

CENTIBOTS

Large Scale Robot Teams

Kurt Konolige, Charles Ortiz, Regis Vincent, Andrew Agno, Michael Eriksen, Benson Limketkai, Mark Lewis, Linda Briesemeister, Enrique Ruspini

*Artificial Intelligence Center
SRI International
Menlo Park, CA*

Dieter Fox, Jonathan Ko, Benjamin Stewart
*Department of Computer Science and Engineering
University of Washington
Seattle, WA*

Leonidas Guibas
*Computer Science Department
Stanford University
Stanford, CA*

Abstract: As part of the DARPA Software for Distributed Robotics Program, SRI International, Stanford University, the University of Washington, and ActivMedia Robotics are designing and implementing a computational framework for the coordination of large robot teams, consisting of at least 100 small, resource limited mobile robots (CentiBOTS), on an indoor reconnaissance task.

Key words: Multirobot mapping and surveillance

1. 101 ROBOTS

The requirements of autonomy, size, low power, and limited computation available on small robots place extraordinary constraints on communication and coordination of robot platoons. Individual robots, executing a limited set of behavioral programs, must cooperate to produce a global behavior that satisfies the goals of the mission. Current robot architectures rely on large, power-hungry subsystems for mobility, communication, and control; furthermore, they do not address the problem of coordinating large numbers of robots.

In the CentiBOTS project, we envisage a system of a large number of mobile robots able to effectively explore, map, and surveil the interior of a building. The name of the project comes from our goal of having 100 robots performing the task. Each robot is minimally capable of self-localization in a map and communication with nearby robots. Some are more specialized for tasks such as mapping or tracking people, with different degrees of performance (and consequent computational and sensor abilities). The robots must function as an effective *team* for accomplishing an extended surveillance mission, with appropriate resilience in the face of environmental challenges. These include a dynamic environment with unexpected events, uncertainty that is inherent in sensing and acting on the physical world, and the need for communication among the robots and/or with a central base station under conditions in which power or bandwidth may be limited, connectivity may be intermittent, and stealth may be a factor.

In the project scenario, the CentiBOTS are deployed as an advanced surveillance team for urban missions. A first set of mapping-capable CentiBOTS will survey an area of interest: build and share a distributed map as well as highlight hazards, humans, and hiding places. They will then combine with a second wave of tracking robots that deploy in an optimal way into the area, configuring themselves to effectively sense intruders and share the information among themselves and a command center. As a large team, the CentiBOTS will be able to reconnoiter a set of buildings faster, more reliably, and more comprehensively than an individual or small set of robots. For example, the team can dynamically form subteams to perform tasks that cannot be done by individual robots. Examples of such tasks might be to measure the range to a distant object, or to use other robots as markers in the building for localization or communication relays. In addition, the team can automatically reconfigure itself to handle contingencies such as disabled CentiBOTS or changing lighting conditions.

The robot teams will collaboratively perform tasks with minimal supervision in dynamic environments. Our major contribution is a distributed robot architecture in which collective behavior is uniquely

adaptive, *fault tolerant*, and *capable*, incorporating the following innovative elements.

- ?? **A collaborative, multi-level architecture, adaptive to new tasks and team organizations, and scalable to very large teams** based on SRI's proven Saphira robot control system and the Distributed Dispatcher Manager (DDM) hierarchical agent framework developed by SRI for the DARPA ANTs program. We are incorporating principles of *collaboration*, derived from our work on structured and dynamic negotiation for the DARPA ANTs program, so that CentiBOTS will be capable of re-organization and re-tasking in response to resource and problem changes in the environment.
- ?? **Optimal distributed map-building and deployment of CentiBOTS for tracking** based on novel distributed spatial reasoning techniques. We extend single-robot probabilistic methods such as Markov sampling and relational dynamic Bayes nets to the multi-robot, distributed case.
- ?? **Large-scale, fault-tolerant communication** building on SRI-developed mobile ad-hoc network protocols that have already been successfully demonstrated on smaller robot teams. The protocol supports the mission specific communication tasks efficiently. Additionally, it alerts the application about decreasing fault tolerance when links break.
- ?? **Robot team interface and monitoring** that provides both robot level attribute-of-interest updating and tracking as well as task and team level goal tracking.
- ?? **Analyzable and predictable behavior** Through systematic experiments with well-defined evaluation metrics, in both the SRI Augmented Reality Robot Simulator as well as demonstrations and experiments, we will show increasing capability of the software solution on a collection of at least 100 COTS mobile robot platforms.

