IROS 2002
2002 IEEE/RSJ International Conference on Intelligent Robots and Systems
September 30 – October 4 2002
EPFL Lausanne, Switzerland

Proceedings Workshop WS9

Robots in Exhibitions

Organizers:
Kai Oliver Arras
Wolfram Burgard

Tuesday, 1 October 2002

Schedule

Time	**Title**, Speaker(s)
9.00	**Welcome and Introduction to the Workshop** Kai Oliver Arras, EPFL, Switzerland Wolfram Burgard, Albert-Ludwigs-University Freiburg, Germany
MORNING	
9.20	**Entertainment Robotics: Examples, Key Technologies and Perspectives** Oliver Barth, GPS GmbH, Stuttgart, Germany
9.40	**The Mobot Museum Robot Installations: A Five Year Experiment** Illah Nourbasksh, Robotics Institute, Carnegie Mellon University, USA
10.00	**Online Robot Teacher Kits for Museum Field Trips** Gerard McKee, The University of Reading, United Kingdom
10.20	**Mobile Robots in Art Museum for Remote Appreciation via Internet** Shoichi Maeyama, Osaka Electro-Communication University, Japan
Coffee Break (10.40 - 11.00)	
11.00	**TOURBOT and WebFAIR: Web-Operated Mobile Robots for Tele-Presence in Populated Exhibitions** Wolfram Burgard, Albert-Ludwigs-University Freiburg, Germany Panos Trahanias, University of Crete, Greece
11.40	**Diligent: Towards a Personal Robot** Rachid Alami, LAAS-CNRS, Toulouse, France
12.00	**Design of and Operational Experiences from Five Museum Robot Installations** Sjur Vestli, MRS Automation, Zurich, Switzerland
12.20	**On the Prospects of Robots in Museums** Andrea Niehaus, Deutsches Museum Bonn, Germany
Lunch Break (12.40 - 14.00)	
AFTERNOON	
14.00	**Demonstrating the Humanoid Robot HERMES at an Exhibition: A Long-Term Dependability Test** Volker Graefe, Bundeswehr University Munich, Germany
14.20	**Experiments at Trade Fairs with Blacky the Robot** Diego Rodriguez-Losada, Universidad Politecnica de Madrid, Spain
14.40	**System Integration, Navigation and Interaction for the Expo.02 Exhibition** Kai Oliver Arras, Björn Jensen, Nicola Tomatis, Roland Siegwart, EPFL, Switzerland
Coffee Break (15.40 - 16.00)	
16.00	**Interactive Exhibition Design: Robots in Exhibitions** Olaf Arndt, Machine-Performance Group BBM, Berlin, Germany
16.30-18.00: Round table discussion	

Table of Contents

9. **On the Prospects of Robots in Museums**
   Andrea Niehaus
   Deutsches Museum Bonn, Germany

10. **Demonstrating the Humanoid Robot HERMES at an Exhibition: A Long-Term Dependability Test**
    Rainer Bischoff, Volker Graefe
    Bunderwehr University Munich, Germany

11. **Experiments at Trade Fairs with Blacky the Robot**
    Fernando Matia, Diego Rodriguez-Losada, Ramon Galan, Agustin Jimenez
    UPM-DISAM, Universidad Politecnica de Madrid, Spain

12. **Design and System Integration for the Expo.02 Robot**
    Nicola Tomatis, Gregoire Terrien, Ralph Piguet, Daniel Burnier, Shamir Bouabdallah, Roland Siegwart
    Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland

13. **A Navigation Framework for Multiple Mobile Robots and its Application at the Expo.02 Exhibition**
    Kai O. Arras, Roland Philippsen, Marc de Battista, Martin Schilt, Roland Siegwart
    Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland

14. **Visitor Flow Management Using Human-Robot Interaction at Expo.02**
    Björn Jensen, Gilles Froidevaux, Xavier Greppin, Antoine Lorotte, Laetitia Mayor, Mathieu Meisser, Guy Ramel, Roland Siegwart
    Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland

15. **Interactive Exhibition Design: Robots in Exhibitions**
    Olaf Arndt
    Machine Performance Group BBM, Berlin, Germany

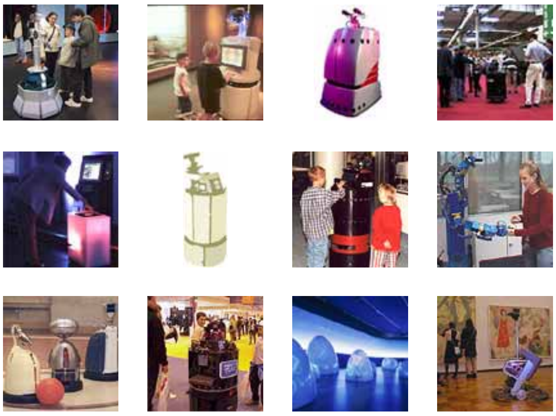

Introduction to the Workshop

Over the past years, the number of robots that have been deployed in museums, trade shows and exhibitions has been grown steadily. The exhibition context has emerged as a new application domain of autonomous robots. At the same time robots have become a new media technology for curators, exhibition makers and artists.

This workshop brings together robotics researchers, practitioners and end-users from various backgrounds to discuss past and ongoing projects, recent developments and visions for the future. Four spin-off companies which have emerged from this area attend the workshop (*Mobot Inc., GPS Ltd., MRS Automation Ltd., Bluebotics Inc.*). Furthermore, we are pleased to welcome two experienced end-users: Olaf Arndt, who was the artistic director of the 72-robot exhibition at the Expo 2000 World Exhibition in Hannover, Germany, and Andrea Niehaus, Director of the Deutsches Museum Bonn, who hosted the early and successful exhibition robot *Rhino* in 1998. The speakers of this workshop unite a seven-year operation experience with more than 90 robots that traveled several thousand kilometers in exhibitions and interacted with hundred thousands of people.

Taxonomy

It is interesting to see the variety of exhibition types, tasks and objectives of these projects. We can structure them according to several criteria: First, there is the exhibition type[1]

- technical museums (*Rhino, Minerva, Museum of Communication Berlin, Hermes)*
- natural science museums (*Mobot museum robots)*
- art museums (*Kapros)*
- trade shows (*Diligent, Blacky)*
- mass exhibitions (*Expo 2000, Expo.02)*
- museums in general (*Tourbot, Webfair)*

Each context poses specific technical, financial, artistic or educational problems. In mass exhibitions, for example, visitor flow, that is, the number of visitors the exhibition can serve per hour is an issue which is much less important in museums. Second, decisive for the role of the robot in the exhibition is the exhibition's message

- robotics-related *(e.g. Expo.02, Hermes)*
- robotics-external *(e.g. Mobot museum robots, Kapros, Expo 2000)*

In the latter case, the robot is the medium of this message. Appearance and interaction have to be designed under this perspective. It might have an educational task and the robot-specific properties are not the primary concern. In the former case, that is, the robot itself is part of the message or *is* the message, more playful interaction scenarios can be appropriate. So far, robots have done the following tasks in exhibitions

- tour-giving *(Rhino, Minerva, Mobot museum robots, Museomobile, Expo.02)*
- entertainment and animation *(Museum of Communication Berlin, Diligent, Blacky)*
- education *(Mobot museum robots, Museum of Communication Berlin)*
- picture taking *(Expo.02)*
- tele-presence *(Kapros, Tourbot, Webfair)*
- interactive art object *(Expo 2000)*
- demonstrations *(Hermes)*

The projects further differ in the number of robots

- single-robot events: *(Rhino, Mobot museum robots, Minerva, Diligent, Blacky, Kapros, Hermes, Museomobil)*
- multi-robot events: *(Expo 2000 with 72 robots, Expo.02 with 10 robots, Museum of Communication Berlin with 3 robots)*

Finally, the duration of the deployement varies

- several days *(Rhino, Diligent, Blacky)*
- several weeks *(Minerva, Kapros)*
- several months *(Expo 2000, Expo.02, Hermes, Museomobil)*
- several years *(Mobot museum robots, Museum of Communication Berlin)*

---

[1] The project names refer to the papers in these proceedings (by number in the table of contents): *Blacky* [11], *Diligent* [7], *Expo 2000* [15], *Expo.02* [12, 13, 14], *Hermes* [10], *Kapros* [5], *Mobot Museum Robots* [3], *Museum of Communication Berlin* [2], *Museomobil* [8], *Rhino, Minerva, Tourbot, Webfair* [6].

Challenges and Opportunities

For the exhibition maker, automous robots are a media technology with new possibilities, implications and constraints. For the roboticist, robots in exhibition encounter a number of challenges:

- *Difficult, inherently dynamic operating environments.* Exhibitions are by their nature crowded, highly dynamic, cluttered, and partly uncontrollable. This has concrete implications on navigation tasks such as collision avoidance, localization, global path planning, and multi-robot coordination as well as on interaction and interface design.
- *Human-robot interaction is a key technology.* The robot's interaction modalities and the way they are employed are crucial for the acceptance and success of an exhibition robot. The interaction should be at the same time entertaining, efficient, educational, and managable.
- *System integration must be addressed completely*. Autonomy with respect to computation, perception, and energy becomes an issue for long-term installations. Safety of the people in the exhibition, of the environment, and of the robot itself (vandalism) may have to be respected during the design of the robot.
- *The robot has to work (really)*. Constant supervision and manual intervention requires resources and is expensive. High degrees of robustness and reliability are mandatory at least for long-term events.

On the other hand, installing robots in exhibitions provides opportunities:

- *Long-term experiments*. Many researchers state a lack of large-scale long-term experiments in today robotics research. In terms of operating duration, interaction intensity and overall travel distance, exhibitions offer opportunities to validate research results never met in research laboratories.
- *Gentle specification profile.* The success metrics of exhibitions is not as severe as that of industrial products. Suboptimal performance may remain unremarked or is not considered as a total failure. Technical problems or failures can in some cases even be a part of the exhibition's message.
- *A new field of application for mobile robots.* Successful exhibition events with robots concretize the vision of robots as an established exhibition technique and innovative media technology usable for museums, curators, and scenographers.

Topics

The fourteen papers in these proceedings give a comprehensive overview of the current state-of-the-art. They discuss

- Enabling technologies for exhibition robots
- Human-robot interaction with individuals and with crowds
- Navigation in highly cluttered and dynamic environment
- Multi-robot coordination and techniques for visitor flow management
- Hardware and safety-related issues for exhibition robots
- Long-term experiments and results
- Robots as a new media technology
- Impact factor from an exhibition maker's point of view: do exhibitions really profit from robots?
- Economic aspects of exhibition robots
- Educational aspects of exhibition robots

We hope that this workshop stimulates the exchange of know-how and perspectives within and across disciplines. We also invite the attendees to visit the “Robotics” pavilion at the Swiss National Exposition Expo.02 which takes place at the same time. The project demonstrates on a unique scale many aspects of robots in exhibitions.

We wish you a fruitful and inspiring workshop and a nice stay in Lausanne.

The organizers

Kai Oliver Arras and Wolfram Burgard

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Entertainment Robotics: Examples, Key Technologies and Perspectives

Birgit Graf, Oliver Barth
Fraunhofer Institute for Manufacturing Engineering and Automation (IPA)
GPS GmbH
Nobelstr. 12, Stuttgart, Germany
Email: birgit.graf@ipa.fraunhofer.de, ovb@gps-stuttgart.de

Abstract

Based on the successful hardware and software architecture of Care-O-bot [7] [9], a new generation of mobile robots has been designed at Fraunhofer Institute of Manufacturing Engineering and Automation (IPA). Three robots have been created to communicate with and to entertain visitors in a museum. Their tasks include welcoming visitors, leading a guided tour through the museum or playing with a ball. The robots have been running in this museum daily since March 25th 2000 without noteworthy problems. In this article the hardware platform of the robots and the key technologies for applying mobile robots successfully in public environments such as navigation and communication skills, safety concept, and handling are outlined. Further the underlying control software of the robots is described. Finally the application of the robots at the 'Museum für Kommunikation' in Berlin is presented and perspectives for future installations of mobile entertainment robots are given.

Keywords: Mobile Robots, Museum Robots, Software Architecture, Navigation, Safety.

1 Hardware Platform

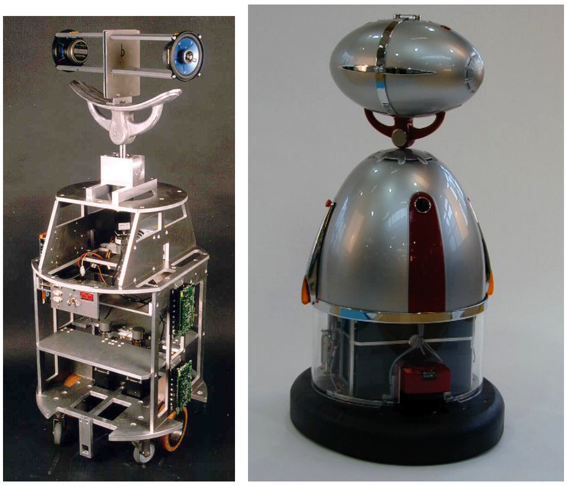

Figure 1. Basic platform and "fully dressed" museum robot
(© Museumsstiftung Post und Kommunikation)

Each vehicle is equipped with two driven wheels (differential drive) including shaft encoders for motion tracking. The robots are able to move at a speed of up to 1.2 m/s. Four castor wheels are further used for keeping the robots upright. A gyroscope is integrated in the robot platforms to track their current orientations.

A 2D laser scanner is attached to the front of each robot. The laser scanner is used for self localization, navigation, and obstacle detection.

Additional safety sensors are a bumper at the bottom of the robots and several infrared sensors which are integrated in the bumper facing upwards. These sensors are used to detect obstacles above the scanning level of the laser scanner. Activating one of the safety sensors as well as pressing either of the emergency stop buttons results in an immediate stop. Besides software restricting the allowed operation area, a magnetic sensor facing towards the ground is used as a secondary system to prevent the robots from leaving their assigned area. This area is bounded by a magnetic band lowered in the ground.

Being equipped with several long lasting batteries the robots are able to move independently for up to ten hours without interruption. For daily operation the robots can be recharged over night.

2 Software Architecture

The control software for the mobile robots is based on the object oriented 'Realtime Framework' and the software library 'Robotics Toolbox', both developed at Fraunhofer IPA. The Robotics Toolbox is an extensive software library, which – in several independent packages – contains modules for implementing all necessary service robot control functions. Furthermore, the use of rapid prototyping methods is being supported by adequate simulation and test environments for all modules.

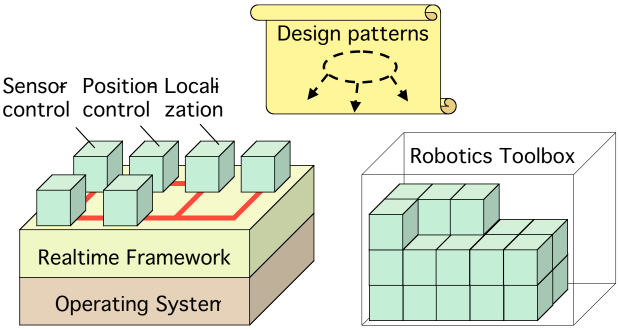

Figure 2. Software Architecture

The Realtime Framework [10] supports the software developer in designing a service robot application. It enables simple and fast integration of single Robotics Toolbox components to an application (Figure 2). The framework provides the structural integration of threads and components (automatic initialisation/deinitialisation, error treatment, etc.). The communication functions of the framework include mechanisms for highly efficient and real-time capable local communication as well as mechanisms for implementation of distributed communication, e.g. for remote diagnosis. The Realtime Framework further presents an abstraction layer for operating system functions and thereby improves the portability of the control software.

3 Robot Features

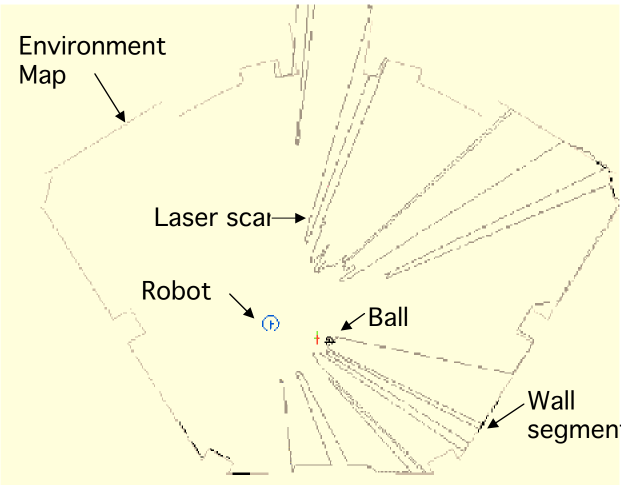

Figure 3. Screenshot of a robot during operation

The following navigation skills have been implemented and tested on the mobile robot platforms:

3.1 Self Localization

Self localization is based on data gained from the wheel encoders (position in x and y) and the gyroscope (robot orientation). However, while using these functions small errors are unavoidable and sum up over time (e.g. 6 degrees of drift per hour for the gyroscope). Therefore the robot's surroundings are modelled in a map (Figure 3). By comparing segments found in the natural environment of the robot (e.g. walls, doors), laser scanner data can be matched to the given map and the robot can correct its position. Information acquired by this method is merged with odometric data using a Kalman filter.

3.2 Robot Motion

Three different types of robot motion planning can be distinguished:

Program controlled navigation: In order to easily specify motion plans for a mobile robot, the "Mobile Vehicle Command Language" (MVCL) has been developed. It allows to write operation programs as simple ASCII files. Operation programs provide the possibility to easily synchronize motion, multimedia and upper axis control commands.

Reactive navigation: In this mode, the current target position for a robot is constantly recalculated in reaction to its environment. Selected objects of a given shape can be detected by the laser scanner (e.g. the ball in Figure 3). The robot then drives to a computed intercepting position.

Preplanned path: If the robot is supposed to move to a certain target position, it will plan the shortest path to this position based on a static map [4].

3.3 Safety concept

One of the most common accidents caused through industrial robots is a person being hit by the robot [1]. For stationary robots the responsibility lies partly with the user – safety measures, as e.g. keeping a certain distance to the robot, must be obeyed. For mobile robots, however, all responsibility lies by the vehicle, therefore the major goal for safe operation should be to prevent a mobile robot from driving into people or from leaving its operation area which might lead to additional incidents as e.g. by a fall down stairs onto people.

For maximum safety a redundant three level safety system has been implemented on Fraunhofer IPA's mobile platforms.

Level one is the laser scanner based collision detection. Whenever an obstacle is detected in the robot's vicinity, the speed of the vehicle is reduced at a degree depending on the distance to the obstacle. If an obstacle or a person gets too close to the vehicle, the robot will stop and wait until the area is clear again.

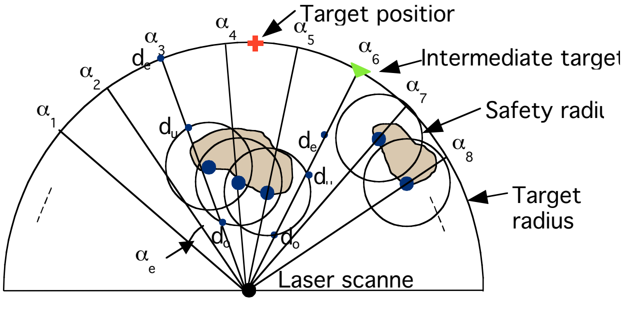

Figure 4. Reactive obstacle avoidance using the "PolarBug" algorithm

The safety module "obstacle detection and surrounding" (Figure 4) is applied in order to avoid unnecessary acceleration and deceleration caused by the collision avoidance. Obstacles detected by the laser scanner are surrounded in advance. The reactive obstacle avoidance algorithm PolarBug [2], based on the VisBug method [6] is being used. This algorithm has been developed especially for obstacle detection with a laser scanner, as well as for fast reaction and navigation in unsteady environments. The major difference to common obstacle avoidance algorithms is the direct processing of the laser scanner data (polar coordinates) which enables a very high efficiency of the algorithm.

Data not only in the planned path of the robot, but all measurements of the laser scanner are evaluated. In case obstacles have been detected between the current position of the robot and a given target, an intermediate position is being calculated which brings the robot around the obstacles as fast as possible. The best free passage is found considering several parameters like e.g. width and depth of passage, deviation of passage from direct line to target and distance of intermediate position to robot and final target position. All relevant factors are joined using a fuzzy logic approach.

Apart from the laser scanner the robot is equipped with a rubber bumper all around the vehicle. Activating the bumper results in an immediate stop. The operation speed of the robot is initially restricted depending on the size of the bumper – so that it can always stop before the bumper is crushed completely. In order to secure the area above the laser scanner, several infrared sensors have been integrated in the bumper facing upwards.

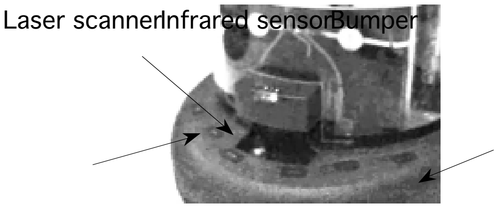

Figure 5. Safety sensors

Thirdly, each robot is equipped with magnetic sensors facing to the ground. They are used as a secondary system to ensure no robot ever leaves its operation area. In the unlikely case of a software failure, by leaving the given operation area and therefore crossing a magnetic band lowered in the ground, an emergency stop will be activated. In addition, each robot is equipped with two emergency stop buttons to deactivate the robots manually.

For applications where the mobile robots move among people in public environments, this safety system has been accepted by the responsible professional association. Furthermore, a CE certification could be acquired for the robots.

3.4 User Interface

Entertainment robots must designed to be used by inexperienced personnel. A joystick with two buttons is the only device necessary to set the robots in operation and to shut them down afterwards. After a robot has been switched on the operator can use the joystick to put the robot in the different start-up modes, such as initial localization and self test. The robot will guide the operator by giving speech output according to its current mode until it starts its default operation mode. For shut down the robot automatically returns to its default rest position before switching itself off.

4 Museum Application

In order to entertain visitors in the recently 'Museum für Kommunikation Berlin' – opened up in March 2000 –with a new technical attraction, three mobile robots have been built and programmed by Fraunhofer IPA [3] [8].

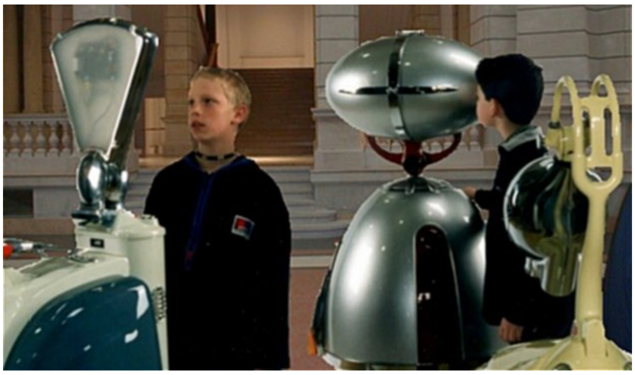

Figure 6. Entertainment robots in the "Museum für Kommunikation" Berlin
(© Museumsstiftung Post und Kommunikation)

4.1 Description of Robots

Each robot has a specific character, expressed through its looks and appearance (driving speed, voice etc.). The robots also differ in what information they give to the museum visitors:

The Inciting: This robot acts as an entertainer. It approaches the visitors and welcomes them to the museum. It moves smooth, but determined at a speed of up to 0.4 m/s. Speech output is further underlined by movement of the robot's head. The robot uses its laser scanner to detect visitors. It looks for features like diameter, shape and distance and then uses fuzzy logic to determine which objects in its surrounding are pairs of legs. The robot distinguish between single persons and groups and uses different sets of welcome phrases for each case. An additional feature is that the robot memorizes the position of persons it has already welcomed for a certain time. During that time it will not welcome people at the memorized positions. Thus it is prevented that the robot welcomes a person several times.

The Instructive: Acting as a guide this robot gives tours in the museum. It moves along straight lines at a speed of 0,3 m/s. The instructive gives explanations about the exhibits of the museum. Moving its head up and down symbolizes the robot looking at the object it is currently talking about. Explanations are further underlined by pictures or video sequences shown on the screen of the robot.

The Twiddling: The child in our "robot family" is, according to its character, unable to speak properly and runs around the museum playing with a large ball. This robot moves rather fast at a speed of up to 0.6 m/s and aims at a ball of a specific size as long as it can detect it. Using its laser scanner it detects the ball by its shape and size, similar to the way The Inciting detects people. This robot can switch between three 'moods'. Depending on the situation it is either happy, grumpy or angry. The 'moods' are expressed by different types of sound output. As long as the robot can detect its ball every now and again it is happy and moves constantly towards it. If it cannot detect the ball for a certain time (for example because a visitor lifted it up) it starts to become grumpy and moves around nervously searching for the ball. If it has not found its ball again after an other period of time it will become angry. The robot then stands still and cries until it detects the ball again.

Apart from performing their standard tasks, the robots are capable to interact with each other as well as with the museum visitors. So if, for example, a robot gets close to one of the others, it will turn towards it to say hello. If The Instructive detects that visitor obstruct its way it will ask them to step aside. If The Twiddling becomes angry, because it cannot find its ball, The Inciting will come to it and ask the visitors to hand the back to The Twiddling.

4.2 Experiences

Since the robots were installed in the museum they travelled more that 1000 kilometres. During all this time no collisions with either visitors or inventory of the museum occurred. The robots also never left their operating area. Thus the robots did at no time present any danger to the visitors of the museum. They usually fulfil their assigned tasks daily without any trouble.

The robots have been well accepted by the visitors of the museum. Children do especially like the ball playing robot. Even children of about 3 years of age enjoy playing with the robot which is with 1.2 meters substantially higher than the children themselves. This proves, that a intuitive interaction with the robots was achieved by IPA's implementation.

Before they were set into operation the robots have been tested in the museum for 2 months. Due to the extensive tests performed during this time the robots' software is now thoroughly debugged and running without any trouble. The only serious hardware problem that occurred was a broken gear axis. The reason was a failure in the material of a commercial gear axis. After months of daily operation a shaft/grain connection became loose on the ball playing robot. This incident occurred on this particular robot, because this one accelerates and decelerates most frequently. The affected connection was modified on all robots.

An inconvenient observation has been made concerning the way visitors of the museum are using the emergency buttons. They tend to press the emergency buttons of the robots for fun. If a button was hit a member of the museum staff has to put the robot back to operation since a key is needed to release the emergency stop. Due to safety regulations the staff members could not be relieved from this duty up to now.

The experiences in the museum show that the implementation of the Fraunhofer IPA can guarantee the following required constraints:

- Elimination of any possible danger for the visitors
- Obstacle detection and avoidance
- Restriction to a given operating area
- Robust design for long operation
- Easy handling for inexperienced personnel
- Operation for up to 10 hours daily

5 Perspectives

Care-O-bot has been designed as a mobile home care system. Based on this platform a group of mobile entertainment robots has been created. Their installation at the 'Museum für Kommunikation' in Berlin proves, that these robots are suited for every day use. Due to the refined way the robots interact with the visitors they are well accepted by them. The positive attitude the visitors develop to mobile robots paves the way for future systems. However, the underlying technological concept is not limited to the given applications. Further functions could be:

- "Personal robot" in private homes ("robotic butler"), robot valet
- Mobil information desk in public areas (shopping malls etc.)
- Safety guard, night watchman
- Robot receptionist in office buildings

Figure 7. Care-O-bot II

Thus development and improvements are going on at Fraunhofer IPA. A new Care-O-bot platform has been build, including a manipulator arm to perform handling tasks (Figure 7).

Generally speaking, the value of robots in entertainment applications depends on the degree of human-machine interaction which can be used. At the moment, robots behaviour is felt to be rather simple, because communication flow is going in one direction only: Machines like tour guides can bring a lot of visual or audible information to humans by display and audio speakers. On the contrary, it is still not possible to talk to machines so that words are properly recognized, not to mention the problems in analysing words and sentences to extract meaningful information. In the long run, to bring a breakthrough to entertainment robotics in widespread applications, input devices like keyboards, buttons or touch screens have to be replaced by audible communication between human and machine.

Another key technology in future applications is manipulation. Having haptic contact to a robots manipulator/ hand is a real sensation for humans, because this kind of interaction is sensed to be very intimate. It is rather easy to imagine some scenarios where haptic interaction is most useful:

- Promotion robots put some give-away articles to visitors.
- Mobile robot servants deliver food and drinks to restaurant visitors.

Unfortunately, haptic interaction has to deal with safety issues. A robot arm that can carry a tablet with food and drinks should be designed for a payload of 2 - 4 kg at least. Taking into consideration 6 degrees of freedom and an arm length of approx. 1m, the manipulators weight will come to be in the range of 20 - 30 kg. It is obvious that such an mechanism could do severe damage to humans if drive or controller malfunctions occur. To bring entertainment robot manipulators into application anyway, some effort is done at the moment:

- Bumpers at the arms hull bring arm motion to a stop when touched
- Sensors like cameras with image processing analyse the robots workspace to prevent contact of robot arm and humans
- Sensors like capacitive sensors or ultrasonic recognize approaching objects to the arms workspace
- Mechanical couplings restrict the torque of robot arm joints to a maximum value

Anyway each solution has its drawbacks, not to mention that there is no guideline to get a certification of the involved institutions at the moment (TÜV and Berufsgenossenschaft in Germany).

6 References

[1] Corke, Peter I: "Safety of advanced robots in human environments. A discussion paper for IARP", 1999.

[2] Graf, B.; Traub, A.; John, D.; Schraft, R.D.: „PolarBug – ein effizienter Algorithmus zur reaktiven Hindernisumfahrung". In Proceedings of AMS 2000.

[3] Graf, B.; Schraft, R.D.; Neugebauer, J.: "A Mobile Robot Platform for Assistance and Entertainment." In Proceedings of ISR-2000, Montreal, pp. 252-253.

[4] Graf, B., Hostalet Wandosell, J. M.: "Flexible Path Planning for Nonholonomic Mobile Robots". In Proceedings of The fourth European workshop on advanced mobile robots (EUROBOT'01), September 19-21, 2001, Lund, Sweden, pp.199-206.

[5] Latombe, J.-C. (1996), Robot Motion Planning, UK: Kluwer Academic Publishers.

[6] Lumelsky, V.J. and Skewis, T. (1990), "Incorporating range sensing in the robot navigation function", IEEE Trans. Sys. Man Cybernet, Vol. 20 No 5, pp. 1058-1068.

[7] Schaeffer, C. and May, T. (1999), "Care-O-bot: A System for Assisting Elderly or Disabled Persons in Home Environments", Proceedings of AAATE-99.

[8] Schraft, R. D.; Graf, B.; Traub, A.; John, D.: "A Mobile Robot Platform for Assistance and Entertainment". In Industrial Robot Journal, Vol. 28, 2001, pp. 29-34.

[9] Schraft, R. D.; Schaeffer, C.; May, T.: "The Concept of a System for Assisting Elderly or Disabled Persons in Home Environments", Proceedings of the 24th IEEE IECON, Vol. 4, Aachen (Germany), 1998.

[10] Traub, A.; Schraft, R. D.: "An Object-Oriented Real-time Framework for Distributed Control Systems", in Proceedings of ICRA-99, pp. 3115-3121, 1999.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

The Mobot Museum Robot Installations: A Five Year Experiment

Illah R. Nourbakhsh
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
illah@ri.cmu.edu

Abstract

This paper describes a long-term project to install and maintain socially interactive, autonomous mobile robots in public spaces. We have deployed four robots over the past five years, accumulating a total operational time exceeding seven years. This document introduces the robots, then focuses on lessons learned from each deployment. Finally, this paper describes how this entire project came to a close, offering a cautionary tale for those who wish to embark on such an effort in the future.

1. Introduction

The history of autonomous mobile robotics research has largely been a story of closely supervised, isolated experiments on platforms which do not last long beyond the end of the experiment. In January 1998, we and others began work on Chips, an autonomous robot intended to be more than a short experiment. Our goal would be to install Chips as a permanent member of the museum staff at the Carnegie Museum of Natural History in Pittsburgh, PA [2].

Shortly thereafter, Mobot, Inc. was incorporated with a charter to improve and extend the Chips technology in a series of robot installations. Following Chips, three more robots have been developed in succession; three of the four operated every day until February 2002. Of the four robots, three are museum robot installations, offering visitors of various exhibit spaces augmented detail regarding the exhibits at hand. Together, these three museum robots have logged more than 2,500 days of operation in *separate*, real-world public spaces.

In striving to deploy autonomous robots in a social niche, we had two high-level goals. First, the robots must be autonomous to the greatest extent possible. Human supervision of a full-time robot is unacceptable. At most, the robots should require only occasional human help, and should request such help explicitly. Even the routine trip to a battery charger should be performed by the robot autonomously.

Secondly, since the robots would be deployed in public, they must have sufficiently rich personalities to achieve compelling and fruitful interaction with humans in their environments. Note that we care not just about ‘compelling’ but about ‘fruitful’ as well. These robots have an educational charter and are therefore justified only if they demonstrate real educational efficacy.

In the end the robots did achieve some measure of educational efficacy as well as long-term robustness, but the social mission of the robots as well as the commercial justification for such robotic endeavours proved to be a challenge beyond our reach. All robots are now off-line, as of April 2002, and so, as the penultimate section of this paper explains, this long-term experiment is now finished.

2. Robot Overview

The three robots described in this paper all share the same basic motive platform (the Nomadic Technologies XR4000 base); the same operating system (RedHat Linux); and the same programming environment (Gnu C++). The first robot of the series, Chips, began work at the Carnegie Museum of Natural History on May 22, 1998 (Fig. 1). Chips operated exclusively in Dinosaur Hall, which contains the large bone collections of *T. Rex* and other massive dinosaurs well as ancillary exhibits focusing on topics such as paleogeology and ancient aquatic life. Chip’s charter was to provide tours in Dinosaur Hall, presenting audiovisual information regarding both the large bone collections as well as the less frequented, smaller exhibits. Until it was taken off-line, Chips operated for just under 4 years, covering a total travel distance exceeding 500 km conservatively.

The second robot, Sweetlips, conducted tours in the Hall of North American Wildlife, also at the Carnegie Museum of Natural History (Fig. 2). This space is comprised of dioramas, where preserved wildlife specimens are shown in naturalistic settings. This portion of the museum has extremely low visitor traffic, so Sweetlip’s charter was to

---

***Figure 1:** Chips the Dinosaur Hall Robot*

attract additional visitors and to then bring the static dioramas to life using high-quality video footage of the wildlife in their natural habitat. Sweetlips operated beginning May 19 1999 and covered a total distance exceeding 185 km autonomously.

***Figure 2:** The Sweetlips robot near a diorama*

The third robot, Joe Historybot, operated in the Atrium of the Heinz History Center (Fig. 3). Its mission was to welcome visitors to this historical museum and to provide both information and a tour of various permanent exhibits placed throughout the atrium. Joe provided historical context in an entertaining multimedia format. The robot also provided tutorials on topics such as speaking English with a Pittsburgh accent. Joe began operations on July 8, 1999 and covered a total distance exceeding 162 km during its total period on-line.

***Figure 3:** Joe Historybot in the Heinz History Center*

3. Lessons Learned

The underlying goals of compelling interaction and maximal autonomy have remained constant throughout the creation of all three robots. However, each succeeding robot was the product of a complete re-design phase based on lessons learned from prior robots. Although some technical aspects remained unchanged, such as the programming language and robot mobility chassis, virtually all else evolved in an effort to improve both the autonomy and interactivity efficacy of each robot.

We are in the unique position of having an established trajectory of real-world interactive social robots. Studying the evolving lessons learned from these installations may prove useful in uncovering information that is valuable to future robot installations. In the following two sections we present such lessons learned, discriminating between the two top-level goals of providing maximal autonomy and producing effective robot-human interaction.

3.1 On Robot Autonomy

The first requirement of a public robot is safety, both for the general public and for the robot itself. At the heart of the matter is the robot's method for avoiding collisions, which must be especially robust, since the robots operate without supervision. It is notable that the collision avoidance code on these robots is by far the least changed code over the course of their creation and installation. The robots use ultrasonic range-finding sensors to detect

***Figure 4:** The fiducial docking landmark for Chips*

obstacles, and move around them reactively, each cycle choosing the appropriate holonomic motion vector to take based strictly on the most recently available sensor data, along with restrictions on how far the robot is allowed to move out of its ideal trajectory.

The obstacle avoidance code is extremely simple, with no explicit mapping or modeling of either the world or the sensors themselves. It is also easy to understand, and because of the lack of internal state, diagnostically transparent [3]. Because of the limited accuracy of sonar at close range, the robots will occasionally become stuck when they approach a wall too closely. Given the infrequency of this failure mode (less than once per week), we feel the increased peace of mind due to convervative motion primitives is worth the price.

Of course there is a great deal more to robot autonomy than safety. A robot must be able to interpret its own behavior, to determine whether or not it is functioning correctly. In order for humans to be confident in the robot's ability to run without supervision, a robot must be able to determine on its own when a failure condition has occured.

Early in the development of the Chips platform, we began using pagers, which the robot was able to signal via electronic mail. The ability to recognize failure and actively request help satisfies near-term requirements for autonomy. Of course the ultimate goal is that the robot never needs to send for help at all, so self-repair becomes a step following self-diagnosis.

Initially, Chips sent for help quickly, giving up as soon as a failure was detected. Soon we began adding diagnostic methods to reset subsystems that were not functioning correctly. This evolved into a general first-level method for diagnostics within our software architecture: each time a task is performed, check the result for validity. If the command failed, then reset the device and try again. Suprisingly, this simple strategy has a commanding effect on a complex robot's failure rate.

In order to achieve true self-reliance, each robot must be able to recharge itself when necessary. This is accomplished using a simple 3D fiducial, aligned with an electrical outlet, that provides both translational and rotational position feedback (Fig. 4). Using this marker, the robots have demonstrated reliable position to an accuracy of 1.5 mm using visual position servoing. The entire docking process, including moving over a distance of 4 m to the outlet and then fine-servoing for the insertion process, takes less than three minutes.

The *retry* method comes into play even in the case of this docking maneuver. If the battery voltage following docking fails to rise, the robot will physically reset by literally backing out of the plug and into the hallway. Then, it will repeat the docking attempt. This policy is effective in most cases because, although the code is deterministic, there is sufficient nondeterminism in the environment that the same software can have dramatically different outcomes when run consecutively.

This general approach is now used for entire classes of robot failure, including but not limited to: battery overcharging and undercharging exceptions; framegrabber anomalies; DVD player errors; encoder value errors; emergency-stop activation errors; etc.

Finally, a critical ingredient for autonomy across all of these mobile robots is the ability to navigate autonomously and with extreme reliability within the space that is served. In all cases, robot navigation is performed through a combination of visual landmark-based navigation and encoder-based position estimation. Furthermore, every environment included a set of allowable travel routes, thus specializing the navigation problem to a route-level travel problem.

For example, Chips made use of a set of high-contrast, high-saturation paper landmarks placed at the end of three of its four travel hallways. As robots were installed in various environments, the fiducial markers' complexity and expressiveness increased, including edge detection of window-wall boundaries and a variety of color and light fiducials. Furthermore, due to varying lighting conditions especially prevalent in the Heinz History Center (due to large, open windows), we added methods to enable tracking multiple fiducials simultaneously as well as "try again" techniques for re-acquisition of lost landmarks.

To summarize, Chips, Sweetlips and Joe Historybot were able to autonomously navigate public spaces for days at a time, charging themselves as they saw fit. This level of autonomy achieved Mean Time Between Failure values of between 72 and 216 hours, and only with great effort would MTBF ever climb beyond such values for any real duration. Failures would eventually become somewhat stochastic, a tyre failing here, and a light bulb failing there. However, a significant achievement with respect to the project as a whole was that, following the first two years of effort, nearly every robot failure was detectable by the robot itself. The days of robot failure unannounced by the robot itself were quickly over in the course of the experiment. Finally, the total autonomous travel distance for the combined set of robots exceeded 840 km.

3.2 On Interaction and Educational Efficacy

Our second requirement was to deploy robots with compelling interactivity. As the science of Human-Robot Interaction (HRI) is in its infancy, so it is not surprising that the robot interaction component was entirely redesigned with each successive deployment.

An interview with the exhibits maintenance staff of any large museum will drive home an important fact: people are basically destructive. Sometimes this purposeful damage is indeed caused by malicious visitors. More frequently, curious individuals who are trying to better understand the exhibits will cause damage. For example, some individuals attempt to push the robot off course physically to see if it will recover. Others will push any large red emergency stop button to see what happens next.

Figure 5: *Chips attracts visitors of all ages*

Also, what attracts people in a public space varies greatly depending on the context of a particular public space. When in an "entertainment" space, such as a museum, people will be curious and attracted by new and unusual exhibits. To that end the physical appearance of a robot is extremely important as a visual hook. But two other characteristics product even better results: motion and awareness. When the robot is in motion, it draws great attention from nearby people. To capitalize on this aspect of human behavior we found it useful to have some robots exhibit limited prosody during delivery of long, static presentations.

