

Internalized Plans for Communication-Sensitive Robot Team Behaviors

Alan R. Wagner, Ronald C. Arkin

*Mobile Robot Laboratory, College of Computing
Georgia Institute of Technology, Atlanta, USA, {alan.wagner, arkin}@cc.gatech.edu*

Abstract

Autonomous teams of robots operating in a dynamic, adversarial environment stand to benefit from using all available resources. But how can knowledge be used to construct a plan that does not interfere with the robots ability to react to its environment? In this research we distill abstract representation into a plan usable by a reactive behavior-based architecture. This plan is then exploited to enhance the performance of a team of robots tasked with maintaining communications while performing reconnaissance. Utilizing multiple plans in serial and in parallel is shown via simulation to be a promising method for increasing mission performance. We conclude that the utility of these internalized plans warrants further investigation as a method for imbuing reactive agents with a priori knowledge.

1. Introduction

Reconnaissance, like surveillance, is a dynamic, distributed sensory problem in an adversarial environment demanding a strongly cooperative solution to achieve a goal [13]. These problems require agents capable of coordinated sensing, processing, and communication [7].

Purely reactive robotic architectures, such as those advocated by Brooks [3], allow a robot to operate in dynamic environments but abandon traditional planning and knowledge representation. Yet a variety of information—such as terrain and logistics advice—is available that could aid robotic reconnaissance missions. Clearly, if information from these sources could be accessed in a manner that did not hinder the robotic team's reactive performance it would assist the mission.

Several hybrid deliberative/reactive architectures have been developed in an attempt to address the shortcomings of purely reactive robotic architectures [1,5,12]. Other approaches attempt to morph higher-level plans into forms suitable for reactive behaviors [8, 9, 10].

Payton also delineates a method for using information in a reactive architecture [14]. In this method, a priori

knowledge becomes an enabling resource for decision-making. From his perspective, traditional plans are artificially abstracted from knowledge that often results in over- or under-specification of a mission's objectives. By minimizing symbolic abstraction, one can develop a plan for action that may be used directly by a reactive agent. Payton brands this type of plan an internalized plan. These internalized plans differ from traditional plans in their lack of symbols and their tight representational coupling to the needs of a reactive robot. Moreover, the plans are used only as advice, injecting world map or other types of knowledge only at the discretion of the robot.

In this research, we employ Payton's notion of internalized plans to fortify reactive control with a priori map knowledge. We then use this knowledge to coordinate the communication and coverage objectives of a team of robots performing a reconnaissance mission in a simulated urban setting. This hybrid deliberative/reactive approach is integrated with Arkin's behavior-based motor schema architecture [2] using the *MissionLab* [11] behavior specification software.

2. The Architecture

The goal of this research is to produce a team of heterogeneous mobile robots capable of coordinated self-sustaining communication while simultaneously satisfying mission requirements. It is hypothesized that by utilizing plans as resources a robotic team's performance will improve on specific communication and coverage metrics in a reconnaissance task. As mentioned above this system was implemented on a hybrid deliberative/reactive architecture employing Arkin's motor schemas. In the following sections we detail the methods used for the creation and utilization of map knowledge as internalized plans for coordinated team behavior.

2.1 Generating Internalized Plans

Prior to a mission, a uniform cost search algorithm [15] is employed to generate cost values for each non-obstacle cell of a grid mesh. Because the distance from any cell to a neighboring cell is assigned a cost of

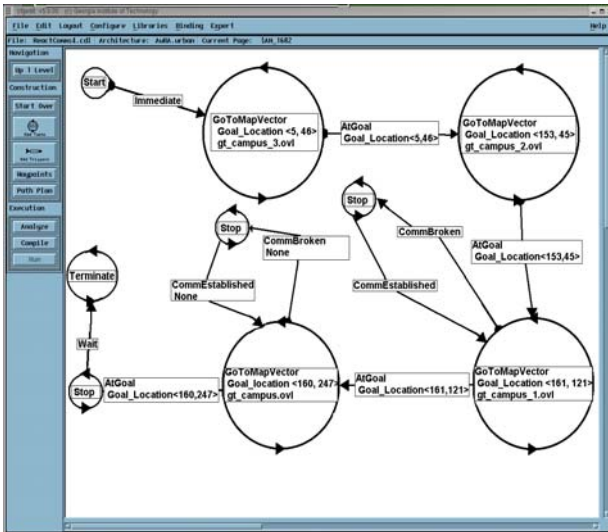


Figure 1: The mission specification for the leftmost robot in the serial plans experiment (Section 2.4).

either 1 or $\sqrt{2}$, the values associated with a location represent an approximation of the distance to a goal location rather than the exact straight-line distance from one point to another. The grid mesh for which the cost values have been generated constitutes an internalized plan.

