In *Proceedings, IEEE International Conference on Robotics and Automation (ICRA-2002), Washington DC, May 11-15, 2002.*

Huey, Dewey, Louie, and GUI — Commanding Robot Formations

Jakob Fredslund (chili@daimi.au.dk)*

Maja J Matarić (mataric@cs.usc.edu)†

Abstract

We present a graphical user interface (GUI) for easy, intuitive control of a group of mobile robots moving in formation, and we give a short description of the general, local, distributed formation algorithm underlying the GUI. In spite of the high level of abstraction of the GUI, it captures the full functionality of the algorithm. The GUI has been tested with robots running in a simulator, and the formation algorithm itself has also been tested with real robots. We show results from experiments in both scenarios.

1 Introduction

Perhaps future Superbowl halftime shows will be performed by robots. Imagine being the halftime show director in charge of hundreds of mobile performer robots. You know what looks good, you have an idea of what physical shapes the performers should form. But you are a choreographer, not a roboticist; you have no detailed knowledge of how to communicate with robots in their language, and it is already difficult to express shapes and formations linguistically. How could you best explain to them what you want?

In this paper we present a graphical user interface (GUI) that allows easy, real-time control of formations of mobile robots, and we briefly describe the general, local, distributed formation algorithm underlying the GUI. This algorithm has been tested both in simulation and on real robots. The GUI and a robot simulator are run simultaneously; the user manipulates a chain of robot nodes in the GUI using the mouse to form a geometric shape, and then downloads it to the robots running in the simulator in a window next to the GUI. Immediately, the robots start rearranging themselves and quickly establish the new formation. A switch between any two formations is possible, and the GUI allows the user to shape precisely the kinds of formations that the algorithm captures. The same GUI could easily be used with real robots, e.g., through the World Wide Web.

When only one mobile robot is involved, an abstract, natural, and intuitive interface enables non-expert humans of controlling the robot, whereas a lower, more programming-like level of abstraction provides for greater sophistication, precision, and expressiveness, but is far less understandable and approachable. This is further emphasized when a whole group of mobile robots is controlled: an abstract, natural means of interaction that works well with only one robot might prove inadequate when controlling more than one at the same time. A different approach, perhaps an even more high-level intuitive interface, might be needed when taking the leap from one to multiple robots.

Our control interface is tailored to one specific task: the control of robot formations. It is purely graphical; there is no programming, only click-and-drag. Our GUI cannot be used to control box pushing or soda-can collecting robots, but it is immediately understandable and requires no prerequisites of the user other than some familiarity with a computer. Since the 'real' robots are run in a simulator window next to the GUI, the user gets immediate, continuous feedback and can follow the workings of the commands issued through the GUI.

2 Related Work

Control of a group of robots as a whole is still a rather unexplored field. Usually, only one robot is controlled at a time through 'alternative' (not programming-like), intuitive, easy-to-use interfaces, typically on the Web. A variety of different interfaces has been suggested.

In [10, 11], a heterogeneous team of three robots (a hovering, semi-autonomous helicopter and two autonomous mobile robots on the ground) are controlled via a GUI to perform a surveillance task. All robots are equipped with wireless video links, radio communication, and GPS. For example, a ground robot can be told to follow the helicopter, thus providing the user more detailed footage of the area overlooked by the helicopter. Also, it is possible in the GUI to identify an area for a robot to patrol.

[3] describes a system where a human user can control a group of Khepera robots as one unit. The robots perform an assembly task of pushing two boxes together that match in shape. Through a GUI, the user can control the robots at individual level or at group level, and potentially divide the robots into sub-groups. The higher the level of control, the less information from the robots flows back to the user; thus, irrelevant details are hidden.

In [9], Web users can choose from a discrete list of locations and command the mobile robot Xavier to go there

*Computer Science Dept., Aarhus University, Denmark

†Computer Science Dept., University of Southern California, Los Angeles, CA 90089-0781, USA

and tell a joke or take a picture. An important feature is the feedback to the users who can see what the robot currently sees. In [8], Web users control a Khepera robot in a maze, but it is possible to get video feedback both from a camera onboard the robot, and from an overhead camera. Furthermore, there is a mirror in the maze allowing the robot (and hence the user looking through its eyes) to see itself.