But the single most successful way for a robot to attract human interest is for the robot to demonstrate awareness of human presence. Interactions between humans and complex machines are typically initiated by humans. When a robot deliberately faces a person and says "Hello," he or she is almost always both surprised and enthralled.

In contrast to entertainment venues, more utilitarian spaces such as large shopping centers and office buildings elicit far less pronounced reactions. In these spaces people tend to have an agenda; a schedule. They rush about from appointment to appointment and have little time to be side-tracked by new and entertaining machines.

One very important lesson learned from our experimentation is that attracting humans, while itself quite challenging at times, is far easier and very much unrelated to the skill of retention. Museum exhibit designers tend to make their exhibits more interactive, often even taking on the characteristics of a conversation, in an attempt to retain the visitor long enough to successfully pass on an educational nugget. An exhibit may pose a question requiring the visitor to lift a panel or push a button to hear the correct answer, for instance.

We have found that these techniques for audience retention are equally valid for HRI [1]. Chips presents long (two minute) video clips at various locations throughout its tour path. As our robots evolved, so did their level of interactivity. For instance, Sweetlips includes the human observer in the process of choosing a tour theme based on their interests. Joe goes further, answering many different classes of questions and even asking humans several questions in a game-like format.

Because of a robot's particular sensory and effectory strengths, dialogue is multimodal and not necessarily verbal. Thus, while the human may be pushing buttons or using a touch screen, the robot may be responding with spoken words, music, graphics, video, text, physical gestures and motion.

We have learned several lessons from such robotic dialogue design. Firstly, there often will be a crowd of people around the robot rather than a single person. Together with background noise from the environment, this will make it difficult or impossible for some to hear the robot's responses if they are purely verbal. Therefore, all robot responses should be multimodal, including not only written text (e.g. captioning) but also graphics and video content.

Second, long presentations, even movies, are guaranteed to drive audiences away. Instead, short responses combined with interactive questions are most effective at extending the *time on task*. This parallels normal human-human interaction: the best conversations are dialogues, not monologues.

A final lesson learned with respect to HRI involves the psychological effect of creating anthropomorphic robots. There are strong social rules governing appropriate behavior between humans (though these rules vary between cultures and segments of society), and there are other behavior patterns that people follow when interacting with machines and computers. A robot that looks somewhat human and has a rudimentary personality falls somewhere between these modes.

The majority of people treat a robot as part human, part machine, clearly following some modified form of human interaction. Often they will treat the robot like a human by default, dodging its path and verbally responding to it naturally. If they become interested in some features of the robot, or want to investigate how it works, however, they will begin to treat the robot like a machine, ignoring social decorum by refusing to get out of its way and standing rudely in its path to elicit a reaction.

We theorize that humans use whichever social mode is most convenient for their short term goals. Fortunately, people will also often accomodate a robot that behaves in a socially unconventionally manner (were it a human).

A second avenue of exploration involves the use of affection in designing robot behavior. The main reason for a utilitarian robot to display emotion is that humans expect and respond to them in somewhat predictable ways. People have a strong anthroporphism urge and tend to attribute internal state to anything that behaves appropriately. People are also strongly conditioned to react to the emotions displayed by another person. These are powerful tendencies that robots could exploit.

These reactions are entirely behavioral. People cannot discern the true internal state of another human or robot. Their responses are thus entirely dependent upon perceived behavior. Chips and Sweetlips used sophisticated internal mood state machines that would change over the course of the day, affecting the behavior of the robot in the small and in the large. But since visitors to a museum only interact with a robot for a short period of time, the long-term mood shifts were moot.

For this reason, Joe Historybot uses no such internal mood representation and, instead, has a more transparent set of affective reactions to simple stimuli. For instance, stand in front of Joe unflaggingly and it would blurt out, “This isn’t the Parkway and you’re not PennDOT!”

Table 1: *Educational concept questions: success rates before and after robot tours*

Question	<before>	<after>
All dinosaurs lived during the same time period	50%	92%
All dinosaurs were huge animals	50%	72%
Other animals lived on the Earth with dinosaurs	50%	76%
All dinosaurs were carnivorous	48%	80%
All scientists agree on how to put dinosaur bones together	40%	76%
All bones in Dinosaur Hall are real	36%	52%

The eventual goal of each interactive robot is to transfer information to human visitors. To test the educational efficacy of Chips, outside evaluators were invited to conduct an educational study [5]. The formal focus of this project was to answer the question: *Is Chips an effective vehicle to educate visitors in Dinosaur Hall when compared with a docent? Effectiveness is defined as being accessible, educational, entertaining and appealing to a broad range of visitors.*

The evaluators chose two methods for collecting data: Robot Observation studies and Questionnaires. Results of the observation studies and questionnaire forms were analyzed with respect to the four effectiveness objectives identified by the team: accessibility, educational efficacy, entertainment and apeal to a broad audience range.

Quantitative measurements of accessibility provide evidence for a general conclusion, that visitors will tend to staty with Chips for a shorter total duration but may return later, whereas visitors tend to follow a docent for the entire tour loop. The team found that 40% of visitors remained with the docent tour for 30 minutes or more, whereas only 4% of visitors remained with Chips for the same duration. However, 74% of visitors remained with Chips for between 5 and 15 minutes. The most significant differences between Chips and docents based on questionnaire results involved tour speed and sound level. Chips’ overall speed was viewed more favorably while docents were rated as much easier to hear. Interestingly, these results agreed among both adult and youth age groups.

Educational efficacy was measured by asking adults and children knowledge-testing questions both before and after robot-led tours. The questions for adults and the success rates before and after robot tours are shown in Table 1. These results were extremely pleasing in that they establish a quantitative educational efficacy for the robot tour guide.

In summary, the interactivity of our robots has evolved along four axes: engagement, retention, dialogue and anthropomorphic/affective qualities. Although this field of research is quite young, it is already clear that there remainds great plasticity in the human-robot interaction model: human biases and bigotry regarding robots are not yet strong and nor fixed. We have an opportunity, as robot designers, to create not just robot behavior but the default human behavior that will lead to the most fruitful possible human-robot interaction of the future.

4. The End of an Experiment

Although launched with much fanfare, this series of robotic experiments came to an end almost silently. Several factors came together simultaneously to bring about this end. Any museum is in the business of image management. Add to this the extreme conservatism that is necessarily part of running most museums, and you have a formula for shunning unnecessary technological advances. That two robots were able to become docents at the Carnegie Museum of Natural History is itself somewhat miraculous and is due in no small part to the vision of Jay Apt, former astronaut and director of that museum during Chips’ deployment. With his retirement from CMNH, the robots, his brainchildren, were left without a champion among the museum executives.

Indeed, the educational study demonstrated real educational efficacy; but this is merely a necessary and not sufficient reason for a museum such as CMNH to spend more than $10,000 each year paying for labor and parts (primarily new high-quality batteries every 3 months) to keep full-time robots running.

The second deciding factor was purely economic. The museum market is small, with only a handful of museums that are large enough to purchase the hardware and audio-visual content for a $200,000 robot system. Taken together with the extremely slow sales cycle for such decisions in the museum market, any company entering this market would have to front-end costs for several years in order to build a sales pipeline. Mobot, Inc. began this process and was able to generate several orders; however, the sort of funding required for a 4-year pipeline-building process was nary available in the post-dot com slump and, therefore, Mobot was often relegated to odd jobs in order to make recoup its monthly expenditures. Such shifting focuses cannot take a small company to success, and so the economic pall that cast a gloom over the internet extended to Mobot as well.

Finally, the ability of people to adjust to just about any situation should not be underestimated. Directors and museum staff will grow accustomed, even to a 300 pound robot moving in their midst daily. After a while, the robot is no longer a curiosity, and its removal is only a logistical relief. Without an active champion, there is really no reason for a piece of high technology to stay in a staid museum, and so the decision will eventually be made: remove the robot and thereby cut superfluous costs. Only a robot that is truly dynamic and expressive, changing every day, can overcome this barrier. That is, unfortunately, close to an AI-hard level of behavioral complexity.

5. Conclusions

Over the course of the past five years, we have built three museum tour guide robots that have each interacted on a daily basis with the public, autonomously and without direct human supervision. While this has been done before [4,6,7] our robots are unique in their completely unsupervised free-roaming obstacle avoidance, and in their mission to entertain and inform the generally public in documented, educationally effective ways. We have learned many useful lessons in attempting to meet the above challenges. Perhaps the most striking is that it is indeed possible to deploy robots like these in public places for long stretches of time. The robots are now off-line, after almost five years of operation, and so this experiment is now finished.

Acknowledgements

This workshop paper draws significant sections directly from an earlier paper by Clayton Kunz, Thomas Willeke and myself. The original Natural History museum robot concept is due to Jay Apt and Red Whittaker. The following companies provided invaluable help throughout these projects: Nomadic Technologies, Magic Lantern, Maya Design, Redzone Robotics. Thanks also to co-creators of these robots: Judi Bobenage, Ben and Eva Carter, Vinton Coffman, Sebastien Grange, Iacovos Koumas, Ronald Lutz, Roland Meyer, Carolyn O’Brien, Todd Simonds, Alvaro Soto and David White.

References

[1] Nielsen, J. *Usability Engineering.* Boston, MA: Academic Press, 1993.

[2] Nourbakhsh, I., Bobenage, J., Grange, S., Lutz, Rl, Meyer, R. and Soto, A. An Affective Mobile Educator with a Full-Time Job. *Artificial Intelligence*. 114 (1-2), pp. 95-124. October 1999.

[3] Nourbakhsh, I. Property Mapping: A simple technique for mobile robot programming. In *Proceedings of AAAI 2000.* July 2000.

[4] Pyxis Corporation. www.pyxis.com.

[5] Roehrig, M. and Stockdale, M. *An Interactive Mobile Educator in the Carnegie Museum of Natural History’s Dinosaur Hall: Measuring the Effectiveness of Chips as an Educational Tool in the Museum Setting.* Program Evaluation, Carnegie Mellon University and Carnegie Institute of Art.

[6] Sarcos Corporation. www.sarcos.com

[7] Thrun, S., Beetz, M., Bennewitz, M., Burgard, W., Cremers, A.B., Dellaert, F., Fox, D., Haehnel, D., Rosenberg, C., Roy, N., Schulte, J., and Schulz, D. Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva. *International Journal of Robotics Research,* 19(11):972-999, 2000.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Online Robot Teacher Kits for Museum Field Trips

Gerard McKee

The University of Reading, Department of Computer Science, Reading, UK

Gerard.T.McKee@Reading.ac.uk

Abstract

Online robots offer a novel and interesting tool for exploring museums. They offer the potential to bring people into contact with exhibits at remote locations in practically no more than the time it takes to start up a web browser. The creation of an interesting and useful experience for the user depends not only on the robotics technology itself but also on the way this technology is integrated with existing museum facilites and newer multimedia technologies to deliver a comprehensive service to users. This paper explores one particular application of online robots, not as a technology solely in itself but as one tool in a range of facilities, to create such an experience. The application specifically demands a wide range of facilities, since its purpose is pedagogical. The application in question is the preparatory support, in the form of teacher kits, for field trips to museums. The ideas can be generalised to other scenarios, particularly for parents preparing for museums visits. The paper discusses three aspects of the application. First, ways for preparing for trips to museums. Second, the potential enhancements provided by online robots. Third, the implications of the latter for the wider online and in-musuem facilities for getting the full benefits of this technology. The paper concludes with a proposal-for-discussion for an educational challenge for online robot systems in museums.

Introduction

Online robots have now a recognized presence both within and outside the robotics community [1,2]. Online visits to museums are one application of these systems [3] and opens the opportunity to apply other forms of robotics and vision technology to enhance the user experience and awareness of the remote envirnment [4]. It is important at this stage to begin seriously thinking about the design and development of comprehensive environments to support a range of applications, so that the utility of online robots can be illustrated beyond proof of concept demonstrations. One of the areas of application is for educational field trips to museums. This paper poses the question of how best to exploit this online robot technology for such applications. The field trip scenarios is particularly interesting since it raises a number of issues beyond simple demonstration or interaction with a robot moving around a museum environment. In particular, it poses the question of how the field trip fits into the wider education of the students, how the online robot systems can be incorporated within this setting to support the teacher, and how in turn the museum itself, the target of the field-trip, can integrate online robot systems with its educational infrastructure. Thus it brings together matters of pedagogy, matters of technology, and matters of management. This paper aims to address these issues in a preliminary manner. The goal is to open a discussion among roboticists, educators and museum personnel about the way in which online robot technology can be exploited in this application scenario.

The remainder of the paper is organised as follows. The following section provides a discussion of modes and methods for preparing for a field trip, exploring their advantages and disadvantages. A set of requirements that address the potential for the involvement of online robots in preparatory work for field trips is presented. Section 3 discusses the role that online robot systems can play in this preparatory work, particularly what they bring to the scenario and how this integrates into an overall toolkit. Section 4 explores the implications of incorporating online robot systems in museum environments in order to support the role(s) identified in section 2 and to create a comprehensive toolkit for teachers. Finally, section 5 provides a summary and conclusions.

2. Preparing For A Field Trip

The purpose of this section is to enumerate some of the ways in which a teacher can prepare for a field trip. To focus the diccussion it is at least assumed that the teacher wishes to visit the museum for the primary objective of viewing a subset of the museums exhibits, and those exhibits are related in some way to the taught material that has been presented to the students or that will be presented following the trip. We at least assume a model whereby the teacher has some preparatory kit that they can use to prepare themselves for the visit, so that they can make the trip as beneficial and productive as possible for the students. Our question is what items could go into that kit.

We can start with the notion of an empty kit. Empty in the sense that the visit is made on spec. The museum is assumed to belong to a category that should have exhibits relevant to the taught material, but the teacher does not know what specific exhibits. This is a poor basis for a field trip, but possibly a good basis for a general school outing. We can take this scenario as our baseline. In general, however, we expect the teacher to have some background knowledge of the museum exhibits.

One of the methods for getting information about the exhibits, and the museum in general, is to talk either to museum personnel directly or to colleagues or friends who have visited the museum. The former, in particular, can provide information leaflets, and can inform the teacher of special facilities and tours for student groups. This same information could also be obtained from the Museum web site – we assume that there is such a web site in place. However, this web-based information, and even just the discussion with the museum personnel, is second hand. There is no direct experience of the museum prior to the school trip.

A better approach would indeed be for the teacher to personally visit the museum prior to the student trip. In this way the teacher would gain first-hand information about the exhibits and could take the time, while at the museum, to organise an itinerary and even to preview it. This approach would also allow the teacher to gather more background informaton about the exhibits, possibly from the museum booklets, and to identify areas or gaps in their knoweldge that can be overcome prior to the student trip by referenceing the relevant information in books or online information via the Internet. This seems an ideal scenario, but is not workable for museums that are some distance away, nor is it usually possible to make such a preparatory trip shortly prior to the school trip. Indeed, it may be some years since the teacher actually visited the museum in person.

Failing the opportunity to visit the museum personally the teacher may makeuse of online tours. These tours typically involves the online visitor being led through a guided tour of the musuem by way of diagrams and images. This can be to a larger or lesser extent interactive. A simple format involves one or more plan-view maps of the environment offering a number of icons that can be selected with a pointer. When selected, an image taken from the specified location will appear either in the current browser window or a separate window. Dynamic web pages and the use of Java applets can often provide a more dynamic feel to the tour, incorporating small video sequences and perhaps three-dimensional models. However, the experience is still very passive. For example, it is often difficult to get a good sense of localisation and orientation within the real environment from the snapshot images. The result is that in cases where these environments are not put together well the full excitement of a visit to the museum is not conveyed. In short, since the experience can be so disconnected and passive, it can be very difficult for the teacher to calibrate themselves to the museum environment.

The above set of cases do not exhaust the range of facilities that mesuems may offer to teachers to support field trips, nor the full range of educational support for such trips. However, it does enough to point to ways in which online robot systems could be employed as an additional resource for preparing field trips. The most important conclusion in this regard is the opportunity for the teachers to be involved more actively in the preparation process by gaining experience through prior visits to the museum site. Since it will often not be realistic to have the teacher personally visit the museum, online robots provide a method for them to get something close to a similar experience. By at least having the ability to drive a robot around the museum site, using vision of course as the main mode of sensory feedback, they can calibrate themselves to the museum environment. By actively exploring the remote environment in this way they will more easily build a model of the museum environment, allowing them to relatively easily localise and orient themselve within the environment when they finally do visit it with the students. This, however, can only be the starting point. The challenge for the online robots community, and the museum, is to not only allow them to drive a robot around the museum, but also to integrate that activity with a set of other tools that will provide the teacher with a rich source of background material to support the educational experience that they ultimately wish to provide for the students. The essential objective must be to place the school teacher in control of the creative process involved in putting this educational experience together.

3. Online Robot Technology

This section aims to provide some insight into the role that online robots can play in the prepration of school trips by students. A simple, though challenging, scenario is described and then extended to incorporate a more flexible, richer, toolkit. The starting scenario is a passive, automated online robot tour guide. This comprises a robot system that takes a well defined tour of the exhibits. In fact one can envision a number of such automated tours emphasising different subsets drawn from the full set of exhibits. One can envision the robot following a route that takes it successively to well-defined points from which it can view each of the exibibits. There may be a number of points form where each single exhibit can be viewed, and the robot can take up these points in a well defined sequence. In addition, from each point it may visually inspect, or view, different aspects of the exhibit. These viewing positions and the corresponding camera parameters (zoom, focus) can be scripted a priori as per the routes. This basic scenario is a challenge in itself, involving path planning, navigation and localisation. And although the viewing may be via one or more cameras, the visual sensor data may play only a minor role in the navigation and localisation of the robot. The video signal, however, can be relayed to the Internet audience to form primary modality in their experience of the museum tour. The tour in this basic form essentially comprises stationary points and transitations between these points. Each of these can be tagged with text or audio descriptions that can be relayed to the online visitor via audible speech, text or graphics, or a combination of these and more. The set of tours could be run daily and some of the more popular can be repeated a number of times on the same day. As long as the robot doesn't move aggressively about the environment, the robot should not interfere with the onsite visitors.

This scenario, although basic in terms of its interactivity, presents some significant challenges. It is within reach of current robotics, Internet and multimedia technology. We can call this the passive robot tour guide. The teacher, similar to general users, will be able to join a tour at the start or anywhere enroute. The tour, in this form, will provide them with a good preparation for the school trip. However, what if they do not want to follow any of the prescribed tours, or indeed, if at some point they would like the robot to stop so that they can dwell for some moments on a particular exhibit. In particular, what if they want to delve deeper into the background for that exhibit, picking out finer details, its historical roots, the people behind the content it displays, and so forth. The passive tour guide will in this regard be quite limiting. Indeed, what if the teacher wants to put a tour together that does not match an existing tour?

If we consider the latter first, one approach would be to present the teacher with a palette of exhibits that they can choose to visit online. The selected exhibits would be passed to a tour planner that can assemble a route map for the robot to follow. The tagged information, prepackaged, can be attached automatically to the appropriate states and transitions, and replayed as per above during the live tour. This will at least allow the teacher to preview the tour. During the selection of the exhibits the teacher can also delve into the background for each exhibit. In this way, the teacher can put together a scheduled tour in a more liesurely and thoughtful manner. Indeed, when the teacher actually visits the site we can envision the robot leading the tour of the exhibits, following this pre-planned tour. This in itself could be a valuable service that museums could offer teachers.

If we were to stop at this point we can see what appear to be significant benefits for the teacher from preparation through to implementation of the school trip. However, one can envision a still further level of interaction whereby the teacher can direct, or even drive, the robot around. What are the benefits of this additional capability over the above? In the scenario depicted above the teacher is at the mercy of the tagged information provided by the musuem and, indeed, the preplanned exhibit states and transitions. There is also, still, a further lingering 'passive' element, in that the teacher has to pre-plan the tour based on a given palette of exhibits. We can see three further extensions of that model. The first, as suggested, is the provision for the teacher to drive the system around. The main benefit of this is the ability for the teacher to craft their own tour. They can select individual exhibits to visit, script the inspection of each exhibit to their own requirements, and attach the tagged information appropriately. Not all teachers will avail themselves of these facilities, but they may at least be more prepared to avial themselves of tours pre-crafted by other teachers, as an alternative to the more mechanical tours that may otherwise be available if teachers were not actively invovled in the enterprise. Hence, we can envision the accummulation of a library of educational tours, building towards a rich educational repository.

We can envision this scenario working in the context where the teacher first creates a ‘draft’ tour and then perturbs, or knocks it into a shape, more in tune with his needs. The second extension, however, would offer the teacher the ability to craft the script on the fly. This essentially means that the teacher can use the robot to explore the museum as they please, taking in the individual exhibits in some preliminary sequence and then integrate a selected subset into a tour. We can describe the robot in this context as an ‘online robot browser’ since it effectively allows the teacher to browse the museum exhibits in preparation to and during the construction of the required tour. As before, the tour can be previewed and further edits can be made during this preview stage. The teacher may then avail of the tour in three ways. In the first, the teacher can lead the school children through the tour, dispensing with any further participation by the robot. In this context the preparation work with the robot will have helped the teacher ‘actively’ script the required tour and relevant background material. In short, it will have provided an educational experience for the teacher, who can then pass on that experience directly. In the second, the teacher can allow the robot to lead the tour. This provides the opportunity for the teach to keep a watchful eye on the overal progress of the tour and the children’s participation. In the third, the robot can accompany the tour. This can be beneficial in a number of ways, most notably the local provision of additional visual aids. For example, to present short video sequences. Each of these scenarios places different requirements on the robot system.

The third extension follows from the previous two. It involves the provision of online tools which allow the teacher to create new material that they can tag onto the tour. This appears to be a straighforward appeal for compositon tools, and indeed it is. However, the tools must provide not only composition and editorial support, but also access to a wide range of resources. Effectively the aim is to allow the teacher to draw in material from a wide variety of educational resources. The provision of these to the robot system in the context of the tour must then be integrated with the facilities at the museum. Hence, there is a major technology integration challenge that is directed more at the musuem resources than at the browser interface. This is an issue that needs study.

In summary, online robot facilites within a musuem environment can provide the basis for a rich creative and educational experience for the teacher, which can then be passed on in ways to the students. This is an important model for education that fits with the aspirations of educators. The teacher remains pivotal throughout, even when they put themselves in the background. In short, onlibe robots allow the teacher to exercise control from their point of view while creating an experience of the student which to the latter appears as an open ended, unstructured experience. The whole aim of preparation is not, ultimately, to script, but to be prepared.

4. Creating the Museum Field Trip Toolkit

This section will look briefly at the integration of mobile robot systems within museum environments to support the type of educational toolkit described above. The first point to make, however, is to emphasis the role of online robots as one component in a range of facilities for supporting the teacher. One of the challenges is to identify how the online robots can complement these other tools, which include leaflets, booklets, and online web resources, including the museum’s own web site. The picture presented in the previous sections, for example, needs to be verified. Will there by unanticipated ways in which online robot systems may be used? In exploring their integration with these other tools, will they motivate new forms of tools unconnected with either? These are open questions. In short, how will the online robots affect, change and in particular enhance, the preparatory toolkits that museums, and indeed schools, offer to teachers.

The immediate prospect of installing online robots in museum environments creates paticular technological problems, which are indeed being addressed. These include the networking and computing structures resident in the museum. It will, for example, include in many cases a partial enhancement of the environment with video technology so that one can not only see through the eyes of the robot, but also see the robot. Wireless technology will have an important role to play in creating very flexible, and reconfigurable, intelligent environments in which to allow the robot to operate and to monitor its operations. There are on-the-ground safety issues as well that need to be addressed in the construction of the robot system, since they will be sharing a space with human beings. The robots, for example, must in some way follow the general temper of the environment. In an art-like exhibit area they should perhaps move about slowly, quietly. In a more science-oriented exhibit area they should perhaps be noisy and moving around ‘energetically’. The teachers must be able to interact with the robot at some point, and for a period of time have sole control of the robot system; other teachers will have to wait their turn. Perhaps there is an opportunity here for teachers to support each other, probably via the medium of a chat room. When in control of the robot they must also be aided by online systems that support visualisation of the robot’s environment and have controls, both in software and hardware, that prevent damage to the robot and to the exhibits. Their task won’t simply be to move about and avoid objects, but to get to exhibit locations. To what extent, then, will the teacher really need to ‘drive’ the robot or ‘select’ the target location and let the robot find its own way. In other words, how far does the teacher really need to be in control. These are issues currently being addressed in research. Video tracking technology might have an important role to play here. The robots will have as well to avoid unintended contact with humans. They must not just avoid humans, but do so in an unobtrusive manner. They must, effectively, blend in.

Beyond these on-the-ground features, the museum must support, either locally or remotely, a repository from which background material can be retrieved, and into which teachers can deposit additional materials. A cornerstone of the facility must be a set of tools that allow the teachers to craft their tours. These facilities require good support in networking and multimedia technology. It is important as well to determine to what extent these facilites are an annex of the museum’s computing framework or are seamlessly integrated with its existing facilities. Indeed, to what extent can online robots motivate a fresh look at the development of museum-wide networking and multimedia technology. In sum, there are a whole set of issues here that have been addressed only tentatively within the context of exploiting online robot systems in museum environments. Much needs to be done.

This is indeed a good point to raise the possibility of activities which can allow us to explore the way in which online robots can be used in museum environments. One such possibility is to develop an educational challenge for online robots in museums. The idea would in general be for student projects aimed at developing museum robots. The goal would be to assemble a set of online museum robot case studies. These could provide a useful preliminary method for calibrating the online robot requirements and opportunities for these museum applications. Support for teaching, as layed out above, is one category of application scenario.

5. Summary and Conclusions

This paper has explored an important application of online robot systems, specifically their use as an educational tool in the preparation of school trips to museums. Online robot systems provide the teacher with the opportunity to assemble and preview an itinerary of musuem exhibits for a school trip. When integrated with other tools, including local and online educational resources, the teacher has the materials available to create an interesting, yet open, educational experience for the students. The development and ingetration of online robots, and the exploitation of robotics, networking and multimedia technology, requires that we understand the needs of the teachers. The material presented in this paper is all to brief a snapshop of percieved needs, and are probably only a shadow of real needs. Studies are required to understand the latter. Finally, the online robots community should aim to develop a number of educational challenges that motivate the development of ‘museum robots’.

References

[1] K. Goldberg, The Robot in the Garden, *The MIT Press*, 2001.

[2] K. Goldberg and R. Seigwart, Beyond Webcams: An Introduction to Online Robots, *The MIT Press*, 2002.

[3] J. A. Fryer, Remote-control experiments using a networked robot, in *SPIE Newsletter on Robotics and Machine Perception, Special Issue on Networked Robotics*, Vol. 5, Issue 1, 1996, p. 12.

[4] D. Schultz, M. Moors, W. Burgard, A. B. Cremers, A Statistical Approach to Tracking Multiple Moving People with a Mobile Robot and its Application to Improved Tele-Presence, in *ICRA 2002 Workshop on Educational Applications of Online Robots*, Washington, USA, 2002.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Mobile Robots in Art Museum for Remote Appreciation via Internet

Shoichi Maeyama,

Osaka Electro-Communication University
Neyagawa Osaka 572-8530, Japan
E-mail: maeyama@isc.osakac.ac.jp

Shin'ichi Yuta, Akira Harada

University of Tsukuba
Tsukuba Ibaraki 305-8573, Japan
yuta@roboken.esys.tsukuba.ac.jp

Abstract

We report mobile robots moving around an art museum in general public. Mobile robots are physical agents in our remote viewing system via Internet. Mobile robots serve their sensing and action as avatars for remote viewers, who are ordinary persons on the WWW. This system is a tool to obtain human behaviors in an art museum for the KANSEI special project. We discuss experimental results at the Tsukuba art museum and lessons learned from trials working the robots in the art museum as public space.

1 Introduction

Recently, popular technical keywords are IT: Information Technology, Internet and robotics. There are some researches connecting Internet with robots as a physical agent [1]. We are also working for the project related with these keywords. In our "KANSEI[1] special project", we are developing remote viewing system of an art museum using a mobile robot via Internet with a WWW browser. Our previous system has been already demonstrated at the IROS2000 conference[2].

At first, we briefly introduce our KANSEI special project. Then, we explain our remote viewing system by mobile robots in the art museum. Finally, we discuss that effectiveness of our system, open problem on mobile robot in art museum and possibility in new application of robotics and IT.

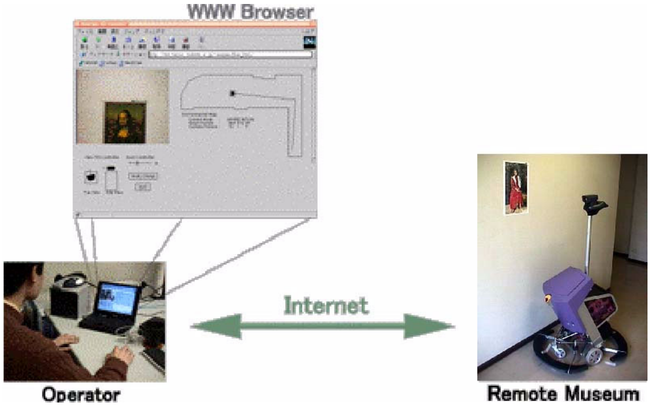

Figure 1: Remote viewing of an art museum

2 Background - KANSEI special project

The University of Tsukuba has promoted a 5-year (1997.4-2002.3) research plan about the special project on "Modeling the Evaluation Structure of KANSEI." This project aims to analyze the appreciation attitude of human interest in works of art, and ultimately to construct a model for the evaluation structure of KANSEI. This project is unique in that the system employs a mobile robot in an art museum, which works as an eye for people remote from the museum [3]. The robot walks

---

[1] KANSEI is an originally Japanese word, which means human sensibility, emotion, feeling, and something difficult to explain by logical reasons.

around the museum in place of the remote visitor, who can move the TV camera on the robot by remote control and view anything in the museum. The above project tackles a number of challenging problems, including those of remotely controlling a mobile robot via the Internet and the construction of a robot that can be interactively operated in an actual museum without causing problems to local visitors. The goal of our challenge is Figure 1, a remote viewing system that enables ordinary people at home or in the office to remotely view works of art in a museum by manipulating the vision of the robot using an ordinary personal computer connected to the Internet.

In this project, we have initially utilized the system as an experimental tool, recorded the positions and postures of the robot working as an avatar, obtained computer images through the robot, and gained an understanding of the behavior of remote visitors. It is not that the robot moves around independently delivering images, but that its behavior is controlled interactively by the remote visitor who is able to view images in real time on the terminal. The data acquisition method of this evaluation system has many advantages over other conventional methods. In the first place, quantitative viewing behavior can be assessed without placing any particular requirements on the viewer. In the second place, remote viewing is accomplished via the Internet using a Web browser, resulting in easy acquisition of an unspecified number of the general public as viewers. In the third place, the remote visitor can appreciate authentic art, which cannot be compared to copies held in a virtual museum. There is also no need to make 3D models of a virtual museum: our system responds instantly to changes in layout in the actual museum. Finally, in the fourth place, it is possible to survey the interactions between the remote visitors and the local visitors while the museum is open to the public. Comparison of the interactions that take place when the museum is closed to the public may enable research on the influence of local visitors on the remote visitor.

Our remote viewing system thus makes it easy to acquire behavioral data on remote visitors. The realization of problem-free robot operation, however, is technically difficult, for the following reasons.

1. Development of the avatar robot
2. Operation of a robot in the Internet environment, where there are random time-delays in data transmission
3. Operational interface using a Web browser as the standard GUI

The aim of our study is to solve technical problems associated with the above points by robotics.

3 Remote viewing on the Web

3.1 Basic strategy

What we want to obtain is a robotic device (avatar) that can move around a museum and can be manipulated by people, who are in fact in remote places, looking at live images on their terminals. However, the Internet of today is not suited for quick delivery of live dynamic images. Also, using these images, it is quite difficult to give the viewer a feeling of being in an actual place due to the narrow bandwidth and limited human interface. In order to make feasible the system's use from a range of connection terminals, the images must at present be sent in slide-show form. In our study, we aim at completion of a basic teleoperation system using live static images, which can be operated as easily as possible by ordinary people. Concerning basic teleoperation systems, an autonomous mobile robot that is directed by spoken instructions has already been built. However, our view is that it is important to present an interactively decision-making environment for the remote visitor each time the live image is switched to a different new one. This is because we also wish to store data on how the remote visitor responds to the switched image, and makes each of their decisions. Thus we have followed the line of developing interactive operations based on static but live images. We have confined the autonomous properties of the robot to self-protective movements such as detection of obstacles and avoidance of collision. Therefore, an absolute position pointing method was chosen as the operational method since it is able to cope with random time-delays.

Our basic strategy for the construction of a remote viewing system can, therefore, be summed up as follows.

1. Operation via absolute position pointing, which is reasonably immune to random time-delays.

2. Autonomous self-protection behavior for the robot-self.

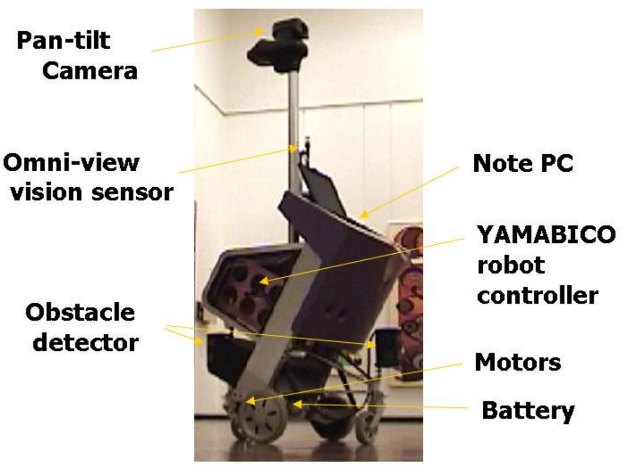

Figure 2: Mobile robotic avatar "**KAPROS** "

3.2 Mobile robotic avatar

We first describe the development of a robot as an avatar. Our avatar does not have to be in perfect humanoid shape. The avatar for our remote viewing system is equipped with limited functions that substitute for human eyes and feet. As for eyesight, humans are endowed with stereo color vision. However, a single color TV camera will be sufficient for the avatar, since the remote visitor will be employing a conventional computer display without any additional equipment. As for movement, human beings are bipedal, but since the floors of museums are usually flat, we chose a wheel drive mechanism for the avatar instead of feet for better moving efficiency. Further, the avatar must co-exist with actual humans in the museum, so the moving speed of the avatar must be equal to the normal walking speed of a human in order not to disturb normal visitors. This is one of the fundamental requisites for the avatar. In respect to scale, the avatar must be able to pass through corridors in the same way as ordinary visitors. Considering these requisites, we have developed a novel robot, which we call "**KAPROS** ", as the avatar (Figure 2). We assumed viewing of ordinary sized pictures exhibited at a height of 140 cm, and set up a TV camera on the avatar at that height.

The following requirements were set when determining the system composition for driving the robot.

1. Autonomous behavior control of the mobile robot

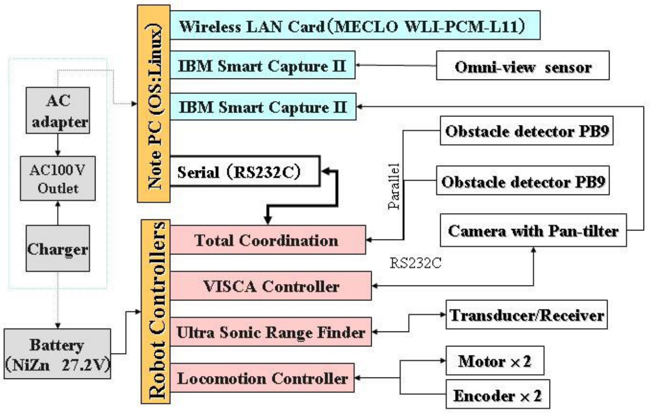

Figure 3: System configuration of the **KAPROS**

2. Wireless communication with the server and the other robots
3. Processing of multimedia information such as images and voices
4. Interface installed on the robot for program development and execution

Figure 3 illustrates the overall design of the controller in the **KAPROS** . It is roughly divided into two parts: one part comprising a notebook PC (IBM ThinkPad 235), and one part comprising the YAMABICO robot controllers[4]. The former part takes charge of communication with the server, the human interface, and the processing of multimedia information (the capturing and display of images and voices). The latter part takes charge of information processing which requires feedback within a very short time period, for example, control of robot movement, recognition of position, detection of obstacles, etc. Communication between the notebook PC and the YAMABICO robot controller is executed based on the RS232C. Since the robot is also equipped with a wireless LAN, it can communicate with the server through an IP connection. In addition, the programming environment for development of the YAMABICO robot controller is also installed on the PC. Concerning the sensor, the robot is equipped with ultrasonic range sensors and obstacle detection sensor "PB9" produced by HOKUYOU AUTOMATIC CO.,LTD. for safer behavior control of the robot. Two TV cameras are also mounted. One is pan, tilt and zoom controllable camera with normal lens for appreciating art works. The other is omni-view sensor, which can capture 360[deg] scene in one image, to give helpful information for moving around without collision through remote control by ordinary persons.

3.3 Tele-driving method and the GUI

This subsection is concerned with the operational method of the robot **KAPROS** working inside a museum. The robot is manipulated by remote control via the Internet using a Web browser. The operation on the GUI as seen in the Figure 4 is roughly divided into the operation for movement and the operation of TV camera for viewing the works of art.

The most difficult problem associated with the tele-operation of a mobile robot via the Internet is the random time-delays, which are likely to occur while issuing operational instructions. More specifically, the problem resides in the uncertainty of time prediction from one operation to the next operation. For this reason, the authors have designed the driving method of the robot by pointing the absolute position in the museum site. Once determining the unique position, this information does not change with elapse of time. We designed the robot guiding system from the present position to the place of destination based on positional information given a short time intervals so that the robot can determine its execution autonomously. However, the destinations towards which the robot would be directed had to be observable in the monitor screen in order to secure interactive operation with the panoramic live image generated by omni-view sensor. The actual interface was designed so as to indicate the ground position on the floor in the live image. This allows easy and intuitive operation by non-specialists [5], as well as interactive operation in response to live images at the same time. The visual information can be converted to instructions in accordance with the absolute position based on the positional data of the robot at the time when the data was programmed in, assuming that the ground seen from the pixel position of the live image is flat and level. Besides the indicated live image, the system manages the *Robot State Data* whichi includes time, own position, sensor data, file name of captured image and so on, corresponeding to each image.

Figure 4: The GUI for remote viewing via Internet with a WWW browser

To appreciate art works, remote visitors only click on the normal image indicated in the upper area on the GUI shown in Figure 4. Then, the camera posture is automatically controlled by pan-tilter table of camera for locating the clicked position to the center of the next live image. Zoom level can be also changed by click of the icon.

Small adjustments of each parameter can be requested by click on each icon.

3.4 Autonomous behavior of the robot

To cope with random time delay, some motions must be done by autonomous functions of the robot. We used the position based tele-operation scheme. Therefore, the motion control related with dynamics such as acceleration and velocity is autonomously controlled by robot-self. And, obstacle avoidance is also managed by robot-self. The obstacles are detected by the obstacle detector PB9. The PB9 can set the detectable area. We set the two areas. The first one (far area) is that the robot can rotate (spin) but not go forward. The second one (near area) can not permit both motions. In the regular cases, the robot detects the obstacles in the first area at first and the robot stops own motion immediately. Then, remote visitor can change the moving direction by rotation of the robot body.

Figure 5: Photograph of the **KAPROS** robot at the Tsukuba art museum

4 Trials at the art museum

4.1 Connecting to the art museum

Finally, we connect to the Internet. However, some art museums do not have a facility of the Ethernet for connecting to the Internet. So, we set up the route to the Internet from the art museum by ISDN. Communication between the robot and the dial-up router is done by the wireless LAN. Connection from the museum to the web server in the University of Tsukuba is achieved by ISDN through dial-up routers. Then, the web server connects our remote viewing system to the Internet. Now, a temporal ISDN phone can be available anywhere in not expensive prices. The web server is permanently placed at the University of Tsukuba. Therefore, we can connect our system to the Internet at almost all museums in less trouble.

4.2 Experiments and demonstrations at the Tsukuba art museum

We have constructed a remote viewing system mentioned the section 3 and have carried out many experiments on seven exhibitions of art and design at the Tsukuba art museum since February of 2000. Figure 5 shows that the first experiments on the graduation exhibition for students of art school in the University of Tsukuba (Feb. 14 - Mar. 3, 2000). On the 2nd of November in 2000, we also demonstrated the remote appreciation from the IROS'00 conference site in our presentation of the regular session. Furthermore, we tried to appreciate 3D art works as well as 2D pictures in the exhibition of statues and sketches (July to August, 2001) shown in Figure 6.