Software was developed that generates internalized plans based on a list of obstacles, if any exist, and the goal. Currently the system operates on both *MissionLab* style overlay formats [6] and MITRE specified OPAR formats [4]. A 250 by 250 meter grid, at 1 by 1 meter resolution, internalized plan can be completed in approximately 20 seconds using standard 1Ghz PC hardware. Upon completion the software generates a resource file that can be stored in a database or shipped along with the robot executable for immediate use on the robot itself. This resource file may contain any number of plans based on different sets of constraints (obstacles, communications, etc). Finally, using the *MissionLab* behavior specification system (Fig. 1), a robot configuration is generated that incorporates the internalized plan behaviors. These behaviors prescribe which plan to use and the manner in which to use it.

2.2 Using Internalized Plans

During mission execution, the robot uses its current location in space to index an internalized plan that produces a movement vector towards the minimal adjacent grid cell. This imposes a requirement of robot localization to use the plan effectively. Fortunately for the outdoor environments for which this DARPA project is intended, GPS can provide this information reasonably accurately. The resulting vector advises the robot to move in only one of eight possible directions. Although this produces sub-optimal advice it greatly




Figure 2: A vector field from an internalized plan

reduces the practical and computational complexity of determining the grid cost values. Remember also that the plan is only advice that is injected into the actual robot controller that contains the collection of other active behaviors, such as *move-to-goal*, *avoid-obstacles*, *noise*, etc. This tends to smooth the digitization bias present in the internalized plan. Figure 2 depicts a typical internalized plan.

2.3 Plans in Parallel

As discussed by Payton [14], internalized plans offer an advantage in that several plans can be utilized simultaneously. We incorporate multiple plans in two ways: in parallel and in serial. Figure 3 depicts the use of parallel internalized plans. Serial plans are discussed in Section 2.4.

Parallel plans consist of more than one internalized plan, each with the same goal location but differing in the constraints on which the plan is based. Moreover selecting a plan (or fusing them) among several in parallel occurs in a single time step. A parallel plan is created by first generating the least constrained plan for that goal-environment pair. Next, additional constraints are added in the form of obstacles to the internalized plan and the plan generation algorithm is rerun. This process can be continued for as many constraints as desired. Determining which vector to output can be achieved in a number of ways. One choice is to select vectors from the most constrained internalized plan that affords a viable route to the goal. This strategy is extremely simple and easily implemented, but rather inflexible. Alternatively cost functions with tunable gains can be created that represent each constraint. This method can greatly increase the systems flexibility and could allow for advice vectors that represent a mission's dynamically evolving constraints. Because plans with additional

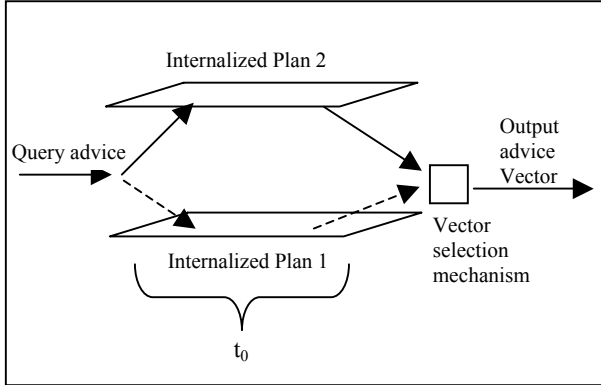


Figure 3: Plans used in Parallel: Advice is queried from plan 2, which results in an output vector. The output occurs in time t_0 .

constraints contain all of the known obstacles from plans with fewer constraints, the robot is never led into an obstacle occurring in a less constrained plan. Moreover, determining if a plan has a route to the goal is accomplished instantly as any plan that does will return a finite cost value. This selection mechanism can either be used statically or dynamically. Static selection chooses the most constrained plan in advance and then follows that plan to completion. Dynamic selection reevaluates which internalized plan should be selected and utilized at each time step.