The museum guide robots in [2, 12] have two interfaces: one onboard the robot, and a Web-based one. Through the onboard interface, museum guests can select from different guided tours by pressing buttons on a touch-screen. On the Web, users are allowed to hand-pick a series of exhibits by clicking them directly in a map, and the robot makes a shortest path going by them all. At the same time, the user can monitor the location of the robot in the Web interface through live video and a Java applet (giving more rapid updates than the video link).

Finally, there is a variety of software environments for multiple robot control. E.g., TeamBots [1], Player [6], and MissionLab [7] are all software packages and interfaces that allow users to control or program robots at varying levels of abstraction. With Player, the user writes the control code, for example in C, communicating with the robot through the robot server that handles low-level control of sensors and actuators. TeamBots is a collection of Java packages that provide the basis for user designed robot control programs. Using the specialized classes and their methods, the user gets access to and control over the robots' sensors and actuators. MissionLab supports object-oriented, graphical construction of control programs using groups of basic behaviors as higher-level building blocks.

3 GUI and formation algorithm

We wanted the GUI to the formation algorithm to be easy to use for any user, and at the same time take full advantage of the general features of the algorithm. Before we describe the GUI, we briefly outline the principles of the underlying formation algorithm to provide the reader with an understanding of the features and limitations of the GUI. We refer to [4, 5] for details of the algorithm.

3.1 Formation algorithm

We have devised a simple, general formation algorithm that uses only local sensing and control: each robot is autonomous and has no information about the position of other robots, except what it can sense by itself locally.

Each robot has a unique ID number that it broadcasts regularly as a heartbeat message. Robots' IDs are detectable by other robots. One robot is the *conductor*, the leader in front of the formation; any robot can take on this role. The conductor broadcasts the current formation number along with its ID. Consequently, all robots know at all times the total number of robots participating (N) and what formation they should do.

Except for the conductor, each robot follows one designated neighboring robot, its *friend*, and keeps a certain distance and angle (with respect to its heading) to it[1]. All robots have at most one follower, except the conductor who can have one or two. This means that the robots are connected in a *chain of friendships*; the chain is always sorted by ID. For some formations, the conductor is the robot that currently holds the lowest ID, for others it is the one with the median ID. The latter group of formations are called *centered formations*, since in these, the conductor is the robot in the center of the chain of friendships, thus having two followers (one on each side). In the non-centered ones, the conductor (at the one end of the chain) only has one follower. If the conductor fails and stops sending heartbeat messages, the robot with the now median or lowest ID becomes the new conductor.

By the formation number and its own ID, each robot R knows its role in the formation: whether it is the conductor, and if not, which friend to look for. By a simple geometric calculation based on N and its rank (called its *lessThanMe*) in the chain of friendships, R also knows the angle to keep to its friend (see Figure 1). When the conductor starts moving, the only stable configuration of the whole group is when all robots find the same heading as the conductor, in a kind of self-organization. Thus, the conductor 'drags' the formation into place just by going its own way, without broadcasting its heading. If a robot dies and stops sending heartbeat messages, the remaining robots realize that N has changed, and so each of them might adapt the angle it is keeping to its friend, or find a new friend, or even take over the conductor duty.

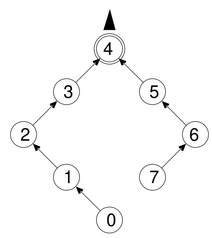


```
n = ⌈N/4⌉
if lessThanMe < ⌊N/2⌋ - n
  θ = 45
else if lessThanMe < ⌊N/2⌋
  θ = -45
else if lessThanMe > ⌊N/2⌋ + n
  θ = -45
else if lessThanMe > ⌊N/2⌋
  θ = 45
```

Figure 1: Calculating the friendship angle in a diamond with 8 robots. The diamond is a centered formation, so the robot with the median ID ($lessThanMe = \lfloor N/2 \rfloor$) is the conductor. Arrows indicate angle to friend, filled triangle shows formation heading.

