Intra-Logistics with Integrated Automatic Deployment: Safe and Scalable Fleets in Shared Spaces

H2020-ICT-2016-2017
Grant agreement no: 732737

DELIVERABLE 7.4
Second demonstration

Due date: month 33 (September 2019)
Deliverable type: DEM
Lead beneficiary: UoL

Dissemination Level: PUBLIC

Main author: Martin Magnusson (ORU)

1 Introduction

This document reports on the outcome of the second live physical ILIAD project demonstration on October 16, 2019. The demo coincided with the second stakeholder meeting with invited industry representatives, at Orkla Foods' facilities in Örebro, Sweden.

The goal of this milestone demonstration — MS3 — was to show that *"the system can run in simulation and in a real environment, exhibiting the ability to self-deploy, pick and palletize goods, perform online integrated task allocation, motion planning and coordination, account for learned environment dynamics and human activity"*. As shown in this report, we demonstrated missions executed with a fleet of three robots at a real-world end-user site, demonstrating key technologies for deployment, picking and palletising, task allocation with changing requirements posted online during operation, motion planning and coordination taking into account learned human motion patterns. In addition, we showed mutual communication of intents, stronger integration between (in particular) perception and safe manipulation, initial implementations of automatic map quality assessment and qualitative human–robot spatial interaction, and improved versions of automatic plastic cutting, people tracking.

In this report we briefly cover the new components that were demonstrated at MS3, compared to what was shown already at MS2 in 2018 (see Deliverable D7.3).

2 Stakeholder meeting

The milestone demonstration was performed as part of a stakeholder meeting with invited industry represesentatives. Figure 1 shows some impressions from the meeting and demonstrations.

The guests from regional industries and business press were first given a popular-science overview of the progress and technical contents of the project. Orkla Foods presented ILIAD and its uses from an end-user perspective, and industrial partners Kollmorgen Automation and Logistic Engineering Services placed the project in the context of current industry developments worldwide. After the general overview, the participants were given a live demonstration of the capabilities of the ILIAD robot fleet, inside the ambient-temperature warehouse at Orkla Foods, as detailed in the next section. After the demo, the stakeholder meeting concluded with coffee and chatting with the guests and the developers, with snacks from the factory.

Figure 1: Images from the second ILIAD stakeholder meeting.

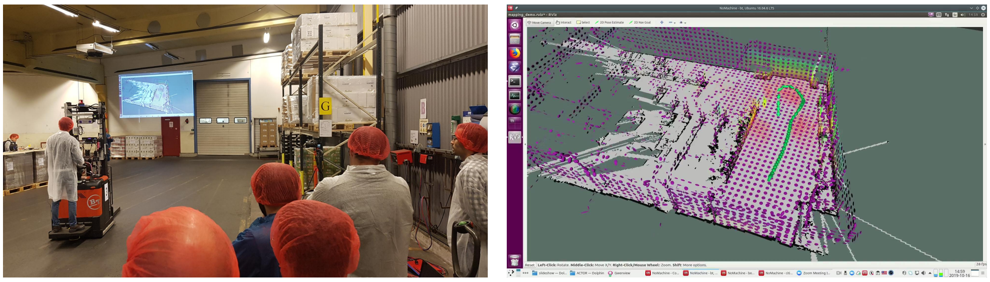

Figure 2: Mapping during deployment. *Left:* live view. *Right:* sensor view and the resulting maps.

3 Milestone demonstration

In the set of demonstrations, we showcased the current implementation of core components for deployment and operation — from when the first ILIAD robot is unpacked at a new site, to fleet coordination, navigation and people awareness, and object picking and unwrapping.

3.1 Mapping and localisation

We demonstrated 3D mapping (Deliverable D1.2) by driving one of the trucks equipped with a 3D lidar, as seen in Figure 2. Moving obstacles are automatically filtered from the map, and a 2D obstacle grid map is extracted from the fully-3D map to be used for further navigation and motion planning.

