Andrzej Chmielniak, M.Sc. M.E. (1)
Artur Dubrawski, M.Sc. M.E. (2)
Barbara Siemiątkowska, Ph.D. (2)

(1) Faculty of Mechanical, Power and Aerospace Engineering, Warsaw University of Technology.
(2) Institute of Fundamental Technological Research, Polish Academy of Sciences.

A distributed system for control and management of teams of mobile robots[1]

Abstract

This paper discusses a distributed control system designed for management and steering of mobile robot teams. A vast majority of mobile robots' applications relies on using single robots only. In fact, many tasks may be executed more effectively, quickly and more precisely by the teams of robots. The purpose of our research is to build a system for steering and monitoring heterogenous groups of indoor mobile robots, which would efficiently use computing resources of the networked onboard and stationary computers. The robots should communicate and cooperate in planning and realization of their tasks, but at the same time they should retain as much autonomy as possible. Although the system is in a fairly early stage of development, it is already possible to demonstrate its capabilities on selected sample tasks.

1. INTRODUCTION

A majority of mobile robotic applications consider implementations of single robots, arranged to achieve their particular tasks individually. In practice, however, there is a substantial area of possible applications in which teams of robots may either perform better than just bunches of the individual agents, or even, in some cases, the teamwork may make the solution conceivable.

The purpose of the research described in this paper is to build a distributed control system for steering and monitoring teams of mobile robots which can perform various tasks in indoor environments. The robots should cooperate in planning and realization of their tasks, but at the same time they should retain as much autonomy as possible. In the system there are only a few centralized processes, such as inter-robot communication control, sensory data acquisition, and the operator's interface.

The system will enable an efficient coordination of teams composed of several, possibly heterogenous vehicles. The heterogenity may occur as well at the level of locomotion capabilities (wheeled or legged platforms, various kinematic and dynamic characteristics), as at the levels of intelligence and autonomy (different sensor modalities, different computing power onboard, etc.).

[1] This work has been supported by the research grant KBN No. 8T11A01312, sponsored by the State Committee for Scientific Research, Republic of Poland, and cooperatively pursued at the Faculty of Mechanical Power and Aircraft Engineering, Warsaw University of Technology and the Institute of Fundamental Technological Research, Polish Academy of Sciences.

The final system should be sufficiently flexible in order to become useful in a variety of possible applications and team configurations, but yet compact enough to remain practical.

The main design assumptions consider intra-team (robot-to-robot) and human-operator-to-team communication issues. Some robots may be controlled with their own onboard computers, and communicate with the other members of the team and with the operator through Ethernet (radio or cable) links. The other vehicles, as well as the human operator, will rather rely on stationary computers and use serial radio links (the robots) or console (the operator) to get in touch with the others. In the context of the system, each robot is seen as a software agent, composed of one or more processes running on one or more computers simultaneously. The system contains also an inter-agent communication program, whose structure and the mode of operation is similar to that of the electronic mail. The other main part of the designed system is the operator's interface, which should provide the operator with high-level control abilities (like the team task specification - the individual robot's activities should be determined automatically) and with an access to the task execution monitoring functions. The operator's interface may also have a distributed, networked implementation.

All the required flexibility and compactness should make it possible to achieve a variety of profits from using many robots. For instance, one of the predicted synergies in navigation will arise in a team-based world modeling, in which the common map of the workspace, as a result of aggregation of data collected by the members of the team carrying different sensor modalities and/or placed in different places, should result in a better model than with the same number of vehicles working alone. We predict even more important outcome from mutual position tracking and mutual relocation of the robots. In the following section of this paper we will give a broader discussion of team tasks considered in our project. Although the discussed system is in a fairly early stage of development, it is already possible to demonstrate its certain capabilities on sample tasks such as manoeuvring in formations or transporting a relatively long beam in a cluttered indoor environment by the team of two mobile robots. This will be described in detail in a section on experiments in the later part of this article.

Among members of our experimental team there are three mobile robots *Pioneer-1* and one *B-14* made by Real World Interface, Inc. (USA), but the multi-robot system should be sufficiently flexible to work properly with any type of mobile platforms, including walking machines. We do not impose any limit to the number of robots in a team now, but for practical reasons the control and monitoring processes should be distributed over a network of host computers. In our case they are three PCs running Linux and X-Window system, on which Motif libraries are used for building operator's interface (one of these computer resides onboard *B-14*). There is also one Sun SparcStation LX with an attached SparcClassic X terminal running Solaris 2.6 and X-Window, and a couple of PCs with Microsoft Windows NT and 95 installed. It is possible to use software simulators of the robots, as well as real vehicles, as the agents which perform tasks and communicate through the distributed control system. Communication between the computers goes through a local area network (Ethernet under TCP/IP). Some robots (like our *B-14*) may have onboard PCs and use their computing power for self control and for running control processes of other members of the team. Other vehicles in our team are connected with (perhaps stationary) computers through serial radio links.

