How to Win RoboCup@Work?

The Swarmlab@Work Approach Revealed

Sjriek Alers[1], Daniel Claes[1], Joscha Fossel[1], Daniel Hennes[2], Karl Tuyls[1], and Gerhard Weiss[1]

[1] Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands
[2] European Space Agency, ESTEC, Keplerlaan 1, Noordwijk, The Netherlands

Abstract. In this paper we summarize how the Swarmlab@Work team has won the 2013 world championship title in the RoboCup@Work league, which aims to facilitate the use of autonomous robots in industry. The various techniques that have been combined to win the competition come from different computer science domains, entailing learning, (simultaneous) localization and mapping, navigation, object recognition and object manipulation. While the RoboCup@Work league is not a standard platform league, all participants used a (customized) Kuka youBot. The youBot is a ground based platform, capable of omnidirectional movement and equipped with a five degree of freedom arm featuring a parallel gripper.

1 Introduction

RoboCup@Work is a recently launched competition which aims at flexible robotic solutions in work-related scenarios. The leagues vision[1] is to "foster research and development that enables use of innovative mobile robots equipped with advanced manipulators for current and future industrial applications, where robots cooperate with human workers for complex tasks ranging from manufacturing, automation, and parts handling up to general logistics".

In contrast to the well developed robotic solutions deployed in common mass-production environments, RoboCup@Work targets smaller companies in which flexible multi-purpose solutions are required, which are not yet available in industry. Example tasks are finding and acquiring parts, transportation to and from dynamic locations, assembly of simple objects etc. From these industrial goals various scientific challenges arise. For example, in perception, path planning, grasp planning, decision making, adaptivity and learning, as well as in multi-robot and human-robot cooperation.

The competition is relative new and started in 2012. Due to the fact that the it is still in a startup phase and not much of source code of last years event was made public, no extensive resources regarding RoboCup@Work competition where available to start from. From this year on all teams agreed to release their codebase so that for next events the teams can improve their capabilities and new teams are able to catch up with the competition.

---

[1] `http://robocupatwork.org/`

S. Behnke et al. (Eds.): RoboCup 2013, LNAI 8371, pp. 147–158, 2014.

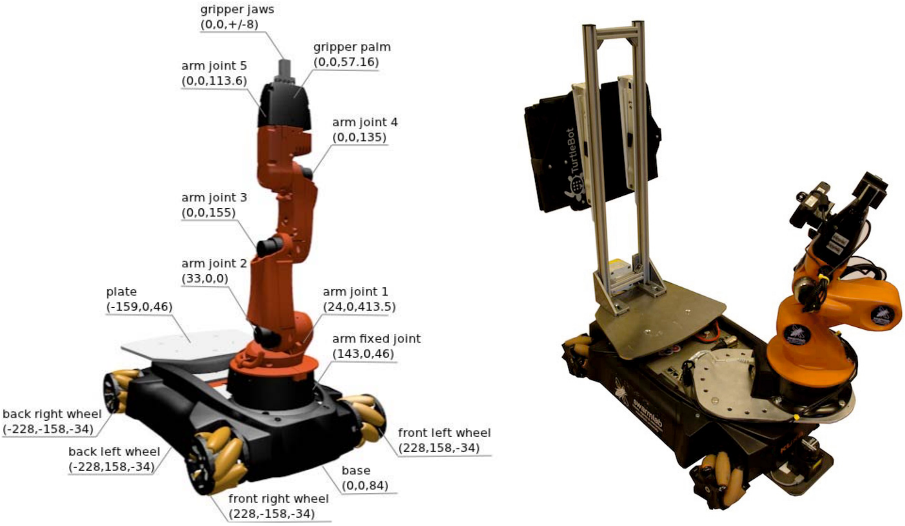

(a) (b)

Fig. 1. (a) CAD model of a stock youBot. (b) Swarmlab@Work modified youBot.

