

“© 2006 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Robotics for Urban Search and Rescue

Gamini Dissanayake, Jonathan Paxman, Jaime Valls-Miro, Oliver Thane and Hue-Tuan Thi
Centre for Autonomous Systems, University of Technology, Sydney, Australia
{gdissa,jpaxman,jaime.vallsmiro}@eng.uts.edu.au

Abstract – This paper describes a team of robots that are designed for urban search and rescue applications. The team CASualty consists of four tele-operated robots and one autonomous robot. A brief description of the capabilities of the robot team is presented together with the details of capabilities of the autonomous robot HOMER. In particular, the software architecture, user interface, strategies used for mapping, exploration and the identification of human victims present in the environment are described. The team participated in an international competition on urban search and rescue (RoboCup Rescue) held in Bremen, Germany in June 2006 where HOMER was placed second in the autonomy challenge.

I INTRODUCTION

In an urban search and rescue scenario, detecting the locations of survivors and then recovering them from a collapsed building is one of the biggest challenges faced by emergency response personnel. The environment can be unstable and difficult to negotiate while survivors trapped need to be rescued within a short time frame. Use of a robot or a team of robots to assist human rescuers in such situations is one of the areas where robotics research can be of great benefit to humanity. While significant progress has been made and a variety of robots have been used during many recent disasters, much research is needed to achieve the objective of deploying a team of autonomous robots for urban search and rescue.

The RoboCup Rescue competition is a forum where the latest research on urban search and rescue can be evaluated in a competitive environment. The rescue arenas constructed to host these competitions are based on a template developed by the U.S. National Institute of Standards and Technology (NIST) [1]. The arenas simulate a collapsed building with three distinct grades of complexity. A maze of walls, doors, and elevated floors provide various tests for robot navigation and mapping capabilities. Variable flooring, overturned furniture, and problematic rubble provide obvious physical obstacles. Intuitive operator interfaces and robust sensor fusion algorithms are essential in order to reliably negotiate the arenas and locate victims. Arena used in the 2005 competition in Osaka, Japan is shown in Fig. 1.

RoboCup Rescue competition aims to promote research so that [1]

“When disaster happens, minimize risk to search and rescue personnel, while increasing victim survival rates, by fielding teams of collaborative robots which can:

- *Autonomously negotiate compromised and collapsed structures*
- *Find victims and ascertain their conditions*

- *Produce practical maps of their locations*
 - *Deliver sustenance and communications*
 - *Identify hazards*
 - *Emplace sensors (acoustic, thermal, hazmat, seismic, etc,...)*
 - *Provide structural shoring*
- ...allowing human rescuers to quickly locate and extract victims.”*



Fig. 1 The arena used in the 2005 RoboCup competition held in Osaka

This paper is organised as follows. Section II describes the robot team deployed by the Centre for Autonomous Systems (CAS) in the June 2006 RoboCup Rescue competition held in Bremen, Germany. Sections III –VI describe various aspects of the autonomous platform HOMER and its capabilities. Discussion and conclusions are presented in section VII.

II ROBOT PLATFORMS

A Caster



Fig. 2 CASTER II

CASTER-II (seen in Fig. 2), is based on the Yujin Robotics DT-3 locomotion platform. CASTER-II uses a tank-like track arrangement that is steered differentially using tracks. The rubber tracks are ribbed to improve traction over rough terrain. However, unlike most tanks, CASTER-II consists of two articulated sections. The rear section is similar to a conventional tank. The front section consists of a triangular track path. This allows CASTER-II to traverse tall obstacles such as stairs.

The sensor package on this robot consists of a number of cameras, including a thermal camera (FLIR Thermovision A10), for situational awareness and victim identification. Three dimensional mapping capability is provided through a time-of-flight range imager (CSEM Swiss Ranger 2). A robot arm is used to carry some of the sensors so that they can be easily manoeuvred.

B Redback

Redback is a lightweight, low cost, advanced mobility platform based on the Tarantula RC vehicle is shown in Fig. 3. These robots complement the abilities of CASTER. Redback is fitted with sensors and a computer that allows it to be driven remotely, generate 3D maps of its environment and perform semi-autonomous actions, allowing them to be used with minimal need for operator attention.



