Designing SANDRA: An Autonomous Tour Guide Robot for the University of Technology, Sydney

Michael Behrens, Marc Carmichael, Mitesh Patel
University of Technology, Sydney, Australia
{michael.j.behrens, marc.g.carmichael, mitesh.patel}@eng.uts.edu.au

Abstract

This paper describes how a team of final year mechatronic engineering students developed an autonomous robotic system intended to act as a tour guide during events such as University open days and explores the opportunities this project presented to extend their knowledge and skills. The specifications of the project required the system to localise and navigate autonomously within a known environment while avoiding collisions with any people or obstacles not included in the prior area map. In addition to these requirements, the system needed to locate humans as potential clients, approach and greet them, offer directions and if required take the guest on a guided tour of the university. While taking the subject Advanced Robotics the students were able to develop a functional first prototype of the system and carry out initial tests. Following the completion of the subject a small number of the students opted to continue working on the project developing a second prototype using the knowledge gained and further enhancing their learning experiences. While this project mainly involved integrating existing well known algorithms, software and hardware, it provided an excellent opportunity to enhance the mechatronic engineering skills of the students involved.

1 Introduction

Recent years has seen the introduction of autonomous tour guide robots into the public realm. Examples include *Rhino* at the Deutsches Museum in Germany [Burgard et al., 1998], *Minerva* at the Smithsonian's National Museum of American History [Thrun et al., 1999], and *RoboX* at the Swiss national exhibition Expo.02 [Siegwart et al., 2003]. As well as fulfilling the functional requirements of providing guided tours, autonomous tour guide robots needs to face the challenges of safe, reliable navigation in crowds and interacting with people in an intuitive and appealing manner [Burgard et al., 1998]. This paper details the development of a new autonomous tour guide robot named the Student Autonomous Navigating and Directing Robotic Assistant (SANDRA) by a group of final year students enrolled in the BE Mechanical and Mechatronic Engineering Undergraduate Program at the University of Technology, Sydney (UTS).

The robot development occurred in two distinct phases. The first phase, the development of the first prototype, occurred during the UTS subject Advanced Robotics (AR) in the spring semester of 2007. The objectives of this subject are to provide students with the essential skills necessary to develop robotic systems for practical applications. Hence, part of the AR subject which is offered to final year mechatronic engineering students is to develop a robotic system. The projects offered vary from year to year but mostly involve the development of some form of robotic system and are designed to test and develop the student's problem solving and teamwork.

The second phase, the development of the second prototype, began at the conclusion of the AR subject. Five of the nine original students decided to continue developing the SANDRA system during the 07/08 summer break and throughout the autumn semester of 2008. This work did not count for course credit and was simply undertaken to gain experience and to enjoy the challenge. An additional objective of this phase was to improve the hardware so that the next group of AR students would be able to utilise it without modification.

2 First Prototype

The first prototype of SANDRA was developed by a team of nine students. This team was further divided into three sub teams with three students in each with a different component of the SANDRA system assigned to each group. The three components were Localisation and Mapping (LAM) [Behrens et.al., 2007], Path Planning and Obstacle Avoidance (PPOA) [Nguyen et. al., 2007] and Human Robot Interaction (HRI) [McQuilty et. al., 2007] and are explored in detail below.

2.1 Hardware

The mobile robot was based on a Pioneer 2DX robot produced by Activmedia. This provided a ready built mobile robot with all of the low level sensing and control electronics required for basic mobility already implemented. The Pioneer 2DX is a differential drive robot with a castor wheel for stability. It has an array of sonar sensors and wheel encoders. This sensor array was augmented with a SICK laser range finder with a viewing angle of 180 degrees and capable of measuring distances up to 80m. Having such a complete system available meant that most of the hardware required for the LAM and PPOA teams was already implemented, allowing the students to focus on the development and implementation of the control and functional algorithms. The only hardware modification which was made by the LAM group was to replace the old and un-reliable battery connection system with a more reliable system. A new and larger battery pack was sourced to allow the system to operate for a useful length of time.