2. RELATED RESEARCH

There is a growing number of publications related to the research and applications of mobile robot teams. Robot formations issues are discussed in [1] and [2], where results obtained merely with computer simulations are shown. In this paper we present formation control experiments conducted with real vehicles. Military-oriented applications of mobile robot teams are presented in [3]. Planning optimal paths for robots in a team has been described in [8]. Some missions for robots equipped with grippers have been specified in [3] and [4]. Implementation of flock behaviors such as foraging, grazing, and herding has been presented in [2], while in [11] there were described successful applications of automatic control of flocks of animals. An interestingly large number of robot teams rely on behavioral control schemes (cf. [4], [5], [6] and [10]), which is sometimes motivated by observations of animal behavior. In our system at least some parts of the control architectures will follow behavioral approach. In [9] there was specified a distributed control system for a pair of manipulators arranged on a mobile platform. Papers [3] and [5] are concerned with operator's role in controlling teams of robots. Inter-robot communication issues and team-wise data collection strategies were the topics of [3], [5], [9] and [10]. An interesting work on theoretical aspects of visibility maintenance for a single indoor robot has been published in [7]. In the near future we will attempt to extend that work towards teams of surveillance robots.

3. TASKS FOR THE TEAMS

The tasks of a team of mobile robots may be classified according to various criteria. One of them is a possibility of performing the task with a single robot:

- tasks which are achievable by a single robot,
- tasks which are unattainable for a single robot, but achievable by a team.

The next criterion is related to an order of improvement of the task realization effectiveness, related to the number of cooperating robots. The effectiveness may be measured with time necessary to perform a given task, consumed energy, obtained quality, and so on, including combinations of the individual criteria. According to such criterion we may distinguish tasks for which the effectiveness is:

- increasing more quickly, than the number of robots in a team,
- increasing proportionally to the number of robots,
- increasing more slowly, than the number of robots,
- independent on the number of robots,
- decreasing with the number of robots.

Some tasks may change their place on the list above in dependence on the number of robots in a team. It may occur for instance when the robots' workspace has limited dimensions, and too many robots form a crowd of agents which disturb each other in achieving their goals.

Depending on the types of robots in use and on their principal duties, we may distinguish, select and take a closer look on some particular applications of mobile robotic teams.

Exploration

The team may have an initial (a priori) map of the environment and each robot knows its initial location on the map. The first job for the team is to divide the workspace and to assign operation zones for individual robots. Next, the robots begin to explore their zones and perhaps they find discrepancies between the a priori map and the sensory-based observations. In such cases they mark the locations and characteristics of the findings. Then, depending on the particular situation, they update their maps accordingly or alternatively call some other member(s) of the team, which perhaps carry unique sensors (such as cameras), to interpret the findings and to verify and confirm the suggested map updates. Such a team-based exploration scenario may be straightforwardly extended towards exploring unknown environments, or towards natural landmark selection (cf. [2] and [6]).

Surveillance

Presumably, the team knows the a priori map of the workspace. At some moment of time some new (or intruding) objects, not yet represented in the map may appear, and some of them may be moving. The main goal for the team is to tell these objects from the expected ones, identify them to as much extent as possible, inform the operator about them, and perhaps begin to follow and track the intruder(s). Like in the case from previous paragraph, the robots begin their job by dividing the workspace into a set of individual, maybe overlapping, surveillance zones. One can imagine that in particular setups there will be some areas of a larger probability (or alternatively frequency) of intruders' appearance than it is elsewhere. Taking also into account a variety of the team members' capabilities, such a team-based surveillance task may be seen as an exciting optimization problem. One of the most important objectives might be to allocate limited team's resources in a way, which would minimize a probability of not detecting the intruding objects. It involves designing team-wise and individual strategies, which may lead to optimal solutions by assigning requested patrol frequencies to the particular places in the workspace. Then, the team (and its control system) would be responsible for planning, executing, coordinating and monitoring actions of the individual members in order to meet the requirements, and to adapt to changing situations. One of the evaluation criteria in this kind of surveillance task may be based on the level of a visual (or in general: sensory) contact with the intruders. Strategies which provide continuous and redundant contact with intruders, given certain limitations, should be rewarded. Of course the best strategy for a given team may vary from one environment to another. We suspect there is a room for some self-adaptation skills, developed with machine learning techniques, which would give a team a certain amount of flexibility in accommodation to varying workspace configurations.

