Improvements of a Heterogeneous Rescue Team

Fares Alnajar
5902037
f.alnajar@student.uva.nl

Hanne Nijhuis
0364568
h.nijhuis@student.uva.nl

January 29, 2009

Abstract

In this paper we describe a new approach to make use of a heterogeneous robot team for the RoboCup Urban Search and Rescue (USAR) competition. We will show improvements that could be achieved by using both a flying and a driving robot as opposed to only two robots operating from the ground. In most of the previous work, only homogeneous teams where used to for the task of exploring the area, finding the victims and reporting their position and status (the most important goals of the USAR competition). Having a heterogeneous team increases the possibilities in defining new strategies, as each of the robots has a different set of features (in terms of sensors and actors) and therefor a different view of the environment. The robots also have their own disadvantages, but they should work together to try to overcome these. By integrating the abilities of the two robots we can gain added value to the teams performance.

1 Introduction

The goal of our research is to show the improvements that could be made when deploying a team of robots in a disaster-area to search for victims. We will show the advantages of having a heterogeneous team in comparison to a team with robots of only one type. More specifically we will use a team consisting of an aerial and a ground-vehicle and compare this to a team of two ground-vehicles. In this section we will elaborate on the relevance of this research to the real world and to the RoboCup Rescue competition to which it should be directly applicable. More details on this competition will be provided to give a better insight in the goals to be achieved and the situations to be considered.

1.1 Relevance

In situations where a disaster like an earthquake has occurred, searching for survivors in the area could be dangerous due to (partly) collapsed buildings that are unstable. It could also be difficult for humans to search in such a collapsed building if the available room is to small to crawl through. In such situations robots could be deployed to search the area and hopefully supply some useful information on the location and status of possible survivors. These robots could be operated by humans (by remote control), but if they are able to explore (semi) autonomously one could deploy a whole team of robots simultaneously to cover a bigger area. To investigate in the possibilities of using robots in these cases and to experiment with their behaviors, a new competition was added to the well known worldwide RoboCup competition as described in the next subsection.

1.2 Rescue League

RoboCup is an international robotics competition founded in 1993. The aim is to develop autonomous robots with the intention of promoting research and education in the field of artificial intelligence. The competition started out with the focus on building a team of robots which are able to play soccer and even defeat the humans in playing this game by 2050. In 2001 a new competition was added to this world-wide annual robotics event, focusing on the deployment of robots in disaster areas. This was partly by demand of the community, with the disastrous earthquake of 1995 in Japan as a motivator. Our research is applicable to this Rescue League competition (which is subdivided into a real and a virtual league with the virtual being split up in agent simulation and virtual robots).

Because deploying actual robots in actual disaster-situations is a very complex task, the *virtual robots* competition is done in a virtual world. This provides the opportunity to simplify the problem by focusing on only a few of the many research-subjects involved. In a virtual world for example, there is no wind in outdoor situations (unless you want to implement it yourself) and sensors only produce noisy output if you want them to (or 'perfect' output if you prefer). Because of this, the competitors can focus on strategies for behaviors and co-operation without having to deal with technical details on noise etc.

1.3 USARSim

Since 2006 there have been annual world-competitions in a virtual environment called USARSim. It's based on the Unreal Tournament 2 engine[1] and provides the ability to have robots operate in a 3D world with the laws of physics (like gravity) already implemented.

A server to provide this 3D world should be initiated with a specific map. For our research we use the `CompWorldDay1` map which supplies a large outdoor- as well as indoor-environment (an overview of the outdoor area we operate in can be seen in Figure 1(c)). Many obstacles like cars, buildings and construction are present, but the sky is fairly empty. This will give the aerial robot the opportunity to fly around without the need for 'obstacle avoidance'.

Robots are recreated in the virtual world, based on real machines that are used in the non-virtual competition. The Air Robot and the PIONEER 2-AT (P2AT) are depicted in figure 1(a) and 1(b).

