Robotic Search-and-Rescue

An integrated approach

Julian de Hoog[1], Stephen Cameron[1], and Arnoud Visser[2]

[1] Oxford University Computing Laboratory, Parks Road, Oxford OX1 3QD, UK
[2] Universiteit van Amsterdam, 1098 SJ Amsterdam, the Netherlands

1. Introduction

Recent advances in robotics and computer science mean that it is now possible to use teams of robots for many real-world applications. One very important application is robotic search-and-rescue. Robots are highly suitable for search-and-rescue tasks since they may be deployed in dangerous and toxic environments without putting human responders at risk.

Initial uses of rescue robots, such as during rescue efforts following the 2001 collapse of the World Trade Centre in New York, have highlighted the fact that many improvements are still required if teams of robots are to provide extensive assistance in search-and-rescue scenarios. While much effort is going into development of more robust and mobile robot platforms, it is also necessary to develop advanced methods for multi-robot coordination and practical user interfaces for humans to control the robots.

Here we describe a unified approach developed jointly by the Oxford University Computing Laboratory and the University of Amsterdam that integrates advanced techniques from a variety of fields such as mapping, localization, exploration, communication, navigation and human-robot interface design. All research reported here is performed using the USARSim simulator[3].

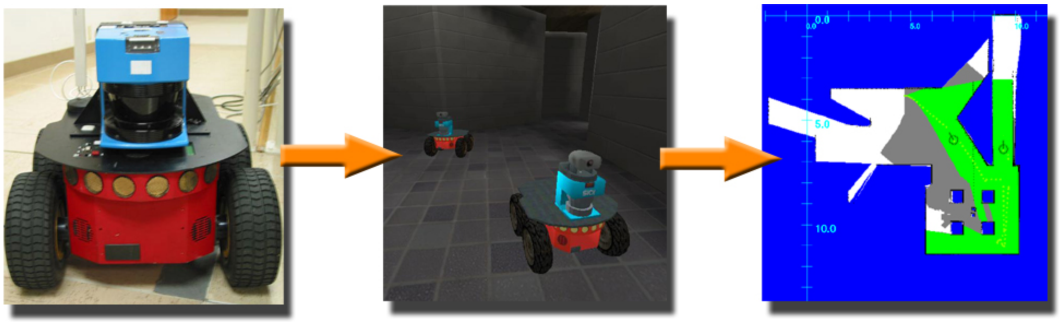

Fig. 1. The simulation process: a real P2AT robot is modeled in the simulator. Teams of simulated P2AT's are then used to test exploration algorithms.

---

[3] Available for download at http://sourceforge.net/projects/usarsim. Based on a commercial game engine, this simulator uses state of the art techniques for simulating physics and rendering graphics. It is fully configurable, allowing for independent development of simulated robot models, sensors and environments.

2. The Integrated Approach

Mapping and Localization. Our mapping is based on a manifold approach that uses scan-matching and has an underlying graph structure. An initial position estimate is determined using the inertial navigation sensor, but subsequent mapping and localization efforts use only measurements from the laser range scanner. Current laser scans are compared with recent laser scans. Once the differences exceed a particular threshold, a new node is created in the graph along with a new link representing displacement. The graph structure may easily be converted into an occupancy grid with standard rendering techniques.

Exploration and Communication. We use a frontier-based exploration approach. As a robot explores, it simultaneously maintains two occupancy grid based maps. The first is based on the maximum sensing range r_{max} of the laser range scanner and the second is based on a more conservative safety distance r_{safe}. "Frontiers" are the boundaries between the safe space and the open space. Robots can then choose which frontier is most suitable for exploration by evaluating a combination of distance to that frontier (i.e. the *movement cost*), potential area to be explored beyond that frontier (i.e. the *information gain*), and likelihood of being able to communicate with the human operator from the frontier (i.e. the *communication likelihood*). This latter value is estimated by maintaining a table of distance - signal strength value pairs and extrapolating to the location of interest.

Navigation. Once an optimal assignment of robots to frontiers has been determined, a path may be planned for each robot to its assigned frontier. This is performed by convoluting the obstacles in the occupancy grid with the shape of the robot using a Gaussian convolution kernel, and using breadth-first search to determine whether a path exists. Way-points along the calculated path are used to guide the robot to its goal.

Human-Robot Interface Design. Control of even a single robot can require extensive operator resources, so proper interface design is crucial to a successful robotic search-and-rescue effort. In our approach the human operator may dynamically assign individual robots a variety of behaviours, from complete teleoperation to complete autonomous exploration. As a result attention may be paid to areas or situations that require it most.

3. Future research

The probability of communication success is built into the exploration algorithm. However, in certain situations significant areas of interest may be out of range of the human operator so communication remains a major issue. To solve this we hope to implement a *territorial exploration* approach: while some robots explore the far reaches of the environment, others are dynamically assigned particular territories within which they behave as information relays. Multi-hop communication has been applied with success elsewhere, but generally not to more than four robots at a time. We hope to experiment with greater numbers of robots.

Simulation research must proceed in step with hardware development and engineering research. As more robust and mobile robot sensors and models become available, control interfaces for these must be developed and incorporated into the already existing framework. In the long term, simulation research reported here is only of use if it is similarly successful in the real world. Applying our integrated approach to a team of real robots would provide an excellent opportunity for validation and evaluation.