Vision-based People Detection Utilizing Reflective Vests for Autonomous Transportation Applications

Henrik Andreasson, Abdelbaki Bouguerra, Todoy Stoyanov, Martin Magnusson and Achim Lilienthal

Abstract— This paper presents a single camera system which can be used to detect people wearing reflective vests. The system has been evaluated in indoor and challenging outdoor environments with direct sunlight into the camera. A monochrome CMOS sensor equipped with an IR bandpass filter, a fish-eye lens and a set of IR-LEDs are utilized to take a pair of image snapshots, one with and one without the IR-flash. These pairs are then processed to detect areas with significant changes in the intensity values, which are likely to correspond to reflective materials such as reflective vests.

I. INTRODUCTION

People detection is an important task which has to be addressed on the way towards safe autonomous machines capable of operating in a shared workspace with humans. It has also received attention from driver assistant technology community to automatically detect pedestrians by, for example, Volvo and Subaru. Both companies use a combination of range data with vision; Volvo uses radar and a single camera, whereas Subaru uses stereo cameras. It is worth noting that these car-systems are not explicitly used for people detection but rather for more generic obstacle avoidance and for automatic speed control.

To perform people detection, many different sensor modalities are used and commonly combined. One of the most utilized sensor is laser scanners, see for example [1], due to its high precision and accurate distance information which makes segmentation and clustering problem much more straightforward than pure camera based solutions. Cameras are on the other hand less expensive, provide high density data but have difficulties to handle a large variety of light conditions. There exists a large amount of work that deals with sensor fusion (laser scanner and camera) to detect people [2]. Another example of a sensor which has been used for people detection are thermal cameras [3], which utilize the emitted heat of humans.

The main difference compared to the related work mentioned above is that in this work we exploit a restriction imposed on the people present in the vehicle's environment to enhance their visibility, namely that they must wear reflective vests, see Fig. 1. With the system proposed in this paper, the reflective material of the vest is detected by processing a pair of images, one taken with an IR-flash and the other taken directly afterwards without active illumination. By processing this pair of images, the system is able to detect persons despite difficult and varying illumination settings ranging from bright sunlight to dark environments.



Fig. 1. A reflective vest. Please note that there is a large variety of reflective vests and there is no standard pattern in which the reflective areas are arranged. It is also common to add reflective markers directly into clothes, an example of a jacket with inbuilt reflective material can be seen in Fig. 7.

II. RELATED WORK

The core principle of the proposed people detection system is to take images with and without the IR-flash and process them as a pair. Therefore the technique presented does not directly compare to other work mentioned here.

To give another viewpoint on safety and people detection, this paragraph discusses the regulations that determines how automated guided vehicles (AGV) can operate in the vicinity of humans. Modern AGVs are allowed to have noncontact obstacle detection systems, laser scanners, instead of bumpers. These systems are regulated by safety standards including US ASME B56.5 and the British EN 1525 / 1526. In the specification EN 1525, the object size and reflectivity that need to be detected to comply with the standards are specified. Based on the performance of the sensor, the sensor is classified for a specific max safety distance, for example, the SICK \$300 has a max safety distance of 2 meters, whereas the SICK S3000 has a 4 meters safety distance. One example of the requirement of the object size is for the sensor to detect a 600 mm cylinder lying perpendicular towards the sensor with a diameter of 200 mm. This volume is intended to represent a fainted person lying flat on the ground. Worth to note is that the safety standards do not distinguish between obstacles and humans. This distinction is, however, very important since a person lying on the ground may be merely a small bump on the path for large vehicles such as the vehicle shown in Fig. 8. These standards are only valid where the flat floor assumption holds and, hence, it is far from being applicable to outdoor environments in uneven terrain or, for

H. Andreasson, A. Bouguerra, T. Stoyanov, M. Magnusson and A. Lilienthal are with AASS, Dept. of Technology, Örebro University, Sweden, firstname.lastname@oru.se

example, in heavy rain/snowfall.

One shortcoming of research performed on people detection is that the case of a fainted person (person lying flat down on the ground) is typically not considered, many authors only consider persons standing upright [1].

To detect people using only visual information are commonly done by utilizing visual features in combinations with with machine learning techniques [4]. In principle, the problem of people detection and object detection using cameras is very similar. However, to list a complete survey of one of the most active topics in computer vision is beyond the scope of this paper. A survey about pedestrian detection using a single camera is presented by Enzweilier et al. [5].

Worth mentioning and also related to our work is the use of reflectivity values from other active sensors, which emit signals such as light. For example, the state-of-the-shelf navigation system for AGVs utilize reflectors as artificial landmarks in the environment. Most laser scanners provide reflectivity values in addition to the range and bearing values which could be used to detect reflective areas. Essentially the same holds for time-of-flight (TOF) cameras which use active illumination to determine range data and also return intensity values. The key difference here is to combine the camera images with and without the emitted light. The most similar case to the previous mentioned sensors above would be to only utilize data from the image taken with the flash (and not data without the flash).

One notable difference to systems which are used to assist drivers in cars is the field of view the sensor has to cover. A loading / unloading scenario frequently comprises sharp turns, reversing etc. compared to driver assistance systems for cars, which need to observe a relatively narrow cone in front of the car, a crucial requirement for our application scenario is a very large field of view (FOV). This imposes that the system should be able to detect people with a low resolution.

III. DETECTION SYSTEM

The system works by comparing two images, one taken with the flash I_f and one taken without I_{nf} . Since the system is mounted on a moving vehicle straightforward background subtraction methods are not applicable here. Instead small interesting sub regions in the images, called keypoints, are used in order to relate the two images I_f and I_{nf} to each other. The method of extracting keypoints is selected based on its ability to detect reflective areas, i.e., local peaks in the intensity values. Each keypoint, please note that keypoints are only extracted from the I_f image, is then tracked in the image without flash I_{nf} . Since I_n and I_{nf} are taken in short succession, the displacement of the tracked keypoints between the two frames should be relatively small, even if during a quick turn or a bumpy ride. Please note that the processing sees the two images as a pair. Based on how the keypoints can be tracked, the system utilizes a window around the keypoint to determine if the intensity change is above a predetermined value. If so the system reports this area to contain reflective material. An overview



Fig. 2. An overview of the data flow within the system. A squared boxes indicates the data type and a box with rounded corners a method.

diagram can be seen in Fig. 2 and an example detection result can be seen in Fig. 3. One important criterion is to perform the detection at "real-time", which for computer vision applications sometimes is referred to be at least 10 Hz [6].

A. Hardware

The system consists of a standard of-the-shelf monochrome CMOS sensor, IDS imaging USB UI-1228LE, with a resolution of 752 \times 480 pixels. A band pass filter with a center wavelength of 852 nm with a full width-half max of 10 nm is mounted between the lens and the sensor. The lens is a fish-eye type with an approximate FOV of 180 degrees. Next to the camera 8 IR LEDs with a wavelength of 850 nm are placed around the camera, see Fig. 4. The emission characteristic of the LEDs is such that the emission reaches its maximum in the direction normal to the LED and falls down to 50% at an angle of 60 degrees to normal direction. This assures a wide and relatively uniform illumination coverage of the camera's FOV. The LEDs are mounted in a way that directs weaker illumination to the top and bottom of the image. These directions are less important for people detection since they typically correspond to partial views of the vehicle itself (when looking downward) and the sky (when looking upward). Further on, the camera images are rectified to only contain a panorama where the top and bottom parts of the image are removed.

One difficulty of having a large FOV is that object size decreases very rapidly in the image coordinates with larger distance, i.e. only a small amount of pixels might represent the object to be detected. For cameras the width x given



Fig. 3. An example result of the detection system mounted on the autonomous wheel loader, shown in Fig. 8. The top figure shows an image taken with the flash and the bottom one without. The purple circles indicate detections. The sun can be seen at the top right of the image. The filled white circle at the bottom left is a lens artifacts that occurs due to direct sunshine into the sensor. Additional artifacts can be seen along the line between this artifact and the sun.



Fig. 4. The standard CMOS monochrome camera equipped with a band pass filter of near IR and a fish-eye lens. Around the camera lens the 8 IR-LED diodes are clearly visible.

in pixels at which an object appears is given by x = kX/Z, where X is the width of the object, Z is its depth and k is a camera constant depending on the resolution of the camera (higher resolutions correspond to larger values of k) and the lens (a large FOV corresponds to a smaller k). To illuminate a large FOV the diodes have to spread the light to cover the FOV of the camera. The active illumination intensity E decrease with distance to the object by $E \propto Z^{-2}$, where Z is the distance to the illuminated object. This makes distant objects not only to be visible with low spatial resolution in terms of few pixels but also to have low active illumination intensity.

An example of an image pair obtained with and without flash can be seen in Fig. 3.

B. Image Pair Snapshot Collection and Unwrapping

All the post processing is done using a pair of images $I_p = \langle I_f, I_{nf} \rangle$ consisting of one image taken with flash I_f and an image without flash I_{nf} . The time *t* between the images is kept as low as possible. The unwrapping of the fish-eye image to a panorama, a virtual cylinder, is done to limit the area of interest since the top part and the bottom part of the fish-eye image only contain the sky and vehicle respectively.

C. Keypoint Extraction

Since the feature selection is done only in the image taken with the flash, only positive intensity peaks have to be detected. The key point extraction method used is the STAR algorithm by Konologie et al., which is a slightly modified version of CenSurE [7].

The key property we require from the keypoint detector for the proposed system properties is to detect blob-like features over different scales, under the assumption that a reflective vest can be well represented as a single blob at longer distances. The scale invariant property is important to be able to detect intensity peaks of different sizes. Further on, the usage of a blob-like feature detector makes the center of the detected keypoint to be in the center of the intensity peak. Note that the STAR keypoint detector was modified slightly to only respond to positive peaks. This means that a peak with bright center will be detected but not negative peak, a center with a darker area. An example figure containing extracted features, marked as circles can be seen in Fig. 5. There exist a large variety of different keypoints detectors [8], in this case the constraints given by the speed, blob-like detector and scale invariance limited our selection. For example, detectors that were found useful albeit much slower were, for example, the keypoint detector used in SIFT [9]. In general, keypoint detectors that search blobs work better to detect reflective vests than detectors which find corners. In the context of our application scenario, corner detectors required a much higher spatial resolution to detect reflective vests at medium of long distance.

D. Keypoint Tracking and Intensity Check

Many detected keypoints do not originate from reflective vests but are simply areas which are bright and therefore contain much light reflected from the sun, see Fig. 3. To check each keypoint for changes in intensity, each keypoint is tracked in the next frame without the flash using the Lucas-Kanade tracker implementation [10]. Since the tracker only needs to track a few keypoints the computational requirements are kept small. The selection of a "keypoint" tracker serves two purposes. First, the tracker naturally keeps the spatial layout of the features by only searching at a vicinity of the keypoints, called temporal persistence, secondly it is enough to detect keypoints at one image instead of two avoiding the search for correspondences between the two keypoint sets. Since the images in each pair I_p are taken in very short succession (the second images is taken between 1-2 ms after the first one) the additional benefits of various



Fig. 5. An example of keypoint extraction in an outdoor setting. The green circles indicates extracted features where the diameter of the feature corresponds to the scale of the feature. Green circles are classified as regions where there is no reflective material. The purple color indicates an substantial intensity change between the corresponding features in the flash/no flash image pair. The radius of the purple circles indicates the amount of intensity difference. It can be seen that the detected changes are induced by the back lights of two wheel loaders, which indeed consist of reflective material. The intensity change could also originate if the brake lights were altered in between the collection of the image pair I_p . The bucket of the wheel loader is seen in the lower part of the image.



Fig. 6. Keypoint extraction is performed only within a bounding box of the image I_f . This allows features to be tracked even in case of quick rotational movements. Please note that the bounding box constraints are only used when extracting the features and not while tracking the features.

invariant properties such as rotation, affine and intensity are not important and the problem can instead be addressed as a tracking problem. Note that a prediction could be made of the movement of the features by utilizing motion estimates of the vehicle, however, this is not currently utilized.

If the tracker finds a suitable match then an additional evaluation is done to check whether the intensity change inside a window w_i around the tracked keypoint is below a threshold t (in the evaluation t = 10 was used), we considered that the keypoints correspond to an area without reflective material, see Fig. 5. Note that the tracker window w_t is set to be larger than w_i (in the evaluation $w_i = 8$ and $w_t = 16$), which makes the illumination changes to be sensitive to small areas while the tracking utilize more spatial information. If the tracker is unable to track the keypoint in the other images could either occur from very large changes in the intensity or that the keypoint is on or outside the border of the image. To avoid extracting keypoints lying on the border of the image I_f which might not be visible in I_{nf} due to quick moments of the vehicle, only features which are located within a distance from the image boundary are used, see Fig. 6.



Fig. 7. An example of an indoor image pair in a warehouse. On top is the image with the flash and below the image without the flash. Worth notifying is that the reflective pattern of the reflective vests differs significantly between different types of reflective vest. This can be seen in the top figure at the rightmost person, who is wearing a jacket with embedded reflective markers in the fabric.

IV. EXPERIMENTS

The reflective vest detection system has been tested in an indoor warehouse environment and outdoors on a gravel loading production site and a parking lot. As expected, the indoor scenario is less challenging, see Fig. 7. However, despite the simplicity of the scene shown in Fig. 7 there are locations, for example, when passing large open front doors with sun glare where, for example, plain intensity thresholding would fail. During the outdoor experiments the camera unit was mounted close to the roof of a wheel loader, see Fig. 8.

Hand-labeled ground truth data is used to evaluated the system using a sequence of images. Since there is no comparable system to our knowledge which works on a similar principle the system is evaluated against the Histogram of Oriented Gradient (HOG) descriptor approach available in the OpenCV library [11]. The results are shown in Fig. 9, as precision-recall curves for both methods.

All other presented results are shown using only a few, but representative snapshots of the system. Further on, all presented figures and results use the exact same parameters independently of the data set to show the generality of the system. It is also worth mentioning that the result presented is from a single frame system, meaning that no information from previous detections are maintained which could be very beneficial to remove spurious outliers.

A. Ground truth evaluation and comparison with HOG

A set of 120 consecutive images was labeled from a sequence with direct sunshine into the camera, one similar example image can be seen in Fig. 3. In the sequence three persons were present at different distances (close, medium and far/medium). The detected reflective areas in Fig. 3 would correspond to close and medium distance. The HOG evaluation was preformed using the images without the flash I_{nf} , also the HOG based detection had no limitation in terms of computational time, instead the search parameters controlling the search area (stride and padding) was adjust to



Fig. 8. Sensor unit mounted on the vehicle.



Fig. 9. Precision recall graph comparing HOG detector approach with the proposed method denoted 'reflex' in the plot.

improve the results giving a computational time of approximately 3 seconds per image. Each evaluation was repeated for a set of different thresholds for both systems to create precision-recall curves shown in Fig. 9. To avoid multiple detections of the same person influences the results only one true positive was allowed per region.

The results shows that the proposed methods works comparatively well. The HOG detector approach only occasionally managed to find the closest person in the test. The false negatives in this test occurs to a large extent because of large distances to the reflective areas and, in some cases, that the reflective areas were overlapped with lens artifacts due to direct sunshine. False positives occurred mainly from lens artifacts.

B. Fainted Person

As discussed in the Sec. II, one important requirement of the safety standard is to be able to detect people lying flat on the ground. To test the camera system proposed in this paper a reflective vest was placed directly on the ground. As can be seen in Fig. 10, the system did not have any problem detecting the reflective areas of the vest in this setup.

C. Maximum Range

To test the maximum range the system can cope with persons were detected while they were moving away. When the system no longer can see the person the distance to



Fig. 10. The top figure shows an image taken with the flash and the bottom one without. Two persons are visible in the images and a reflective vest lying on the ground representing a fainted person, which would be very difficult to spot in range data alone.

a nearby object (measured by a laser scanner) is used to estimate the maximum distance the system can handle. Based on this principle the maximum range of the system was estimated to be approx. 20 meters, in outdoor clear conditions with sunshine, see Fig 11. The experiment was done with a slightly different hardware setup, where an IR pass (Kodak Wratten) was used instead of the narrow band pass filter. This setup will therefore be more sensitive to background illumination.

D. Snowfall

Some of the data sets were collected in snowy and windy conditions. This can mainly be seen as a slightly bent lines in the image. The system proved to handle the snowfall well. The reflective areas are still detectable, however, as can be seen in Fig. 12, the snow sometimes gives false positives. It is also evident that the snowflakes which are closest to the camera are the ones which gives rises to the strongest responses. This would be one example where a filtering approach could provide a remedy for these spurious false positives.

V. CONCLUSIONS AND FUTURE WORK

We have presented a system that addresses the problem of detecting persons on a site where machines and people are working in close proximity. The system assumes that each persons is wearing a reflective vest, which is obligatory for many workplaces. Currently, the proposed camera works on the basis of a double flash/no flash snapshot and does not exploit information from previous frames, which would help to decrease spurious readings from snow. One important



Fig. 11. These four pictures are from the sequence in which the maximum distance at which a person can be detected was investigated by observing when the person could not detected any longer. The pile is located 20 meters away. During this setup the camera / wheel loader was slightly tilted.

characteristic of the developed system is that it can run without any other sensory input to provide data at high rate (> 10 Hz), the output could also be useful as additional safety information to assist human drivers.

Our future work will include primarily an extensive evaluation with hand labeled ground truth labeled data, which will allow to compare different methods and parameter values. It could also be worth to investigate different sensor setup to extend the maximum detection range. To obtain range, or relative location estimate to the detected regions, geometrical constraints would be needed preferable from an additional range sensor.

Another ongoing work is the classification of the returned reflective areas due to their shape to reduce the false positives originating, for example, from snow, reflective markers or lens artifacts visible in direct sun light. Finally, we will implement a mechanism that tracks reflective blobs over a sequence of frames. This should decrease the amount of spurious false positive and false negative readings.