2 Swarmlab@Work

Swarmlab[2] is the research laboratory of the Department of Knowledge Engineering at Maastricht University that focuses on designing Autonomous Systems. The general research goal and mission of the lab is to create adaptive systems and robots that are able to autonomously operate in complex and diverse environments. The Swarmlab@Work team [1] has been established in the beginning of 2013, consisting of 5 PhD candidates and 2 senior faculty members. Since then the team has won the @Work competitions of the 2013 RoboCup German Open, and the 2013 RoboCup world championship. The team's mission is (a) to apply Swarmlab research achievements in the area of industrial robotics and (b) identify new research challenges that connect to the core Swarmlab research areas: autonomous systems, reinforcement learning and multi-robot coordination.

In the remainder of this section we introduce the Swarmlab@Work robot platform that was used for the 2013 RoboCup@Work world championship. Especially, we describe the necessary hardware modifications that were made to adapt this standard platform to the winning configuration.

2.1 youBot Platform

The youBot is an omni-directional platform that has four mecanum [8] wheels, a 5 DoF manipulator and a two finger gripper. The platform is manufactured

---

[2] `http://swarmlab.unimaas.nl`

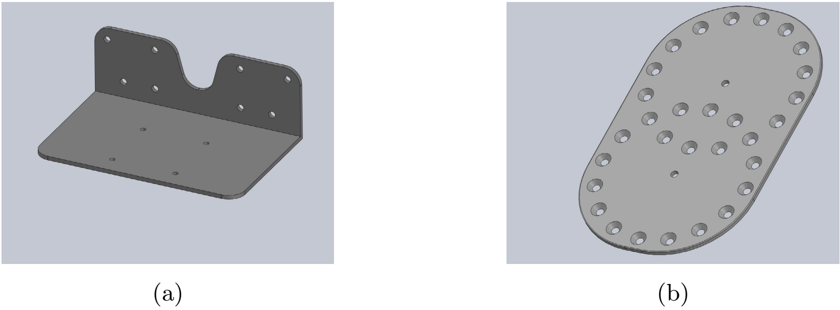

(a) (b)

Fig. 2. (a) a CAD model of the laser mounts, (b) a CAD model of the arm extension plate

by KUKA[3], and is commercially available at the youBot-store[4]. It has been designed to work in industrial like environments and to perform various industrial tasks. With this open-source robot, KUKA is targeting educational and research markets. Figure 1a shows a model of the stock youBot.

The omnidirectional base is approximately 580x380 mm in size and it's height is 140*mm*. It weighs 20 kg and it can carry a payload of 20 kg. Due to the fact that the platform has four mecanum wheels it has the possibility to move in every direction with a maximum speed of 0.8 m/s. The chassis is build up of 3*mm* thick stainless steel plates. Inside the base, a mini ITX Atom based PC-Board with 2 GB RAM is located together with a 32 GB SSD Flash drive. This computer enables onboard processing and control of the robot. Inside the base also the different motors and motor-driver boards are located as well as the power charging and power-distributing unit. The robot is powered by a 24Volt, 5 Ampere maintenance free rechargeable lead acid batteries. All hardware modules of the robot communicate internally over real-time EtherCat [9].

The youBot comes with a 5-degree-of-freedom arm that is made from casted magnesium, and has a 2-degree-of-freedom gripper. The arm is 655 mm high, weights 6.3 kg, and can handle a payload of up to 0.5 kg. The working envelope of the arm is 0.513 m^3, and is is connected over EtherCat with the internal computer, and has a power consumption limit of 80 Watts. The gripper has two detachable fingers that can be remounted in different configurations. The gripper has a stroke of 20mm and a reach of 50 mm, it opens and closes with an approximate speed of 1 cm/s.

In order to meet the requirements we demand from the youBot platform to tackle the challenges, we made a number of modifications to the robot. In this paragraph we describe which parts are modified and why these modifications are a necessity for our approach. Figure 1b shows the modified Swarmlab@Work youBot setup.