A novel addition, first demonstrated at MS3, was the map scoring tool, which aims at automatically highlighting problem areas of the map, if any — such as clutter, or registration errors. One use case for this method is to assist a deployment engineer checking to verify if the map constructed in the previous stage is consistent and useful. Figure 3 shows example output of a prototype implementation of this tool. The colors can guide attention to potential problem areas of the map. In the bottom right image of Figure 3, please note that the clutter has been effectly removed, and the wall near the bottom of the excerpt is indicated as highly likely. Also, a door, which shows in the map as a small inset in the wall near the right edge of the image, is scored as less likely. While this is not an error, but a real and permanent feature of the map, its appearance is close to what would be expected from a registration error and thus warrants closer inspection. A more recent implementation and report has since been published online.[1]

3.2 Task planning

Similar to MS2, we demonstrated automated task planning when a simulated order is received. Given a product list for the order, the OPT Loading software provided by ACT Operations Research plans how each object should be stacked on each pallet for efficiency and stability. The fleet management layer then sends tasks to the robots in the fleet, going to pick pallets and products in order, to fill the pallet as planned.

A new addition that was shown at MS3, over what was shown already at MS2, was dynamic task re-allocation on the fly, as requirements change. This can happen as unexpected delays or faults are encountered, so that the tasks of the fleet must be rescheduled

[1] See paper at https://arxiv.org/abs/2004.08794 and code at https://github.com/tkucner/rose.

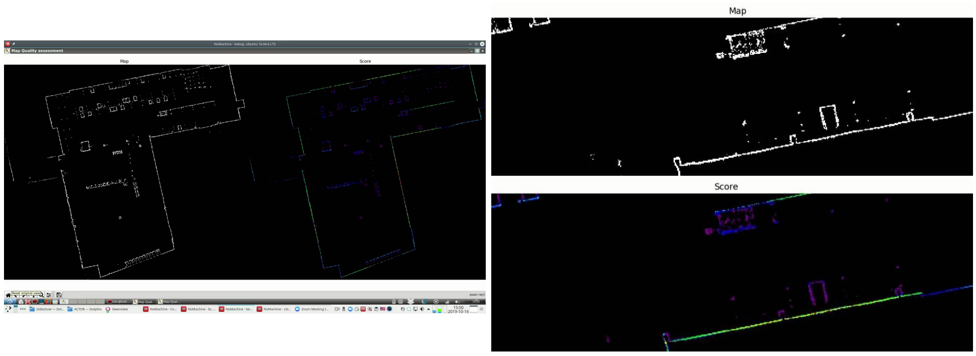

Figure 3: Prototype of the map quality assessment tool. *Left:* Complete map of the warehouse, and per-pixel map quality scores. *Top right:* Zoom-in on a smaller region of the map. Free space is black, and walls and obstacles are white. *Bottom right:* Per-pixel map quality scores. Brighter (yellow/green) parts are "likely" and dimmer (blue/violet) parts are "suspicious", as they do not support the main structure extracted from the map.

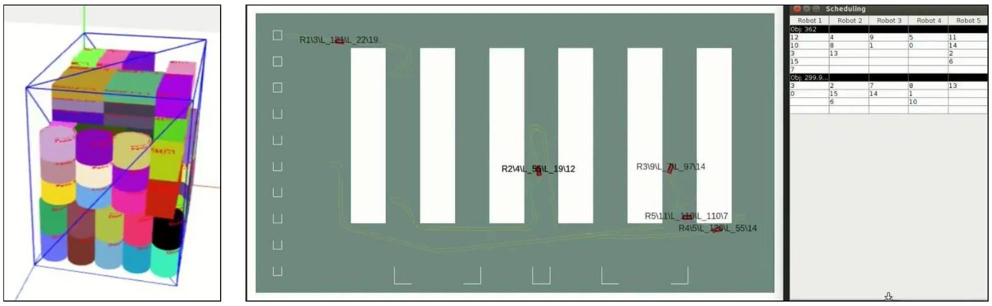

Figure 4: Task allocation. *Left:* Pallet planning, given a set of products to go in an order. *Right:* Robots fulfilling the order. The two tables shown to the right indicate a sequence of two task schedules. The second one has been created automatically, following an unexpected delay of one of the trucks.

or shifted to other robots. This part was shown in simulation, in order to better demonstrate the capabilities in an environment with more robots, and a larger warehouse, than what was physically available at the demo site. Figure 4 shows an excerpt of this demo.