REFERENCES

- L. E. Navarro-Serment, C. Mertz, and M. Hebert, "Pedestrian detection and tracking using three-dimensional ladar data," *The International Journal of Robotics Research*, May 2010.
- [2] G. Gate, A. Breheret, and F. Nashashibi, "Fast pedestrian detection in dense environment with a laser scanner and a camera," in *VTC Spring*, 2009.
- [3] C. Dai, Y. Zheng, and X. Li, "Pedestrian detection and tracking in infrared imagery using shape and appearance," *Comput. Vis. Image Underst.*, vol. 106, pp. 288–299, May 2007.



Fig. 12. These four pictures are from a larger sequence, which was recorded under conditions of snowfall with wind. From this it can be seen that snowflakes can cause errornous readings especially when they are close (the top figure).

- [4] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: A benchmark," *IEEE Conference on Computer Vision and Pattern Recognition (2009)*, pp. 304–311, 2009.
- [5] M. Enzweiler and D. M. Gavrila, "Monocular pedestrian detection: Survey and experiments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, pp. 2179–2195, 2009.
- [6] D. Beymer and K. Konolige, "Real-time tracking of multiple people using continuous detection," in *Proc. ICCV Frame-rate Workshop*, 1999.
- [7] M. Agrawal, K. Konolige, and M. R. Blas, *CenSurE: Center Surround Extremas for Realtime Feature Detection and Matching*. Springer, 2008, vol. 5305, pp. 102–115.
- [8] K. Mikolajczyk and C. Schmid, "Scale and affine invariant interest point detectors," *International Journal of Computer Vision*, vol. 60, no. 1, pp. 63–86, 2004.
- [9] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," Int. J. Comput. Vision, vol. 60, pp. 91–110, November 2004.
- [10] J.-Y. Bouguet, "Pyramidal implementation of the lucas kanade feature tracker description of the algorithm," 2000.
- [11] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal of Software Tools, 2000.

Performance Indicators for Robotics Systems in Logistics Applications

Wolfgang Echelmeyer, Alice Kirchheim, Achim L. Lilienthal, Hülya Akbiyik, Marco Bonini

The transfer of research results to market-ready products is often a costly and time-consuming process. In order to generate successful products, researchers must cooperate with industrial companies; both the industrial and academic partners need to have a detailed understanding of the requirements of all parties concerned. Academic researchers need to identify the performance indicators for technical systems within a business environment and be able to apply them.

In service logistics today, nearly all standardized mass goods are unloaded manually with one reason for this being the undefined position and orientation of the goods in the carrier. A study regarding the qualitative and quantitative properties of goods that are transported in containers shows that there is a huge economic relevance for autonomous systems. In 2008, more than 8,4 billion Twenty-foot equivalent units (TEU) were imported and unloaded manually at European ports, corresponding to more than 331,000 billion single goods items.

Besides the economic relevance, the opinion of market participants is an important factor for the success of new systems on the market. The main outcomes of a study regarding the challenges, opportunities and barriers in robotic-logistics, allow for the estimation of the economic efficiency of performance indicators, performance flexibility and soft factors.

The economic efficiency of the performance parameters is applied to the parcel robot -a cognitive system to unload parcels autonomously from containers. In the following article, the results of the study are presented and the resultant conclusions discussed.

I. INTRODUCTION

Globalization increases the transportation of goods, implying that efficient processes for goods transportation would be a competitive advantage for companies to have. One possibility to create efficient processes is the automation of process chains, but automation within logistics chains faces one of its largest challenges in the handling of randomly packed goods. If it is not possible to develop a system that handles these tasks using existing methods and technology, research projects are often initiated to tackle the problem. The subsequent transfer of research results to market-ready products is frequently a time and cost consuming process. If researchers keep in mind the performance indicators for technical systems in a business environment, it is possible to achieve this purpose with less effort. This paper therefore focuses on performance indicators for autonomous systems in loosely structured environments within logistics.

The paper begins by identifying the unloading of goods from containers as a task within logistics, still requiring a high amount of manual labor and also being a highly challenging scenario for automation. Subsequently, a study about the qualitative and quantitative properties of transported goods reveals the economic relevance of employing cognitive systems for this task. The key information from a study about potentials, requirements and challenges for robotics in logistics are summarized to derive performance parameters for cognitive robotics systems in logistics applications. These statements are used to obtain performance parameters, which are applied afterwards to the parcel robot – a cognitive system for unloading cuboid goods. The conclusions are summarized in the final part of the paper.

II. AUTOMATION IN LOGISTICS CHAINS

Many different definitions exist for logistics with respect to the angle of view [1,2]. All definitions commonly agree that logistics is the planning and implementing of material flow and information flow within a supply chain.

This paper sets a focus on service logistics. "Service logistics includes the complete planning, controlling, realization and testing of all institution internal and overlapping flow of goods and personnel" [3]. Typical tasks for the material flow within service logistics are palletizing/de-palletizing, commissioning and loading/unloading of carriers. Among these tasks the automation of unloading is the most challenging due to the undefined position and orientation of goods in the carrier.

Typically, goods in containers can be classified by their packaging. They are either transported on pallets, have random packaging, or have standardized packaging. Today, forklifts usually unload palletized goods, further automation cannot promise more economic efficiency. Goods with random packaging may have different dimensions (up to the

Manuscript received July 25, 2011. This work is supported by EU FP7 in the area of Cognitive Systems and Robotics (ICT-2009.2.1) under Grant Agreement 270350.

Prof. Dr.-Ing. Wolfgang Echelmeyer is professor for material flow / logistics at the University of Reutlingen, Alteburgstraße 150, 72762 Reutlingen, Germany, phone: +49 7121-271-3076; e-mail: wolfgang.echelmeyer@reutlingen-university.de

Dr.-Ing. Alice Kirchheim is senior researcher at Orebro University, Studentgatan 1, 70182 Örebro, Sweden

Prof. Dr. Achim L. Lilienthal is professor at Orebro University Studentgatan 1, 70182 Örebro, Sweden

Hülya Akbiyik is researcher at Reutlingen University, Alteburgstraße 150, 72762 Reutlingen, Germany

Marco Bonini is researcher at Reutlingen University, Alteburgstraße 150, 72762 Reutlingen, Germany

size of the container) and different weights (up to the maximum payload of containers) so that automation may not be economically efficient because individual solutions would be necessary. On the contrary, goods with standardized packaging are often consumable goods; their packaging is within a certain bandwidth regarding both their dimension and weight, making automation a good prospect [4]. A study has been conducted to validate these assumptions about the qualitative and quantitative properties of transported goods [5]. The core statements are communicated in the following section.

III. TRENDS IN THE CONTAINER MARKET

Global economy goes hand in hand with world trade and builds the foundation of container traffic. There is a positive relationship between three factors, since an increase of the international economy of 3 % leads to a 6 % increase in world trade and a 9 % increase in container traffic [9]. The growing needs in container shipping caused by an increase in the volumes handled worldwide are justified, considering the recent positive development of global economy after the global financial crisis in 2009.

The container traffic is mostly handled within three major trade routes: the Transatlantic route between North America and Europe, the Transpacific route between North America and Asia and the Europe-Far East route between Europe and Asia. These are the major routes within the container market, since they connect the countries of the triad with 85% of the worldwide container traffic occurring within these routes, as shown in the figure 1 [10].



Figure 1: Trade routes [10].

In order to gain new information about the recent developments in the container market, a study was made among the market leading shipping companies. The study was conducted on European container traffic with its focus on the contents of containers arriving in European ports and the process of unloading of containers.

The main ports within Europe are in Nordrange namely Antwerp, Bremerhaven, Hamburg, Amsterdam and Rotterdam, whereas the study includes the European ports shown in figure 2.



Figure 2: European ports.

The target of the study was receiving internal data concerning the contents of containers, which includes the different types and quantities of goods transported within containers to European ports. The survey was conducted through the use of interviews via telephone and e-mail. All data from the three shipping companies was then analyzed. These companies are within the first ten shipping companies in the world ranking. Apart from the internal political decisions of the interviewed companies, other companies were unable to participate in this study due to the technical limitations of their systems. As a result, more than 13 % of total European imports could be analyzed in detail.

The method used to generate the qualitative and quantitative statements was based on the information regarding the various kinds and quantities of enterprises. Based on this information, the next analysis step took different types of packaging and geometrical forms into account. Beginning with the available literature and interviews with experts, as well as personal experiences and assumptions, the piece goods were given weights and volumes. The quantitative statements arise from the total amount of goods and the previous qualitative analysis.

The following transport packaging is mainly used for the transport of goods in containers and load carriers:

- · barrels, cylinders, canisters
- euro palettes
- boxes, cartons
- · heavy-load carriers, low-load carriers
- sacks, bags
- large bags
- sheds
- rolls, coils

A thorough examination of the aforementioned transport packaging, regarding its relevance for the process of automatic unloading, shows that only sacks and bags, cartons and tires can be considered for this purpose. These correspond to the cuboid, cylindrical and free geometrical forms. After analyzing the imported goods, about 63.8 % of the goods in question, those that have been imported in containers, can be automatically unloaded in European ports. The key result of the report stated that 46.7 % of these goods come in boxes of different sizes. Goods in sacks are made up 15.1 %. The proportional amount of cylindrical forms is limited to 0.05 % [5]. These results are shown in figure 3.

Proportional Distribution of Piece Goods



Figure 3: Main results of the study [5].

Most boxes arriving in Europe come from the Far East, whereas free-shaped goods are mainly imported from South American countries.

The need and feasibility for automatic unloading of containers will be more than justified, considering the imported volumes in European ports measured in Twenty-foot equivalent units (TEU). In 2008, 8.4 billion TEU arrived in European ports that were dedicated for automatic unloading [10]. The potential for such systems is growing with the increasing trade volumes due to globalization and labor division. Within the study it has been attempted to express the amount of imported TEU in the amount of piece goods. Knowing that this analysis is based on several assumptions, there was potentially more than 331,000 billion piece goods unloaded manually [5].

The previous section has identified the economic relevance for automatic systems for unloading containers. A study was conducted because the success of employing such systems in companies is influenced by further factors. The objective is to point out the potential, requirements and challenges for employing robotic systems as an alternative to manual labor.

IV. PERFORMANCE INDICATORS FOR AN AUTONOMOUS SYSTEM

In 2007, an online survey was conducted, interviewing parties involved in the *robotic-logistics* market [6]. The study is continuously updated and will be published again in 2011. Most of the participants are from logistics companies, technology suppliers or research institutions. Within this section, core statements were extracted from the study, which form the basis for obtaining performance indicators for cognitive robotics in logistics applications.

The study states with respect to:

1) Implementing robots: The main reason not to implement robots is the lack of economic efficiency (50%). Further reasons would include products not being market-ready (23 %) and suboptimal basic conditions for robots (23 %).

2) Automation and implementation of robots: More than half of the participants (54%) stated that automation is important for their institution. 67% of the participants plan to implement robots within the next five years or further investigate the idea in more detail. The optimal path to automation should be achieved step by step until reaching complete automation (54%).

3) Barriers, challenges and risks: The participants stated that economic efficiency, cost-benefit ratio and inflexibility are the primary barriers, challenges and risks. Furthermore, a lack of know-how has been highlighted.

4) Places of implementation for robots: Primary places for robot implementation is the unloading/loading and commissioning of goods. One third of the participants evaluated adaptability as major criteria for the implementation of robots.

5) Potential from the implementation of robots: Participants mainly stated three potentials, which are the improvement of process quality, the combination of robots with further technologies (e. g. RFID) and the efficient planning of material flows.

The study shows that the most important point for companies is economic efficiency, which has an influence on the decision of whether to implement automation. Therefore, economic efficiency is one performance indicator. Because the economic efficiency of a system is mainly dependent on the machine performance and its costs, machine performance is a second performance indicator. The flexibility of the system is a third performance indicator, derived from the statement regarding the barriers, challenges and risks. The forth performance indicator is a collection of soft factors including increase of process quality and ergonomics of the working area. So, as a result of the study at least four key performance indicators can be derived:

1) Economic efficiency: The economic efficiency is calculated via a comparison of cost method and a net present value method. For details regarding these methods the interested reader is referred to Götze et al. [7].

2) Engine Performance: The minimum engine necessary to operate the system can be calculated from the parcels per year and per incoming buffer.

3) Flexibility: It is not possible to measure the flexibility of the system in numbers, however estimation of the percentage of goods that can be handled automatically gives guidance.

4) Soft factors: Ergonomics of the working area, safety, process quality.

The first two performance indicators are described in detail in this paper. Three scenarios of the cognitive system parcel robot will be developed exemplarily, so that the performance indicators economy efficiency and engine performance can be analyzed and evaluated in the context of a possible application at logistics companies. Within the scenarios, the efficiency of the system is evaluated depending on engine performance (cycle time), meaning that both performance indicators are covered.

The other two performance indicators, flexibility and soft factors, are difficult to validate and evaluate. On the one hand, this is due to the fact that it is not possible to make statements about the impact of the development of more flexible piece goods on system cost and performance (cycle time). On the other hand, information and figures about costs of down times are held strictly confidential by the companies, so that a realistic evaluation is difficult to result in accurate statements. Therefore, the analysis within this report is based and focused on the first two performance indicators.

V. EVALUATION OF PERFORMANCE INDICATORS FOR AN AUTONOMOUS SYSTEM

The performance indicators are evaluated for an application field within logistics that is still challenging for researchers as is the unloading of goods from containers. This is challenging due to the undefined position and orientation of the goods within the restricted workspace, the very low salary level for workers and the requirements for the short period of time for unloading a container. As a cognitive robot (parcel robot) is already available on the market for the specific case of parcels, performance indicators are applied for the unloading of parcels from containers.

The parcel robot, shown in figure 4, is a 4-axis kinematic system able to unload containers with a rate up to 500 parcels per hour. Cuboid goods in the range of 200 mm to 600 mm for each edge length and a weight of up to 31.5 kg are suited to be automatically unload through this system. The possibility to extend the parcel robot up to 6-axis kinematic, with modular grippers and a telescopic belt conveyor, gives the system the necessary flexibility in order to reach all parcels positions [8].



Figure 4: Parcel Robot

The evaluation is restricted to the first two performance indicators economic efficiency and performance. Three typical types of companies were chosen for the evaluation. The values of the parameters necessary to apply the performance indicators are the number of parcels per year and per incoming buffer (P), the operation time in hours per day (O) and the costs of a worker per hour (C_w) . By using the number of parcels per year, incoming buffer and the operation time per day, the minimum parcels to be unloaded automatically per hour can be calculated by using an average number of 250 working-days per year. Furthermore, the interest on capital is set to 6 %, the cost for the system (C_8) is set to \notin 250,000 and the operating costs for the system (C₀) is set to \notin 4.30 per hour for all three companies. To assure comparable results, all companies have been chosen such that the parcel robot has to unload 330-480 parcels per hour.

All companies are representative within the area of service logistics. Company 1 is located in West Germany has a value of 1,200,000 parcels per year and incoming buffer, operation time is 10 hours and the cost of a worker per hour is \in 39. Company 2 represents a company in East Germany, where the costs for workers is far below those in West Germany. The number of parcels per year and incoming buffer is 1,500,000 parcels, the operation time is 18 hours and the costs of a worker per hour is \in 10. Company 3 is located in Canada. The number of parcels is 2,250,000 parcels, the operation time is 20 hours and the costs of a worker is \in 40. An overview of all scenarios is given in Table 1.

 TABLE 1

 Scenarios for evaluation of performance indicators

Parameter	Scenario1	Scenario 2	Scenario 3		
P [pcs]	1,200,000	1,500,000	2,250,000		
0 [h]	10	18	20		
$C_W[\epsilon/h]$	39	10	40		
$C_s[\epsilon]$	250,000	250,000	250,000		
$C_o[\epsilon/h]$	4.3	4.3	4.3		
I [%]	6	6	6		

The net positive value analysis in Figure 4 shows the difference in pay back period in the three scenarios.



Figure 5: Net positive value analysis.

For company 1, the first scenario, the investment is paid back and becomes profitable in 3 years, for company 2 in 15 years and for company 3 in 1 year and 6 months.





Figure 6: Break Even Point in number of parcels I.

Figure 7: Break Even Point in number of parcels II.

The breakeven analysis in Figures 5 and 6 is shown in the number of parcels. It indicates that the breakeven point is reached at 3,487,000 parcels handled per buffer in the first scenario, 10,022,000 parcels in the second scenario and 3,151,000 parcels in the third scenario.

VI. CONCLUSION

Several conclusions can be drawn from the evaluation of performance indicators in the previous section:

1) The convenience of automatic systems increases rapidly with the number of parcels suited to being unloaded automatically. For this reason companies in the third scenario can benefit from higher profits from an automatic system than companies in the first scenario. Companies in the third scenario have a payback period of 1 year and 6 months, which is lower than the target payback period of 2 years that is considered to be the cut-off in the decision of whether to make the investment. Companies in the first scenario have a payback period of 3 years, which is beyond this limit, but for companies in this scenario the system could be modified in order to save 10-20% of the initial investment. This would shift the payback period very close to the cut-off of the two years.

2) The convenience of automatic systems increases substantially with the growth of personnel costs. For this reason companies in scenario 2 don't seem to be suitable for using automatic systems; the extremely low personnel cost in this scenario pushes the payback period far beyond the cut-off of two years (15 years in scenario 2).

3) As mentioned in point 1, automatic systems could be modified in order to save 10-20 % of the initial investment. This saving could be translated into a 10-20 % reduction in the breakeven point in the number of parcels, but is only true if the modifications to the system will not reduce the number of parcels that are suited to be automatically unloaded.

ACKNOWLEDGMENT

Supported by the European Commission through the Seventh Framework Programme for Research and Development under the Grant Agreement number 270350.