Figure 6: Remote viewing 3D art works by using the **KAPROS** robot

4.3 Experimental results

Our remote viewing system was built through overcoming technical challenges as follows:

1. Instruction by absolute position pointing method to cope with random time delay
2. Autonomous execution of sensor reactive behavior and motion control related with dynamics
3. Updating live images without blinking by a refined JAVA programming
4. Firewall-free command sending through CGI
5. Almost all museum connectable with the WWW server of the University of Tsukuba via ISDN.
6. Appearence design of the robot suitable for an art museum

On the above points, we could confirm the technical feasibility of our system and save data about appreciation attitude of the remote visitors. However, the system remains some problems such as slow update cycle of live images, only one viewer available at the same time and disappearing own avatar on the viewer's GUI. Therefore, it has not been enough comfortable for remote visitors, yet.

5 Lessons learned

Not so populated art museum is a good place for the first mission of mobile robots working in public space. Because behavior of people visiting in the art museum is quite gentle to everything including robots and there are small numbers of children. Usually, a child is a natural enemy of robots. There are some persons to observe the all floors for helping visitors and guarding art works. Therefore, if robots become enough reliable, no additional people for robots maybe are needed. Consequently, I think that the art museum has a good nature for the first stage to test the robot working in general public. But, of course, we need agreements with artists/designers of the exhibition and organizer of the art museum. In our case, our project is joint project with professors of art and design. So, we did not meet so many troubles.

We are surprised at response of local visitors about the robots. We supposed that some complaints from local visitors against the robots. But, there are not any complaints. Maybe, visitors understand the robot as a kind of art works. So, the robot moving around museum is not so anything strange. It means that appearance design is very importance factor of the robot working for the art museum. In another important factor, the sound noise should be reduced as lower level as possible.

The robot controlled via Internet in art museum gives additional value for art museum and artists. It can be considered as a new method of propaganda. Remote visitors can appreciate real art works via Internet. But, remote viewing via Internet does not give enough satisfaction by technical reasons such as the limited human interface and time delay of data transmission. Therefore, if the exhibition is attractive for remote visitors, they want to go out to the real museum.

Remote viewing system can be extended to new application of 3D viewing on the Web. The current 3D viewing on the Web needs to prepare the 3D data of the contents. But, remote viewing system does not need the stored data. The robot approaches the real 3D objects on demand via Internet and can send the image from various view points. That is a kind of new application combined with robotics and IT.

6 Conclusions

We have described our remote viewing system by tele-operation of the mobile robots working as avatars of remote visitors via the Internet using a WWW browser. We introduced the background which is the 5 years joint project with professors of art and design, implementation of overall system, and experimental results at the Tsukuba art museum from 2000 to 2002. Finally, we discussed lessons learned from many trials using mobile robot in the art museum as an example of public space.

In future, our system will be extended to remote viewing 3D art works and multiple mobile robots in the same floor. Then, we will discover new application combied with robotics and IT.

References

[1] "ROBOTS on the WEB," IEEE Robotics and Automation MAGAZINE, vol.7 No.1 and No.2, 2000.

[2] S. Maeyama, S. Yuta and A. Harada, "Experiments on a Remote Appreciation Robot in an Art Museum," Proceedings of 2000 IEEE/RSJ International Conference on Intelligent Robot and Systems, pp.1008-1013, 2000.

[3] A. Harada, "Modeling the Evaluation Structure of *Kansei* Using Networked Robot," Proceedings of International Conference on Advanced Intelligent Mechatronics, pp. 48-55, 1997.

[4] http://www.roboken.esys.tsukuba.ac.jp

[5] T. Sekimoto, T. Tsubouchi and S. Yuta, "A Simple Driving Device for a Vehicle - Implementation and Evaluation -," Proceedings of IEEE/RSJ International conference on Intelligent Robots and Systems, pp.147-154, 1997.

[6] S. Maeyama, S. Yuta and A. Harada, "Remote Viewing on the Web using Multiple Mobile Robotic Avatars," Proceedings of 2001 IEEE/RSJ International Conference on Intelligent Robot and Systems, pp.-, 2001.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

TOURBOT and WebFAIR: Web-Operated Mobile Robots for Tele-Presence in Populated Exhibitions

Wolfram Burgard[1] Panos Trahanias[2] Dirk Hähnel[1] Mark Moors[3]
Dirk Schulz[3] Haris Baltzakis[2] Antonis Argyros[2]

[1]University of Freiburg, Computer Science Department, Germany

[2]Foundation for Res. and Technology - Hellas (FORTH) and University of Crete, Greece

[3]University of Bonn, Department of Computer Science, Germany

Abstract

The current paper presents techniques that facilitate mobile robots to be deployed as interactive agents in populated environments, such as museum exhibitions or trade shows. The mobile robots can be tele-operated over the Internet and this way provide remote access to distant users. Throughout this paper we describe several key techniques that have been developed in the relevant projects. They include robust mapping and localization, people-tracking and advanced visualizations for Web users. The developed robotic systems have been installed and operated in the premises of various sites. Use of the above techniques, as well as appropriate authoring tools, has resulted in drastic reduction in the installation times. Additionally, the systems were thoroughly tested and validated in real-world conditions. Such demonstrations ascertain the functionality and reliability of our methods and provide evidence as of the operation of the complete systems.

1 Introduction

Mobile robotic technology and its application in various sectors is currently an area of high interest and research in this field promises advanced developments and novelties in many aspects. More specifically, applications of mobile robotic technology in public spaces can be found in a field that we can informally term “robots in exhibitions”. In this context, robots can offer alternative ways for interactive tele-presence in exhibition spaces.

Two recent EC-IST funded projects, namely TOURBOT (www.ics.forth.gr/tourbot) and WebFAIR (www.ics.forth.gr/webfair) address the above goal. TOURBOT started January 2000 and ended successfully February 2002; it pursued the development of an interactive tour-guide robot able to provide individual access to museums' exhibits over the Internet. The results of TOURBOT were demonstrated through the installation and operation of the system in the real environment of the three museums that participated in the TOURBOT consortium as well as other interested organizations. WebFAIR started December 2001 and ends May 2004. WebFAIR builds on TOURBOT results and attempts to extend relevant developments to the more demanding environments of trade shows. Additionally, WebFAIR introduces tele-conferencing between the remote user and on-site attendants and employs a multi-robot platform, facilitating thus simultaneous robot control by multiple users.

The motivation for pursuing TOURBOT was twofold, put forward by researchers in the robotics field as well as in the museum community. Evidently, from the robotics vantage point, the research and technical challenges involved in developing this application was the main driving force. Museum curators and organizers were fascinated by the innovative concept of TOURBOT and the idea to offer novel services to their visitors. The successful course of TOURBOT and the vision to introduce corresponding services to the taxing case of trade fairs, resulted in launching WebFAIR. The latter, currently under development, was additionally endorsed by experts in the organization and promotion of large trade shows.

In this paper we present highlights of the techniques developed in the above mentioned projects. They cover various aspects of robots that are deployed in populated environments and hence have to interact with people therein. Among them is a feature-based technique for mapping large environments, a method for tracking people with a moving mobile robot, and an approach to filter spurious measurements coming from persons in the environment while the robot is mapping it. Furthermore, we describe new aspects of the user interfaces. Among them are a speech interface for on-site users and a flexible web-interface with enhanced visualization capabilities for remote users. Additionally we report on the demonstration events that took place in the framework of TOURBOT and argue on the drastic reduction of the system set-up time that was achieved.

2 Related Work

Over the last decade, a variety of service robots were developed that are designed to operate in populated environments. Example cases are robots that are deployed in hospitals [25], museums [7, 35, 48], trade-fairs [38], office buildings [2, 44, 1, 24], and department stores [13]. In these environments the mobile robots perform various services, e.g., deliver, educate, entertain [40] or assist people [39, 29].

Recently, a variety of methods have been developed that estimate the positions of persons in the vicinity of the robot or generate actions given knowledge about a person's position or activity [26, 45, 28, 30, 6]. The TOURBOT and WebFAIR systems apply sample-based joint probabilistic data association filters to estimate the positions of multiple persons in the vicinity of the robot.

Creating maps with mobile robots is one of the key prerequisites for truly autonomous systems. In the literature, the mobile robot mapping problem is often referred to as the *simultaneous localization and mapping problem (SLAM)* [10, 12, 31]. Approaches to concurrent mapping and localization can roughly be classified according to the kind of sensor data processed and the matching algorithms used. For example, the approaches described in [43, 10, 12, 31] extract landmarks out of the data and match these landmarks to localize the robot in the map being learned. The other set of approaches such as [32, 20, 47] use raw sensor data and perform a dense matching of the scans. All these approaches, however, assume that the environment is almost static during the mapping process. Especially in populated environments, additional noise is introduced to the sensor data which increases the risk of errors during the mapping process. To cope with these problems, our system includes a feature-based technique for simultaneous mapping and localization. Additionally, it uses a people tracking system to identify spurious measurements and to consider them appropriately during the mapping process.

In addition, a variety of Web-based tele-operation interfaces for robots has been developed over the last years. Three of the earlier systems are the Mercury Project, the "Telerobot on the Web", and the Tele-Garden [17, 18, 46]. These systems allow people to perform simple tasks with a robot arm via the Web. Since the manipulators operate in prepared workspaces without any unforeseen obstacles, all movement commands issued by a Web user can be carried out in a deterministic manner. Additionally, it suffices to provide still images from a camera mounted on the robot arm after a requested movement task has been completed. The mobile robotic platforms Xavier, Rhino and Minerva [44, 7, 48] could also be operated over the Web. Their interfaces relied on client-pull and server-push techniques to provide visual feedback of the robot's movements; this includes images taken by the robot as well as a java-animated map indicating the robot's current position. However, their interfaces do not include any techniques to reflect changes of the environment. 3D graphics visualizations for Internet-based robot control have already been suggested by Hirukawa et al. [23]. Their interface allows Web users to carry out manipulation tasks with a mobile robot, by controlling a 3D graphics simulation of the robot contained in the Web browser.

The TOURBOT and WebFAIR systems use video streams to convey observed information to the user. Additionally, they provide online visualizations of their actions in a virtual three-dimensional environment. This allows the users to choose arbitrary viewpoints and leads to significant reductions of the required bandwidth.

3 Feature-based Mapping

In order to navigate safely and reliably, an autonomous mobile robot must be able to find its position within its environment. For this purpose, the creation and maintenance of suitable representations of the environment is necessary. Two alternative mapping techniques have been developed, that produce occupancy grid maps and feature maps, respectively. The former is suitable for use with discrete (Markov-based) localization approaches [8, 15, 27], while the latter facilitates the use of continuous (Kalman filter based) localization techniques, as well as hybrid approaches [4].

The feature-based mapping algorithm utilizes line segments and corner points which are extracted out of laser range measurements. At first, a variant of the Iterative-End-Point-Fit algorithm [33] is used to cluster the end-points of a range scan into sets of collinear points. Corner points are then computed at the intersections of directly adjacent line segments [5]. During mapping, the pose of the robot is estimated via a hybrid localization approach, namely a switching-state-space model [4]. At each (discrete) state, an Extended Kalman Filter (EKF) is used for accurate pose estimation. The success of any Kalman filtering method for localization tasks heavily depends on the correct data association. If features are matched in a wrong way, then any filter can diverge with the effect that the mapping process fails. Our robot utilizes the method described in [4] which is based on a dynamic programming string-search algorithm. The algorithm exploits information contained in the spatial ordering of the features. Additionally, the dynamic programming implementation furnishes it with computational efficiency.

To close loops during mapping, the algorithm interleaves localization and mapping just like other techniques which

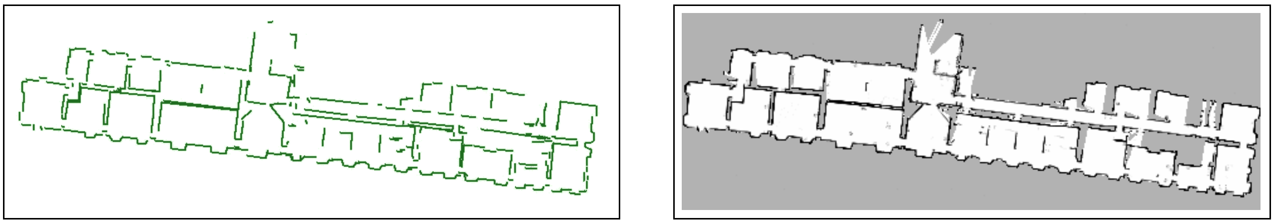

Figure 1. Line feature map (left) and occupancy grid map (right) of an exhibition site generated by the robot

rely on the popular EM-algorithm [49]. During the E-step, our algorithm uses the EKF to provide a maximum a-posteriori estimate of the robot pose given all available measurements; in the M-step the mapped features are re-calculated. This procedure is iterated until convergence is achieved (no significant changes are made to the map features) or a maximum number of iterations is reached. The left image in Figure 1 shows a typical map of an exhibition site resulting from this process. During mapping the robot could successfully close several cycles.

To perform several navigation tasks, such as path planning and obstacle avoidance, the TOURBOT and WebFAIR robots employ occupancy grid maps [34] and apply the probabilistic algorithms described in [7, 11]. The right image in Figure 1 shows a typical occupancy grid map that is learned from the same data and used for the navigation while the robot is giving tours.

4 People Tracking

Tour-guide robots by definition operate in populated environments. Knowledge about the position and the velocities of moving people can be utilized in various ways to improve the behavior of tour-guide robots. For example, it can enable a robot to adapt its velocity to the speed of the people in the environment. It can also be used by the robot to improve its collision avoidance behavior in situations in which the trajectory of the robot crosses the path of a human. And of course, being able to keep track of people is an important prerequisite for human-robot interaction.

The TOURBOT and WebFAIR systems apply sample-based joint probabilistic data association filters (SJPDAFs) [41] to estimate the positions of people in the vicinity of the robot. A set of particle filters [19, 37] is employed to keep track of the individual persons in the vicinity of the robot. The particle filters are updated according to the sensory input and using a model of typical motions of persons. The approach computes a Bayesian estimate of the correspondence between features detected in the sensor data and the different objects to be tracked. In the update phase it then uses this estimate to update the individual particle filters with the observed features.

The features are extracted from range data obtained with two laser-range finders. These two sensors, which are mounted at a height of 40 cm, cover the whole surrounding of the robot at an angular resolution of 1 degree. To robustly identify and keep track of persons, the robot extracts different features. Persons typically generate local minima in the distance profile of the range scan. To distinguish people from static objects that produce similar features, our robot additionally considers the changes in consecutive scans in order to distinguish between static and moving objects. To avoid that the robot loses track of a person when it is occluded by other persons or even objects in the environment, the robot computes occluded areas. The information about occluded areas is particularly useful for the computation of the correspondences and for the updates of the particle filters in situations in which the corresponding feature is missing. The whole process is described in detail in [41].

Figure 2 shows a typical situation, in which the robot is tracking up to four persons in its vicinity. As can be seen from the figure, our approach is robust against occlusions and can quickly adapt to changing situations in which additional persons enter the scene. For example, in the lower left image the upper right person is not visible in the range scan, since it is occluded by the person that is close to the robot. The knowledge that the samples lie in an occluded area prevents the robot from deleting the corresponding sample set. Instead, the samples only spread out, which represents the growing uncertainty of the robot about the position of the person.

5 Mapping in Dynamic Environments

Learning maps with approaches as described in Section 3 has received considerable attention over the last two decades. Although all approaches possess the ability to cope with a certain amount of noise in the sensor data, they assume that the environment is almost static during the mapping process. Especially in populated environments, additional noise is introduced to the sensor data which increases the risk of localization errors or failures during data association. Additionally, people in the vicinity of the robots may appear as objects in the resulting maps and

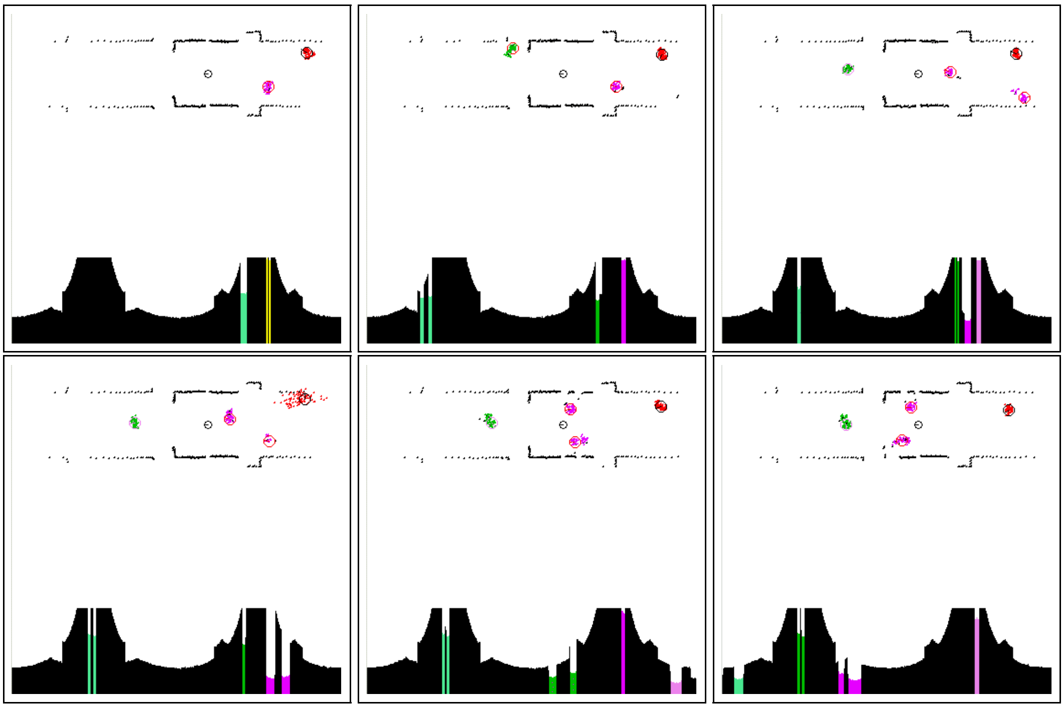

Figure 2. Tracking people using laser range-finder data.

therefore make the maps not usable for navigation tasks. Our mapping system, therefore, is able to incorporate the results of the people tracking process during mapping [22]. This leads to several desirable advantages. First, by incorporating the results of the people tracker, the localization becomes more robust. Additionally, the resulting maps are more accurate, since measurements corrupted by people walking by are filtered out. Compared to alternative techniques such as [50] our approach uses a tracking technique and therefore is able to predict the positions of the person's even in situations in which the corresponding features are temporarily missing.

To avoid spurious objects in the map coming from beams reflected by persons, a bounding box is computed for each sample set of the people tracker. According to our technique only such beams whose endpoint does not lie in any of the bounding boxes are integrated. To cope with the possible time delay of the people tracking process we also ignore corresponding beams of several previous and subsequent scans before and after the person was detected. During the generation of the grid map one generally can be more conservative, because the robot usually scans every part of the environment several times.

Figure 3 shows maps of the Byzantine and Christian Museum in Athens that were recorded with and without incorporating the results of the people-tracker into the mapping process. Both maps actually were generated using the same data set. While the robot was gathering the data, up to 20 people were moving in this environment. The left image shows the endpoints of the laser-range data after localization. Obviously, a corresponding grid map would be useless, since it would contain many spurious objects that might have a negative effect on several standard navigation tasks such as localization and path planning. The right image of Figure 3 shows the Map resulting from our approach. As can be seen from the figure, our robot is able to eliminate almost all spurious objects so that the resulting map provides a better representation of the true state of the world.

6 The Web Interface

In addition to interacting with people in the exhibitions, a main goal in our projects is to establish tele-presence over the internet. Compared to interfaces of other systems such as Xavier, Rhino and Minerva [44, 9, 42], the web interface of the TOURBOT system provides enhanced functionality. Instead of image streams that are updated via server-push or client-pull technology, it uses a commercial live streaming video and broadcast software [51] that provides continuous video transmissions to transfer images recorded with the robot's cameras to the remote user. Additionally, web-users have a more flexible control over the robot. They can control the robot exclusively for a fixed amount of time which generally is set to 10 minutes per user. Whenever a user has

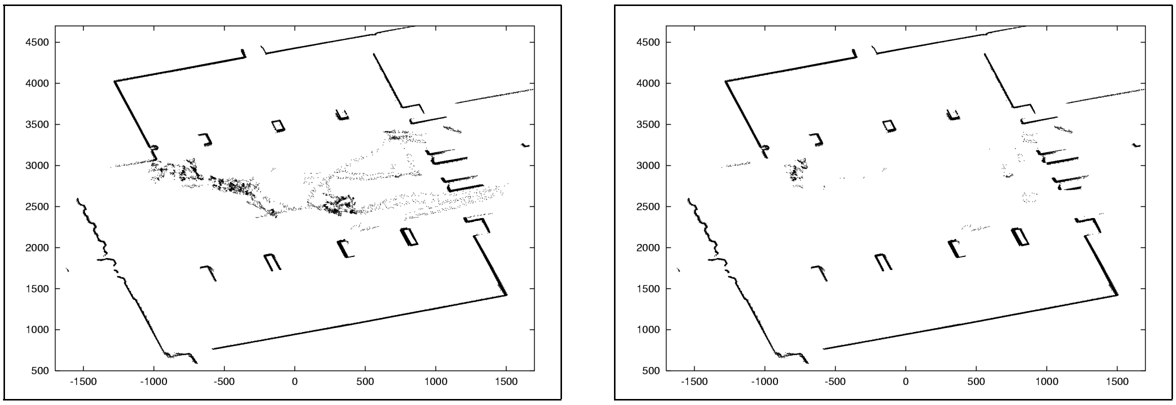

Figure 3. Maps of the Byzantine and Christian Museum in Athens created without (left) and with (right) people filtering.

You are in control.
180 seconds remain.

St Demetrios
Attributed to
Frangos Katelanos
Mid-16th century
0.95 x 0.69 m
T 143

The patron saint of Thessaloniki is portrayed in frontal pose against a gold and light blue ground, with the nominative inscription preserved at the top: Ο ΑΓΙΟΣ ΔΗΜΗΤΡΙΟΣ. Attired in military uniform, he holds a spear in the right hand

and a sword in the left, while a helmet and shield hang from his shoulders. His halo and the icon's frame are decorated in relief with vegetal motifs.

The stylistic and iconographic traits of the icon recall works by Frangos Katelanos, the leading representative of the so-called School of Northern Greece, which flourished in parallel with the Cretan School of monumental painting in the sixteenth century.

Figure 4. Web interface of the TOURBOT system for exclusive control over the robot.

control over the robot, he/she can direct it to arbitrary points in the exhibition. Also, the user can select from a list of predefined guided tours. Furthermore, the user can direct the robot to visit particular exhibits in the exhibition. At each point in time, the user can request a high-resolution image grabbed with the cameras maximal resolution. This way, the interface combines the properties of previous systems. In addition to that, it also allows to control the pan-tilt unit of the robot. Thus, the user can look in arbitrary directions at every point in time. Finally, it offers complex navigation tasks. For example, the user can request the robot to move around an exhibit in order to view it from all possible directions. The control page of the interface is depicted in Figure 4. The left side contains the predefined tours offered to the user. The center shows the live-stream as well as a Java applet animating the robot in a 2D floor-plan. This map can also be used to directly move the robot to an exhibit or to an arbitrary location in the exhibition. Between the map and the live-stream, the interface includes control buttons as well as a message window displaying system messages. The right part of the interface shows multi-media information about the exhibit including links to relevant background information.

7 Enhanced Visualizations

Once instructed by a Web user, the robot fulfills its task completely autonomously. Since the system also operates during opening hours, the robot has to react to the visitors in the museum. This makes it impossible to predict the robot's course of action beforehand. Therefore, it is highly important, to visualize the environment of the robot and the moving people therein, so that the web user gets a better understanding of what is going on in the museum and why the robot is carrying out the current actions.

A typical way of providing information to the users is video streams, recorded with static or robot-mounted cameras. This, however, has the disadvantage of limited perspectives and high bandwidth requirements. For these reasons, we developed a control interface, which additionally provides the user with a virtual reality visualization of the environment including the robot and the people in its vicinity. Based on the state information received from the robot and our tracking algorithm, our control interface continously updates the visualization. Depending on the level of detail of the virtual reality models used, the Internet user can obtain visualizations, whose quality is comparable to video streams. For example, Figure 5 shows two sequences of visualizations provided during the installation of the system in the Deutsches Museum Bonn in November 2001 along with images recorded with a video camera

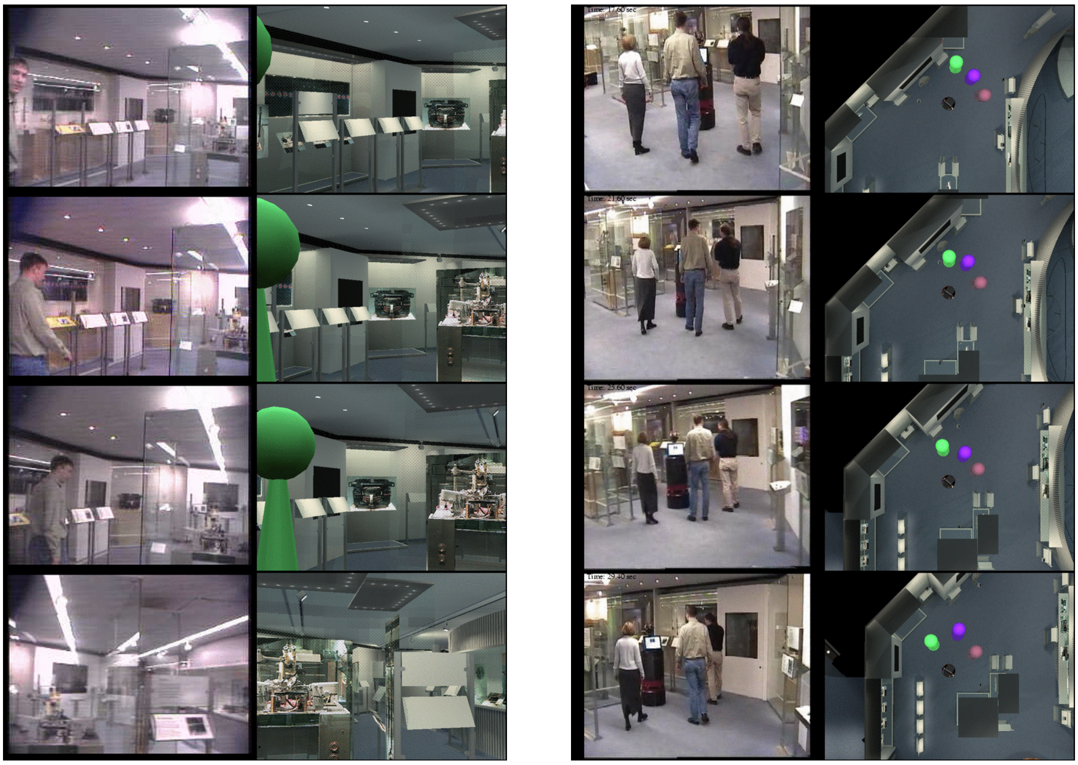

Figure 5. The enhanced 3D visualization allows arbitrary view-points. The left sequence shows the real and the virtual view through the robot's cameras. The right images show the robot guiding three people through the museum and a bird's eye view of the scene.

and with the robot's on-board camera. Within the graphics visualization, people are shown as avatars. As can be seen, the visualization is almost photo-realistic and the animated avatars capture the behavior of the people in the scene quite well.

Compared to the transmission of video streams, the graphics-based visualization highly reduces the bandwidth requirements of the control interface. TOURBOT's standard web interface used a single video stream to transmit images of 240 by 180 pixels in size with a frame rate of about 5 Hz. This still required a bandwidth of about 40kBit/s. Compared to that, the graphics-based visualization only needs about 1kBit/s to achieve the same frame rate, if we assume that 7 people are constantly present in the robot's vicinity. It has the additional advantage, that the bandwidth requirement is independent of the image size. The graphics-based solution, therefore, allows for more detailed visualizations. Beyond the bandwidth savings, the graphics-based visualization offers an increased flexibility to the Internet user. Virtual cameras can be placed anywhere and the viewpoints can even be changed at run-time, as illustrated in the right image sequence of Figure 5. Our current prototype implements these ideas. It uses Open Inventor models of the robot and of the environment for the 3D rendering. On start-up, the control interface connects to the robot via TCP/IP and after downloading the model, the visualization component receives state information from the robot and starts rendering the scene accordingly.

8 The Speech Interface

To enhance the communication with users in the museum, the robots are equipped with a speaker-independent speech interface. We employ a commercially available speech system [36] that detects simple phrases. The input of the user is processed and the parsed phrase is used to generate corresponding actions. To improve the recognition rate, the software allows the definition of contexts, i.e., sets of phrases that are relevant in certain situations. Depending on user input or depending on the task that is currently carried out, the system can dynamically switch between the different contexts. The current system includes 20 different phrases, that can be used to request information about the robot, the exhibition site, or even the time and the

Figure 6. Person interacting with Albert during a Hannover trade fair demonstration.

weather. In several installations in populated environments we figured out that the overall recognition rate is approximately 90%. Figure 6 shows a scene in which a person interacts with the robot Albert during the Hannover trade fair in 2001. Here the person asked several questions about the robot and requested information about the time (*who are you?*, *where are you from?*, *what are you doing here?*). Depending on the input of the user the robot can dynamically generate speech output. The text to be spoken is converted into audio files that are directly sent to the sound card.

9 System Installation and Demonstration

In the framework of the TOURBOT project a number of demonstration trials was undertaken in the premises of the participating museums. More specifically, the TOURBOT system has first been developed and fully tested in the laboratory environment. Following that, and in order to acquire performance data from actual museum visitors, the system has been installed and demonstrated in the three museums of the consortium. These demonstrations were combined with relevant events in order to publicize and disseminate the results of the project to professionals and the broader public. Factual information of these events is as follows:

- Foundation of the Hellenic World, Athens, Greece, May 28–June 2, 2001. Exhibition: “Crossia, Chitones, Doulamades, Velades - 4000 Years of Hellenic Costume.” The exhibition area comprised 2000 square meters. During the trial the robot operated approximately 60 hours covering a distance of 14 kilometers. More than 1200 web users observed the exhibition through TOURBOT. A typical situation, in which the robot Lefkos guides visitors through the museum is shown in Figure 7.

Figure 7. Robot Lefkos operating in the exhibition of the Foundation of the Helenic World.

Figure 8. Robot Rhino operating in the Deutsches Museum Bonn.

- Deutsches Museum Bonn, Bonn, Germany, November 6–11, 2001 (see Figure 8). Exhibition: “Part of the permanent exhibition, highlighting scientific achievements that were awarded the Nobel Prize.” The exhibition area in which the robot moved comprised about 200 square meters. The system operated about 60 hours, covering a distance of 10 km. Approximately 1900 web visitors had a look around the museum via the robot.

- Byzantine and Christian Museum, Athens, Greece, December 3–7, 2001 (see Figure 9). Exhibition: “Byzantium through the eyes of a robot.” The exhibition area comprised about 330 square meters. During the trial the robot operated 40 hours, covering a distance of 5.3 kilometers. The number of web users was small in this trial, due to the following fact. Since the

Figure 9. Robot Lefkos operating in the Byzantine and Christian Museum.

Figure 10. Robot Albert interacting with a person at the Heinz Nixdorf MuseumsForum. This picture is curtesy of Jan Braun, Heinz Nixdorf MuseumsForum.

first day of the trial at the Byzantine and Christian Museum, a large number of (on-site) visitors were coming to the exhibition. This forced the TOURBOT team to the decision to devote significantly more time of the system to on-site visitors as opposed to web visitors.

Additionally, TOURBOT was installed and operated for a longer period of time (Oct. 2001–Feb. 2002) at the Heinz Nixdorf MuseumsForum (HNF) in Paderborn, Germany (see Figure 10). This was in the framework of the special exhibition at HNF ”Computer.Gehirn” (Computer.Brain) with a focus on the comparison of the capabilities of computers/robots and human beings. Recently (June 2002), TOURBOT was introduced for one week in the Museum of Natural History of the University of Crete, Heraklion, Greece.

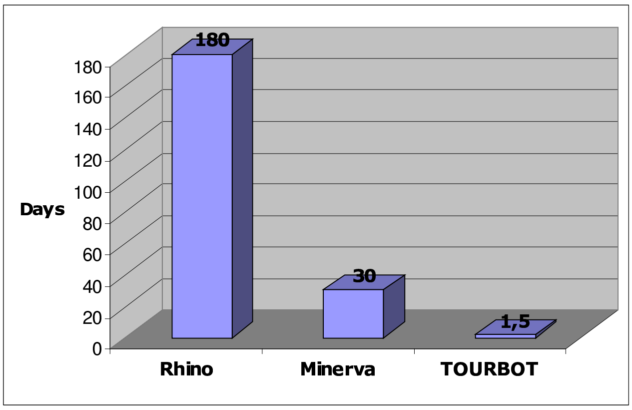

Figure 11. Time required to install the different tour-guide systems Rhino, Minerva, and TOURBOT.

9.1 Installation Time

The large number of test installations of the TOURBOT system required sophisticated tools for the setup of the overall system. Obviously, the most crucial part is the generation of the navigation map. However, based on the techniques described above, the overall mapping process could in all cases be accomplished within several hours. To avoid that the robot leaves its desired operational space or collides with obstacles that cannot be sensed, we manually create a second map with artificial obstacles. These artificial obstacles are fed into the collision avoidance module [7] and thus prevent the robot from moving into the corresponding areas.

A further time consuming process is the generation of the multimedia-content that is presented to the user for each exhibit. The TOURBOT system includes a generic Multimedia database including html-pages, images, audio, and video sequences. Material in the database can be changed and/or edited using available software tools. Furthermore, the robot is equipped with a task specification that defines where the designated exhibits are and which content has to be presented.

Most of the multimedia information pertinent to the exhibits can be obtained directly from the exhibition sites, since pictures, text and other relevant material are often already contained in existing Web presentations. The whole setup can therefore be accomplished in less than two days. This is an enormous speed-up compared to previous tour-guide systems. Figure 11 shows the time required to install the Rhino and Minerva systems [7, 48] in comparison to that of the TOURBOT system. As can be seen, the TOURBOT system requires significantly less time than Rhino and Minerva. Our experience with tour-guide robots in exhibition sites suggests that three-dimensional models of exhibitions’ premises are generally not available. The automatic generation of such models with the mobile robot itself is a subject of ongoing research [21].

10 Conclusions

The goals set for by the TOURBOT and WebFAIR projects are in-line with on-going activities towards the development of fully autonomous robots that operate in populated environments. The mentioned projects aim at the development of interactive tour-guide robots, able to serve web- as well as on-site visitors. Technical developments in the framework of these projects have resulted in robust and reliable systems that have been demonstrated and validated in real-world conditions. Equally important, the system set-up time has been drastically reduced, facilitating its porting in new environments. Current research extends the navigation capabilities of the robotic systems by addressing obstacle avoidance in the cases of objects that are not visible by the laser scanner [3], 3D mapping [21], mapping in dynamic environments [22], predictive navigation [14], and multi-robot coordination [16]. Moreover, in the context of the above projects additional issues are addressed that consider (a) how to adapt this technology in order to fit the long-term operational needs of an exhibition site, (b) how to evaluate the robotic system in terms of its impact to the main function and objectives of the exhibition site (financial impact, accessibility, marketing and promotion, impact on visitor demographic, etc.), and (c) how to evaluate the content and educational added value to museum and exhibition visitors, and generate a feedback to the technology developers in order to improve in the future the robotic avatars and adapt further to the needs of the users.

11. Acknowledgments

This work has partly been supported by the by the IST Programme of Commission of the European Communities under contract numbers IST-1999-12643 and IST-2000-29456. The authors furthermore would like to thank the members of the IST-project TOURBOT for helpful comments and fruitful discussions.

References

[1] K. Arras and S. Vestli. Hybrid, high-precision localisation for the mail distributing mobile robot system MOPS. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 1998.

[2] H. Asoh, S. Hayamizu, I. Hara, Y. Motomura, S. Akaho, and T. Matsui. Socially embedded learning of office-conversant robot jijo-2. In *Proceedings of IJCAI-97*. IJCAI, Inc., 1997.

[3] H. Baltzakis, A. A., and P. Trahanias. Fusion of range and visual data for the extraction of scene structure information. In *Intl. Conf. on Pattern Recognition, (ICPR)*, 2002.

[4] H. Baltzakis and P. Trahanias. Hybrid mobile robot localization using switching state-space models. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.

[5] H. Baltzakis and P. Trahanias. An iterative approach for building feature maps in cyclic environments. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.

[6] H. Bui, S. Venkatesh, and G. West. Tracking and surveillance in wide-area spatial environments using the Abstract Hidden Markov Model. *Intl. J. of Pattern Rec. and AI*, 2001.

[7] W. Burgard, A. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2), 1999.

[8] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 1996.

[9] W. Burgard and D. Schulz. Robust visualization for web-based control of mobile robots. In K. Goldberg and R. Siegwart, editors, *Robots on the Web: Physical Interaction through the Internet*. MIT-Press, 2001.

[10] J. Castellanos, J. Montiel, J. Neira, and J. Tardós. The SPmap: A probabilistic framework for simultaneous localization and map building. *IEEE Transactions on Robotics and Automation*, 15(5):948–953, 1999.

[11] F. Dellaert, W. Burgard, D. Fox, and S. Thrun. Using the condensation algorithm for robust, vision-based mobile robotlocalization. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1999.

[12] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. In *ICRA'2000 Workshop on Mobile Robot Navigation and Mapping*, 2000.

[13] H. Endres, W. Feiten, and G. Lawitzky. Field test of a navigation system: Autonomous cleaning in supermarkets. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 1998.

[14] A. Foka and P. Trahanias. Predictive autonomous robot navigation. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.

[15] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo Localization: Efficient position estimation for mobile robots. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 1999.

[16] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3), 2000.

[17] K. Goldberg, S. Gentner, C. Sutter, J. Wiegley, and B. Farzin. The mercury project: A feasibility study for online robots. In K. Goldberg and R. Siegwart, editors, *Beyond Webcams: An Introduction to Online Robots*. MIT Press, 2002.

[18] K. Goldberg, J. Santarromana, G. Bekey, S. Gentner, R. Morris, J. Wiegley, and E. Berger. The telegarden. In *Proc. of ACM SIGGRAPH*, 1995.

[19] N. Gordon, D. Salmond, and A. Smith. A novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F*, 140(2):107–113, 1993.

[20] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the IEEE Int. Symp. on Computational Intelligence in Robotics and Automation (CIRA)*, 1999.

[21] D. Hähnel, W. Burgard, and S. Thrun. Learning compact 3d models of indoor and outdoor environments with a mobile robot. In *Fourth European workshop on advanced mobile robots (EUROBOT'01)*, 2001.

[22] D. Hähnel, D. Schulz, and W. Burgard. Map building with mobile robots in populated environments. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.

[23] H. Hirukawa, I. Hara, and T. Hori. Online robots. In K. Goldberg and R. Siegwart, editors, *Beyond Webcams: An Introduction to Online Robots*. MIT Press, 2002.

[24] I. Horswill. Polly: A vision-based artificial agent. In *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI)*. MIT Press, 1993.

[25] S. King and C. Weiman. Helpmate autonomous mobile robot navigation system. In *Proc. of the SPIE Conference on Mobile Robots*, pages 190–198, 1990.

[26] B. Kluge, C. Köhler, and E. Prassler. Fast and robust tracking of multiple moving objects with a laser range finder. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2001.

[27] K. Konolige and K. Chou. Markov localization using correlation. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1999.

[28] E. Kruse and F. Wahl. Camera-based monitoring system for mobile robot guidance. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1998.

[29] G. Lacey and K. Dawson-Howe. The application of robotics to a mobility aid for the elderly blind. *Journal of Robotics and Autonomous Systems (RAS)*, 23:245–252, 1998.

[30] S. M. Lavalle, H. H. Gonzalez-Banos, G. Becker, and J.-C. Latombe. Motion strategies for maintaining visibility of a moving target. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 1997.

[31] J. Leonard and H. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In *Proc. of the Ninth Int. Symp. on Robotics Research (ISRR)*, 1999.

[32] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.

[33] T. Michael and T. Quint. Sphere of influence graphs in general metric spaces. *Mathematical and Computer Modelling*, 29:45–53, 1994.

[34] H. P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, pages 116–121, 1985.

[35] I. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto. An affective mobile educator with a full-time job. *Artificial Intelligence*, 114(1-2), 1999.

[36] http://www.novotech-gmbh.de/.

[37] M. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association*, 94(446), 1999.

[38] D. Rodriguez-Losada, F. Matia, R. Galan, and A. Jimenez. Blacky, an interactive mobile robot at a trade fair. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.

[39] C. Schaeffer and T. May. Care-o-bot - a system for assisting elderly or disabled persons in home environments. In *Assistive technology on the threshold of the new millenium*. IOS Press, Amsterdam, 1999.