We investigated the use of statically selected parallel plans. In this case, a parallel plan consisted of one internalized plan depicting the placement of physical obstacles in the environment and one internalized plan further constrained by communication attenuating areas. For the results reported in this paper, the experimenter visually selected potential communications attenuating areas such as alleys, parked cars, etc. Ideally, in future work a selection mechanism will be based on actual automated terrain analysis.

2.4 Plans in Serial

Another way to exploit multiple internalized plans is to utilize each plan in serial. In this case each plan may maintain a different goal location and a perceptual trigger determines which plan to use during different time steps. Figure 4 graphically depicts a serial plan. Perhaps the most intuitive example of plans in serial is a contingency plan. Regarding this communications-related research, a simple contingency plan might consist of a location for the team to meet at if communications fail or degrade, or before moving on to another section of the environment. Perceptual triggers such as being located at a specific point or recognition of a communications failure, for example, transition the robot to query advice from a different internalized plan.

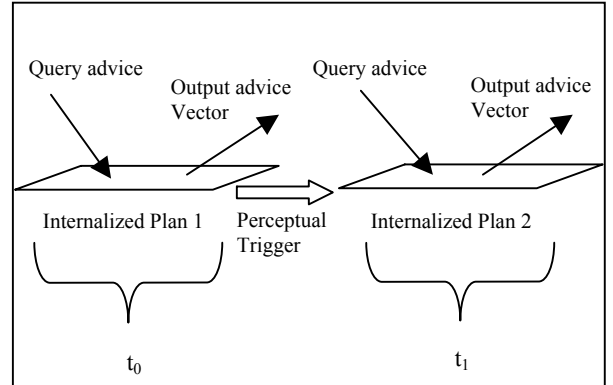


Figure 4: Plans used in Serial: Advice is first queried from plan 1. Some perceptual trigger causes the transition from plan 1 to plan 2. Thereafter advice is queried from plan 2. The output from each plan occurs in during a unique time step.

The use of seven different plans in serial was investigated. Each robot utilized a unique subset of these seven plans to improve coverage while also maintaining communication. Initially every robot queried a plan that guided them to unique locations that would maximize the coverage of a large (250m by 50m) rectangular area. Upon reaching the location a perceptual trigger invokes another plan that directs all robots back to a common point located at the entrance to a street. Once communication is re-established the robots utilize different plans to maximize coverage as they traverse a narrow pathway to the goal.

3. Performance Metrics

We defined several metrics that characterize performance in real world reconnaissance operations. As maintaining robust communication connectivity is a primary focus of this study, funded under DARPA's MARS Vision 2020 program, the amount of time the team operated as a single communicating network, the amount of time each robot is connected to another robot, and the overall connectivity was logged in our current simulation studies. Ideally there would only be one connected network within the team. Whether full connectivity is desired will depend on other mission constraints, such as the desirability of coverage, relative to the importance of maintaining communication.

Coverage is another important criterion for evaluating reconnaissance missions. In general, it is assumed that the probability of discovering a target increases with increased coverage of the terrain. For this study, it is further assumed that each robot had the ability to sense a three by three meter area in all directions from a point location, excluding obstacles.

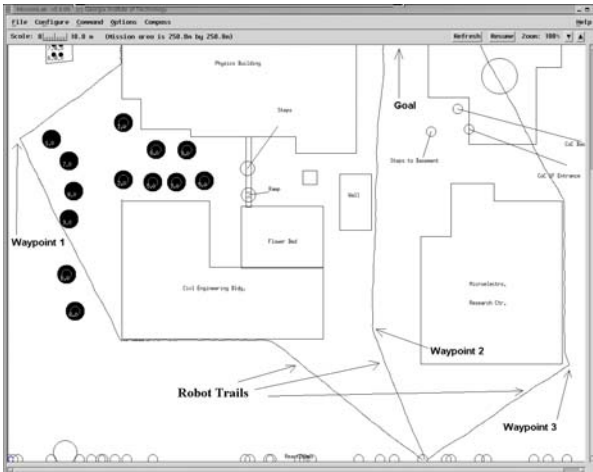


Figure 5: Reactive behaviors guide the robots through the environment. Each robot first goes to a waypoint and then proceeds to the goal (not shown).