In our implementation of this algorithm, based on the sensors of our physical Pioneer robots, each robot uses a laser scanner for distance measurements, a radio link for communication, sonars for close-up collision avoidance, and a panning color-blob detecting camera for identify-

[1]The term *conductor* is analogous to *leader* in the literature. Each robot follows its *friend*, so all robots (except the boundary cases) serve as "local leaders", and all are also followers (except the conductor).

ing its friend's ID; speciel custom-made helmets with two brightly colored stripes encode and display the robots' IDs. The only way of telling one robot from another is by the color helmets. Odometry on the robots is not realiable, and so they cannot keep track of their positions in real-world coordinates. The only reliable position information is a robot's relative position to its friend, and that is sensed locally. Thus, each robot knows only the (relative) position of *one* other robot.

The forward-pointing lasers have 180-degree fields of view, so no robot can keep its friend 'in sight' if the friend is behind it. This results in a restriction of the type of formations that our algorithm captures: they are precisely all shapes that can be folded from the chain of friendships, keeping either the middle or the end of the chain in front, with no robots having to look back for their friends. For example, with upwards heading, a shape like $\cap$ is possible, whereas a $\cup$ is not.

For formal experimental evaluation of the formation algorithm, we refer the reader to [4, 5]. In this paper we only show some empirical results with real robots. For the GUI experiments, we used an improved version of the algorithm that has very high stability for all friendship angles; i.e., after a formation is established, it is maintained in over 90% of the remaining trial duration if no obstacles are present, on average. We used four ActivMedia Inc. Pioneer2 DX robots (named Huey, Dewey, Louie, and Donald) with the SICK LMS200 laser, sonars, and the Sony PTZ camera[2], and with the Player [6] software. Player is a server and protocol that connects robots, sensors, and control programs across the network.

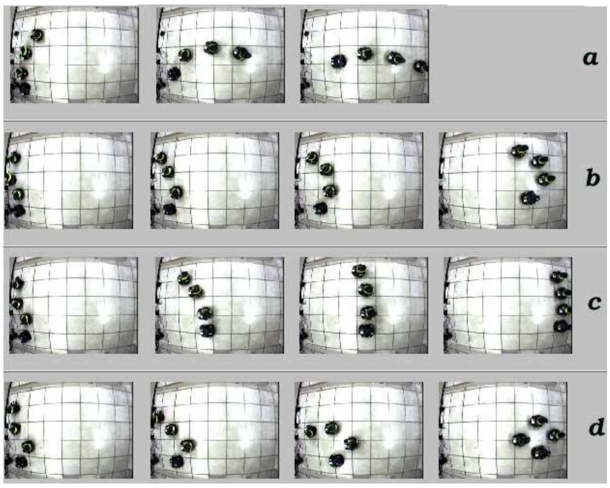

Figure 2: Overhead camera images, order is left to right. Establishing a column (**a**), wedge (**b**), line (**c**), and diamond (**d**).

In Figure 2, overhead camera images are shown from four experiments where four robots establish the basic formations line, wedge, diamond, and column (due to lack of space, the column was not completely straight in the end). Initially the robots were placed in the correct order with respect to their IDs but otherwise arbitrarily and with random headings. Figure 3 shows a switch from diamond to line. The diamond unfolds as the back robot swings out and catches up, moving to its friend's side as it changes its camera angle from -45 to +90 degrees.

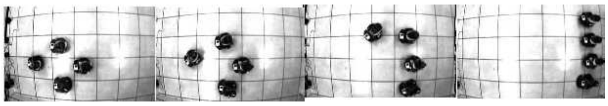

Figure 3: Sequence of overhead camera images of a switch from diamond to line, order is left to right.

In the experiment of Figure 4, two robots had to negotiate a wall that was only in the path of one of them (Huey). Looking far ahead, Huey alerted the other one, Louie, with a simple radio message that it was about to make a swerve, and so Louie could make a so-called swerve of solidarity in time to let Huey keep its position. Hence the robots maintained the line formation while still evading the wall.