The project started in July 2002. By the time of the workshop, we will have completed our first demo at 6 months, and have results to report on the complete integration of distributed mapping and surveillance, ad-hoc network communication, and team formation and execution. As of the current writing (Dec 1, 2002), we have major portions of the system in place. Recent results are posted to the website www.ai.sri.com/centibots.

2. ROBOTS AND INFRASTRUCTURE

Given the nature of the mission and the large number of robots, a critical component of the project is the robot platform, sensor suite, and software programming base. The robots must have enough capability to be able to perform mapping, localization, communication and tracking functions, while at the same time they are subject to the conflicting demands of low power, simplicity, modest computational load, and small size. In contrast to swarm-based robotics, where individual robots have almost no ability for independent action, we equip each robot with the capability to localize and perform some tracking or mapping function. Pioneer and AmigoBot robots from ActivMedia, and robot software from SRI and ActivMedia, form the core of the infrastructure. The basic robot types are as follows.

Robot class	#	Computer	Sensors	Capabilities
Pioneer II DX/AT	6	PIII EBX	LRF, sonars	map, detect, track
AmigoBot	10	VIA Epia	Stereo vision, sonars	detect, track (form, range)
	60	VIA Epia	Mono vision, sonars	detect, track (color, bearing)
	14	VIA Epia	Doppler radar	track (motion, no light)

All robots are equipped with 802.11b wireless links, and are capable of localizing in a map, through the use of sonars or LRF sensors. The larger (30 lb), more capable Pioneer II DX/AT robots have a powerful computer, and a laser range finder for mapping and people-tracking. These robots perform the distributed mapping task, and assist in detection and tracking of people.

The smaller (6 lb) AmigoBots are used as detection and tracking robots for finding the targets, and for surveillance of intruders. They have low-power VIA Epia processors, and are able to survive for long periods of time in low-power mode (6-12 hours), while still performing communication and surveillance on wakeup. The mix of sensors on the AmigoBots reflects differing environmental conditions, capabilities, and power requirements.

The base software for the robots, Saphira/Aria, is a joint project of SRI and ActivMedia. Saphira/ARIA is a modern 3-level (behaviors, sequencing, strategy) robot control architecture with an extensive library of modular capabilities, including probabilistic localization, map-building, optimal realtime path planning, and visual tracking, developed by members of the project team. We have extended the core architecture to be network-aware, so that each robot becomes a member of an ad-hoc mobile network (Figure 1).

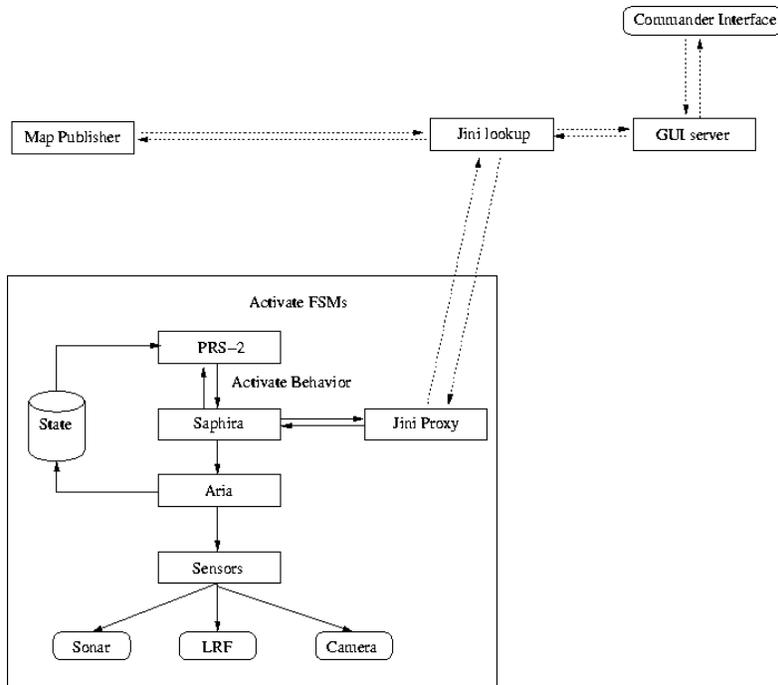


Figure 1. Robot control architecture, showing an individual robot, and its connections to the robot net.