Guarding places

This kind of task may be illustrated with a scenario of restricting an access to some dangerous place such as a mine discharge zone or a contaminated area. Presumably, the team knows this place, and is supposed to automatically decide about the locations and surveillance zones for each of its members, so that each entry to the guarded area will be sufficiently secured. In case of detecting any intruder the operator should be immediately informed. Depending on the type of danger present in the area, some parts of it may be inaccessible or too dangerous for some or all team members. The guarding strategy must to take it into account, and maintain the required safety margins for each of the robots.

Driving moving objects

Here we consider driving a human, an animal, a robot or some other moving object or their group to a desired place. It may be realizable under assumptions that the driven object(s) behavior is to some extent predictable, and that the robots repel them in some way. Also, the workspace is presumably known in advance. The team of robots is supposed to pursue a strategy similar to that of cowboys driving a cattle home, and to finish the job in a specified time. A spectacular implementation of a single robot controlling a flock of little ducks was reported in [11]. It is however clear that in sufficiently complicated cases a lone robotic cowboy just can't have it done, so an extension of the cited work would be in order. Also, it is not hard to conceive an implementation of a game of two driving teams playing their strategies against each other.

Transporting items

Let us assume that the individual robots have limited operating ranges in the workspace, for instance because of finite battery capacities, limited ranges of radio links and particular kinematic constraints. Then it is possible that a task of transporting some item across the workspace will require more than one robot to complete. In such a case the control system will come up with a relay strategy, in which each of the employed robots will be in charge of planning and executing their individual motions between the established item passing areas.

Sometimes it is necessary to employ more than one robot to carry the transported thing, because it is too large or too heavy for a single vehicle, or if there are certain limitations due to the construction of grippers or the shape of the item. In such cases the robots' trajectories must be planned together, and their movements have to be thoroughly controlled. A pair of robots carrying a long beam in a cluttered environment may serve as an example of such a transporting task. It was discussed in [5] and [6] for a case of a pair of robots with rotating heads. In our experiments we use a pair of *Pioneer-1* robots which do not possess such a nice feature, and a really precise coordination of the vehicles' movements is required in order to successfully carry the beam through a narrow doorway.

4. HARDWARE

As it has been already mentioned, we use simple mobile platforms *Pioneer-1* for the described project purposes. Their on-board micro-controllers are used for basic routines such as motion control, sensory data acquisition and communication with a host computer through a radio modem. *Pioneers* are equipped with odometers and with seven standard Polaroid ultrasonic range sensors each. Each of them can also carry a *LightRanger*, that is a 3-D scanning laser range finder made by HelpMate Robotics, Inc. (USA)[2]. Due to the micro-controller limitations, the navigational capabilities of these robots are practically limited to a point-to-point mode of operation. The mobility of a *Pioneer-1* is a subject of non-holonomic constraints due to their two-wheel differential drive mechanism with one caster wheel in the rear.

Another member of our robotic team, *B-14*, is a far more advanced device. It has an Intel Pentium-Pro based onboard computer running Linux, and connected to the world through an Ethernet cable (to be upgraded to a radio Ethernet link). It is equipped with rings of 16 Polaroid ultrasonic range sensors and as many infrared proximity sensors, as well as with two rings of bumpers mounted around its barrel-like casing. It also carries a color camera with a frame grabber, and speech synthesis and understanding software and hardware, which have not been used in this project yet. The upper part of the casing forms a rotating turret, and the mobile base contains batteries and a synchro drive mechanism. *B-14*'s fairly accurate kinematic setup allows for pursuing precise navigation with coordinated turns and other manouvres.

Our robots will be tentatively equipped with grippers. Also, we plan on mounting a stationary camera overlooking the robots' playground.

There are a few stationary computers which could be used for our system purposes. Among them there are two PCs running Red Hat Linux and X-Window with Motif libraries used for building operator's interface, one Sun SparcStation LX with an X-Terminal running Solaris 2.6 and X-Window, and a couple of PCs with Microsoft Windows NT and 95 installed. The hardware is wired and communicates through Ethernet local area network under TCP/IP protocol.

5. THE OPERATOR

In the distributed control and management system the operator interacts through an interface which allows for task type selection and specification, map management (selection, downloading, saving), and monitoring the status of the task realization and the states of each of robotic team members. In the standard mode of operation there is no need for controlling the individual robots, and the

---

[2] Provided generously by Dr. Ryszard Sawwa and his collaborators from the Industrial Institute of Measurements and Automation (PIAP) in Warsaw.

operator communicates with the team as a whole. A direct influence on the individual agent's behavior is limited to commanding an emergency stop and to clearing that command.