(a) The Air Robot

(b) The P2AT ground robot

(c) Outdoor area of `CompWorldDay1`

Figure 1: Images from the USARSim environment

1.4 Performance Metrics

To be able to compare the performance of the participating teams some metrics needed to be developed. The initial 2006 metrics were specified like this:

$$S = \frac{V_{ID} \times 10 + V_{ST} \times 10 + V_{LO} \times 10 + t \times M + E \times 50 - C \times 5 + B}{(1+N)^2}$$

Where V_{ID} is the number of victims identified, V_{ST} is the number of victims for which a status was reported, V_{LO} is the number of properly localized victims, B is an optional amount of bonus points rewarded by a referee for additional information on victims, t is a scaling factor for the accuracy of the map, M is the points assigned by a referee for the quality of the map, E is the points assigned by a referee for the exploration efforts and C the amount of collisions between a robot and a victim.

Over the years these metrics have changed[3] and it will be difficult to take all factors into account while defining an approach for the heterogeneous team. This has to do with the limitations of the aerial robot considering its sensors. Building a map of the world while exploring is very difficult when relying only on a camera (as supposed to having a laser scanner to 'draw' obstacles and free space). Since this map is used in the assessment of the teams performance, but it is unclear how to rate this for the area explored by the flying robot, we chose to leave the points for exploration out in our research. The points for the map are based only on the accuracy of the map of the ground-robot.

1.5 Related work

Experiments with a heterogeneous team have been done before by several competitors of the Rescue League. The 2nd place of the 2008 competition was actually rewarded to a heterogeneous team (one aerial and one ground robot) [7], but these were both operated manually (by one person switching between the robots). We take this to be evidence for being able to show great improvement with a heterogeneous team, because even though there were only two robots, the team still ended up in 2nd place (competing with much larger teams with autonomous ground-robots). Our approach incorporates semi-autonomous behavior as explained in section 3.1. This has also been done before, for example by the Jacobs team (from the Jacobs University in Bremen), but without communication between the robots. We will gain added value by implementing a way of allowing the robots to share their knowledge of the world they operate in. In this 'sharing of knowledge' we are inspired by a principle as introduced by the Freiburg team for fusing local maps of all robots into one global map [5]. Multi-robot teams have shown to have a higher exploration effort, and by communicating this knowledge to each other they are able to make sure new areas are explored as much as possible (known areas are avoided. A periodic synchronization among the team members is a key demand to ensure that the shared information is valuable and there is no huge difference in localization.

1.6 Outline

In the following section we will explain the approach we took on showing improvements by using a heterogeneous team. Next we will cover the methods we used to implement this approach in section 3. We will discuss the experiments and their results in section 4 to round of with a conclusion and discussion (section 5) and the possibilities for future work (section 7).

2 Approaches

We discuss several approaches to make the heterogeneous team operate (semi) autonomously. To confine the enormous amount of things one should take into account with developing these approaches, we focus on outdoor situations only. It will be a lot easier to operate the Air Robot in an open sky and there will be less complex path planning (like in a cubicle office for instance) for the ground robot.

2.1 Exploration

The main advantages of the Air Robot are its speed and its ability to move around in 3 dimensions, making it easier to avoid obstacles (you can fly over them instead of driving around them). This would make the Air Robot ideal for exploring a large area quickly. Furthermore the ability to explore from a high position provides a higher probability of having an unobstructed view on victims.

Unfortunately the pay-load of the robot is very limited, so in real-life situations we can not supply the robot with heavy gear. This limits not only the amount and type of sensors the robot could carry, but also the hardware and with that the computational power available on board. Luckily we don't have to worry about these real-world problems or the rules as they apply for the actual competition, but for this research to be useful we should try to stay as close as possible.

2.2 Victim localization

Since the main goal of the Rescue League is to find victims in a disaster-area we want to use the Air Robot to search for them, because this can be done much faster from the sky than by a slow ground robot. The biggest problem however, is that the Air Robot has fewer options to accurately estimate its own location (due to a smaller sensor suite). In the current competition performance metrics victims should be localized with an accuracy of at least 250 cm to get any points at all. We have tested the accuracy of the Air Robot for its estimated location and found that it would be a problem to rely on that for the victim detection. The Air Robot uses its inertia sensor (acceleration sensor) and therefor has an accumulated error on the estimation of its location.