Fig. 3 Three Redbacks in various stages of development


C HOMER

HOMER (Fig. 4) is a lightweight differential drive robot designed to autonomously operate in relatively flat terrain. The 2006 competition included pitch and roll ramps of up to 10° elevation. The principal mapping sensor is a Hokuyo URG-04LX laser range finder, and obstacle avoidance is achieved with the laser and additionally two built-in sonar range sensors. Victim identification is achieved via a Thermoteknix Miricle thermal infrared camera, and two optical cameras – a Point Grey Flea, and a Point Grey Dragonfly2. The Flea is co-mounted with the thermal camera on a pan-tilt unit constructed from two Megarobotics AI-1001 modules, while the Dragonfly2 is attached to a Vstone hyperboloidal mirror for omnidirectional vision.


All computation for mapping, localization, exploration, navigation, obstacle avoidance and victim identification is performed on-board the robot with a Toshiba Libretto U100 computer. The Libretto has a very small form factor, measuring 210x160x30mm, and weighing less than a

kilogram. It contains a 1.2 GHz Pentium M processor, with 1024MB RAM and on-board USB and IEEE1394 ports for device interfacing.

HOMER's front wheels are differentially driven by two Maxon RE-35 motors, and a custom built motor controller board. The rear wheels are Kornylak omni wheels. These provide excellent omnidirectional motion without the usual drawbacks of caster wheels – particularly when changing direction. This main focus of this paper is on the development of the HOMER.



(b) Front view: Two sonar sensors as well as the optical (right) and thermal (left) cameras mounted on a pan-tilt unit are visible.



(b) Side view: The omni-directional back wheels and a panoramic camera (not used during the competition) visible.

Fig. 4 Two views of the HOMER

III MAPPING

The ability to generate a map, while navigating in an unknown environment is one of the most essential capabilities of an autonomous robot. The principal mapping sensor in the HOMER is the Hokuyo URG-04LX laser rangefinder. This is a new low cost, low power and lightweight range sensor. It produces range scans at 10Hz with an angular field of view of 240 degrees, and a maximum range of 4m. This sensor is ideally suited to mapping in an indoor, tightly constrained environment.

Maps are produced using a new algorithm "Map Referenced Iterative Closest Point" (MRICP). This method incorporates Iterative Closest Point (ICP) [2], an established general purpose algorithm for matching clouds of points in two or three-dimensions according to a least squares criterion. In MRICP, a subset of an existing occupancy grid map is converted to a point cloud and matched with the most recent laser scan in order to produce an accurate

estimate of current position. The current laser scan is then used to update the occupancy grid based on Bayes rule using a model of the laser range sensor. This method generally produces better maps than standard sequential ICP scan matching, and in particular it obviates the need for explicit loop-closures over short distances (up to 10-20m). As larger loops are included, this method needs to be extended by detecting loop closures explicitly, or by a hierarchical method which integrates small map patches into a larger map.

The mapping strategy was initially designed with a totally flat environment in mind. At the 2006 Rescue competition, pitch and roll ramps of up to 10° elevation were introduced for the first time. Ramps introduce particular difficulties for robots which exclusively use horizontal laser scans for mapping. Firstly, if the laser is mounted too low, the top of a ramp may be detected as an obstacle and incorporated into the map even though it is in fact traversable. Secondly, laser scans taken from an inclined position may not accurately match with scans taken from a fully horizontal position. The first issue was addressed by raising the laser position sufficiently so that ramps do not impede the laser view. The MRICP method worked robustly enough that the second issue did not impact the laser alignment substantially, and the map produced was still quite navigable. Fig. 5 shows the map produced during the competition run including pitch and roll ramps. Mismatches due to the ramps are visible as multiple lines corresponding to a single wall. Fig. 6 shows a more accurate map generated in a larger flat area during laboratory testing. There are plans to include a tilt sensor based on a two-axis accelerometer in the sensor package to deal with some of the issues associated with the presence of the ramps.



Fig. 5 Occupancy grid map generated by HOMER during the competition in Bremen, Germany. The shade of grey is proportional to the probability of occupancy (white = 1, black = 0). The area mapped is approximately 7x5m.



Fig. 6 Larger occupancy grid map generated by during laboratory testing. The area mapped is approximately 12x8m.

IV NAVIGATION

In the search and rescue scenario, the autonomous robot needs a high level planner that makes certain that all unknown areas are explored. Furthermore, the robot also needs a reactive strategy for avoiding nearby obstacles as path planning for exploration is typically done at a coarser resolution to minimize the computational resources required. This section describes the algorithms used for these two tasks.

