

Frontier-Based Exploration Using Multiple Robots

Brian Yamauchi

Navy Center for Applied Research in Artificial Intelligence

Naval Research Laboratory

Washington, DC 20375-5337

yamauchi@aic.nrl.navy.mil

1. ABSTRACT

Frontier-based exploration directs mobile robots to regions on the boundary between unexplored space and space that is known to be open. Previously, we have demonstrated that frontier-based exploration can be used to map indoor environments where walls and obstacles may be in arbitrary orientations. In this paper, we show how frontier-based exploration can be extended to multiple robots. In our approach, robots share perceptual information, but maintain separate global maps, and make independent decisions about where to explore. This approach enables robots to make use of information from other robots to explore more effectively, but it also allows the team to be robust to the loss of individual robots. We have implemented our multirobot exploration system on real robots, and we demonstrate that they can explore and map office environments as a team.

1.1 Keywords

Mobile robots, exploration, map-building, multi-agent coordination, multi-robot teams

2. INTRODUCTION

Multirobot teams can explore an unknown environment faster than a single robot, but a central question is how to coordinate the behavior of multiple robots, especially in an environment whose structure is unknown. In previous research [10], we have demonstrated that frontier-based exploration can be used to explore office buildings and to build occupancy grids that represent the spatial structure of those environments. In this paper, we show how frontier-based exploration can be extended to multiple robots.

Autonomous Agents 98 Minneapolis MN USA
0-89791-983-1/98/5

Our approach is cooperative, decentralized, and robust to individual failures. Information about the world is shared, but each robot maintains its own map, and makes its own decisions about where to navigate. As a result, the robots in the team can continue to explore even if one or more of the members is disabled. We have implemented our system using real mobile robots, and we have tested this system in an indoor office environment.

3. FRONTIER-BASED EXPLORATION

The central question in exploration is: *Given what you know about the world, where should you move to gain as much new information as possible?* Initially, you know nothing except what you can see from where you're standing. You want to build a map that describes as much of the world as possible, and you want to build this map as quickly as possible.

The central idea behind frontier-based exploration is: *To gain the most new information about the world, move to the boundary between open space and uncharted territory.*

Frontiers are regions on the boundary between open space and unexplored space. When a robot moves to a frontier, it can see into unexplored space and add the new information to its map. As a result, the mapped territory expands, pushing back the boundary between the known and the unknown. By moving to successive frontiers, the robot can constantly increase its knowledge of the world. We call this strategy *frontier-based exploration*.

If a robot with a perfect map could navigate to a particular point in space, that point is considered *accessible*. All accessible space is contiguous, since a path must exist from the robot's initial position to every accessible point. Every such path will be at least partially in mapped territory, since the space around the robot's initial location is mapped at the start. Every path that is partially in unknown territory will cross a frontier. When the robot navigates to that frontier, it will incorporate more of the space covered by the path into mapped territory. If the robot does not incorporate the entire path at one time, then a new frontier will always exist further along the path, separating the known and unknown segments and providing a new destination for exploration. In this way, a robot using frontier-based exploration will eventually explore all of the accessible space in the world.

3.1 Evidence Grids and Laser-Limited Sonar

We use evidence grids [7] as our spatial representation. Evidence grids are Cartesian grids containing cells, and each cell stores the probability that the corresponding region in

space is occupied. Initially all of the cells are set to the prior probability of occupancy, which is a rough estimate of the overall probability that any given location will be occupied. Evidence grids have the advantage of being able to fuse information from different types of sensors.

We use sonar range sensors in combination with a planar laser rangefinder to build our robot's evidence grid maps. Sonar sensors are cheap and widely available, but specular reflections often significantly degrade their accuracy. When a sonar pulse hits a flat surface at an oblique angle, it may reflect away from the sensor. As a result, either the sensor detects nothing, or it senses objects that, like reflections in a mirror, appear to be much farther away than the nearest surface.

These reflections could cause difficulties for frontier-based exploration, not only due to inaccuracies in the map, but also because specular reflections often appear as large open areas surrounded by unknown territory. As a result, the robot could waste a great deal of time trying to reach non-existent frontiers.