[40] R. Schraft and G. Schmierer. *Serviceroboter*. Springer Verlag, 1998. In German.

[41] D. Schulz, W. Burgard, D. Fox, and A. Cremers. Tracking multiple moving objects with a mobile robot. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.

[42] D. Schulz, W. Burgard, D. Fox, S. Thrun, and A. Cremers. Web interfaces for mobile robots in public places. *IEEE-Magazine on Robotics and Automation*, 2000.

[43] H. Shatkay. *Learning Models for Robot Navigation*. PhD thesis, Computer Science Department, Brown University, Providence, RI, 1998.

[44] R. Simmons, R. Goodwin, K. Haigh, S. Koenig, and J. O'Sullivan. A layered architecture for office delivery robots. In *Proc. of the First International Conference on Autonomous Agents (Agents)*, 1997.

[45] S. Tadokoro, M. Hayashi, Y. Manabe, Y. Nakami, and T. Takamori. On motion planning of mobile robots which coexist and cooperate with human. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1995.

[46] K. Taylor and J. Trevelyan. A telerobot on the World Wide Web. In *Proceedings of the 1995 National Conference of the Australian Robot Association*, 1995.

[47] S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *Journal of Robotics Research*, 20(5):335–363, 2001.

[48] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *Journal of Robotics Research*, 19(11), 2000.

[49] S. Thrun, W. Burgard, and D. Fox. A probabilistic approach to concurrent mapping and localization for mobilerobots. *Machine Learning and Autonomous Robots (joint issue)*, (31/5), 1998.

[50] C.-C. Wang and C. Thorpe. Simultaneous localization and mapping with detection and tracking of moving objects. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.

[51] http://www.webcam32.com/.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Diligent: Towards a Personal Robot

Rachid Alami*
LAAS/CNRS,
7 avenue du Colonel Roche, F-31077 Toulouse Cedex 04, France
Rachid.Alami@laas.fr

Abstract This paper presents our project to develop a "Personal robot" called "Diligent". We first discuss the main issues involved in personal robotics and more particularly, some of the key features implied by what we call "Human-friendly navigation". Then, we present the current state of Diligent. Diligent is already able to navigate safely and repeatedly in a relatively large, and relatively dynamic environment. It is used almost daily in our laboratory and has been operated remotely over the Internet in several occasions. It has also been deployed in an exhibition. We conclude by discussing some lessons drawn from experience.

1 Introduction

The development of personal robots is a new center of convergence and a motivating challenge in robotics research. One key aspect is "added" to the "standard challenge" of autonomous robots: the essential role of the "human in the loop". This has numerous consequences. Two of them are of particular interest for us:

1. the robot should be able to operate in an environment which has been essentially designed for human, and
2. the robot will have to interact with human.

The human-centered theme is currently investigated in several areas which are rather different. The spectrum of developments range from humanoids[18] to wearable computing and sensing, human augmentation, telepresence, smart rooms or even intelligent objects[7]. In our case, we aim to contribute to this domain through the "narrow" point of view of robot navigation.

This paper presents our project to develop a "Personal robot" called "Diligent". We first discuss the main issues involved in personal robotics and more particularly, some of the key features implied by what we call "Human-friendly navigation". Then, we present the current state of Diligent. Diligent is already able to navigate safely and repeatedly in a relatively large, and relatively dynamic environment. It is used almost daily in our laboratory and has been operated remotely over the Internet in several occasions. It has also been deployed in an exhibition. We conclude by discussing some lessons drawn from experience.

2 Human-Friendly Navigation

Diligent's essential task is navigation[1]. Our target, within the framework of personal robotics, is to endow Diligent with "human-friendly navigation".

This notion has still to be precisely defined. In fact, we envision our project itself as a mean of incrementally defining and investigating the main issues it involves. However, we can already mention the following issues:

- The ability to operate in an environment open to the public with very little or no adaptation. Indeed, "human-friendly navigation" implies safe (for the human and the robot), robust and reasonably efficient (for the task itself) navigation in such an environment.
- "Human-friendly navigation" means also a mobile robot which is easy to instruct (map and locations learning), easy to use, easy to interact with.
- Another aspect should be the ability of the robot to analyze and react efficiently to human motion around it. For example, it may have to make decisions like "follow", or "pass", or "leave room" to a person or a group of persons. In other words, the robot should "comply" with human activities in a given environment.
- One should also consider more subjective issues such as human understanding and human acceptance of the robot behavior.

We envision such a robot as an extension and a new challenge for autonomous mobile robotics.

---

[*] A number of colleagues have participated to this project, in alphabetical order: Igor Belousov N.V.Keldysh Institute for Applied Mathematics, Michel Devy, Sara Fleury, Jean-Bernard Hayet, Matthieu Herrb, Félix Ingrand, Maher Khatib, Frédéric Lerasle, Javier Minguez (University of Zaragoza), Philippe Moutarlier, Benoit Morisset, Thierry Siméon

[1] Depending on the context, we will endow it with a set of complementary tasks that will heavily rely on navigation abilities

3 Diligent control architecture

Figure 1: Diligent: a Nomadic XR4000.

Diligent hardware is a standard XR4000 Nomadic robot (Fig. 1) equipped with a SICK laser range finder, a pair of cameras mounted on a pan&tilt platform and a tactile screen.
Diligent control architecture is based on the LAAS[2] architecture [1, 2] which has been ported on Linux.

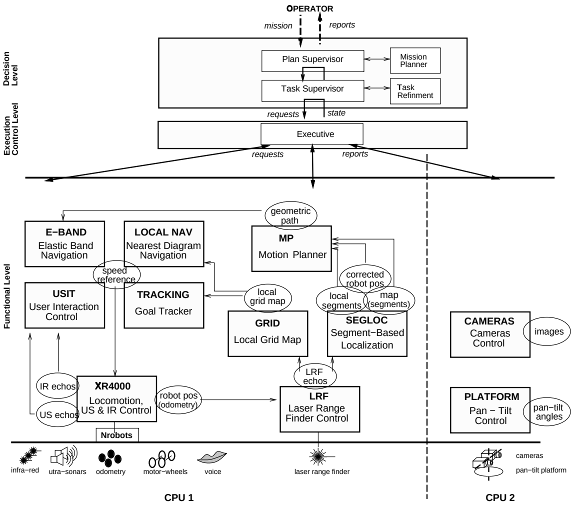

Figure 2: Diligent's architecture

We briefly recall the main ideas underlying this architecture. It involves three hierarchical levels (Fig. 2) having different temporal constraints and manipulating different data representations.

- **The functional level** is composed of a set of "modules" that embed all basic built-in robot action and perception capacities. In the next section, we will present and illustrate the use of the most important ones.
- **The execution control level**, or Executive, controls and coordinates the execution of the functions distributed in the modules, according to the task requirements specified by the decision level.
- **The decision level** which includes the capacities for producing the task plan and supervising its execution, while being at the same time reactive to events from the other levels. It may embed several layers, according to the application; their basic structure is a planner/supervisor pair that enables to integrate deliberation and reaction.

In our point of view, such an architecture is particularly adapted to and should provide a convenient framework for developing human-robot interaction scenarii. Indeed, it is plan-based: this means that the supervisor is built on an explicit reasoning on the tasks and on the robot capacities to achieve them in a given context.
Plan-based architecture allows also to exhibit the robot goals and current plans; this can serve as a basis for developing interaction and/or cooperation schemes. We have already successfully experimented such a similar paradigm for inter-robot cooperation ([4]).

4 A robust navigation system

We briefly present and illustrate here-below the basic bricks with which we have built the current version of Diligent navigation system: localization, motion planning and motion execution primitives as well as a robot supervisor which controls the overall system.

4.1 Segment-based Localization

The localization method relies on a map of obstacle edges incrementally built by the robot itself from laser range data. Map building and localization procedures use an Extended Kalman Filter to match the local perception with a previously built model [16]). Diligent localization module, called SEGLOC, provides three main localization modes:
- The Re-Location Mode is used when the robot is lost.
- The Continuous Tracking Mode is used when a good probabilistic estimate of the robot position is available. It relies both on robot position variance and on a predicted local aspect.
- The Simple Tracking Mode is used when there is not enough information in the environment to produce a new position (e.g. in a corridor).

Fig. 3 shows a localization experiment in a map of about 400 segments previously built by the robot with the

---

[2] LAAS: LAAS' Architecture for Autonomous Systems.

help of a human instructor. The localization process in this relatively large model is performed in less than 120 milli-seconds. We run it at 1 Hertz.

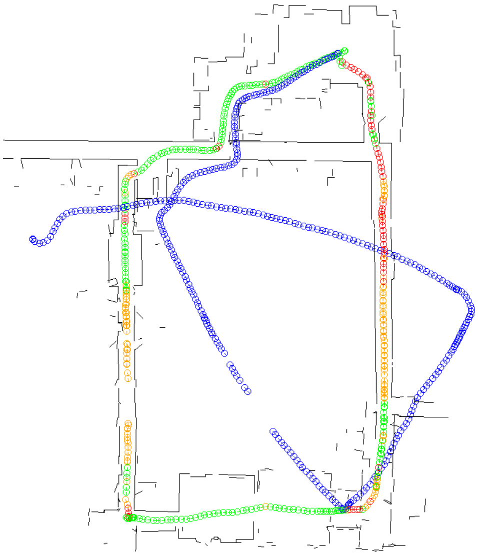

Figure 3: Comparison of position estimation obtained from XR4000 odometry alone (dark circles - missing positions are due to temporarily communication break down) and from the SEGLOC module (grey circles). Dimensions : $45x60m^2$.

4.2 The Motion Planner

The main purpose of this module is to plan feasible paths for the robot. It is able to compute trajectories for holonomous as well as non-holonomous robots [14]. One of its peculiarities, is that it accepts requests which allow the supervisor to control and to adapt motion planning activities to the current task needs. Indeed, the supervisor can select the level of discretization, the type and the source of obstacles to be taken account by the planner (from models built by the other modules), the shape and the kinematic constraints of the robot. The MP module is used in conjunction with the elastic band (E-BAND module): it produces an initial path which is dynamically adapted during execution.

4.3 Elastic Band for Plan-execution

The *elastic band* is a method to dynamically modify a trajectory in order to take into account variations in the obstacle layout between the model used during path planning and the actual sensor data. The method used here was first proposed by S. Quinlan and O. Khatib [17] and further developed and implemented by M. Khatib [12]. The principle is to build a flexible path between the current robot position and the goal, described by a sequence of configurations in free space.

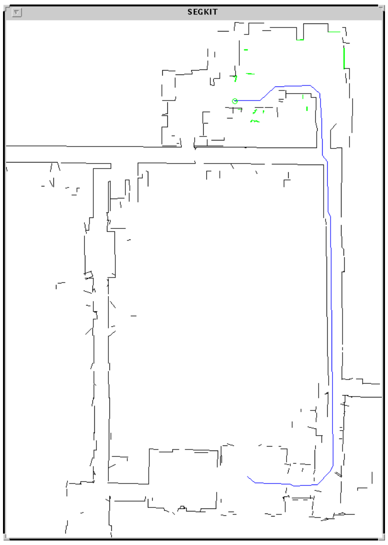

Figure 4: A map of the lab (45 x 60 m) and a path produced by the Motion Planner.

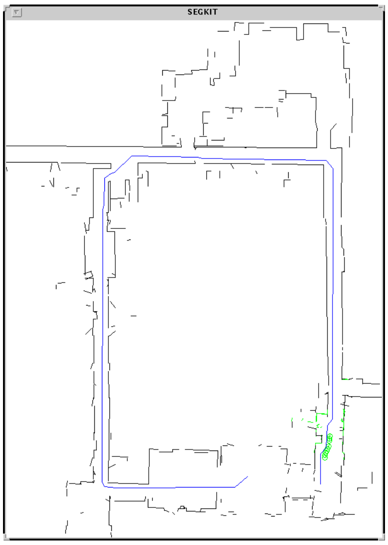

Figure 5: Another path to same goal after having been blocked by a (partially) closed door.

Connexity between these configurations is maintained by a set of internal forces that also optimize the global shape of the path. External forces are associated with obstacles and are applied to all configurations in order to maintain the path away from obstacles (Fig. 6).
The computation time of an elastic band is proportional to its length. In order to maintain the real-time capabilities of the robot, the E-BAND module limits the portion of the trajectory that is elastic. If the band retracts beyond a minimum length, new points are added toward the goal along the trajectory computed by MP. Once it has reached the goal, the band retracts until it

Figure 6: A example of a motion execution based on the elastic band.

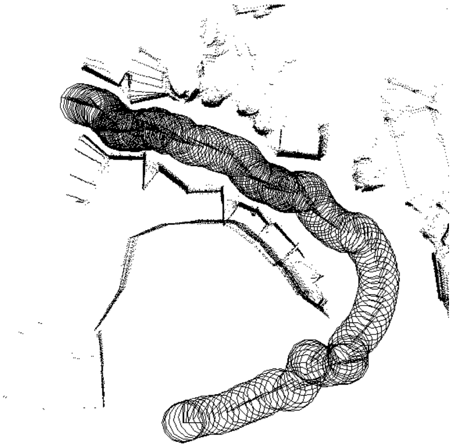

Figure 7: A typical run of ND

only contains a single configuration, indicating that the execution is completed.

In order to take into account moving obstacles, if, during execution, the path is blocked the execution of the band is stopped and a watchdog timer is armed. When the watchdog expires, a failure status is reported to the supervisor. If the next configuration becomes free again before the expiration of the timer (because the obstacle moved), the execution resumes.

4.4 A reactive obstacle avoidance

Diligent is also endowed with a reactive motion execution scheme: the Nearness Diagram method (ND)[11]. ND works in two steps. First, it extracts, from the sensory data, a description of free obstacles regions. It selects one of them (that is usually the closest to the goal location) and evaluates the robot security (based on the distance to obstacles).

For ND, all the possible goal locations and obstacle configurations (relatively to the robot) are described by five general situations. Thus, with the obtained information, ND selects one situation and applies the corresponding law to compute a motion command.

The ND approach is well-suited to deal with constrained and dynamic environments. We intend to use it in a very cluttered environment or in densely populated areas.

4.5 The Robot Supervisor

The robot supervisor is programmed in Propice (formerly PRS) [10, 9]. It provides an interface with all the robot modules thus allowing to program robot control. All supervisor activities (task refinement, control, display, dialog…) are programmed using goal directed and/or situation driven procedures.

The main activities, performed by the Robot Supervisor, in its current version are: Mission Management, Robot Localization Control, Robot Navigation Control, Human/Robot Interaction Control. These different activities interact depending on the current context.

Mission management: This activity involves the management of the set of missions that can be incrementally allocated to the robot through several means. In the current implementation, missions involve essentially navigation tasks. Mission management allows to add or delete missions through textual or through direct human/robot interaction. It is able to postpone, restart or re-plan missions depending on execution contingencies.

Robot Localization Control: This activity runs and controls the localization system (see 4.1). It permanently monitors the "quality" of robot localization and implements different levels of reaction depending on the threshold reached by the robot position uncertainty. The first level induces a change in the re-localization modalities. The second level induces an active strategy (e.g. a momentary stop with rotations in order to find landmarks). The third level corresponds to the inability of the robot to solve the problem by itself; it postpones any autonomous motion tasks and asks for human help.

Robot Navigation Control: Robot navigation tasks are implemented through the use of a "classical" plan-and-execute paradigm based on a previously learned map. However, this paradigm is enhanced by several features. First, trajectory execution, based on the control of an elastic band, allows an effective robustness to contingencies, from local obstacle avoidance, to adaptation to moving and transient obstacles as well as adaptation to significant re-localization updates. Second, if the band is completely blocked for a period of time, a new plan is searched, taking into account an update of the learned map. The loop is repeated iteratively until the robot reaches its goal or the planner finds no path, or an "external" event entails the postponement or the cancellation of the navigation mission.

Human/Robot Interaction Control: This aspect will progressively take an important part in the project. This activity is currently limited at the supervisor level to the control of the function described below.

5 Human-robot interaction

5.1 Various control interfaces

The robot supervisor offers (1) a programmer interface through the Propice programming environment, (2) a textual user interface through a TCP/IP connection and (3) a user and public interface through a compliant motion mode.

Interfaces (2) and (3) allow the user to add or delete navigation missions (e.g. "GO-TO Station-N"). Besides, interface (3) allows the user or any person from the public to set the robot in a compliant mode. The human can then "push" the robot in a compliant but safe (no collision) manner. The direction of motion is detected using the ring of infra-red sensors while the ultra-sonic sensors are used to create repulsive forces in order to avoid contact with obstacles. This mode is used in two situations: (1) when a human encountering the robot wants to move it aside and (2) to provide to an operator a comfortable and fast mean to guide the robot to a given place for example for allowing a robust re-localization.

Another interface feature is the use of the speech synthesis system to indicate the main changes in robot state as well several displays.

5.2 WEB-based navigation control

In the last few years, many systems for WWW-based robot control have been developed. A list of active systems providing free access through Web browsers is presented on the NASA Telerobotics Web-page[19]. Concerning mobile robots, the first systems that have been demonstrated are "Khep On The Web"[13], the Blimp from Berkeley[6] and Xavier[21].

The main characteristic of our system, in its current state, is the ability to control the robot at task level. The operator selects the goal and the robot is entirely autonomous. It plans on-line and adapts its execution to the actual context. However, the operator can request a stop at any moment and can control the camera orientation.

Through this connection the user gets the map of the environment the current planned path, the execution trace, as well as some augmented reality markers on the image allowing to better understand what the robot is doing: projection of the laser segments and of the path on the image. The operator may also request a reconstruction in java-3D of the current view of the map[8].

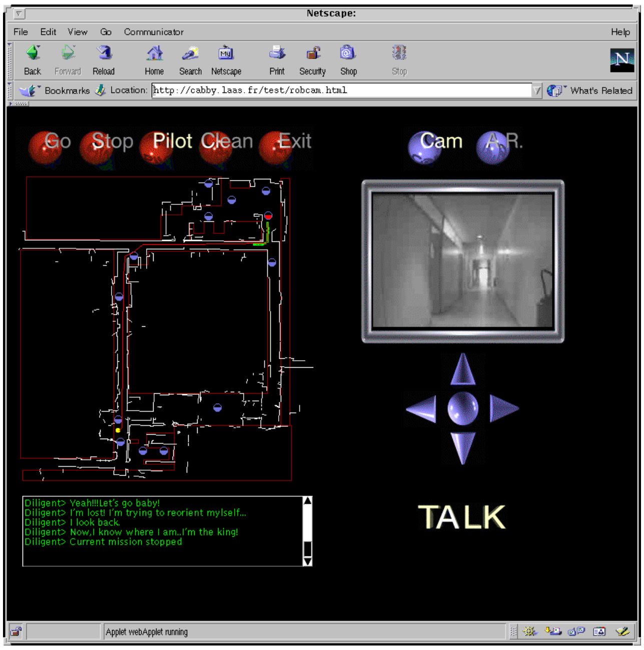

Location: http://cabby.laas.fr/test/robcam.html

```
Diligent> Yeah!!!Let's go baby!
Diligent> I'm lost! I'm trying to reorient mylself...
Diligent> I look back.
Diligent> Now,I know where I am..I'm the king!
Diligent> Current mission stopped
```

Applet webApplet running

Figure 8: Diligent WEB page

The URL is not open to the public. We open it to colleagues and partners. Indeed, Diligent has already been controlled successfully from several places (France, England, Spain, Germany, Japan).

6 Results and learned lessons

6.1 Diligent in the lab

The basic function that we wanted to demonstrate is the ability of the robot to navigate "safely" and repeatedly in a learned in-door environment from one place to another.

We can say that we have succeeded in this first step. Diligent is used almost daily in our laboratory. It performs "long range" navigation tasks in the robotics lab and in the corridors and halls of the building (a tour corresponds to more than 200 meters) see figure.

The robot is also sufficiently robust to serve as a a tool for other developments: other sensor modalities, new reactive motion generators, human-robot interaction developments. In such case, the robot simply "transports" the new developments: it is asked for example to go to some place where the experiment will began. having this, the developer does not have to worry with the robot itself nor to use a joystick to take it to a desired place or to get it back to its station.

Diligent is capable of long run robust navigation in an environment which can be reliably modeled by horizon-

tal laser scans. The navigation is robust to substantial differences between the actual environment state and the learned map. The various figures presented in this paper illustrate such capabilities. Besides, Diligent demonstrated effective autonomy in planning and executing its navigation tasks, end re-planning in case of permanent obstacle or “non-cooperative” human (we call “non-cooperative” human, a human who blocks completely the path of the robot and who does not leave place after a fixed amount of time - several seconds). Fig. 4 and Fig. 5 show a situation where the robot has been blocked at a door and has decided to re-plan.
As mentioned above, it has been also controlled at task-level by remote operators who used Diligent to “visit” our laboratory.

6.2 Diligent outside the lab

However, there is a great difference between running a robot in a “protected” lab environment and leaving the robot “alone” in an environment open to the public. We have deployed Diligent in two public areas: a conference at the “Cité de l’Espace”(Figure 9), and more particularly the SITEF 2000 exhibition (Figure 10).

Figure 9: Diligent at the Cité de l’Espace

SITEF is a an industrial fair. It is not a permanent or a long-term exhibition. The hall where we had to operate the robot was composed or numerous stands belonging to various companies. The configuration of the stands was under the responsibility of each company which means that it was never completely stable.
In spite of that and of the fact that the hall was from time to time really crowded, the segment-based localization proved to be very robust (Figure 11). This is due to the the ability of the system to match segments (corresponding to walls or furniture) even they are partially occluded. Redundancy allows also to have an large part of the environment that is completely occluded.
The 2D laser range data are used for localization and for obstacle avoidance purposes. While for the second, it is reasonable (even though not really enough) to place the sensor at a low height (40 cm), localization is better performed at a height above tables and chairs, and even perhaps above the size of a man.

Figure 10: Diligent at the SITEF 2000

6.3 Human-robot interaction

Diligent in its current version has limited human-robot capabilities (see above). At the SITEF, we welcomed the public at the LAAS exhibition stand where interested persons where invited to assign tasks to the Diligent and to supervise its behavior through its Web interface. Visitors at the stand where given information on what the robot was able to do.
But most visitors simply encountered the robot while it navigates in the exhibition hall. Because people often like to play with robots and to block them, the robot seemed to perform only reactively. Consequently, most visitors did not notice that after long detours and numerous local avoidance and even global re-planning steps, Diligent finally succeeded in reaching a desired position. Another aspect that was almost not noticed (it is even considered as trivial by most visitors), is the fact that the robot was able to permanently "know" its position in the hall.
A funny observation was also made. In order to stop the robot or to make it perform a detour, the visitors often occluded the camera and did not notice that obstacle avoidance was exclusively performed through laser range data.

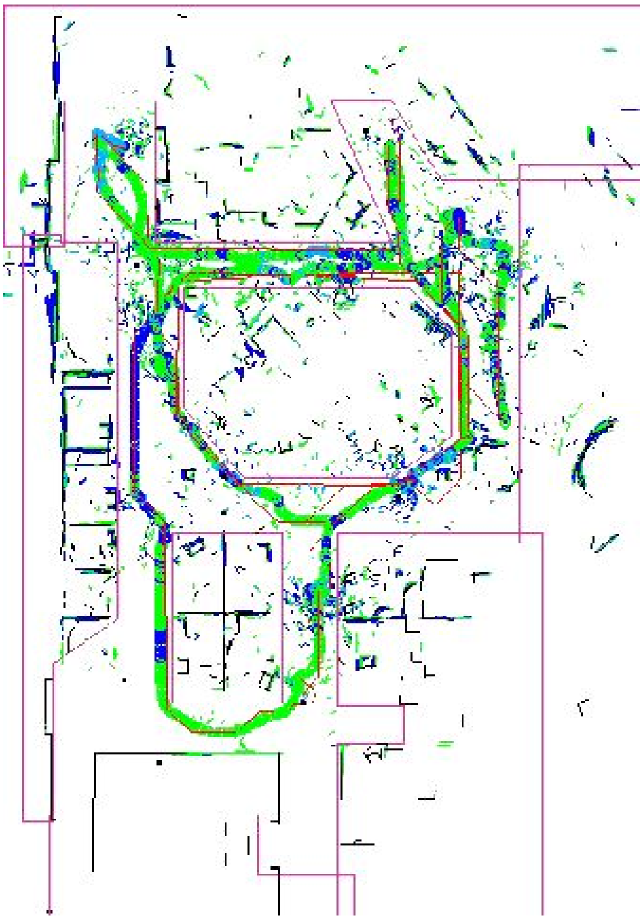

Figure 11: Robot trajectories in a SITEF exhibition hall (60x120m

6.4 Comments and future work

There are numerous extensions that should be developed in order to reach a higher level of autonomy and usefulness of mobile robots in public areas.

We are working on some of them and have obtained some preliminary results that are not yet integrated in the robot.

One way to deploy the robot in large environments is the construction and use of a tolopogical map. We have obtained some preliminary results towards this goal[20], based on the incremental construction of a Voronoi diagram directly from range data.

We are also developing motion planning methods that take explicitly into account dynamic obstacles. We have developed and we will integrate soon a methodology for computing the maximum velocity profile for a planned trajectory. The profile is computed considering the robot and environment dynamics as well as the constraints of the sensing apparatus.

The main advantage is that the robot that moves with the computed velocity profile can assure from its side that it would not collide onto any of the numerous moving objects that could intercept its future trajectory. Indeed, with this computed profiles, the robot's immobility before collision is guaranteed. The methodology has also been incorporated onto a motion planner for a non holonomous robot[5].

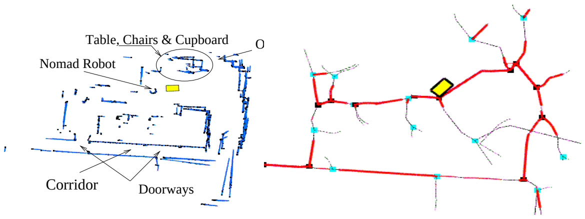

Figure 12: *A topological model bult by Diligent in our lab*

Now, concerning the navigation itself, there is a strong limitation in the localization and obstacle modeling capabilities of our robot. There is a clear need for a 3D modeling of the environment in order to detect all obstacles (stairs, holes,..). This is mandatory for effective deployment of robots in public environments.

This has been "solved" in our running system by adding "virtual obstacles" in the map generated by the robot in order to make it avoid obstacles that it cannot perceive or simply to forbid it to enter some given areas.

Another key aspect that we plan we intend to tackle is Human-robot interaction. We propose to consider it as an incremental and interactive problem solving process. But of course, these developments will be possible only if we also implement perceptual primitives allowing the detection and localization of human, human motion detection, human recognition, human gesture or movement interpretation.

References

[1] R. Alami, R. Chatila, S. Fleury, M. Ghallab, and F. Ingrand. An Architecture for Autonomy. *International Journal of Robotics Research*, 17(4):315–337, April 1998.

[2] R. Alami, R. Chatila, S. Fleury, M. Herrb, F. Ingrand, M. Khatib, B. Morisset, P. Moutarlier, and T. Siméon. Around the lab in 40 days... In *IEEE ICRA, San Francisco, (USA)*, 2000.

[3] R.Alami, I.Belousov, S.Fleury, M.Herrb, F.Ingrand, J.Minguez, B.Morisset Diligent: towards a human-friendly navigation system *2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2000)*, Takamatsu (Japon), 31 Octobre - 5 Novembre 2000, Vol.1, pp.21-26.

[4] R. Alami, S. Fleury, M. Herrb, F. Ingrand, and F. Robert. Multi Robot Cooperation in the Martha Project. *IEEE Robotics and Automation Magazine: Projects funded by the Commission of the European Union*, March 1998.

[5] R. Alami, T. Siméon, K. Madhava Krishna On the influence of sensor capacities and environment dynamics onto collision-free motion plans *2002 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2002)*, Lausanne.

[6] The Blimp. *vive.cs.berkeley.edu/ blimp, currently off-line*, 1999.

[7] J. Canny and R. Alami. Workshop on Personal Robotics. In *IEEE ICRA, Detroit, (USA)*, 1999.

[8] G.J. Clapworthy I. R. Belousov, J.C. Tan. Teleoperation and java3d visualisation of a robot manipulator over the world wide web. In *Proc. of the IEEE Int. Conf. On Information Visualisation IV'99, London, U.K.*, july 1999.

[9] F. F. Ingrand, R. Chatila, R. Alami, and F. Robert. PRS: A High Level Supervision and Control Language for Autonomous Mobile Robots. In *IEEE ICRA, St Paul, (USA)*, 1996.

[10] F. F. Ingrand, M. P. Georgeff, and A. S. Rao. An Architecture for Real-Time Reasoning and System Control. *IEEE Expert, Knowledge-Based Diagnosis in Process Engineering*, 7(6):34–44, December 1992.

[11] J.Minguez, L.Montano, T.Simeon, R.Alami Global nearness diagram navigation (GND) *2001 IEEE International Conference on Robotics and Automation (ICRA'2001)*, Seoul (Korée), 21-26 Mai 2001, pp.33-39.

[12] M. Khatib. *Contrôle du mouvement d'un robot mobile par retour sensoriel.* PhD thesis, Université Paul Sabatier, Toulouse, December 1996.

[13] Khepera. *http://KhepOnTheWeb.epfl.ch*, 1999.

[14] J. P. Laumond, P. E. Jacobs, M. Taix, and R. M. Murray. A motion planner for nonholonomic mobile robots. *IEEE Transactions on Robotics and Automation*, 10(5):577–593, 1994.

[15] R. Luo and T. Chen. Remote supervisory control of sensor based mobile robot via internet. In *IROS'97 Conference*, 1997.

[16] P. Moutarlier and R. G. Chatila. Stochastic Multisensory Data Fusion for Mobile Robot Location and Environment Modelling. In *Proc. International Symposium on Robotics Research, Tokyo*, 1989.

[17] S. Quinlan and O. Khatib. Towards real-time execution of motion tasks. In R. Chatila and G. Hirzinger, editors, *Experimental Robotics 2*. Springer Verlag, 1992.

[18] M. Tanie and G. Giralt (Eds). Europe-japan workshop on human-friendly robotics. In *Tokyo (Japan)*, November 1999.

[19] NASA Telerobotics Web-page. *http://ranier.oact.hq.nasa.gov/telerobotics_page/realrobots.ht* NASA, 1999.

[20] D.Van Zwynsvoorde, T.Simeon, R.Alami Building topological models for navigation in large scale environments *2001 IEEE International Conference on Robotics and Automation (ICRA'2001)*, Seoul (Korée), 21-26 Mai 2001, pp.4256-4261.

[21] Xavier. *vwww.cs.cmu.edu/afs/cs.cmu.edu/Web/People/Xavie* 1999.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Design of and Operational Experiences from Five Museum Robot Installations

Sjur J. Vestli

MRS Automation GmbH
Geeringstrasse 67, 8049 Zürich, Switzerland
s.vestli@ieee.org

1. Abstract

This paper describes the background of the robot systems used at the exhibitions: “Der Gehilfe - vom Dienstboten zum Service-Design”/“AL SEU SERVEI!, del majordom al servidor virtual” and the permanent exhibition at the SiemensForum in Zürich, Switzerland.We describe the design and implementation of the control software which was especially designed to be simple and robust. This has enabled commissioning of the robot systems by the museum staff (and not by expensive robot designers/service personnel). Furthermore details on how the robot systems have operated is provided.

2. Introduction

When designing the exhibition “Der Gehilfe - vom Dienstboten zum Service-Design” at the “Museum für Gestaltung Zürich” and the “Kunsthalle Krems” [1] the exhibition makers wanted to have a live robot integrated as part of the exhibition[1]. Eventually the author was approached and a system designed around the SmartROB-2 mobile robot platform was proposed. The SmartROB-2 mobile robot platform is an improvement of the mobile robot kit described in [3], and it is commercially available from MRS Automation GmbH.

As a results of the discussions between the robot designer and the exhibition designers a robot system that navigates around the exhibition and at pre-defined points (waypoints) reads out text relevant to the exhibition items close by. In addition to fulfilling the above functional requirement the following aspects must also be considered by the robot designer (more-or-less order of importance):

- The robot system may not bump into anything or anybody
- The robot system must be as cheap as possible
- The robot system must operate autonomously (i.e. without human intervention) for 1 full exhibition day (maintenance etc. only in the evenings)
- The robot system must be installable at a new exhibition site by non-technical personnel (the exhibition was designed from the outset to be movable and presented by various museums)
- The order of the texts to be spoken must be changeable, since, due to the specifics of the exhibition rooms of the various museums, the exhibition could not be kept identical at all the venues
- The texts themselves must be changeable

Based on these requirements a robot system based on the SmartROB-2 platform was proposed which will be described in detail subsequently. The robot system, as a part of the exhibition, was installed, and used during the following exhibition venues (see also figure 1 for a visualization of this robot system). The robot was given the name “museomobil” by the exhibition designers.

- February to May 2000, Museum für Gestaltung Zürich, Ausstellungsstr. 60, CH-8005 Zürich (9 weeks, 3 days/week: 10h, 3 days/week: 7h).
- December 2000 to February 2001, Kunsthalle Krems, Franz Zeller Platz 3, A-3500 Krems-Stein (8 weeks, daily, 8 hours per day)
- May to August 2001, Museum für Kunst und Kulturgeschichte, Hansastraße 3, D-44137 Dortmund (12 weeks, 4 days/week: 7 h, 1 day/week: 10h, 1 day/week: 5 h)
- April to June 2002, MUSEU DE LES ARTS DECORATIVES, Palau Reial de Pedralbes, Av. Diagonal, 686, E-08034 Barcelona (10 weeks, 5 days/week 8h, 1 day/week: 5h). Exhibition ran under the name: “AL

---

[1] The exhibition, compatible with its theme, already had current robots for service tasks on display. However, the suppliers of these robots did not allow the exhibition makers to “show them live”.

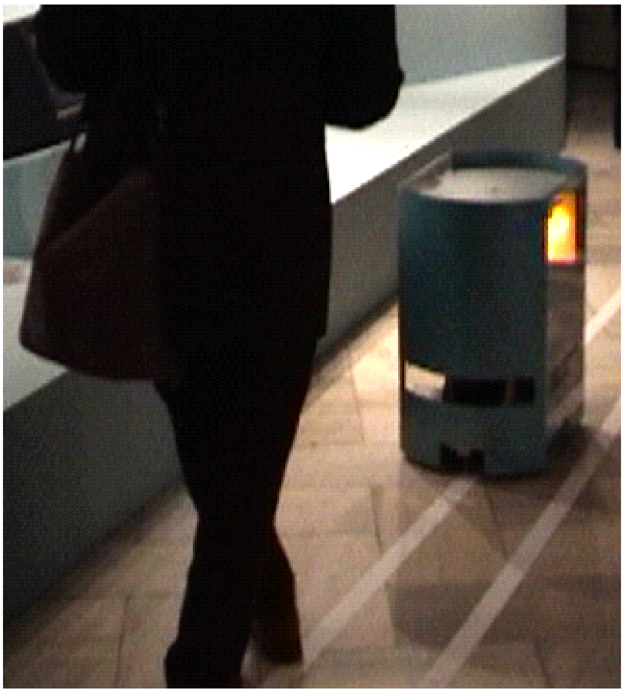

Figure 1. *The “museomobil” at the opening of one of the exhibitions. The guide wires can be seen as white stripes. The conductive part of the guide was realized using “copper tape”*

SEU SERVEI!, del majordom al servidor virtual”

As a direct result of the exhibition in Zürich the robot designer was approached by Siemens in Zürich, Switzerland to design a similar robot system integrated in the permanent exhibition at the SiemensForum in Zürich. In particular the robot system should be integrated into the part of the exhibition which shows the Automation and Drives (A&D) business segment of Siemens. This robot system has been operational since November 2000 at the Siemens-Forum, Freilagerstrasse 28, CH-8047 Zürich (see also figure 2 for a visualization

Compared to the system used at the exhibition “Der Gehilfe - vom Dienstboten zum Service-Design” a number of improvements was required:

- Whereas the batteries of the original design was charged using an external charger and connecting the charger was done by hand in this case the charging was to be automatic
- The actions of the robot was to partly controlled by an external PLC[1] system supplied by the A&D group of Siemens
- The robot was to have a more interesting design (provided by Siemens) which meant that additional digital IO had to be provided by the robot controller

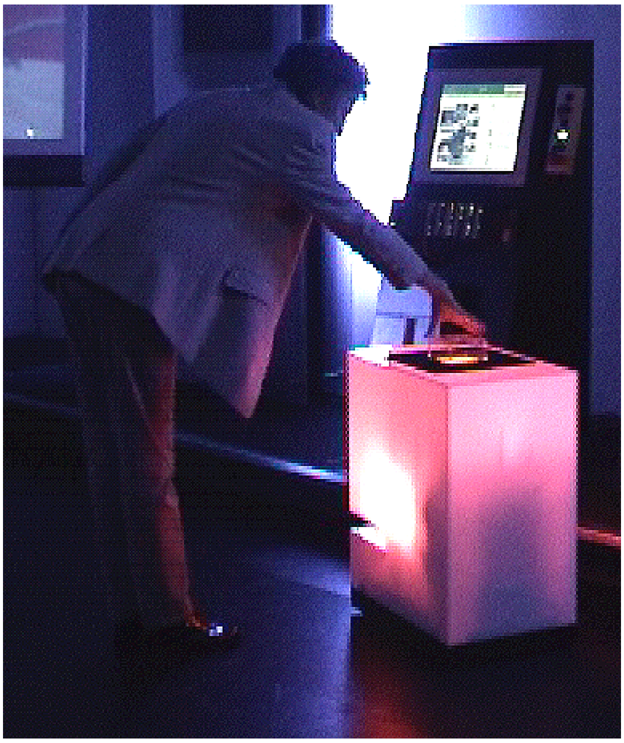

Figure 2. *“Kasimir” interacting with the curator of the “SiemensForum” in Zürich.The guide wires are invisible, as they are put beneath the double flooring.*

3. Detailed Systems Description

Based on the constraints described above the SmartROB-2 platforms was selected as a basis for the developments. For the navigation along the fixed course on inductive wire was chosen, a system still widely used in the AGV industry. The on-robot antenna and signal processing electronics generates a bi-polar analogue signal that is proportional to the displacement of the guide wire from the centre of the antenna. The guidance wire consists on one single loop which goes via all the waypoints. At the waypoints a separate frequency is present along the guide wire. The robot system moves in only one direction along the guidance wire. With the differential drive system employed in the SmartROB-2 a simple PID controller can be designed to follow the wire using as control output the rotational velocity of the robot.

A SICK PLS Laser Scanner, programmed with a fixed protective field and triggering a “stop” of the robot with its digital output was employed to ensure that the “no-collision” requirement was met. The digital signal from the SICK is read by software and based on this signal the translational velocity component of the robot.

With this scheme, installation essentially consists of laying the guide wire and the signal wires. A task that was accomplished by the personnel of the museums/exhibitions without onsite support of the system supplier[2].

---

[1] Programmable Logic Controller, an embedded computer specifically designed for industrial automation.

Since both low cost and high flexibility is demanded of the spoken texts it was decided against integrating any special hardware for replay of spoken text (apart from conventional active loudspeaker as can be connected to any desktop PC). Instead the real-time capability of the existing robot controller and one free digital to analogue (DA) converter (connected to line in of the active loudspeakers) was exploited. Under the XO/2 operating system [2] in use, installing a high frequency (10s of kHz) task that reads from an audio file and generates the corresponding signal on the DA converter. Since no HW is available to decode highly compressed formats a “simple” format like WAV was selected. Obviously with this scheme it is not possible to generate hi-fi quality sound, furthermore extensive storage space is needed to store the audio data. Since all the hardware is already available on the robot controller this functionality could be implemented without causing higher system costs. Changing the texts to be spoken means recording and encoding the texts as WAV files and writing these onto the file system of the robot controller. Physically writing the file system (which also determined the order in which the texts are played) must be done by the robot designer. As this is only necessary once for each exhibition this was deemed acceptable[1].

The texts are associated with the aforementioned waypoints and the waypoints are detected by the presence of a separate frequency along the guide wire. As the number of different frequencies that can be separated is limited one particular frequency pattern may be available at more than one waypoint. By internally to the robot software keeping track of where the robot has been the waypoints are uniquely identified. When the waypoint has been identified a separate process is started to replay the audio file associated to this point (or indeed take any other action).

The standard SmartROB-2 kit has a on board battery supply of 24V and 12 Ah (gel lead acid). Initial experiments showed that this would be adequate for approximately 5 hours operations. Whereas 10 hours continuos operation is required for the museomobil the specifications were relaxed so that one battery change per day was allowed. Care was taken so that such a battery change can be done quickly by the museum guards (or other museum personnel).