2.2 Software

All the control was based on software running in Linux. The Kubuntu distribution was used as support was readily available from postgraduate students in the UTS ARC Centre of Excellence for Autonomous Systems (CAS). Software written to control SANDRA was connected to the hardware through the robotic middleware software Player, version 2.0.4. This is an open source software project designed to be used on mobile robotics. This package was selected as it was taught in the subject Advanced Robotics. It also has a wide selection of drivers available including drivers for the Pioneer 2DX and SICK laser range finder which facilitated simple hardware integration. The player software package also has a number of pre-made mapping and localisation drivers available. These drivers did require significant setup and tweaking to optimise them for operation in the desired area however it allowed for a fairly robust final solution as lots of the development has already been carried out. Player also has a large online support community.

2.3 Localisation and Mapping

The specifications defined for this section were to develop a sub-system capable of localising the mobile robot within the foyer areas of buildings 1 and 2 of the UTS city campus. While in operation the localisation sub-system provides location information to both the PPOA and HRI sub-systems. The assignment scope allowed the system to rely on a pre-generated area map to simplify the design effort. The students were required to generate this area map in a format suitable for use by the onboard control system. The maps were generated using data from the SICK laser range finder gathered as the robot was manually driven around the operational area. The data was logged to allow different methods to be tested quickly and easily.

The first mapping technique investigated by the students was the Map Reference Iterative Closest Point (MRICP) driver written by Tarek Taha for use in the 2006 RoboCup Rescue competition [Dissanayake et. al., 2006]. The students initially selected this driver as it was designed to generate a map in real time. Unfortunately it was found that while the MRICP worked well in small confined areas similar to the environment which it was originally designed for, it gave poor performance when it was used in large open areas. This poor performance was attributed to the lack of features in the foyer of UTS which caused the algoritim to become ‘lost’ very easily. Figure 1 shows a typical map which was able to be generated using the MRICP algorithm before the algorithm failed.

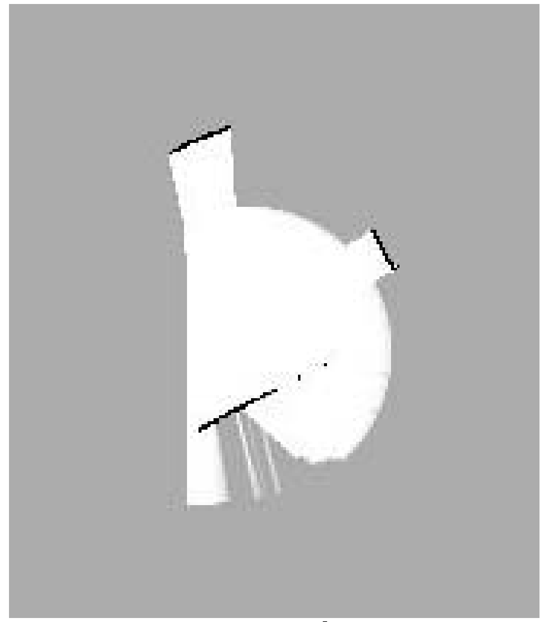

Figure 1: Map generated using MRICP

When it was found that the MRICP algorithm was not suitable the students investigated the suitability of another algorithm, PMAP which is supplied with Player. While this algorithm does not provide an optimal solution, as it requires offline data processing and manual stitching of map sections, it was readily available and was able to be implemented in the required timeframe. This was achieved using PMAP_TEST, a utility provided with Player which uses the PMAP algorithm along with logs of robot laser and odometry data to generate detailed maps offline. Processing data offline allows the use of a more robust brute force style of map generation. Figure 2 shows a typical map generated by PMAP_TEST. For this map section, the robot was driven in a loop around the top right quadrant. While this method was much more successful at producing results, it was still possible for the algorithm to become lost or misaligned over distance. To resolve this issue the area was mapped in small sections and manually stitched together to form a complete map. The final map was created with resolution of 0.1m/pixel.