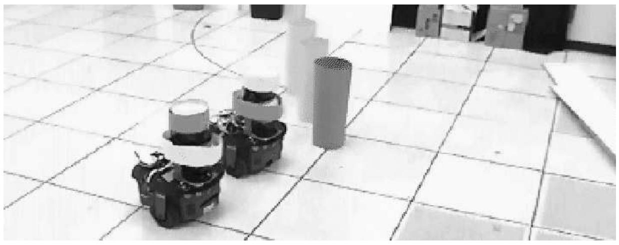

Figure 4: Two robots in a line, negotiating a wall.

3.2 GUI

Since our algorithm enables each robot to follow a friend at any angle between -90 and +90 degrees (thus not looking back), the group as a whole is capable of a great variety of customized formations besides the typical, basic, symmetrical ones (line, column, diamond, wedge). However, describing a customized formation in terms of angles and distances would be tedious and un-intuitive. Therefore, we have devised a graphical user interface in which the user gives a graphic representation of the formation he[3] wants the robots to do.

The core of the GUI is the chain of friendships. When started, the GUI displays a group of robot symbols (*nodes*) in the line (abreast) formation (see Figure 5). The nodes are

[2]Figure 4 shows Huey and Louie with the color helmets encoding their IDs.

[3]Throughout the paper, for convenience, we will write 'he' for the anonymous user.

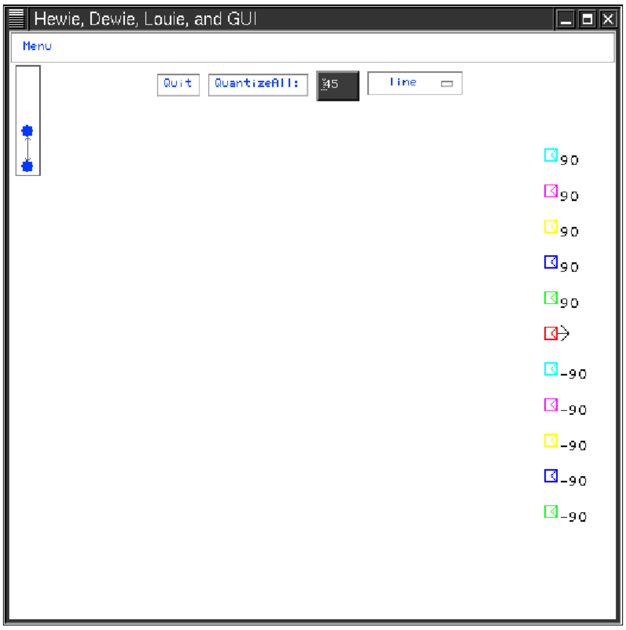

Figure 5: The GUI, here run with 11 nodes.

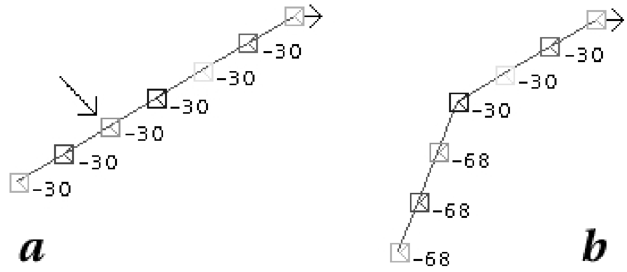

Figure 6: GUI: In formation **a**, the node marked by a large arrow is dragged downwards. Preserving its distance to its friend, it will move along an arc, centered in this friend. Its immediate follower, and its follower, etc., will describe the same motion as a rigid body. The result is formation **b**).

similar to the robots running in the simulator[4] window next to the GUI window. Also, next to each node is displayed the angle it is keeping to its friend. The conductor, the fixed point of the structure, is marked with an arrow indicating the formation heading. Meanwhile, in the simulator, the robots are started and begin establishing the line formation. Within a minute (depending on the number of robots), the robots have established the line, and the simulator window and the GUI window present similar displays. The robots maintain their current formation until commanded otherwise.

The user can send commands from the GUI. With the mouse, he can click and drag each individual node to change its friendship angle. By doing so, the node's cadre of followers, if any, moves like a rigid body attached to the node being dragged (see an example in Figure 6). The GUI imposes the restriction that no robot can keep an angle of more than ±90 degrees to its friend, and so all shapes that the user is allowed to fold out of the chain of friendships in the GUI can actually be established by the running robots in the simulator.