The robot systems must be operated by non-technical personnel. Hence, no complex startup scheme can be tolerated, nor can a lengthy startup. This has been achieved by using traditional embedded system design techniques. The whole robot system boots from on-board flash containing a statically pre-linked image. Additionally this flash also contains a file system where the text to be spoken is located. In this way the system is operational within 30 seconds and the only interaction between the personnel and the robot is the on/off switch (which is hidden so as to avoid tempering by the visitors) and (in the case of the robot for the “Der Gehilfe - vom Dienstboten zum Service-Design”) the battery connector. The signal generator is permanently switched on. The feedback to the operator indicating that the robot is operational is that it centers itself over the guide wire and, if un-obstructed, starts driving along the guide wire.

For Kasimir such operation was undesirable. The exhibition in the SiemensForum is not permanently open but guided tours by the museum personnel are given on request. Furthermore, there is no personnel resources available that can ensure that the batteries are always charged. Therefore the system was enhanced with a docking unit for charging the batteries. Also, controlled by the external PLC, the robot is not continuously moving but started (it stops automatically) and sent to a given waypoint. At one of these waypoints the contacts for the charger is located.

The operation of Kasimir is as follows: 1) move to waypoint a and stop, 2) request a special cube to be placed on top of robot (done by visitor), 3) move to waypoint b and stop, 4) external vision system tries to identify the cube and plays a video depending on which cube was picked by the visitor, 5) move to start position (where docking mechanism for the charger is) and stop. To ensure correct operation the PLC can query the robot for the presence of the cube (mechanical switch), the position of the robot along the guide wire.

The navigation methods and general operation of these robot systems contrasts with what is found in the academic literature, which tend to be far more complex (see for example [4]). Although we agree that our system is limited in terms of its navigation we believe that free-ranging laser-navigating systems are currently too costly, both in terms of hardware costs and installation costs, for many installations.

4. User Feedback & Lessons Learned

The museomobil has, at its 4 venues, covered a distance of 677 km (assuming 1883 hours running time and an average speed[2] of 0.1 m/s). During this whole period there have been no reports of fundamental malfunctioning of the control software.

The reported issues have been:

- The quality and loudness of the replayed texts are not adequate. This has been reported also by a number of visitors

---

[2] With one exception, and the robot designer always wanted to go to Barcelona.

[1] Experience also showed that mechanical wear parts (wheels, bearings, gears, etc.) needed replacement, and this service could easily be combined with reprogramming of order and content of the texts (if necessary).

[2] The un-obstructed speed of the robot is programmed to be 0.2 m/s, the estimate of an average speed of 0.1 m/s caters to stopping because of obstruction and battery change. Unfortunately no more accurate data is available.

- The battery capacity should be higher to allow 10h operation without recharging
- The battery lifetime (number of charge/discharge cycles) should be higher. In Barcelona the robot system failed at the end of the exhibition because of this and due to logistics issues it could not be resolved.
- The wear parts should be (a bit) more robust. At the end of the exhibition in Zurich the behaviour was reported as “not very good” (however no service call was requested). The inspection afterwards showed damaged bearings.

Much to our surprise we have not been confronted with complaints on the lack of “intelligence” and flexibility. This can be explained by the fact that the expectation to the robot system in the exhibition is not very high (no special mention of the presence of a museum guide robot in the promotional material of the exhibitions). Also the robot was designed into the exhibition and in the press clipping it is rarely mentioned and as such is not perceived so much as “A Robot” but is a part of the exhibition.

There has also been no, known to us, complaints on the presence of the, rather ugly, guide wires.

In the case of Kasimir the distance travelled is not known, but believed to be significantly smaller. Partly the reported issues have been similar. In particular the battery lifetime has not been satisfactory, this might be due to the different duty cycle/operation of the robot.

It has been mentioned previously that the frequencies indicating the waypoints is not unique. This occasionally causes the robot to “believe” that it is at another position than it, in reality, is at.

In the case of the museomobil it recovers at one unique location along the loop. The texts have been spoken at “the wrong places”, however this happens, in the worst case, one full loop before the process recovers, and there has been no reports of this (it might still have happened, nobody noticed).

In the case of Kasimir this is known to happen every now and then. Also here the robot recovers by the end of the loop, however, the demo and interaction with the PLC is in such cases less than impressive.

There has also been one reported case of a broken guide wire (after 1 year of operation). Upon inspection it was impossible to detect any crack or other source of bad contact and after disconnecting and reconnecting the signal generator the system worked as designed.

5. Conclusion and Outlook

By relying upon well proven navigation methods and the simplest of obstacle avoidance scheme we have installed museum guide robots at five different sites. In general the response from the users have been positive, although a few reliability issues still remain.

The reliability issues can, we believe, be mastered by using well proven techniques (e.g. unique transponders for position), however, due to financial reasons, there has been no opportunities to verify this.

The museomobil has been a demonstration of concept and apart from costs of replacement parts and expenses incurred at one installation there has been no financial transactions involved.

Kasimir has been a commercial projects and, by focusing on low system cost, has at least covered its expenses.

It would be better in many respects to replace the inductive navigation by laser navigation. In industrial AGV projects this has been done by MRS Automation GmbH. Until now there have been no opportunities to test this navigation in museum environments.

6. Acknowledgements

The author would like to thank Prof. Dr. G. Schweitzer for initially supporting these efforts whilst the author was still at the Institute of Robotics, ETH Zürich, Switzerland, A. Tönnies of the SiemensForum, The staff of Museum für Gestaltung Zürich, Kunsthalle Krems, Museum für Kunst und Kulturgeschichte, Dortmund MUSEU DE LES ARTS DECORATIVES, Barcelona and F. Wullschleger for support and the workshop organizers for their patience.

7. References

[1] N.N.: “DER GEHILFE, Vom Dienstboten zum Service-Design”, with contributions by: Erika Keil, Oliver Kielmayer, Birgit Mager, Uli Marchsteiner, Ramón Reichert, Gerhard Schweitzer, Robert Trappl / Paolo Petta, Peter Stalder, Andreas Volk, Klaus Weyh. Museum für Gestaltung Zürich und Kunsthalle Krems 2000, ISBN 3-907065-84-0.

[2] Brega. R.: A Combination of System Software Techniques Aimed at Raising the Run-Time Safety of Complex Mechatronic Applications. Dissertation ETH Nr. 14513, Zürich, 2002.

[3] Siegwart R.: Grasping the interdiciplinarity of Mechatronics, IEEE Robotics and Automation Magazine, June 2001, pp 27-34.

[4] Tomatis N., Terrien G., Piguet R., Burnier D., Bouabdhalla S., Siegwart R.: Design and System Integration for the Expo.02 Robot, Work-shop on Robots in Exhibitions, IEEE/RSJ International Conference on Intelligent Robots and Systems, Lausanne, Switzerland.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

On the Prospects of Robots in Museums

Dr. Andrea Niehaus, Deutsches Museum Bonn,
Ahrstraße 45, 53175 Bonn, Germany

Introduction

Access to exhibits is a crucial issue in museums that is recently approached under new, technological perspectives. This development particularly took place in technical museums. In the last years the introduction of media-technologies could be observed, such as media- and Web-presentation and graphical animations. More and more Web-technologies are employed giving access to selected exhibits. The Web-visitors sometimes even has the opportunity to have a virtual tour through some or all departments. In a very few cases, high-end technologies, such as virtual reality or robotics, are employed in museums. This contribution will focus on the prospects of robots in museum, an emerging field of imparting knowledge.

Robots in Museums within the changing concepts of imparting knowledge

Collecting (exhibits), retaining (exhibits) and imparting (the knowledge on them) are the central functions of a museum. The first two functions have remained unchanged until now, but the way of imparting knowledge to visitors underwent a considerable change. This new situation is due to social and technological innovation. The concept of imparting knowledge has changed from the museum of the 19th century focusing on dry and pitiless instruction to an environment of active and easy learning. Museums have to consider this, if they want to remain interesting places to see for their visitors. Museum robots can be excellent tools in the framework of new concepts in imparting knowledge. Examples from museum's practices will elucidate this issue.

Implementation

Possessing such an excellent tool is one thing, implementing it into a running museum programme is another. Introducing a robot in a museum is a conceptual challenge in the first place whereas the complex social framework and economic implications should not be underestimated. Social aspects include: the concept of the robot itself (its tasks in the museum), its interaction with the public (emotional address, appropriate behaviour in interactions with visitors) and its social integration in the museum (human tour guides vs. robot tour guides). Particularly the latter is an important issue that may be decisive for the museum management to buy a museum robot. The conceptual aspects of the robot also play an important role in the acquisition of financial means for its purchase. Potential sponsors are easily convinced to contribute to such a project if the implementation of the robot in the museum is attractive, achieving a successful balance between conceptual, social and technical aspects (such as a high reliability and a user-friendly system).

Conclusion

A short insight into these issues and their interdependencies will be provided to draw some first conclusions on robots in museums as exciting and fascinating new approaches in forward-looking concepts of imparting knowledge.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Demonstrating the Humanoid Robot *HERMES* at an Exhibition: A Long-Term Dependability Test

Rainer Bischoff and Volker Graefe

The Institute of Measurement Science
Bundeswehr University Munich, 85577 Neubiberg, Germany
Bischoff@ieee.org, Graefe@unibw-muenchen.de
http://www.unibw-muenchen.de/hermes

Abstract

HERMES, a laboratory prototype of a humanoid service robot, served in a museum, far away from its home laboratory, for more than six months up to 18 hours per day. During this period the robot and its skills were regularly demonstrated to the public by non-expert presenters. Also, HERMES interacted with the visitors, chatted with them in English, French and German, answered questions and performed services as requested by them. Only three major failures occurred during the 6-months-period, all of them caused by failures of commercially available modules that could easily be replaced.

Key to this success was the dependability that had been designed into HERMES. We introduce the concept of dependability and describe the design strategies that have led to a high degree of dependability of our robot. To be accepted by society and to be entrusted with important or even critical services, future service robots must be similarly dependable as today's cars or telephones. We argue that true dependability of complex intelligent robots can only be achieved by actually building and integrating prototypes and subjecting them to long-term tests with outsiders and away from their home laboratories. In fact, by demonstrating HERMES in the museum, at trade fairs and in TV studios we have learned valuable lessons, especially regarding the interaction of a complex service robot with unknown humans.

1 Introduction

Exhibitions offer excellent opportunities for studying and evaluating a robot's communication skills and dependability under real-world conditions, especially if the robot is exposed to the public, and allowed to interact with it, for extended periods of time. However, to have a chance of surviving such a long-term test at an exhibition without annoying failures, a robot must be much more dependable than a typical research robot in a laboratory.

This requirement is probably the reason that, to the best of our knowledge, only two research groups have ever undertaken long-term experiments with their robots interacting with strangers outside their own laboratories. One, the museum tour guide, Sage, installed by the group of Nourbakhsh at the Carnegie Museum of Natural History in Pittsburgh [Nourbakhsh et al. 1999], and two, the entertaining robots of Fraunhofer IPA [Graf et al. 2000], still working in the entry hall of the telecommunications museum in Berlin. Both projects accumulated valuable experience in non-expert operation in a crowded environment well over a year. During the World Exposition 2000 in Hannover, 72 mobile robots (size 1.6 to 4.5 meters) were constantly moving freely on a surface of 5000 m^2 with speeds up to 0.25 m/s while reacting to the presence of visitors and coordinating themselves in relation to each other [BBM Expo 2000]. Unfortunately, up to date we have not become aware of any scientific report on this experiment. Similar tests were carried out by Thrun and Burgard with the robots RHINO [Burgard et al. 1999] and MINERVA [Thrun et al. 2000], albeit under the supervision of experts and only for a few days. Long-term experiments with mobile robots in their respective institute environments were carried out by [Simmons et al. 1999] at the Robotics Institute (CMU, Pittsburgh) with the robot XAVIER, one of the first mobile robots controllable via a Web interface, and by a research group at the Institute of Robotics (ETH, Zürich) with a mobile mail distribution system called MOPS [Tschichold et al. 2001]. Commercially available robots that do not possess complex interaction interfaces, but are nonetheless easy to operate and have been exposed to a general public, are the Helpmate robot [King, Weiman 1990], that was installed in dozens of hospitals world-wide, and a cleaning machine equipped with a Siemens Corporation navigation system, still working in a supermarket in the Netherlands [Endres et al. 1998].

There might be other groups that have been carrying out similar experiments, but the fact that those experiments have not been reported at major conferences shows that integration and dependability issues as well as long-term experiments are not yet considered important and interesting problems, neither in the robotics research community nor by the funding agencies or bodies. Also, the projects listed above focused primarily on navigation and more or less simple human-robot communication (more complex in case of MINERVA and RHINO). – We wonder if service or personal robots will ever become valuable servants of our future society if not more robots are fielded for extended periods of time with a richer set of functionalities, a higher level of human-robot interaction and in realistic settings.

As we pointed out before, dependability is crucial for a robot to be able to serve at an exhibition, and also for future personal and service robots to be accepted by society. “Dependability” is a system concept that integrates such attributes as reliability, availability, safety, confidentiality, integrity, and maintainability [Laprie 1992]. The goals behind the concept of dependability are the abilities of a system to deliver a service that can justifiably be trusted and to avoid failures that are more frequent or more severe, and outage durations that are longer, than is acceptable to the user(s).

Figure 1: Humanoid experimental robot *HERMES*; 1.85 m x 0.7 m x 0.7 m; mass: 250 kg

Our society largely depends on infrastructures that are controlled by embedded information systems and the dependability concept has been widely employed for such systems. Although future service and personal robots are supposed to become an important part of our future society, dependability aspects have been largely neglected by researchers. However, dependability is needed especially for these types of robots because they are intended to operate in unpredictable and unsupervised environments and in close proximity to, or in direct contact with, people who are not necessarily interested in them, or, even worse, who try to harm them by disabling sensors or playing tricks on them.

It is one aim of this paper to raise the awareness for research on integration and dependability, and for long-term experiments. There is no other way to increase the dependability of service robots in the long run.

2 Designing for Dependability

In our opinion the dependability of a robot is not something that can be added on after the robot has been designed and built. Rather, it must be designed into the robot and, specifically, it emerges from the following design strategies:

1. Learning from nature how to design reliable, robust and safe systems
2. Providing natural and intuitive communication and interaction between the robot and its environment
3. Designing for ease of maintenance
4. Striving for a tidy appearance

These design strategies have guided us in the design and construction of our humanoid robot *HERMES* (Figure 1). They are explained in greater detail in the sequel.

Learning from nature. According to the classic approach, robot control is model-based. Numerical models of the kinematics and dynamics of the robot and of the external objects that the robot should interact with, as well as quantitative sensor models, are the basis for controlling the robot’s motions. The main advantage of model-based control is that it lends itself to the application of classical control theory and, thus, may be considered a straightforward approach. The weak point of the approach is that it breaks down when there is no accurate quantitative agreement between reality and the models. Differences between models and reality may come about easily; an error in just one of the many coefficients that are part of the numerical models can suffice.

Organisms, on the other hand, are robust and adapt easily to changes of their own conditions and of the environment. They never need any calibration, and they normally do not know the values of any parameters related to the characteristics of their “sensors” or “actuators”. Obviously, they do not suffer from the shortcomings of model-based control which leads us to the assumption that they use something other than quantitative measurements and numerical models for controlling their motions. Perhaps their motion control is based on a holistic assessment of situations for the selection of behaviors to be executed. Possibly robotics could benefit from following a similar approach.

Following this line of argumentation we strongly believe that sensing in general should be based on the senses that have proved their effectiveness in nature. Therefore, vision – the sensor modality that predominates in nature – is also an eminently useful and practical sensor modality for robots. Also, tactile sensing and hearing may greatly improve a robot’s safe operation as shown by nature.

Providing natural and intuitive communication and interaction. Any person who might encounter a service robot needs to be able to communicate and interact with it in a natural and intuitive way. Therefore, the communication interface has to be designed in such a way that no training would be required for any person who might get in contact with the robot. This can be achieved if the human-robot communication resembles a dialogue that could as well take place between two humans.

Designing for ease of maintenance. The first step to make a complex system dependable is to make its components reliable. Moreover, we believe that only a robot that needs little or no maintenance and that can be easily repaired (if ever needed) will be accepted as a co-worker, caretaker or companion.

Striving for a tidy appearance. It is a matter of personal experience that, especially in research environments, robots often fail because of broken cables and unreliable connections. Such robots often look very cluttered with cables criss-crossing each other, and circuitry and connectors hidden under bundles of wires. This not only makes visual inspection difficult, but it may also be taken as an indication that those who built and maintain the robot have placed little emphasis on a systematic design. Although software is not visible, the observer wonders whether the structure of the robot's software might resemble the layout of the robot’s wiring.

3 The Humanoid Robot *HERMES*

In designing our humanoid experimental robot *HERMES* we placed great emphasis on modularity and extensibility of both hardware and software [Bischoff 1997].

3.1 Hardware

HERMES has an omnidirectional undercarriage with 4 wheels, arranged on the centers of the sides of its base. The front and rear wheels are driven and actively steered, the lateral wheels are passive. The manipulator system consists of two articulated arms with 6 degrees of freedom each on a body that can bend forward (130°) and backward (-90°). The work space extends up to 120 cm in front of the robot. Currently each arm is equipped with a two-finger gripper that is sufficient for basic manipulation experiments.

Main sensors are two video cameras mounted on independent pan/tilt drive units in addition to the pan/tilt unit that controls the common “head” platform. The cameras can be moved with accelerations and velocities comparable to those of the human eye.

HERMES is built from 25 drive modules with identical electrical and similar mechanical interfaces yielding 22 degrees of freedom. Each module contains a motor, a Harmonic Drive gear, a micro-controller, power electronics, a communication interface and some sensors. The modules are connected to each other and to the main computer by a single bus.

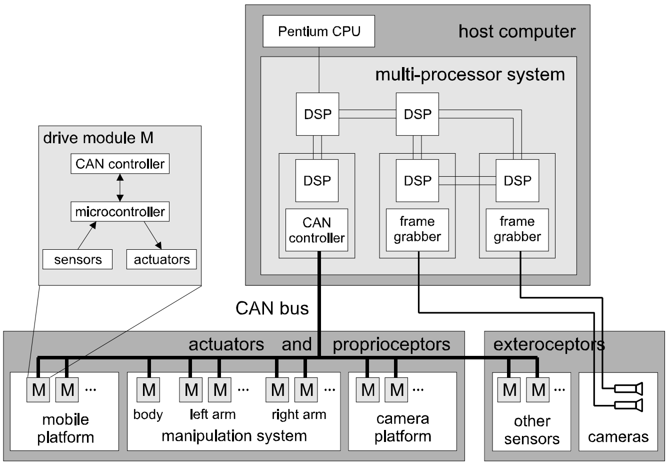

Figure 2: Modular and adaptable hardware architecture for information processing and robot control.

A hierarchical multi-processor system is used for information processing and robot control (Figure 2). The control and monitoring of the individual drive modules is performed by the sensors and controllers embedded in each module. The main computer is a network of digital signal processors (DSP, TMS 320C40) embedded in a ruggedized, but otherwise standard industrial PC. Sensor data processing (including vision), situation recognition, behavior selection and high-level motion control are performed by the DSPs, while the PC provides data storage, Internet connection and the human interface.

3.2 Software and System Architecture

Seamless integration of many – partly redundant – degrees of freedom and various sensor modalities in a complex robot calls for a unifying approach. We have developed a system architecture that allows integration of multiple sensor modalities and numerous actuators, as well as knowledge bases and a human-friendly interface. In its core, the system is behavior-based, which is now generally accepted as an efficient basis for autonomous robots [Arkin 1998]. However, to be able to select behaviors intelligently and to pursue long-term goals in addition to purely reactive behaviors, we have introduced a situation-oriented deliberative component that is responsible for situation assessment and behavior selection.

Figure 3 shows the essence of the situation-oriented behavior-based robot architecture as we implemented it. The situation module (situation assessment & behavior selection) acts as the core of the whole system and is interfaced via “skills” in a bidirectional way with all other hardware components – sensors, actuators, knowledge base storage and MMI (man-machine, machine-machine interface) peripherals.

These skills have direct access to the hardware components and, thus, actually realize behavior primitives. They obtain certain information, e.g., sensor readings, generate specific outputs, e.g., arm movements or speech, or plan a route based on map knowledge. Skills report to the situation module via events and messages on a cyclic or interruptive basis to enable a continuous and timely situation update and error handling.

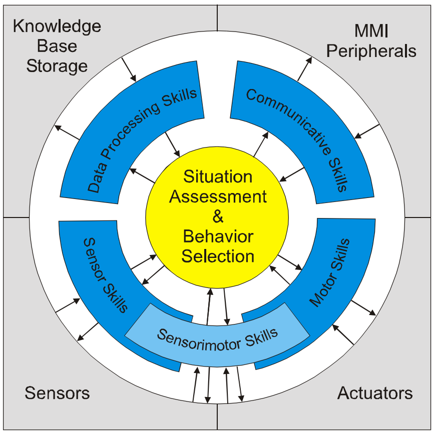

Figure 3: System architecture of a personal robot based on the concepts of situation, behavior and skill.

The situation module fuses via skills data and information from all system components to make situation assessment and behavior selection possible. Moreover, it provides general system management (cognitive skills). Therefore, it is responsible for planning an appropriate behavior sequence to reach a given goal, i.e., it has to coordinate and initialize the in-built skills. By activating and deactivating skills, a management process within the situation module realizes the situation-dependent concatenation of elementary skills that leads to complex and elaborate robot behavior. For a more profound discussion of our system architecture which bases upon the concepts of situation, behavior and skill see [Bischoff, Graefe 1999].

Several of the fundamental concepts developed at our Institute were implemented in *HERMES* and contribute to its remarkable dependability: e.g., an object-oriented vision system with the ability to detect and track multiple objects in real time [Graefe 1989] and a calibration-free stereo vision system [Graefe 1995]. Also, the sensitivities of the cameras can be individually controlled for each object or image feature, and several forms of learning assure adaptation to changing system parameters as well as working in new environments from scratch. Moreover, a speaker-independent speech recognition for several languages and robust dialogues form the basis for various kinds of human-robot interaction [Bischoff , Graefe 2002].

4 Experiments and Results

Since its first public appearance at the Hannover Fair in 1998 where *HERMES* could merely run (but still won “the first service robots’ race”!) quite a number of experiments have been carried out that prove the suitability of the proposed methods. Of course, we performed many tests during the development of the various skills and behaviors of the robot and often presented it to visitors in our laboratory. The public presentations made us aware of the fact that the robot needs a large variety of functions and characteristics to be able to cope with the different environmental conditions and to be accepted by the general public.

In all our presentations we have experienced that the robot’s anthropomorphic shape encourages people to interact with it in a natural way. As presented in the preceding sections, *HERMES* possesses several other promising features inside and outside that makes it intrinsically more reliable and safer than other robots. One of the most promising results of our experiments is that our calibration-free approach seems to pay off, because we experienced drifting of system parameters due to temperature changes or simply wear of parts or aging. These drifts could have produced severe problems, e.g., during object manipulation, if the employed methods relied on exact kinematic modeling and calibration. Since our navigation and manipulation algorithms only rely on qualitatively (not quantitatively) correct information and adapt to parameter changes automatically, the performance of *HERMES* is not affected by such drifts.

In the sequel we concentrate on demonstrations that we performed outside the familiar laboratory environment, namely in television studios, at trade fairs and in a museum where *HERMES* was operated by non-experts for an extended period of time. Such demonstrations, e.g., in television studios, subjects the robot to various kinds of stress. First of all, it might be exposed to rough handling during transportation, but even then, it should still function on the set. Second, the pressure of time during recording in a TV studio requires the robot to be dependable; program adaptation or bug-fixing at the location is not possible. *HERMES* has performed in TV studios a number of times and we have learned much through these events. We found, for instance, that the humanoid shape and behavior of the robot raise expectations that go beyond its actual capabilities, e.g., the robot is not yet able to act upon a director’s command like a real actor (although sometimes expected!). It is through such experiences that scientists get aware of what “ordinary” people expect from robots and how far, sometimes, these expectations are missed.

Trade fairs, such as the Hannover Fair, the world’s largest industrial fair, pose their challenges, too: hundreds of moving machines and thousands of people in the same hall make an incredible noise. It was an excellent environment for testing the robustness of *HERMES’* speech recognition system.

Last but not least, *HERMES* was field-tested for more than 6 months (October 2001 - April 2002) in the Heinz Nixdorf MuseumsForum (HNF) in Paderborn, Germany, the world’s largest computer museum. In the special exhibition “Computer.Brain” the HNF presented the current status of robotics and artificial intelligence and displayed some of the most interesting robots from international laboratories, including *HERMES*.

We used the opportunity of having *HERMES* in a different environment to carry out experiments involving all of its skills, such as vision-guided navigation and map building in a network of corridors; driving to objects and locations of interest; manipulating objects, exchanging them with humans or placing them on tables; kinesthetic and tactile sensing; and detecting, recognizing, tracking and fixating objects while actively controlling the sensitivities of the cameras according to the ever-changing lighting conditions. *HERMES* was able to chart the office area of the museum from scratch upon request and delivered services to *a priori* unknown persons (Figure 4). In a guided tour through the exhibition *HERMES* was taught the locations and names of certain exhibits and some explanations relating to them. Subsequently, *HERMES* was able to give tours and explain exhibits to the visitors. *HERMES* chatted with employees and international visitors in three languages (English, French and German). Topics covered in the conversations were the various characteristics of the robot (name, height, weight, age, ...), exhibits of the museum, and actual information retrieved from the World Wide Web, such as the weather report for a requested city, or current stock values and major national indices. *HERMES* even entertained people by waving a flag that had been handed over by a visitor; filling a glass with water from a bottle, driving to a table and placing the glass onto it; playing the visitors' favorite songs and telling jokes that were also retrieved from the Web (Figure 5).

5 Lessons Learned

We found it interesting to observe how *HERMES*, actually just a laboratory prototype despite its designed-in dependability, survived the daily hard work far away from its "fathers", where no easy access to repair and maintenance was available, and how it got along with strangers and even with presenters who did not know much about robot technology. In fact, we were surprised ourselves that it performed so well. During 6 months of operation (lasting up to 18 hours a day during video recordings for documentation purposes) only one motor controller, one drive motor and one audio amplifier ceased to function, all of them commercially available and easily replaceable. According to the museum staff, *HERMES* was one of the few robots at the show that could regularly be demonstrated in action, and among them it is considered the most intelligent and most dependable one. This statement is supported by the fact that the museum staff never called for advice once the initial setup was done. We had expected to give much more support and wondered how often we would have to travel from Munich to Paderborn (a six-hour-drive, one way) to help. Actually, we only were in Paderborn for setting up the robot for the exhibition, for presenting and documenting our research work during the first two weeks after the exhibition's opening and for 4 days of documentation work in December.

Preparing the robot for the exhibition was indeed fun, but also a lot of work: it made us realize that many operational details had never been documented before, such as powering the robot on and off, charging the batteries, starting the main program and testing functionality. Now they had to be written down in a manual for non-experts, i.e., people with little engineering background. Actually, the museum staff had insisted on having such a reference guide, but as a matter of fact, it shared the fate of most reference manuals in the world: it was almost never looked at, because people rather like to try out how things work instead of studying manuals, which makes the need for safe behavior even more evident.

Being afraid that the robot might come back to our university in pieces, we had made an effort to finish many of the laboratory's research projects before sending *HERMES* to the museum. Actually, such time pressure helped to speed up work on algorithms and implementation details.

Although we knew that thorough testing is only possible in different environments with numerous different people interacting with the robot, we had never before really been able to do so over an extended period of time. This exhibition gave us the opportunity, and eventually, it proved that our concepts and approaches (as presented

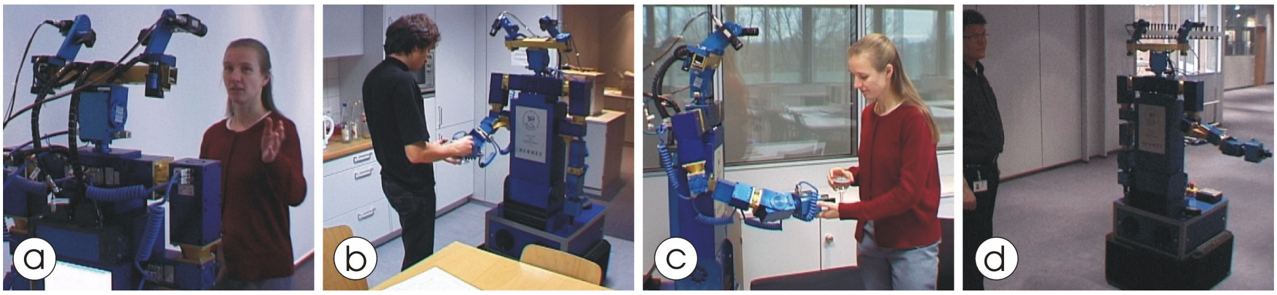

Figure 4: *HERMES* executing service tasks in the office environment of the Heinz Nixdorf MuseumsForum: (a) dialogue with an *a priori* unknown person with *HERMES* accepting the command to get a glass of water and to carry it to the person's office; (b) asking a person in the kitchen to hand over a glass of water; (c) taking the water to the person's office and handing it over; (d) showing someone the way to a person's office.

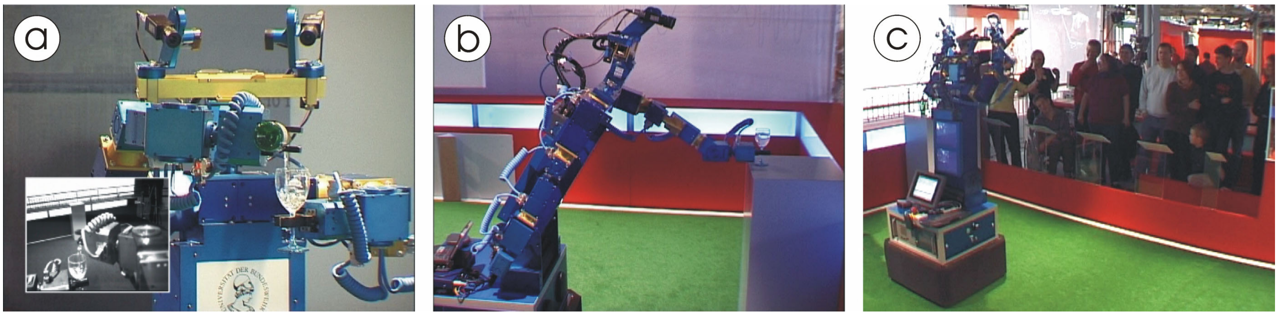

Figure 5: *HERMES* performing at the special exhibition "Computer.Brain", instructed by natural language commands: taking over a bottle and a glass from a person (not shown), filling the glass with water from the bottle (a); driving to and placing the filled glass onto a table (b); interacting with the visitors (here: waving with both arms, visitors wave back!) (c)

in chapters 2 and 3) were correct. Consequently, to *really* see the robot working in a completely different environment, operated by non-experts for over 6 months, was certainly the most valuable experience of this long-term experiment. Some behaviors worked much better in the new environment than in our institute, others worse. For example, navigation worked much better on the one hand because the floor was not as reflective as our institute's floor. On the other hand, the overall lighting conditions were rather poor and in the actual exhibition area it was almost too dark to navigate by means of vision. Although a large part of the exhibition featured red and yellow walls and a grey floor, it was very difficult for our monochrome vision system to distinguish between walls and floors. A color vision and a higher dynamic range of the cameras would certainly be desirable for our robot.

Especially children liked interacting with the robot. Surprisingly enough, the robot could understand the children's high voices and sometimes not fluently spoken phrases. They even hugged the robot, albeit under close supervision of the staff, without being afraid of breaking something, and, much more important, being afraid of being hurt by such a massive chunk of moving metal. Adults, on the other hand, faced the robot with all due respect.

Some people pushed the robot's emergency button that was clearly visible in the back of the robot, and expected something to happen. Since the emergency button only disconnects the motors from the power but not the computers, a lengthy reboot procedure was not required. The staff just had to pull up the emergency button again to restart the robot. We know now that the state of the emergency button should be monitored by the robot in order to react adequately to such a situation.

The funniest interaction for most of the visitors and the staff alike resulted from touching the tactile bumpers placed around the robot's undercarriage. The robot was programmed to stop moving and to say "Ouch". This simple "emotion" made most of the people smile, and kept them touching the bumpers more than once. On the other hand, behaviors that the developers considered more impressive, such as navigation and manipulation, were taken for granted. The interaction capabilities on top of assumed (normal) behavior is what most people are interested in. Certainly, this does not simplify the robot scientist's work since his robots obviously have to "compete" with the well-known robots from science fiction movies.

According to a museum press release, more than 80.000 visitors had been attracted by the special exhibition "Computer.Brain" which was 30.000 more than had been hoped for. The maximum capacity of the museum was reached on several days, leading to long waiting lines. This tremendous success is certainly due to the highly interactive character of the exhibition. Of the 330 exhibits 52 were interactive, the most spectacular ones being robots. The overall exhibition's media presence was remarkable with 18 independent broadcasts in televison (not counting reruns) and 11 in the radio, in addition to an uncountable number of newspaper articles. Taking media presence as an important indicator for successful and well recognized work, our project was indeed quite successful: to our knowledge *HERMES* was featured to a larger extent at least 6 times in TV, twice in radio and 18 times in newspaper articles (most of them during the two weeks after the exhibition's opening).

6 Summary and Conclusions

HERMES, an experimental robot of anthropomorphic size and shape, interacts dependably with people and their common living environment. It has shown robust and safe behavior with novice users, e.g., at trade fairs, television studios, at various demonstrations in our institute environment, and in a long-term experiment carried out at an exhibition and in a museum's office area. The robot is basically constructed from readily available motor modules with standardized and viable mechanical and electrical interfaces. Due to its modular structure the robot is easy to maintain, which is essential for system dependability. A simple but powerful skill-based system architecture is the basis for software dependability. It integrates visual, tactile and auditory sensing and various motor skills without relying on quantitatively exact models or accurate calibration. Actively controlling the sensitivities of the CCD cameras makes the robot's vision system robust with respect to varying lighting conditions (albeit not as robust as the human vision system). Consequently, safe navigation and manipulation, even under uncontrolled and sometimes difficult lighting conditions, were realized. A touch-sensitive skin currently covers only the undercarriage, but is in principle applicable to most parts of the robot's surface. *HERMES* understands spoken natural language speaker-independently, and can, therefore, be commanded by untrained humans.

In summary, *HERMES* can see, hear, speak, and feel, as well as move about, localize itself, build maps and manipulate various objects. In its dialogues and other interactions with humans it appears intelligent, cooperative and friendly. In a long-term test (6 months) at a museum it chatted with visitors in natural language in German, English and French, answered questions and performed services as requested by them.

Although *HERMES* is not as competent as the robots we know from science fiction movies, the combination of all before-mentioned characteristics makes it rather unique among today's real robots. As noted in the introduction, today's robots are mostly strong with respect to a single functionality, e.g., navigation or manipulation. Our results illustrate that many functionalities can be integrated within one single robot through a unifying situation-oriented behavior-based system architecture. We also believe that our simple design strategies, such as modularity, calibration-free control and truly human-like interaction, would enable other researchers, too, to build similarly dependable robots. Our results suggest that testing a robot in various environmental settings, both short- and long- term, with non-experts having different needs and different intellectual, cultural and social backgrounds, is enormously beneficial for learning the lessons that will eventually enable us to build dependable personal robots.

7 Literature

Arkin, R. C. (1998): Behavior-Based Robotics. MIT Press, Cambridge, MA, 1998.

BBM Expo (2000): Future Intelligence – Frequently Asked Questions Available at http://www.bbm-ww.de /expo2/english.zip (last accessed July 21, 2002).

Bischoff, R. (1997): *HERMES* – A Humanoid Mobile Manipulator for Service Tasks. Proc. of the Intern. Conf. on Field and Service Robotics. Canberra, Australia, Dec. 1997, pp. 508-515.

Bischoff, R.; Graefe, V. (1999): Integrating Vision, Touch and Natural Language in the Control of a Situation-Oriented Behavior-Based Humanoid Robot. IEEE Conference on Systems, Man, and Cybernetics, October 1999, pp. II-999 - II-1004.

Bischoff, R.; Graefe, V. (2002): Dependable Multimodal Communication and Interaction with Robotic Assistants. Proceedings 11th IEEE International Workshop on Robot and Human Interactive Communication (ROMAN 2002). Berlin, September 2002 (to appear).

Burgard, W.; Cremers, A. B.; Fox, D.; Hähnel, D.; Lakemeyer, G.; Schulz, D.; Steiner, W.; Thrun, S. (1999): Experiences with an interactive museum tour -guide robot. Artif. Intelligence, 114(1-2):3-55.

Endres, H.; Feiten, W.; Lawitzky, G. (1998): Field test of a navigation system: Autonomous cleaning in supermarkets. Proc. IEEE Int. Conference on Robotics and Automation, ICRA 1998, pp. 1779-1784.

Graefe, V. (1989): Dynamic Vision Systems for Autonomous Mobile Robots. Proc. IEEE/RSJ Int. Works. on Int. Robots and Systems, IROS '89. Tsukuba, pp. 12-23.

Graefe, V. (1995): Object- and Behavior-oriented Stereo Vision for Robust and Adaptive Robot Control. Intern. Symp. on Microsystems, Intelligent Materials, and Robots, Sendai, pp. 560-563.

Graf, B.; Schraft, R.D.; Neugebauer, J. (2000): A Mobile Robot Platform for Assistance and Entertainment. Proceedings of 31st International Symposium on Robotics, ISR-2000, Montreal.

King, S.; Weiman, C. (1990): Helpmate autonomous mobile navigation system. Proceedings of SPIE Conference on Mobile Robots, Vol. 2352, Boston, pp. 190-198.

Laprie, J. C. (1992): Dependability: Basic Concepts and Terminology. Springer-Verlag, 1992.

Nourbakhsh, I.; Bobenage, J.; Grange, S.; Lutz, R.; Meyer, R.; Soto, A. (1999): An Affective Mobile Educator with a Full-time Job. Artificial Intelligence, 114(1-2), pp. 95-124.

Simmons, R.; Fernandez, J.; Goodwin, R.; Koenig, S.; O'Sullivan, J. (1999): Xavier: An autonomous mobile robot on the web. Robotics and Auotmation Magazine, 1999.

Tschichold, N., Vestli, S., Schweitzer, G. (2001): The Service Robot MOPS: First Operating Experiences. Robotics and Autonomous Systems 34:165-173, 2001.

Thrun, S.; Beetz, M.; Bennewitz, M.; Burgard, W.; Cremers, A. B.; Dellaert, F.; Fox, D.; Hähnel, D.; Rosenberg, C.; Roy, N.; Schulte, J.; Schulz, D. (2000): Probabilistic algorithms and the interactive museum tour-guide robot minerva. Intern. Journal of Robotics Research, Vol. 19, No. 11, pp. 972-999.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Experiments at Trade Fairs with Blacky the Robot

Fernando Matia, Diego Rodriguez-Losada, Ramon Galan, Agustin Jimenez

UPM-DISAM, Jose Gutierrez Abascal 2, E28006 - Madrid, SPAIN
{*matia,drodri*}*@etsii.upm.es*

Abstract

Blacky has been tested in exhibition-like contests, more precisely, in two trade fairs and one competition event. Presently, its main task is as tour-guide and entertainment, making an special emphasis in artificial intelligence techniques for human-robot interaction capabilities, a crucial point for robot acceptance by humans. The robot has really worked in long-term experiments, where system integration and safety issues have been taken into account. The navigation algorithms, as well as the lessons learnt, are described in the paper, focussing on robot movements in an indoor, populated, complex and low structured environment. The market point of view is also analysed with the help of an exhibition organizer.

1 Project goals

The main goals of this work are threefold:

- To design a mobile robot able to operate in complex environments, guaranteeing real working and safety.
- To implement an artificial intelligence reasoning system added to the robot to give a feeling of intelligence, both in dialogues and in robot consciousness of the situation.
- To have at the end of the project a local demonstration at UPM-DISAM laboratories running frequently at scheduled times. Meanwhile, tests are being done at real fairs and exhibitions.

Principal uses of the robot are:

- Tour-guiding: the user can select a programmed tour or ask for a customized one. In both cases the robot decomposes the problem in reaching specific targets.
- Tele-visit: a remote user can have a feeling of the exhibition site navigating through the robot. He/she can ask for high level tour-guiding commands, or tele-operate the robot.
- Entertainment: people present at the exhibition site can enjoy the behavior of the robot while it moves around without a predefined task to do. This point is extremely important, because it provides an additional reason to be commercially feasible.

2 Description of the robotic system

2.1 The mobile robot

A MRV4 platform from Denning Branch, Inc. was used for the experiments. It is 70cm diameter and 140cm tall, having a ring of 24 sonars. The 3 wheel synchro-drive system is equipped with optical incremental encoders, and no heading movement is possible. Communication with the drive system and the sonars hardware is performed through an ISA card inside the on board Pentium II PC. Radio Link is used for wireless Ethernet connection.