In order to reduce the effect of specular reflections, we have developed a technique we call *laser-limited sonar*. This technique combines the advantages of sonar and laser range sensors. Sonar sensors are effective at determining that large volumes of space are clear of obstacles, but they suffer from specular reflections. Laser range sensors are effective at accurately determining the positions of obstacles. However, the inexpensive, triangulation-based laser rangefinders commonly found on mobile robots are limited to detecting obstacles within a plane.

The standard evidence grid formulation assumes that each sensor reading is independent of every other sensor reading. In reality, this is not the case, and we take advantage of this. We use a laser rangefinder in combination with the sonar sensors, and if the laser returns a range reading less than the sonar reading, we update the evidence grid as if the sonar had returned the range indicated by the laser, in addition to marking the cells actually returned by the laser as occupied.

As a result, evidence grids constructed using laser-limited sonar have far fewer errors due to specular reflections, but are still able to incorporate obstacles detected by the sonar below (or above) the plane of the laser. It is possible that obstacles undetected by the laser may cause sonar specular reflections which will be added to the map. However, in practice, we have found that laser-limited sonar drastically reduces the number of uncorrected specular reflections from walls and other large obstacles, which tend to be the major sources of errors in evidence grids built using sonar.

Figure 1 compares the evidence grids constructed using raw sonar (Figure 1a) and laser-limited sonar (Figure 1b). Cells representing open space are represented by whitespace. Cells representing occupied space are represented by black circles. Cells representing unknown territory are represented by small dots.

Both of these grids were built by a robot positioned in an alcove at one corner of a large open area. This open area is visible in the lower right portion of both grids. The grid constructed using raw sonar includes three specular reflections from the adjacent walls. Using laser-limited sonar, these reflections are eliminated.

With a narrow field-of-view laser rangefinder, such as the one currently mounted on our robot, it is necessary to sweep the robot's sensors to obtain a complete picture of the robot's surroundings. With wide field-of-view laser rangefinders, such as the ones we plan to install on our robot in the near future, the robot does not need to stop and sweep its sensors. Instead, laser-limited sonar can be used to build evidence grids while the robot remains in motion.

3.2 Frontier Detection

After an evidence grid has been constructed, each cell in the grid is classified by comparing its occupancy probability to the initial (prior) probability assigned to all cells. A value of 0.5 was used in all of the experiments described in this paper. In general, this does not need to be an accurate estimate of the actual amount of occupied space. A prior probability of 0.5 works fine in environments where only a small fraction of the total space is occupied, as well as in far more cluttered environments.

Each cell is placed into one of three classes:

- **open:** occupancy probability < prior probability
- **unknown:** occupancy probability = prior probability
- **occupied:** occupancy probability > prior probability

The only time the occupancy probability will be exactly equal to the prior probability is when it has never been sensed. Different thresholds could be used to determine which cells are open or occupied, but the simple thresholds above work well in practice.

A process analogous to edge detection and region extraction in computer vision is used to find the boundaries between open space and unknown space. Any open cell adjacent to an unknown cell is labeled a frontier edge cell. Adjacent edge cells are grouped into frontier regions. Any frontier region above a certain minimum size (roughly the size of the robot) is considered a frontier.

Figure 2a shows an evidence grid built by a real robot in a hallway adjacent to two open doors. Figure 2b shows the frontier edge segments detected in the grid. Figure 2c shows the regions that are larger than the minimum frontier size. The centroid of each region is marked by crosshairs. Frontier 0 and frontier 1 correspond to open doorways, while frontier 2 is the unexplored hallway.

3.3 Navigating to Frontiers

Once frontiers have been detected within a particular evidence grid, the robot attempts to navigate to the nearest accessible, unvisited frontier. The path planner uses a depth-first search on the evidence grid, starting at the robot's current cell and attempting to take the shortest obstacle-free