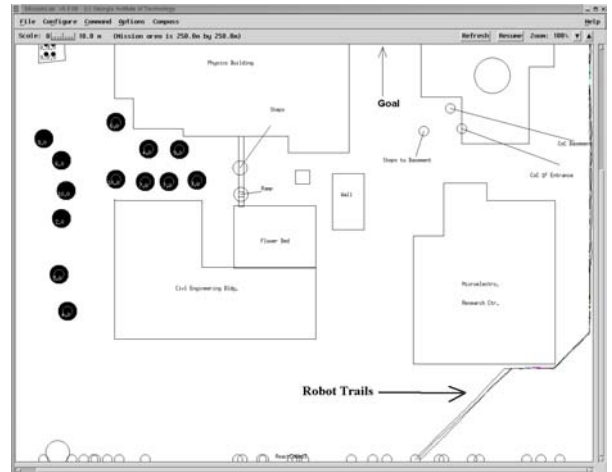


Figure 7: Plans are used in parallel. In this case the most constraining plan avoids communication-attenuating locations by circumventing the buildings.

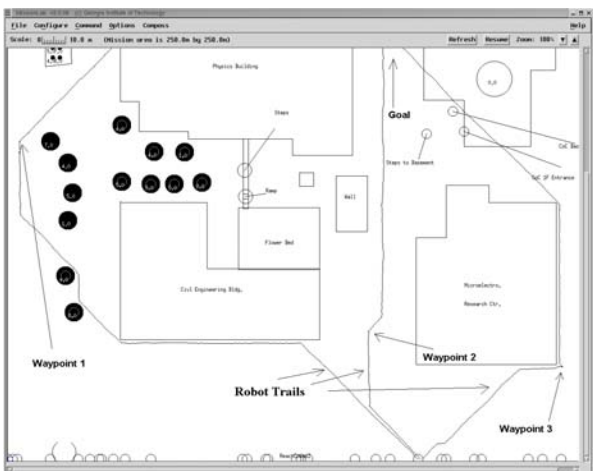


Figure 6: Navigation using an internalized plan constructed with only obstacle knowledge guides the robots. As in the previous experiment, each robot proceeds to a waypoint before going to the goal.

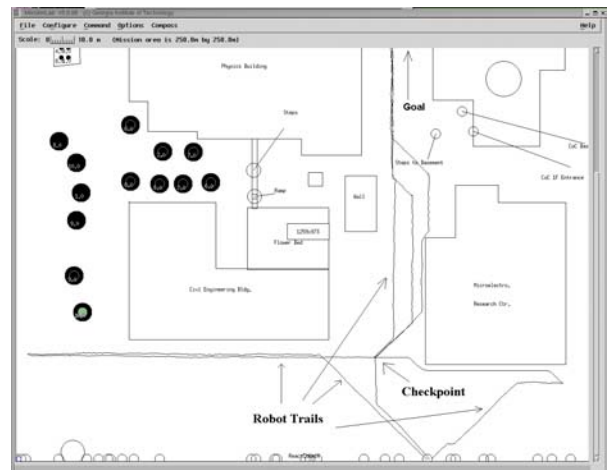


Figure 8: Serial plans first cover the initial open area before regrouping and following the street located in the center to the goal.

Time to complete the mission and the overall distance traveled are also considered for metrics, as successful reconnaissance may correlate to timeliness in mission completion and the total distance traveled is an indirect measure of energy expended.

4. Experiments

To evaluate the system four experiments were performed on a simulated section of the Georgia Tech Campus. Three simulated robots were tasked with navigating from a start location to a goal location approximately 250 meters away. A start position was randomly chosen along the x (bottom) axis. Figures 5-8 depict the environment and the routes taken during a typical run of each experiment. For the experiments shown in Figures 5 and 6 the robots were required to

stop at intermediate waypoints (shown in the figures) to assure initial dispersion. A detailed network attenuation model has also been integrated into the *MissionLab* system for accurate simulation of networking failures. This model currently reduces the communication signal strength when a robot is occluded by terrain. Thirty trials per experiment were run. All experiments included a small amount of random noise behavior (output vector gain of .11).

To establish baseline performance initially the use of purely reactive behaviors was first examined (Fig. 5); where each robot's behavioral specification consisted of only reactive *GoTo* behaviors. This behavior produces a single **instantaneous** vector resulting from the summation of the attraction to the goal (or the designated waypoints), repulsion from obstacles and the small contribution from the noise behavior. This behavior is purely reactive. As stated earlier, each

robot in this experiment proceeded to a single intermediate point before moving onto the goal.