2.3 Sharing knowledge

Each member of the team (the ground robot and the aerial one) makes an estimation of its position. It is mentioned before that the error made by the aerial robot is greater than the one made by the ground robot. To unify the estimation 'localization' of the two robots, the estimations of a specific point by the two robots are computed and used to find out the shift the aerial robot makes with respect to the ground one.

There are two ways to achieve this 'unification'. In the first one, the ground robot computes this shift and sends the 'correction' to the aerial robot for the next estimations. But this method requires a consistent connection with the aerial robot which is not always available in the disaster situations, and requires the two robots to be close to each other. This can be solved by storing the shift 'correction' at the ground robot, and in every new victim to head to, the ground robot tests if there is still modification to send to the aerial one (not sent due to communication problem). If so, the ground robot modifies the victim position itself, and wait till the communication is back to retransmit the 'correction'.

The first method is useful when the aerial robot makes some kind of mapping and it is important to know its exact position. But as the aerial robot has mainly one task which is reporting the victims, it doesn't matter for it to have this piece of information.

The second method is simpler. It keeps the shift vector at the ground robot only, so every time there is 'unification', the ground robot computes the shift vector and stores it. This is the method we use in our experiments.

2.4 Overview of our Strategy

In our final design we have a heterogeneous team of two robots, ground robot and aerial one. Due to its fast movement and higher degree of freedom, we used the aerial robot for exploring the environment and searching for victims. This robot is tele-operated as it is difficult to behave autonomously (because of possible crashes), and fact that it is navigating in three dimensions instead of two. One drawback of the aerial robot is its poor localization; we can depend on it to provide an approximated location of any victim. The ground robot is slower than the aerial one, but more accurate in localization. The aerial robot sends the approximated locations of the victims to the ground one, which is equipped also with victim sensors. So as the ground robot gets close to the victim, it can detect its exact position. The positions sent by the aerial robot guide the ground robot to the victim places, and then it makes further investigation for the victim in its local area. One case this approach can fail is when the localization error of the aerial robot gets so large, so the victim locations it gives may mislead the ground robot. This is due to the fact that the localization error of the aerial robot is accumulated. To overcome this, a gradual 'correction' of the aerial robot localization is achieved. At some points the locations estimated by the two robots are computed, these estimations are used to figure out how much the aerial robot estimation differ from the ground robot one. When the ground robot takes the next victim location, it changes the location by the current shift vector it has, and heads to this new location. The new victim position changed by the ground robot is not guaranteed to be perfect, but it is more accurate than the initial location sent by the aerial robot. The longer the correction is made, the less accurate the new positions are.

3 Methods used

In this section we explain the following-behavior of the ground robot to navigate the environment for detecting new victims. Also we discuss how the two robots update the estimation shift made by the aerial robot.

3.1 Following

The `follow-behavior` (as used by the ground robot) is composed of more than one strategy or 'motion' depending on the local environment around the robot (obstacles, victims, and/or other teammates). This behavior is built on the motions used for the autonomous explorations of UvARescue [4]. The motion is a set of rules used by the robot to navigate the environment. In each time one motion is active, and the robot behavior switches from one motion to another, depending on the robots situation. Here we give a description of these motions.

Following This is the main motion of the `follow-behavior`, in which the robot first requests the position of the next target (the victim positions). Then, it plans a path from its current position to the target taking into consideration the information it has of the environment so far (for example the obstacles). After the robot drove a specific distance (4 meters) it re-plans the path again, because it gets more information about the environment and a better path could exist. It may seem more accurate to re-plan the path after a shorter distance (like 2 meters), but this re-planning costs the robot additional computation time, so it's better to exploit the current path before going for another planning and additional computation time.