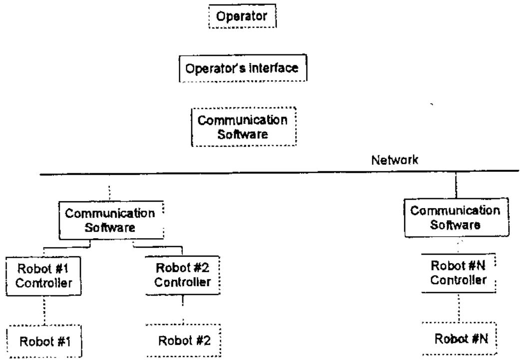

Fig. 1: Architecture of the distributed system for control and management of teams of mobile robots.

The operator has a set of tools for monitoring team's activities and achievements. These tools use information on the current states reported by each of the team members, and video images provided with on-board and stationary cameras. The team is responsible for reporting certain events (such as detection of an unexpected object in the workspace) to the operator, and for requesting operator's assistance or commands whenever desired.

It is necessary to provide the operator with a friendly graphic interface. At the current stage of development there is already one such program available. It is running under Red Hat Linux with XFree86 X-Window environment, and is based upon Lesstiff (Motif compatible library available at no charge for Linux/XFree86 platforms). With this software, the operator can watch activity of each robot upon request, thanks to a distributed and networked concept of X-Windows.

6. THE TEAMWORK

Work in a team requires coordination and communication between the processes controlling the robots, and the operator's interface. All these processes may run simultaneously on one or more networked host computers. The information transmitted between the operator's console and the robots consists of task execution requests, task execution status reports, world map updates, and the individual robots' state reports. Transmissions between the robots contain selected data on their own kinematic states (position, orientation, velocities), sensory-based local world models, and information necessary to share the duties and to coordinate execution of the subsequent tasks. The data collected by various team members and with other components of the system (such as stationary cameras for instance) can be aggregated in order to increase reliability and confidence of the constructed shared global maps of the workspace.

For a single computer inter-process information transfer there is implemented a messaging mechanism based upon a concept of shared memory. For communication between the processes running on different machines there is implemented an information distribution system similar to that of electronic mail. On each computer taking part in the game, there is set up a mail server taking care of the specified processes which may be run there. Each process is supposed to inform the server about its existence, and then it can access the assigned mail box(es), and it can also send messages to all other or to specified processes in the system, addressing them with their names. Each process can have more than one mail box of a specified size associated with it. Each message is marked with a timestamp, and the outdated messages may be removed from the boxes to make space for the new incoming mail. The basic idea of the messaging system is to make the actual pattern of distribution of processes among the computers in the network as transparent for the participating processes as possible. This makes it easier for a programmer to code and run new processes without spending much time on getting into communication details.

It is worth noting that when two or more robots equipped with the same type of ultrasonic range sensors operate in the same area there is a risk of cross-talk and corruption of range measurements. With our system it is however possible to synchronize sensing cycles of the individual robots and thus to avoid such problems.

7. EXPERIMENTS

A typical question encountered when implementing a multi robot system is how to resolve a colliding paths problem (Fig. 2). In our system the solution is implemented as a behavior in each robot's control algorithm. First, each of the robots tries to predict possible colliding paths and the time to the nearest collision. Then, depending on the particular situation, the robots coordinate their behaviors so that one of them accelerates while the other one breaks in order to keep the necessary clearance near the estimated collision point. We verified experimentally, that in case of *Pioneers* the sonars are too unreliable, and the collision resolver module must rely only on the internal estimates of each robot's pose, based upon odometry and model-based relocation.

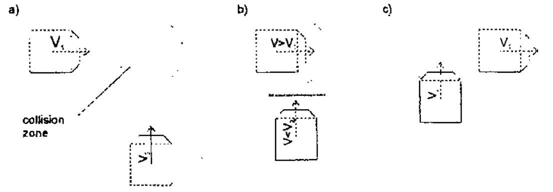

Fig. 2: Cross-path behaviors.

Another typical team-wise task for the robots is their movement in close formations. In our tests we have experimented with extended lines and lines ahead, and with manoeuvres needed to orderly switch from one to the other type of formation (Fig. 3). A typical illustrative experiment begins with a group of robots arranged in a line ahead executing a sequence of straight motions and coordinated turns. Then the formation type changes into an extended line, and again the team moves along straight and curved paths and turn. The robots should maintain their locations in the formation and keep the required separations by the way of coordination their movements upon the incoming kinematic data describing states of the individual members of the team. In the illustrative example shown in Fig. 3, the leading robot is teleoperated, and its partner's job is to keep the order of the formation.