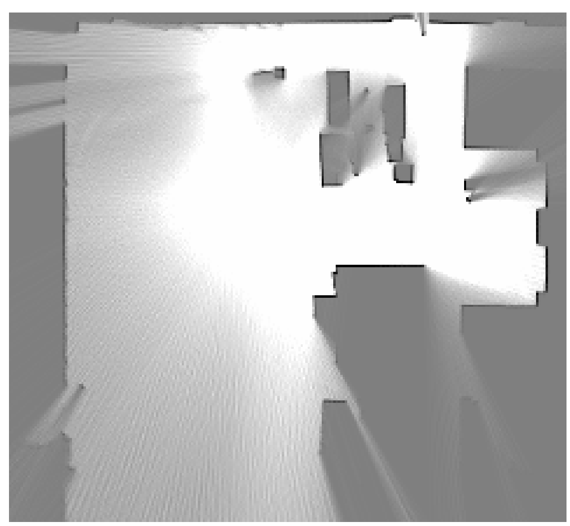

Figure 2: Map generated using PMAP_TEST

Localisation was achieved using the Adaptive Monte-Carlo Localisation (AMCL) driver. The AMCL driver is a particle filter based localisation algorithm supplied with Player. It utilises the robots laser range finder, odometry and the pre-made map of the operating area to estimate the position of the robot. The students found that it worked best when given a relatively accurate estimation of the starting position. An initialisation routine was developed to take an initial estimate of position and drive in a figure of eight pattern until the location estimate converged. Once converged, the position estimate was able to be maintained even when tested in moderate crowds. Figure 3 shows a visualisation of the AMCL driver during initialisation using the playerv utility. The large red circle indicates the uncertainty in the robot position which will collapse to a point as the robot localises, the smaller dark red circle is position of the most likely position hypothesis and the red dots are the particle positions which will also converge as the robot localises. The black sections indicate areas which will be detected as occupied by the laser range finder. This imformation is extracted from the pre-generated map of the operating environment.

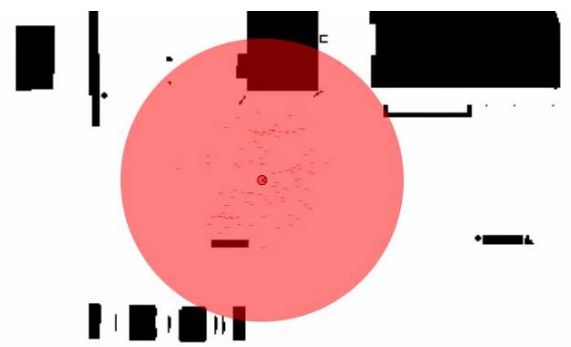

Figure 3: AMCL visualisation during initilisation in the UTS foyer of building 1

Objectives Achieved

The students implementing the LAM were able to create a system which could localise itself within the operating area if given a reasonable estimate of the starting location. It could maintain a location estimate while driving anywhere in the operating area even when it encountered un-expected obstacles such as people. The mapping system was adequate as it allowed a map of the area to be generated. However, it was a very labour intensive task.

Areas for Improvement

Several areas were identified for further development. Firstly, if the mapping routine could be refined and automated this would allow the robot to be released at night prior to the event to generate a new map which includes the locations of semi-permanent obstacles such as tables, stalls or bins. Secondly, a second laser range finder mounted towards the rear of the robot could provide more information to the AMCL algorithm and assist in localising in heavy crowds. Finally, as the update frequency of the AMCL algorithm is quite slow, it would be useful to implement a technique which uses the odometry of the robot for short term localisation and only uses the AMCL to perform corrections and updates to the estimated position.