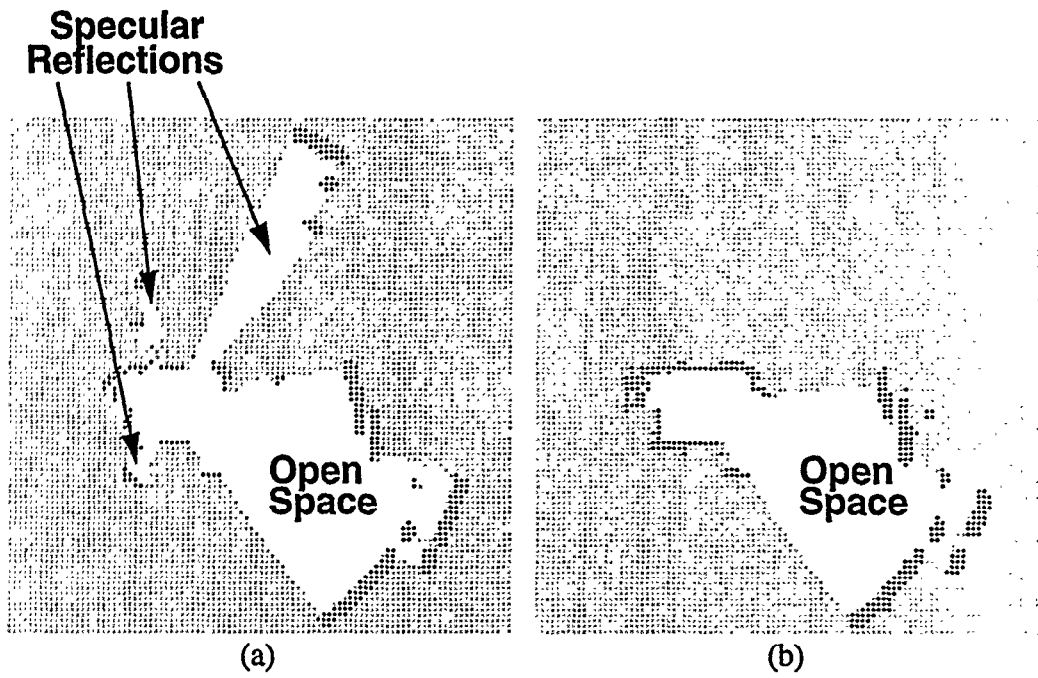


Figure 1. Comparison of evidence grids built using (a) raw sonar and (b) laser-limited sonar