For perceiving the environment, two Hokuyo URG-04LX-UG01 light detection and ranging (LIDAR) sensors are mounted parallel to the floor on the front

---

[3] `http://kuka.com`

[4] `http://youbot-store.com/`

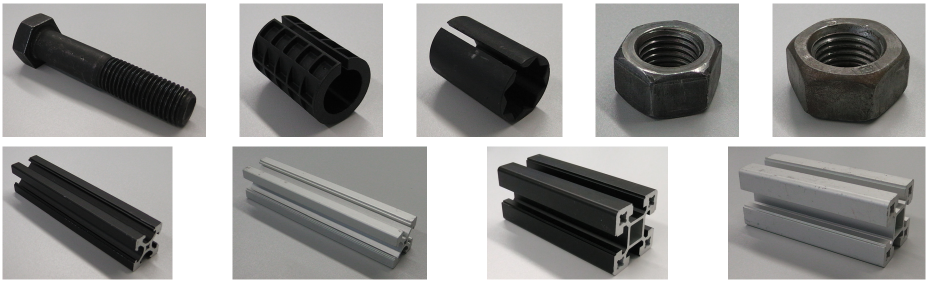

Fig. 3. Manipulation objects, from top left to bottom right: M20_100, R20, V20, M20, M30, F20_20_B, F20_20_G, S40_40_B, S40_40_G

and back of the robot. For mounting these LIDAR sensors a special laser mounting bracket was designed, see Figure 2a. In contrast to the thin aluminium stock brackets, these custom brackets are constructed from 4mm thick steel in order to prevent deformation, reduce vibration and to ensure constant horizontal alignment.

To extend the reach of the robot-arm in respect to the chassis, we designed an extension plate of 5 mm thick aluminium, the CAD model of this plate is shown in Figure 2b. This plate can extend the arm towards the bounds of the chassis, and is designed to be a multi-purpose extension for the youBot arm. This means that the plate can be mounted onto the chassis in angular steps of 22.5 degrees, placing the arm in the optimal positions for the given tasks. By using this plate we can gain an additional reach of the robot arm of 12.5 cm. A secondary advantage of this plate is that also the arm can be mounted onto this plate in steps of 22.5 degrees. This feature is very useful to optimize the turning position of the arm regarding its dead angle. While mounting this plate, we elevated the arm position approximately 2.5 cm, which also results in a reach advantage of about 2 cm. The maximum reach is achieved when the arm is stretched horizontally, i.e. the second joint is aligned such that the rest of the arm is parallel to the ground. To grasp an object from above, the fourth joint has to be perpendicular to the ground. In this configuration, the fingers extend further down than the mounting point of the arm by 7.5 cm. Our configuration mounts the arm at about 15.5 cm, i.e. the optimal picking height is at 8 cm, which is close to the 10 cm as specified as height of the service areas where the objects have to be picked up and placed.

In order to detect and recognize manipulation objects, an ASUS Xtion PRO LIVE RGBD camera is attached to the last arm joint. This camera is mounted, so that it faces away from the manipulator, as can be seen in Figure 7a. The main idea of this rather odd mounting position is that we want to use the RGB-D data of this camera, which is only available for distances larger than ~0.5m. This is why we use a special pre-grip pose, in which we scan the area and try to look at the objects from above. A useful byproduct of this mounting technique is that grasped objects do not block the view of the camera, which typically happens with manipulator facing cameras. The downside is that it prevents the

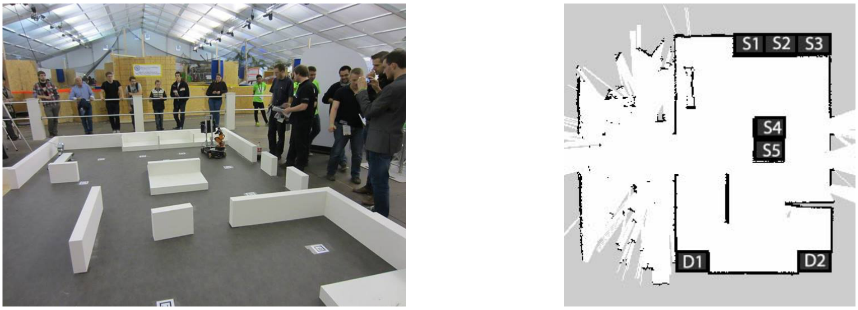

(a) (b)