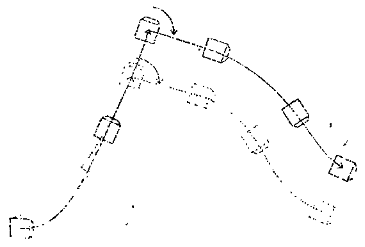

Fig. 3: Ordered manoeuvers in a formation.

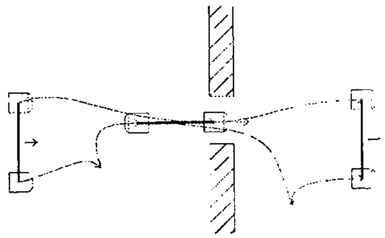

Fig. 4: Transporting a long beam through a narrow passage.

The third type of a sample task considered in this paper is transporting a long beam in a cluttered environment using a pair of *Pioneers*. The ends of the beam lie freely on the cargo areas of the robots' tops. In empty spaces the robots travel in parallel, trying to keep the beam perpendicular to their trajectories (Fig. 4, left). To get through a narrow passage such as a doorway, the robots must carefully reconfigure the formation so that they do not loose the beam (Fig. 4, center), and then get back to the original setup (Fig. 4, right). This type of a task requires a very focused and accurate coordination of the vehicles' movements, and it can serve as an excellent test for a distributed control system for teams of mobile robots.

8. CONCLUSION

It is clear that there exist many tasks at which the use of robotic teams may be desirable or advisable, in comparison to employing sets of individual robots. We strongly believe that development of practical, well defined, flexible, easy to reconfigure and easy to operate control systems is one of the most important prerequisites of an eventual success of the multi-robot approaches. Introduced in this paper, the distributed system for control and management of teams of mobile robots is in a quite early stage of development and validation, but it is already possible to show it working on a few selected sample tasks. The particular design considerations, such as using computing resources distributed over a LAN, make the described system capable for controlling teams of many cooperating heterogenous robots, which can work in large and spread indoor workspaces.

REFERENCES

[1] ALPTEKIN G., *Geometric Formation With Uniform Distribution And Movement In Formation Of Distributed Mobile Robots*, 1996, http://dubhe.cc.nps.navy.mil/~yun/student/gokhan.

[2] ARKIN R. C., ALI K. S., *Integration of Reactive and Telerobotic Control in Multi-agent Robotic Systems*, Proc. Third International Conference on Simulation of Adaptive Behavior, (SAB94), Brighton, England, 1994.

[3] ARKIN R. C., BALTCH T, *Cooperative Multiagent Robotic Systems*, AI-based Mobile Robots: Case Studies of Successful Robot Systems, MIT Press, 1996.

[4] BALCH T., BOONE G., COLLINS T., FORBES H., MacKENZIE D., SANTAMARIA J. C., *Io, Ganymede and Callisto - a Multiagent Robot Trash-collecting Team*, AI magazine, No. 16(2), 39-53, 1995.

[5] BARNES D. P., AYLETT R. S., CODDINGTON A. M., GHANEA-HERCOCK R. A., *A hybrid approach to supervising multiple co-operant autonomous mobile robots*, ICAR, Monterey, 1997.

[6] GHANEA-HERCOCK R., BARNES D. P., *Coupled behaviors in the reactive control of cooperating mobile robots*, Adv. Robotics, Vol. 10, No. 2, 161-177, 1996.

[7] LAVALLE S. M., GONZÁLES-BANOS H. H., BECKER C., LATOMBE J.-C., *Motion Startegies for Maintaining Visibility for a Moving Target*, Proc. of the 1997 IEEE International Confernce on Robotics and Automation.

[8] LAVALLE S. M., HUTCHINSON S. A., *Optimal Motion Planning for Multiple Robots Having Independent Goals*. IEEE Trans. On Robotics and Automation, Vol. 14, No. 6, December 1998.

[9] LEANGLE T., LUETH T. C., REMBOLD U., *A distributed Control Architecture for Autonomous Robot Systems*, Modelling and Planning for Sensor-Based Intelligent Robot Systems, 1995.

[10] SUN J., NAGATA T., KUROSU K., *Cooperative Behavior of a Schedule-Based Distributed Autonomous Robotic System*, Journal of Robotic and Mechatronics, Vol. 6, No. 2, 162-166, 1994.

[11] VAUGHAN R., SUMPTER N., FROST A., CAMERON S., *Experiment in automatic flock control*, Proc. 6th International Symposium on Intelligent Robotic Systems, SIRS 98, 47-54, Edinburgh 21-23 July 1998.