REFERENCES

- [1] P. R. Murphy and D. Wood, *Contemporary Logistics*. Prentice Hall, New York, 10th edition, 2010.
- [2] P. R. Murphy and D. Wood, *Logistics Systems: Design and Optimization*. Springer US, New York (USA), 1st edition, 2010.
- [3] H. Baumgarten. (2011, June 10). [Online]. Available: http://logistics.de/downloads/04/84/i_file_45167/Jahrbuch%202002% 20-%20Dienstleistung.pdf
- [4] A. Kirchheim, Verfahren zur Erkennung von sackförmigen Stückgütern für die automatische Entladung in logistischen Prozessen. Mainz Verlag, Aachen (Germany), 2011.
- [5] H. Akbiyik, A. Kirchheim and W. Echelmeyer, "Latest Trends in the Container Market", in *Bremer Value Reports*, Bremen, vol. 3, 2009

- [6] N. Pfeffermann, W. Echelmeyer and B. Scholz-Reiter, *RoboScan'07*. Springer-VDI-Verlag, Düsseldorf (Germany), 2007.
- U. Götze, D. Northcott and P. Schuster, *Investment Appraisal:* Methods and Models. Springer Verlag, Berlin (Germany), 1st edition, 2007.
- [8] ThyssenKrupp System Engineering. (2011, June 10). [Online]. Available: http://www.thyssenkruppsystemengineering.com/cms/website.php?id=/587/299/en/productrang e/service/roboticlogistics/parcetrobot htm
- e/service/roboticlogistics/parcelrobot.htm
 [9] B. Lemper and A. Hader, "Entwicklungstendenzen des seewärtigen Welthandels", in *Schriftenreihe der Deutschen*
- Verkehrswissenschaften Gesellschaft e.V., Kiel, 8th Conference, 2001.
 [10] MDS Transmodal, Supply and Demand in the Container Trades: Global and Regional Trends and Forecasts, Chester, 2007.

Logistics for Coordinated Distributed Mobile Sensors and the Quest for Performance Measures

Frank Ehlers, Member, IEEE

Abstract—At present dramatic advances are seen in sensor and multi-sensor fusion techniques, in networking and multiagent system design, and in capabilities of robotic platforms in terms of robustness and endurance. Applications for autonomous surveillance appear on the horizon, especially wanted for hazardous or dangerous environments. However, whilst for each subsystem measurements of performance exist (e.g. receiver operating characteristics for sensors, data rate for networks, energy efficiency for a robotic platform), we have to find also a systematic approach on how to construct measures of performance for an entire system of distributed sensing robots. Having such a methodology available would allow the comparison of complete systems which are designed with different emphasis on sensor quality, collaboration tactics and platform capabilities, finally resulting in profound investment decisions. In this paper, we discuss potential similarities between the aforementioned task of organizing teams of sensing robots in a surveillance task and the workshop topic of teams of transport robots in intra-logistics. The common focus is on the construction of performance measures for the autonomously self-organized team of (transporting or sensing) robots.

I. INTRODUCTION

RECENTLY, an interdisciplinary initiative was started. called the Distributed Mobile Sensors Interest Group (DMSIG), which has brought together a team of international researchers, with the objective to promote collaboration among them in the area of Distributed Mobile Sensor Networks for Surveillance Applications, with a current focus on active surveillance systems, e.g. multistatic sonar, radar, or laser scanners, whereby application fields can be indoor, outdoor and underwater. This collaboration is achieved through regular meetings, participation in special sessions at conferences (e.g. [1]), publications in journals (e.g. [3]) and the analysis of common testbeds. For the latter topic, the analysis of common testbeds, as a first step a questionnaire has been sent to the participants and determined the focus of the group to be on the *coordinated* distributed mobile sensors. As a second step a roadmap has been generated to allow the participants to work towards common results. This roadmap is outlined in Section V of this paper.

While generating the roadmap, it became obvious that, without having a method to construct performance measures, the generation of common results is impossible [1]. A discussion of ideas towards generating performance measures in the DMSIG is given in Section II.

Since the focus of the DMSIG is on the *coordination* aspects within a team of multiple robots, a similarity to the objectives of the IROS 2011 Workshop on Metrics and Methodologies for Autonomous Robot Teams in Logistics (MMART-LOG) [4] is already visible. Roughly speaking, by changing 'transportation tasks' into 'information gathering tasks' commonalities between the two initiatives, DMSIG and MMART-LOG, can be further examined on a theoretical level.

Also from the practical or operational point of view, there are similarities between the equipment used in both fields, the Automated Guided Vehicles (AGVs) in MMART-LOG and, for example, the Autonomous Underwater Vehicles (AUVs) as a potential part of testbeds in DMSIG. In Section III, we describe the potential usage of AUVs in underwater surveillance applications in more details. Then, in Section IV, we show that in an operational environment, it might be necessary to program the AUVs to follow specific paths, i.e. to take away degrees of freedom from their autonomous behaviors, which, at the end, makes them more similar to AGVs that act in a structured environment. The reason for not allowing the AUVs to perform autonomously multiobjective optimization within the objective function calculated by a large number of practical constraints (as further discussed in Section IV) is the necessary linkage to human decision makers in military operations who have to know at any point in time with sufficient precision where the AUVs are and what their next actions will be. For example, in the framework of underwater communication this necessary information cannot be transmitted always, hence only pre-planning can generate a feasible concept of use for the surveillance applications. A more general application with communication constraints exists when covert receivers are used which should only very rarely give the target a chance to estimate their positions by intercepting communication signals. In this case of covert receivers, the necessity for pre-planning seems to be especially interesting from the point of view of how to enable coordinated team work. An example from sports is the 'no-look pass' which members of a soccer or basketball team practise in training before they use it in the actual competition [5].

II. THE NEED FOR PERFORMANCE MEASURES

For surveillance systems, from a procurement strategy point of view, a quantitative statement on potential improvement is necessary to justify investments for the introduction of unmanned systems. Three major thrusts to work on distributed sensor network based robotic teams are easily listed: data fusion from distributed multi-sensor

¹Manuscript received June 1, 2011.

F. Ehlers is with the Research Department for Underwater Acoustics and Marine Geophysics (FWG) at the Bundeswehr Technical Center for Ships and Naval Weapons, Maritime Technology and Research in Kiel, Germany (phone: 0049-431607-4279; fax: 0049-431607-4150; e-mail: frankehlers@ bwb.org).

measurements clearly improves the overall receiver operating characteristic, design and maintenance tools of multi-agent teams are available and guarantying disruption tolerant services, the robustness and persistence of robotics platforms is dramatically increasing [1].

An open question is how Distributed Mobile Sensor Networks can be compared, i.e. how to generate procurement decisions for the system with "better" quality (or capability). One systematic way of defining the statement "better" could be to place this research in the framework of game theory for non-cooperatively acting teams. A team of targets is challenging different surveillance teams and their success rate is compared. The expertise needed in this case is on how to organize, evaluate and predict outcomes of such competitions. The role of learning mechanisms for distributed robot systems is then to ensure that optimal parameter settings for the challenged systems are used.

The DMSIG could be a forum capable to initiate collaborations in form of potential future competitions and joint definitions of quality measures for 'smart sensors on cleverly teamed smart platforms'.

III. APPLICATION IN UNDERWATER SURVEILLANCE

In this section, we first introduce the 'sonar equation' which describes sensing performance of active sonar systems. Then, we discuss with the help of the sonar equation why coordinated teams of platforms produce a much higher overall system performance than non-cooperating single platforms. We shortly mention that, for the underwater surveillance, hardware exists to perform autonomous sensing, and conclude this section with the statement that the implementation of the coordination in an efficient manner is needed to fully exploit the high overall performance with a limited amount of participating platforms.

A. Sonar equation

The sonar equation combines in logarithmic units (i.e., units of decibels relative to the standard reference of energy flux density of rms pressure of 1 μ Pa integrated over a period of one second) the following terms:

SE = (S - TL1 - TL2) - (NL - AG) + TS

which define signal excess (SE) where: S: source energy flux density at a range of 1m from the source; TL: propagation loss for the range separating the source and the target (TL1) and the target and the receiver (TL2); NL: noise energy flux density at the receiving array; AG: array gain that provides a quantitative measure of the coherence of the signal of interest with respect to the coherence of the noise across the receiving array; TS: target strength whose value strongly depends on the aspect of the target to the source receiver pair, if the target is a long thin cylinder [6].

For the description of active sonar, the sonar equation has to be applied for the sound path from the source to the target where the received level plus the target strength (TS) is reflected to the receiver. Especially interesting with respect to the control of receiver platforms are the parts of the sonar equation, which depend on the target position (TL, TS) and on target orientation and velocity (TS). In this paper, we only refer to the sonar equation in the noise limited case. A similar formulation is proposed for the reverberation limited case.

B. Coordinated platforms

Looking at the terms of the sonar equation and taking into account that potential targets often have a cylindrical shape, the added-value by using coordinated platforms can be seen at the TS term: If the coordinated platforms can achieve to see the target (independently on its maneuvers) sufficiently frequently in high TS regions, the stealth capability of the target has been taken away. In other words, the aim of the coordination is to generate and hold tracks of a target which cannot be detected by a single platform on its own. Data fusion and tracking algorithms for this scenario are described in [6].

C. Equipment for Autonomous Underwater Surveillance

A survey in the Internet shows that there are plenty of AUV platforms available to provide essential capabilities for an autonomous underwater surveillance. For example, in [8], 112 AUV types from 53 suppliers are listed. This information is from October 2010. The number of AUVs in the field is constantly increasing.

In [9], a system concept for the autonomous multistatic surveillance system is described. Figure 1 is sketching the application of this concept.



Figure 1: Geometry in an example environment (left). Simulation results of teamed AUVs trapping a target: For a slow target (top right) and a faster target (bottom right) [5].

D.Requirements

However, underwater surveillance areas are very large, whilst sensor ranges are rather small. Hence, to cover an entire area with a complete network of uncoordinated AUVs would be too expensive. Therefore, we have to learn how to use the Distributed Sensor System efficiently. In the case of underwater surveillance, we have to learn how to exploit the added-value, as explained above in terms of the TS value. I.e, we need to implement the capability to autonomously generate teams with the available platforms. And we have to be able to compare the performance of different implementations of team behaviors.

IV. LINK TO LOGISTIC PROBLEM

A. Constraints in the underwater surveillance planning

Looking in more details at the capabilities of the AUVS and at the operational constraints, it turns out that the AUVs have to stay within geographical limits. They also have to keep distances to each other. They have to be maintained, also in terms of battery capacities in order to achieve a persistent surveillance. Furthermore, minimal distances to communication devices have to be ensured. Other constraints are generated by the sensors itself which limit the maneuverability of the AUVs. This list of constraints can be easily extended. Of particular importance in this paper are the constraints generated by the need to perform a coordinated surveillance without constantly sending coordinating messages between team members.

	Sensor Teams	Logistic Teams
Objective	collecting information	collecting goods
Adaptability	target behavior unknown	ordering behavior unknown
Constraints	maintenance & communication	maintenance & space management
Uncertainty	own sensor information & positions and plans of team members	delayed information & plans of team members

Table 1: Summary of comparion 'Sensor Team' vs. 'Logistic Team'

B. Information gathering instead of transportation

As discussed before, we might think about a team of sensing robots as if they were a team of AGVs [10] which have the task to 'gather information', instead of transporting goods. We can further interpret the task of 'information gathering' as maximizing the likelihood of receiving high Signal-to-Noise-Ratios at the receivers. In a recent work [7], this has been formulated as an "optimization and search" problem: It has been shown that the likelihood function of the Signal to Noise Ratio (SNR) of the target contact for the ith receiver at the mth sonar ping can be calculated as a function of the terms of the sonar equation (see Section III A). Hence it is possible to choose the best receivers from a given set of receivers [7]. In [2], it has been shown that also the setup of the acoustic sources in the active sonar scenario can be found by optimizing the information predicted to be gathered from a tracked target.

Since these formulations are linked to the actual tracking and data fusion algorithm, which is constantly fed with measurements, and also explicitly contain the target strength term (TS) of the sonar equation, we can be sure that these formulations are able to exploit the 'added-value' of the team work automatically.

Instead of just selecting setups from a small set of possible setups as in [2] and [7], these formulations have to be extended to determine the best sensor positions and movements at future time steps whereby current positions and movements of sensors are given as input into these new algorithms.

When extending the approaches in [2] and [7], answers for particular problems can already be given, like for example in Figure 2 where the the solution of the task described in Figure 1 is plotted for the case that the AUVs keep formation while following the non-cooperative target.



Figure 2: Example for formation flight (as an example for a solution with maximal logistic constraints) and the mathematical solution for the resulting non-cooperative game, generated with the LevelSet toolbox [12].

If we want to implement a realistic underwater surveillance as described before, detailed concepts have to be developed [5]. The question remains open how to implement these concepts, further research is necessary. The framework of the DMSIG seems to be a perfect fit for further investigations.

V. THE DMSIG ROADMAP

For the DMSIG, we aim to generate a framework to allow collaboration between researchers in the field of Distributed Mobile Sensors. Therefore, we have written a roadmap (as part of a recipe, which is outlined in this section).

A. Ingredients to generate a team of coordinated distributed sensing robots

- Sensors which quality is geometry dependent
- Platforms which are spatially distributed and unmanned
- Communication witch is limited for both control & data exchange
- Targets which have a stealth design and are faster

than platforms

- Surveillance area which is a 2D grid whereby sensor quality measures are mapped into cells of this grid
- Constraints which can be started to be listed as
 - not enough sensors being available to cover the surveillance area entirely
 - only a limited period of time given to find targets
 - etc.
- B. Tools to enable collaborative research
 - software testbed which simulates a surveillance scenario,
 - performance measure in order to compare different team behaviors (including the way they are implemented, see [5] for one example of an implementation) by 'competition',
 - systematic approach: There are many free parameters in the testbed. As a first step, leading to the roadmap below, we keep the communication settings and the parameters describing the platforms capabilities fixed, and we elaborate only on the effect of a changing sensor performance for the overall team performance.

C. Instructions (Roadmap)

- 1. Build software testbed with common interfaces
- 2. Select a reasonable number of sensors, reasonable parameters for the bandwidth of communication networks, for the maneuverability of platforms, etc.
- 3. Within DMSIG give a specific setup to each participant (where on the quality of sensors differs).
- 4. Each participant calculates the performance of a non-coordinated behavior.
- 5. Each participant uses "internal competition" to generate the best performing team behavior.
- 6. Collect performance limits from each participant.

D. Aim of the Roadmap

By following the roadmap, the participants of the collaborating researchers can jointly work on finding answers for the following questions:

- Can 'coordination' compensate for having lower sensor performance?
- When and how to implement 'coordination?
- How much 'money' can be saved by using cheaper sensors and boosting the performance by the efficient usage of a 'coordinated' sensor team?

VI. SUMMARY AND CONCLUSION

In this paper, we have discussed similarities between the workshop topic "Metrics and Methodologies for Autonomous Robot Teams in Logistics" (MMART-LOG) and the field of research discussed in the Distributed Mobile Sensors Interest Group (DMSIG).

As an example for a potential field of collaboration

within the DMSIG, we outlined the task of underwater surveillance performed by Autonomous Underwater Vehicles (AUVs). There, we noted that many practical and operational constraints have to be maintained, which, at the end, results in the statement that, in order to avoid problems while handling a multi-AUV team for underwater surveillance, it is recommended to take away many degrees of freedom from a single AUV behavior, letting the maneuverability of the AUVs become similar to the maneuverability of Automated Guided Vehicles (AGVs), used in logistics.

As in logistics, where e.g. many small robots are used to move heavy goods, also in the underwater surveillance only the team behavior is enabling the capability to track and pursuit of stealthy and fast targets.

For the DMSIG, a roadmap has been developed to enable collaboration between researchers and to, in the future, build upon joint results. Since there are fundamental similarities between DMSIG and MMART-LOG, the MMART-LOG workshop gives the opportunity to discuss potential future joint steps ahead.

ACKNOWLEDGMENT

The author thanks all participants of the Distributed Mobile Sensors Interest Group (DMSIG) for their support.

References

- DMSIG participants et al, Prepared Proposal for a Workshop at the ICRA 2012 on "Smart sensors on cleverly teamed smart platforms and how to measure their performance"
- [2] F. Ehlers,"Particle filter for information management in an autonomously operating activated distributed sensor network," Proceedings of OCEANS'07, Aberdeen, June 2007.
- [3] F. Ehlers, D. Sofge, M. Chitre, J. Potter, Call for Papers for a Special Issue of the International Journal of Distributed Sensor Networks: "Distributed Mobile Sensor Networks for Hazardous Applications"
- [4] Call for Papers for the IROS 2011 Workshop "Metrics and Methodologies for Autonomous Robot Teams in Logistics" (MMART-LOG)
- [5] F. Ehlers, A Multi-agent Concept for Multi-sensor Surveillance, Waterside Security Conference WSS2011, Marina di Carrara, Italy, Nov. 2011.
- [6] M. Daun, F. Ehlers,"Tracking Algorithms for Multistatic Sonar Systems", EURASIP Journal on Advances in Signal Processing, 2010.
- [7] D. Orlando, F. Ehlers, Advances in Multistatic Sonar, Book Chapter in "Sonar Systems", INTECH 2011.
- [8] www.AUVAC.org (Oct. 2010)
- [9] Frank Ehlers, A Multi-Agent Concept for Multistatic Sonar, Workshop SDF, INFORMATIK 2010, Leipzig 2010.
- [10] D. Sun, A. Kleiner, C. Schindelhauer, Decentralized Hash Tables for Mobile Robot Teams Solving Intra-Logistics Tasks, Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada, May 2010.
- [11] A. Kleiner, Mapping and Exploration for Search and Rescue with Humans and Mobile Robots, Thesis, 2007.
- [12] I. Mitchell, A Toolbox of Level Set Methods (Version 1.1), UBC CS TR-2007-11, June 2007.

User-Specified Performance Metrics for Autonomous Robots in Warehouse Logistics

Lutz Frommberger, Torsten Hildebrandt, Bernd Scholz-Reiter

Abstract-Autonomous robots rely on a large variety of different skills that all contribute to solving a concrete problem. For any of these skills, performance measures exist. It is not always clear, however, how expressive these measures are with regard to the overall performance of the robotic system in a particular task. In this paper, we discuss performance metrics for autonomous robots on an example from warehouse logistics. We point out that individual metrics for lower-level subtasks might have limited value for assessing the whole system from a logistics perspective, also depending on the task to be solved. Thus, we argue for formalized user-defined performance measures on a high level of abstraction that allow a human operator to assess the things that really matter for the problem at hand. We exemplify this by showing results from an ongoing case-study with a surveillance robot in a simplified warehouse setting.