In the next set of experiments, the use of communications-insensitive internal plans was examined. Three experiments were conducted involving different proportions of using plans as advice. In the first experiment a 0.66 gain adjusted reactive vector was summed with a 0.34 gain adjusted planning vector. A second experiment summed a 0.34 gain adjusted reactive vector with a 0.66 gain adjusted planning vector. The final experiment consisted of only using the planning vector. These plan vectors simply reflected the location of physical barriers on a map. Although the start, intermediate, and endpoints for this experiment are identical to those of the earlier experiments, the actual route taken by the robot is in part determined by the robot's plan rather than the instantaneous potential field. The performance on all metrics was expected to be about the same whether using purely reactive behaviors or communications insensitive plans.

Next the use of plans in parallel (Fig. 7) was examined. Here two plans were developed: One described paths to the goal circumventing only the physical obstacles, the second described paths to the goal circumventing both the physical obstacles and also the areas likely to attenuate communication. Upon beginning the mission the robot determines if the most constrained plan can be implemented, where a constraint might refer to communication attenuation, fuel reserves, or terrain coverage. If so, it uses this plan. If not, it uses the next less constrained plan. We tested only the case in which the more constrained of the two plans could be utilized.

Finally, the use of plans in serial (Fig. 8) was examined. A succession of plans was selected so as to afford improved terrain coverage. The robots initially employ a plan that guides them in different directions to ensure coverage. Upon reaching a predefined area, a perceptual trigger transitions the robot to use a different plan. This plan directs each robot toward a predefined checkpoint. From here they move through a narrow (~25 meter) street. There is no centralized robot-to-robot coordination enforced, rather each robot operates independently. Again the robots use their internalized plans with spread goal locations to maximize coverage while also exploiting specially defined perceptual triggers to maintain communication links to the other robots (Fig. 1). These triggers query *MissionLab's* network attenuation model to determine if a robot can communicate with another robot. If so the robot proceeds toward its goal. If not, the robot stops and waits until communication is re-established. As currently implemented this behavior is somewhat deadlock prone, yet it suffices as a proof of concept for multiple internal plans used in serial. Additional behavior specifications, such as one that causes the

robot to wander when communication is lost or a simple timeout mechanism, would quickly solve any deadlock problems.

5. Results

It was conjectured that by utilizing plans as resources a robotic team's performance would improve on communication and coverage metrics in a simulated reconnaissance task. The results demonstrate that internalized plans can be used to improve the team's ability to communicate and that additional resources (e.g. reliable knowledge) improve performance when integrated smoothly into a behavioral controller. Figures 9 and 10 compare the results obtained using various proportions of plan advice. Figures 9 and 11 compare the percent time each robot team maintained a single complete mobile ad hoc network. Figures 10 and 12 compare the mean area covered per experiment. Finally, Figures 11 and 12 compare the use of additional map resources in the form of serial and parallel internalized plans when compared to reactive behaviors alone and plans without additional map resources.

As indicated by figure 9, no significant difference in the ability to communicate ($0.51 < p < .98$ for all comparisons) exists when using communications insensitive plans. Some significant differences in coverage, however, do result (fig. 10). Specifically, using reactive behaviors alone results in significantly greater coverage ($p < .04$ in all cases). This increase in coverage is likely the result of differences in obstacle navigation. Plans guide the robots through obstacle fields in a near optimal manner. Reactive behaviors, on the other hand, approach an obstacle until repulsive forces guide them around the object. Thus, the robot may loiter around the obstacle gaining small bits of coverage before finding a path past object. No significant difference in coverage was observed when varying the proportion of planning included ($.08 < p < .97$).

Although not directly comparable, a significant (vs. reactive $p = 0$ for both) and drastic difference occurs when communication-sensitive internalized plans are incorporated (fig. 11 and 12). The communications planning in serial experiment strives to improve coverage while also maintaining adequate communications via checkpoints and perceptual triggers. Although a significant reduction in communication performance occurs ($p = 0$) when plans are used in serial vs. parallel, this difference represents the mission specification and not the use of internal plans per se.