Figure 1: *Blacky, the robot*

The platform also has a horizontal rotating laser called *LaserNav*, communicated with the on-board PC through a serial port. It can detect, measure the angle and identify up to 32 different bar coded passive targets. The laser target landmarks are located at the height of the sensor (120 cm), so they are frequently obstructed by people. Auto-amplifying loudspeakers are also used on Blacky's top for voice synthesis. Additionally, the cover of the LaserNav is provided with funny eyes and mouth, as well as a pirate hat that completes its personality.

2.2 The navigation system

Robot navigation architecture is composed of two levels. A low-level controller serves as interface to the robot hardware, providing collision avoidance. Several reactive behaviors implementing simple patterns of movements are used by a supervisory system to carry out high level tasks.

Additionally, a virtual corridor map, combined with simulated perception and an extended Kalman filter (EKF) for localisation, are used to overcome the lack of perception. The main components of the navigation system are:

- The no collision low level controller, which is in charge of collision avoidance, being able to reduce the speed and to stop the robot, if necessary. It also includes dead reckoning and control loop timing.
- The reactive behaviors, that use the low level controller to command movements to the robot, depending on both the robot position and the incoming readings from proximity sensors. The implemented behaviors are *follow corridor*, *go to point*, *escape from minimum*, *border by the right or the left* and *intelligent escape*, although only two of them are actually being applied at exhibitions: *follow corridor* which moves the robot along a corridor towards a defined direction, and *intelligent escape*, used for making oral presentations, that tries to emulate intelligent movement achieved by humans, while speaking to crowds.
- The virtual corridors model that is used to cope with the big forbidden areas that the robot is not able to percieve. This model constraints robot's movement within these corridors. Since the robot cannot sense the virtual corridor, sensor readings must be simulated so that reactive algorithms can be applied. A good robot position estimation is needed to achieve a good performance.
- The task planner, which defines plans in a similar way to directions given to humans: *walk along this corridor in this direction*, *take this other corridor in this direction*, etc. The execution of the plan is accomplished by the supervisory control.
- The localization module, based on encoders and *LaserNav* measures, due to the poor behaviour of sonars for this subtask. A continuous localization algorithm is provided, rather than an absolute one. The applied EKF requires the manual introduction of an initial rough position and orientation.

2.3 The control architecture

The adopted control architecture corresponds to a hybrid scheme reactive/hierarchical:

- The low level controller for collision avoidance is typical of the Subsumption architecture.
- Reactive behaviors running in parallel with continuous localization is typical of layered reactive control.
- The supervisory control which monitors reactive behaviors and takes into account the information of the virtual map, is the hierarchical component of the architecture.
- The use of a virtual map is an addition to this system architecture.
- Voice synthesis is an actuator at both reactive (to ask for free way) and high levels (making oral presentations and guided tours).

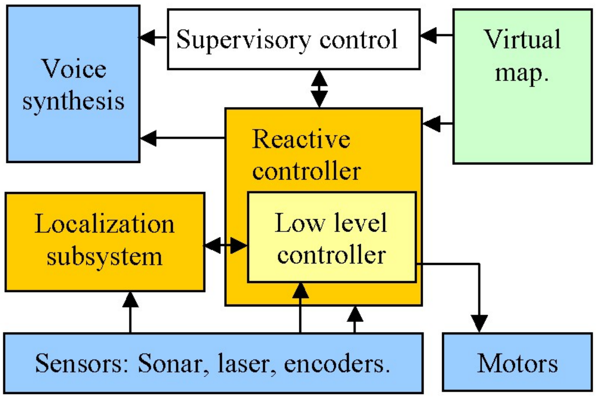

Figure 2: *Architecture scheme*

Asynchronous multithread control is used for modules coordination and synchronization.

2.4 Voice interaction

Robot moves autonomously in a populated public environment while it interacts with the people via voice, being human-robot interaction a key point. Acceptance of the robot by present people is crucial, and is achieved by means of using speech synthesis and a nice face. People needs a point to look at, and a way to communicate with the robot (microphone, keyboard, touch screen, etc).

An on board laptop computer was used for voice synthesis, using IBM ViaVoice. Pre-defined codes for greetings, welcome, self presentations, goodbye, etc., produce a random selection of predefined sentences. Furthermore, voice synthesis is proved to be an effective aid for navigation. *Ask for free way* codes are used by the supervisory control with different levels such as *polite*, *insistent* or *insulting*. Oral presentations based on text files can be synchronized with movement. Presentations are automatically interrupted to ask for free way, resulting in a high degree of intelligence and autonomy appearance.

2.5 System integration

The system has a distributed architecture in which efficient communication between computers is extremely important. Different operating systems coexist, while different software modules interact. The remote user can use his/her own PC without any modification, typically with Windows operating system, Internet connection and a web browser. This user connects him/herself to the system through its Internet Server Provider (ISP). The local part of the system is a set of PCs acting as web server, and an Ethernet radio link. The mobile part is composed of a mobile platform with two PCs (one with Linux operating system and other with Windows). The control PC is connected to the local area network via an antenna.

3 Exhibition sites

The kind of sites in which the robot operates are indoor environments of relatively high complexity: rooms, corridors, stands, objects, panels with text, people moving around, etc. The main features that the robot design must cope with are:

- Walls, not always vertical and with right-ended corners. Inclined walls must be foreseen in the collision avoidance system.
- Different type of floors that affect the locomotion system, such as of wood, floor tile, or carpet, as well as with slipperiness, slight inclination and roughness. Corridor limits are defined by thin plastic or metallic pieces (difficult to be detected by robot sensors), or just a change in floor color.
- Ceiling, that can be used to place landmarks, may vary its height from what is normal in a home or office to what is more usually found in castles or pavilions, where signs, security cameras and light focus are hanged from it.
- Stands, that are normally built of vertical poles that support an upper structure for signs. Front sides are usually open, while back ones use to be closed with walls (panels). Nevertheless this is not a rule to be strictly followed, since stands with no walls and no poles may be found, and stand limits can be defined following any strategy such as a different color of the floor carpet.
- It is also usual to have stairs in an indoor environment, a kind of obstacle that is not usually seen by 2D sensors and must be taken into account to avoid robot collisions.
- Illumination especially affects vision navigation systems. Artificial (structured) light is the most convenient one, since it guarantees homogeneous illumination, but it is also usual to have incoming sun light through windows.
- It is frequent to find in exhibitions objects of different sizes and heights placed around corridors and stands, that may not be always correctly mapped by the sensors.
- Crowded environments of people implies high sound level and influences the robot's movement often blocking its path and originating an important problem to navigation efficiency.

Figure 3: *Lack of structure in a fair*

Thus, trade fair and exhibition-like environments present several perception problems which produce that perception, map building or localization are not feasible with only proximity sensors information.

On one hand, attractiveness of the robot invite people to surround it, making it impossible to percieve around. Sonars beam divergence and the relatively low number of them, implies that very few people are necessary to block completely the sight of the robot. This problem could be partially solved with the use of othet kind of proximity sensors such as laser range finders.

On the other hand, stands of a trade fair rarely have physical walls, as appears in conventional fair maps, being usually built of aluminum structures with sometimes steps on the floor, a platform, floor colors, hanged posters, furniture, fences, etc, as commented above. None of these items are visible to the robot proximity sensors, becoming an important problem. Furthermor, manual definition of their existence is not practical at all, since stands are dynamically reconfigured.

An extra problem is that the huge dimensions of the typical environments locates navigational reference objects (e.g. the back wall of a stand) out of the sensors range. *Coastal planning* is not a suitable solution as all the *invisible objects* are close to the possible references. Automatic occupancy map building is not possible at all, and manual map building is not practical either, using only proximity sensors.

4 Experiments and discussion

Blacky the robot, has worked successfully in actual environments on 3 occasions, accumulating 2 weeks of intensive use:

- INDUMATICA: it is a trade fair organised annually by UPM students, in which companies pay to have a stand and present their job offers. Public are both students and professors. In March 2001, Blacky made its public presentation during 1 week.
- CYBERTECH: it is a robot contest organised by UPM, in which students present their prototypes and compete. In April 2001 Blacky entertained UPM students for 2 days, while they were attending the exhibition.
- Madrid for the Science: it is a trade fair organised annually at IFEMA, in which universities and research centres present their activities to the youngest. In May 2001, Blacky interact during 3 days, for first time with normal citizens, mainly children attending the fair with their parents.

In order to run the system, some adaptation to the exhibition sites were necessary:

Figure 4: *INDUMATICA fair*

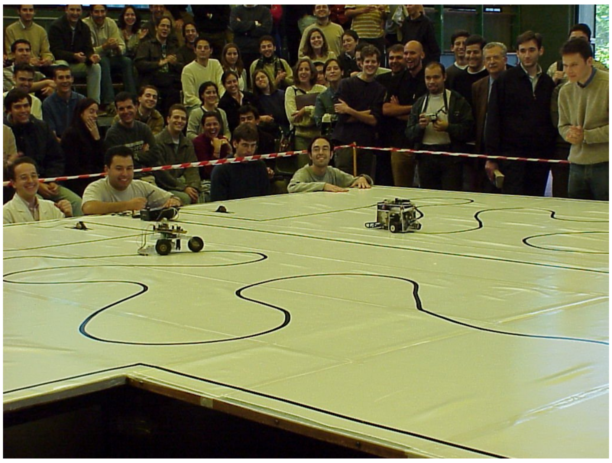

Figure 5: *CYBERTECH contest*

Figure 6: *Madrid for Science fair*

- A computing area with Internet access was used for the host computer. This area was located at UPM-DISAM stand, near to the place where the robot was navigating. A wireless link transmitted Ethernet signal to the robot.

- A small corner reachable by the robot was reserved for robot's batteries recharging.

- A set of artificial landmarks were strategically distributed along the site. Scientific exhibition organizers are usually open to modify the environment without any restriction.

Exhibition-like sites are appropriated places to test robot navigation algorithms. Autonomy is demonstrated by reaching the desired target, localizing itself in a complex low-structured indoor environment, where objects are moved continuously, while avoiding dynamic obstacles. Experiments demonstrated that the low level controller ability of avoiding collisions was extremely useful for crowded environments.

Simple reactive behaviors with voice synthesis assistance was a good solution for traveling along fair corridors. Humans reaction to robots requirements for free were clearing of the way, getting bored after robot AI based reasoning system insisted for some time. Adults reacted leaving space for the robot while children intentionally blocked it's path or sat on the floor in front of it.

The combination of virtual corridors maps with the perception simulation were a good solution to the some of the fair environment previously commented.

The localisation system worked acceptably, although it failed in several occasions, being the most critical part of this work. The use of EKF produces lost of geometric information that should be reconsidered. The size of the environment is also a problem for landmark placing and mapping.

System integration has also been addressed in order to achieve a robust platform, and to guarantee safety. Real work in events provides more useful conclusions than laboratory tests. Furthermore, technical problems with the robot brought important contributions. One person was in charge of supervising the robot while it was moving, acting as intermediate for dialogues, and startup and shutdown the robot when batteries needed to be recharged.

Only one time he had to press the emergency stop button on-board the robot due to a localisation error combined with a lack of perception. Multithread and distributed software implementation run robustly. Neither system hanged nor communication problems appeared.

Face orientation was more important than it seemed a priori, since it determines robot's intention of movement direction. Speech synthesis might be used combined with smooth movements of the robot indicating also its intention, although this may be completely ignored especially by the youngest visitors who see this experience as a game. The friendliness and personality of *Blacky* gained high attention.

The robot is also an excellent advertisement distributor. People always obey robot orders so they will be delighted to take any letter, paper or advertisement over it under robot's instructions.

5 Future working lines

5.1 Perception

Future research is being focused in increasing perception habilities through the use of new sensors with three dimension peception capabilities. Robot operation needs sensors for navigation, environment modeling, localization, security and remote operation:

- Navigation: proximity sensors such as ultrasounds and infrared, already supplied by the robot, are enough for obstacle avoidance. RangeFinder lasers may also be used for this reason, providing more accuracy.

- Environment modeling: proximity sensors are also needed for any kind of map (geometric or occupancy grid based).

- Localization: this is one of the most important issues in robot navigation. Many sensors are available for this purpose, and sensor fusion is a good option. Among them, we can mention: odometry, electronic compasses, rotatory lasers for detecting reflecting landmarks, color and b/w cameras for signs and marks in walls and ceilings, RangeFinder lasers for walls and reflecting beacons, proximity sensors, global positioning systems, and magnetic sensors for resetting position when passing through certain areas.

- Security: trade fair environments have high complexity. In case of failure of the navigation sensors, sensor redundancy must be supplied for emergency stop. Different sensors must be used: proximity sensors (infrared or RangeFinder) for detecting the presence of floor, stairs, very low objects, and objects hanging from the wall or ceiling. Guide following marks on the floor could also be used for detecting working area limits. Special sensors such as temperature, smoke and intruder detection sensors, could also be used in the case of reusing the robot for other applications such as vigilance.

- Remote operation: on-board cameras may be used for transmitting on-site images to the remote user. Different possibilities are available: color cameras, omnidirectional cameras, or fixed cameras on the walls.

5.2 Navigation

Research must focus on automatic artificial and natural landmarks identification and mapping. Landmarks like lines in the middle of the corridors are considered to be conceptually more correct, so the robot can always be able to *see* them. Possible errors are eliminated, as reference line placing is immediate and mapping can be done automatically.

Present technologies for navigation should be redesigned in the case of changing the navigation philosophy. The virtual corridors based approach could be a good solution, provided that it is complemented with additional information. Present navigation issues must take into account all information provided by both the sensors readings and the exhibitor organizers (maps).

In the first case, the robot can use its perception capabilities to autonomously map the environment (for the regions in which this is possible). In the second case, exhibition companies use to design a CAD map of the environment with the stands distribution. This one is usually a qualitative map that could also be obtained by visual inspection.

Nowadays our work focusses on using fuzzy logic and a probability/possibility Kalman filter to fuse and manage both kind of information, by using geometric and topological approaches.

It is also interesting to expand the interaction abilities of the robot as a main goal to the overall success of the system. Head orientation and facial expression capabilities are considered important for future developments.

5.3 Exploitation

From the point of view of CACSA, an organizer of Science exhibitions company, trade fairs mean a new field of application for mobile robots. This company considers this kind of robots of high impact factor from the market point of view. These robots are also a good way to introduce new communication technologies in daily life, allowing tele-presence.

Several additional uses of the robot, proposed by CACSA are vigilance, cleaning or lost children localization. Furthermore, they agree that benefits from the robot can be obtained through the different users of the system:

- The remote visitor can pay for a special ticket to visit remotely the exhibition. Visitors could be normal citizens attending a leisure fair, or business people attending a commercial one.

- The robot is an exhibit of itself, so the exhibition organizer could use the robot as an added service (as power source, light, phone line, Internet connection, etc) for its clients. The robot should be programmed in such a case to guide visitors to those companies that have paid this extra service.

Acknowledgments

This work is funded by Spanish Minister of Science and Technology (URBANO project: DPI2001-3652-C02-01) and the European Union 5th Framework R+D Programme (WebFAIR: IST-2000-29456).

References

[1] J. Borenstein, B. Everett and L. Feng. "Navigating Mobile Robots: Systems and Techniques." A. K. Peters, Ltd, Wellesley, MA, 1996.

[2] R.A. Brooks. "A Robust Layered Control System for a Mobile Robot". IEEE J. Robotics and Automation. Vol. 2 (1986) 14-23.

[3] J.L. Crowley. "Mathematical Foundations of Navigation and Perception for an Autonomous Mobile Robot". In Proceedings of the RUR, Amsterdam, Dec. 1995.

[4] D. Gachet, M.A. Salichs, J.R. Pimentel, L. Moreno and A. De la Escalera. "A software Architecture for Behavioral Control Strategies of Autonomous Systems". In Proceedings of the IECON 1992.

[5] D. Rodriguez-Losada, F. Matia, R. Galan and A. Jimenez. "Blacky, an Interactive Mobile Robot at a Trade Fair". In Proceedings of the 2002 IEEE Int. Conf. on Robotics and Automation, ICRA'02, Washington, USA, May 2002.

[6] N. Roy, W. Burgard, D. Fox and S. Thrun. "Coastal Navigation – Mobile Robot Navigation with Uncertainty in Dynamic Environments". In Proc. IEEE Conf. Robotics and Automation (ICRA), May, 1999.

[7] S. Thrun, M. Bennewitz, W. Burgard, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. "MINERVA: A second generation mobile tour-guide robot". In Proc. of the IEEE International Conference on Robotics and Automation (ICRA), May, 1999.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Design and System Integration for the Expo.02 Robot

N. Tomatis[†‡], G. Terrien[†], R. Piguet[†], D. Burnier[‡], S. Bouabdallah[‡], R. Siegwart[‡]

[†]BlueBotics SA
PSE-C
CH-1015 Lausanne
n.tomatis@ieee.org

[‡]Autonomous Systems Lab, EPFL
Swiss Federal Institute of Technology Lausanne
CH-1015 Lausanne
r.siegwart@ieee.org

Abstract

This paper presents the genesis of the Expo.02 robot. RoboX the tour guide robot has been built from the scratch for this project based on the experience of the Autonomous Systems Lab. The production of 11 of those machines has been guaranteed by a spin-off of the lab: BlueBotics SA. The goal was to maximize the autonomy and interactivity of the mobile platform while ensuring high robustness, reliability and performance. The result is an interactive moving machine which can operate in human environments and interact by seeing humans, talking to and looking at them, showing icons and asking them to answer its questions. Here, the complete design of mechanics, electronics and software is presented first, followed by the statistics about the first two month of operation.

1. Introduction

Within the Expo.02, the Swiss National Exhibition, the *Robotics* exhibition takes place in Neuchâtel, where the main thematic is *nature and artifice*. *Robotics* is intended to show the increasing closeness between man and machine. For this the visitors interact with ten autonomous, freely navigating tour guide robots, which present the exhibit going from industrial robotics to cyborgs on a surface of about 320 m^2.

The task of a tour guide robot is to be able to move around autonomously in the environment, to acquire the attention of the visitors and to interact with them efficiently in order to fulfill its main goal: give the visitors the pre-defined tour. The environment is known and accessible, but a general approach requiring no environmental changes is better suited for a commercial purpose. For the same reason a fully-autonomous and self-contained robot is preferable. Furthermore such a machine is required to have a long live cycle and a high mean time between failure (MTBF), which minimizes the need of human supervision and therefore the maintenance costs.

2. Related Work

The tour-guide robot task can be subdivided in two separate issues, which are navigation and interaction.

Navigation: A limited number of researchers have demonstrated autonomous navigation in exhibitions or museums [5], [12], [14], [8] and [15]. Furthermore, most of these systems have still some limitations in their navigation approaches. For instance *Rhino* [5] and *Minerva* [14] have shown their strengths in museums for one week, 19 kilometers and two weeks, 44 kilometers respectively. However, their navigation has two major drawbacks: it relies on off-board resources, and due to the use of raw range data for localization and mapping it is sensible to environmental dynamics. *Sage* [12], *Chips*, *Sweetlips*, *Joe* and *Adam* [15], use a completely different approach for permanent installations in museums: the environment is changed by adding artificial landmarks to localize the robot. This approach performed well, as shown with a total of more than half a year of operation and 323 kilometers for *Sage* [12] and a total of more than 3 years and 600 kilometers for *Chips*, *Sweetlips*, *Joe* and *Adam* [15]. However their movements, but for *Adam*, are limited to a predefined set of unidirectional safe routes in order to simplify both localization and path-planning. Another permanent installation which is operating since March 2000, is presented in [8]. Three self-contained mobile robots navigate in a restricted and very well structured area. Localization uses segment features and a heuristic scheme for matching and pose estimation. Another exhibition where *Pygmalion*, a fully autonomous self-contained robot was accessible on the web during one week [1] has shown its positive characteristics but, due to the unimodal characteristic of the used Extended Kalman Filter, the robot can still lose track if unmodeled events take place.

Interaction: Human-centered and social interactive robotics is a comparatively young field in mobile robotic research. However, several experiences where untrained people and robots meet are available. The analysis of the first public space experience with *Rhino* [5] underlines the importance to improve human-robot interfaces in order to ease the acceptance of robots by the visitors. In [14] *Minerva* attracted visitors and gave tours in a museum. It was equipped with a face and used an emotional state machine with four states to improve interaction. The *Mobot Museum Robot Series* [12] and [15] focused on the interaction. Robustness and reliability were identified as an important point for the credibility of a public robot. The permanent installation at the *Deutsches Museum für Kommunikation* in Berlin [8], uses three robots which have the task to welcome visitors, offer them exhibition-related information and to entertain them.

The system presented here is designed to offer enhanced interactivity with complete autonomous navigation in a completely self-contained robot and without requiring changes of the environment. Furthermore it is intended to work permanently with minimal supervision.

3. Design and System Integration

The typical environment of an exhibit, which is highly dynamic, and the visitor experience expected with such a robot impose various constraints on the design and control. This leads to a specification of the mobile platform that can be summarized as follows:

- Highly reliable and fully autonomous navigation in unmodified human-environments crowded with hundreds of humans.
- Bidirectional multi-modal interaction based on speech (English, German, French and Italian), facial expressions and face tracking, icons (LED matrix), input buttons and robot motion.
- Safety for humans, objects and the robot itself at all time.
- Minimal human intervention and simple supervision.

The esthetic of the robot has been designed in collaboration with artists, industrial designers and scenographers. The result of the design of both hardware and software is RoboX: a mobile robot platform ready for the real world (figure 1).

Given the above mentioned specifications the mechanical, electronic and software design are now presented.

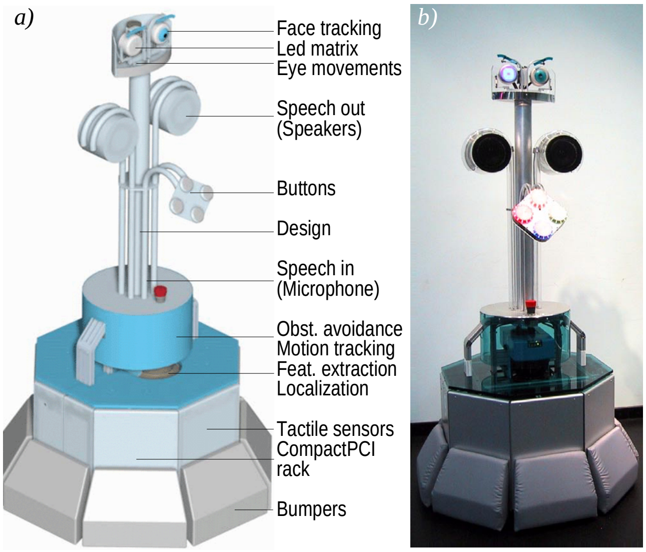

Figure 1: a) Functionality of the tour guide robot RoboX. b) An image of RoboX 9.

3.1 Mechanical Design

The navigation base (lower part of the robot) consists mainly in the batteries, the CompactPCI rack with two control computers, the laser range sensors (two SICKs LMS-200), the bumpers and the differential drive actuators with harmonic drives. The base (figure 2) has an octagonal shape with two actuated wheels on a central axis and two castor wheels. In order to guarantee good ground contact of the drive wheels, one of the castor wheels is mounted on a spring suspension. This gives an excellent manoeuvrability and stability to the 1.65 m height robot.

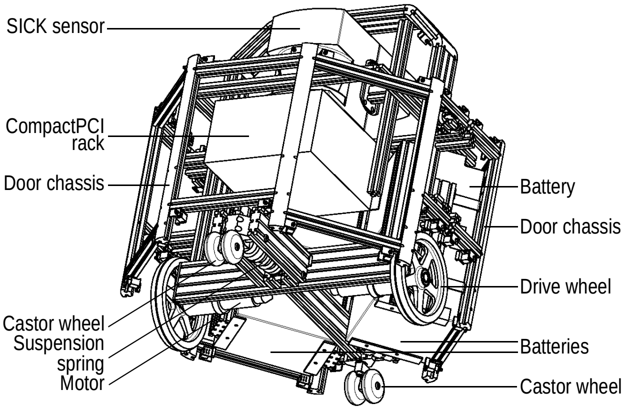

Figure 2: Mechanical design of RoboX base.

The upper part of the robot incorporates the interaction modules. The face includes two eyes with two independently actuated pan-tilt units and two mechanically coupled eyebrows. The left eye is equipped with a color camera, which is used for face tracking. The right eye integrates a LED matrix for displaying symbols and icons. The eyebrows further underline eyes expressions by means of a rotational movement. Behind the face, a gray scale camera pointing to the ceiling is mounted for localization purpose.

The main input device for establishing a bidirectional communication with the humans are four buttons which allow the visitors to reply to questions the robot asked. The robot can further be equipped with a directional microphone matrix for speech recognition even though this remains very challenging in the very noisy environment of an exhibition.

3.2 Electronic Design

The control system (figure 3) has been designed very carefully by keeping in mind that the safety of the humans and the robot has to be guaranteed at all time. It is composed of a CompactPCI rack containing an Intel Pentium III card and a Motorola PowerPC 750 card. The latter is connected by the PCI backplane to an analogue/digital I/O card, a Bt848-based frame grabber, an encoder IP module and a high bandwidth RS-422 IP module. Furthermore a Microchip PIC processor is used as redundant security system for the PowerPC card (figure 3).

The navigation software runs on the hard real-time operating system XO/2 [4] installed on the PowerPC. This processor has direct access to the camera looking at the ceiling, the two SICK sensors, the tactile plates and the main drive motors. It communicates with the interaction PC through Ethernet via an on-board hub.

The interaction software is running under Windows 2000 on an industrial PC. This allows using commercial off-the-shelf (COTS) software for speech synthesis and recognition, and makes scenario development easier. The PC has direct access to the eye camera, the eyes and eyebrows controller, the input buttons, the two loudspeakers and the microphone.

The robot (both CPUs) is connected by a radio Ethernet to an external computer only for supervision, in order to track its status at any time on a graphical interface.

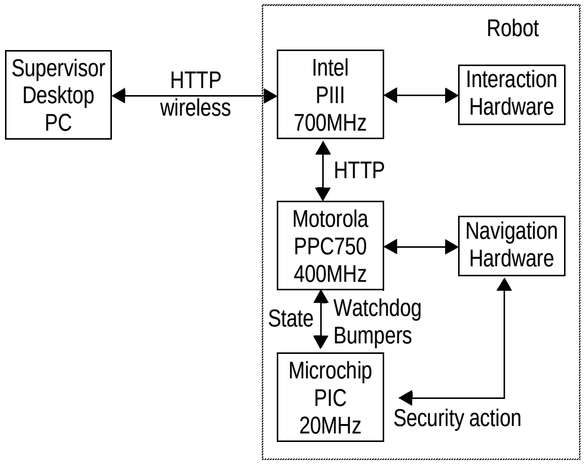

Figure 3: Simple scheme of the electrical design

3.3 Software Design

As explained in the section above, the robot is composed of both an Intel Pentium and a Motorola PowerPC systems. The software has been firstly designed without taking into account this fact based on the functionality which was to be developed. However, as soon as the implementation started, the objects have been assigned to one of the two distributed systems. For hardware related objects (mainly sensor drivers) the choice was obvious. For the others, their relevance to safety has been evaluated: due to the hard real-time characteristics of XO/2, all the time-critical objects in relation with the security have been implemented on the PowerPC. Objects requiring COTS components have been implemented on the Windows machine because of their wider availability (f.e. MBrola for speech out, small FireWire camera in the eye for the face tracking, ...).

The resulting object distribution is represented in figure 4. In the following part of this section each component of figure 4 is briefly presented starting with the interaction system followed by the navigation. A complete description of the interaction of RoboX can be found in [10]. Its navigation system is presented in [3].

Interaction

Scenario Controller: It is the central object of the interaction subsystem, which accesses all the other objects. A *scenario* is a sequence of tasks from all modalities (speech, face expression, motion, LED matrix, etc.). A sophisticated tour-guide scenario consists of several small scenarios which are played by the scenario controller.

People detection: It permits to detect movements of objects around the robot by means of the laser scanners. By assuming a static environment, these moving objects are either humans or other robots. The moving objects are then tracked by means of a *Kalman Filter*.

Speech Out: By using software permitting either text-to-phonemes-to-speech or directly text-to-speech, this object permits the robot to talk in four languages (English, German, French and Italian). Furthermore files of format wav and mp3 can be played.

Buttons Controller: This controls the main input device for the interaction between the robot and the humans. Four inductive buttons with different colors are used in combination to questions from the speech out to close the interaction loop with the robot.

LED Matrix: The LED matrix is in the right eye. Its controller permit to show icons and animations.

Eyes Controller: The eyes can be moved independently. The controller has a set of predefined expressions, which can be directly played.

Face Tracking: The color camera in the left eye is used to track skin colored regions. The approach is based on [9]. This permits, in combination with the eyes controller, to track a face on the image and with the movement of the eyes.

Navigation

Odometry Driver: Calculates the position and uncertainty of the robot based on the wheels rotations.

Speed Controller: Regulates the speed defined by the obstacle avoidance with a PID controller accessing the encoders and updates the odometry.

Localization: Uses a new approach [2] based on an *Extended Kalman Filter* [6] to correct the odometry with exteroceptive sensors (laser scanners, CCD camera).

Obstacle Avoidance: Calculates a collision free path by initializing the path with a *NF1* function [11] and using the *Elastic Band* approach [13] to dynamically adapt it. Furthermore it guarantees that the robot can stop before collision at any time with the *Dynamic Window* approach [7].

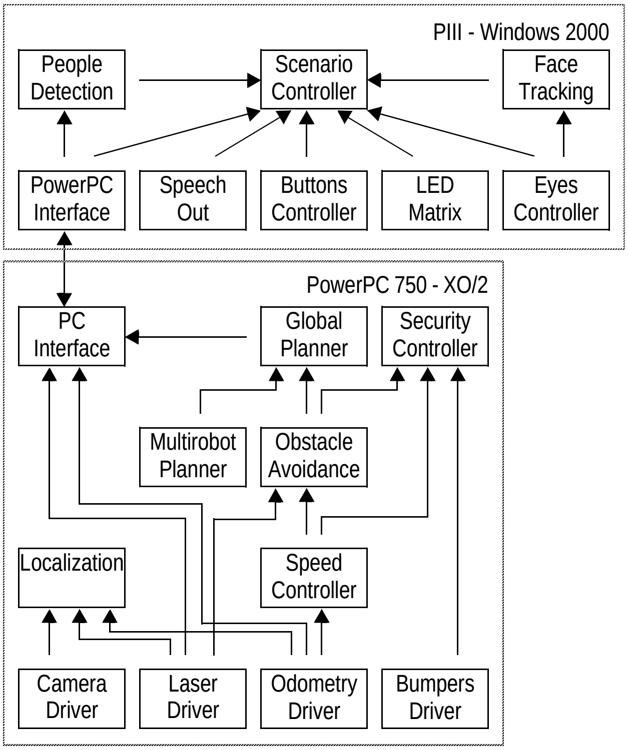

Figure 4: Object distribution of the software on the distributed embedded system.

Multirobot Planner: Synchronizes the movement of the robots to avoid to have many robots at the same place.

Global Planner: Plans the navigation of the robot on the map level, by defining via points which permit to reach the goal point within the graph representing the map.

Security Controller: Guarantees that the robot cannot become dangerous even in case of failure by supervising the safety-critical software and sensors. Due to the importance of this issue for a robot sharing the environment with humans, the next section presents the security system in details.

4. Security

In this section the involvement of the security issue in the design of the whole system is pointed out in more details.

All the software which relates to the movement of the robot is defined as safety critical. In order to guarantee the security of both the users and the robot itself, safety is on three levels: the operating system, the software implementation and the redundancy of the hardware.

4.1 Operating System

All navigation software is implemented on the PowerPC which is operated by XO/2, a deadline driven hard real-time operating system [4]. Due to its characteristics XO/2 guarantees:

- *Safety*: nothing bad happens.
- *Progress*: the right thing do (eventually) happen.
- *Security*: things happens under proper supervision.

Static safety is guaranteed by the strong-typing characteristic of Oberon-2, the language used under XO/2. Many errors are already found at compile-time instead of run-time. Furthermore, index-checks, dynamic type systems and especially the real-time compatible garbage collector guarantee dynamic safety by forbidding almost any memory-management related errors.

The deadline driven scheduler is in charge of progress: it guarantees that each task is executed within the predefined deadline. Of course this is possible only if the constellation of the tasks running on the PowerPC requires less than 100% of CPU. For this, the duration of each tasks has to be known. Admission tests are performed at each installation of a new real-time task to guarantee their feasibility. As soon as the progress of all real-time tasks is guaranteed, the CPU is scheduled between the non-real-time tasks depending on their priorities.

Each error causes a *system trap* which is under complete control of the operating system. The system knows exactly where the error toke place (which line in the code), who called this part of the code up to the task currently running. This is very helpful for debugging, but it is even more important for security because for each task an *exception handler* can be defined. The actions which have to take place in such a case can therefore be properly defined.

4.2 Software Security

Tasks whose failure could cause injuries to objects or people required a special attention during design. Software watchdogs are therefore implemented for the speed controller, the obstacle avoidance and the bumpers driver (figure 4). Failure of one of these tasks is detected by the security controller which then either restarts the failed task or stops the robot and sends an e-mail to the maintenance. This permit to centralize the control of the security and to ask to a single object if a defect is disturbing the system. Furthermore the security controller generates a watchdog signal on a digital output permitting to know if both the operating system and the security controller are still running.

4.3 Hardware Redundancy

The above mentioned software permits to have a consistent control system running on the PowerPC. However, this isn’t enough to guarantee the security of the robot and its neighborhood. Even in case of failure of the electronics or problems on the operating system of the PowerPC the robot must remain un-dangerous. For this the robot has a third processor: a Microchip PIC (figure 3). The software running on it checks the watchdog generated by the security controller, awaits acknowledgements from the security for each bumper contact and controls that the pre-defined maximal speed is never exceeded. If one of these conditions is not respected the redundant security software running on the PIC safely stops the whole system and set it in emergency mode (acoustic alarm).

5. Experiments

At the time of writing 64 days of operation, from May 14 to July 17, are available for statistics. Each day from six to ten freely navigating tour-guide robots have given tours for 10.5 hours, from 9:30am to 8:00pm, on the surface of the exhibit which is approximately 320 m^2.

5.1 Definitions

Failure: A failure is any kind of problem which required a human intervention. The only exception is for the emergency button, which can be pressed and released also by visitors, and, due to logging difficulties, for situations where the robot remains blocked somewhere because it is to near to an object. In the latter case the staff can displace the robot by a switch which de-connects the motors from the amplifiers and allows to move the 115 kilograms robot easily.

Uncritical: Uncritical failures are those which does not stop the task of the robot. For example, a failure consisting in a robot which stops sending an image to the supervisor is not critical for the tour the robot is giving to the visitors.

Critical: Critical failures stop the robot until the human intervention is performed. An example is the failure of the scenario controller or of the obstacle avoidance.

Reboot: Critical failures requiring a reboot of either the Pentium or the PowerPC are treated separately only because they require more time before operation.

5.2 Results

After 64 days of operation the robots served more than 283’000 visitors for a total of 5’290 hours of operation. In order to do this job, they travelled more than 1’250 kilometers for a total moving time of more than 3’730 hours meaning that the mean displacement speed is 0.094 meters per second. As it can be seen in table 1, the uncritical failures represent only a small portion of the total amount of failures (10.9%). Furthermore they do not disturb the operation of the robot. They are therefore not treated in the following analysis which will focus on the critical and reboot failures of the whole robot first and then of the PowrePC.

Run time	5’293 h
Movement time	3’736 h
Travelled distance	1’259 km
Average speed	0.09 m/s
Failures (total / critical / uncritical)	2’097 / 1’869 / 228
Critical failures (PC / PPC / HW)	1’641 / 163 / 65
Visitors	283’319

Table 1: *Two months of operation. After more than 5’000 hours of operation the RoboXes have travelled more than 1’250 kilometers and served more than 280’000 visitors.*

As it can be seen in figure 5, the beginning of the exposition in the middle of May showed that some work was still to be done. The software running on the PC was very unstable due especially to some errors in treating the list of the tasks running into the scenario controller.

The mean time between failure (MTBF) of the whole robot (PC, PowerPC and hardware) during the first week was less than one hour (figure 6). This has been improved and is now around seven hours, which means that during one day with 10 robots, the staff has to perform a total number of interventions which is between 10 and 20. The type of interventions goes from the simple double-click to restart an application (typical intervention on the PC) to the change of an motor amplifier (very rare, it happened four times until now, two of them due to a motor defect). After the first three weeks, the MTBF already doubled. Some errors were found after some weeks of operation, some other come for the first time after one month. The chance of having thousands hours of operation permits to improve the software and hardware to a level which is simply un-achievable with smaller projects.

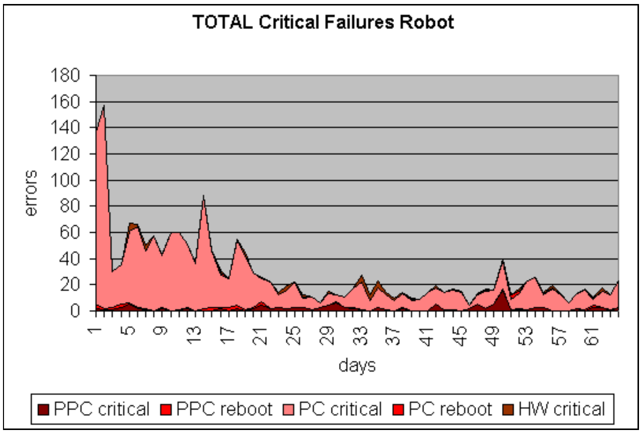

Figure 5: Due to many delays in the development, the software was still in the test phase at the beginning of the exposition. The first four weeks represent a huge improvement in the stability of the software, especially on the PC side.

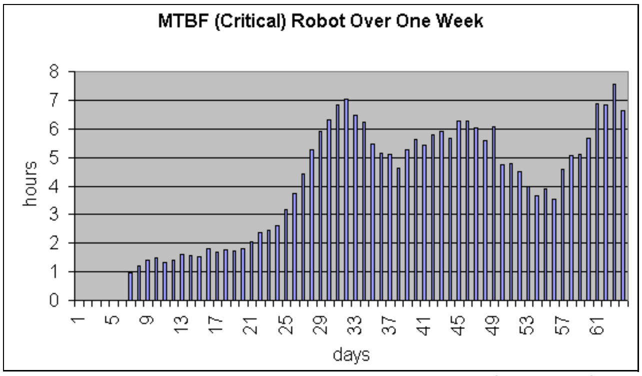

Figure 6: The mean time between critical failure of any kind (PC, PowerPC, hardware). The improvement has been constant exponential during the first four weeks, where the most important errors have been found. The current errors, which are rare, are more difficult to find.

Another interesting chart is in figure 7, where all the critical failures coming from the navigation software (PowerPC) are shown. During the first three weeks, errors in the safety-critical tasks were treated by the security controller, but could sometimes require a reboot in order to restart the trapped task. This has been partly corrected allowing for much faster intervention in case of failure. Critical failures in figure 7 contains also error which have nothing to do with the implementation. For example, failures of the localization system are sometimes requiring human intervention. The peak of 17 critical failures on day 50 in figure 7 is due to a new person in the staff, which handled the robots without using the switch permitting to de-connect the motors from the amplifiers. This caused huge errors in the odometry and therefore some failures of the localization system. This is also the cause of the loss of MTBF of the robot between day 43 and day 57 in the chart of figure 6.

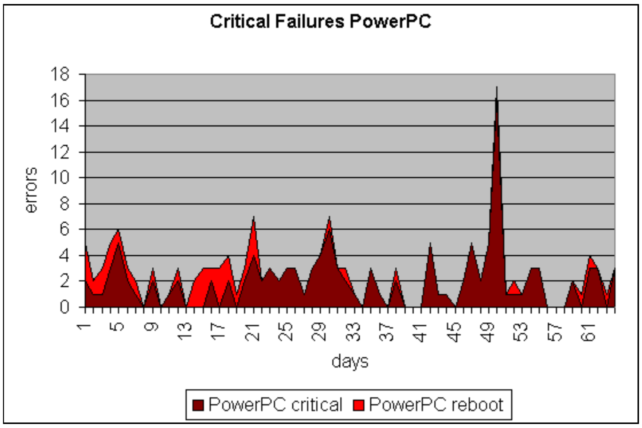

Figure 7: The critical failures of the PowerPC (navigation system). Some of the critical errors require the reboot of the PowerPC. The peak of day 50 is due to bad manipulations of the robot by an untrained member of the staff.

The MTBF for the PowerPC (figure 8) was already at the beginning of the exposition at another level with respect to the rest of the software with values of 20 hours after one week and between 50 and 60 in the last two weeks. Without the day 50 problem the MTBF would be over 50 hours since around day 40. This better result is due to the characteristics of the XO/2 operating system which has been developed for embedded systems focusing on the robustness and safety, and due to the navigation software which is evolving since more than four years at the Autonomous Systems Lab in contrast to the interaction software which has been developed only for this application starting in late year 2000.