3.3 Human-aware motion planning and coordination

We demonstrated that the trucks can plan paths in the environment and *coordinate* their motion so that they automatically give way in sections where their paths overlap spatially.

New to MS3 was that we demonstrated the integration of *flow-aware* mapping, motion planning, and coordination. We also demonstrated coordination of two different truck platforms: one large Toyota BT SAE200, and one small Linde CiTi truck (both of which can also be seen in Figure 1). An example of this is shown in Figure 5. The fleet can learn statistical human motion patterns from observation while operating, given input from the human tracking module (see Section 3.5). In ILIAD, we have developed two flow-aware map representations: STeF-map (which can be used to learn and predict periodic patterns in a discretised motion space) and CLiFF-map (with models locally continuous motion patterns as stationary distributions), and motion planning algorithms that take the "flow cost" of these representations into account when planning how to reach a point. Figure 5

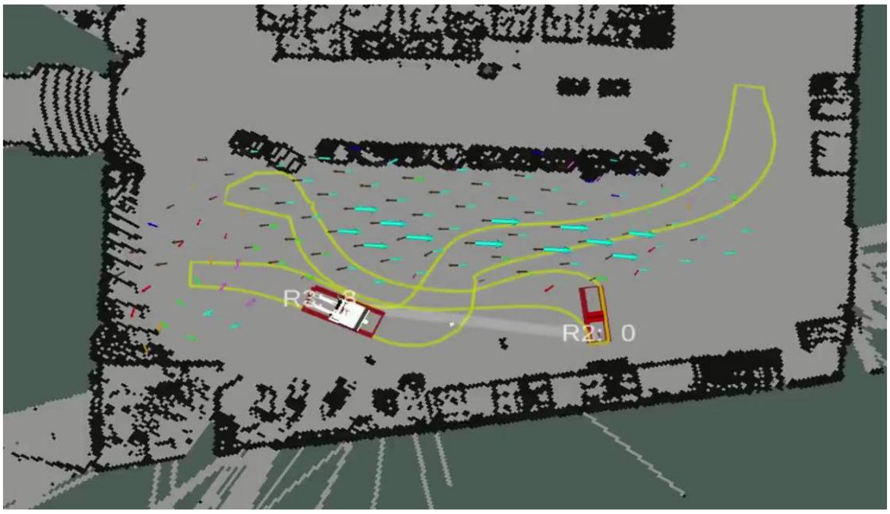

Figure 5: Human-aware motion planning and coordination. The arrows in the map indicate statistical flow patterns that have been learned from observations using sensors mounted on one of the trucks. The yellow path envelopes show the planned paths of the two robots R2 and R3 (including the area they will sweep while turning). R3 plans to go from left to right in this image, and R2 from right to left. Both robots plan paths that, to the extent possible, avoids the expected flow of people as well collisions with other robots and obstacles. The arrow pointing from R2 to R3 indicates that R2 will not enter the region where the paths intersect until R3 has passed. These constraints are computed on the fly by the coordinator module.

shows a (black/grey) geometric navigation map from the warehouse, overlaid with a flow map (coloured arrows) that has been constructed from motion data observed in the warehouse. Also shown are the flow-aware motion plans of two robots, trying to avoid the expected flow as much as possible.

3.4 Manipulation

The manipulation demonstration was shown via video link from Pisa, as the arms will not be physically mounted on a robot for fully integrated manipulation until the final Milestone 4 demo.

The two main additions to the MS3 demo over last year's demo was (1) more advanced plastic cutting, also taking into account objects whose position is not exactly known in advance, and objects with uneven shapes such as the one in Figure 6, and (2) integrating ORU's object detection module with the manipulation planning and execution from UNIPI.

Also demonstrated via live video link, from the Technical University of Munich, was the "safe motion unit" (SMU) which enforces biomechanically safe motion of the robot arm in the vicinity of people. Also here, a new addition compared to the MS2 demo was integration of computer vision so that the SMU is activated automatically when a human enters the robot arm's reachable space. See Figure 7. This system has been reported in D4.2.

