/hour. They may stop, turn, about face, reverse or continue maneuvering.
4. Mobile targets have detection and lethality zones of 10 m diameter (5 m radius), subject to building occlusion. Any UGV entering the lethality zone of a mobile target will be deemed to have been detected and damaged.
5. Mobile targets must be secured to complete the challenge. To secure a mobile target, a disruptor and sensor UGV must simultaneously view it. The target must then be continuously viewed and tracked by both UGVs for a period of 15 seconds.¹
6. Static targets have an activation zone of 5 m diameter (2.5 m radius), subject to building occlusion. Any UGV entering the activation zone of a static target will cause it to detonate. Static targets have a lethality zone of 20 m diameter (10 m radius), subject to building occlusion. If detonated, any UGV, target or non-combatant within the lethality zone will be deemed to have been damaged.
7. To secure a static target, one or more sensor UGVs must locate and classify it, and the team leader must communicate this information to judges. A disruptor UGV must then approach the target to within 2 m of its activation zone (within 4.5 m of the target) and view it for at least 30 sec.
8. During one or more phases of the challenge, at least one UGV will be lost to enemy “sniper” action.

Several rules were simplified in this paper to decouple the problem and highlight the comparison between coordination strategies. As noted above, we argue that these

¹ The maximum sensor range of the UGVs is not specified in the competition rules and is determined by the specifics of the detection system. In this paper, a maximum sensor range of 50 m is assumed.

simplifications do not affect the base set exclusivity of the coordination methods (see Section 3). The following rules were modified as noted:

1. The environment is assumed to be known a priori. If this is not the case, the application of replanning methods is straightforward, though could lead to decreased performance.
2. Non-combatants are not present in the environment. Non-combatants would primarily affect the actions of the disruptors once a target is acquired. Thus, they can be decoupled from the search task.
3. All-to-all communication is allowed between the vehicles through Wi-Fi. Active and/or passive connectivity management techniques [7] could be applied if this assumption were violated.
4. Random failures due to “sniper” action and other failures are not examined. These failures would decrease the likelihood of mission success, but they would obfuscate the analysis of failures due to target proximity.

5 Search and Secure Scenario

We implemented the Phase 2 Scenario (see Figure 1) from the MAGIC information document [3] in the Player/Stage simulator [5]. Player/Stage provides realistic simulation of the sensing and locomotion capabilities of 2D wheeled vehicles. In addition, Player/Stage simulates realistic communication between vehicles in multi-robot coordination scenarios. The simulator uses randomized noise, along with sensor and actuation models, to determine the sensor values and the movements of the robots. In the context of simulated experiment design, the noise helps to perturb the modeled factors and contributes to the base set robustness of the methods (see Section 3).

The 200 m \times 200 m Phase 2 environment contains five targets, three stationary and two mobile. The mobile targets patrol around the building as shown in the diagram at a speed of 4 km/hr. Any searcher or disruptor that comes within 5 m of a mobile target is disabled and deemed unable to function (i.e., it cannot move, detect, or secure targets). If any searcher or disruptor comes within 2.5 m of a stationary target, the target detonates and disables all robots within 10 m.

The searchers start in the northwest corner of the map and must locate and secure all five targets to successfully complete the mission. The searcher vehicles are modeled as having a maximum speed of 10 km/hr, which is the imposed speed limit from the competition. The vehicles are assumed to be car-like with the appropriate non-holonomic constraints. If two vehicles collide, they are also considered to be disabled; however, this event rarely occurs due to low-level collision avoidance.

The collision avoidance between robots is provided by a lattice planner, which was implemented specifically for this domain and based on work by Ferguson et al. [13]. The lattice contains a regularly spaced number of points in the $[x, y, \theta]$ space of each robot. Each robot plans from its current location to a lattice point




Fig. 1 Phase 2 environment from the MAGIC competition guidelines [3]. A heterogeneous team of UGVs must locate and secure all mobile and stationary objects of interest in the environment.

near its goal location using A^* search. The lattice is built online to contain obstacle regions for the current positions of the other robots. The resulting planner is run as new information becomes available to avoid inter-robot collisions. Thus, if a robot moves into the path of another robot, the path is replanned to avoid a collision.