3. COMMUNICATION ARCHITECTURE

CentiBOT teams must operate with little or no infrastructure and present a challenging scenario for information operations. Networks are formed in an **ad-hoc** fashion, and information exchanges occur via the wireless networking equipment carried by the individual CentiBOTS. While the CentiBOTS are executing their mapping and search mission, fluctuations in the network topology occur when an individual moves or when wireless transmissions are blocked by building features, distance, or interference from other RF transmissions.

In spite of such dynamically changing conditions, the team's CentiBOTS must maintain close communication with one another. We therefore anticipate a requirement for self-configuring, self-sustaining dynamic networks coupled with a location-independent flexible addressing architecture for effective network communication.

The CentiBOT teams are highly collaborative in nature with a requirement for time-critical communication. However, the transmission range of each node is limited in order to preserve its battery power. Hence, the CentiBOT team is organized into a **Mobile Ad-hoc Network (MANET)**, wherein messages are exchanged directly between members of the team or may be forwarded via other members to extend the range. Since communication bandwidth is a scarce resource in a MANET, it is important that the routing protocol be efficient in terms of overhead.

3.1 The TBRPF Routing Protocol

SRI has developed a protocol called **Topology Broadcast based on Reverse-Path Forwarding (TBRPF)** [Bellur and Ogier 1999, Ogier et. al. 2002] to manage the network multihop routing while the topology is changing. TBRPF is an efficient proactive, link-state routing protocol designed for mobile ad-hoc networks, which provides hop-by-hop routing along shortest paths to each destination. Each node running TBRPF computes a source tree (providing paths to all reachable nodes) based on partial topology information stored in its topology table, using a modification of Dijkstra's algorithm. To minimize overhead, each node reports only 'part' of its source tree to neighbors. This is in contrast to other protocols in which each node reports its 'entire' source tree to neighbors.

TBRPF uses a combination of periodic and differential updates to keep all neighbors informed of the reportable part of its source tree. Each node also has the option to report additional topology information (up to the full topology), to provide improved robustness in highly mobile networks. TBRPF is extremely agile in that a change in the up or down status of links is quickly detected, and alternate routes are immediately computed. The proof of correctness and pseudo-code for TBRPF as well as examples illustrating its operation can be found in [Bellur and Ogier 1999, Ogier et. al. 2002].

TBRPF operates transparently on each node of the ad hoc network, providing a standard IP stack interface to network-based applications. The robots and command center send and receive messages as if being connected through a common local area network. We display current topology information at individual nodes and an aggregated picture at the command center (Figure 2). In future implementations of CentiBOTS, we plan to support multicast messages suitable for more efficient intra-team communication.

The network interfaces are commercially available IEEE 802.11b PCCards and USB network interfaces operating in the 2.4 GHz frequency band, using the Direct Sequence Spread Spectrum (DSSS) modulation, and provides up to 11 Mbps data transfer rates with a maximum range of approximately 1000 m line of sight. We configure the cards to use **Ad-hoc** mode, rather than to be dependent on fixed infrastructure elements.

Aiming at deployment of as many as 100 nodes, we will operate one of the largest mobile ad hoc networks known today. Yarvis et al. [Yarvis et al. 2002] have reported an ad hoc sensor network for interactive voting applications with configurations of 24, 48, and 91 nodes. However, they placed and fixed the nodes over a rectangular grid, and their experiment lasted for one hour. In our CentiBOTS project, all robots are mobile and their mission tasks take several hours.

3.2 Distributed Directory Service

For higher level reasoning, the robots need to obtain the current status of other robots on the mission. We are developing a distributed directory service to provide robots with such information. Given the scarce bandwidth and the inherently unpredictable network structure, the distributed directory service must carefully adjust the amount of messages that keep status information up to date at the expense of getting possibly stale data.