2.4 Path-Planning and Obstacle Avoidance

The specifications defined for this section were to develop and implement algorithms to find the shortest traversable path from the current location, provided by the LAM sub-system, and the desired position, provided by the HRI sub-system. They were also responsible for generating algorithms for avoiding people and other unmarked obstacles while traversing this path. The PPOA team chose to extend the scope of their assignment by writing their own algorithms from scratch instead of using the available pre-built drivers.

The operation of the path planning algorithm developed by the PPOA team can be described as follows. The map is divided into a grid of equal square sections with a side length of 0.5 meters. Using a search based on the A* algorithm [Goto et. al., 2003] the shortest path is computed between the current location and the goal. This path consists of a string of un-occupied interconnected squares. While this algorithm was sufficient to fulfil the requirements of the AR subject there was insufficient time to locate and resolve all of the software bugs in the code. As a result, occasionally the planner would malfunction and fail for no apparent reason. In addition to this, due to the setup of the cost function, the selected path tended to follow closely to walls even when traversing large open areas.

The obstacle avoidance algorithm was closely tied to the path planning algorithm. If an obstacle was detected the map section containing the obstacle was temporarily marked as blocked and a new path planned to the goal. This method was adopted to prevent the robot doubling back while avoiding an obstacle and worked providing there was sufficient room to traverse around an obstacle and that the users did not seek to confuse the robot. Unexpected obstacles were detected using the sonar sensor array to ensure that the system did not collide with obstacles which cannot be detected by the laser. .

Objectives Achieved

The outcome of the PPOA section was influenced by the groups' decision to develop the navigation algorithms from scratch. This decision, coupled with the flexibilty of the project specification, enabled the group members to focus their learning to areas which they wanted to enhance. As a result, their knowledge and ability relating to software development for player based control systems has been improved to a greater extent than could have been achieved by simply implementing the pre-built navigation drivers avaliable in player. As a concequence of this decision the finished product is not as reliable and robust as a system implemented using the navigation drivers available in player.

The PPOA group was able to produce two algorithms which were operational under ideal conditions. The path-planning algorithm was able to successfully determine a path in the majority of situations. The system was also able to avoid un-expected obstacles under good operating conditions. The two algorithms developed where successfully integrated with the LAM section. Figure 4 shows the robot system used to intergrate these two components.

Figure 4: Pioneer 2DX used for LAM and PPOA development.

Areas for Improvement

As a result of the decision to focus their learning on software development, the algorithms developed by the PPOA team did not provide the optimal navigation solution for the mobile system. Both algorithms would require extensive debugging and further development to achieve the same level of robustness as the available Player drivers which was outside the scope of the AR subject. For these algorithms to be useful a number of issues needed to be addressed.

The cost function used during path selection needed to be refined to encourage the planner to select a natural path to the destination. This would enable SANDRA to behave in a manner similar to a human tour guide while leading visitors to their destination. In addition to this, a number of code bugs exist which can cause the planner to get stuck 'inside' a wall and be unable to reach the destination. These bugs needed to be identified and resolved.

The obstacle avoidance routine is very basic and could be greatly improved in a number of ways. As the obstacle avoidance routine operates by temporarily marking the corresponding map sections as blocked when an obstacle is detected, it can easily surround itself with virtual obstacles if operating in a crowded environment or with a user who is trying to confuse the robot. There are few ways to resolve these issues without developing a whole new approach to obstacle avoidance. Another important improvement which should be implemented is to fuse the data from both the laser and sonar sensors when detecting obstacles. This method would benefit from increased accuracy and range afforded by the laser allowing obstacles to be detected and reacted to earlier while still utilising the sonar sensors to detect obstacles invisible to the laser.

An alternative solution identified, was to fulfil the project requirements using the navigation drivers available with Player. These drivers were implemented in the second prototype.