Fig. 4. (a) 2013 arena during BNT run with extra obstacles. (b) Map of the arena. Annotated with source (S1-S5) and destination (D1 and D2) areas.

use of visual servoing. However, we found that the arm can be steered precisely enough to make this disadvantage negligible.

The base computer is upgraded from the stock ATOM based architecture to an Intel i7 CPU and is powered by a dedicated 14.8V / 5 Ah Lithium Polymer battery pack that is charged and monitored by an OpenPSU power unit. For cooling we mounted an additional fan in the base of the robot. The base computer is supported by an i5 notebook, which is mounted on a rack at the backside of the robot.

For safety reasons the robot is equipped with an emergency stop button, that stops all robot movement, without affecting the processing units, so when the stop is released the robot can continue its movement without having to re-initialize it again.

In this section we described the hardware modifications needed to optimally execute the required RoboCup@Work tasks, in the next section these tasks will be describes in more detail, as well as the methods that we used to solve the specific tasks.

3 RoboCup@Work

In this section, we introduce the various tests of the 2013 RoboCup@Work world championship. Additionally, we sketch the different capabilities that we developed and explain briefly how these techniques are combined to tackle the different tests. Note that while the tests are slightly different for various levels of difficulty, we only give a description for the difficulty chosen by Swarmlab@Work. For a more detailed description we refer the reader to the RoboCup@Work rule book[5]. In the following the term service area refers to a small table from which typically objects have to be grasped, while destination area refers to a small table on which objects have to be placed. Figure 5a shows our robot performing a grasp in one of the service areas. A list of manipulation objects used in 2013

---

[5] `http://robocupatwork.org/resources.html`

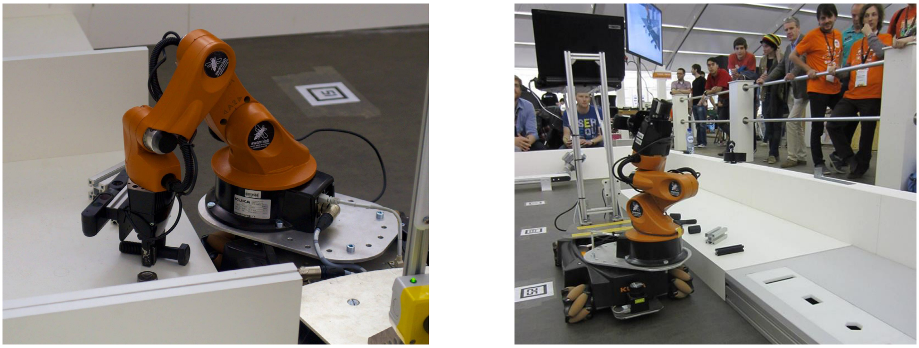

(a) (b)

Fig. 5. (a) Grasping the M20_100 in a service area. (b) PPT setup.

is given in Figure 3. In the arena, there are five service areas and two destination areas as annotated in Figure 4b. A picture of the 2013 arena is shown in Figure 4a, in which two extra obstacles are placed.

Basic Navigation Test. The purpose of the Basic Navigation Test (BNT) is testing mapping, localization, path planning, obstacle avoidance and motion control. The robot has to autonomously enter the (known) arena, visit multiple way-points, and exit the arena again. Each way-point has to be accurately covered by the robot for a given duration. Additionally, obstacles may be spawned randomly in the arena, where a penalty is given if the robot collides with an obstacle.