Figure 12 presents results for the mean area covered by communication-sensitive planning, communication-insensitive planning, and using reactive behaviors

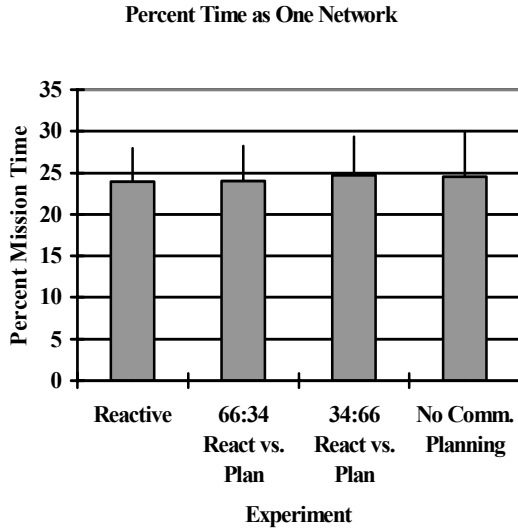


Figure 9: Percent time as a single network per experiment. The lack of additional resources in the form of internalized plans results in no significant improvement.

Mean Total Area Covered

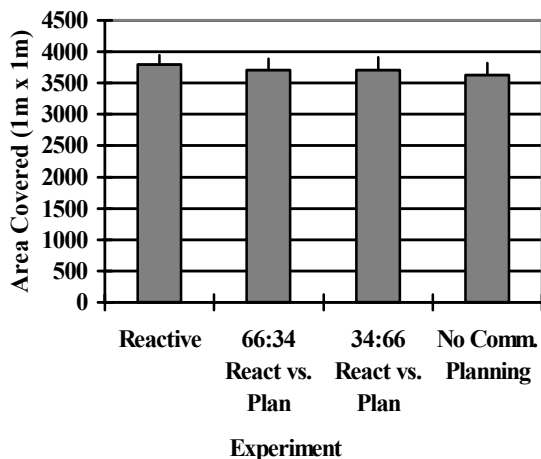


Figure 10: Percent time as a single network per experiment. The use of additional resources in the form of internalized plans results in a drastic improvement in communications.

alone. The reactive and communications-insensitive planning experiments result in the greatest coverages. In these cases, the robots essentially diverge pursuing their own directions relative to the goal unconstrained by communications concerns. In the communications planning in parallel experiment, a route for each robot is generated that avoids communication-attenuating areas while navigating to the goal. This route is the same for all robots, thus resulting in near perfect communications and poor coverage. The communications planning in serial experiment somewhat alleviates this by guiding the robots to explore specific areas of the environment. Therefore, although we cannot claim that serial plans improve

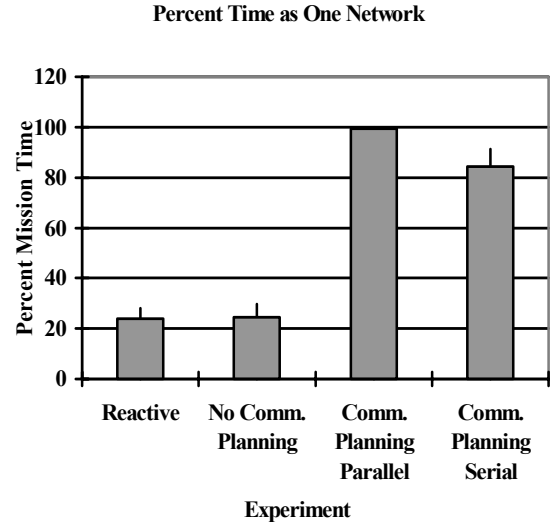


Figure 11: Mean Total Area covered per experiment. The use of additional resources in the form of internalized plans results in less coverage although serial plans aid coverage when compared to a single parallel plan.

Mean Total Area Covered

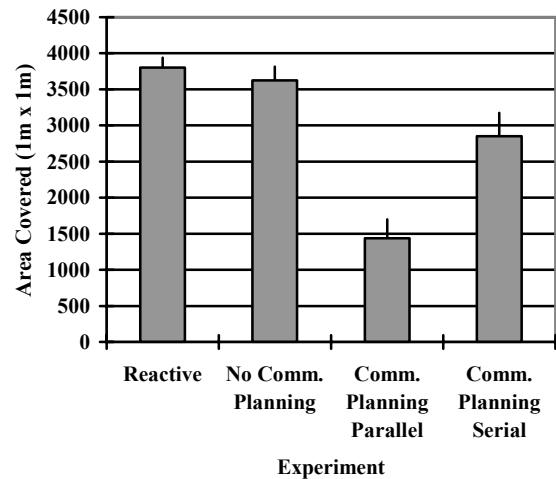


Figure 12: Mean Total Area covered per experiment. The use of additional resources in the form of internalized plans result in less coverage although serial plans aid coverage when compared to a single parallel plan.