I. INTRODUCTION

For a logistics company running a warehouse, the use of robots is an economic decision: They are employed if it either pays off monetarily or yields other benefits. In a warehouse context, potential savings can be due to reduced storage times in the warehouse, shorter traverse paths, higher throughput of goods, etc.. Looking at these parameters, it can easily be evaluated if the use of a robotic system pays off in contrast to not using it. This comparison, however, can only be performed after the robotic system has been implemented but a company would want to answer the question of costeffectiveness *before* paying for implementing a new system.

In this paper, we particularly look at surveillance robots. To complete their mission, such robots have to perform several different tasks, such as mapping and localization of the environment, navigation within the warehouse, cooperation of robot teams, identification of goods, detection of storage processes, etc.. For each of these tasks, performance measures can be defined, and strategies can be adapted to optimize these measures. The crucial point here is that the named tasks are not independent. On the one hand, better performance in one task will increase the performance of another task (for example, better object recognition might have a positive effect on mapping). On the other hand, the opposite can be the case, and an improvement in one task can lead to a loss of performance in another; exploration vs. exploitation is a classical example here. Thus, optimizing the overall performance under cost constraints results in a multidimensional optimization problem where the impact of a single system component is hardly tractable.

We discuss the problem of evaluating the performance of a surveillance robot system in a concrete scenario in the field of warehouse logistics. The questions posed are of a general kind, however, and apply to many more problems where multiple (or single) robots are employed in a logistics context.

In the following, we first describe the warehouse scenario we use as an ongoing example in this paper. In Section III, we investigate performance measures both from a logistic and a robotic point of view and exemplify shortcomings of common benchmarks. Then, we argue for user defined abstract performance descriptions in Section IV in order to focus on the relevant information for the task at hand. To exemplify this, Section V reports on an evaluation of process detection within a simplified warehouse scenario before we end with a summary.

II. SCENARIO DESCRIPTION

In this paper, we investigate the problem of measuring the performance of a robotic system by examining a concrete project dealing with the use of surveillance robots (see Figure 1) in warehouse logistics in order to enable semiautomatic logistic optimization. We address the problem of understanding so-called chaotic or random-storage warehouses, characterized by a lacking predefined spatial structure, that is, there is no fixed assignment of storage locations to specific goods. Thus, storage processes are solely in the responsibility of the warehouse operators and basically not predictable-goods of the same type may be distributed over various locations. This makes it a hard problem for people aiming at understanding the in-warehouse processes, e.g., different storage patterns. Knowledge about these storage patterns is, however, of critical importance to understand what is going on in the warehouse, and this is a prerequisite to enable an operator to optimize the warehouse.

In this scenario, we have a solid set of background knowledge, for example, we know that storage processes always follow the same schemata involving *functional zones* within the warehouse. But the concrete details—for example, the concrete locations of these zones—are unknown and can even change over time. Using autonomous robots as a minimally invasive means to observe in-warehouse processes in order to support a logistic optimization of a chaotic

Lutz Frommberger is with the Cognitive Systems Group, Dept. of Mathematics and Informatics, University of Bremen, Germany, email: lutz@informatik.uni-bremen.de

Torsten Hildebrandt and Bernd Scholz-Reiter are with the Dept. of Production Planning and Control, University of Bremen, Germany, email: hil@biba.uni-bremen.de, bsr@biba.uni-bremen.de

Work carried out in this paper was supported by the APF focus project program of the University of Bremen, Grant Nr. 011 600, and the collaborate research centers SFB/TR 8 "Spatial Cognition" and SFB 637 "Autonomous Logistic Systems", funded by the German Research Foundation (DFG).



Fig. 1. A surveillance robot in a warehouse scenario setup

warehouse was proposed in [14]. The information provided by the robot shall allow for posing queries about observed spatio-temporal activities (such as "How often have goods been relocated within the storage zone?") as well as about regions in space (e.g., "Which areas in the warehouse have been used as a buffer zone?"). Answering such queries is an important step towards logistic optimization.

The goal of this project is to specify and detect high-level logistic processes inside the warehouse and to make them available to build a logistic simulation that can then be used for optimization purposes. This optimization is not within the scope of this paper; we restrict ourselves to the process detection based on observations from robots.

III. PERFORMANCE MEASURES

Both logistic systems as well as the robot systems are complex systems and their design involves the proper choice of various design options. Expressed more formally it requires the choice of specific design options for a logistic (robotic) system out of a certain "design space" $d_L \in D_L$ ($d_R \in D_R$). Examples of such design choices for the robotic system can be which mapping algorithms to use, the object recognition to employ, or how many robot units to use. Design choices for a warehouse system are, e.g., the order picking policy used, or the layout of storage areas and whether to use fixed assignments of certain kinds of goods to specific storage locations.

These design options determine the performance characteristics (including cost) of the system. Using $p_L \in P_L$ $(p_R \in P_R)$ to denote the performance vectors of the logistic (robotic) system, we can define a function to derive the mapping from the design space to the "performance space" $f_L: D_L \to P_L$ ($f_R: D_R \to P_R$). To actually determine this mapping and find the potentially complex relationships between design options and performance will typically require simulation, or even physical experimentation for both types of systems.

The semi-automatic, robot-supported optimization of warehouse systems intended in this project can now be seen on an abstract level as a function $g_{\text{opt}} : P_R \to D_L$. The use of the robotic system with performance characteristics $p_R \in P_R$ allows the surveillance of the warehouse system and enables a logistics expert to simulate the system, and find new, potentially better design choices $d_L \in D_L$ for the logistic system.

Therefore we get logistic performances before (d_L) and after (d'_L) the robot-supported optimization. Whether d'_L is now better, and therefore the use of the robotic system and subsequent optimization is successful, depends on the preferences of the company running the warehouse, as the problem is a multi-objective decision problem. Let $\succ: P_L \times$ $P_L \to \{0,1\}$ denote the preference (or utility) function of this company, which is 1 iff the first performance vector is preferred to the second, or 0 otherwise. Then the use of the robot-supported optimization was successful, if $f_L(d'_L) \succeq_u f_L(d_L)$, with $d'_L = g_{opt}(f_R(d_R))$.

A. A logistic view on performance

In general measuring the performance of a logistic system is difficult [5]. Depending on the level of abstraction, time horizon, as well as potential recipients, relevant information differs. Usually an internal and external view on logistic performance can be defined [19], reflecting different conflicting goals of running a logistic system. The external, customer-oriented view focuses on logistic quality, i.e., short delivery times and high adherence to delivery dates, whereas the internal view tries to quantify logistics efficiency trying to ensure cost-efficient operation (examples include capacity utilization, stock levels, turnover frequency). Thus, logistic performance is inherently multi-dimensional. There are, however, attempts to aggregate various performance indicators in hierarchic performance measurement systems [25].

Logistic systems are complex, highly dynamic systems, so insights into the operation and assessing the impact of changes to such systems are often not feasible using analytic approaches, thus requiring simulation models. Using such models allows to optimize certain aspects of the system, such as the order picking strategy in a warehouse [6], either manually by domain experts or automatically using simulation-based optimization [11].

Assessing the success of such optimization is usually evaluated focusing on some low level performance measure and its improvement. Example measures for optimizing warehouse operations are throughput of picking orders, or the average/maximum time till fulfillment of picking orders. This, however, is only a simplification, and performance measures can rarely be treated in isolation. Therefore, to decide which of several options to prefer depends on a user's preferences (denoted with \succ in the previous section) in order to compare otherwise unrelated performance measures (such as costs, and, e.g., order fulfillment time) with each other. Without knowing this preference function only clearly inferior options can be ruled out, i.e., those worse in every dimension, which are therefore Pareto-dominated.

In our scenario, these logistic measures only refer very indirectly to the performance of the observing robot, as it is only providing information for the optimization process and, thus, is not directly involved into the operation of the warehouse. It only provides information that is then used for warehouse optimization later on. To judge the quality of the observations, we need to consider different benchmarks.

B. A robotics view on performance

For the given warehouse scenario, the robots have to follow a complex behavior that tackles many different robotics tasks: The environment needs to be mapped and the robots have to localize themselves in the generated map; the environment needs to be explored continuously because of its dynamics; objects have to be tracked and to be identified; activities happening to those objects have to be recognized and categorized. For all these subtasks, reasonable performance measures exist. In the following, we provide an overview by giving a (non-exhaustive) set of examples.

1) Localization and Mapping Quality: The general approach to evaluating the success of mapping and localization is to calculate an absolute error over the distance between prediction and actual position of the robot or features in the map (the ground truth). For examples, see [20] or [7]. The main problem arising here is that the ground truth might not be easily available. For situations where the precise robot position cannot be retrieved easily for comparison, more semantically driven approaches have been proposed (e.g. [9]). Recent work focuses more on challenges emerging from robot use in the field. For example, [28] proposed suitable benchmarks for 6-D outdoor SLAM. The measures mentioned up to now are designed for grid based maps; for topological maps, an approach is to define an error measure of how well a generated road graph fits into the geometric map, e.g. [26].

2) Multi-robot coordination and exploration: The performance metric for coordinated multi-robot exploration usually is the time in which they manage to finish a coverage task, that is, to map the whole environment (e.g., [4], [24]). Coordination mechanisms are evaluated implicitly with regard to coverage performance. Often, evaluation stops at the point of complete coverage. When it comes to repeated coverage, as in sweeping tasks or our warehouse surveillance task, evaluation usually becomes anecdotal [2] or performance is proven by analytic assurances [10].

3) Object recognition: Object recognition¹ is a field that is intensively benchmarked. To rate the detection success,

methods are usually evaluated over standardized object libraries such as the "Columbia object image library" (COIL) [18] that has been used in various seminal approaches (e.g., [23], [3]), or the Caltech 101 or Caltech 256 data sets [13]. While this ensures very comparable results, it has been pointed out that results relying on these benchmarks can be highly misleading and might lead into very wrong research directions [21]. For many specific tasks, more specific benchmarks are developed, for example for pedestrian detection [8].

4) Activity recognition: Activity recognition or process recognition tasks are usually evaluated counting the detection rate against a human-annotated ground truth. While there had been no standard metric to compare such systems 5 years ago [17] and metrics borrowed from related fields often failed to provide meaningful insights [27], recent developments have led to quite some progress with regard to performance metrics [27]. Especially in the field of human activity recognition, standard benchmarks have been set up [12], and even an "Activity Recognition Challenge" is running this year [1].

C. Shortcomings

We have seen that a multitude of metrics exist for evaluating the performance of a robotic system with regard to specific subtasks. In a robotics application in logistics, however, several of these subtasks contribute to the overall performance. One might be tempted to try to increase the performance of the robotic system by optimizing the performance in any subtask. Of course, this does not work as hoped for, as the subtasks are not independent. In the following, we give two examples.

1) Mapping Quality vs. Exploration: Mapping of the environment is an important component of most robotics applications, as it is needed for navigation and localization of the machine. Assessing the quality of a generated grid map is not a hard problem (see Section III-B.1). But apart from the sensors and algorithms used, the mapping quality is also affected by the navigation speed of the mapping robot, higher speed might easily lead to poorer results. But the question is up to which point a high-quality grid map is needed for the task at hand. Topological approaches with local navigation controllers that rely on less exact data have shown great success and robustness as well [16, e.g.].

Faster navigation and, following from that, a slight decrease in mapping quality will pay off in faster exploration of the environment and this creates more opportunities to detect activities. In surveillance tasks, this might be more important than a detailed and exact grid map, so maximizing the performance of the mapping system might be the wrong idea at some point.

2) Activity Frequency vs. Activity Diversity: When surveying the warehouse, we would like to detect at the best *all* activities involving the movement of goods that happen in there—which is not realistic, unless we place a stationary robot everywhere. The robots have to move within the warehouse to find activities and, thus, they will miss some. Let us assume our measure would be to count the percentage

¹Object recognition is not a serious issue in the given warehouse scenario, as we rely on uniquely identifiable features such as RFID tags.

of detected activities over time. Then, the best strategy according to this measure would be that a robot would remain at a location where many activities are going on (perhaps the entrance, where all goods will be located when entering the warehouse). This extreme case would create a very clear picture of this one location (a *hotspot*), while other locations remain unexplored. In this case, we will miss important information. For example, it will remain unclear whether some storage location is not used at all or rarely used—and under which circumstances. This distinction might be a relevant detail for the optimization purpose.

IV. RETRIEVING ESSENTIAL INFORMATION

A. Describing processes on different levels of abstraction

The difficulty when assessing different performance measures in the overall scenario is how relevant the performance metrics are for the overall process. On the lower sensory level, this question is hard to answer. As argued for in Section III-C, an improvement with regard to one measure might even lead to a performance drop for the overall system. The important question is: How relevant is the measured behavior for the targeted outcome?

To be able to utilize the acquired sensory input coming from the robot in a reasonable way, the sensory data is subject to interpretation, leading to different levels of abstraction for the perceived data. Measured distances are transformed into a map, dynamic changes within the map are interpreted as activities, certain patterns of activities have a specified meaning as an abstract concept-for example, they can be interpreted as a redistribution of the good within the warehouse. Such higher-level concepts allows an expert in the logistics domain a much clearer assessment of how relevant the measured parameters are. While, for example, it is doubtful what an increased range of sensor perception will contribute to the overall success of the system, the percentage of correctly detected redistributions in a warehouse can much easier be related to its importance for the overall performance. Clear abstract descriptions of high-level processes especially become important when human experts are included in the assessment process.

In the framework given in Section III, changing one or more options in the flat vector of design options (d_R) in the hope to improve the overall logistic performance P_L usually results in cumbersome trial-and-error experiments with a following long simulation phase. Instead of doing so we argue for having (limited) set of design options that are integrated to contribute to higher level concepts that become part of the performance space P_R .

B. User-defined Queries for Individual Performance Assessment

Depending on the task at hand, *different* processes become relevant. While it might once be the redistributions we are interested in, it might also be the information about whether different goods are moved out of the warehouse together. Thus, it is reasonable that an operator can specify relevant in-warehouse processes according to his needs. For



Fig. 2. The idealized warehouse scenario used for evaluation of the symbolic process detection approach

the described warehouse surveillance task, Kreutzmann et al. show how this can be done by using a formal language [15], namely linear temporal logic (LTL) [22]. For example, without going into details in LTL, a *redistribution* operation was specified like this:

$$\operatorname{Redistribution}_{G,L_i,L_j} = \operatorname{at}(G,L_i) \wedge \operatorname{in}(L_i,S) \rightarrow (1)$$

$$\diamond \left(\operatorname{at}(G,L_j) \wedge \operatorname{in}(L_j,S)\right)$$

This defines that we have a redistribution when a good G is observed at a location L_i (at (G, L_i)) and this location lies within the storage zone S (in (L_i, S)) and at some point in time later (specified by the modal operator \Diamond) G is observed at another location L_j (at (G, L_j)) which also lies within the storage zone (in (L_j, S)).

More abstract processes specified like this can then also be used for evaluation by judging the success with regard to exactly this process with a metric specified by a human operator, that is, it becomes part of the performance space P_R of the robotic system. Various different specifications of processes and measures can be specified and evaluated concurrently (e.g., having both measures for exploration and exploitation) such that it is, similar to the utility function \succ , the system operator's preference which measures to optimize in order to achieve a better overall performance.

In the following evaluation, we will introduce the concept of *histories* of goods, a concept that is not really relevant for the understanding of the warehouse but has been specified as a measure to integrate different facets of the robot's performance.

V. AN EXAMPLE EVALUATION WITHIN THE WAREHOUSE SCENARIO

In this section we report on a first experimental evaluation performed for the LTL-based process recognition approach, as described in [15]. The goal was to measure the success of the proposed process detection mechanism in an idealized warehouse scenario set up in the lab (see Figure 2).

In the experiment, cradles were moved around in the warehouse between pre-defined zones while the robot was



Fig. 3. Evaluation results showing the rate of correct history detections. Diagram taken from [15].

steered within this scene and mapped the position of the cradles which it could identify by unique optical tags attached to them. The processes to be detected had been specified in LTL, as described in the previous section. Observations made by the robot's sensory system were abstracted and translated to symbolic expressions (like at(...) or in(...)), and a process was detected if we can find a model (based on these symbolic observations) that satisfies the corresponding formal specification.

The question was now what to measure when evaluating the approach. An obvious metric would be to count the percentage of correctly detected in-warehouse processes. However, this would veil the information which processes had been detected well and which not. A fairly high percentage of correctly detected observation could eventually camouflage the fact that one particular process had hardly been correctly observed at all. So the evaluation considered correctly observed histories of goods. The history of a good is the sequence of different processes that involves one particular good. Looking at histories ensures that the flow of goods over the whole warehouse is monitored and assessed. A correctly determined history ensures that all kinds of processes can be recognized and the whole spatial range of the warehouse is covered. Histories can also be formulated as LTL formulas. A history is correctly identified if temporal order and number of processes match the ground truth.

Figure 3 shows results from the experiment under different amounts of background knowledge. In particular, the success was evaluated with, without, and with only partial knowledge of the functional regions. Also, the performance of the clustering algorithm that translates measured positions of goods automatically to a smaller, discrete set of qualitative abstract locations was compared to a clustering algorithm where the cluster centers had been (reasonably) pre-defined. The results reveal the major problem of the implementation so far: with automatic clustering, detection rates of histories significantly dropped. Especially, this problem became evident for the case of only knowing two of the functional regions (entrance and outlet), which is a most realistic assumption in this scenario. With pre-defined clustering, the proposed algorithm could infer the knowledge about regions very well (detection rate only dropped from 82.9% to

75.4%), but with automatic clustering it dropped from 58.1% to 43.1%. The used clustering algorithm was known to be quite straight-forward, but it was assumed that its quality was sufficient. The given experiment disproved this assumption by evaluating its impact on high-level process detection. Thus, the evaluation approved one of the main goals of the approach: to show that complete knowledge about the whereabouts of the functional regions is not really necessary and can be deduced by formal reasoning—but only with a sufficiently high quality of the clustering method. Improving this quality is a reasonable next step to improve the overall performance of the system.