2.5 Human Robot Interaction

The specifications defined for the HRI section were to enable the robot to communicate with users from the general public in a natural and intuitive manner. The group were required to develop a method of gathering user input and interpreting the users' desires. They also were responsible for enabling the robot to express emotion and communicate its requirements to the users. Unfortunately this team encountered difficulty working together effectively as a team and as a result, did not achieve all the desired outputs.

The HRI group proposed a system which utilised an expressive face and a voice synthesis algorithm for robot to human communication which was to be coupled with touch screen display for human to robot communication. However, they chose only to focus their development effort on the communication channels from the robot to the human user leaving the implementation of the touch screen display as future work.

The purpose of the expressive face was to allow the robot to convey emotions naturally to the human users. One application of this ability is to display an increasing level of annoyance if people are preventing the robot from carrying out its duties. The face also imparted 'human' qualities to the robot which encouraged members of the public to interact with SANDRA. While this effect was not expected by the students working on this section, the implementation of the face produced a noticeable difference in the way that the public interacted with SANDRA. Prior to implementing the face, most people were very apprehensive regarding the robot and many would shy away if approached. In contrast, people were much more comfortable interacting with the robot when it was equipped with the face. It was also more successful in attracting attention. In light of these effects, the HRI group recommended that SANDRA be given as many human characteristics as possible.

It was found that a surprisingly simple representation of a human face was effective in encouraging interaction from members of the public. The face was consisted of two spherical webcams simulating eyes and two rows of Light Emitting Diodes (LEDs) which allowed either a smile or frown to be displayed. These components were mounted on a flat sheet of transparent acrylic. The face assembly was mounted on the robot using a simple aluminium frame as shown in Figure 5.

Another component the HRI group determined necessary to communicate efficiently with human users was speech capability. This would allow at least half of the conversation to be conducted through the verbal medium, hence reducing the amount of information that the user is required to read from the display. The group developed a Mac based voice synthesis system which was capable of saying set phrases, for example 'welcome to UTS. Would you like assistance?" Unfortunately due to time constraints this system was never implemented on the robot. The implementation was also hampered by the fact that all other development had occurred on PCs and communication between Mac and PC proved to be too difficult to implement in the time available.

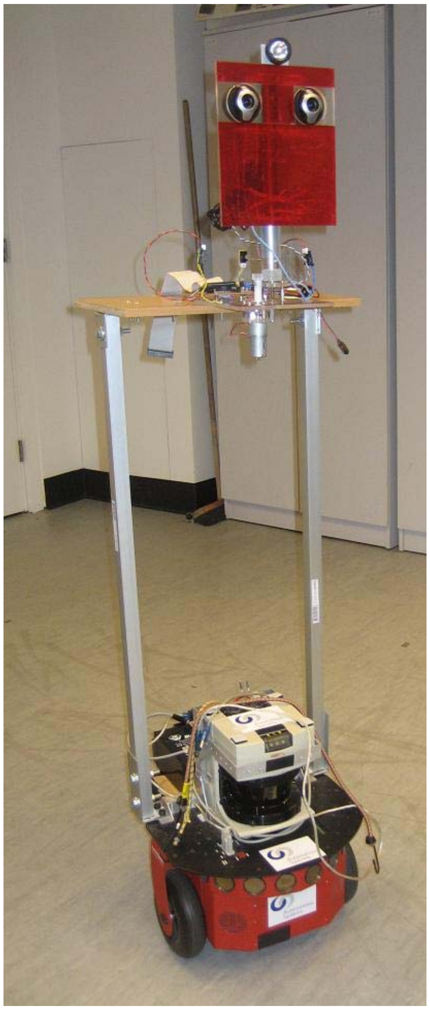

Figure 5: First Prototype

Objectives Achieved

The HRI team was able to develop a basic system for robot to human interaction. However, due to time constraints, it was not integrated with either the LAM or PPOA groups work.