When done folding the chain, the user saves the new formation and gives it a name. By choosing 'download' from the top-left menu (see Figure 5), the data of the formation are broadcast to the robots in the simulator and they begin re-arranging themselves to establish the new formation. In the GUI, there is a pull-down menu that initially lists a few basic formations (line, column, diamond, wedge, circle); the newly created formation's name is added to this list. A quick way to make the robots do, e.g., a circle is to choose the circle from this pull-down menu: the nodes in the GUI transform into a circle, and when they are done, the user can press 'download' to command the robots to do the same – see Figure 7. This can be done with a new formation too, once it has been saved.

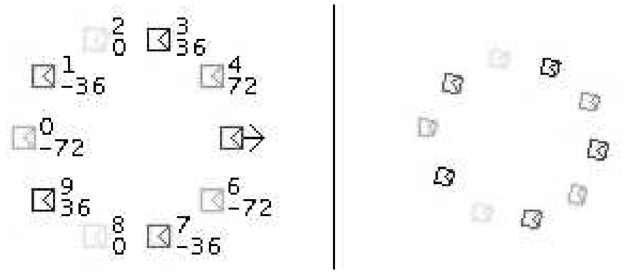

Figure 7: A 10-robot circle; GUI nodes on the left, simulated robots on the right. Nodes display ID and friendship angle.

Regarding the appearance of the nodes, the user can toggle between different features in the menu. The nodes can be displayed as boxes (similar to the simulated robots) or as filled circles; each node can have its ID and/or its friendship angle shown next to it; and a gray line can be drawn through all nodes to clarify the chain structure. Finally, by dragging up and down the top one of the two circles shown in the top left of the GUI, the user can scale the size of the whole formation by setting the inter-robot distance.

When switching between a centered and a non-centered formation (or vice versa), the conductor duty shifts between the median-ID robot and the lowest-ID robot. The robots between these two consequently have to find new friends on the opposite side of the old ones, and this makes such switches a bit more involved than switches between two centered or two non-centered formations. To make sure the finding of new friends goes smoothly, any such switch

[4]Using software Player and Stage [6, 13], developed jointly at the USC Robotics Research Labs and HRL Labs and freely available under the GNU General Public License from `http://playerstage.sourceforge.net`.

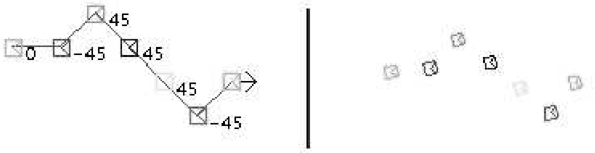

Figure 8: GUI: A customized, non-centered zigzag formation.

goes through the line formation. When a line is established, the conductor duty changes, and the involved robots simply pan their cameras 180 degrees since their new friend is now directly opposite their old one (to their right instead of their left, or vice versa). Then the switch from the line to the intended formation can take place. Figure 8 shows on the left a GUI-designed, non-centered zigzag formation, and on the right the formation that the robots actually formed (see figure 9) when the zigzag was downloaded (image **j** from figure 9).

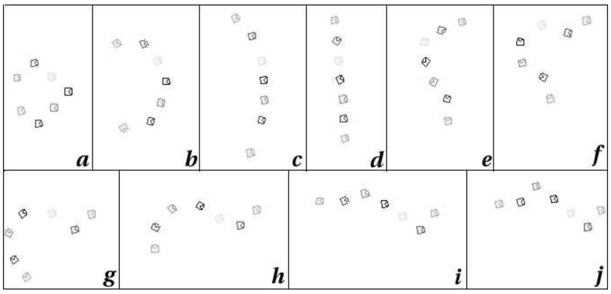

Figure 9: Simulator: A switch from a centered (incomplete) diamond with 7 robots to a non-centered zigzag-formation; in image **d**, the intermediate line is established and the top robot becomes the conductor, taking over from the middle robot.