The distributed directory service provides a query interface through which a robot can perform various searches for other robots and computers, their locations and their assets. Each robot maintains a local table with all information it gathers about the current state of the network. In adjustable time intervals, the robot updates its local copy of the table with entries kept in a network-wide table, and also submits its own current information to the network table.

We chose the JavaSpaces™ technology to implement the distributed directory service. JavaSpaces is a Jini™ service that provides a high-level

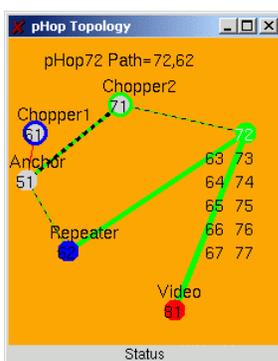


Figure 2. Dynamic topology display screen shows current nodes and links.

tool for creating collaborative and distributed applications in Java. The JavaSpaces model is different from techniques like message passing and remote method invocation. A space is a shared, network-accessible repository for objects. Processes use the repository as a persistent object storage and exchange mechanism; instead of communicating directly, they coordinate by exchanging objects through spaces.

4. MULTIROBOT MAPPING AND LOCALIZATION

Coordinated exploration of an unknown environment is one of the most fundamental problems in multi-robot coordination. We propose a novel, distributed approach that addresses this problem in its most general way. Key features of our approach are the consideration of limited communication between robots, no assumptions about relative start locations of the robots, and dynamic assignment of processing tasks. We apply efficient, statistical methods to determine hypotheses for the relative locations of robots. To achieve maximal robustness, these hypotheses are verified before maps are merged. Once robots know their relative locations, they form *exploration clusters* so as to coordinate their actions. Furthermore, our approach dynamically assigns processing tasks and roles to robots, thereby avoiding the dependency on a central server.

4.1 Communication and Coordination Architecture

Our distributed approach to mapping and exploration is enabled by pair-wise relations between robots. Each pair of robots can have four different types of interactions: none, hypothesis generation, hypothesis verification, and coordination. At each point in time, the state of the multi-robot system can be summarized by a graph structure where the nodes are individual robots and edges represent the current interaction between two robots (see Figure 3). We will now briefly discuss the different types of interactions.

1. **No interaction:** The robots are not within communication range (no arc between nodes in Figure 3).
2. **Hypothesis generation** (dotted edges): The robots can communicate but don't know their relative locations. In this stage, one of the two robots receives sensor data from the other robot and estimates their relative location using its own map. Which of the two robots adopts the estimation role depends on available computational resources.
3. **Hypothesis verification** (dashed edges): Robots can communicate and verify a location hypothesis determined in the hypothesis generation phase. This is done by moving to a point at which the robots try to meet. If the robots don't meet at the expected location, the hypothesis is rejected and

they continue with the hypothesis generation phase. Otherwise, the robots can establish their relative positions, combine their maps, and coordinate their exploration efforts.

4. **Coordinated exploration** (solid edges): Once the robots determined their relative locations, they can share their maps and perform coordinated exploration. A nice feature of this interaction type is transitivity, *i.e.* if robot i and j can share their maps, and robot j and k can share their maps, then all three robots can build a combined map. Hence, robots in this interaction mode form *exploration clusters* in which they can coordinate their actions (indicated by the gray areas in Figure 3). Each cluster determines one robot responsible for data combination and robot coordination (dark nodes). All information is frequently spread throughout the cluster, and direct communication between all robot pairs within the cluster is not required.

4.2 Technical Approach

To implement the individual parts of our architecture, we rely on existing, well established techniques whenever possible. The key components are:

?? **Individual mapping and exploration** [Yamauchi 1998, Thrun et al. 2000, Gutmann and Konolige 2000].