6 Coordination Strategy

The key idea behind the proposed coordination strategy is that the team starts with a guaranteed clearing schedule and attempts to repair this schedule if failures occur. If a guaranteed schedule is no longer possible, the team resorts to an efficient search schedule. A modification of the G-GSST algorithm was used to generate guaranteed schedules for the search and secure scenario. First, the environment was discretized into 32 cells (shown in Figure 2) using a region growing technique. The largest dimension of the cells was limited to 50 m to correspond to the UGV sensor range. Next, G-GSST was run to find a clearing schedule on the discretized map.

After a clearing schedule was found, the locations of the disruptors were determined at each step using a sequential minimization of the average distance of each searcher to a disruptor. Each disruptor determines what one-step move it should




Fig. 2 Polygonal representation of MAGIC environment (left) and discretization using a region growing technique (right).

make to minimize the metric and then shares that information with the other disruptors. The other disruptors consider the move fixed when planning their own moves. Note that this disruptor coordination strategy does not require explicit coordination between the robots (only information sharing), and it is decoupled from the generation of the search schedule.

When a target is located during the search, the searcher calls for the nearest disruptor to come secure the target. If that searcher is not available, the next furthest disruptor is called until an available disruptor is found. The disruptor moves to the current location of the searcher, rather than the centroid of the cell. Thus, the discretization can be coarser than the necessary detection radius and still allow for target securing. After the target is secured, the clearing schedule continues.


When robots come too close to a target, they are disabled, and the remaining team replans to determine if a clearing schedule can still be executed from their current locations. If a clearing schedule is not found, the team executes FHPE+SA to maximize the efficient search objective function (i.e., the discounted probability of locating a slowly moving random target). This strategy is then executed until either all targets are secured (mission success) or all disruptors are disabled (mission failure).

7 Simulated Missions


Figure 3 shows snapshots of the Player/Stage simulator as a seven robot team (two disruptor UGVs and five sensor UGVs) successfully locates and secures all targets in the Phase 2 environment.²

² Multimedia extensions showing playback from the Player/Stage simulation for various runs are available at the following URL:


<http://robotics.usc.edu/%7egeoff/iser10/>



(a) 0 sec: the searchers start in the northwest corner of the map. (b) 2 min 24 sec: the first mobile target is secured.



(c) 5 min 25 sec: the first stationary target is secured. (d) 9 min 54 sec: the second stationary target is secured.



(e) 13 min 15 sec: the second mobile target is secured. (f) 14 min 18 sec: the final target is secured.

Fig. 3 Snapshots of Player/Stage simulation of robots executing a MAGIC search and secure task. Dark areas may contain targets; light areas are cleared. Searcher UGVs are shown in blue, disruptor UGVs in yellow, and targets in red.

A number of simulated trials were run to determine the robustness of the proposed solution. If a robot is disabled (due to getting too close to a target), the rest of the team replans to determine if a clearing schedule can be found from their current position. If a clearing schedule is found, they begin execution as before (considered a soft failure). If one is not found, they execute FHPE+SA to maximize probability of detecting the remaining targets (considered a hard failure). If all disruptors or searchers are disabled, the mission is considered aborted. Table 1 gives a summary of these robustness results.

Table 1 Robustness trials for MAGIC 2010 scenario. The cases with no failures and soft failures result in a guaranteed clearing schedule. The hard failures result in a completion without a clearing guarantee. Averages are over 200 trials; errors are one standard deviation.