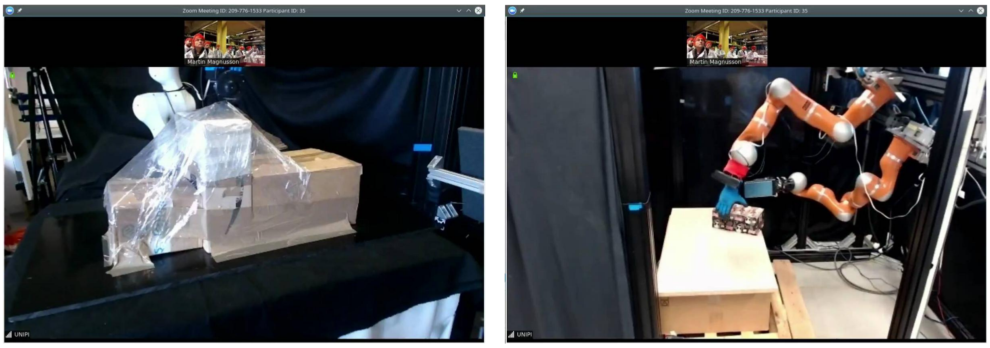

Figure 6: Manipulation demo with live streaming over video link. *Left:* Automatic cutting the plastic wrapping around a set of products, in preparation of picking. *Right:* Picking an object from an unknown position, showing integration with computer vision-based object detection.

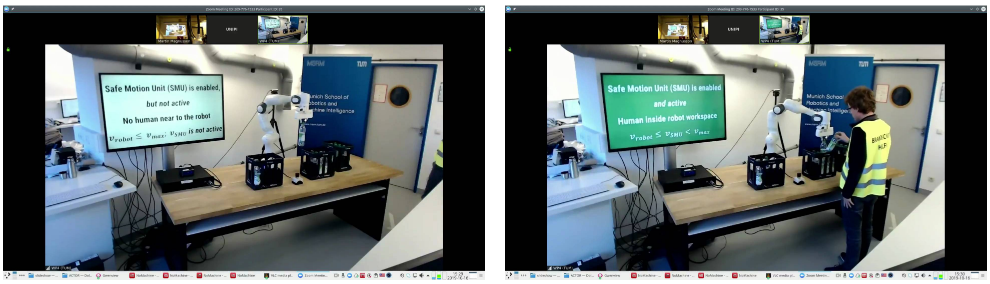

Figure 7: Safe motion unit (SMU) for maximising speed and safety in robot manipulation close to humans. As a person enters the reach of the arm, its velocities are constrained to safe limits.

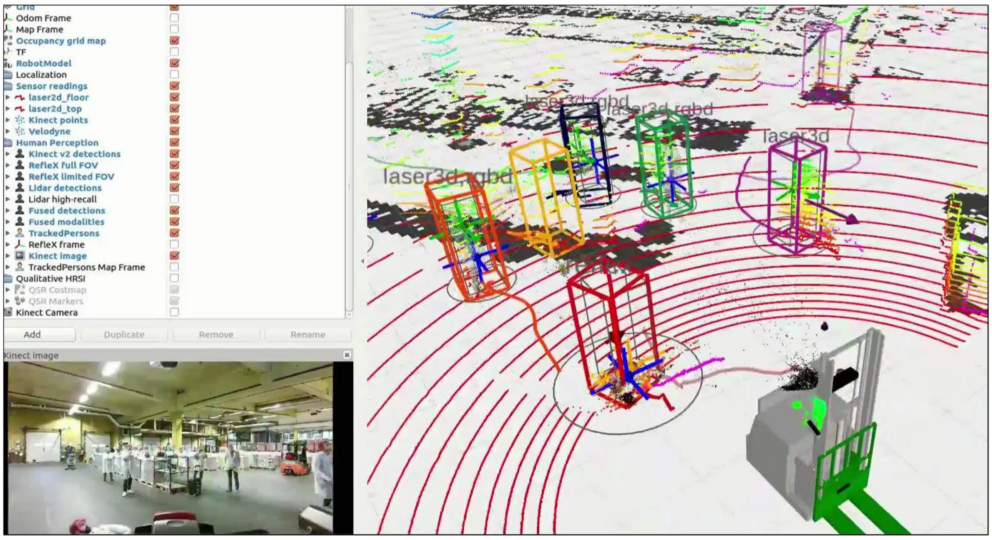

Figure 8: Example output of the ILIAD people detection and tracking framework. Detected people are circumscribed with boxes, including coloured tracks for their past motion and arrows for their predicted future direction of motion. Each detected person is also labelled with the sensor modalities contributing to the detection. The different sensors cover different fields of view and ranges.