Areas for Improvement

Initially the face was positioned towards the front of the robot and approximately 2 meters above the ground. This caused the robot to become very unstable as the centre of gravity was too far forward and too high. As a result the robot would tend to fall forwards whenever it slowed down or stopped. These effects were reduced shortly after the end of semester by moving the face towards the rear of the robot and reducing the overall height to approximately 1.5 meters. In addition to making the system more stable the height reduction also made the robot less intimidating.

Future work on this section should include the development and implementation of an effective system for gathering user input. This should be accompanied by a well designed and intuitive user interface. The speech synthesis algorithm should also be integrated with the remainder of the system. Finally, the system should be given a humanoid appearance to encourage members of the public to interact with the system.

3 Second Prototype

Following the completion of the Advanced Robotics subject and the subsequent completion of the first prototype a number of the original group members agreed to continue development of the project during the summer break and the autumn semester of 2008. The students agreed to undertake the work as it presented an enjoyable opportunity to gain further experience in robotic system development. Since there was no course credit associated with this work and there was significantly more time available, the students were able to focus on the aspects of the hardware development which had been impractical to complete while undertaking the AR subject. The objectives of this work where to produce a system suitable for demonstration at the mid-year open day in 2008. It was also intended for the platform to be available to future AR students in a stable form to allow them to develop additional features as a component of their course work.

Most improvements implemented in the second prototype were related to improving the HRI capability of the system. This was achieved through a combination of designing and fabricating a humanoid frame and housing, designing and fabricating a new expressive face and sourcing and assembling the hardware necessary to implement a touch screen based Graphic User Interface (GUI). Some time was also devoted to replacing the path planning and obstacle avoidance algorithms developed for the first prototype with the appropriate Player drivers.

3.1 Hardware

The majority of the hardware development that occurred during the construction of the second prototype was related to HRI.

A new light weight but rigid aluminium frame was designed using SolidWorks and then fabricated in the university workshops. This frame replaced the rough frame developed for the first prototype allowing the head to be positioned at approximately the same height. It also included rigid mounting points for the primary laser range finder and mounting points for a secondary optional rear facing laser and was designed with built in shoulders to give the robot a humanoid form. An area of thin aluminium sheeting was included to allow circuit boards and other electronic components to be mounted neatly.

The whole frame was encased in a plastic shell to further enhance the humanoid shape and to prevent any undesirable tampering with internal components while in operation. The shell also provided a mounting point for the cosmetic arms and smoothly connected the outside of the top plate of the pioneer with the perimeter of the shoulder plate on the frame. The shell was fabricated in sections by vacuum moulding. Each mould was roughly cast in plaster of Paris which was then formed to the desired shape. Once the moulds were created they were brought into the university and the plastic panels were fabricated by vacuum moulding. Following this the panels were connected and reinforced with strips of aluminium, painted and assembled on the Robot. Both the assembly and the casting/shaping processes were very labour intensive.

Another aspect of the hardware which was upgraded was the expressive face. Drawing inspiration from the Minerva autonomous tour guide robot [Thrun et al., 1999], it was decided that the robot's head needed to be upgraded from the very simple design which was built for AR. The new head incorporated space for two Logitech Sphere webcams which would act as eyes, an 8x40 bi-colour led array positioned to be the mouth and servo driven eyebrows. The shell of the head was developed using SolidWorks. It was made in two halves by a rapid prototyping machine situated in the design department. The use of the rapid prototyping technique allowed the shape to be quite complex, however, as the material was usually only used for making architectural models it proved to be quite weak and brittle. Fortunately any breaks can be repaired to better than their original strength with super glue. The head assembly was mounted on a deep grove ball bearing which acted as a neck allowing a servomotor to turn the head from side to side.