coverage, we do maintain that a suitably constructed behavioral specification that utilizes internalized plans can result in improved coverage and improved communications. Moreover, in all cases the coverage statistics assume that the robot can only sense (or provide coverage for) in its immediate 3 by 3 meter vicinity. This is probably an underestimate. If we assume larger areas of detection visibility the differences in coverage will likely become less

Table 1: Results on performance metrics relevant to reconnaissance*

	Time (time steps)	Distance (meters)	Time connected (time steps)	Percent time fully connected (time steps)
Reactive Mean:	66.6	834.7	1435.7	21.0
Standard Deviation:	11.6	80.2	180.1	2.8
66:34 React. Vs. Plan	49.5	797.0	1408.8	20.6
Standard Deviation:	4.3	71.4	173.3	2.4
34:66 React. Vs. Plan	45.3	800.0	1372.6	20.6
Standard Deviation:	4.0	71.2	146.2	2.4
Plan No Comms:	53.8	820.8	1347.5	20.4
Standard Deviation:	8.7	71.5	128.0	2.7
Plan Comms Parallel:	60.6	1017.8	2051.9	98.9
Standard Deviation:	13.0	313.1	343.7	0.8
Plan Comms Serial:	116.8	915.6	3560.9	76.3
Standard Deviation:	23.7	69.3	317.5	7.0
t-Test Reactive Vs. 66:34 React. Vs. Plan	P = 0	P = 0.06	P = 0.56	P = 0.56
" " Vs. 34:66 React. Vs. Plan	P = 0	P = 0.08	P = 0.14	P = 0.56
" " Vs. Plan No Comms	P = 0	P = 0.4786	P = 0.0339	P = 0.4013
t-Test 66:34 Vs. 34:66 React. Vs. Plan	P = 0	P = 0.87	P = 0.39	P = 1.0
" " " Plan No Comms	P = 0.02	P = 0.20	P = 0.12	P = 0.76
t-Test 34:66 Vs. Plan No Comms	P = 0	P = 0.26	P = 0.48	P = 0.76
t-Test Plan Comms. Parallel Vs. Plan Comms. Serial	P = 0	P = 0.0897	P = 0	P = 0

*N = 30 for all trials. Confidence Interval = 95%. P values are double sided. P = 0 when $t > 3.014$ or $P < 0$.

significant. Future experimentation should lend proof to this claim.

Table 1 depicts the robot team's performance on the other reconnaissance metrics. Regarding the use of communications-sensitive plans versus reactive behaviors, the time required to complete a mission increases when the plan guides each robot to explore more terrain. Distance traveled increases when using communication-sensitive planning. This is a reflection of the plans themselves. The time each robot is connected via communications to other robots increases with communications-sensitive planning. Finally, the percentage time as a fully connected network reflects the general trend in connectivity already observed: communications-sensitive planning helps maintain communications.

Table 1 also provides the results from hypothesis testing of reactive behaviors only vs. reactive planning mixes. For most metrics, different proportions of communications-insensitive planning do not make a significant difference. The sole exception is mission time. In this case, adding some amount of planning decreased the time to complete a mission significantly. On the other hand, solely plan-driven missions required significantly more time than reactive plan mixes.

6. Conclusions and Future Work

We conclude that additional resources in the form of communications-sensitive internalized plans aid performance on the studied communications metrics. These additional resources help the robot select a path that will necessarily improve these metrics. It should be noted that there is no attempt to supplant reactive

control with a priori planned control. Rather the robot is merely offered additional resources for use in determining its trajectory. As the paths produced by the reactive-behaviors-only and the communications-insensitive planning experiments were based on similar resources, it is not surprising that they had similar results. Yet, the fact that these two dissimilar methods generated nearly identical results is in itself revealing. This seems to indicate that when confronted with the same environment, the internalized plan method selects a path exceedingly similar to that of purely reactive methods.

On the other hand, when some degree of planning was included, the time required to complete a mission decreased significantly versus using either reactive behaviors or planning alone. This decrease in mission time is likely due to the combination of the plan shuttling the robot through obstacles and the reactive behaviors smoothing digitization bias.

The experiments employing parallel and serial plans illustrate two potential methods for maximizing performance on the specified communication and coverage metrics. Strictly speaking these experiments cannot be directly compared to the other experiments—the paths chosen were based on different resources. Even so, as methods for generating additional resources for the robot, the parallel and serial techniques discussed above are shown to be effective.