A Local Navigation

Local navigation is achieved by applying the well known Vector Field Histogram (VFH) method [3]. The VFH method uses a two-dimensional Cartesian histogram grid as a world model. The world model is updated with data about local obstacles - in this case, data from the URG laser and two onboard sonars are fused to produce local proximity estimates. VFH provides a robust local navigation method which reliably avoids collisions, but has limitations in very tightly constrained spaces. A modified version of VFH algorithm is implemented on the HOMER in order to improve performance in cases where the goal point is behind the robot, or out of the angular range of the range sensors.

B Exploration

Exploration requires the robot to be driven towards unexplored regions of the environment. Using the current occupancy grid provided by the mapping module, the boundaries between explored and unexplored regions, known as frontiers, are first extracted. A global planner based on the wavefront expansion method is used to generate a set of way points the robot must follow to reach the nearest frontier. The way point generation also attempts to make the robot stay in the middle of open spaces and minimize turning if at all possible. The global path is recalculated at regular intervals as the occupancy grid gets updated. At all times, the robot is in the look-out for victims in the vicinity. As soon as a victim is detected, a new path

to the victim is recalculated and the platform driven closer to the potential victim for its identification. Once that is accomplished, the autonomy module drives the robot back in search of places and victims in the rescue arena not yet explored.

V VICTIM IDENTIFICATION

The victim identification strategy has two components; firstly the sensor data is processed to produce signs of life; secondly these signs of life are mapped to an occupancy grid (where the occupancy represents the probability of a victim in that cell) to localize different victims with respect to the map generated by MRICP.

Four signs of life expected in the simulated victims placed in the competition arena were examined during the development of the HOMER victim identification strategy.

A Body Temperature

Heat is one of the most reliable means of detecting the presence of human life, due to the very specific temperature range of human skin. The most reliable use of heat detection is achieved with a thermal camera such as the Thermoteknix Miricle camera mounted on HOMER. A camera enables thermal readings to be filtered and clustered so that only regions of appropriate size and shape are considered. An elevation filter is also applied in order to minimize false-positive readings from lights or spectactors.

B Alternative Signs of Life

Skin colour can be extracted by considering histograms of skin and non skin images from the optical camera. Skin detection using Gaussian Mixture Models is, in fact, quite reliable and robust to variation in skin colour and texture. However skin detection is much more reliable for actual skin rather than the painted dummies used in Robocup Rescue.

Movement is another sign of life which can be found by thresholding differential images taken from a stationary camera, or by optical flow techniques.

A directional microphone can be used to detect the sound of a person talking or baby crying. By recording the sound for a short period and filtering the signal over a narrow vocal range containing human formants, the presence of a voice may be recognized.

C Victim Map

By combining the sensor model of a victim detection sensor with the occupancy grid map produced by MRICP, a 'victim map' containing spatial probabilities of the presence of a victim is generated. When enough positive sensor readings are found in a particular location, a victim is registered in the map, and the navigation component is instructed to move to a closer proximity. When the robot achieves close proximity to a victim a 'snapshot' of all of the sensors is sent back to the user interface and the operator is alerted. The robot then waits for the operator's instruction to resume the search.

In the Bremen competition, the thermal camera was the principal victim sensor used. Skin detection was unreliable,

as very little skin was visible on most victims and because skin was painted rather than genuine human skin (for obvious reasons).

Fig. 7 and 8 show the visual and thermal images for the two victims HOMER detected during the it's successful final competition run.




Fig. 7 Visual (left) and thermal (right) images of the first victim. Note that the two cameras are not aligned. The victim is seen as a bright region in the top right corner of the thermal image




Fig. 8 Visual (left) and thermal (right) images of the second victim (a heated baby doll).

VI SOFTWARE ARCHITECTURE

As the CASualty team consists of five robots, a distributed software architecture is essential so that each robot can collect data from the local sensors and make the required decisions based on the level of autonomy of the platform. The raw/processed data is sent back to the remote central operators unit, where a number of tasks are performed automatically, such as managing and manipulating the maps provided by the platforms, placing victims and landmarks in the global map.

Given the large amount of simultaneous data available to the operator, as well as the need to effectively control the remotely-operated and semi-autonomous platforms at the same time, a great deal of effort has been dedicated to present the user with an integrated user-friendly GUI [4]. The left half of the interface (see Fig. 9) contains collective information: maps in one tab, and victim snapshots in another. The right half contains robot specific information: There is one tab for each robot, including the main camera view, an unfolded omni-camera view and a thermal overlay. There is also a variety of robot-specific status and control available in each tab.