5 Guaranteed Searchers, 2 Disrupt.	Percentage	Average Completion Time
Completion (no failure)	72.5%	687.41 ± 39.48 sec
Completion (soft failure)	8.0%	755.69 ± 57.75 sec
Completion (hard failure)	9.5%	854.42 ± 313.84 sec
Unable to complete	10.0%	-
6 Guaranteed Searchers, 2 Disrupt.	Percentage	Average Completion Time
Completion (no failure)	76.5%	643.44 ± 43.83 sec
Completion (soft failure)	9.0%	706.44 ± 125.93 sec
Completion (hard failure)	7.5%	885.07 ± 327.73 sec
Unable to complete	7.0%	-
6 Guaranteed Searchers, 3 Disrupt.	Percentage	Average Completion Time
Completion (no failure)	75.5%	634.86 ± 43.27 sec
Completion (soft failure)	14.0%	644.39 ± 53.91 sec
Completion (hard failure)	8.0%	698.06 ± 180.06 sec
Unable to complete	2.5%	-
5 Guaranteed, 1 Efficient, 2 Disrupt.	Percentage	Average Completion Time
Completion (no failure)	44.0%	503.60 ± 94.96 sec
Completion (soft failure)	32.0%	705.16 ± 77.55 sec
Completion (hard failure)	14.5%	821.28 ± 403.01 sec
Unable to complete	9.5%	-

The results in Table 1 show completion times and percentages for different failure modes using a varying number of robots. The proposed coordination strategy achieves as high as a 97.5% success rate for mission completion with 6 searchers and 3 disruptors. Failures occurred in approximately 25% of trials, and the majority of failures eventually led to mission success using one of the two recovery strategies. The average completion times show that completion without failure was faster than completion with a soft failure, and completion with a hard failure had a high variance (ranging from fast to slow completions).

In addition, Table 1 shows results where one searcher is allocated to optimize a random target model using FHPE+SA while the other five searchers execute a clearing schedule (a similar idea as proposed in prior work [6]). The reallocation of an efficient searcher leads to a reduction in the mean completion time, but a large number of recoverable failures occur when the efficient searcher is disabled. Based

on these results, the user may be tempted to run an efficient strategy using FHPE+SA that does not guarantee clearing. However, it is important to note that few of the strategies that resorted to FHPE+SA led to mission success. In addition, running FHPE+SA from the start results in a mission failure in this scenario for all three varying robot numbers. The high failure rate of FHPE+SA likely results from the tendency for searchers to spread out, which taxes the disruptors to respond quickly. Thus, the efficient strategy is a “more risky” strategy than the clearing strategy. These results further confirm the trends in prior work [6], and they validate the hypothesis that coordination strategies for the MAGIC competition that are designed to locate worst-case targets are more robust than model-based strategies for locating average-case targets.

8 Conclusions

This paper has provided simulated experimental analysis of coordination techniques for the MAGIC 2010 competition. A scenario from the competition announcement was analyzed in the Player/Stage simulator. It was shown that modifications of G-GSST [6] and FHPE+SA [8] provide robust and effective solutions to many of the coordination problems posed in the MAGIC competition. The algorithms were also extended to allow for the coordination of a heterogeneous team of robots performing different, but interrelated, tasks of searching and securing.

The proposed solution to the MAGIC problem allows for search and secure with a heterogeneous team of searchers and disruptors, and was able to achieve mission success as much as 96% of the time. In addition, failure modes were explored that allowed for recovery of the clearing schedule and resorted to efficient search as necessary. The results in this paper demonstrate that riskier strategies that utilize a non-adversarial model of the target tend to fail more often than more conservative strategies for finding worst-case targets. This trend coincides with insights from our prior work [6] as well as ongoing work in the search literature [18]. Simulated experiments allowed for this analysis, which would not have been possible based on the few robot trials from the competition itself.

In addition to the analysis of the MAGIC scenario, we have provided a preliminary discussion of the necessary experimental design requirements for running simulated experiments with scenarios taken from robotics competitions. Such experiments have the potential to determine the average-case and worst-case performance of the systems that will be used (or have been used) in competitions. However, careful experimental design is necessary to ensure that the simulated environments accurately model the characteristics of the physical world necessary for a meaningful comparison.

Future work includes the simulation and analysis of a large number of systems for upcoming and past challenge scenarios, as well as the design of more comprehensive requirements for accurate simulated comparison. In addition, actual data from the competitions can be utilized within the simulation to increase the realism of the

simulated experiments (e.g., real-world sensor data used offline to guide the planning algorithms). Through such analysis, we can hope to gain insight into the most repeatable challenge solutions, rather than simply the winning challenge solutions.