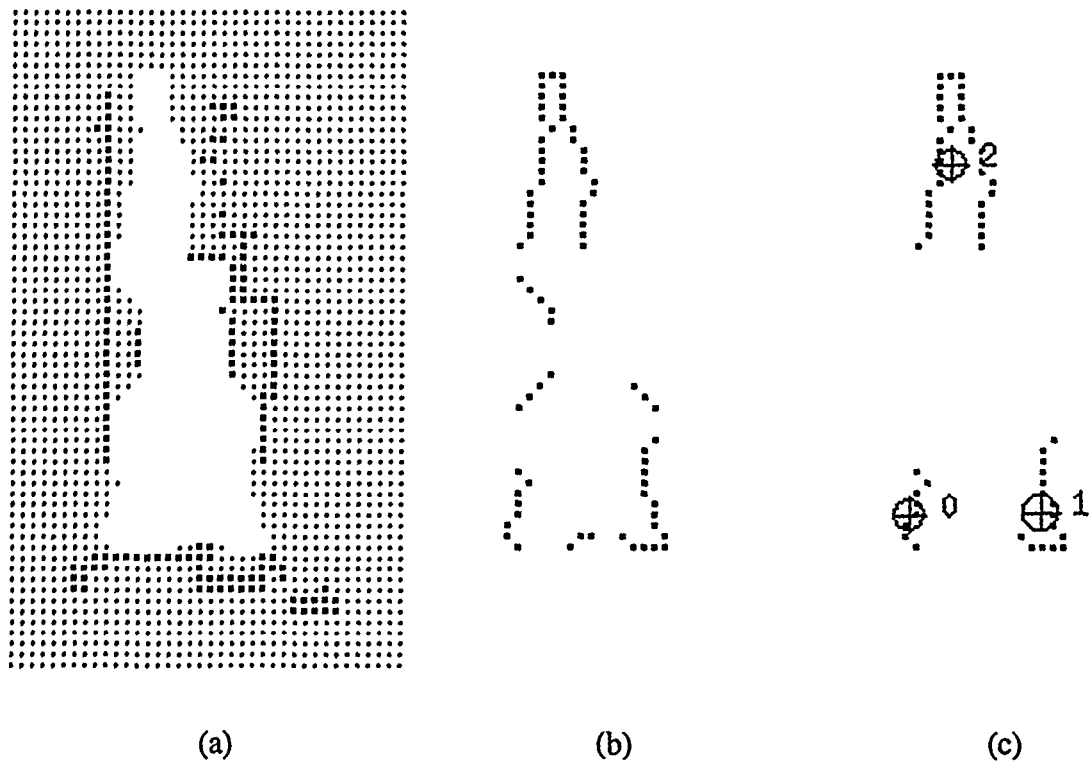


Figure 2. Frontier detection: (a) evidence grid, (b) frontier edge segments, (c) frontier regions

path to the cell containing the goal location. For path planning, the grid is treated as a graph with a node at the center of each open cell, and each node is connected to the nodes of all adjacent open cells.

While the robot moves toward its destination, reactive obstacle avoidance behaviors prevent collisions with any obstacles not present while the evidence grid was constructed. These behaviors allow the robot to steer around these obstacles and, as long as the world has not changed too drastically, return to follow its path to the destination.

When the robot reaches its destination, that location is added to the list of previously visited frontiers. The robot performs a 360 degree sensor sweep using laser-limited sonar and adds the new information to the evidence grid. Then the robot detects frontiers present in the updated grid and attempts to navigate to the nearest accessible, unvisited frontier.

If the robot is unable to make progress toward its destination for a certain amount of time, then the robot will determine that the destination is inaccessible, and its location will be added to the list of inaccessible frontiers. The robot will then conduct another sensor sweep, update the evidence grid, and attempt to navigate to the closest remaining accessible, unvisited frontier.

4. MULTIROBOT EXPLORATION

Figure 3 provides an overview of our multirobot exploration strategy. Each robot has its own global evidence grid that represents its knowledge about the environment. Whenever a robot arrives at a new frontier, it sweeps its sensors and constructs a local evidence grid representing its current surroundings. This local grid is integrated with the robot's global grid, and also broadcast to all of the other robots.

Each robot stores the local grids received from other robots. When a robot arrives at a new frontier, it integrates these local grids with its global map, along with the new local grid it constructs at the frontier.

Two evidence grids can be integrated in a straightforward way. A log odds representation is used, as described by Moravec [6]. Using this representation, independent probabilities can be combined using additions rather than multiplications. The log odds probabilities of each pair of corresponding cells are added and the sum is stored in the corresponding cell of a new grid. These probabilities are normalized so that the log prior probability is equal to zero, so that cells with no information in either grid remain equal to the prior probability in the new grid. Since each cell requires only a single addition, this integration can be done in real-time. Two grids with 65,536 cells (256 x 256) can be integrated in 1-2 seconds on a Sparcstation 20.

This approach has the advantage of being both cooperative and decentralized. All of the information obtained by any robot is available to each robot. This allows robots to use the data from other robots to determine where to navigate.

Based on this information, a robot can determine which areas have already been explored by other robots, and then choose to explore another region. A robot can also discover that a frontier detected by another robot is nearby, and decide to investigate.

While information is shared, control is independent, and this allows the team to be robust to failures of individual robots. If a robot becomes disabled, it will no longer report new information, but the remaining robots will still continue to explore the entire environment. Since this system is completely asynchronous, robots are never waiting for other robots, so the loss of one or more robots will never bring the team to a halt.