VI. SUMMARY

In this paper, we have investigated how to evaluate the performance of a surveillance robot system in order to gather information for logistic optimization a warehouse scenario. Several options exist to improve the robot's behavior in certain subtasks, and all of the subtasks have individual performance measures. The impact of different design options to overall logistic performance measures is hard to understand, as they mutually affect each other. We have exemplified shortcomings of current robotics metrics with regard to their contribution to a complex system and argued for an integrating approach where users can specify performance metrics on a higher level of abstraction such that they can be related to the overall performance more easily. Finally, we have exemplified this by showing evaluation results from an ongoing case study that was able to identify shortcomings in one particular part of the system by evaluating higher-level concepts.

REFERENCES

- Activity recognition challenge. http://www. opportunity-project.eu/challenge.
- [2] M. Ahmadi and P. Stone. A multi-robot system for continuous area sweeping tasks. In *Proceedings of the IEEE International Conference* on *Robotics and Automation (ICRA)*, pages 1724–1729, 2006.
- [3] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24:509–522, 2002.
- [4] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 476–481, 2000.
- [5] Garland Chow, Trevor D. Heaver, and Lennart E. Henriksson. Logistics performance: Definition and measurement. *International Journal* of Physical Distribution & Logistics Management, 24(1):17–28, 1994.
- [6] R. Dekoster, T. Leduc, and K. Roodbergen. Design and control of warehouse order picking: A literature review. *European Journal of Operational Research*, 182(2):481–501, 2007.
- [7] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1322–1328, 1999.
- [8] P. Dollar, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: A benchmark. In *Proceedings of the IEEE Computer Society Conference* on Computer Vision and Pattern Recognition, pages 304–311, 2009.
- [9] T. Duckett and U. Nehmzow. Mobile robot self-localisation and measurement of performance in middle-scale environments. *Robotics* and Autonomous Systems, 24(1-2):57–69, 1998.
- [10] Yehuda Elmaliach, Noa Agmon, and Gal Kaminka. Multi-robot area patrol under frequency constraints. *Annals of Mathematics and Artificial Intelligence*, 57:293–320, 2009.
- [11] M. Fu. Simulation optimization. In Proceedings of the 2001 Winter Simulation Conference, pages 53–61. IEEE Press, 2001.

- [12] L. Gorelick, M. Blank, E. Shechtman, M. Irani, and R. Basri. Actions as space-time shapes. In *In ICCV*, 2005.
- [13] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. Technical report, California Institute of Technology, 2007.
- [14] Torsten Hildebrandt, Lutz Frommberger, Diedrich Wolter, Christian Zabel, Christian Freksa, and Bernd Scholz-Reiter. Towards optimization of manufacturing systems using autonomous robotic observers. In Proceedings of the 7th CIRP International Conference on Intelligent Computation in Manufacturing Engineering (ICME), June 2010.
- [15] Arne Kreutzmann, Immo Colonius, Lutz Frommberger, Frank Dylla, Christian Freksa, and Diedrich Wolter. On process recognition by logical inference. In *Proceedings of the 5th European Conference on Mobile Robots*, September 2011.
- [16] Benjamin Kuipers. The spatial semantic hierarchy. Artificial Intelligence, 119:191–233, 2000.
- [17] D. Minnen, T. Westeyn, T. Starner, J. Ward, and P. Lukowicz. Performance metrics and evaluation issues for continuous activity recognition. *Performance Metrics for Intelligent Systems*, 2006.
- [18] S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (COIL-20). Technical report, Department of Computer Science, Columbia University, New York, 1996.
- [19] Peter Nyhuis and Hans-Peter Wiendahl. Fundamentals of Production Logistics: Theory, Tools and Applications. Springer, 2008.
- [20] S. Oore, G. E. Hinton, and G. Dudek. A mobile robot that learns its place. *Neural Computation*, 9(3):683–699, 1997.

- [21] Nicolas Pinto, David D. Cox, and James J. DiCarlo. Why is real-world visual object recognition hard? *Computational Biology*, 4(1), 2008.
- [22] Amir Pnueli. The temporal logic of programs. In Proceedings of the 18th Annual Symposium on Foundations of Computer Science (FOCS), pages 46–57, 1977.
- [23] M. Pontil and A. Verri. Support vector machines for 3D object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(6):637–646, 1998.
- [24] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the National Conference on Artificial Intelligence* (AAAI), pages 852–858, 2000.
- [25] Rien Ploos van Amstel and Guido D'hert. Performance indicators in distribution. *The International Journal of Logistics Management*, 7(1):73–82, 1996.
- [26] Jan Oliver Wallgrün. Qualitative spatial reasoning for topological map learning. Spatial Cognition and Computation: An Interdisciplinary Journal, 10(4):207–246, 2010.
- [27] J.A. Ward, P. Lukowicz, and H.W. Gellersen. Performance metrics for activity recognition. ACM Transactions on Intelligent Systems and Technology (TIST), 2(1):6, 2011.
- [28] O. Wulf, A. Nüchter, J. Hertzberg, and B. Wagner. Benchmarking urban six-degree-of-freedom simultaneous localization and mapping. *Journal of Field Robotics*, 25(3):148–163, 2008.

Future of Material Handling – modular, flexible and efficient

Kai Furmans, Christoph Nobbe, Melanie Schwab

Abstract – This article discusses the limiting factors for material handling systems and the consequences. We also present some guidelines for the design of future generations of material handling systems which allow to overcome these limitations. As an example two new systems, which have been consistently developed to fulfill these guidelines, are introduced.

I. INTRODUCTION

The demands on material handling systems are changing. Rather than level of automation, cost effectiveness and maximum throughput, other characteristics are more important today: flexibility, reconfigurability and high availability. In every manufacturing business the diversity of products is dramatically increasing. The introduction of hybrid engines in the automobile industry alone doubled the range of products in some of the companies, for example.

Due to virtual development, the time needed to introduce a new model in the market has been shortened rapidly, also shortening the product life cycle itself. These shortened planning horizons and smaller amortization periods result in smaller investments. The customer request for more individual products leads to mass customization, which leads to smaller quantities in almost every sector of the consumer market.

If one looks at different press reports during the economic crisis, one can clearly see a demand for flexibility. At the end of 2009 it was reported that the situation for the component suppliers in the car sector is getting worse (München, dpa-AFX). In the beginning of 2010 you could read: Car production in Germany further decreasing (Frankfurt/Main, ddp), whereas in the middle of the year it was reported, that the global car production will reach the same level as before the crisis (Frankfurt, dpa-AFX).

Manuscript received July 15, 2011. This work was supported in part by the Department of Science, Research, Culture and Arts, Baden-Wuerttemberg, Germany (AZ 32-729.85/127)

Kai Furmans is Professor at Karlsruhe Institute of Technology (KIT) and head of the Department for Material handling and Logistics (IFL). Kaiserstrasse 12, 76131 Karlsruhe, Germany. (e-mail: kai.furmans@kit.edu)

Christoph Nobbe is with the Department for Material handling and Logistics (IFL) at the KIT, (phone: +49-(0)1721-608 48627, fax: +49-(0)1721-608 48629, e-mail: christoph.nobbe@kit.edu)

Melanie Schwab is with the Department for Material handling and Logistics (IFL) at the KIT, (e-mail: melanie.schwab@kit.edu)

For most systems it was common to run for many years in the same configuration, whereas now it is often needed to change the layout and readjust it to new requirements after a couple of years - sometimes even faster. Although the industry has implemented hundreds of successful automated material handling systems in the past, there are signs that business is or will be getting more difficult. For instance, a survey in Germany showed that users of automated material handling systems or are planning to "deautomate" significantly, since the systems are too inflexible in relation to perceived changes in the requirements and processes [1]. Although this study is now somewhat dated, we believe its underlying message is still valid.



Fig. 1: "Reduction in automation in several production areas, according to ISI" [1]. Statistics indicate the percentage of respondents who were planning or who had already implemented a reduction in automation.

According to Furmans, Schoenung and Gue [2] the inflexibility is due to the fact that an adaptation of automated systems is very costly and risky. The reason for this is, among other things, that the functionality needed to provide the logistical services is spread over various levels and disciplines. The total function of such a system is also provided by interacting software and hardware, which are themselves spread out over many different levels and layers. An automatic adjustment of the structure to new tasks and load models, or at least an easy adaptability, is not intended with these systems. With a change in the procedures today, all these levels are typically modified by several different experts.

To cope with the fast changing demands of the markets, new material handling systems need to be more flexible and alterable. An approach, that allows to reconfigure the system "on the fly", without the loss of operational readiness, costly start-up periods and tests, is the only economical solution. Scalability in size, throughput and in the nature of the transported goods is therefore essential, in order to efficiently adapt to changing performance requirements. Ultimately, the setup, adoption and usage must be so simple that new automated systems can compete in flexibility and reusability with non automated systems. In order to achieve this goal, such systems must have characteristics similar to those in modern computing systems. Installation time and complexity has to be reduced dramatically, in order to be cost efficient [2]. Material handling systems should be set up as a plug and play solution. Once the system is physically configured, everything is done. Components can be added by simple insertion. To operate resource efficient, the system needs to be economical (therefore the ratio of weight and load needs to be better, transportation routes need to be shorter and unnecessary work needs to be eliminated). Through easy reusability of the system fewer resources for manufacturing the equipment are used [3]. To meet these demands it is advisable to build the system with small scale units, which can be combined in a very variable manner to form a complete system. This way it would be possible to react flexibly on different demands in throughput and layout. Equipped with a decentralized control, the number of participating units in a system could be easily changed without the need for reprogramming.

II. KARIS – AN INNOVATIVE SYSTEM WHICH CAN MEET THE NEW DEMANDS

KARIS (a German acronym for Kleinskaliges Autonomes Redundantes IntralogistikSystem which can be translated as small scale, autonomous, redundant material handling system) represents a realization of such a standardized, scalable, decentralized, reusable and resource efficient material handling system. The system consists of identical basic elements which are interacting and cooperating with each other to form a flexible system. A basic element is modularly constructed, and can therefore be adapted to different specifications. It is organized in three layers, of which the lower one includes the holonomic drive unit in order to give a maximum movability. The element is able to reach speeds up to 2m/s and transport loads up to 100kg (250kg are planned). The middle layer contains the safety relevant sensors and provides space for a lifting unit, allowing the element to adapt its height to dock on different levels of the peripheral installations or to

raise the work piece directly to working height. The top layer of the element is fitted with a multidirectional conveying unit, for a wide range of applications.



Fig. 2: "KARIS element"

A. Set-up

A special feature of KARIS is the free navigation, which is not relying on a dedicated infrastructure for navigation. An arrangement of landmarks such as radio transmitter, reflectors, magnets, floor markings or RFID-tags, is not needed. The corners and faces which are present in every environment, such as walls, columns or machinery and equipment, are used to navigate. For startup, an initial map is captured with the integrated laser scanners while controlling the device manually. This raw version then is enriched with additional information of the systems environment. Examples are footpaths, restricted or hazardous areas. Other information can be provided as well. Connection points to existing conveying installations, such as an automated storage and retrieval system or other areas, are defined as transfer points and therefore serve as interfaces for the material flow. If the material flow needs to be changed, due to new models for example, transfer points can be changed flexibly to meet the new demands. It is only necessary to transmit their new coordinates, which can be done during normal operation. The decentralized approach allows to add and to remove elements from the system and therefore adapting the actual throughput. Since the path planning and the coordination among each other is done by the elements themselves, there is no timeconsuming reprogramming of the system. This enables the operator to use more elements in high-times to cushion the peeks in the needed performance.

B. Cost efficient

The standardized module structure of a basic element allows the compensation of the breakdown of single units, by just replacing it with a working one. This ensures the redundancy for the system. Building costs of the elements can be reduced by the high quantities of standardized elements, since no customizations are needed. The fact that the availability of the system is part of the system structure allows to have a lower availability for the single elements, which will reduce the building costs further compared to today's AGVs.

C. Application fields

The application opportunities for these new Material handling systems like KARIS are promising. Not only is it possible to replace existing systems, it is also possible to find new areas of use for automated systems. A standard task is the transport of different cases one by one without having to install a cost intensive material flow system. This way it is possible to bring automation to fast changing productions and to smaller companies, which would normally rely on manual labor. Since it is possible to share routes with workers, no extra space for conveyors or streets for forklifts are needed. The material can be brought directly to the workstation, or the KARIS-element itself can serve as a flexible workstation.

In a picking area the standard procedure today is, that a worker has to push a cart, stacked with handhelds, printers or other electronic equipment, boxes and the orders which can easily sum up to over 100kg per cart. If a pick & pack strategy is used, the KARIS-element can travel alongside the worker, providing him with all the needed information, the order can be picked directly into the shipping box on the element and as soon as the order is complete it is directly transferred to shipping. Bigger orders can be processed by using more than one element providing the worker with more shipping boxes. Since the worker does not have to bring the boxes to a drop off point and new elements provide the workers with the needed boxes, a lot of time of walking around can be saved. Another strategy is that the elements maneuver freely to the picking point in the area and the workers can travel on fixed routes to serve the waiting elements. In both cases the work load for the person is drastically lightened, which is an important factor since the average age of the workforce is increasing worldwide.

III. GRIDFLOW: AN APPLICATION FOR HIGH-DENSITY STORAGE SYSTEMS

In the last years, new shuttle based systems pushed into the market to bring flexibility and energy efficiency to storage systems. Even though these systems are developing in the right direction, they do not pose a real flexible and cost efficient alternative to the existing systems, since the costs for the high rise racks are still high and it is difficult to change the layout. Often, to keep pace with the fast changing requirements, the best way is to rent a storage building as it is needed. In these cases it is not possible to install an expensive storage solution. Flexible systems with very little installations are needed that can also be dismantled when the facility is no longer in use. For many companies, therefore, the only option is to buy or rent a forklift and store the pallets on the factory floor in block storages. These storages for pallets do not contain aisles and therefore offer a very high density which makes them surface and cost efficient. The access is however limited to loading units at the edges. To retrieve a pallet from within the block several others have to be picked up, moved away, buffered elsewhere and restored after the retrieval of the desired pallet.

A. Idea of the GridFlow system

To avoid this limited access Gue and Kim [4] propose puzzle-based storage systems. These automated storage systems base on the 15-slide puzzle (see Fig. 3) and consist of a grid of loading units and conveyors with at least one empty location. Each location is equipped with a conveyor and thus each pallet can be moved to all 4 directions in the horizontal level at any time, if the neighboring location is empty. Taylor and Gue analyzed the retrieval behavior of puzzle based storage systems [5], [6]. Gue and Furmans showed that such a system can also be operated by a decentralized control [7].

2	7	10	15		
5	13	6	12		
14			8		
11	(4)	9	3		

Fig. 3: "15-puzzle" [4]

Due to the conveyors the system flexibility to layout changes is limited and requires a big investment. We therefore present the concept GridFlow replacing the conveyors with robots (automated guided vehicles) (see Fig. 4).



Fig. 4: "GridFlow system with pallets and robots"

Alfieri et al. [8] propose a similar system for racks moved by automated guided vehicles.

The robots make the GridFlow system very flexible to layout and throughput changes. Furthermore the system does need very little infrastructure and is easy to put into operation.

A system with one open location (escort) and one robot

offers the highest possible density (n-1)/n, where *n* is the number of location. By increasing the number of escorts and robots, retrieval times of the system can be decreased while the density decreases as well. Thus the system can be adjusted to the required throughput while ensuring the highest possible density for this case.

To determine the number of escorts and robots for a certain throughput, we need to calculate the retrieval times depending on the number of open locations, the number of robots, the layout of the grid and the position of the Input-/Output-Points (I/O-Points).

We will here consider a system with one escort, one robot and one I/O-Point.

B. Modeling the system

The model consists of a grid which is *n* locations wide and *m* locations long and holds a total capacity of C = n x*m* locations. Each location can hold one loading unit. A location within the grid is identified by two coordinates, were the first coordinate stands for the horizontal axis and the second one for the vertical axis. Both axes start in the lower left corner with position (1, 1) (see Fig. 5).

The position of the retrieval unit e.g. the loading unit to be retrieved is given by (i, j) and that of the I/O point by (k_i, k_j) .



Fig. 5: "Model of the GridFlow system"

C. Basic movements

All movements considered here start and end with the vehicle being in the escort.

3-move: A 3-move swaps the position of the escort with the position of an adjacent loading unit (see Fig. 6).





described by Gue and Kim with conveyors. In a system with one robot it requires nine robot movements and is thus called a 9-move (see Fig. 7).



15-move: The 5-move described by [4] to move a loading unit by one location in a direction opposite to the location of the escort becomes a 15-move in this system. In the beginning of this move the escort and the target location are on different sides of the retrieval unit. Therefore the 15-move will be done in the same direction as the preceding move (see Fig. 8).



D. Retrieval process

As a robot will always end a retrieval in the I/O-Point location, and the escort will move with the robot we assume the next retrieval to start with the escort and the robot being in the I/O-Point location.

In the first phase the escort has to be moved from the I/O-Point to the retrieval unit. The target position of the escort depends on the arrangement of retrieval unit and I/O point within the grid as the number of 9-moves to be performed in phase 2 is maximized (see [4]). The number of 9-moves is maximized as with those moves the retrieval unit can be moved with less robot movements than with 15 moves.