Various tools are available to the user for designing new formations. By right-clicking and dragging, the whole formation can be rotated around the conductor, instead of just one single node around its friend. The 'quantize all' button changes all angles to the closest multiple of the *quant* parameter. This parameter is particularly useful for aligning the robots precisely. E.g., if a segment of nodes have angles 42°, 46°, and 44°, respectively, to their friends, quantizing the nodes with a quant of 45 will align all of these angles to precisely 45°. Applying the same quant value of 45 to angles of, e.g., 6°, 11°, and 94°, will quantize them to 0°, 0°, and 90°, respectively. The quant is displayed next to the button (in Figure 5 it is 45), and it can be set arbitrarily by the user. Pressing `shift` and left-clicking on a node quantizes that node only.

4 Discussion

As explained, the algorithm (and hence the GUI) allows all shapes where 1) no robots have to look behind them for their friends, and 2) either the middle or one end of the chain of friendships is in front. This still leaves room for highly specialized formations, like the one in Figure 9.

Building the GUI around the chain of friendships yielded a very flexible and versatile formation design tool. Folding the structure into a desired shape is fast and intuitive. Even with the restrictions, surprisingly varied formations are possible. Imagine a future Superbowl halftime show in which hundreds of micro-air vehicles (MAVs) spell out messages in the air above the stadium. Such vehicles would have to be small, and hence their sensing capabilities would probably be limited, making them a good possible application for our GUI and local formation algorithm.

In addition to being a design tool for new formations, the GUI offers an easy means of formation control to a commander of the robots. By the pull-down menu with its list of basic formations (and possibly new ones), it is straightforward to command the robots to perform a switch between any two formations. The GUI shows a small animation of how the robots – ideally – would carry out the switch; how each individual robot would gradually change its friendship angle and slide to its correct, new position. Figure 10 demonstrates a switch from a customized formation, shown on top, to the line; in the simulator screenshot below, the robots were initially in this customized formation. Their traces show how they performed the switch.

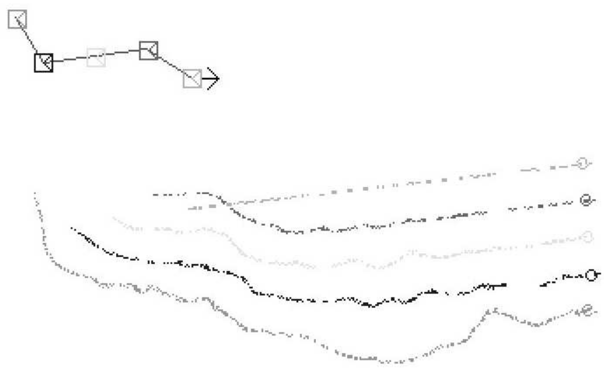

Figure 10: A switch from a customized 5-robot formation (shown in the GUI on top) to the line (simulation with traces, shown underneath).

5 Conclusions and future work

A formation is a spatial phenomenon; it is more intuitive to give a graphical description of a formation than a list of angles and distances. There is no way of representing more intuitively the formation in, e.g., Figure 9**j** than the picture itself; Figure 9**j** is its own best representation.

For this reason we have devised a graphical user interface that enables users to design formations and have a simulated group of robots establish the formation immediately, in real-time. The user needs no prerequisites in mathematics, robotics, or football marching band conducting. Switches between any two formations is also possible and easily commanded through the GUI.

In any control scenario, feedback is important: it is desirable to be able to judge the progress after issuing a command. With direct remote control of robots (like in [8]), the effect of a command is immediately visible, whereas in systems with some local intelligence and a more abstract level of control (like [9, 2, 12]), the (initial) effect of a command might be more obscure. In our system, with the simulator running next to the GUI, the robots start rearranging themselves immediately after a formation change command is given, and the user can see how each robot gradually changes its angle towards the new, right one. Another important point, also made in [8], is that our GUI captures all possible operations of the underlying algorithm. Thus, even with control at a very high level of abstraction, the user still feels 'in command' of all the individual robots, and no functionality is abstracted away.

We are currently working on generalizing the GUI, so that one can save newly created formations in a database and retrieve them for later experiments. The formations are stored in a general way, so that line segments where n robots have the same angle to their friends are represented as a segment of length n/N, with N being the total number of robots (this is where the 'quantize' function comes in handy). Thus, a formation designed originally for $N_1 = 10$ nodes would be applicable to a group of, say, $N_2 = 20$ robots. Of course, this introduces a potential problem if N_2 is not a multiple of N_1; in fact a problem without a 'correct' solution. In this case some interpretation of the user design has to be made, but since there is no way of knowing what abstract shape he had in mind, any choice would do. Alternatively, the GUI could suggest a few generalized shapes and let the user choose among them.