Basic Manipulation Test. In the Basic Manipulation Test (BMT), the robot has to demonstrate basic manipulation capabilities such as grasping and placing an object. For this purpose three target and two decoy objects are placed on a service area, and the robot has to single out the three actual target objects and move them to any other area. Implicitly this also tests wether the object recognition works sufficiently well. Since navigation is not a part of this test, the robot may be placed at any location before the test starts.

Basic Transportation Test. The Basic Transportation Test (BTT) is a combination of the BNT and BMT. Here, certain objects have to be picked up from various service areas and transported to specific destination areas. Similar to BNT and BMT both decoy objects and obstacles are placed on the service areas and in the arena, respectively. The robot has to start and end the BTT outside of the arena. In addition to the challenges individually posed by BNT and BMT, the robot now has to approach multiple service and destination areas within a certain time window, hence for optimal performance determining the optimal route and payload configuration is necessary.

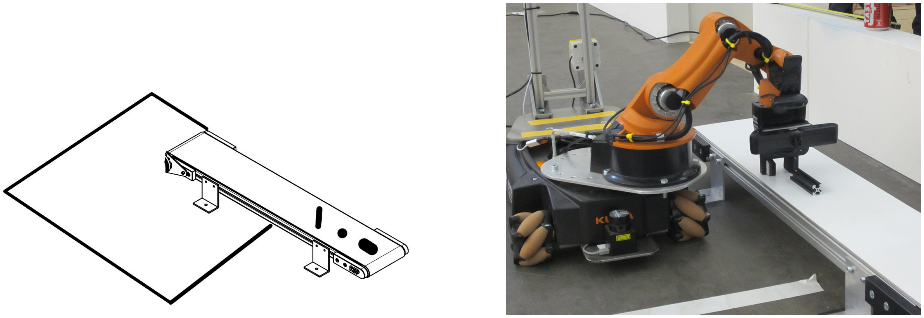

(a) (b)

Fig. 6. (a) Schematic of the conveyor belt setup. (b) Our CBT approach.

Precision Placement Test. The Precision Placement Test (PPT) consist of grasping objects (see BMT), followed by placing them very accurately in object-specific cavities. Every object has one specific cavity for vertical, and one specific cavity for horizontal placement. Each cavity has the same outline as the object plus 10 percent tolerance. Figure 5b show the PPT destination area with some sample cavities.

Conveyor Belt Test. In the Conveyor Belt Test (CBT) the robot has to manipulate moving objects. For this purpose the robot has to autonomously approach the conveyor belt, grasp the object(s) which pass by, and place them on its rear platform. Due to the increased base difficulty, no decoy object have been used in the 2013 competition in this test. Figure 6a show the conveyor belt, which moves with approximately 3 cm/s.

4 Approach

In this section, the different techniques used to tackle the above mentioned tests are explained. We developed different modules for many different capabilities, e.g. basic global navigation, fine-positioning, object recognition, inverse kinematics of the arm, etc. We also explain our various recovery methods. By combining these capabilities in state-machines we are able to solve the tasks specified above.

Mapping and Localization. One of the most crucial capabilities of an autonomous agent is to localise itself efficiently in a known environment. To achieve this, we use gmapping [5] to build a map of the arena beforehand. The map of this years arena is shown in Figure 4b. After the map is recorded it can be used by AMCL [3] for efficient global localization.

Navigation. Another necessary capability of the robot is to navigate in the known environment without colliding with obstacles. The map created with

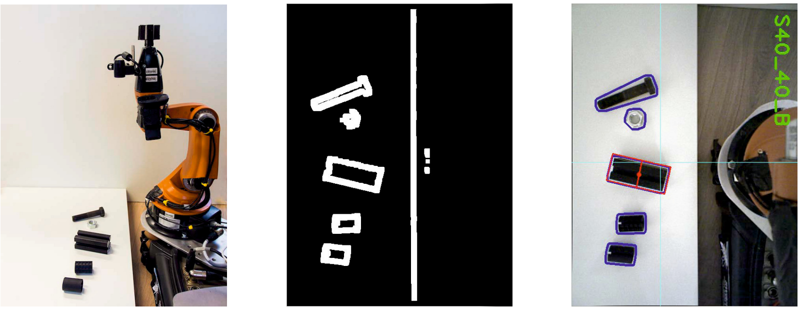

(a) (b) (c)