The second phase starts with a 3-move of the retrieval unit by one location. Next the vehicle will perform as many 9-moves as possible and then clear the remaining distance with 15-moves. According to this we developed the algorithm for the optimal retrieval time as follows:

Algorithm for retrieval of one loading unit with one robot and one escort:

(1, f): position of the retrieval unit after the last step

1: \geq if $(l \geq k_i \text{ and } |i-k_i| \geq |j-k_i|)$ then 2: \geq Move the escort to (i-1, i)3: \geq Do a 3-move to the left 4: $\flat i i$ ($j \ge k_i$ and $|i \cdot k_i| \ge |j \cdot k_i|$) then 5: |>|>> Do 9-moves down and to the left until $l = k_i$ and $(l = i - (|j - k_i| + 1) \text{ or } l = k_i)$ 6: |>|>> Do 15-moves to the left until $l = k_l$ 7: \>\>\> break 8: \>\> else 9: |>|>> Do 9-moves up and to the left until $j' = k_i$ and $(l = i - (|j - k_i| + 1) \text{ or } l = k_i)$ 10: |>|>> Do 15-moves to the left until $l = k_i$ 11: \>\> break 12: \triangleright else if $(i < k_i \text{ and } |i - k_i| \ge |j - k_i|)$ then 13: \geq Move the escort to (i + 1, j)14: \geq Do a 3-move to the right 15: |>|> if $(j \ge k_i)$ then 16: |>|>|> Do 9-moves down and to the right until $l = k_l$ and $(l = i + (|j - k_i| + 1) \text{ or } l = k_i)$ 17: |>|>|> Do 15-moves to the right until $l = k_i$ 18: \>\>\> break 19: \>\> else 20: |>|>>> Do 9-moves up and to the right until $j = k_i$ and $(l' = l + (|j - k_i| + 1)$ or $l = k_i$ 21: |>|>> Do 15-moves to the right until $l = k_i$ 22: $\langle \rangle \rangle \rangle$ break 23: \flat else if $(j > k_i \text{ and } |i - k_i| < |j - k_i|)$ then 24: \geq Move the escort to (i, j-1)25: \geq Do a 3-move down 26: $\geq if (i \ge k_i)$ then 27: |>|>>> Do 9-moves to the left and down until $l = k_i$ and $(j = j - (|i - k_j| + 1) \text{ or } j = k_j)$ 28: |>|>> Do 15-moves down until $J = k_i$ 29: \>\> break 30: \>\> else 31: |>|>>> Do 9-moves to the right and down until $l = k_l$ and $(j' = j - (|j - k_j| + 1) \text{ or } j = k_j)$ 32: |>|>>> Do 15-moves down until $f = k_i$ 33: \>\>\> break 34: \triangleright else if $(j < k_i \text{ and } |i - k_i| < |j - k_i|)$ then 35: \geq Move the escort to (i, j+1)36: \>\> Do a 3-move up 37: $\geq if (i \ge k_i)$ then 38: |>|>>> Do 9-moves to the left and up until $l = k_i$ and $(j = j + (|i - k_i| + 1) \text{ or } j = k_i)$ 39: |>|>> Do 15-moves up until $l = k_i$ 40: \>\>\> break 41: \>\> else 42: |>|>>> Do 9-moves to the right and up until $l = k_i$ and $j = j + (|i - k_i| + 1)$ or $j = k_i$ 43: |>|>>> Do 15-moves up until $f = k_i$ 44: \>\> break

If movements are performed according to this algorithm

we can calculate the retrieval time (RT) in time units by (1). One time unit represents the time needed to perform one robot movement. We verified this formula by simulating various scenarios with self-developed Visual Basic Code.

$$RT(i, j) = \begin{cases} 18 \cdot |\mathbf{i} - \mathbf{k}_{i}| + 6 \cdot |\mathbf{j} - \mathbf{k}_{j}| - 16, & |\mathbf{i} - \mathbf{k}_{i}| > |\mathbf{j} - \mathbf{k}_{j}| \\ 0, & |\mathbf{i} - \mathbf{k}_{i}| = |\mathbf{j} - \mathbf{k}_{j}| = 0 \\ 24 \cdot |\mathbf{i} - \mathbf{k}_{i}| - 10, & |\mathbf{i} - \mathbf{k}_{i}| = |\mathbf{j} - \mathbf{k}_{j}| \neq 0 \\ 6 \cdot |\mathbf{i} - \mathbf{k}_{i}| + 18 \cdot |\mathbf{j} - \mathbf{k}_{j}| - 16, & |\mathbf{j} - \mathbf{k}_{j}| < |\mathbf{i} - \mathbf{k}_{i}| \end{cases}$$
(1)

The expected retrieval time E(RT) can be calculated as the average of the retrieval times of all grid locations (2).

$$E(RT) = \frac{1}{C-1} \sum_{i=1}^{n} \sum_{j=1}^{m} RT(i, j) \quad (2)$$

E(RT) depends on the grid layout which is defined by the aspect ratio of the grid and on the location of the I/O-Point. To give system design rules we have to obtain the aspect ratio and I/O-Point location resulting in the lowest expected retrieval time.

We analyzed different scenarios and found the I/O-Point to best be located in the middle of the longer side of the grid in all of these scenarios. Fig. 2 gives an example of the expected retrieval times depending on the location of the I/O-Point for a grid capacity of 1600 locations. Every bar gives the expected retrieval time for the I/O-Point being in this location (see Fig. 9).



Fig. 9: "Expected retrieval times in dependence of the I/O-point location"

The analysis of the aspect ratio is aggravated by the fact that for every grid capacity there is only a discrete number of possible aspect ratios. Therefore we are not able to compare every aspect ratio for every grid capacity.

However, data from grid capacities smaller than 2000 loading units suggest that an aspect ratio of 2:1 results in the lowest expected retrieval time. For a given capacity, we therefore recommend to use an aspect ratio that is as close to 2:1 as possible.

There are grid capacities which result in very high aspect ratio and in prime number cases in a one-dimensional layout. Those extreme cases can easily be avoided in industry application.

We derived the minimum expected travel time for each

capacity smaller than 2000 and analyzed the behavior of the expected retrieval time in dependence of the grid capacity. The expected retrieval time of applicable cases tends to form a linear function for capacities bigger than 500 loading units (see Fig. 10).



As we are only able to derive the best location of the I/O-Point and the best aspect ratio from experiments, we can not give design rules but design guidelines. Those have to be validated by a mathematical optimization of the expected retrieval time which will be part of our further research.

The GridFlow system might open up opportunities in different fields of material handling providing the possibility for fast and easy setups and reconfigurations. The sorting and buffering process in cross docking centers might be an example. Robots could transport the loading units to buffering stages or directly to the outgoing ramp giving the opportunity to automate a process which is today done manually due to the needed flexibility.

Further research needs to work on different application scenarios and their evaluation. Especially strategies for the cooperation of multiple robots need to be developed and quantified.

IV. CONCLUSION

Ever faster changing requirements put today's material handling systems to their limit. The research community agrees on guidelines for new material handling systems. New approaches have to be decentralized, modular and scalable to offer the needed flexibility.

To preserve the competitiveness of our industry we need these new systems as well as efficient strategies to use them. The presented systems KARIS and GridFlow are two systems designed according to the mentioned rules and can work as good examples for future material handling systems. With these approaches new installations do not need to be oversized from the beginning. They could grow according to the company's growth. The evolution of new systems has just begun and there are various opportunities research can adopt and continue.

However, we have to make sure that the developed systems are economically interesting. To decide on this, it is possible to calculate monetary criteria such as investment costs, personal costs or surface costs. Especially in comparison with today's system, the aspect of flexibility is very important but hardly measurable. This is why we need to define how much this flexibility is worth to us and the research community should think about ways to measure it.

REFERENCES

- [1] Fraunhofer ISI, Zukunftsstudie des Fraunhofer ISI, VDI-
- Nachrichten, 35:3, 2001
 K. Furmans, F. Schönung, K. R. Gue, *Plug-and-Work Material Handling Systems*, 2010 International Material Handling Research Colloquium, Milwaukee, USA, 2010
- [3] C. Nobbe, D. Berbig, VDI-IFL Sommerseminar 2010, Wege zu einer verantwortlichen Ressourcenverwendung in der Logistik, Ressourcen länger und flexible nutzen, ISSN Nr. 2190-6599
- [4] K. R. Gue, B. S. Kim, *Puzzle-Based Storage Systems*, Naval Research Logistics, 54(5):556-567, 2007
- [5] G. D. Taylor, K. R. Gue. *Retrieval Time Performance in Puzzle-Based Storage Systems*. In A. Johnson and J. Miller, editors, Proceedings of the 2010 Industrial Engineering Research Conference, 2010.
- [6] G. D. Taylor, K. R. Gue. Design and Performance of Multi-Level Puzzle-Based Storage Systems. In Proceedings of the 2009 International Conference on Value Chain Sustainability, 2009
- [7] K. R. Gue, K. Furmans. Decentralized Control in a Grid-Based Storage System, In T. Doolen and E. Van Aken, editors, Proceedings of the 2011 Industrial Engineering Research Conference, 2011
- [8] A. Alfieri, M. Cantamessa, A. Monchiero, E. Montagna. *Heuristics for puzzle-based storage systems driven by a limited set of automated guided vehicles*, Journal of Intelligent Manufacturing, published online, 2010.

Mobile Robotics in Education and Research of Logistics

Ulrich Karras, Dirk Pensky, and Octavio Rojas

Abstract—In the area of discrete manufacturing use of socalled AGVs (automated guided vehicles) is state of the art. Mobile systems connect workplaces, storage areas as well as incoming and delivery zones achieving lean production logistics. This publication focuses on the need for providing effective educational systems and developing innovative control methods to increase the use of mobile robots in this professional field of application. To fulfil these tasks, we present the mobile mechatronic learning system and research object Robotino offered by the industrial education company Festo Didactic. Besides introduction of motivating learning systems for future production logistics experts, we document supported options how to realize new concepts for more efficient transportation execution. The presentation is based on three practical samples covering different levels of complexity.

I. INTRODUCTION

In industry mobile robots are used as autonomous transport systems to make logistics processes more flexible and efficient. However, it turns out that the mobile robot systems in this kind of applications are guided and controlled by supervisory systems and are still far away to act autonomously as intelligent mobile systems. There are numerous scientific papers about multi-agent planning and scheduling methods but there exist no reliable implementation on mobile robot systems in industry.



Fig. 1. The mobile mechatronic learning and research system Robotino

Manuscript received July 8, 2011.

U. Karras, Priv.-Doz. Dr. rer. nat., is consultant for Festo Didactic GmbH & Co. KG in Denkendorf, Germany.(e-mail: Ulrich-Karras@t-online.de).

D. Pensky, Dr.-Ing., is product manager at Festo Didactic GmbH & Co. KG in Denkendorf, Germany. (phone: +49-711-3467-1334; fax: +49-711-346754-1334; e-mail: pens@de.festo.com).

O. Rojas is sales manager at Festo Mexico in Mexico City, Mexico. (e-mail: Octavio.Rochas@mx.festo.com). One goal of this paper is to address this industrial problem to the education and research of logistics processes in mobile robotics. Our implementation platform will be the mobile robot system Robotino [1] that has gained worldwide acceptance for a variety of educational and research institutions.

This article focuses on three logistics applications performed using Robotino and providing different levels and aspects of possible learning and research scenarios up to the well known robotics competition at the RoboCup event [2].

II. THE MOBILE MECHATRONIC LEARNING SYSTEM AND RESEARCH OBJECT ROBOTINO

A. System Design

Robotino is a mobile mechatronic learning system and research object developed by the industrial education company Festo Didactic, Denkendorf, Germany, in cooperation with the Robotics Equipment Corporation (REC), Munich, Germany. Worldwide, educational and research institutions use this fascinating and motivating system which features a compact and robust design as well as expandable hardware, see Fig. 1.

The three drive modules of the Robotino are integrated in a sturdy, laser-welded stainless steel chassis. The chassis is protected against collision by means of a rubber protective guard with integrated switching sensor. The robot's dimensions are 370 mm in diameter and 210 mm in height with an overall weight of approximately 11 kg. With its omni-directional drive, Robotino can move forwards, backwards and sideways in any direction and also turns on the spot. Three industrial DC motors with optical shaft encoders and gears with interchangeable pinions permit speeds of up to 10 km/h. The chassis contains nine infrared distance sensors. Analogue inductive and optical reflective sensors are also available for the Robotino to sense for example an aluminium strip or a coloured line. Additionally, Robotino is supplied with a colour webcam with JPEG compression. The compressed webcam image can be transmitted to an external PC via the wireless LAN for image evaluation or used as a live camera image. Power is supplied via two 12 V lead gel batteries which permit a running time of up to two hours.

A wide range of accessories is available such as sensors, e.g. optical, inductive, cameras, gyroscopes and laser range finders, plus handling devices, e.g. electrical grippers, arms and lifting devices. Moreover, users are able to integrate their own custom devices into Robotino by making use of various interfaces.

B. Control System Architecture

The Robotino features a high-performance embedded PC based on the PC 104 form-factor and an AMD Geode processor with a real-time Linux kernel. Operating system, programs and data are stored on an interchangeable compact flash card. The embedded PC provides several interfaces such as Ethernet, Wireless LAN, USB, RS232, and VGA port.



Fig. 2. Structure of I/O board EA09 for Robotino

The Robotino server, a real-time Linux application, forms the heart of the controller. It controls the Robotino drive units and can communicate with external applications in three different ways:

- An open Linux library of C++ basic functions is available for direct, on-board programming of the embedded PC under Linux.
- A TCP/IP communication interface is provided for communication with the control computer via wireless LAN. The user is able to write C++ applications for wireless control of Robotino on the basis of a Windows C++ function library.
- Robotino View, an interactive graphical programming system, communicates over wireless LAN directly with the robot system, with no compiling or downloading to the controller.

The activation of the three motors is performed via an additional board with digital and analogue inputs and outputs which is connected to the embedded PC via a serial connection. The control board provides the following additional interfaces for custom expansion, see Fig. 2:

- 8 analogue inputs 0 10 V, 50 Hz
- 8 digital inputs and outputs (24 V, short circuit proof and overload proof)
- 2 relays for additional actuators

C. Programming

Robotino View, developed by REC, supports the development of control sequences based on the standardised sequential function chart language of IEC 61131 [3]. Users benefit from a navigation module providing functions for

moving and defining travel paths, as well as from comprehensive functional module libraries e.g. mathematic and logic functions, see Fig. 3. Custom functional modules can be developed by programming in C++ or Lua, a powerful interpreter language [4]. For the communication between several Robotinos and to external controllers, links via object linking and embedding for process control (OPC), TCP/IP or user datagram protocol (UDP) are provided.

There are numerous application programming interfaces (APIs) available that enable the user to program Robotino in well established high level languages as C++, .Net Framework, JAVA, LabVIEW, MATLAB and the open source robot operating system ROS [5].

In comparison to other available mobile robot systems for education Robotino intensively focuses on the use as an AGV in production environments. Based on this standardized platform users are able to carry out different applications in logistics processes easily by developing their own handling and sensor components or by employing available accessories.

III. LOGISTICS TRAINING ARENA

This training set, called Logistics Training Arena, arose from following question: What are the fundamentals in order to accomplish an autonomous intelligent transport system? Many discussions with teachers and industrial companies document that even beginners should be able to solve following problems:

- Navigation consisting of localization and path finding to target positions
- 2. Interfacing to external stations, e.g. stack magazines, buffers, conveyor belts or charging stations
- 3. Communication to external control systems

For navigation the programming system Robotino View already includes all basic modules. Interfacing requires at least two important steps, the integration of handling devices and high accurate positioning at external stations

The accuracy in positioning must be at least ± 1 mm. This requires additional structure elements in the working area, e.g. metallic tapes that can be detected by inductive sensors. Of course, numerous other more sophisticated solution concepts exist but for beginners this method is an appropriate first step.

To build up communication between the AGV and the supervisory control system communication protocols shall be used well known in production industry. For this purpose one of the provided data exchange protocols OPC, TCP/IP or UDP are suitable.

The Logistics Training Arena consists of a small training platform with dimensions of the transport area of 0.9 m to 1.8 m with Robotino and some basic logistic elements as flat storage areas, magazine, conveyor belt, work pieces, and a charging station for the robot system, see Fig. 3.



Fig. 3. Logistics Training Arena for Robotino

The programming and control system Robotino View running on an external PC or locally on the embedded PC guides Robotino to fulfil tasks like transporting work pieces from one logistic element to the other. Based on OPC Robotino communicates with a PLC (programmable logic controller) to trigger the conveyor belt. As an option, the Logistics Training Arena can be expanded by connection to production stations and workplaces to build up a production environment, see Fig. 4.



Fig. 4. Logistics Training Arena coupled to production stations

Users of this learning system give us positive feedback regarding effectiveness in qualifying and motivating beginners at school and vocational school level. The main focus is on programming a single AGV in a known environment.

IV. PROLOG - THE LOGISTICS LEARNING FACTORY

When the logistics term is taken in consideration, all participants within the supply chain must contribute actively on the process. Therefore not only the mobile mechatronics system Robotino works as a dynamic link among the different warehouses and manufacturing sites, but also intermediate warehouses, delivery sites and software tools to find the optimal path as well as real time communication system play a dynamic role which will provide accurate and precise to and from the supply chain management system. The ProLog (Production Logistics) learning factory consists of three main areas [5]:

1. The manufacturing sites; where different types of solutions and pieces are produced see Fig. 5. On those sites, mechatronics systems work together divided in small production stations where each of them performs a certain process. Those cells together build a production process which adds value to the product.



Fig. 5. Production line in the ProLog learning factory

2. The commissioning station where the final products from different manufacturing sites are sorted and selected. This robot cell receives the products coming from different manufacturing sites just in sequence (JIS) and palletizes them according the incoming order and delivery process. Depending on the order the robot picks and places them on the pallet. The data transmission is made through an RFID device (Radio-Frequency Identification) which writes the new status in the tag of the pallet so that this pallet can be identified in the next delivery process, see Fig. 6.



Fig. 6. Commissioning station in the ProLog learning factory

3. The intermediate warehouses area where the orders (pallets) are buffered and then on time delivered by two mobile robots, Robotino. Within the ProLog learning factory, realistic experiments and practical, relevant training on mobile robotics in logistics are possible. Two Robotinos are equipped with a telescopic fork and operate like forklift trucks to reach all positions in the warehouse area, see Fig. 7.



Fig. 7. Warehouse area with Robotino in the ProLog learning factory

The mobile robots which work as a "milk - runner" or "water - spider" within the system not only moving products, or pallets but also searching for the optimal path to get parts from one point to another or to organize the warehouse, the delivery process. Of course every time the mobile robot makes a movement, all info such as part number, delivery number, product characteristics, etc. is transferred from one point to the next.