The final hardware upgrade performed for the second prototype was to replace the onboard computer system. The computer included with the pioneer 2DX was too old to serve any useful purpose. It was replaced with another, more powerful system with the same form factor allowing it to be replaced with very little modification. The upgraded computer has a 3GHz processor and 2GB of RAM which will allow image processing to be carried out on board in the future. The system uses an 8.4” touch screen display for user input and output and a small set of stereo amplified speakers for audio output. The computer is powered from the pioneers onboard batteries using a power supply designed for use in vehicles. It allows a single switch start-up and shutdown with no other user input required. It also prevents deep discharge of the batteries by disabling the standby current to the PC after being switched off for 2 hours.

3.2 Software

The software upgrade for the second prototype included changing the Linux operating system from Kubuntu to Ubuntu. This was change was carried out because Ubuntu provided a simple user setting required to shut down the computer gracefully when the power button was pressed where Kubuntu did not. In other respects the two distributions are very similar so all the software that was developed for the first prototype could be re-used on the second prototype.

The same version of Player was re-installed on the PC even though a newer version had become available. This decision was made to ensure that the software from the previous prototype could be reused and the time spent dealing with compatibility issues could be minimised. New drivers were incorporated into the player system to allow the control of the pan-tilt units in the webcam eyes and to do speech synthesis. The path planning algorithm from the first prototype was replaced with the player driver wavefront planner and the obstacle avoidance driver ND was used. These are both plug-in software drivers developed in player, so they were able to be implemented with little effort and proved to be robust.

The control software for the head was implemented on an Atmel ATmega128 micro-controller. It controls the position of the head, eyebrows and the expression on the led array mouth. A separate power supply was designed for the head as the led mouth draws a large amount of current. The mouth circuitry was later re-developed to to reduce its power consumption and increase the robot’s battery life.

The microcontroller accepts commands in the form of character strings from the control system over a RS232 serial link. It interprets the received strings and generates the appropriate actions such as smile, frown, move neck or move eyebrows Routines were also developed for calibration of the head to allow the head to be zeroed facing forwards. All the calibration of the head is done from the onboard PC. Figure 6 shows a photo of the final robot used in a recent demonstration.

Figure 6: SANDRA demonstration at UTS Open Day

3.3 Events/testing

To date, SANDRA has been demonstrated at two public events and has been used by two groups of 2008 Advanced Robotics students while completing their course work.

The first event was the annual Asquith Girls High School (AGHS) Science Fair at which Michael Behrens was one of the keynote speakers. At this event SANDRA attracted a lot of attention and was able to demonstrate basic mobility functions and expressions. The second event was the UTS open day in August 2008 which occurred a few weeks later. A number of students from AGHS who attended the open day commented that seeing SANDRA was one of the main reasons they had come after seeing it at the science fair or hearing about it from their friends. Although the GUI was still under development at the date of the open day, preventing the robot from conducting completely autonomous tours, it was still able to demonstrate its facial expressions and exhibit its sensing capabilities.

4 Discussion and Future Work

While the majority of this project involved integrating existing well known algorithms, software and hardware, it provided an excellent opportunity to enhance the mechatronic engineering skills of the students involved. The majority of the students who contributed to this project as a component of their course work, including both the current AR students and the authors of this paper, have indicated that they enjoyed working on this project immensely and that much of this can be attributed to the fact that the project included a hardware implementation rather than simply theory and simulation.

At the completion of the second prototype a number of areas for further development were identified. These included improving the operation time, implementing an intuitive GUI, installing a wireless emergency stop, enabling the system to operate in crowded environment and programming the webcams to track the users face.

Work towards enabling the system to operate in crowded environments and the development of the emergency stop system is currently being undertaken by a group of AR students from the spring semester of 2008. The preliminary results of their work show a marked improvement in performance when operating in a crowded environment. During a recent test run, SANDRA was able to navigate through a dense crowd which was not anticipating an interaction with a robot. Another group of AR students have implemented a face tracking algorithm using the 'eye' webcams. Preliminary results are also promising for this project with the system successfully able to track the face of the closest user. This gives the illusion that the system is watching you and assists in creating an engaging user experience.