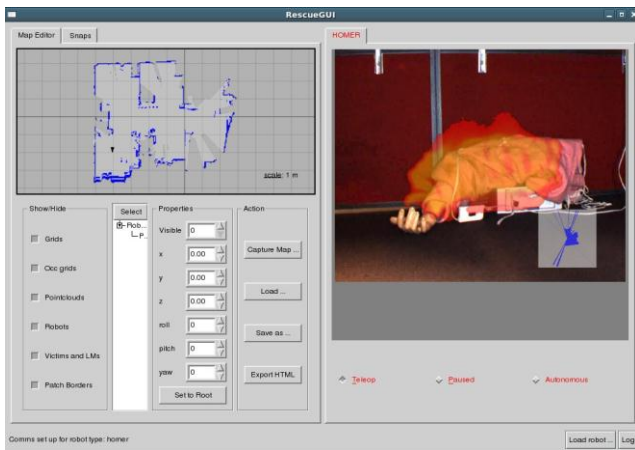


Fig. 9 Screen shot from RescueGUI. The current global map can be viewed on the left, while the current robot (HOMER) view can be seen on the right, including camera view, thermal view (yellow/red overlay) and laser view (blue/grey inset).

The interface has also the ability to accept or reject victims found by the autonomous platform, and to record logs of the various sensor data, including victim/landmark snapshots. Distributed middleware software architecture, ORCA [5], has been employed to connect the various components that make up the system. This is an open-source (orca-robotics.sourceforge.net) framework for developing component-based robotic systems, which takes advantage of the flexibility of the ICE (www.zeroc.com) middleware package for inter-component communications.

VII DISCUSSION AND CONCLUSIONS

The Robocup Rescue competition is an excellent testing ground for robotic platforms and algorithms. The competition enables practical research outcomes in autonomous mapping, navigation and decision making to be tested against other research groups from around the world.

In the Bremen 2006 competition, CASulaty achieved second place in the Autonomous Robot event, and the semi-finals of the multi-robot mixed initiative event. In the final run of the autonomous event, HOMER successfully navigated through approximately eight metres of arena over two ramps, and correctly identified two victims along the way.

It is important however, not to lose sight of the real-world application towards which this research effort is directed. The objective is for a team of robots to be able to be

deployed in genuine disaster scenarios in order to survey the area, detect the presence or otherwise of human victims and report this data back to an incident commander.

The development of autonomous solutions enables a single operator to be able to take charge of multiple robots in a coherent way.

An important next step is for some of the autonomous capabilities of flat-terrain robots to be transferred to more mobile platforms. This will be one of the objectives of the CASualty team for next year.

ACKNOWLEDGMENTS

The authors wish to acknowledge the work of Waleed Kadous and Raymond Sheh from University of New South Wales. Although they were mainly concerned with the teleoperated robots, CASTER and Redback, they led the software design effort and developed much of the user interfaces.

The authors also wish to acknowledge the contributions of the following. Tarek Taha, Cindy Leung, Nawid Jamali, Rudino Salleh and John Zaitseff participated in the development of the CASualty team. HOMER was designed and built by Nathan Wilson, Matthew Rozyn and Michael Andrews. HOMER hardware/software was gradually refined by Ray Clout, Weizhen Zhou, Jaime Valls Miro and Oliver Thane. The team in Bremen consisted of Claude Sammut, Waleed Kadous and Raymond Sheh from University of New South Wales, and Jonathan Paxman and Oliver Thane from University of Technology, Sydney.

This work is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales Government

REFERENCES

- [1] <http://www.isd.mel.nist.gov/projects/USAR/competitions.htm> (accessed on 30 June 2006)
- [2] P. Besl and N. McKay, "A method for Registration of 3--D Shapes", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol 14, No. 2, pp 239 -- 256, 1992.
- [3] J. Borenstein, Y. Koren, "The vector field histogram: fast obstacle avoidance for mobile robots". *IEEE Trans Robotics and Automation* Vol 7, No 3, pp. 278 -- 288, 1991
- [4] M. Kadous, R. K. Sheh and C. Sammut, "Effective user interface design for rescue robotics", *Proc. of Human Robots Interaction Conf.*, 2006.
- [5] A. Brooks, T. Kaupp, A. Makarenko, S. Williams and A. Oreback, "Towards component-based robotics", *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pp. 163 -- 168, 2005.