The ProLog learning factory covers all state of the art industrial communication levels:

- Factory level Order entry and enterprise resource planning (ERP)
- Plant level Supervisory control and data acquisition (SCADA) plant data collection
- Field level PLC, robot controller and RFID device

The logistics market has been always demanding "hands on" training systems where all people related to this field could improve their skills in addition to build better and optimized methods which contribute to the process optimization. Additionally, the use of mobile devices within logistics process has spread rapidly therefore those systems must provide accuracy, speed, and reliability in operation as well as in data transfer besides the opportunity to add more and more duties to them through open sources platforms.

V. ROBOCUP - FESTO LOGISTICS COMPETITION

The mission of RoboCup is the promotion of robotics and AI research, by offering publicly appealing, but formidable challenge. One of the effective ways to promote engineering research, apart from specific application developments, is to set a significant long term goal. When the accomplishment of such a goal has significant social impact, it is called the grand challenge project. Building a robot to play soccer game itself do not generate significant social and economic impact, but the accomplishment will certainly considered as a major achievement of the field. We call this kind of project as a landmark project. RoboCup is a landmark project as well as a standard problem. For 2050 the ultimate dream of RoboCup is that a robot team will beat the world best human soccer team.

In industry mobile robots are used as guided autonomous transport systems controlled by a supervisory system. Such systems are still far away to act autonomously as intelligent mobile robots to make logistics processes more flexible and efficient. This problem is closely related to the research about multi-agent planning methods [6]. There are numerous scientific papers about this subject. However, there are many questions open if one wants to apply these theoretical results to a set of autonomous robot systems. The value of classical task planning in soccer and rescue leagues is rather limited; if it used at all, these planners have to deal with constraints that are quite different from those in most industrial domains.

Our approach is to put it in the framework of an innovative new RoboCup competition, the Festo Logistics Competition (FLC), in order to encourage young research teams to develop new creative implementations of artificial intelligence methods to solve autonomously logistic transportation problems being more industrial oriented than soccer games.

A. The Logistics Competition

We envisage a kind of hardware-in-the-loop simulation method, i.e. there is a flexible simulated production hall with integrated real mobile robot systems having the task to create an efficient material flow to provide a high rate of deliveries in due time [7], [8].

Each team consists of three autonomous working robots. The production machines are represented by industrial read/write devices with an industrial signal light and the product components are represented by RFID tags with well defined ID-numbers carried on a pallet, see Fig. 7.



Fig. 7. Competition court of the Festo Logistics Competition at RoboCup

The signal light shows the status of the machine – idle, running, waiting, out of order status. The behavior of the machines is controlled by an external PLC. The production process is simulated in the way that a robot places a pallet with an RFID tag representing a subcomponent in the valid working area of the read/write device and the ID of the tag is changed after a predefined processing time.

There are various challenges:

- The locations of the machines are fixed but the machine type may vary. Thus, the teams have to develop a strategy to identify the locations of the various machine types as fast as possible.
- For the final production step the product to be finished requires at least 3 subcomponents which have to be produced at different machines.
- There are at least three delivery gates but by random only one gate is active.
- Consumed material should be recycled.
- At random times express goods are provided which might be handled in a predefined time slots.
- Obstacles may appear in terms of robots of the opponent team which might block paths and the access of certain machines.
- The main challenge of course is that the robot team autonomously creates a most efficient logistics process taking into account all challenges mentioned above.

The FLC competition uses the concept of our common industrial platform Robotino. The rules only admit that teams can create their own concept of sensors and their own pallet carrier. There are several reasons for this decision:

- 1. Focus on the main task to develop intelligent autonomous behavior.
- 2. Standard equipment reduces costs and saves developing time of well known equipment.
- 3. Restriction to common platforms is an important requirement from the industrial viewpoint of support and maintenance.
- 4. Common platforms provide fair chances to new teams to join the competition.

B. Experiences

In 2010 the FLC runs as a demonstration competition at the international RoboCup event in Singapore for the first time with 9 teams from Germany, Swiss, Hungary, Egypt, Singapore and Korea. This year 2011 the FLC runs at the international RoboCup event in Istanbul with 15 teams. In addition, this competition is also available at other RoboCup events, e.g. the German Open.

The FLC task has been shown to be highly challenging but teams never gave up and spent much time on research and development in order to be one of the bests. The present results have shown that there is a huge potential of improvements and optimization of presented solutions.

C. Scientific Challenges

The approach of FLC provides an outstanding research platform to make important progress in this challenging task because of a rather generic factory structure and interface structure between robot and factory production device. The actual rulebook is focussed on a rather simple 3-stage production process which can be easily changed related to product complexity and production processes, see Fig. 8.



Fig. 8. 3-stage production process of the FLC

At present there is only one component on the pallet. If we admit various components on a pallet then we would immediately generate a much more complex material flow opening the door for additional strategies to proceed. This would also imply that the robots need robot arms to handle the different components. A further step could be to integrate an ERP system as we have already considered in the ProLog learning factory in section IV.

As well the actual rules avoid dynamic appearance of obstacles which can be easily created by admitting teams play directly against each other. With other words: One team plays the role of production teams and the other plays the role of dynamic obstacles in order to block the access to important production machines.

It is well known that planning and scheduling methods may be very well also analyzed by simulation. In fact, teams of the FLC competition are provided with a simulation tool in order to prepare strategic procedures. However, industry is mainly interested in a reliable implementation on a common hardware platform.

In industrial logistic applications mobile robot systems are controlled by a supervisory system yielding movement commands to the robot controllers. A task oriented communication between supervisory system and mobile robots is by no means accepted in industry since no reliable solutions are available on the market. As well, autonomous docking procedures or interactions with external devices may work in some laboratories but not at all in an industrial environment.

VI. CONCLUSION

Highlighting the significance of mobile robots for production logistics we derive the necessity for a standardized learning and research system. The mobile robot platform Robotino by Festo Didactic combines all functions and properties to be used for qualification as well as for developing new methods to accomplish autonomous transportation systems. To prove this we present three applications introduced at the schools level up to the university level. Starting with the programming of a single AGV, taking the next step to synchronize two mobile robots in the same area and share tasks between these systems, the Festo Logistics Competition at RoboCup is the most challenging application research institutes around the world working on.

REFERENCES

- R.-C. Weber, M. Bellenberg, Robotino Manual, Festo Didactic GmbH & Co. KG, Denkendorf, Germany, 2010
- [2] ROBOCUP 2011, RoboCup Demonstration Event Festo Logistics Competition, see: http://www.robocup2011.org, January 2011
- [3] M. Bellenberg, Robotino View 2 Manual, Festo Didactic GmbH & Co. KG, Denkendorf, Germany, April 2010
- [4] R. Ierusalimschy, L.H. de Figueiredo, W.C. Filho: Lua An Extensible Extension Language. Software: Practice and Experience 26(6) (1996) 635–652.
- [5] C. Verbeek, Robotino View API Documentation, Robotics Equipment Corporation GmbH, see: http://doc.openrobotino.org/documentation/_ OpenRobotinoAPI/1/doc/rec_robotino_com/annotated.html, Januar 2011
- [6] S. Schober, ProLog Manual, Festo Didactic GmbH&Co.KG, Denkendorf, Germany, 2009.
- [7] M. de Weerdt, A. ter Mors, and C. Witteveen, Multi-agent planning an introduction to planning and coordination. Technical report, Delft University of Technology, 2005. URL http://www.st.ewi.tudelft.nl/ mathijs/publications/easss05.pdf.
- [8] M. Riedmiller, U. Karras, Simulation of Logistics in RoboCup Competition, SIMPAR 2010, Workshop Proceedings, see: http://www.simpar.org, November 2010
- [9] C. Deppe, U. Karras, Festo Logistics Competition 2011 Rulebook, Version 1.006.2011, Festo Didactic GmbH & Co. KG, Denkendorf, Germany, April 2011

Simultaneous Localization and Dynamic State Estimation in Reconfigurable Environments

Gian Diego Tipaldi

Daniel Meyer-Delius

Maximilian Beinhofer

Wolfram Burgard

Abstract—The majority of existing approaches to mobile robot localization assume that the world is static, which clearly does not hold in most real-world application domains. In this paper we present a probabilistic approach to global localization in reconfigurable environments, where the robot pose and the environment state are jointly estimated using a Rao-Blackwellized particle filter. The environment is represented as a spatial grid and a hidden Markov model is used to represent the occupancy state and state transition probabilities of each grid cell. The HMM parameters are estimated offline using the EM algorithm. Experimental results show that our model is better suited for representing reconfigurable environments than standard occupancy grids. Furthermore, the results show that the explicit representation of the environment dynamics can be used to improve localization accuracy in reconfigurable environments.

I. INTRODUCTION

An accurate model of the environment is essential for many mobile robot navigation tasks. Although the environment generally is dynamic, most existing navigation approaches assume it to be static. They typically build the map of the environment in an offline phase and then use it without considering potential future changes. There are robust approaches that can handle inconsistencies between the map and the actual measurements. However, a largely inconsistent model can lead to unreliable navigation or even to a complete localization failure. Moreover, some areas can be *semi-static* (e.g., furnitures can move, cars can be parked) and this information can be used by the robot to improve its navigation performance.

In this paper we consider the problem of modeling a mobile robot's environment taking the dynamics of the environment explicitly into account. We present a probabilistic model that represents the occupancy of the space and characterizes how this occupancy changes over time. We then show how this information can be used to jointly estimate the pose of the robot and the configuration of the environment during a global localization task.

The environment is represented as a two-dimensional grid where each cell uses a hidden Markov model (HMM) to represent the belief about the occupancy state and state transition probabilities. Our model, called *dynamic occupancy grid*, is a generalization of a standard occupancy grid. Figure 1 illustrates the fundamental difference between these



Fig. 1. Bayesian network describing the dependencies between the states of a cell c and observations z in standard and dynamic occupancy grids.

two models: while occupancy grids characterize the state of a cell as static, our representation explicitly models state changes.

In addition to the explicit representation of the environment dynamics, the HMM framework provides efficient algorithms for estimating the model parameters. Given that, we can use a Rao-Blackwellized particle filter (RBPF) to jointly reason about the robot pose, which represents the sampled part of the filter, and the occupancy probability of a cell, represented in the analytical part of the factorization.

The contribution of this work is a global localization approach that estimates the pose of the robot and, at the same time, explicitly infers how the state of the world changes over time. To the best of our knowledge, this is the first approach to address this problem. Previous attempts either focused on how to filter spurious observations due to dynamic objects or addressed the easier problem of pose tracking. We describe our model and how the representation can be updated as new observations become available. We further propose a local map representation that is able to forget changes in a sound probabilistic way, using the mixing times of the associated Markov chain, and to minimize memory requirements.

We evaluate our approach in simulation and using realworld data. The results demonstrate that our model can represent dynamic environments more accurately than standard occupancy grids. Furthermore, we show that it outperforms standard Monte Carlo localization in complex real world environments with consistent changes over time.

II. RELATED WORK

Most mobile robot navigation systems rely on a map of the environment built beforehand in an offline phase. To deal with subsequent changes in the environment, sensor measurements caused by dynamic objects are usually filtered out. Robust approaches rely on probabilistic sensor models that identify the measurements inconsistent with the map. For example, Fox *et al.* [1] use an entropy gain filter, while Burgard *et al.* [2] use a distance filter based on the expected distance of a measurement. Despite the success of these

All authors are with the University of Freiburg, Dept. of Computer Science, D-79110 Freiburg, Germany.

This work has been partially supported by the European Commission under contract numbers FP7-231888-EUROPA and FP7-260026-TAPAS. Also by the German Research Foundation (DFG) within the Research Training Group 1103.

techniques, they discard valuable information about the environment. Instead of filtering out inconsistent measurements, Montemerlo *et al.* [3], use them for people tracking while localizing the robot. They show that the state of dynamic objects in the environment can be used to increase the robustness of the pose estimation process. Motivated by this idea, we utilize all sensor measurements to keep the map of the environment up-to-date while simultaneously localizing the robot within the updated map.

Orthogonal to the work on localization in dynamic environments, many authors have addressed the problem of modeling such environments. Hähnel et al. [4], for example, combine the EM algorithm and a sensor model that considers dynamic objects to obtain accurate maps. The approach of Anguelov et al. [5] computes shape models of nonstationary objects. They create maps at different points in time and compare those maps using an EM-based algorithm to identify the parts of the environment that change over time. Wolf and Sukhatme [6] propose a model that maintains two separate occupancy grids, one for the static parts of the environment and the other for the dynamic parts. Biber and Duckett [7] propose a model that represents the environment on multiple timescales simultaneously. For each timescale a separate sample-based representation is maintained and updated using the observations of the robot according to an associated timescale parameter. In our work, we use the gridbased representation described by Meyer-Delius et al. [8] to represent dynamic environments. Besides being able to continuously adapt to changes over time, this model provides also an explicit characterization of the dynamics of the environment that can be learned from data.

Whereas most of the work on mapping dynamic environments assumes that a good estimate of the robot's pose is available, most of the work on mapping where the pose of the robot is not available (i.e., SLAM) assumes that the environment is static. Only few authors address the problem of jointly estimating the pose of the robot and the state of a dynamic environment. Avots et al. [9], for example, use a Rao-Blackwellized particle filter to estimate the pose of the robot and the state of doors in the environment. They represent the environment using a reference occupancy grid where the location of the doors is known, but not their state (i.e., opened of closed). Petrovskaya and Ng [10] propose a similar approach where instead of a binary model, a parametrized model (i.e., opening angle) of the doors is used. Similar to these approaches, we also use a Rao-Blackwellized particle filter to estimate the pose of the robot and the state of the environment. In contrast to their methods, however, we estimate the state of the complete environment, and not only of small, specific areas or elements. Additionally, we also learn the model parameters from data.

III. DYNAMIC OCCUPANCY GRID

Occupancy grids (as they were introduced by Moravec and Elfes [11]) are a regular tessellation of the space into a number of rectangular cells. They store in each cell the probability that the corresponding area of the environment is



Fig. 2. State transition probabilities of the parking lot of the University of Freiburg. The left and right images correspond to the distributions $p(c_t = f | c_{t-1} = f)$ and $p(c_t = o | c_{t-1} = o)$ respectively. The darker the color, the larger the probability for the occupancy to remain unchanged.

occupied by an obstacle. To avoid a combinatorial explosion of possible grid configurations, the approach assumes that neighboring cells are independent from each other.

Occupancy grids rest on the assumption that the environment is static. As mentioned above, they store for each cell c of an equally spaced grid the probability p(c) that c is occupied by an obstacle. To probabilistically model how the occupancy changes over time in dynamic environments, the approach relies on an HMM (see Rabiner [12]) to explicitly represent both the belief about the occupancy state and state transition probabilities of each grid cell as illustrated in Figure 1.

An HMM requires the specification of a state transition, an observation, and an initial state distribution. Let c_t be a discrete random variable that represents the occupancy state of a cell c at time t. The initial state distribution or prior $p(c_{t=0})$ specifies the occupancy probability of a cell at the initial time step t = 0 prior to any observation.

The state transition model $p(c_t \mid c_{t-1})$ describes how the occupancy state of cell c changes between consecutive time steps. We assume that the changes in the environment are caused by a stationary process, that is, the state transition probabilities are the same for all time steps t. These probabilities are what allows us to explicitly characterize how the occupancy of the space changes over time. Since we are assuming that a cell c is either free (f) or occupied (o), the state transition model can be specified using only two transition probabilities, namely $p(c_t = f | c_{t-1} = f)$ and $p(c_t = o \mid c_{t-1} = o)$. Note that, by assuming a stationary process, these probabilities do not depend on the absolute value of t. Figure 2 depicts transition probabilities for the parking lot at our faculty. The darker the color, the larger the probability for the corresponding occupancy to remain unchanged. The figure clearly shows the parking spaces, driving lanes, and static elements such as walls and lampposts as having different dynamics. The "shadows" in the upper left and lower right areas of the maps were mostly caused by maximum range measurements being ignored.

The observation model $p(z \mid c)$ represents the likelihood of the observation z given the state of the cell c. The observations correspond to measurements obtained with a range sensor. In this paper, we consider only observations obtained with a laser range scanner. The cells in the grid that are covered by a laser beam are determined using a ray-tracing operation. We consider two cases: the beam is not a maximum range measurement and ends up in a cell (a hit) or the beam covers a cell without ending in it (a miss). Accordingly, the observation model can also be specified using only two probabilities: p(z = hit | c = f)and p(z = hit | c = o). We additionally take into account the situation where a cell is not observed at a given time step. This is necessary since the transition model characterizes state changes only for consecutive time steps. Explicitly considering this no-observation case allows us to update and estimate the parameters of the model using the HMM framework directly without having to distinguish between observations and no-observations. The concrete observation probability for a no-observation does not affect the results as long as the proportion between the two remaining probabilities remains unchanged.

From the discussion above it can be seen that standard occupancy grids are a special case of dynamic occupancy grids where the transition probabilities $p(c_t = f | c_{t-1} = f)$ and $p(c_t = o | c_{t-1} = o)$ are 1 for all cells *c*.

A. Occupancy State Update

The update of the occupancy state of the cells in a dynamic occupancy grid follows a Bayesian approach. The goal is to estimate the belief or posterior distribution $p(c_t | z_{1:t})$ over the current occupancy state c_t of a cell given all the available evidence $z_{1:t}$ up to time t. The update formula is:

$$p(c_t \mid z_{1:t}) = \eta \, p(z_t \mid c_t) \, \sum_{c_{t-1} \in \{f, o\}} p(c_t \mid c_{t-1}) \, p(c_{t-1} \mid z_{1:t-1}) \,, \ (1)$$

where η is a normalization constant. Exploiting the Markov assumptions in our HMM, this equation is obtained using Bayes' rule with $z_{1:t-1}$ as background knowledge and applying the theorem of total probability on $p(c_t | z_{1:t-1})$ conditioning on the state of the cell c_{t-1} at the previous time step t-1. Equation (1) describes a recursive approach to estimate the current state of a cell given a current observation and the previous state estimate. This approach corresponds to a discrete Bayes filter. The structure of our particular HMMs allows for a simple and efficient implementation of this approach. Note that the map update for standard occupancy grids is a special case, where the sum in (1) is replaced by the posterior $p(c_t | z_{1:t-1})$.