Acknowledgements The authors thank Joseph Djughash, Benjamin Grocholsky, and Matthew Aasted from Carnegie Mellon University as well as Gaurav Sukhatme from the University of Southern California for their insightful comments.

References

1. B. Alspach. Searching and sweeping graphs: A brief survey. *Matematike*, 59:5–37, 2006.
2. F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte. Optimal search for a lost target in a bayesian world. *Field and Service Robotics*, 24:209–222, 2006.
3. Defence, Science and Technology Organisation, Land Operations Division. Multi Autonomous Ground-robotic International Challenge: Guidelines, June 2009. <http://www.dsto.defence.gov.au/MAGIC2010/>.
4. F. V. Fomin and D. M. Thilikos. An annotated bibliography on guaranteed graph searching. *Theoretical Computer Science*, 399:236–245, 2008.
5. B. P. Gerkey, R. T. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proc. Int. Conf. Advanced Robotics*, 2003.
6. G. Hollinger, A. Kehagias, and S. Singh. Improving the efficiency of clearing with multi-agent teams. *Int. Journal of Robotics Research*, 29(8):1088–1105, 2010.
7. G. Hollinger and S. Singh. Multi-robot coordination with periodic connectivity. In *Proc. IEEE Int. Conf. Robotics and Automation*, 2010.
8. G. Hollinger, S. Singh, J. Djughash, and A. Kehagias. Efficient multi-robot search for a moving target. *Int. Journal of Robotics Research*, 28(2):201–219, 2009.
9. N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6(2):325–368, 1997.
10. A. Kolling and S. Carpin. Pursuit-evasion on trees by robot teams. *IEEE Transactions on Robotics*, 26:32–47, 2010.
11. H. Lau, S. Huang, and G. Dissanayake. Optimal search for multiple targets in a built environment. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2005.
12. H. Lau, S. Huang, and G. Dissanayake. Probabilistic search for a moving target in an indoor environment. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2006.
13. M. Likhachev and D. Ferguson. Planning long dynamically feasible maneuvers for autonomous vehicles. *Int. Journal of Robotics Research*, 28(8):933–945, 2009.
14. N. Megiddo, S. Hakimi, M. Garey, D. Johnson, and C. Papadimitriou. The complexity of searching a graph. *Journal of the ACM*, 35(1):18–44, 1988.
15. S. C. W. Ong, S. W. Png, D. Hsu, and W. S. Lee. POMDPs for robotic tasks with mixed observability. In *Proc. Robotics: Science and Systems Conf.*, 2009.
16. T. Parsons. Pursuit-evasion in a graph. In Y. Alavi and D. Lick, editors, *Theory and Applications of Graphs*, pages 426–441. Springer, Berlin/Heidelberg, 1976.
17. N. Roy, G. Gordon, and S. Thrun. Finding approximate POMDP solutions through belief compression. *Journal of Artificial Intelligence Research*, 23:1–40, 2005.
18. J. Strom, R. Morton, K. Reilly, and E. Olson. Online probabilistic pursuit of adversarial evaders. In *Proc. ICRA Workshop on Search and Pursuit/Evasion in the Physical World*, 2010.
19. S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L.-E. Jendrossek, C. Koelen, C. Markey, C. Rummel, J. van Niekerk, E. Jensen,

- P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, and P. Mahoney. Stanley: The robot that won the darpa grand challenge. *Journal of Field Robotics*, 23(1):661–692, 2006.
20. C. Urmson, C. Baker, J. Dolan, P. Rybski, B. Salesky, W. L. Whittaker, D. Ferguson, and M. Darms. Autonomous driving in traffic: Boss and the urban challenge. *AI Magazine*, 30(2):17–29, 2009.
 21. M. Vieira, R. Govindan, and G. S. Sukhatme. Scalable and practical pursuit-evasion with networked robots. *Journal of Intelligent Service Robotics*, 2(4):247–263, 2009.