

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

Wolfram Burgard¹ Panos Trahanias² Dirk Hähnel¹ Mark Moors³
Dirk Schulz³ Haris Baltzakis² Antonis Argyros²

¹University of Freiburg, Computer Science Department, Germany

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

³University 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

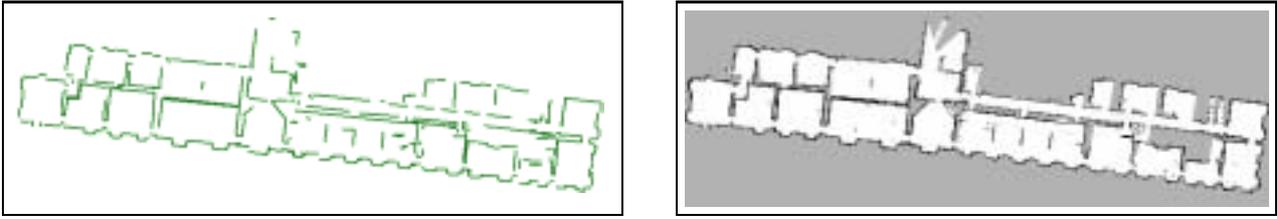


Figure 1. Line feature map (left) and occupancy grid map (right) of an exhibition site generated by the robot

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 recalculated. 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 (SJPDFs) [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

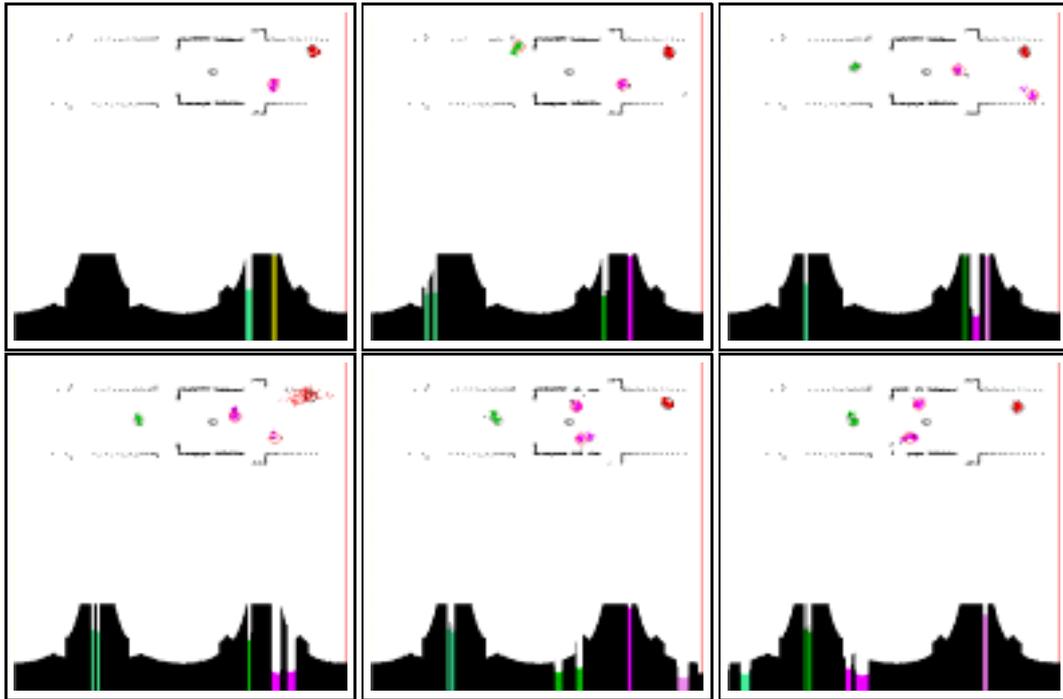


Figure 2. Tracking people using laser range-finder data.

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

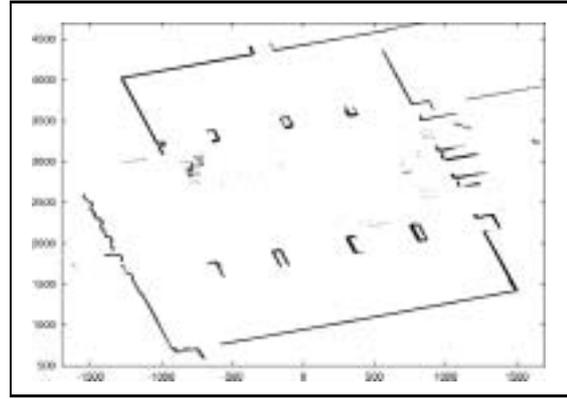
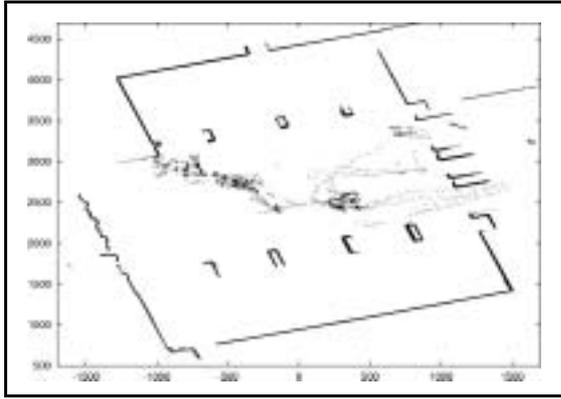