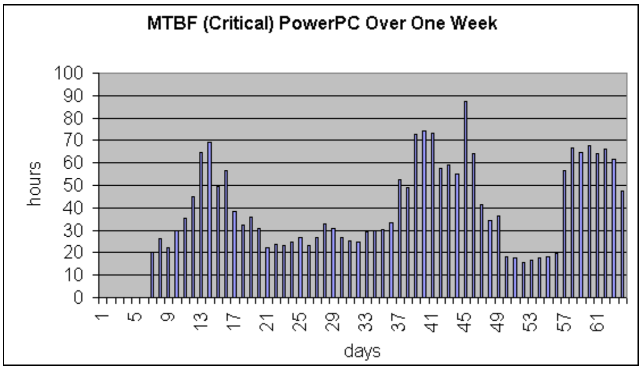

Figure 8: The MTBF during the first week was of about 20 hours. By neglecting the problem encountered during day 50, the MTBF would be over 50 hours starting from day 40.

Hardware failures (figure 9) are due to some uncritical design errors at the beginning (design of doors), to some motor-amplifier problems and to the temperature which was up to 35° in the exhibit between day 33 and day 40 (the SICKs do not like this!). This also showed a lack of the security approach, which did not take into account a possible failure of the laser scanners. When this occurred the obstacle avoidance continued to receive the last available scan from the driver causing a collision to the next object.

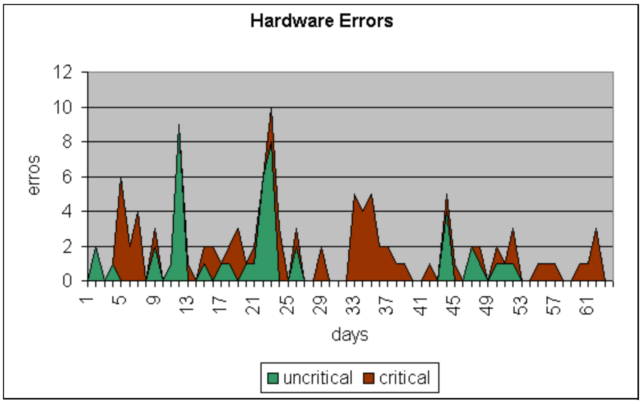

Figure 9: Hardware problems also cause critical failures. Four motor amplifiers have had some problems. The block of errors between day 33 and day 40 is due to a very-good-weather week with temperature up to 35° in the exhibit.

6. Conclusions and Outlook

This project represents a milestone in the field of mobile robotics: for the first time tour-guide robots are produced (11 robots) and used for long time (five months) as real products instead of prototypes as in former projects. The paper presents their characteristics first, then goes into details about the mechanical, electrical and software design. The security issue is faced seriously for ensuring security of the humans and the robot itself all the time. In the experiments section the results of the first 64 days of operation of the *Robotics* exposition are presented and analyzed focusing on the amount and type of failures which occurred to the robots.

References

[1] Arras, K. O., N. Tomatis, B. Jensen, and R. Siegwart (2001). "Multisensor On-the-Fly Localization: Precision and Reliability for Applications." Robotics and Autonomous Systems 34(2-3): 131-143.

[2] Arras, K. O., J. A. Castellanos, and R. Siegwart (2002). Feature-Based Multi-Hypothesis Localization and Tracking for Mobile Robots Using Geometric Constraints. IEEE International Conference on Robotics and Automation, Washington DC, USA.

[3] Arras, K. O., R. Philippsen, M. De Battista, M. Schilt, and R. Siegwart (2002). A Navigation Framework for Multiple Mobile Robots and its Application at the Expo.02 Exhibition. Workshop on Robots in Exhibitions, IEEE/RSJ International Conference on Intelligent Robots and Systems, Lausanne, Switzerland.

[4] Brega, R., N. Tomatis, K. Arras, and R. Siegwart (2000). The Need for Autonomy and Real-Time in Mobile Robotics: A Case Study of XO/2 and Pygmalion. IEEE/RSJ International Conference on Intelligent Robots and Systems, Takamatsu, Japan.

[5] Burgard, W., A. B. Cremers, et al. (1999). "Experiences with a Interactive Museum Tour-Guide Robot." Artificial Intelligence 00(1999): 1-53.

[6] Crowley, J. L. (1989). World Modeling and Position Estimation for a Mobile Robot Using Ultrasonic Ranging. IEEE International Conference on Robotics and Automation, Scottsdale, AZ.

[7] Fox, D., W. Burgard, et al. (1997). "The Dynamic Window Approach to Collision Avoidance." IEEE Robotics & Automation Magazine: 23-33.

[8] Graf, B., R. D. Schraft, et al. (2000). A Mobile Robot Platform for Assistance and Entertainment. International Symposium on Robotics, Montreal, Canada.

[9] Hilti, A., I. Nourbakhsh, B. Jensen, and R. Siegwart (2001). Narrative-level Visual Interpretation of Human Motion for Human-robot Interaction. IEEE/RSJ International Conference on Intelligent Robots and Systems, Maui, Hawaii.

[10] Jensen, B., G. Froidevaux, X. Greppin, A. Lorotte, L. Mayor, M. Meisser, G. Ramel and R. Siegwart (2002). Visitor Flow Management using Human-Robot Interaction at Expo.02. Workshop on Robots in Exhibitions, IEEE/RSJ International Conference on Intelligent Robots and Systems, Lausanne, Switzerland.

[11] Latombe, J.-C. (1991). Robot motion planning. Dordrecht, Netherlands, Kluwer Academic Publishers.

[12] Nourbakhsh, I., J. Bodenage, et al. (1999). "An Effective Mobile Robot Educator with a Full-Time Job." Artificial Intelligence 114(1-2): 95-124.

[13] Quinlan, S. and O. Khatib (1993). Elastic bands: connecting path planning and control. IEEE International Conference on Robotics and Automation.

[14] Thrun, S., M. Beetz, et al. (2000). "Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva." International Journal of Robotics Research 19(11): 972-99.

[15] Willeke, T., C. Kunz, et al. (2001). The History of the Mobot Museum Robot Series: An Evolutionary Study. Florida Artificial Intelligence Research Society (FLAIRS), Florida.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

A Navigation Framework for Multiple Mobile Robots and its Application at the Expo.02 Exhibition

Kai O. Arras, Roland Philippsen, Marc de Battista, Martin Schilt, Roland Siegwart

Autonomous Systems Lab
Swiss Federal Institute of Technology Lausanne (EPFL)
CH–1015 Lausanne, Switzerland

Abstract

This paper presents a navigation framework which enables multiple mobile robots to attain individual goals, coordinate their actions and work safely and reliably in a highly dynamic environment. We give an overview of the framework architecture, its layering and the subsystems reactive obstacle avoidance, local path planning, global path planning, multi-robot planning and localization. The latter receives particular attention as the localization problem is a key issue for navigation in unmodified and difficult environments. The framework permits a lightweight implementation on a fully autonomous robot. This is the result of a design effort striving for compact representations and computational efficiency.

The experimental testbed is the 'Robotics' pavilion at the Swiss National Exhibition Expo.02 where ten fully autonomous robots are interacting with more than half a million visitors during a five-month period.

1. Introduction

Navigation responds to three questions: 'where am I?', 'where am I going?' and 'how do I get there?'. A navigation framework has the task to offer the simplest possible interface to these questions, hiding their complexity to the user or the application layer.

At the inside, navigation deals with various constraints on different time scales and levels of abstraction. A common approach to structure the problem is a three-layered architecture which consists in a planning layer, an execution layer and a reactive behavior layer [6, 23] (figure 1).

- The *planning layer* decides how to achieve high-level goals using a model of the environment. Typically, the model encodes the environment topology in form of a graph or an occupancy grid. Planning takes places under constraints of task-specific cost functions and limited resources.
- The *execution layer* subdivides a plan into executable subplans, activates and deactivates behaviors and supervises their completion.
- The *reactive behavior layer* interfaces the robot's sensors and actuators. In mobile applications, it acts as a position controller under constraints of a dynamic environment, the robot shape, vehicle kinematics and dynamics.

Controls, abstracted sensor data and status information flow vertically between layers. Typically, controls such as plans and subplans are passed to lower levels and information such as status codes and termination flags are passed to higher levels. In case of failures in a layer (e.g. path blocked), requests for revised controls are sent to higher layers (e.g. replanned path). Time scale and abstraction increase from bottom to top, model fidelity and real-time concerns increase from top to bottom.

We adopt this three-layered architecture here as it accommodates deliberative and reactive behaviors, allows for constraint distribution over the layers and embodies an intuitive way of increasing abstraction from bottom to top. It is further suitable for many applications including manipulators and multiple robots [6, 23, 24].

The application we envisage are ten fully autonomous

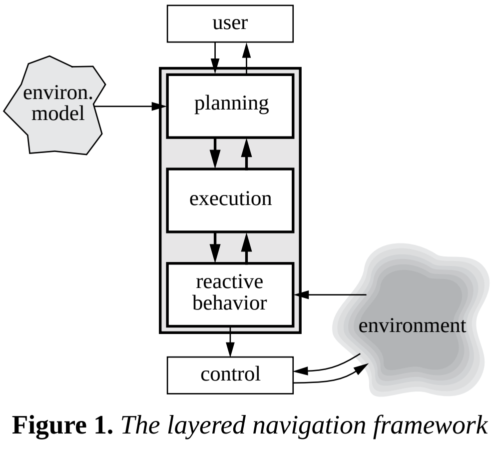

Figure 1. *The layered navigation framework*

mobile robots deployed in a mass exhibition with up to 500 visitors per hour. Their tasks include tour giving, entertainment and picture taking of visitors. They share the same space and the same goals for the tours (e.g. showcases) and operate in an unmodified environment. The paper presents how the architecture was adapted to suit our needs, and the choices we made for its components: obstacle avoidance, local and global path planning, multi robot coordination and localization.

2. The Navigation Framework

There are applications for mobile robots which require high degrees of autonomy. A mass exhibition with ten mobile robots is such a case. Today we (still) face limited computational resources in embedded systems for real-time. While radio-linked off-board hardware might be an alternative for a single robot, it would result in prohibitive bandwidth costs for multiple robots, if, for instance, raw sensor data for localization were transmitted. System complexity and reliability issue further confirm our intent for fully autonomous robots and decentralized concepts. Therefore, we look out for compact representations and computational efficiency.

2.1 Environment Model

Our approach to environment modeling is feature-based using geometric primitives such as lines, segments and points (sometimes called landmarks). The environment topology is encoded in a weighted directed graph with nodes and edges between the nodes. Neither for global path planning nor for localization we use a free space model like occupancy grids. The advantage of this choice is compactness: in indoor environments, a map of this type (features plus graph) requires typically around 30 bytes per m^2. Further, scaling to 3d is polynomial, whereas grid maps scale exponentially.

The graph has two types of nodes: *station nodes* and *via nodes*. Station nodes correspond to application-specific (x, y, θ)-locations in space with a meaning. Examples from Expo.02 include: showcase with industrial robot, tour welcome point or location to trigger picture caption. Via nodes have two tasks. First, they correspond to topology-relevant locations like doors or corridor-crossings. Thereby the graph models the environment topology. Second, in environments with large open spaces, they might further provide topological redundancy by locations with favorable traversability. Favorable, for instance, with respect to visitor flow criteria or other specific requirements from the application.

The map further contains so called *ghost points*. Ghost points are (x, y)-positions in the world reference frame which act as invisible barriers. If the environment contains forbidden areas undetectable for the robot's sensors (e.g. staircases, glass doors, exits, etc.) ghost points are used to prevent the robot to go there by injecting them into the sensor data as virtual readings (see also section 2.4 and [12]).

The Expo.02 environment covers a surface of 315 m^2 and has 17 places of interest for the robots[1]. The map contains 44 segments on 44 lines, 17 station nodes, 37 via nodes and 20 ghost points (figure 2). Its exact memory requirement is 8 kbytes or 26 bytes per m^2. The ghost points (not shown in fig. 2) are at the entrance (bottom of fig. 2) and the exit (top) of the circulation area.

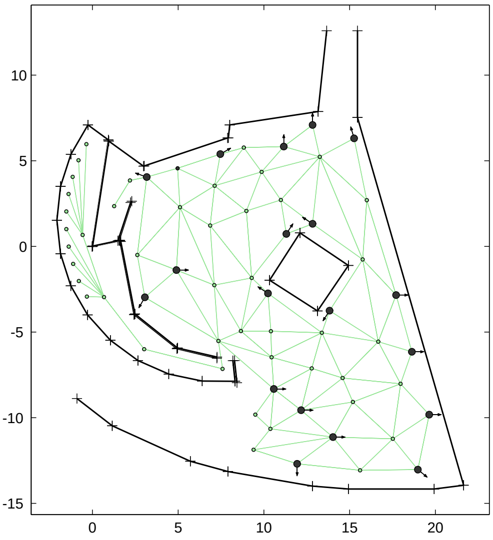

Figure 2. *The Expo.02 map*

2.2 Global Path Planning

With a topological graph, global path planning becomes a graph search problem for which a number of algorithms exist. From simple depth-first search with fixed costs to dynamic programming techniques and probabilistically learned cost functions for edge-traversability [15]. Global path planning in our case uses a priority-first search [22] and costs assigned to edges and nodes [27]. In a single-robot context costs are fixed, for multi-robot planning the costs of nodes are variable and depend on the distance to other robots. In the multi-robot case, besides paths, goals are shared as well, and must be negotiated among the robots. See section 2.5.

At Expo.02 we give visitors the opportunity to choose their next station of a tour by themselves. This closes the

---

[1] This number is subject to change as improvements with respect to visitor flow and scenographical criteria are made.

first loop to the environment which is asynchronous and has a cycle time in the order of 0.01 Hz (figure 6). Path planning in the graph of figure 2 is a matter of a few milliseconds in our implementation.

2.3 Command Queue

Paths from the global planner consist in a list of nodes. In the execution layer, a command queue passes the list to the behavior layer in a node-by-node manner. For via nodes it actives a (x, y)-position-only variant of the obstacle avoidance, in case of the last list node, the full pose (x, y, θ) must be met. Near the last node in the list (typically a station node), localization gets deactivated (10 cm in our implementation). This is because localization causes the robot pose estimates to be noisy such that for any position controller the goal cannot be reached. The reached condition for via nodes is treated on this level. It is satisfied when the robot enters a large disk around the node (2 m in our implementation).

2.4 Local Path Planning and Obstacle Avoidance

For mobile robots, the primary role of the reactive layer is that of a position controller. A number of constraints must be accounted for on this level: vehicle shape, vehicle kinematics, vehicle dynamics and, of course, environment dynamics which in the case of a mass exhibition is extensive. Since purely reactive obstacle avoidance methods usually suffer from local minima problems, we divide the task into a reactive and a path planning sub-layer [7, 21, 2].

For the reactive sub-layer we rely on the idea of the dynamic window approach (DWA) [11]. The method uses a simple model of the vehicle dynamics (maximal acceleration / deceleration) and can – with the appropriate extension [21, 2] – take into account an arbitrary robot shape. In comparison to the original version of the DWA our approach differs in the following points:

- Working with differential drive robots, the objective function trading off speed, heading and clearance are calculated in the actuator phase space (v_l, v_r) instead of the Cartesian (v, w)-space. This models the acceleration limits of the vehicle physically more properly.
- As in [21] and [2], we account for polygonal robot shapes. The robot shape is not hard-coded in our implementation and can be specified at boot time.
- Instead of using the distance to collision as a clearance measure, we use time to collision. This solves a singularity when the robot is turning on the spot (any collision would seem instantaneous because the distance travelled seems zero). It also means the robot will choose more clearance when travelling at higher speeds.
- Ghost points from the global map are taken into account. After a global-to-local transform they are injected as virtual sensor readings.

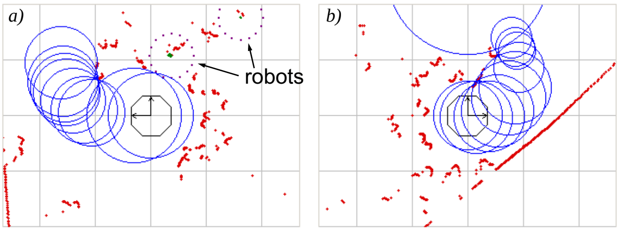

Figure 3. *Two situations from Expo.02 which show how the elastic band finds a smooth path around people. In a) there are two robots virtually blown up with ghost points.*

The dynamic window is part of the time- and safety-critical software of our robot. We therefore install this process as a deadline-driven real time task with a 10 Hz frequency. Special attention was paid to optimize its execution time to be short and predictable. For reasons similar to those mentioned in [21], we use a look-up table for the clearance measure.

The second sub-layer is a path planner which operates locally as it relies on sensory data without memory. A modified elastic band [20] is employed which generates smooth trajectories around obstacles (figure 3) and uses an NF1 navigation function [17] for initialization. Although the NF1 always yields topologically correct solutions (within its scope and if a solution exists), it generates unsmooth trajectories with the tendency to graze obstacles. The initial plan, generated by the NF1, continuously evolves towards a smoother curve.

Updates of the elastic band are implemented in a non-time critical thread which runs at several Hz. As soon as the elastic band "snaps", replanning is initiated. At Expo.02, this takes place typically in the order of 0.1 Hz.

For path planning and evolution of the elastic band, the robot is assumed to be circular and heuristics are used to ignore some sensor readings. This results in simplified and speed-up implementations. The simplifications are acceptable because the DWA ensures the dynamic, kinematic, and geometrical constraints.

The modifications of the DWA and the elastic band are described in more detail in [19]. At the lowest level finally, the speed controller, also installed as a real-time task, runs at a 1 kHz frequency (figure 6).

2.5 Multi-Robot Planning

For multi-robot planning we distinguish goal coordination and path coordination.

For the paths, environment dynamics from visitors is important. Visitors who play with a robot easily deviate the vehicle from its path and provoke path collisions where at planning time no collision had occurred. We therefore face the problem of replanning paths for all robots on-the-fly. This has to happen in real-time since we do not want the robots to stop and wait each time we detect a path deviation somewhere.

For this we employ a potential field approach where graph nodes receive costs proportional to the distance to a robot. Each robot assigns weights to those nodes which are in its current plan. The resulting graph superimposes the weights of all robots. Using this graph, the global planner delivers cost-optimal paths that contain nodes which are currently unused by other robots.

Theoretically, there is no guarantee on collision freeness with potential fields. One can construct (pathological) situations where it is cheaper for two robots to use the same node at the same time. More sophisticated methods [17] can provide a guarantee but are computationally unfeasible in our application. The advantage of this technique is its efficiency. It enables multiple robots to adapt their plans quickly and requires a minimum of shared information: a list of elements (i, w_i), with i being the node identifier and w_i its weight – a matter of a few bytes. Regarding collision freeness, a difficulty of the Expo.02 application – its high degree of dynamics – turns out to be an advantage. A population of robots that interacts with people so intensely is continuously in motion. Even if congestions might occur, they do not persist since the distribution of weights is in motion as well.

Coordination of the goals is necessary since a limited number of shared goal locations is to be allocated to multiple robots. Path coordination cannot avoid that several robots choose simultaneously the same station node since robots with the same goal would, due to the lack of redundancy, insist to go there regardless the costs.

To give visitors the choice of their next tour station the robot makes a proposition which is based on the currently unoccupied stations, the list of stations included in the tour and the stations already visited. The selected station is then reserved for this robot. More details on goal negotiation and implications for visitor flow can be found in [16].

2.5.1 The Robot-Sees-Robot Problem

Multi-robot coordination so far presented, depends on the knowledge of the robot positions. Even with a reliable and accurate localization, this creates critical interdependencies. Robots shall therefore be able to see each other on a raw data level. This, however, is not straight-forward with platforms of the same mechanical design which all measure at the same height since, at this height, the true vehicle size will be underestimated from the sensor readings.

Our approach is to mount two retro-reflecting beacons onto the vertical profiles in the blind zone between the two Sick laser scanners (figure 7). The Sick LMS 200 sensors provide an intensity signal which allows to easily extract the reflector information. We then use ghost points to artificially create a virtual robot contour at the extracted reflector positions (see also figure 3). Thus, obstacle avoidance provides an anytime fall-back solution.

2.6 Localization

Localizing a robot robustly in a mass exhibition environment is certainly a challenge. In former exhibition projects, localization was based on off-board hardware [8, 26] or environment modifications [18]. In our earlier work we employed features and an extended Kalman filter (EKF) [3]. This is also the approach for the three museum robots described in [13]. However, a robot doing (single-hypothesis) pose tracking can loose its track as the inherent data association problem is ignored. With our new localization technique introduced in [1], we address the data association problem and extend the conventional EKF localization approach to a global localization technique.

Unlike POMDP or Markov approaches [8, 26] where locations[1] are generated before they get evaluated by the exteroceptive sensors (as a grid or as particles), our approach to localization turns this process around: locations are generated as a direct consequence from sensory information. Features tell us *when* and *where* to place a location hypothesis. This allows to maintain always as many hypotheses as necessary and as few as possible.

The technique for hypothesis generation is a constrained-based search in an interpretation tree [14, 10, 9]. This tree is spanned by all possible local-to-global associations, given a local map of observed features $L = \{l_i\}_{i=1}^{p}$ and a global map of model features $G = \{g_j\}_{j=1}^{m}$. Furthermore, besides track formation, we present an algorithm for track splitting under geometric constraints. It relies on the same idea as hypothesis generation (search in an interpretation tree), forming thus a consistent framework for global EKF localization.

---

[1] We use the terms *location*, *position* and *pose* interchangeably. They denote all the full (x, y, θ) vehicle pose.

We briefly outline the approach[1]: The search space for hypothesis generation is the space of all possible associations of the observed features l_i and the modeled features g_j. The search space has the structure of a tree with p levels and $m + 1$ branches [14]. p is the number of observed features in L, m the number of modeled feature in G. The extra branch (called star branch) allows correct associations in the presence of outlier observations (false positives) and thus accounts for environment dynamics and map errors. During tree traversal, statistically feasible *pairings* $p_{ij} = \{l_i, g_j\}$ are sought given all uncertainties associated to the features. A pairing says that the observed feature l_i and the modeled feature g_j denote the same physical object in the environment (g_j is called an 'interpretation' of l_i). Although the problem is of exponential complexity, the geometric constraints reduce enormously the space to be explored. The constraints can be classified into two categories:

2.6.1 Location Independent Constraints

Unary constraint. We accept the pairing p_{ij} if l_i and g_j are of the same type, color, size or any other instrinsic property. Examples: l_i and g_j are both (x, y)-point features, or the length of the observed segment l_i is smaller or equal than the length of the modeled segment g_j.

Binary constraint. Given a valid pairing p_{ij} we will accept the pairing p_{kl} only if the two local features l_i and l_k are compatible to the two global features g_j and g_l. Examples: l_i and l_k are lines with the angle φ_{ik} between the lines. Then, the pairing p_{kl} is considered compatible if the angle φ_{jl} is the same. With point features, for instance, the distances $l_i - l_k$ and $g_j - g_l$ must correspond.

2.6.2 Location Dependent Constraints

The above tests do not involve the robot position L_h. Once this is known, a further class of constraints can be applied.

Visibility constraint. This constraint only applies to model features. It tests whether g_j is visible from the robot position L_h. Example: lines or segments can always be seen only from one side. If the robot is behind a wall, one of the two lines modeling the wall is invisible. With sensor specific parameters, the visibility constraint rejects features which are not detectable, for instance, because they are farer away than a maximal perception radius.

Rigidity constraint. A pairing p_{ij} is considered compatible if l_i and g_j, transformed into the same coordinate system given L_h, coincide (are at the same position). This is what happens in the matching step of any EKF localization cycle. Usually, g_j is transformed into the frame of l_i.

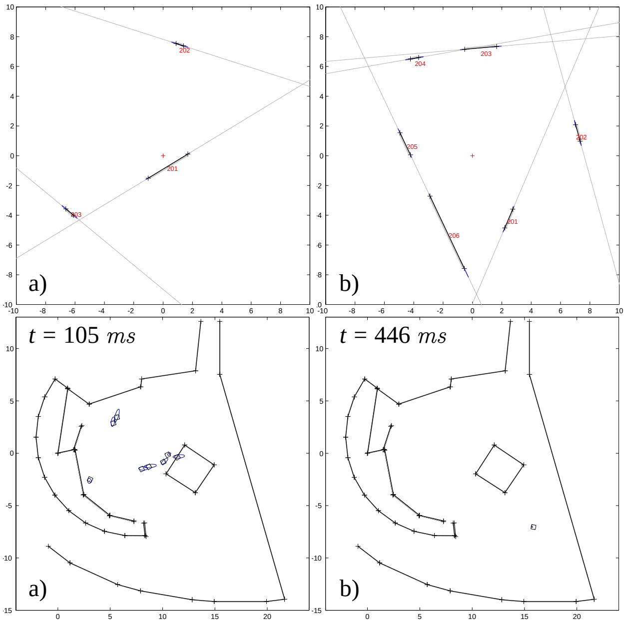

Figure 4. *Hypothesis generation. Given the local maps in a) and b), hypotheses are generated at locations where the local map 'fits' into the global map. In a) there are 15 hypotheses (with their 95% error ellipse); the location is ambiguous. In b) there is a single hypothesis; the robot is instantaneously localized. 't' denotes the execution time.*

Extension constraint. A pairing p_{ij} is considered compatible if l_i and g_j, transformed into the same coordinate system given L_h, fully overlap. Example: an observed segment l_i must be fully contained in the transformed g_j seen from the location L_h.

These constraints allow to discard whole subspaces (subtrees) from the search each time when an incompatible pairing is found at the root of such a subtree. With the uncertainties associated to local and global features, all decisions make use of the Mahalanobis distance on a significance level α.

Tree traversal is implemented as a recursive back-tracking search algorithm described in [9, 1]. The strategy is to first find a minimal number of valid pairings with location independent constraints such that a location estimate can be determined in order to apply location dependent constraints, too. Each time when the algorithm reaches the bottom of the tree, that is, the end of a branch where all observed features l_i could have been assigned to a model feature g_{j_i} or to the star branch, we have a valid robot location hypothesis. The pairings which support the hypothesis are put together in a *supporting set* $S_h = \{\{l_1, g_{j_1}\}, \{l_2, g_{j_2}\}, \ldots, \{l_p, g_{j_p}\}\}$ and thereby constitute a location hypothesis $h = \{S_h, L_h\}$. All hypotheses together form the set of robot location hypothesis $H = \{h_i\}_{i=1}^{n}$.

---

[1] Please refer to [1] for a more complete presentation.

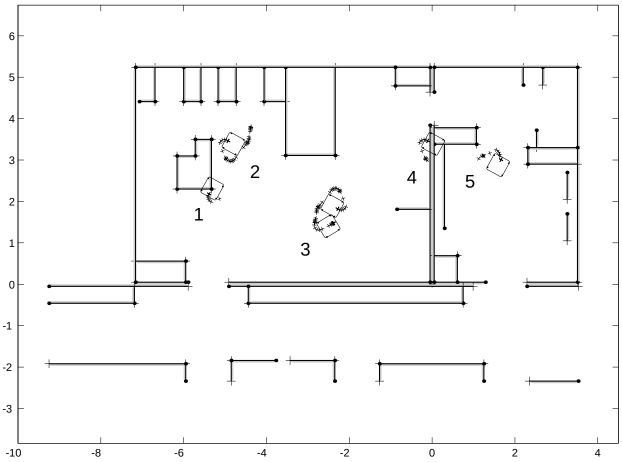

Figure 5. *Multi-hypothesis tracking. Starting with five hypotheses, track #3 turns out to be the true one after the last track (#2) was rejected at a distance of 1.89 meters.*

2.6.3 Estimating the Robot Location from S_h

With the supporting set, the (x, y, θ)-pose of the robot is not yet known. This is what the extended information filter (EIF) does. Given a supporting set with all associated uncertainties, it estimates the robot location and its covariance in the least square sense. The difference between the EIF and the EKF is that the former is the batch estimator formulation of the latter (which is recursive). This is needed, because, during hypothesis generation, there is no a priori knowledge on the robot location which formally means that the state prediction covariance, usually called $P(k|k+1)$, is infinite. With the EIF, this can be properly expressed as $P^{-1}(k|k+1) = 0_{3\times3}$ since covariance matrices are represented in the information matrix form, that is, by their inverse.

Figure 4 shows two examples of hypothesis generation in the Expo.02 environment. With multiple discrete hypotheses, to be localized is simply expressed as having a single hypothesis. For line extraction we use the method described in [4]. Extraction times are around 20 ms. Localization cycle time is in the order of 10 Hz.

2.6.4 Multi-Hypothesis Tracking

The main reason for lost situations during tracking is incorrect data association. This occurs typically when there are several statistically feasible pairing candidates for an observation. Choosing the closest one (the most widely applied strategy – called nearest neighbor standard filter), leads to filter inconsistency and mostly to filter divergence if it was the wrong one. Therefore, besides uncertainties in the value of measurement, robust pose tracking must also account for uncertainty in the origin of measurement [5].

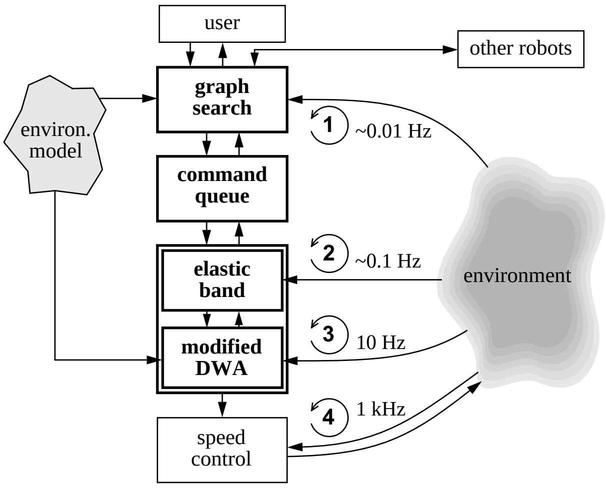

Figure 6. *Layering and control loops with their precise (3, 4) and typical (1, 2) cycle time.*

We look for an algorithm which re-generates hypotheses during tracking as soon as there is no guarantee anymore that the correct association of an observation can be done. This property has algorithm 2 in [1] which, given a predicted location, a local and a global map, splits up into multiple offspring hypotheses if statistical compatibility with several supporting sets can be established at that location. It has the identical structure than the algorithm for hypothesis generation but employs location dependent constraints *only* and does *not* recur with a refined position estimation. In this manner the algorithm finds all supporting sets in the vicinity of the initially predicted location L_h.

After each hypothesis has been tracked, there are three cases: *(i)* hypothesis confirmation, *(ii)* hypothesis rejection and *(iii)* hypothesis splitting. Track rejection takes place when the predicted location is not supported anymore by location dependent constraints on the level α. When track splitting occurs, their locations get newly estimated and the best one is taken. ‘Best’ in a goodness-of-fit sense, expressed by the joint Mahalanobis distance.

Figure 5 shows an experiment how the robot converges towards the true location after a short trajectory. Note that by geometry only (or geometric falsification respectively), false tracks get rejected quickly. No free-space information is needed. The execution times for tracking are around 10 ms per hypothesis.

Finally, figure 6 shows the resulting decomposition of the architecture, the four control loops and their cycle times. The DWA sublayer requires the position of the ghost points which explains its connection to the environment model. Localization can be seen as a behavior as it is switched by the execution layer and connected to the robot’s sensors. It is however invisible since there is no direct connection to other components or behaviors. It acts in the background continuously correcting the odometry. Thereby, all other framework components can consider odometry as being 'perfect'. For multi-robot coordination, robots are connected on the level of the planning layer.

3. Implementation

In view of the requirements from the Expo.02 application and former experience in robot design and system integration in our lab, the decision was taken to construct a robot from scratch. The outcome is shown in figure 7 and is described in detail in [25]. For navigation, a PowerPC G3 at 380 MHz serves as main CPU running the deadline-driven real-time operating system XO/2. The robot is fully autonomous and does not rely on off-board resources. The only exception is a central server for multi-robot planning. The approach described in section 2.5 does not require this at all, but with the current implementation using HTTP, opening a connection takes more time than data transmission itself. Using a central unit, the number of communications could have been minimized (see also [16]).

4. Operation Experience and Discussion

From the opening of the exhibition on May 15th, 2002 until July 17th, 2002, 283,319 persons visited the 'Robotics' exhibition (4427 per day, 422 per hour). Assuming an average visit duration of 15 minutes (typical for mass exhibitions), around 105 persons share the 315 m^2 circulation area with ten robots. In other words, the robots encounter heavy environment dynamics, mostly benign but also hostile (people who try to outwit the robots). The overall travel distance during this period was 1,259 km.

Safety. It was never observed that a robot was the cause of a dangerous situation (e.g. with small children, elderly or handicapped people). Collisions with visitors occur but are typically provoked by the visitors themselves. The lack of additional sensors close above floor (IR or ultrasonic), like the robot in [8], is easily bearable with the combination of tactile plates and soft bumpers. Blocked-situations due to bumper contact can also be handled by the interactive part (robot expressing friendly menaces).

Reliability. The division of obstacle avoidance into a purely reactive part with high model fidelity and a planning part with local scope is a powerful conjunction. Very often, groups of visitors form a *U*-shaped obstacle or leave only small 'holes' for passage. The robots have no difficulty to escape from such situations, and due to the fact that we account for Robox' true octagonal shape, narrow passages are efficiently used (figure 3). Further, the elastic band generates good-looking trajectories. The concept for the robot-sees-robot problem works well in general. The employed model however (blowing up a robot contour at the detected reflector position) turned out to be an oversimplification in certain situations.

Figure 7. *RoboX, the robot built for Expo.02. It sports two SICK LMS 200, a PowerPC G3 at 380 MHz and a Pentium III at 700 MHz.*

The reliability of localization was a surprise in view of the environment dynamics and the fact that we use laser data only. A fall-back solution with lamp features extracted by a camera looking to the ceiling was prepared but never used. Lost situations do occur but are, so far observed, mostly caused by people. Sometimes a lost situation is the consequence of a failure of an other component or has an unknown cause. Examples of the first type include staff member and visitors who push or rotate the vehicle. Global localization as described in section 2.6 is then very useful. Often, the robot can be instantaneously relocalized within all people, enabling it to resume the current operation. The specific geometry of the Expo.02 environment is helpful here since it contains few symmetries.

Besides the specific qualities of our approach described so far, we believe that the use of geometric features for navigation is also a very appropriate choice, particularly for highly dynamic environments. During feature extraction, sensor readings are sought which satisfy a spatial model assumption (from a line or a corner). Thereby, the extraction process acts as a filter saying which reading is to be taken for localization and which one is to be ignored. This filter relies typically on sound regression techniques and works independently on whether the robot is localized or not (opposed to the 'distance filter' in [26]). A group of people standing around the robot does not produce evidence for the line extraction. Spurious segments, for example on line-like objects carried by people, can occur and are treated by the star-branch in the localization algorithm.

At the time of this writing, two functionalities are not yet operational. Multi-hypothesis tracking (single-hypothesis tracking is used instead) and path coordination. For multi-robot planning, goal negotiation is running since July 1st, 2002. It avoids successfully that several robots choose the same station nodes. Experiments with the algorithm for path coordination have been done and were promising. Multi-hypothesis tracking is also subject of ongoing work. The objective is to be able to stay localized even in the situations described above (robot gets pushed or rotated).

5. Conclusions

The scale of the Expo.02 project is a unique opportunity to validate techniques and to gather long-term experience on various levels. All conclusions at this point are preliminary as the exhibition continues until October 20th, 2002. The framework presented in this workshop paper (the architecture and the components), meets all initial requirements very well. We described how the three-layered architecture was adapted to our application and outlined the environment model, global path planning, local path planning, reactive obstacle avoidance, multi-robot planning and localization. Striving for compact representations and efficient algorithms, we obtained light-weight implementations in terms of memory requirements and computation times. This allows Robox, the platform for Expo.02, to be a truly autonomous robot in a truly challenging application.

References

[1] Arras K.O., Castellanos J.A., Siegwart R., “Feature-Based Multi-Hypothesis Localization and Tracking for Mobile Robots Using Geometric Constraints,” IEEE Int. Conf. on Robotics and Automation, Washington DC, USA, 2002.

[2] Arras K.O., Persson J., Tomatis N., Siegwart R., “Real-Time Obstacle Avoidance for Polygonal Robots With a Reduced Dynamic Window,” IEEE Int. Conf. on Robotics and Automation, Washington DC, USA, 2002.

[3] Arras K.O., Tomatis N., Jensen B., Siegwart R., “Multisensor On-the-Fly Localization: Precision and Reliability for Applications,” Robotics and Autonomous Systems, vol. 34, issue 2-3, pp. 131-143, February 2001.

[4] Arras K.O., Siegwart R.Y., “Feature Extraction and Scene Interpretation for Map-Based Navigation and Map Building,” Proc. of SPIE, Mobile Robotics XII, Vol. 3210, 1997.

[5] Bar-Shalom Y., Li X.-R., Multitarget-Multisensor Tracking: Principles and Techniques, ISBN 0-9648312-0-1, 1995.

[6] Bonasso R.P., Firby R.J., Gat E., Kortenkamp D., Miller D., Slack M.,“Experiences with an Architecture for Intelligent, Reactive Agents,” Journal of Experimental and Theoretical Artificial Intelligence 9(2), 1997.

[7] Brock O., Khatib O., “High-Speed Navigation Using the Global Dynamic Window Approach,” IEEE Int. Conf. on Robotics and Automation, Detroit, USA, 1999.

[8] Burgard W., Cremers A.B., Fox D., Hähnel D., Lakemeyer G., Schulz D., Steiner W., Thurn S., “Experiences with a Interactive Museum Tour-Guide Robot,” Artificial Intelligence 00(1999): 1-53, 1999.

[9] Castellanos J.A., Tardos J.D., *Mobile Robot Localization and Map Building: A Multisensor Fusion Approach*, Kluwer, 1999

[10] Drumheller M., “Mobile Robot Localization Using Sonar”, IEEE Transactions on PAMI, 9(2), p. 325-32, 1987.

[11] Fox D., Burgard W., Thrun S., “The Dynamic Window Approach to Collision Avoidance,” IEEE Robotics and Automation Magazine, 4(1), pp. 23-33, March 1997.

[12] Fox D., Burgard W., Thrun S., “A Hybrid Collision Avoidance Method For Mobile Robots,” IEEE Int. Conf. on Robotics and Automation, Leuven, Belgium, 1998.

[13] Graf B., Schraft R.D., et al., “A Mobile Robot Platform for Assistance and Entertaiment,” International Symposium on Robotics, Montreal, Canada, 2000.

[14] Grimson W.E.L., Lozano-Pérez T., “Localizing Overlapping Parts by Searching the Interpretation Tree”, IEEE Transactions on PAMI, 9(4), p. 469-82, 1987.

[15] Hu H., Brady M., “Dynamic Global Planning with Uncertainty for Mobile Robots in Manufacturing,” IEEE Transactions on Robotics and Automation, 13(5):760-7, 1997.

[16] Jensen B., Froidevaux G., Greppin X., Lorotte A., Mayor L., Meisser M., Ramel G., Siegwart R., “Visitor Flow Management using Human-Robot Interaction at Expo.02,” Workshop on Robots in Exhibitions, IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Lausanne, 2002.

[17] Latombe J.-C., *Robot motion planning*, Kluwer, 1991.

[18] Nourbakhsh I., Bodenage J., Grange S., Lutz R., Meyer R., Soto A., “An Affective Mobile Robot Educator with a Full-time Job,” Artificial Intelligence 114(1-2): 95-124, 1999.

[19] Philippsen R., Jensen B., Siegwart R.,“Smooth and Efficient Obstacle Avoidance for a Tour Guide Robot”, submitted to IEEE Int. Conf. on Robotics and Automation, Taiwan, 2003.

[20] Quinlan S., Khatib O., “Elastic bands: connecting path planning and control,” IEEE Int. Conf. on Robotics and Automation, 1993.

[21] Schlegel C., “Fast Local Obstacle Avoidance Under Kinematic and Dynamic Constraints,” IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Victoria BC, Canada, 1998.

[22] Sedgewick R., *Algorithms*, 2nd ed., Addison-Wesley, 1988.

[23] Simmons R., Goodwin R., Haigh K., Koenig S., O’Sullivan J., “A Layered Architecture for Office Delivery Robots,” First Int. Conf. on Autonomous Agents, 1997.

[24] Simmons R., Smith T., Dias M.B., Goldberg D., Hershberger D., Stentz A., Zlot A., “A Layered Architecture for Coordination of Mobile Robots,” in *Multi-Robot Systems: From Swarms to Intelligent Automata*, Kluwer, 2002.