Avoid Teammate This motion is called when there is potential risk of collision between two robots (the distance between them is less than

1 meter). Such cases are subtle and should be dealt carefully. As the other robots are not fixed, and as the robot is trying to avoid its teammate, there is a possibility that that teammate is trying to avoid our robot (if it's not tele-operated), so cooperation among the team members is a key issue so there won't be misunderstanding and collision. One robot should allow the other to continue moving in its way. As there is no clear prioritization of one robot over the rest, any prioritization is fine. In our implementation we used the one in [4], which depends on the robot IDs, so the robot with highest ID gives the way to the other robots. There are four states the robot can have while trying to avoid a teammate:

- The teammate is facing the robot from the front side: in this case the robot should turn right or left to get out of the way.
- The teammate facing the robot from the back side: the same as the previous state.
- The teammate is in front of the robot but not facing it: the robot waits till the teammate move a way
- The teammate is behind the robot but not facing it: the robot keeps its normal behavior (return the control to the 'following' motion).

Avoid Victim This motion is used when the robot gets closer to the victim less than 1 meter. After detecting the closest part of the victim, the motion keeps the robot away of the victim. When the robot is far (more than 2 meters), the control is returned to the following motion.

As we can see the flowing of control from one motion and another depends on the situation of the robot, this is depicted in figure 2.

3.2 Updating location estimation shift

We still need to know how to find out the difference in estimation between the team members. One way to do so is by visually tracking the aerial robot (see section 6 for details). But in this way the aerial robot should be always in the field of view of the ground one. Another method is by using RFID tags, this method used originally by [5]. An RFID tag is a device which can be incorporated into an object for the purpose of tracking and identification using radio waves, these tags can hold information which is stored and accessed by robots. At some points the aerial robot dispenses RFID tags, and stores in them their estimated locations. When the ground robot reads the content of one of these tags, it compares the tag position (which is estimated by the ground robot) and the read position (which is

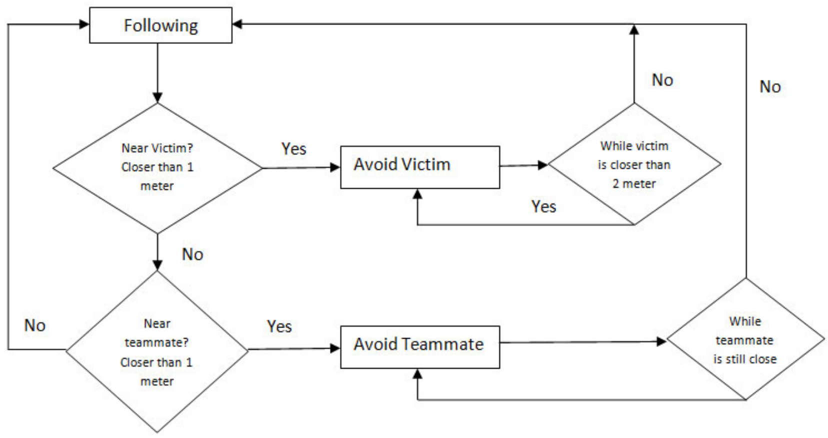

Figure 2: Diagram of the following behavior

estimated by the aerial robot), then computes the new shift vector and stores it for further victim positions sent by the aerial robot.

The dropping could happen periodically, for example every 60 seconds. It is important to mention that the more tags the aerial robot carries the slower it becomes. So there should be a compromise between the number of carried tags and the speed of the aerial robot.

A different strategy could be dropping a tag on top of the victims, but in this case there should not be huge distance between two victims, otherwise the shift could be too large.

4 Experiments and Results

We tested our heterogeneous team on out-door environment. We used the "P2AT" for the ground robot. The ground robot will visit the places reported by the aerial robot and detect the victim locations accurately. At end of the test run, there were even some victims that the ground robot didn't have time to visit and detect their exact position.