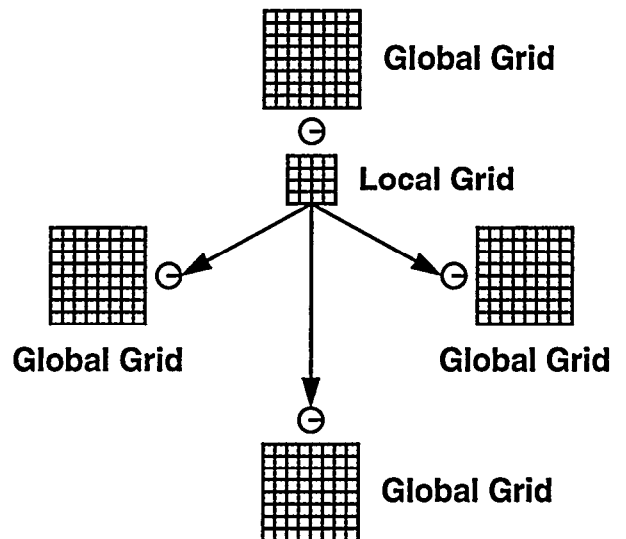


Figure 3. Multirobot exploration overview

This approach also has a few limitations. Since navigation is independent, it may not be optimally efficient. Robots may waste time by navigating to the same frontier. Reactive behaviors will prevent collisions, but it is still possible for one robot to block another, or in rare occasions, for two robots to block each other. In these cases, however, any blocked robots will continue to explore other frontiers, and only the blocked robot will mark its destination frontier as inaccessible, so other robots may still explore that frontier. Also, when each robot scans the room, it will incorporate the other robots as obstacles within its map. However, as these robots move to new positions, their old locations will be marked as clear space when those locations are rescanned in subsequent sensor sweeps.

In future research, we plan to study the costs and benefits of adding additional communication and coordination between robots, but for our initial study, we wanted to test the

performance of a highly decentralized strategy with minimal coordination.

5. EXPERIMENTS

Frontier-based exploration was implemented on two Nomad 200 mobile robots. Each robot was equipped with a planar laser rangefinder, sixteen sonar sensors, and sixteen infrared sensors. A Pentium processor onboard each robot handles low-level sensor processing and motor control. For each robot, a process for frontier-based exploration runs on a separate Sparcstation 20. These processes communicate with each other over an ethernet, and with the robots over a radio ethernet. All communication between processes is implemented using Internet domain sockets, except for the transmission of evidence grids, which are passed as files using NFS.

Experiments were performed in an indoor environment, roughly 50 feet by 50 feet, consisting of three rooms. The robots were started in different locations, and allowed to explore using the multirobot exploration strategy described above,

Each robot was provided with an initial estimate of its position, and dead reckoning was used to update this position as the robot moved through the environment.

The experiments described in this paper use dead reckoning alone for position estimation. However, since dead reckoning accumulates error over time, we have also developed a technique for continuous localization that matches recent perceptions against a stored evidence grid to determine the robot's current position [8]. We have integrated continuous localization with frontier-based exploration in ARIEL [11], and we plan to use ARIEL for multirobot exploration in future research.

Figure 4 shows the sequence of global evidence grids constructed from a typical trial as the robots explore the office environment. The Xs indicate the original positions of the robots, and the crosshairs indicate the positions of the robots at the time the global grid was last updated. Each of the robot has its own global grid, so this figure shows the sequence from only one of the robots. The grids from the other robot are similar.

In Figure 4a, the robots start in separate rooms at opposite ends of the environment. Each robot can detect frontiers in both rooms, but the frontiers in the other robot's room are inaccessible, since reaching those frontiers would require navigating through unknown space. So, at this point, each robot will navigate to the nearest accessible frontier within its own room.

In Figure 4b, the robots meet in the central room. The open, explored space within the global grid is now contiguous, so each of the robots can navigate to any of the frontiers within the environment. Since an open doorway leads to the one unexplored room, both robots navigate into this room, and explore it together. The final global grid constructed by the robots is shown in Figure 4c.

6. RELATED WORK

Most of the exploration systems implemented on real robots have used individual robots. For example, Mataric [5] has developed Toto, a robot that combines reactive exploration, using wall-following and obstacle-avoidance, with a simple topological path planner. However, the reactive nature of Toto's exploration limits its ability to map environments where wall-following is insufficient to explore the complex structure of the world.