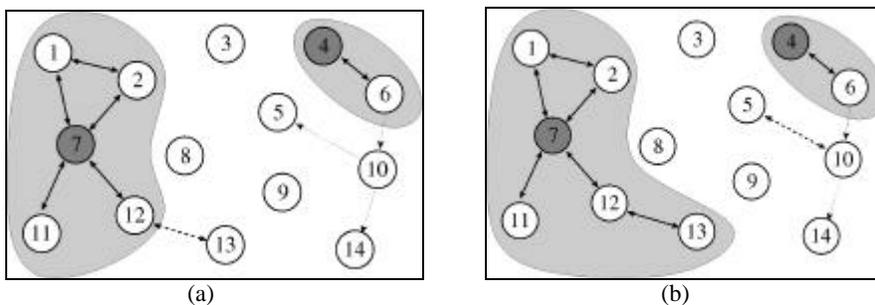


Figure 3. Dynamic communication / interaction graph at two points in time. Note that this graph illustrates different interactions between robots, *not* spatial relations. Robots are shown as circles. Solid edges indicate coordinated exploration of robots; dashed lines indicate that two robots currently navigate to a meeting point so as to verify a hypothesis for their relative positions; and dotted lines show communication between robots without valid location hypotheses. (a) Robots 1, 2, 7, 11, and 12 already established their relative locations and coordinate exploration. For this exploration cluster, robot 7 was chosen to perform data combination and exploration coordination. Robots 12 and 13 do not yet know their relative locations. They are currently moving to a meeting point so as to verify a location hypothesis. Robot 10 can communicate with robots 5, 6, and 14, but no good location hypothesis has been generated so far. (b) Robots 12 and 13 moved to the meeting point and detected each other. As a result, robot 13 is integrated into the exploration cluster. Robot 5 determined a hypothesis for robot 10's relative location. Robot 10 accepted the meeting point and they both move to this location.

?? **Coordinated mapping and exploration** [Burgard et al. 2000, Simmons et al. 2000].

?? **Hypothesis generation and verification** [Fox et al. 2000b, Fox et al. 200b]

A companion poster presents experimental results and further details on the underlying Bayesian estimation techniques.

4.3 Current Mapping Progress

At the moment, we have successfully integrated maps from 5 separate robot mapping runs into a large scale map (Figure 5). The map data was collected offline, by running a mapping robot on 5 different trajectories, and then mixing all the scans as if 5 robots were running simultaneously.

5. DISTRIBUTED COORDINATION AND PLANNING

We briefly discuss the team-level organization of CentiBOTS. This level will become increasingly important as we perform more ambitious

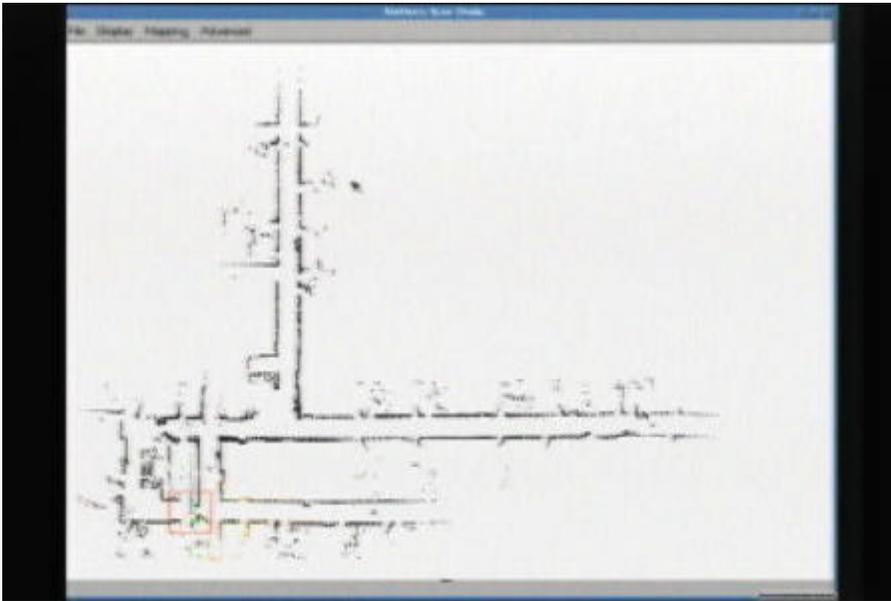


Figure 5. 5 robots mapping from a common breach. The map is about 2/3 constructed at this point, with 4 areas of frontier and 1 area of infill exploration. Map is about 60mx80m.

experiments.

The team level is responsible for decisions involving societal aspects of the robot group, such as negotiations with other robots on the division of responsibility or the allocation of resources. The team level is also designed to respond to changes in the environment that could impact the performance of the group (e.g., a robot that suddenly detects an intruder entering a team member's sector should realize that if that team member is already tracking another intruder, it will need help).