Fig. 7. (a) Pre-grip scan position. (b) Pre-processing of the image. (c) Detected objects, classification and grasp position of the closest object.

gmapping is used for the basic global navigation. The global path is computed by an A* algorithm and is then executed using a dynamic window approach [4] trajectory rollout for the local path planning. This planner samples different velocities and ranks them according to the distance to the goal and the distance to the path, while velocities that collide with the environment are excluded. We tuned the trajectory rollout in such a way that it incorporates the omni-directional capabilities of the robot and also included scores for the correct destination heading, when being in close distance to the goal.

Fine-Positioning. The navigation is very well suited to navigate between larger distances from different positions in the map. But the accuracy of the navigation and localization is not high enough to navigate with high reproducibility to previously known goals. Thus for aligning to those previously known locations another technique is needed, since we want to be as close as possible to the manipulation areas as possible. Also for the basic navigation test, the markers have to be covered exactly, which is not always the case when using only AMCL for localization. Thus, we implemented ICP based scan registration [7] for fine grain positioning. This techniques records a laser scan at a certain position. When the robot is close to this position, the differences between the current laser scan and the registered one are calculated into a direction which is used to steer the robot to the old position with very high accuracy. By using correspondence rejectors we make sure that slight changes in the environment, e.g. caused by humans standing around the robot, do not interfere with the scan registration accuracy.

Object/Hole Recognition. Besides all the navigation tasks, object detection and recognition is crucial to be able to interact with the environment, i.e. picking up objects and placing them in the correct target locations. We use the openCV-library[6] to detect the objects. An adaptive threshold filter is applied to the input

---

[6] `http://opencv.org`

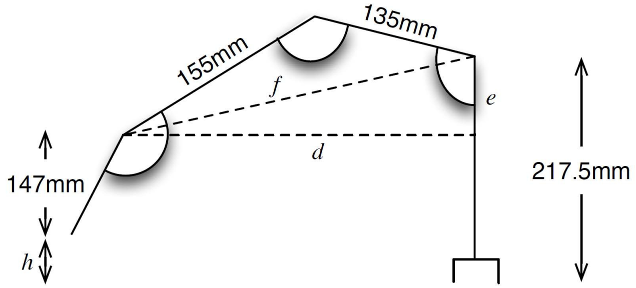

Fig. 8. Simple inverse kinematics: d and h are the grip distance and height, relative to the mount point of the arm. By always gripping from a top down position, e and f can be calculated and by that we can determine all angles for the joints.

image. Afterwards the image is converted into a black and white image and this is used to detect the contours of the objects as shown in Figure 7b. We use various features of the detected objects, e.g., length of principal axis, average intensity and area, and use a labeled data-set that we created on the location to train a J4.8 decision tree in weka [6] for the recognition of the objects. This decision tree is then used for the online classification of the objects. Figure 7c shows the detection in a service area.

Inverse Kinematics of the Arm. In order to manipulate the detected objects, the various joints of the arm have to be controlled such that the objects is grasped correctly. We implemented a simple inverse kinematics [10] module to calculate the joint values for any top-down gripping point that is in the reach of the robot. Since we are gripping from a top-down position, the inverse kinematics can be solved exactly, when we fix the first joint such that it is always pointing in the direction of the gripping point as shown in Figure 8. Then the remaining joints can be calculated in a straight forward manner, by solving the angles of a triangle with three known side lengths, since we know the distance of the grip and also the lengths of all the arm-segments. Since the position-reproducibility of the arm is in sub millimetre order, this proved to be sufficient for performing highly accurate grasp and place trajectories.