Figure 3. Maps of the Byzantine and Christian Museum in Athens created without (left) and with (right) people filtering.



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 camera's 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 continuously 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

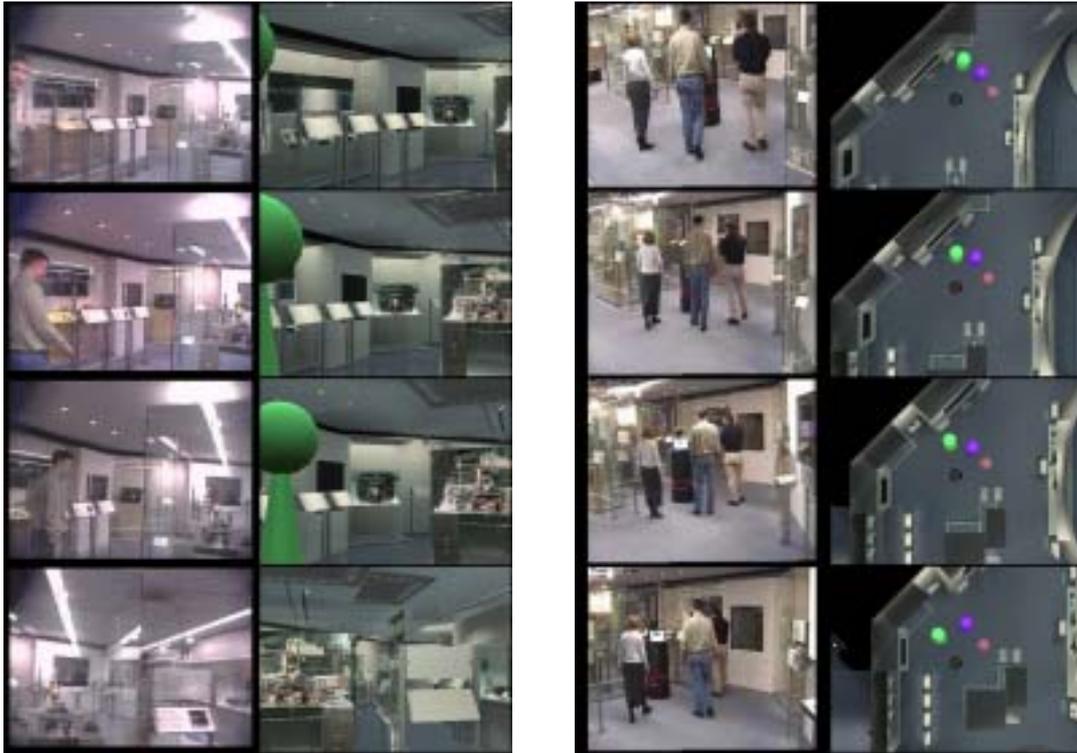


Figure 5. The enhanced 3D visualization allows arbitrary view-points. The left sequence shows the real and the virtual view through the robot’s cameras. The right images show the robot guiding three people through the museum and a bird’s eye view of the scene.

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 Hellenic 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 courtesy 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.

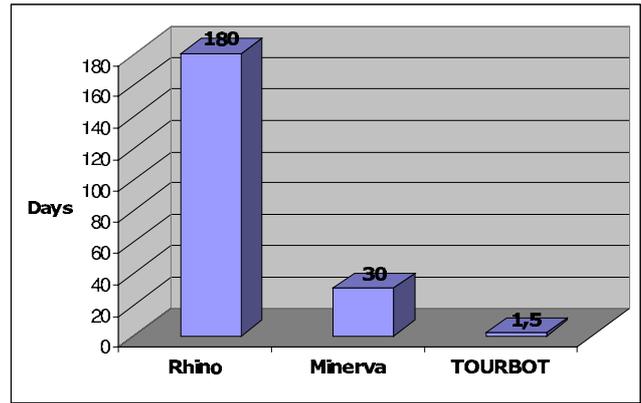


Figure 11. Time required to install the different tour-guide systems Rhino, Minerva, and TOURBOT.

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 robot localization. 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/>.