Thus far we have assumed that the environment meets two criteria: 1) the robot's internal knowledge of the environment accurately reflects the actual environment, and 2) the robot is capable of localizing itself within the environment. Since internalized plans are used as

advice, we suspect that these two criteria are malleable or may not be necessary at all. In future work we intend to explore this possibility.

Other future work will attempt to reproduce these findings using situated physical robots in an actual urban environment. We also hope to investigate merging reactive/deliberate techniques for coping with unexpected obstacles. Finally, we would like to investigate algorithmic methods for determining areas of terrain maps that might be communication attenuating and include the cost functions that best describe these areas.

7. Summary

The extent to which reactive plans improve the communication and coverage of a simulated robotic team in an urban environment has been investigated. Improvements in communication occur when planning includes communication-attenuating factors. Parallel internalized plans may be more appropriate when communications constraints cannot be relaxed. Moreover, parallel plans offer the possibility of adding additional constraint knowledge to a single internalized plan. Serial internalized plans, on the other hand, afford a method for employing different plans in different situations. Although coverage degrades when using communications-focused plans, the use of plans in serial offers hope for improving coverage while also maintaining communications. Therefore we conclude that by including all available knowledge resources, a robotic team's performance can be made to improve on specific communication and coverage metrics in a reconnaissance task.

Acknowledgments

This research is funded under DARPA/DOI contract #NBCH1020012 as part of the MARS Vision 2020 program that is a joint effort of Georgia Tech, UPenn, USC and BBN.

8. References

- [1] R.C. Arkin and T. Balch, "AuRA: Principles and Practice in Review", *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2):175-189, 1997.
- [2] R.C. Arkin, *Behavior-Based Robotics*, MIT Press, Cambridge Mass. USA. 1999.
- [3] R.A. Brooks "Intelligence without Representation", *Artificial Intelligence*, 47:139-159, 1991.
- [4] G. Comparetto, S. Kao, J. Marshall, and N. Schult "OPNET Path Attenuation Routine (OPAR) Description Document", MITRE Tech. Rep., V. 1.0, Aug. 2001.
- [5] J. Connell, "SSS: A Hybrid Architecture Applied to Robot Navigation", *Proc. IEEE Intern. Conf. on Robotics and Automation*, pp. 2719-2724, 1992.
- [6] Georgia Tech Mobile Robot Laboratory, *User Manual for MissionLab* version 5.0, www.cc.gatech.edu/ai/robotlab/research/MissionLab/, January 2002.
- [7] C.P. Diehl, M. Sapharishi, J.B. Hampshire II, and P. Khosla, "Collaborative Surveillance Using Both Fixed and Mobile Unattended Ground Sensor Platforms", *SPIE's 13th Annual International Conference on Aerospace/Defense Sensing, Simulation, and Controls*, Vol. 3713, April, pp. 178-185. 1999.
- [8] R.J. Firby, and M. Slack, "Task Execution: Interfacing to Reactive Skill Networks", working notes, *AAAI Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents*, Palo Alto, CA. pp 97-111. 1995.
- [9] E. Gat, "Integrating Planning and Reaction in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots", *Proceedings of the AAAI*, 1992.
- [10] D. Lyons, and A. Hendriks, "Planning for Reactive Robot Behavior", *Proc. IEEE Intern. Conf. on Robotics and Automation*, Nice, FR. pp 2675-2680. 1992.
- [11] D.C. MacKenzie, *Design Methodology for the Configuration of Behavior-Based Robots*, Ph.D. Diss., College of Computing, Georgia Inst. of Tech., 1997.
- [12] C. Malcolm, and T. Smithers, "Symbol Grounding via a Hybrid Architecture in an Autonomous Assembly System ", In *Designing Autonomous Agents*, pp. 123-144, MIT Press, Cambridge, MA. 1990.
- [13] Lynne E. Parker, "Cooperative Robotics for Multi-Target Observation", *Intelligent Automation and Soft Computing*, special issue on Robotics Research at Oak Ridge National Laboratory, 5 (1), 5-19, 1999.
- [14] D. Payton, J. Rosenblatt, and D. Keirse, "Plan Guided Reaction" In *IEEE Transactions on Systems Man and Cybernetics*, pp. 1370-1382, 20(6), 1990.
- [15] S.J. Russell, P. Norvig, *Artificial Intelligence A Modern Approach*, Prentice Hall, Upper Saddle River NJ, 1995.