[25] Tomatis N., Terrien G., Piguet R., Burnier D., Bouabdallah S., Siegwart R., “Design and System Integration for the Expo.02 Robot,” Workshop on Robots in Exhibitions, IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Lausanne, Switzerland, 2002.

[26] Thrun S., Beetz M., Bennewitz M., Burgard W., Cremers A.B., Dellaert F., Fox D., Hähnel D., Rosenberg C., Roy N., Schulte J., Schulz D., “Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva,” Int. Journal of Robotics Research 19(11): 972-99, 2000.

[27] Tschichold-Gürman N., Vestli S.J., Schweitzer G., “Operating Experience with the Service Robot MOPS,” 3rd European Workshop on Advanced Mobile Robots (Eurobot’99), Zurich, Switzerland, 1999.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Visitor Flow Management using Human-Robot Interaction at Expo.02

B.Jensen, G.Froidevaux, X.Greppin, A.Lorotte, L.Mayor, M.Meisser, G.Ramel, R.Siegwart

Autonomous Systems Lab - Swiss Federal Institute of Technology - CH-1015 Lausanne

Abstract

In this paper we will regard the task of operating a public mass exposition with several autonomous robots at a time. This implies questions regarding human-robot interaction, multi-robot control and interaction management. To enable human-robot interaction while guiding a tour we outline the SOUL environment. Multi-robot and interaction management are regarded with respect to visitor density and visitor flow. Concluding we will present and discuss results from the Swiss national exhibition Expo.02 in the time from 15.05.02 to 17.07.02, corresponding to 5293 hours of total robot operation time up to date and in interaction with 283319 visitors.

1. Introduction

Public space experiences in recent years are proof of a remarkable progress in mobile robotics. This enabled the operation of a public mass exposition with ten autonomous mobile robots at a time during the Swiss national exhibition Expo.02.

Having several identical robots serving as tour-guide and main attraction of an exposition during a five-month period from 15.05.02 to 20.10.02 created a special situation. Men and machine operating in the same space make reliable and safe robot operation is mandatory. Ten and a half-hours operation per day, seven days per week over the exposition period imposed high demands on robotics hardware. In addition to this, visitor flow and fun factor of an exposition are important to operators and financiers of a public mass exposition.

To meet these requirements, the interactive mobile tour-guide RoboX has been recently developed by our lab. Developing the interactive part for the exposition meant always taking into account the demand for visitor flow and entertainment. These criteria translate more or less directly into guided tour and unconstrained interaction. Our solution is embodied in the SOUL (Scenario Object Utility Language) system [8] controlling guided tour and interaction together.

The fact of having several robots at disposal makes them easier and faster available for the visitors, but requires a resource management for the exposition space. The autonomous nature of our robot evokes the question of centralized or distributed system architecture, which we will regard later on.

Figure 1: RoboX interacting with people visiting Expo.02.

Closely related with the multi-robot control, we try to support the natural visitor flow direction from entry to exit by constraining the displacement of the robot.

In general high visitor density and a rapid visitor flow constrain interaction. Since these parameters are external, we seek a system allowing for a maximum of interaction under the current conditions.

Concluding we will evaluate these elements under real world conditions based on experience gained at the Expo.02.

2. Related work

We will look at mobile robot experiences in public spaces, arguing that the mobility of the platform and the direct presence of both human and robot render interaction particularly interesting. We find the importance of improving human robot interfaces [1], to help visitors in interacting with mobile robots. Face and emotional state machines were found useful elements for tour-guide-robots [2]. The Mobot Museum Robot Series [3,4] focused on the interaction. Robustness and reliability was identified as an important part of a public robot. Several experiences with the museum robots showed further that the visitors do not always behave cooperatively with the robot and switch between seeing it as a simple machine or a tour-guide. Another permanent installation is at the “Deutsches Museum für

Kommunikation” in Berlin, where three robots welcome the visitors and invite them to play with a ball [5].

Summarizing, we can state that the development of public robots has to take into account the differences in visitors’ behavior. First of all, the robot needs to sense the presence of visitors in order to react appropriately. We may distinguish if the robot is seeking an interaction or if it is already giving a tour and interacting with someone else [4]. It was further found that the time visitors spend with the robot is not easily predictable or controllable. Some visitors get bored after a couple of minutes with the robot, others spent days with it. During this time the visitors’ behavior changed from collaborative to investigative interaction.

3. RoboX

During Expo.02, the time which visitors can spend with RoboX is rather limited. We decided to use intuitive means of communication in order to use this time as efficiently as possible. The design of the robot should use common features for communication, situating its appearance somewhere between anthropomorphic and machine. The face of RoboX is intended the source of communication helping the visitors to feel more comfortable when communicating with the robot.

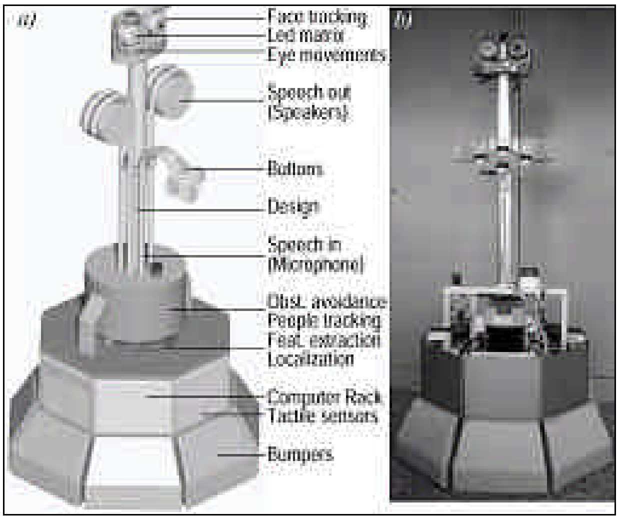

Figure 2: Outline of RoboX elements and photo of the first prototype.

Even though collaborative interaction will mainly take place between one visitor and the robot, we anticipate that a certain audience of other visitors will follow this interaction. For good visibility we constructed RoboX (figure 2) to be of approximately average visitor’s height. Basically, the robot consists of a mobile base with an interactive top, making the face easy to look at.

Two differentially driven wheels located at the center of the robot allow on the spot turns. Two castor wheels, one at its back and one, with a suspension at its front, ensure the stability of the mobile base. Obstacle avoidance and reliable localization [6] ensure that the robot knows at all times its position and does not collide with visitors or parts of the exposition.

As an additional means of security, touch sensitive plates and foam bumpers ensure that the robot stops if running into anything. Two SICK Laser scanners mounted at knee height provide environmental information for navigation and interaction. A camera mounted in one of the robot’s eyes provides additional information for the interaction.

Furthermore, the mobile base houses motor controllers, batteries for 10h autonomy, a PowerPC 750 clocked at 400 MHz dedicated for navigation and obstacle avoidance and a Pentium III running at 700 MHz, 128 MB RAM on Windows 2000 for all interaction tasks. Both computers can communicate with each other over a 10 Mbit/sec local Ethernet and with a central computer over wireless interfaces to allow monitoring the state of the robot for security reasons. Technical details are discussed in [7].

4. Interaction at Expo.02

Interaction of visitors with several robots in a public exposition is a complex task. First of all we will present how interaction between RoboX and a visitor is realized. We will distinguish static and dynamic elements, which help in making each tour of the robot individual. By taking into account dynamic elements, which we will precise later on we aim at giving the robot an aura conscious of its environment.

Since RoboX is giving a tour it will stop at several stations and supply information related to a certain part of the exposition. With the several RoboXs running at the same time we faced the problem of multi-robot coordination to avoid having several robot s intending to go at the same place at the same time.

Finally we will present how parameters like visitor flow and visitor density are taken into account to provide the most of interaction under the current conditions of the exposition.

4.1 SOUL

We will briefly present SOUL, controlling interaction on RoboX. It aims at combining elements of a guided tour with human-robot interaction. The tour the robot is giving presents a certain amount of information on several parts of the exposition. They will change rarely if ever, for the period of the exposition. Henceforth static scenarios can easily represent this information. A scenario is in the SOUL context the succession of robot actions as speaking, moving and similar actions for a limited amount of time.

Intelligent appearance can hardly be achieved by repeating these scenarios over and over again. Therefore we use methods of changing presentation and methods of adaptive behavior to avoid repetition.

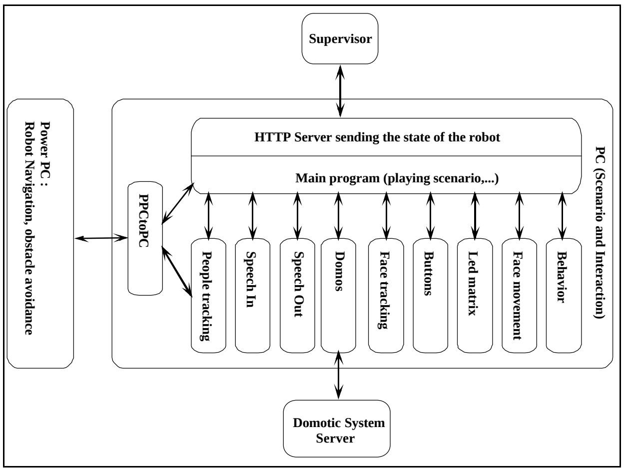

Figure 3: Structure of the interactive system. The supervisor is a separate computer allowing the operator to monitor of the robot's operation.

One way to avoid repetitive behavior is to provide several alternatives of the text and actions presented. Thus changing the method of presentation. The tools available to the SOUL system for creating such scenarios are exhibited in figure 3.

In addition to this permutation approach, we aimed at having a robot responding to a couple of dynamic events, which can occur during a tour. This changes its behavior. Such events can be visitors are blocking the robot or even hitting its bumpers. They are playing with the buttons without being asked to or are pressing the emergency button. The battery of the robot is low or other. From the point of view of interaction one can see these signals as a certain acceptance of the robot by the visitor. From the point of view of a guided tour, however they are exceptions and are treated by SOUL as such. Technically SOUL will interrupt the current scenario and execute a corresponding exception scenario telling the visitor that it is aware of his actions, before resuming the tour. RoboX will treat one exception at a time.

4.1.1 SOUL sensors

RoboX is using several sensors and algorithms to achieve awareness of its environment. Simple switches detect events like visitors pressing the emergency button, the interactive buttons or hitting the bumpers. The obstacle avoidance provides information when visitors are blocking the robot.

In addition the robot is aware of visitor presence in its surrounding by means of face and motion tracking [8].

4.1.2 SOUL expression

There are three interfaces available to communicate with the visitor. To express itself RoboX is using synthesized speech in English, French, German and Italian using Mbrola [9] and LAIPTTS [10].

Figure 4: Three facial expressions. From left to right: happy, surprised, and angry.

The interactive buttons can be illuminated to indicate in which mode they are in (language choice, yes/no, etc.).

For visitors the most expressive part remains the face (figure 4) imitating several grimaces and by means of a small LED display mounted in one of the eyes display symbols and short animations.

4.1.3 Behavior component

Our aim was to create individual tours according to the visitor's action, until yet their action affected the tour only shortly by starting the appropriate exception scenario. With the behavior component presented in [11] RoboX started to accumulate impressions during a tour and to adapt its behavior accordingly.

Here we have to distinguish two main cases in which RoboX uses the expressions. The first case happens to emphasize or illustrate its speech and is controlled directly by the scenario running. In the second case, the expressions are more like the mirror of the subject's emotions. For the representation of this internal state we chose the Arousal-Valence-Stance affect space [12], because of its three dimensional representation which is very intuitive to use. The robot current state is therefore defined as a point in the three-dimensional AVS space (see figure 5). In this space, six basic expressions regions are defined as: sadness, disgust, joy, anger, surprise and fear.

Also, we use the origin of the space as a reference expression that can be considered as a calm state. Of course, other expression regions can be defined in this space. But, we decided to limit ourselves to those seven regions in order not to overwhelm the visitor with many reactions to subtle for our robot expressive capacities.

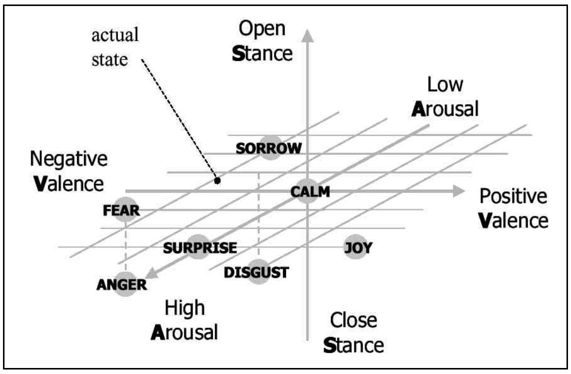

Figure 5: Representation of the six basic expressions and the neutral expression in the AVS space.

	Pitch	Rate	Volume
Fear	high	very fast	medium
Surprise	very high	very fast	very loud
Joy	high	fast	loud
Sorrow	little low	slow	very soft
Disgust	low	very slow	soft
Anger	very low	very slow	very loud

Figure 6: Variation of pitch, rate and volume for the standard expressions.

The internal state is mainly communicated using the synthesized voice, face movements in some cases symbols are shown on the LED screen. Figure 6 shows how the internal state effects the synthesized voice.

4.2 Multi-robot coordination

Figure 7 exhibits the layout of the exposition. Presentation stations are defined near particular objects in the expositions. There are several places where robots welcome visitors, thus tours can start simultaneously. At the time of writing there are fifteen presentation stations all over the exposition space. Finally there are goodbye stations close to the exit.

Each station corresponds to one scenario in the SOUL system, providing visitors with the necessary explanatory or entertaining information. Tours can be created by a succession of several presentation stations. Two stations are except from the tours and are permanently occupied with a dedicated robot. They have tasks of taking pictures from the visitors and presenting a slide show. In these cases the tour consists of one station only.

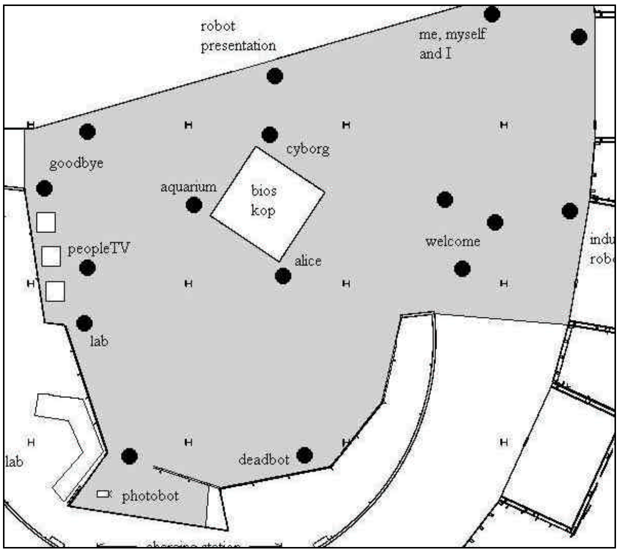

Figure 7: Scheme of the 315 m^2 exposition area with the presentations stations shown.

Working with multiple robots makes resource allocation an important point. In order to avoid having several robots presenting the same object an assignment has to be made at a certain moment.

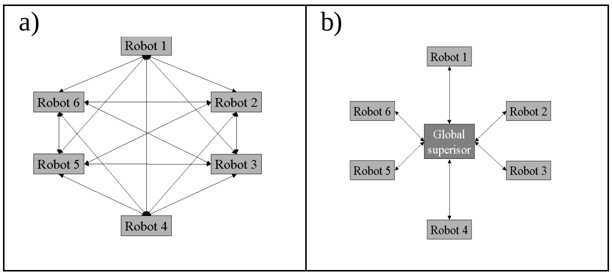

Figure 8: Communication structure a) without central server, b) with central unit.

In the beginning we solved this problem by assigning several stations exclusively to one tour which was operated by one robot all day. The tours were designed to have robots working spatially separated in order to avoid collisions among robots.

With ten robots operating the exposition this was no longer feasible, since it would result in tours of one or two stations only and thus quasi-static mobile robots. Improved obstacle avoidance allowed the robots to see each other and to avoid collisions. This enabled a dynamic assignment of stations to a robot for the duration of its presentation. The station is released thereafter and can be used by other robots.

This is modeled by a list of all stations and their state. Stations are free until reserved by a robot. The robot can chose among the free stations in order to avoid deadlocks. Care has to be taken that robots decide successively to avoid several robots choosing the same goal.

Figure 8 shows two different communication architectures for the assignment process. On the left side communication takes place among the robots only. Even though this uses only intelligence and information present in the robots it requires a complex communication. Each robot has to communicate with all other robots and needs to monitor which robots are currently active. Assuming N robots at hand all reserving one station this results in N·(N-1) communications.

By adding a central instance as shown in figure 8 b) this number drops to N communications. We opted for this solution since it results in a much easier and thus more reliable communication scheme. Technically this global instance could be run on one selected robot, so that the group of robots still can be considered as an autonomous system.

Multi-robot coordination in our case is based on local decisions by each robot. When terminating a presentation the robot will ask the state of all exposition stations from the global instance. This request blocks the global supervisor until the robot reserves a specific station. The decision, which station to reserve is based on the free stations, the list of stations included in this tour and the stations already visited. The first free and unvisited station in the tour list is reserved.

4.3 Visitor density and interaction

Expo.02 was considered a mass exposition with several thousands visitors per day. During the preparation of this project we anticipated up to 500 visitors per hour, which assuming a 15 minutes stay inside the exposition results in 125 visitors which are at the same time enjoying the robots.
Visitor behavior can hardly be anticipated. To ensure a functioning of the exposition even with a lot of visitors on the hand and to provide intensive interaction when viewer visitors are in the exposition, four exposition modes were defined:

1. *Wait for visitor*: with few visitors, so that robots wait for one to come close enough before starting to talk and ask him which station he would like to see.
2. *Visitor's choice*: more visitors, so that the robot can ask permanently whether the visitor wants to go to a station without talking to no one.
3. *Robot's ch*oice: even more visitors, so that the robot will decide what is the next station and go there without asking.
4. *No move*: too many visitors for the robot to move, so that each robot will stay with one station and present it permanently.

The exposition mode is defined manually by the staff. It is included in the data provided by the global supervisor, so every time the robot requests the state of the exposition stations it receives an update of the state and can adapt accordingly. Figure 9 shows how this is taken into account by the SOUL system:

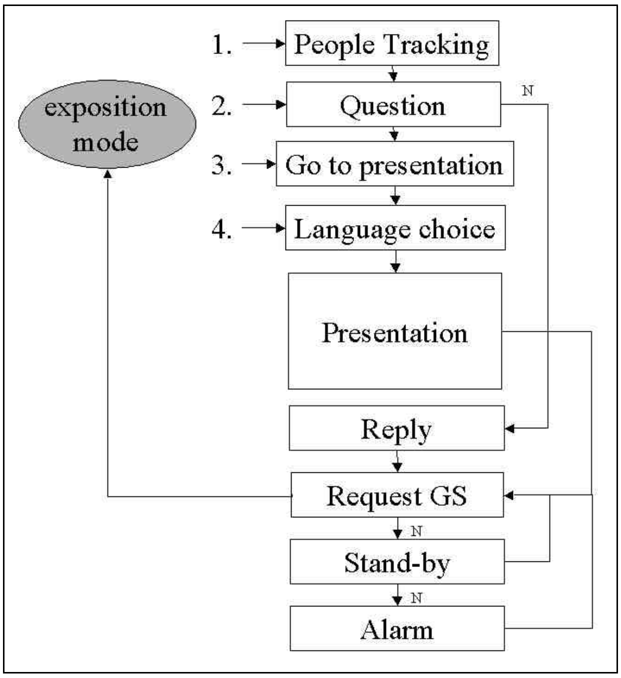

Figure 9: Structure of the SOUL sequence for a typical presentation station.

Depending on the exposition mode the scenario starts either with people tracking (*wait for visitor*), the question "Do you want to see … ?" (*visitor's choice*), the robot moving to the station (*robot's choice*) or directly with the language choice (*no move*).
These blocks are executed successively except if the visitor declines to go to a station. In this case SOUL jumps directly to the reply block commenting in some way the visitor's decision.
The request from the global supervisor is executed either after the reply block or after the presentation of a station. It provides all empty stations at this time the choice is made as explained in the paragraph above.
If no empty station is available and all empty station already have been visited during this tour the robot can not go on. Then it starts one of several stand-by scenarios. These are presentations, which are not located at a specified place in the exposition. The robot talks about itself, sings or makes funny faces. Thus the robot gains time during which a presentation station may be released by another robot.
After the stand-by scenario the robot request once again the exposition state to find a free presentation station. If one is found the next scenario is run. Otherwise the robot continues to play stand-by scenarios and to request the global supervisor until either a presentation station is available or it has run out of stand-by scenarios. In the latter case the global supervisor will give an alarm and the staff can interact. Starting the robot with another tour may solve this problem.
To avoid having several robots giving the same presentation a station remains blocked by one robot until it starts moving on to the next station.

4.4 Visitor flow

We estimated the average visit to 15 minutes in order to meet the visitor flow requirements. Previous test in our lab [8] proved it difficult for the robot to make visitors leave. In general their interest span is not directly related to the duration of a tour.
Visitor flow is channeled by two factors. First of all the number of stations the robot visits. The robot to visits S stations before it executes the goodbye scenario, which is located near the exit. By this proximity we aim at encouraging visitors to leave. The goodbye scenario is special in the way that it resets the list of stations visited during a tour and sets the counter of stations visited back to zero.
Throughout the exposition a tour will always lead visitors closer to the exit. This eases navigation and helps maintaining the visitor flow. Technically this is realized by a list of possible next presentation stations. Each presentation scenario is assigned an individual list, containing only stations to support the direction of the main visitor flow. When requesting exposition state from the global supervisor the robot will seek only stations which it has not yet visited and are closer to the exit than it is currently.

5. Results

In the period from 15.05.02 to 17.07.02 an average number of 4427 people were visiting the exposition every day. The minimal number of visitors was one time 2299 the maximum achieved was 5473. The average number results in a visitor flow of 422 persons per hour on 315 m^2 exposition space with up to ten robots in operation. This corresponds to a load of 84.3% percent of the planned maximal flow of 500 visitors. The maximum flow corresponds to a load of 104%.

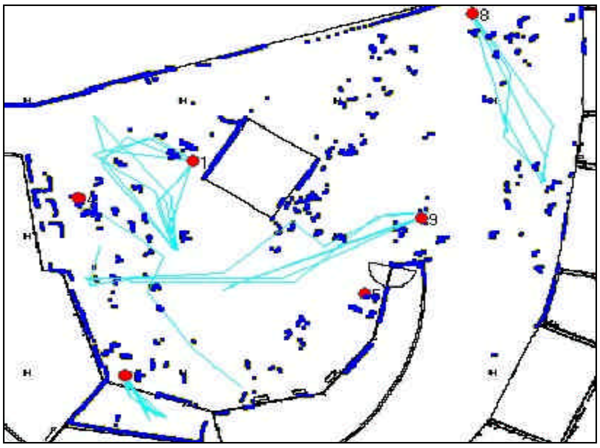

Figure 10: Map of exposition with 6 robots and laser scanner data showing visitors and robots (circles).

The global supervisor system is operational since the 01.07.02. Until yet the exposition mode *visitor's choice* was active approximately 95% the mode *robot's choice* 5% of the time. We experienced ten days with more than 5000 visitors, even in this crowded environment robots managed to move to their goal in a reasonable time, so that the mode *no move* was never used. Up to date the mode *wait for visitor* was never used, since the robots are most of the times surrounded by interested visitors anyway. Figure 10 shows a typical situation.

With currently three stand-by scenarios, alarms of a robot running out of those scenarios occurred approximately once a week. With two additional stand-by scenarios we aim at reducing this rate further.

Visitors stay between 10 and 45 minutes with the robots. We tried to control this by changing the tour length from two to ten stations without noticing an impact an the visitor's stay. People just move on to the next robot or even stay with the current one. Here enhanced environmental information, like motion information of the visitor or face recognition might help creating more convincing scenarios. We found that visitors quit a robot approximately after four stations, which is the actual tour length. The average number of visitors during the 17 days of operation of the global supervisor rose slightly to 4576 per day. This makes it hard to prove a quantitative effect on the visitor flow. However, observation of the crowd shows that visitor appreciated having the choice to go to a station. This adds a little interactive element to the tour.

6. Conclusion

During over 5293 hours of operation, 283319 visitors interacted with the robots in the time from 15.05.02 to 17.07.02. SOUL seems to provide an appealing compromise of a guided tour and unconstrained interaction. For the last two and a half weeks the exposition was running with a multi-robot resource control scheme taking into account the visitor density and supporting visitor flow.

Quantitative parameters like visitor flow and density meet the planning parameters. By enhancing environmental perception aim at creating even more convincing human-robot interaction.

References

[1] Burgard W., A. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, S. Thrun, “Experiences with an interactive museum tour-guide robot”. AI (114), 1999, pp. 3-55.

[2] Thrun S., M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, D. Schulz, “MINERVA: A Second-Generation Museum Tour-Guide Robot”. In: Proc. ICRA, USA; 1999, vol.3; pp.1999-2005.

[3] Willeke T., C. Kunz, I. Nourbakhsh, “The History of the Mobot Museum Robot Series: An Evolutionary study”. In: Proc. FLAIRS 2001, Florida.

[4] Nourbakhsh I., J. Bobenage, S. Grange, R. Lutz, R. Meyer, A. Soto, "An Affective Mobile Educator with a Full-time Job". AI, 114(1-2), October 1999, pp. 95-124.

[5] Graf, B., W. Baum; A. Traub, R. Schraft, "Konzeption dreier Roboter zur Unterhaltung der Besucher eines Museums". In: VDI-Berichte 1552, VDI-Verlag, Düsseldorf, 2000, pp. 529-536.

[6] Arras, K., R. Philippsen, M. Battista, M. Schilt, R. Siegwart “A Navigation Framework for Multiple Mobile Robots and its Application at the Expo.02 Exhibition”. Workshop: Robots in Exhibitions, IEEE/RSJ.IROS, 2002, Switzerland.

[7] Tomatis N., G. Terrien, R. Piguet, D. Burnier, S. Bouabdhalla, R. Siegwart, “Design and System Integration for the Expo.02 Robot”. Workshop: Robots in Exhibitions, IEEE/RSJ.IROS, 2002, Switzerland.

[8] Jensen, B., G. Froidevaux, X. Greppin, A. Lorotte, L. Mayor, M. Meisser, G. Ramel, “The interactive autonomous mobile system RoboX”. In: Proc. IEEE/RSJ.IROS, 2002, Switzerland.

[9] Dutoit, T., et al, "The MBROLA Project: Towards a Set of High-Quality Speech Synthesizers Free of Use for Non-Commercial Purposes". In: Proc. ICSLP'96.

[10] Siebenhaar-Rölli, B. et al, “Phonetic and Timing Considerations in a Swiss High German TTS System”. In: E. Keller, G. Bailly, A. Monaghan, J. Terken & M. Huckvale, “Improvements in Speech Synthesis”, pp. 165-175, Wiley & Sons, 2001.

[11] Mayor L., B. Jensen, A. Lorotte, R. Siegwart, “Improving the expressiveness of mobile robots”. In: Proc. ROMAN(2002), Berlin, Germany.

[12] Ekman P., R. Davidson, “The Nature of Emotion: Fundamental Questions”, Oxford University Press, New York, 1994.

Proceedings of the IEEE/RSJ IROS 2002
Workshop on Robots in Exhibitions
Lausanne, Switzerland, October 2002

Interactive Exhibition Design: Robots in Exhibitions

Olaf Arndt

Machine-Performance Group BBM
Hasenheide 71
D-10967 Berlin, Germany

What's new about robots as being part of interactive exhibitions is the preoccupation with high technology used by non-experts, visitors, users and advertisers who want to create a so-called «immersive» atmosphere surrounding the customer in a sensual and physical way. Immersive surroundings are supposed to lead to an «intelligent space» which – by the use of «smart technologies» – accommodates itself dynamically to the needs of its fitters.

For the simple reason that the required knowledge is to be found beyond the intersections of the traditional disciplines it is obvious to put into action unconventional thinkers already during the planning of such exhibitions.

Such reflections may have led to invite us, the members of the artists' group BBM (Beobachter der Bediener von Maschinen, Observers of Operators of Machines), more often in the last years to contribute specific robotic solutions to bigger content based or thematic exhibitions.

After 10 years of artistic work with machines and simple robotic systems we accepted a commissioned work for the Thematic Area «Knowledge» at the World Exposition «Expo 2000» which took place in Hanover (Germany). In the course of two years and sheltered by the umbrella organization «Centre for Media Arts, ZKM, Karlsruhe» we started an extremely complex-ridden experiment in «collective robotics» connecting 72 autonomous media-robots in a swarm-like acting network. The most important thing in the scenography of this momentous undertaking was not only the technical functioning, but the attempt to pour the complex relations of a networked world into an easy to understand picture. Using the available technologies of robotics and control we wanted to tell a plausible story of the self-determined acting of autonomous entities in the world of information technology. For this purpose and under the creative direction of BBM more than 30 designers, programmers, film makers and authors came together and developed a scenario which was then realized by research institutes, mechanical engineers and suppliers of media-technology.

Learning from the New Flow

In 1998 we successfully sold Expo 2000 a concept based on "autonomous robotics" that established an anti-hierarchical system, the so-called "bottom-sideways" structure (top down is classical management, bottom up means a complex situation management, still according to hierarchies; a third approach invented by Sadie Plant focusses guerrilla-like rhizome structures, the bottom-sideways model).

However, the concept was later changed around by the Fraunhofer Institute, so that our fundamental idea of a self-organising network became a Carousel with "artefacts" chained to each other virtually, machines that permanently went round in a circle: the Theme Park had firmly re-stated its identity as a Mega-Carousel.

Originally we were interested in the ambivalence of autonomy, of self-organisation and of collectivity: these are leading concepts in social movements, but they are also used by engineers when defining "navigating attitudes". One of the focal points for BBM, since 1997, was the way so-called autonomous robots, using vehicles, in staged scenarios, interacted positively or even aggressively with visitors, enjoying free movement throughout the available space, steered by sensors and software. A particularity of these machines is that they are capable of co-operating with each other, and can thus combine into groups, or more precisely, they can simulate them. Moreover, they react to being "accosted".

These are both typical situations in any public space where people gather together and seek to establish contact. Because the machines and the visitors move around in the same space, the analysis of attentiveness, and of visitors' flow-modes, plays a role of outstanding importance in our work.

Ironically, or perhaps because of the in-grown tendency of our society towards ever greater specialisation - seen as an absolute prerequisite if you want to "stay in the market" - BBM found itself able, using the above-mentioned know-how, to act as consultant to the management of a Hypermarket every bit as effectively as the so-called emotioneering specialists (emotion engineers), who think up ways of using space to measure and control the moods, the buying attitudes and the vigilance of visitors.

Such posts are usually filled by "post-heroic" managers, the sort of people who planned the media circus in the Gulf War, or who recruit workers during the build-up phase of a new secret service section: all jobs that require much the same profile.

The filmmaker Harun Farocki collects material concerning this kind of activity, for a film on shopping malls. In an interview for "Jungle World", he tells of a laboratory where they test floor surfaces to find out which have the most favourable influence on buying behaviour. Now, "putting the brakes on the economy" has always been capitalism's biggest problem. Sure, we'll slow down ...some other time!

So, "slowing down" has to be scientifically organised. Carpets and flowerpots, acting as slalom-posts, raise the likelihood of a sale. The watchword "get 'em moving" is what lies behind the design of these semi-public places in the new Shopping-Cities. Farocki calls it "science as the practice of magic". In view of the seamless interweaving of science and urban planning, are we not in the presence of the quintessence, the core, of a capitalistic organisation of space? Is this not magic, are these not the actions of charlatans, of modern shamans?

This reminded us of the famous "gang map of Manhattan", an example architects invented in the 70ies to illustrate complex urban traffic functions. it is two maps with identical plans, air views of the Manhattan silhouettes, positioned congruent one on the other to compare the lines that are drawn inside the silhouette: one is the official rectagonal plan where everything is organised according to the compass, the traffic law, the arithmetical language of modern western city surveillance. "down the avenues" means not only helpful orientation but as well being obedient, acting as a citoyen who likes to live in harmony with all democratic principles.

Totally different does the "gang map" look. the gang map emerged from the analysis of thousands of interviews that social workers had with gang members about their daily social lives as part of a 70ies integration procet. one of the topics they discussed with each of them was how they managed to cross Manhattans different social areas without conflicting all day long with the NYPD. So sociologist and urban planners together drew a plan noting carefully each individual way through the rectangular jungle of official Manhattan. they found gates of huge importance that non-gang inhabitants of the same quarters might not even have notice even though passing them by every day: old rotten fence doors, vast building cellars which where connected with next blocks cellars by a hole that somebody brake in the wall, backyards with stairways that could be "bridged" to the wall of the next yard by using a piece of wood that will be hidden after the illegitimate crossing etc.

The most interesting result for those researchers was that after they noted all this passes and tried to get them congruent with the official plan - there only were crosses on different levels and no parallelity, exactly none of it.

This anecdotes, the experiences of that research we later tried to use as part of the Expo concept.

Expo forced us to hand in a feasible solution of what they called a "visitor management": we had to answer their questions how to "get 3000 people an hour through hall filled with

72 robots"? we voted for a concept with 72 total autonomous entities that where able to go in any direction they like to go, that are able to flock and to catch up by accelerating if they lost their group, and that they were allowed to change their group whenever "they liked" or to form one big group out of two or three former groups. this was exactly the spontaneous guerilla concept giving them all freedom a machine can get by being equipped with motor, batteries, sensors and chips, programming them for all abilities of grouping and set them into a totally open ambient in which they can react to the pressure they feel from visitors. swarming is a survival system and our problem was to simulate mating and hunger, the biggest forces in nature.

Since we couldn’t make them "hungry" because detecting a power source that they could drill in to fill their batteries was technically to complicated and since we couldn’t develop a really neuronal network software because there where to many interferences security advices from the Expo we decided to rely on the visitors behaviour to get the "pressure" we needed for the machine’s "reactions".

But this coincidated with two very serious contradictive decisions Expo took: they organized the "visitor flow" as powerful vectorial stream going straight through the exhibition. one entrance, one exit, 10 minutes time and a merchandising shop on the one end of the hall. second decision was to contract a company for the software development which is on the one hand the biggest German research facility but on the other they never programmed an autonomous machine before. they always build stationary industrial streets like in used in steel factories etc. with lots of complicated details but all fixed to an installation.

So after half a year of heavy communication problems with them that they finally saw as "linguistic" (they meant "philosophical") differences they decided to program a merry-go-round with virtual tracks. again the official system couldn’t made congruent with the idea of inscenating the subversive.

You can easily imagine the result: each time the visitors vector crossed the virtual circle of the merry-go-round we had a chain reaction which ended up in a full stop of 72 machines. their local abilities of decision and restarting where to small.

The philosophical difference was practical now: we always dreamt of a space where machines were in constant movement forming new patterns day by day and reflecting it by a 72 channel media network. the organisation always had nightmares about how they can successfully stop the permanent traffic so that the visitors get a chance to survive Expo.

Why robotics for exhibition designers?

In my view technical engineering is the future of art and design. As far as technical proceedings determine the work of designers, art is going to be the tangible hope of shaping and designing.

Technical engineering as material is going to equal to stone and colour. It thus transforms inevitably and radically the idea of craft and design schools. It may sound like mere presumption but we think, that art will at best conquer the leading role once owned by the universal scholar. There are good reasons for this view. The starting-point is the synthetic reflection of disparate social and thus operational functions in terms of responsibility.

In a process of creation determined by the vertiginous rotations of the market, the professional outsiders, the artists are in a position to see «where, in the white room of technics, one can find the concealed door that the engineer doesn't find» (Peter Weibel). In other words, the qualified designer operating within the limits of his discipline is not able to see the exit of the fully covered chamber of traditional creation principles by which he might leave his self-constructed professional prison.

Perhaps it is not a great step the artist can do, but the result will clearly differ from the possible performances of an academic culture depending on economic guidelines.

Critical reflection is the condition of substantial creation. Without understanding and analysis of the conditions there is no forward-looking styling and design.

So, in the future science will be the agent of creation, just as Eros and Thanatos have influenced art for thousands of years, and as actual understanding of the basic needs has shaped the design of commodities. The fusion of the three poles allows a new view on genetic engineering and the continuation of alchemic handicraft work made by early universal designers and proto-designers.

Nowadays the medieval basic idea of robotics to build an artificial homunculus is filled with new content. It no longer aims at adjusting the space to the limited abilities of the body but to adjust the body to the given conditions of the space. The Cyborg is the brilliant feat to shift the scope of activity of the designer into the inside of the body. Robotics as ancestors of genetic engineering are endogenous design.

Under the circumstances of «biological power» (Foucault) interactivity means survival training.

It is the responsibility of designers to ensure that science is for the benefit of society. Science has therefore to get rid of short-term orientation towards the pursuit of fast profit and the fixation on total control.

Most thrillingly design takes on a substantial meaning in this context: a critical reflection of socially pertinent motor activity resulting in the creation of reactions that are dedicated to men. Its creations adopt utopian character and a concept of education that heeds perspective instead of moral intentions.

Since the fifties the interdisciplinary departments of science and art in the USA are working on the decisive project of robotics. In view of the technical possibility of human mechanization this research is oriented towards the potential of human independence. Here artists are only employed as catalysts of the hard sciences or the engineers that are – not even by using the most expensive tracking systems – not able to locate their concealed door. Adopting the point of view that design must embrace a higher level of values there is a missing link in the fact that the tools do not get into the hands of those who want to change the conditions of life into something that is worth to be mentioned.

Designers play an important part in this process. Their education should ensure that they do not get stuck in quotations, nor stick to bygone styles or to a formal modernism.

It is important for them to learn to apply the «state of the art technology». We don't talk about «know-how» but about «use it!»

It is characteristic of recent and especially of «soft» technologies that they can solely be «decontextualized» when used in an anarchist fashion.

It is a fact that the aesthetic and device-based metastasis of the so-called "information society" absorbs more than 90% of the designers activity and force him to deal with inferior exercises in the two-dimensional cultivation of surfaces. From our point of view it would be very welcome to unmask this as mere mumbo jumbo of market strategists.

Design will be able to take on a relevant function within culture and society after having redressed these deformities.

What is «interactive exhibition design»?

In February 2002 we were working with students in Kolding who were in their third and fourth year of studying «Interaction Design». We tried to explain what a workshop in robotics can achieve within three weeks. Here the focus was on the concept of «embedded technology». We wanted to outline a realistic idea of potentialities and limits of the use of technology.

We did by no means want to disappoint high hopes for the efficiency of contemporary hard- and software but train the ability to use such technologies in order to tell a comprehensible story. We chose the topic «exhibition design».

Which components can be used to communicate a certain content?

What is feasible under the conditions of continuous operation and mass processing, both characteristic for exhibitions using high technology?

We developed seven different scenarios, referring to potential claims of virtual customers. This covered the whole range of banks, business enterprises, museums and NGOs.

We do not simply focus on the subject of narrative technique but consider as well pragmatic demands concerning the guidance of visitors, the operationality and the so-called interactivity.

We wanted the students to develop a prototype of the planned robotic project and, moreover, a paper describing in words and pictures the complete scope of their idea and particular approach to the subject. We accepted printed products as well as webpages, but the students were expected to elucidate their respective choice.

In addition to that the students were encouraged to document the work process on video and then cut the material into a three-minute-film.

Thus several levels of media came into action. The students were not simply asked to choose and develop a particular subject. Soon they realized that a workshop on robotics confronts them from the start with the necessity to build well-working teams and that it is not good enough to find kindred spirits but to split up the job in a professional way.

One particular question raises when a decision has to be made concerning the use of technology: Is there any chance for the visitor to realize that a technical installation in an exhibition is running in an «intelligent» way? How can he detect that the installation depends on his presence and conduct when spitting out data – but does not just spin around or drone «stupidly»?

If a technical installation is supposed to arouse interest its application has to come as a surprise. That is not a matter of its real electric or mechanical complexity but results from the intelligent construction of the story it tells.

Apart from the obvious aspect of engineering (How does my control board work in the «real world»? How does the set run beyond my computer?) the embedding of technology has a social meaning. A narration with many odd words is being told. The future designer therefore does not only have the traditional task to do the layout but becomes a translator too.

In conjunction with the complexity of the necessary knowledge, responsibility grows for an extended dimension of design. Design does not only contain a new model of cooperative work or interdisciplinary work, as people once used to call it.

Whilst claiming that «interaction design» is not only a subordinate and dependent service (design of surfaces) but a forward-looking educational qualification we are aware of accepting the fact that the use of high technology is going to transform society.

Although the place of action called «exhibition» is often regarded as a children's playground of the «just-for-fun-society» the fact is still true that powerful institutions like scientific, army and health research as well as the countless state-owned control organizations are testing or recycling the required technology in exhibitions and leisure parks.

From this point of view the study of robotics in exhibitions seems to be «contagious» in every respect.