Teams are organized hierarchically in order to manage the complexity of the team activity. Within a team, any robot is able to take on either a leadership role or a supporting role in the collaborative activities of the team. Our approach is based on a system called the distributed dispatcher manager (DDM) [Yadgar et al. 2002]. DDM organizes teams hierarchically, which reduces the degree of communication necessary between agents.

DDM is loosely based on a metaphor for task distribution modeled on the activity of a "taxi dispatcher" who assigns taxis to incoming calls (tasks). In DDM, a team leader or dispatch agent distributes a problem and its solution among the robot team. The team is organized in a hierarchical fashion. In DDM, the dispatcher agent is called a zone coalition leader. The zone coalition leader can either be initially assigned or can be dynamically chosen through some voting process. The lower level of the hierarchy consists of individual robots that are grouped together according to particular sectors or capabilities. Each group also has a leader. These group leaders are also grouped according to their sector. Each such group of leaders is associated

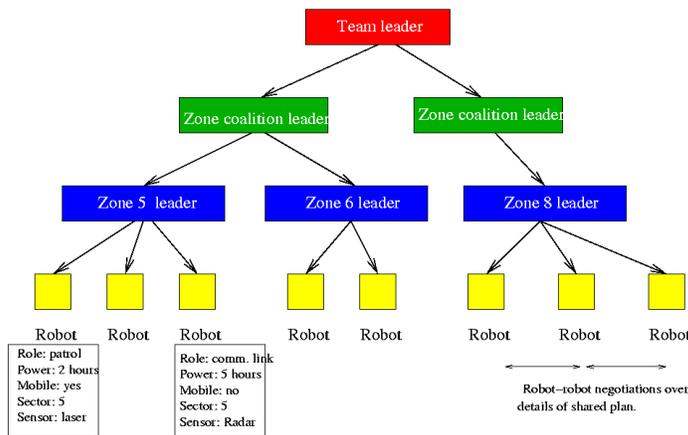


Figure 7. The Distributed Dispatch Manager (DDM).

with a zone group leader. Figure 7 illustrates this structure. Zone group leaders are also grouped according to associated sectors with their own zone group leader. Individual robots are mobile. They may therefore change their group when changing their area. The zone group leaders execute a load balancing algorithm so that if too many tasks appear in a particular sector, agents from other sectors can be diverted from other sectors through negotiations with other zone leaders [Yadgar et al. 2002]. DDM is also *fault tolerant*. If a zone leader becomes disabled, DDM reorganizes itself so that another agent will take its role.

The role of a zone coalition leader is to distribute robots of the appropriate *general* capability to sectors where they are needed. Just as with the taxi dispatcher, it need not concern itself at this high level of the details of *how* the agents will perform their tasks.

REFERENCES

- B. Bellur, and R.G. Ogier. 1999. "A Reliable, Efficient Topology Broadcast Protocol for Dynamic Networks," *Proc. IEEE INFOCOMM '99*, pp. 178–186.
- R.G. Ogier, M.G. Lewis, F.L. Templin, and B. Bellur. 2002. "Topology Broadcast based on Reverse Path Forwarding (TBRPF)," draft-ietf-manet-tbrpf-06.txt, (November) (work in progress).
- M.D. Yarvis, W.S. Conner, L. Krishnamurthy, A. Mainwaring, J. Chhabra, and B. Elliott. 2002. "Real-world experiences with an interactive ad hoc sensor network," *Proc. International Workshop on Ad Hoc Networking (IWAHN)*, pp. 143–151 (August).
- O. Yadgar, S. Kraus and Charles Ortiz. Hierarchical organizations for realtime large-scale task and team environments. *AAMAS 2002*.
- W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. *ICRA 2000*.
- D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3):325-344, 2000.
- D. Fox, S. Thrun, F. Dellaert, and W. Burgard. Particle filters for mobile robot localization. In Doucet et al., eds., *Sequential Monte Carlo in Practice*. Springer-Verlag, 2001.
- J.S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. *CIRA 2000*.
- R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. *AAAI 2000*.
- S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and {3D} mapping. *ICRA 2000*.
- B. Yamauchi. Frontier-based exploration using multiple robots. *Proc. of the Second International Conference on Autonomous Agents*, 1998.