Recovery Methods. Of course, in robotics many things can go wrong, so we include various recovery methods. The navigation can easily get stuck, since the robot is tuned to drive very close to obstacles. Therefore, we implemented a force field [2] recovery. More specifically, the robot moves away from every obstacle that is picked up by the laser scan for a couple of seconds and the navigation is restarted. When the object recognition misclassifies an object, which happens especially with the two silver aluminium profiles, it can happen that the robot tries to pick up object S40_40_G. This is physically impossible, since the gripper cannot open sufficiently wide. Thus, we measure the voltages in the arm, and as

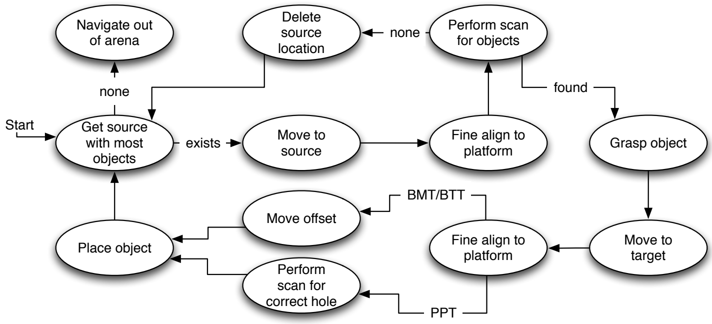

Fig. 9. The simplified global state-machine that is used for the BMT, BTT and PPT competitions, without recovery states

soon as a certain voltage is exceeded, we recover the arm to an upwards position and back up a couple of states.

4.1 State Machine

For the different tests, we combine the various capabilities explained in the previous section in different state machines to complete the tasks. For the BNT, we do not need any of the manipulation techniques, so we have a very small state machine that combines the navigation and the fine-positioning. More specifically, we use the navigation to get close to the position and then the fine-positioning moves the robot exactly to the asked position. If the navigation fails, the force-field recovery is activated. Afterwards it retries to move to the current location. When it fails more than a fixed number of times, the current goal is skipped.

In principle the BMT, BTT and PPT tasks do not differ very much. For instance, the BMT is basically a BTT, in which only one source and one target location are used and the PPT is a BMT in which the objects have to be placed in a specific location instead of anywhere on the platform. Thus for all these other tests (excluding CBT), we created a global state-machine that is shown in its simplified form in Figure 9. For sake of simplicity, the recovery connections are left out.

The general idea of the state machine is to move to the source location which has the most objects that we have to pick up. Then we perform a scanning motion along the platform, while looking for any object that we need. If we do not find anything we continue to the next source location or end, if we have visited all source locations specified. On the other hand, when an object is found, the robot grasps it and moves to the target location specified for this object. If it is a BMT or BTT, we move a certain offset to ensure that the objects are not stacked and place the object. If it is a PPT, we perform a scanning motion, however now looking for the matching hole for the object we are currently transporting. Afterwards, we again move to the current source location and repeat until everything

Table 1. Scores for the different tests

Place	Team	BNT	BMT	BTT	PPT	CBT	FINAL	Total
1	Swarmlab@Work	616	600	600	600	200	337.5	2953.5
2	LUHbots	490	225	500	112.5	75	400	1802.5
3	WF Wolves	211.25	0	37.5	131.25	0	206.25	586.25
4	RoboErectus@Work	243.75	37.5	0	150	0	37.5	468.75
5	b-it-bots	280	0	0	187.5	0	0	467.5

is transported. For almost every state we have a recovery state connected, that kicks in if the arm or the navigation fails. These recovery states are then usually connected to the previous state, which is re-initiated for a number of times and skipped as soon the number of retries are exceeded.

For the CBT, we did not have the facilities to prepare anything beforehand and there was only very little preparation time to solve the task. We came up with a straight forward solution. The robot aligns itself next to the conveyer belt and prepares the gripper such that it is hovering over the belt. An additional camera attached to the side of the arm searches for the object and as soon as the centre of the object is detected within a certain window the gripper starts to close. Figure 6b shows the approach.