The system is currently able to operate for approximately 45 minutes on batteries. This is acceptable for short demonstrations but it requires the batteries to be regularly swapped to achieve a whole day of operation. There would be benefits to be gained by performing a power analysis of the system and identifying ways of optimising the power draw. It would also be worth while investigating the possibility of replacing the heavy and relatively low capacity Sealed Lead Acid (SLA) batteries with a new set of batteries with a higher power density.

A final area which should be developed is the ability to autonomously navigate and map an area. This would form a suitable project for future Advanced Robotics students. This capability will make it easier to deploy the system in different areas. It will also allow the map of the UTS foyer to be updated the night before an open day to include the locations of displays and stalls.

5 Conclusions

This paper presented the development of the first and second prototypes for an autonomous robotic tour guide to be used at the University of Technology, Sydney. It presents the methodology of the incremental system design and a discussion of the students' learning experiences.

The students involved in the development of this system found the project both challenging and interesting and enjoyed the opportunity to put their knowledge into practice. The second prototype has proved to be a capable platform with a significantly enhanced scope for implementing effective human robot interaction. As such, aspects of the project will be presented to the next group of Advanced Robotics students to allow them to continue development. Thus, over time and step by step, the dream will be realised.

Acknowledgements

The authors would like to acknowledge the work of Jordan Nguyen, Shane Williams, Minh Lee, Matt McQuinty, Hani Assad and Jamil Helweh in developing the first prototype. We would also like to thank Shoudong Huang and Terak Taha from the ARC Centre of Excellence in Autonomous Systems for their guidance and support and Richard Moore, John Dennison and Wojciech Wawrzyniak from the UTS workshops and design departments for their assistance in fabricating components for the second prototype.

References

[Behrens et al, 2007] M. Behrens, M. Carmichael, M. Patel. *Welcome Assistant Robot for UTS Part (a), Laser/Sonar based Localization and Mapping.* Faculty of Engineering, University of Technology Sydney, Sydney, 2007.

[Burgard et al., 1998] W. Burgard, A.B. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, S. Thrun. *The Interactive Museum Tour-Guide Robot.* Proc. of the 15th National Conference on Artificial Intelligence (AAAI-98), Madison, Wisconsin, 1998.

[Dissanayake et. al., 2006] Dissanayake, Gamini; Paxman, Jonathan; Miro, Jaime Valls; Thane, Oliver; Thi, Hue-Tuan, "Robotics for Urban Search and Rescue," *Industrial and Information Systems, First International Conference on* , vol., no., pp.294-298, 8-11 Aug. 2006

[Goto et. al., 2003] Goto, T.; Kosaka, T.; Noborio, H., "On the heuristics of A* or A algorithm in ITS and robot path-planning," *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on* , vol.2, no., pp. 1159-1166 vol.2, 27-31 Oct. 2003

[McQuilty et al., 2007] M. McQuilty, H. Assad, J. Helweh. *Greeter-bot.* Faculty of Engineering, University of Technology Sydney, Sydney, 2007.

[Nguyen et al., 2007] J. Nguyen, S. Williams, M. Lee. *Welcome Assistant Robot for UTS: Obstacle Avoidance and Path Planning.* Faculty of Engineering, University of Technology Sydney, Sydney, 2007.

[Siegwart et al., 2003] R. Siegwart, K.O. Arras, B. Jensen, R. Philippsen, N. Tomatis. *Design, Implementation and Exploitation of a New Fully Autonomous Tour Guide Robot.* In Proc. of the 1st International Workshop on Advances in Service Robotics (ASER 2003).

[Thrun et al., 1999] S. Thrun, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, J. Schulte1, D. Schulz. *MINERVA: A Second-Generation Museum Tour-Guide Robot.* Proc. of the International Conference on Robotics and Automation (ICRA '99).