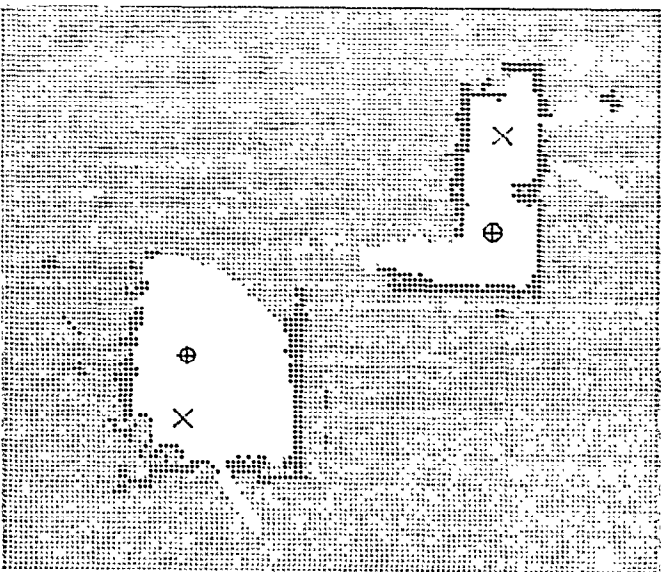
Lee [3] has implemented Kuipers Spatial Semantic Hierarchy [2] on a real robot. This approach assumes that all walls are parallel or perpendicular to each other, and this system has only been tested in a simple environment consisting of a three corridors constructed from cardboard barriers.

Thrun and Bücken [9] have developed an exploration system that builds a spatial representation that combines evidence grids with a topological map. This system has been able to explore the network of hallways within a large building. While this approach works well within the hallway domain, it also assumes that all walls are either parallel or perpendicular to each other, and that they do not deviate more than 15 degrees from these orientations. An implicit assumption is that walls are observable and not obstructed by obstacles. These assumptions make this approach unsuitable for rooms cluttered with obstacles that may be in arbitrary orientations.

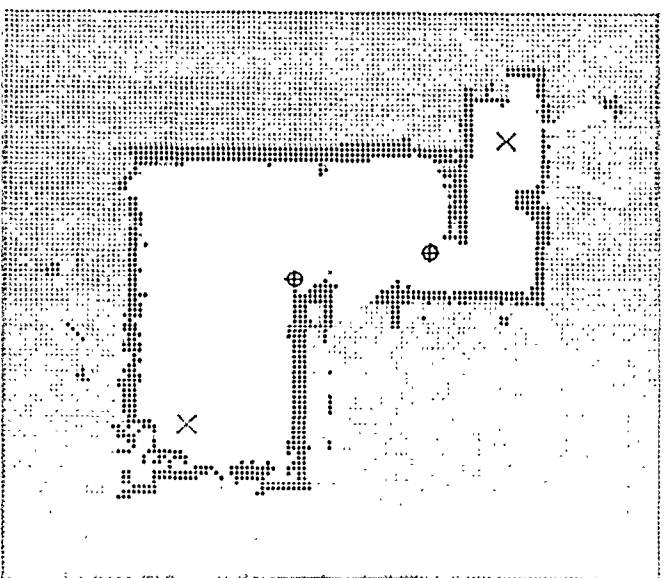
Duckett and Nehmzow [1] have developed a mobile robot system that combines exploration and localization. This system uses a simple reactive wall-following strategy for exploration. For localization, this system uses a self-organizing neural network trained using ART. This network learns to classify regions in space based on input from infrared range sensors. Each of these regions is associated with the corresponding robot position as determined by dead reckoning. By looking at sensor inputs and motor commands over time, this system can reduce the amount of perceptual aliasing that occurs. This robot has only been tested in a small enclosed area (6 meters by 4 meters), so it is unclear whether this approach will scale to larger, more complex, environments.

A technique for multirobot exploration has been developed by López de Mántaras, Amat, Esteva, López, and Sierra [4]. They use simple robots equipped with IR proximity sensors that wander randomly while avoiding obstacles. Any walls the robots detect are modeled as line segments that intersect at right angles. Their approach has the advantage of requiring only simple, inexpensive hardware, but the disadvantage of being limited to environments consisting only of walls that intersect at right angles, and without other non-orthogonal obstacles. In addition, due to the random nature of their robots' exploration, regions of the environment are often left unexplored.