5 Competition

Table 1 shows the final scores for the 2013 RoboCup championship, where "final" refers to a BTT with six objects instead of three.

While scoring highest in most tests, Swarmlab@Work did not perform outstandingly well in the finals. The reason for that is twofold: (a) We are the only team which did not use the back platform to store multiple objects, but instead transport only one object at a time. (b) Due to the lack of gripper feedback we always double check wether an object has indeed been gripped/transported. Both (a) and (b) resulted in insufficient time to finish the final runs.

Additionally, one can see that the average performance over all teams in the CBT is very low. This can be explained by the difficulty of the test, and also the lack of pre-competition preparation time since most teams do not have access to a conveyor belt outside the competition. Videos and further information can be found on our website[7].

6 Future Work

The Swarmlab@Work code will be released on Github[8]. For future competitions we plan on improving the object recognition by adding state of the art techniques, such as deep learning algorithms in order to decrease vulnerability to

---

[7] `http://swarmlab.unimaas.nl/robocup_at_work/robocup-2013/`

[8] `http://github.com/swarmlab`

lighting conditions. We also hope that this will allow for a more perspective stable object recognition. Furthermore, we are currently working on a learning by demonstration approach to allow for learning new arm trajectories on the fly. This might enable grasping various 'more difficult to grasp objects', e.g. glasses. It will also help to cope with future RoboCup@Work tests, which are yet to be introduced, e.g. assembly tasks. Combining the two previously mentioned approaches, we also hope to overcome the poor CBT performance.

Hardware-wise we plan to upgrade the gripper opening speed and distance in order to increase the range of graspable objects that the youBot can manipulate. We also plan on mounting the arm on a rotating extension plate, so that the arm can be turned towards the optimal grasp position. These modifications will unlock challenges with higher difficulty levels, which require amongst others the ability to grasp all available @Work objects - a capability we lacked in 2013. The use of the backplate to collect objects on, and thus enabling to transport more objects efficiently, will also be required in the next RoboCup challenge.

References

1. Alers, S., Bloembergen, D., Claes, D., Fossel, J., Hennes, D., Tuyls, K., Weiss, G.: Swarmlab@work team description paper. Robocup (2013)
2. Chong, N.Y., Kotoku, T., Ohba, K., Tanie, K.: Virtual repulsive force field guided coordination for multi-telerobot collaboration. In: Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA 2001), pp. 1013–1018 (2001)
3. Fox, D., Burgard, W., Frank, D., Thrun, S.: Monte carlo localization: Efficient position estimation for mobile robots. In: Proceedings of the Sixteenth National Conference on Artificial Intelligence, AAAI 1999 (1999)
4. Fox, D., Burgard, W., Thrun, S.: The dynamic window approach to collision avoidance. IEEE Robotics & Automation Magazine 4 (1997)
5. Grisetti, G., Stachniss, C., Burgard, W.: Improved techniques for grid mapping with rao-blackwellized particle filters. IEEE Transactions on Robotics 23, 43–46 (2007)
6. Hall, M., Eibe, F., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: an update. SIGKDD Explor. Newsl. 11(1), 10–18 (2009)
7. Holz, D., Behnke, S.: Sancta simplicitas - on the efficiency and achievable results of slam using icp-based incremental registration. In: IEEE International Conference on Robotics and Automation, pp. 1380–1387 (2010)
8. Hon, B.E.: Wheels for a course stable selfpropelling vehicle movable in any desired direction on the ground or some other base, US Patent 3,876,255 (1975)
9. Jansen, D., Buttner, H.: Real-time ethernet: the ethercat solution. Computing and Control Engineering 15(1), 16–21 (2004)
10. McCarthy, J.M.: An introduction to theoretical kinematics. MIT Press, Cambridge (1990)