To make an objective comparison between the homogeneous team and the heterogeneous one, the competition should be run against other different homogeneous teams (to make sure that the improvement is because of the nature of the team (homogeneous or heterogeneous), but not because of the strategy used by one of the homogeneous teams). And it should be tested many times to eliminate the randomness factor. Unfortunately, due to a lack of resources, we could not conduct such comparison.

5 Conclusion

In this paper we discussed the difference between the homogeneous and heterogeneous teams. We discussed the pros and cons of the two types; we showed that the flying robot is suitable for exploration but it gives non-accurate estimations, while the ground robot is slower but is more accurate. Our heterogeneous team is composed of a tele-operated flying robot which explores the environment and reports the approximated locations of the victims, and a semi-autonomous ground robot that investigates the positions provided by the flying robot and specify the exact locations of the victims. We discussed how gradual correction of position estimation of the flying robot can improve the rescue mission.

6 Discussion

Our team works well in the outdoor environment, but care must be taken when testing the team in indoor environments. The exploration of the aerial robot becomes very restricted due to the small free space and the many obstacles around the robot, which make the possibility of collision higher. Also, victims in indoor environment could be in places can't be reached by the aerial robot, and the ground exploration is required.

It may seem that the added value is due to using aerial robot; in this case it would be more suitable to use a team of two flying robots, so we can explore faster, cover larger spaces and detect more victims. But the victim locations wouldn't be accurate in this case, and we may lose the effort in exploration because of the inaccuracy.

7 Future Work

In this section we discuss three topics of future work: achieving the position estimation update via visual tracking, free space detection, and finally extending the team to more robots. As mentioned before, one method to calculate the estimation difference between the two robots is by visual tracking. The ground robot can estimate the position of the aerial one by computing its relative position. Using the height of the aerial robot, the ground robot can continuously compute the difference and transmit it to the aerial robot. To improve the tracking we can color the aerial orange which makes the tracking easier.

The second option for improvement is using the free space detection to help the ground robot mapping the environment. The main sensors used by the ground robots to localize themselves are the laser sensors, but they can only work for few meters. By using both the visual and the laser sensors, the free space can be recognized for farther areas [6]. The main idea of this improvement is to model the free space ground using the color histogram. The laser sensors provide the ground truth for the training. This model is updated and transmitted to the aerial robot, which has more global view. The aerial robot uses this model to detect the free space areas and report the ground robot about them.

An important field of further work is extending the team size to more than two. Many questions should be answered: how the ground robots communicate with each other? How to divide the work of the ground robots (going to the victims)? What's the difficulty in adding more than one flying robot to explore the environment autonomously (as we can operate only one robot)? And what's the best division (what is the ratio between the ground robots and the aerial ones)?

References

[1] Balakirsky, S., Scrapper, C., Carpin, S., and Lewis, M. (2006). "Usarsim: Providing a framework for multi-robot performance evaluation". In *Proc. of the Performance Metrics for Intelligent Systems (PerMIS) Workshop*, pages 98-102.

[2] S. Balakirsky, S. Carpin, A.Kleiner, M. Lewis, A. Visser, J. Wang and V.A. Ziparo (2007). "Towards heterogeneous robot teams for disaster mitigation: Results and Performance Metrics from RoboCup Rescue". Journal of Field Robotics 24(11-12):943-967, 2007

[3] S. Balakirsky, C. Scrapper, S. Carpin. "The Evolution of Performance Metrics in the RoboCup Rescue Virtual Robot Competition". Proceedings of PerMIS 2007, 91-96

[4] Arnoud Visser and Bayu A. Slamet, "Including communication success in the estimation of information gain for multi-robot exploration"

[5] Vittorio Amos Ziparo, Alexander Kleiner, Alessandro Farinelli, Luca Marchetti, Daniele Nardi, "cooperative Exploration for USARrobots with indirect communication"

[6] Gideon Maillette de Buy Wenniger, Tijn Schmits, "Identifying Free Space in a Robot Bird-Eye View "

[7] Personal communication with the Tech Committee of the 2008 competition, Arnoud Visser, January 29 2009.