This posterior corresponds to a prediction of the occupancy state of the cell at time t based on the observations up to time t - 1. Prediction can be considered as filtering without the processing of evidence. By explicitly considering no-observations as explained in the previous section, the update formula can be used directly to estimate the future state of a cell or estimate the current state of a cell that has not been observed recently.

B. Parameter Estimation

As mentioned above, an HMM is characterized by the state transition probabilities, the observation model, and the initial state probabilities. We assume that the observation model only depends on the sensor. Therefore it can be specified beforehand and is the same for each HMM. We estimate the remaining parameters using observations that are assumed to correspond to the environment that is to be represented.

One of the most popular approaches for estimating the parameters of an HMM is an instance of the expectationmaximization (EM) algorithm. The basic idea is to iteratively estimate the model parameters using the observations and the parameters estimated in the previous iteration until the values converge. Let $\hat{\theta}^{(n)}$ represent the parameters estimated at the *n*-th iteration. The EM algorithm results in the following re-estimation formula for the transition model of cell *c*:

$$\hat{p}(c_t = i \mid c_{t-1} = j)^{(n+1)} = \frac{\sum_{\tau=1}^{T} p(c_{\tau-1} = i, c_{\tau} = j \mid z_{1:T}, \hat{\theta}^{(n)})}{\sum_{\tau=1}^{T} p(c_{\tau-1} = i \mid z_{1:T}, \hat{\theta}^{(n)})}, \quad (2)$$

where $i, j \in \{f, o\}$ and T is the length of the observation sequence used for estimating the parameters. Note that the probabilities on the right-hand side are conditioned on the observation sequence $z_{1:T}$ and the previous parameter estimates $\hat{\theta}^{(n)}$. These probabilities can be efficiently computed using the forward-backward procedure [12].

For more details about this model we refer the reader to the associated technical report by Meyer-Delius *et al.* [8].

IV. SIMULTANEOUS LOCALIZATION AND DYNAMIC STATE ESTIMATION

In this section we describe our approach to simultaneously estimate the robot pose and the dynamic state of the environment. Although on first sight one can see the addressed problem as an instance of the better known simultaneous localization and mapping (SLAM), there are two main differences between them.

The first difference is the absence of a global reference frame in the SLAM problem. No global pose is required and the initial pose of the robot can typically be set freely. On the contrary, we explicitly address global localization as part of the estimation aspect. We have a global reference frame and the initial pose of the robot is unknown and assumed to be uniformly distributed over the whole environment. The second difference regards the dimensionality of the map. In the SLAM problem, the size of the map is not known in advance and can grow unbounded with time. In our problem the size of the map is known and we only focus on estimating the actual configuration of the dynamic objects present in the environments. Despite the differences, the two problems do share the same state space, i.e., the robot pose and the state of the map, and one can exploit the same factorization that made Rao-Blackwellized particle filters a feasible solution to the SLAM problem [13], [14].

In the following we show how this factorization can be exploited and we derive the RBPF that will be used to estimate the posterior $p(x_{1:t}, m_{1:t} | z_{1:t}, u_{1:t-1}, m_0)$ about the trajectory $x_{1:t}$ of the robot and the configuration of the environment $m_{1:t}$, given the observations $z_{1:t}$, the odometry

measurements $u_{1:t-1}$ and the prior over the map m_0 . The key idea is to separate the estimation of the robot pose from the map estimation process,

$$p(x_{1:t}, m_{1:t} \mid z_{1:t}, u_{1:t-1}, m_0) = p(m_{1:t} \mid x_{1:t}, z_{1:t}, m_0) p(x_{1:t} \mid z_{1:t}, u_{1:t-1}, m_0).$$
 (3)

This can be done efficiently, since the posterior over maps $p(m_{1:t} \mid x_{1:t}, z_{1:t}, m_0)$ can be computed analytically given the knowledge of $x_{1:t}$ and $z_{1:t}$ and using the forward algorithm for the HMM. The remaining posterior, $p(x_{1:t} \mid z_{1:t}, u_{1:t-1})$, is estimated using a particle filter which incrementally processes the observations and the odometry readings. Following [15], we obtain a sample of the robot trajectory by incrementally sampling the current pose from the motion model $x_t^{(i)} \sim p(x_t \mid x_{t-1}^{(i)}, u_{t-1})$ and setting $x_{1:t}^{(i)} = \{x_t^{(i)}, x_{1:t-1}^{(i)}\}$. This recursive sampling schema allows us to recursively compute the importance weights using the following equation

$$w_t^{(i)} = w_{t-1}^{(i)} \frac{p(z_t \mid x_{1:t}^{(i)}, z_{1:t-1}, m_0) p(x_t^{(i)} \mid x_{t-1}^{(i)}, u_{t-1})}{p(x_t^{(i)} \mid x_{t-1}^{(i)}, u_{t-1})} = w_{t-1}^{(i)} p(z_t \mid x_{1:t}^{(i)}, z_{1:t-1}, m_0).$$
(4)

The observation likelihood is then computed by marginalization over the predicted state of the map leading to

$$p(z_t \mid x_{1:t}^{(i)}, z_{1:t-1}, m_0) = = \int p(z_t \mid x_t^{(i)}, m_t) p(m_t \mid m_{t-1}^{(i)}) dm_t = \prod_j \mathcal{N}(z_t^j; \hat{z}_t^j, \sigma^2),$$
(5)

were z^j is an individual laser reading and \hat{z}_t^j is the closest cell in the map to the reading, with an occupancy probability above a certain threshold. Note that the *disappearance* of the integral is not an approximation but a direct consequence of using the likelihood field model described in [16].

A. Map Management

As we already mentioned above, the initial pose of the robot is unknown and assumed to be uniformly distributed in the whole environment. This forces us to use a high number of particles, generally above thousands, to accurately represent the initial distribution. Since every particle needs to have its own estimate of the map, memory management is a key aspect of the whole algorithm. In order to save memory, we want to only store the cells in the map that have been considerably changed from the a priori map m_0 , which is shared among the diverse particles. This is done exploiting two important aspects of the Markov chain associated to the HMM: the *stationary distribution* and the *mixing time*.

As the number of time steps for which no observation is available tends to infinity, the occupancy value of a cell converges to a unique stationary distribution π (see [17]). This stationary distribution represents the case where the environment has not been observed for a long time and is represented by our a priori map m_0 . In the case of a binary HMM, the one used in this paper, this distribution is computed using the transition probabilities

$$\begin{bmatrix} \pi_f \\ \pi_o \end{bmatrix} = \frac{1}{p+q} \begin{bmatrix} q \\ p \end{bmatrix}, \tag{6}$$

where for notation simplicity we have

$$p = p(c_t = o \mid c_{t-1} = f)$$

$$q = p(c_t = f \mid c_{t-1} = o).$$

Every time an individual particle observes the state of a cell for the first time, the state distribution of that particular cell changes from the stationary one and the particle needs to store the new state of the cell. In order to reduce memory requirements, only a limited number of cells should be stored and a *forgetting* mechanism should be implemented. This can be done in a sound probabilistic way, by exploiting the mixing time of the associated Markov chain. The mixing time is defined as the time needed to converge from a particular state to the stationary distribution. The concrete definition depends on the measure used to compute the difference between distributions. In this paper we use the total variation distance as defined by Levin et al. [17]. Since our HMMs have only two states, the total variation distance Δ_t between the stationary distribution π and the occupancy distribution p_t at time t can be specified as

$$\Delta_t = |1 - p - q|^t \Delta_0, \tag{7}$$

where $\Delta_0 = |p(c_t = f) - \pi_f| = |p(c_t = o) - \pi_o|$ is the difference between the current state $p(c_t)$ and stationary distribution π . Based on the total variation distance, we can define the mixing time t_m as the smallest t such that the distance Δ_{t_m} is less than a given value ϵ . This leads to

$$t_m = \left\lceil \frac{\ln(\epsilon/\Delta_0)}{\ln(|1-p-q|)} \right\rceil.$$
(8)

In other words, the mixing time tells us how many steps are needed for a particular cell to return to its stationary distribution, that is how many step a particle needs to store an unobserved cell before removing it from its local state and rely on the a priori map m_0 .

V. EXPERIMENTS

We implemented our proposed model and tested it using data obtained with a real robot. We steered a MobileRobots Powerbot equipped with a SICK LMS laser range finder through the parking lot of our faculty. We performed a run every full hour from 7am until 6pm during one day. The range data obtained from the twelve runs (data sets d_1 through d_{12}) corresponds to twelve different configurations of the parked cars, including an almost empty parking lot (data set d_1) and a relatively occupied one (data set d_8). We used a SLAM approach [18] to correct the odometry of the robot and obtain a good estimate of its pose. Range measurements were sampled at about 1 Hz, and the trajectory and velocity of the robot during each run were approximately the same to try to avoid a bias in the complete data set.



Fig. 3. Comparison between dynamic and standard occupancy grids. Shown are the ground truth (top), dynamic occupancy grid (middle), and standard occupancy grid (bottom) maps at two different points in time.

Figure 3 shows a qualitative comparison between dynamic and standard occupancy grids for the parking lot data set. We assumed that the parking lot did not change considerably during a run and used the occupancy grids obtained from every data set with the above-mentioned SLAM approach as ground truth. In the figure, the maps on the left column show the grids after the third run, that is, after integrating data sets d_1 through d_3 . The maps on the right show the grids at the end of the last run, after integrating data sets d_1 through d_{12} . As can be seen, the dynamic occupancy grid readily adapts to the changes in the parking lot. Thus, it constitutes a better representation of the environment at any point in time. Additionally, dynamic occupancy grids provide information about the occupancy probability of areas that have not been recently observed. This appears in the grids in the figure (specially the right column) as light gray areas in the places where the cars most frequently park.

In order to assess the performances of the localization approach, we compared it to standard Monte Carlo localization both in a global localization and pose tracking setting. For each data set, we compared our approach (MCL-D), MCL using the standard occupancy grid (MCL-S), and MCL using the ground-truth map for that specific data set (MCL-GT). We performed 100 runs for each data set, where we randomly sampled the initial pose of the robot. In order to obtain a fair comparison, the same seed has been used to generate the initial pose, as well as to perform all the random sampling processes for each approach. All the approaches have been initialized with 10,000 particles for global localization and 500 particles for pose tracking. They also used the same set

of parameters, an occupancy threshold of 0.6 and a maximum distance of 1m for the likelihood field model.

The results of the global localization experiment are shown in Table I. The table shows the success rate of the global localization, as the percentage of time the filter converged to the true pose, and the residual squared error, with respective variance, after convergence. The success rate is reported relative to the result of MCL on the ground-truth map, in order to have a measure independent of the complexity of the environment. The results show that our approach outperforms the standard MCL on static maps both in terms of convergence rate and accuracy in localization.

Table II shows the results for the pose tracking experiment, where the filter is initialized at the true pose and keeps tracking the robot. The table shows the failure rate, i.e., the percentage of time the robot got lost during tracking, as well as the residual squared error. The results of this experiment show that the performances of the dynamic maps in pose tracking are almost equivalent to MCL with the ground-truth maps, with a failure rate of only 2%.

Both experiments show two important aspects of the problem and of the solution adopted. The first aspect is that the problem is much more complex than global localization since the search space is bigger and deciding if a measurement is an outlier or is caused by a change of the configuration is not a trivial task. Furthermore, analyzing the performances in pose tracking, we see that if the filter is initialized close to the correct solution, i.e., the search is reduced to the correct subspace, it is able to estimate the correct configuration. The second aspect is how the algorithm scales with different amount of change in the environment configuration. In the first four data sets, the parking lot is almost empty and it becomes quite full in the last ones. This is evident, when analyzing the results of MCL on the static maps, since the performance gets worse with an increasing amount of change. On the other hand, the performance of MCL on the dynamic maps is independent from the amount of change, as can be seen from the tables.

VI. CONCLUSIONS

In this paper we introduced a novel approach to global localization in reconfigurable environments using a grid map that explicitly represents how the occupancy of individual cells changes over time. The model is a generalization of standard occupancy grids and applies HMMs to update the belief about the occupancy state of each cell according to the dynamics of the environment. We described how a Rao-Blackwellized particle filter can be used to jointly estimate the robot pose and the configuration of the environment. This was possible thanks to the reduced memory requirements obtained by exploiting the properties of the Markov chains. We evaluated our approach using real-world data. The results demonstrate that our model can represent dynamic environments more accurately than standard occupancy grids and outperforms standard MCL on static maps in both global localization and pose tracking.

Data set	MCL-GT			MCL-D			MCL-S		
	Success	Error ²	σ^2	Success	Error ²	σ^2	Success	Error ²	σ^2
01	100%	0.21	0.36	50%	0.26	0.36	50%	0.26	0.18
02	100%	0.19	0.29	40%	0.10	0.08	33%	0.13	0.09
03	100%	0.13	0.19	80%	0.10	0.29	52%	0.19	0.17
04	100%	0.04	0.03	60%	0.08	0.14	53%	0.15	0.19
05	100%	0.07	0.18	54%	0.07	0.09	35%	0.15	0.18
06	100%	0.02	0.01	87%	0.02	0.02	45%	0.06	0.02
07	100%	0.06	0.08	59%	0.12	0.22	43%	0.14	0.20
08	100%	0.05	0.10	71%	0.03	0.02	28%	0.03	0.01
09	100%	0.02	0.01	53%	0.12	0.22	31%	0.06	0.02
10	100%	0.14	0.28	62%	0.13	0.31	34%	0.30	1.01
11	100%	0.11	0.11	38%	0.15	0.21	26%	0.24	0.29
12	100%	0.19	0.32	20%	0.16	0.14	22%	0.27	0.38
Total	100%	0.11	0.19	52%	0.11	0.18	36%	0.17	0.22

TABLE I

GLOBAL LOCALIZATION EXPERIMENT

Data set	MCL-GT			MCL-D			MCL-S		
	Failure	Error ²	σ^2	Failure	Error ²	σ^2	Failure	Error ²	σ^2
01	0%	0.04	0.01	3%	0.09	0.03	5%	0.18	0.07
02	0%	0.03	0.01	4%	0.08	0.05	24%	0.18	0.10
03	0%	0.04	0.01	2%	0.05	0.04	10%	0.09	0.04
04	0%	0.02	0.01	0%	0.04	0.01	10%	0.08	0.02
05	0%	0.02	0.01	3%	0.03	0.04	13%	0.06	0.02
06	0%	0.02	0.01	2%	0.02	0.01	26%	0.09	0.12
07	0%	0.02	0.01	0%	0.03	0.01	34%	0.07	0.01
08	0%	0.02	0.01	2%	0.02	0.01	35%	0.09	0.15
09	0%	0.02	0.01	4%	0.03	0.01	37%	0.07	0.16
10	0%	0.02	0.01	0%	0.03	0.01	36%	0.09	0.10
11	0%	0.03	0.01	1%	0.05	0.02	42%	0.10	0.05
12	0%	0.03	0.01	5%	0.06	0.01	44%	0.15	0.20
Total	0%	0.03	0.01	2%	0.04	0.02	27%	0.10	0.08

TABLE II Position tracking experiment

REFERENCES

- D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *Journal of Artificial Intelligence Research*, vol. 11, 1999.
- [2] W. Burgard, A. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "Experiences with an interactive museum tour-guide robot," *Artificial Intelligence*, vol. 114, no. 1-2, 2000.
- [3] M. Montemerlo, S. Thrun, and W. Whittaker, "Conditional particle filters for simultaneous mobile robot localization and people-tracking," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2002.
- [4] D. Hähnel, R. Triebel, W. Burgard, and S. Thrun, "Map building with mobile robots in dynamic environments," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.
- [5] D. Anguelov, R. Biswas, D. Koller, B. Limketkai, S. Sanner, and S. Thrun, "Learning hierarchical object maps of non-stationary environments with mobile robots," in *Proc. of the Conference on Uncertainty in AI (UAI)*, 2002.
- [6] D. F. Wolf and G. S. Sukhatme, "Mobile robot simultaneous localization and mapping in dynamic environments," *Autonomous Robots*, vol. 19, no. 1, pp. 53–65, 2005.
- [7] P. Biber and T. Duckett, "Dynamic maps for long-term operation of mobile service robots," in *Proc. of Robotics: Science and Systems* (*RSS*), 2005.
- [8] D. Meyer-Delius, M. Beinhofer, and W. Burgard, "Grid-based models for dynamic environments," Dept. of Computer Science, University of Freiburg, Tech. Rep., 2011.
- [9] D. Avots, E. Lim, R. Thibaux, and S. Thrun, "A probabilistic technique for simultaneous localization and door state estimation with mobile robots in dynamic environments," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.

- [10] A. Petrovskaya and A. Y. Ng, "Probabilistic mobile manipulation in dynamic environments, with application to opening doors," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2007.
- [11] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA), 1985.
- [12] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition." in *Proceedings of the IEEE*, vol. 77 (2), 1989, pp. 257–286.
 [13] K. Murphy, "Bayesian map learning in dynamic environments," in
- [13] K. Murphy, "Bayesian map learning in dynamic environments," in Proc. of the Conf. on Neural Information Processing Systems (NIPS), Denver, CO, USA, 1999, pp. 1015–1021.
- [14] G. Grisetti, C. Stachniss, and W. Burgard, "Improved Techniques for Grid Mapping with Rao-Blackwellized Particle Filters," *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.
- [15] A. Doucet, J. de Freitas, K. Murphy, and S. Russel, "Rao-Blackwellized partcile filtering for dynamic bayesian networks," in *Proc. of the Conf. on Uncertainty in Artificial Intelligence (UAI)*, Stanford, CA, USA, 2000, pp. 176–183.
- [16] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [17] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. American Mathematical Society, 2008.
- [18] G. Grisetti, C. Stachniss, and W. Burgard, "Improving grid-based SLAM with Rao-Blackwellized particle filters by adaptive proposals and selective resampling," in *Proc. of the IEEE Int. Conf. on Robotics* & Automation (ICRA), 2005.