3.5 People detection and tracking

In the demonstration of the people detection and tracking stack used and developed in ILIAD, we now showed full integration of all detectors: including 2D and 3D laser, RGB-D camera, and the Retenua "emitrace" camera for reflex vest detection. Figure 8 shows an excerpt of this demo, also showing the complementariness of the sensor modalities. In this particular example, the emitrace reflex camera is the only one that detects the lying person; the RGB-D and 3D laser, but not the emitrace, detect the persons without safety vests in front of the robot; and only the 3D laser detects far away persons outside of the field of view of the cameras.

3.6 Intention communication and social navigation

In the demo, we showed for the first time the integration of the qualitative human–robot spatial interaction module from Task T3.4. Figure 9 shows a representation of the robot state using "qualitative trajectory calculus" to reason about the relative motion of the robot and the nearest human. Depending on the motion of the robot and the human ("approaching from left/right", "stopping", etc.), the robot will, in the final version, also be able to take appropriate action to produce both legible and collision-free paths.

Finally, we demonstrated mutual communication of intents via, respectively, eye-tracking glasses on a person and a projector on one of the robots. One of the robots is equipped with a projector that displays an arrow on the floor in front of it, indicating its current movement/turning direction. The integration with the eye-tracking glasses is implemented such that when the person is not looking at the robot, the arrow is blinking in order to get the attention of the person; and when the person is observing the robot, the arrow is displayed statically.

This work has been reported in more depth in Deliverable D3.3.

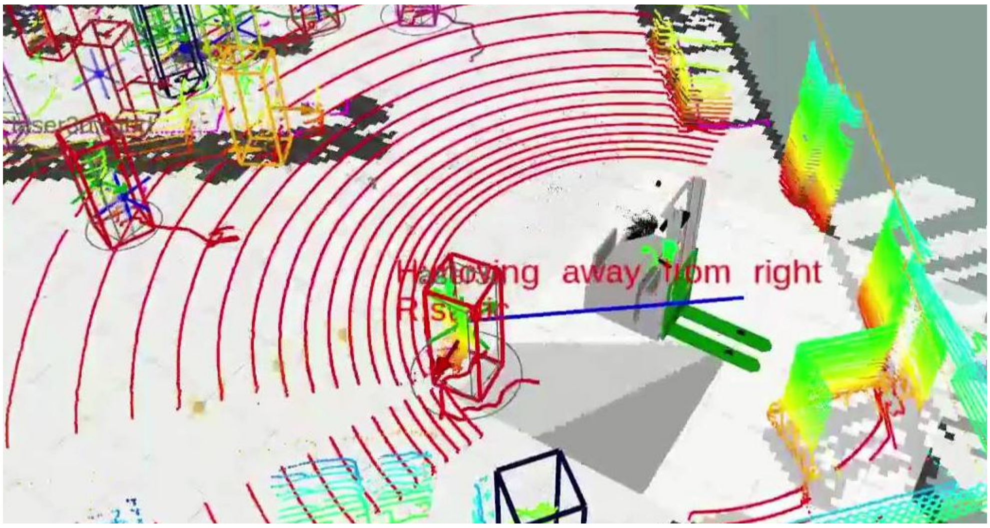

Figure 9: Qualitative human–robot spatial interaction. The shaded sections on the floor mark different costs for moving in different directions, computed from the relative motion of the robot and the detected person. Darker sections are less preferred.

4 Summary

In conclusion, this milestone demonstration showed the most important capabilities of the ILIAD system, most of which were integrated on the same physical fleet, running live at the end-user site; while the manipulation components were demonstrated via live demonstrations in separate locations. At the final Milestone 4 demo, we will demonstrate full integration also of manipulation with the rest of the system, and final version of the remaining components.