As robot teams emerge to entertain us as well as to handle hard or dangerous problems in our everyday lives, easy-to-use tools for human control of groups of robots will be needed. The work on this GUI is part of an effort toward this goal; the GUI serves as an example of an application-specific solution. More information, including video clips of both simulated and real robot formations and of the GUI, can be found at http://robotics.usc.edu/∼agents/projects/formations.html.

References

[1] Tucker Balch, *TeamBots*, www.teambots.org.

[2] Wolfram Burgard, Armin B. Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, Sebastian Thrun, *Experiences with an interactive museum tour-guide robot*, Artificial Intelligence, Vol. 114, No. 1-2, pp. 3–55, 1999.

[3] José Beltrán-Escavy, Akio Nakamura, Shinjiro Kakita, Tamio Arai, Jun Ota, *Human-Robot GUI System With Group Control For Multiple Mobile Robot Systems – graphic interface and comparison with previous joystick-based system*, Intl. Journal of the Japan Society for Precision Engineering, Vol. 32, No. 4, pp. 291–296, December 1998.

[4] Jakob Fredslund, Maja J Matarić, *A General, Local Algorithm for Robot Formations*, Accepted for publication in the IEEE Transactions on Robotics and Automation, to appear in April 2002.

[5] Jakob Fredslund, Maja J Matarić, *Robot Formations Using Only Local Sensing and Control*, Intl. Symp. on Computational Intelligence in Robotics and Automation (IEEE CIRA 2001), pp. 308–313. Banff, Alberta, Canada, July 29 - Aug. 1, 2001.

[6] Brian P. Gerkey, Richard T. Vaughan, Kasper Støy, Andrew Howard, Gaurav S. Sukhatme, Maja J Matarić, *Most Valuable Player: A Robot Device Server for Distributed Control*, Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS 2001), pages 1226–1231, Wailea, Hawaii, October 29 - November 3, 2001.

[7] Douglas C. Mackenzie, Ronald C. Arkin, Jonathan M. Cameron, *Multiagent Mission Specification and Execution*, Autonomous Robots, Vol. 4, No. 1, pp. 29-52, Jan. 1997.

[8] P. Saucy, F. Mondada, *KhepOnTheWeb : One Year of Access to a Mobile Robot on the Internet*, Proc. IEEE Intl. Conf. on Intelligent Robots and Systems (IROS-98) Workshop on Robots on The Web, pp. 23–30. Victoria, Canada, Oct. 13-17, 1998.

[9] Reid Simmons, Joaquin Fernandez, Richard Goodwin, Sven Koenig, Joseph O'Sullivan, *Xavier: An Autonomous Mobile Robot on the Web*, IEEE Robotics and Automation Magazine, Vol. 7, No. 2, 1999.

[10] Gaurav S. Sukhatme, James F. Montgomery, *Heterogeneous Robot Group Control and Applications*, Proc. Association for Unmanned Vehicle Systems, International (AUVSI-99) Conf., Baltimore, Maryland, USA, 1999.

[11] Gaurav S. Sukhatme, James F. Montgomery, Maja J Matarić, *Design and Implementation of a Mechanically Heterogeneous Robot Group*, Proc. Society for Photo-optical Instrumentation Engineers (SPIE) Conf. on Sensor Fusion and Decentralized Control in Robotic Systems II, Vol. 3839, pp. 122-133, Boston, Massachusetts, USA, 1999.

[12] Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B. Cremers, Frank Dellaert, Dieter Fox, Dirk Hähnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulte, Dirk Schulz, *MINERVA: A Second-Generation Museum Tour-Guide Robot*, Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA-99), pp. 1997–2085, Detroit, Michigan, USA, May 10-15, 1999.

[13] Richard T. Vaughan. *Stage: a multiple robot simulator*, Institute for Robotics and Intelligent Systems Technical Report IRIS-00-393, University of Southern California, 2000.