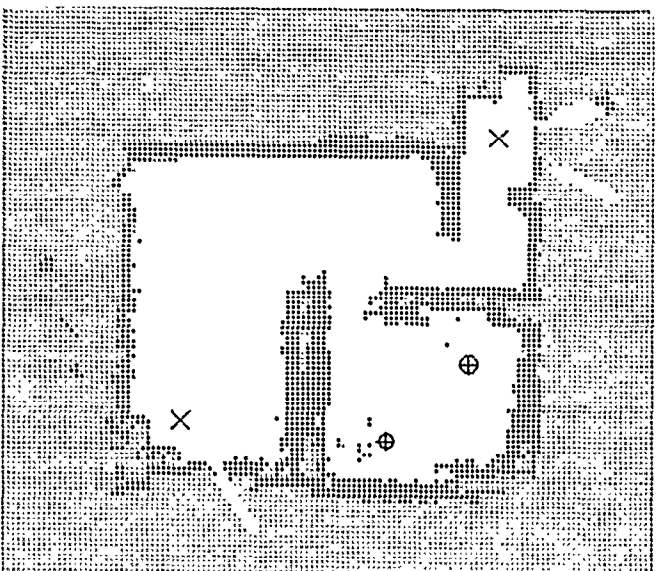
Our approach is unique in a number of ways. Unlike systems that can only handle parallel or perpendicular walls, our grid-based approach can be used in unstructured



(a)



(b)



(c)

Figure 4. Multirobot exploration sequence

environments with obstacles in arbitrary positions. Unlike purely reactive approaches, our system uses frontier-based exploration to direct robots to the areas that are likely to provide the most new information about the world. Our coordination strategy allows multiple robots to share information and explore cooperatively, without any central control, thus enabling the robust and efficient exploration of unknown environments.

7. CONCLUSIONS AND FUTURE WORK

Frontier-based exploration provides an efficient way to explore environments where walls and obstacles may be in arbitrary orientations, and where wall-following or random wandering is unlikely to produce a complete map. In this paper, we have shown how frontier-based exploration can be extended for use with multiple robot teams.

In our approach, robots maintain their own global maps and make independent decisions about where to explore. Robots share information about their perceptions whenever they arrive at a new frontier, and they integrate the information from other robots into their own global map. In this way, robots cooperate and use the information from other robots to guide their own exploration. At the same time, this approach is robust to the loss of individual robots, and does not require any explicit coordination or synchronization between robots. We have implemented this system using real robots, and demonstrated that this approach can be used to map an indoor office environment.

Previously, we have demonstrated that continuous localization using evidence grids can be integrated with frontier-based exploration [11]. This enables robots to maintain an accurate position estimate even when exploring unknown territory. We plan to integrate this capability within multirobot exploration teams, where the information from all of the robots can be used to localize each robot. We also plan to investigate the tradeoffs associated with more explicit coordination between robots, and to conduct quantitative experiments comparing the performance of different multirobot exploration strategies.

8. ACKNOWLEDGMENTS

This work is supported by the Office of Naval Research.

9. REFERENCES

- [1] Duckett, T. and Nehmzow, U. Experiments in evidence-based localisation for a mobile robot. In *Proceedings of the AISB Workshop on Spatial Reasoning in Mobile Robots and Animals* (Manchester, UK, 1997).
- [2] Kuipers, B. and Byun, Y. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47-63, 1991.
- [3] Lee, W. Spatial Semantic Hierarchy for a Physical Robot. Ph.D. Thesis, Department of Computer Sciences, The University of Texas at Austin, 1996.
- [4] López de Mántaras, R., Amat, J., Esteva, F., López, M., and Sierra, C. Generation of unknown environments maps by cooperative low-cost robots. In *Proceedings of Agents '97*, (Marina del Rey, CA, February 1997).
- [5] Mataric, M. Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation*, 8(3):304-312, June 1992.
- [6] Moravec, H. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, 9(2):61-74, 1988.
- [7] Moravec, H. and Elfes, A. High resolution maps from wide angle sonar. In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, (St. Louis, MO, 1985), pp. 116-121.
- [8] Schultz, A. and Adams, W. Continuous localization using evidence grids. To appear in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, (Leuven, Belgium), May 1998.
- [9] Thrun, S. and Bücken, A. Integrating grid-based and topological maps for mobile robot navigation. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, (Portland, OR, August 1996), pp. 944-950.
- [10] Yamauchi, B. A frontier-based approach for autonomous exploration. In *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, (Monterey, CA, July 1997), pp. 146-151.
- [11] Yamauchi, B., Schultz, A., and Adams, W. Mobile robot exploration and map-building using continuous localization. To appear in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, (Leuven, Belgium, May 1998).