Collaborative 3D Reconstruction Using Heterogeneous UAVs: System and Experiments

Timo Hinzmann[1(✉)], Thomas Stastny[1], Gianpaolo Conte[2], Patrick Doherty[2],
Piotr Rudol[2], Marius Wzorek[2], Enric Galceran[1], Roland Siegwart[1],
and Igor Gilitschenski[1]

[1] Autonomous Systems Lab, ETH Zurich, Zurich, Switzerland
{timo.hinzmann,thomas.stastny,enric.galceran,roland.siegwart,
igor.gilitschenski}@mavt.ethz.ch
[2] Department of Computer and Information Science, Linköping University,
Linköping, Sweden
{gianpaolo.conte,patrick.doherty,piotr.rudol,marius.wzorek}@liu.se
http://www.asl.ethz.ch,
http://www.liu.se

Abstract. This paper demonstrates how a heterogeneous fleet of unmanned aerial vehicles (UAVs) can support human operators in search and rescue (SaR) scenarios. We describe a fully autonomous delegation framework that interprets the top-level commands of the rescue team and converts them into actions of the UAVs. In particular, the UAVs are requested to autonomously scan a search area and to provide the operator with a consistent georeferenced 3D reconstruction of the environment to increase the environmental awareness and to support critical decision-making. The mission is executed based on the individual platform and sensor capabilities of rotary- and fixed-wing UAVs (RW-UAV and FW-UAV respectively): With the aid of an optical camera, the FW-UAV can generate a sparse point-cloud of a large area in a short amount of time. A LiDAR mounted on the autonomous helicopter is used to refine the visual point-cloud by generating denser point-clouds of specific areas of interest. In this context, we evaluate the performance of point-cloud registration methods to align two maps that were obtained by different sensors. In our validation, we compare classical point-cloud alignment methods to a novel probabilistic data association approach that specifically takes the individual point-cloud densities into consideration.

Keywords: Collaborative UAV mapping missions · Point-cloud generation · Vision-laser point-cloud alignment · Delegation of heterogeneouse agents

1 Introduction

Field robotics has seen great gains in recent years owing both to robustified robotic platforms and increasing autonomous behaviors and capabilities. In particular, autonomous unmanned aerial vehicles (UAVs) of various classes, utilizing

D. Kulić et al. (eds.), *2016 International Symposium on Experimental Robotics*,
Springer Proceedings in Advanced Robotics 1, DOI 10.1007/978-3-319-50115-4_5

state-of-the-art perceptive sensors and sensing techniques, have proven worth in both large- and small-scale mapping applications, providing a wide array of sensor data. In large-scale mapping scenarios, recent developments in solar-powered, fixed-wing UAV (FW-UAV) technology have enabled extreme long-endurance for low-altitude coverage of vast areas in a compact and hand-launchable form [1]. Finer-resolution mapping on a smaller scale has also been demonstrated using aerial laser scans from autonomous helicopters [2]. Fast and fully autonomous generation of up-to-date maps could potentially be a great advantage for rescue workers looking for missing persons, or in disaster management scenarios, like floods, forest fires, and earthquakes. However, a crucial element to the utility of such operations is the ease of use for, possibly, non-technical operators. Further, no single UAV is a one-fits-all solution for the wide array of sensing data that may be required by end users. In these cases, a robotic team of various actors with mixed, but complementing, capabilities, working together within the context of a collaborative, cognitive framework, on a higher-abstraction, would be particularly impactful.

2 Problem Statement: Collaborative 3D Reconstruction

Point-cloud generation from optical cameras on a large scale from small FW-UAVs is sparse, due to, relatively high flying altitudes and limited image resolution. Contrarily, laser point-clouds generated from low-flying autonomous helicopters are dense, but only cover a small area. Merging these two data types together into a single global map has obvious benefits in the sense of real-world search and rescue or disaster management operations, where a large scale (sparse) map could provide operators with coarse information and a means to select "areas of interest" to send agents for a "closer look". This closer look would provide dense maps of smaller areas which, when merged with the global map, results in a more accurate representation of the environment for both, human operators and collaborating robotic actors. Leveraging the various capabilities of each participating agent in an autonomous manner also requires a higher abstraction of task delegation. In sum, we show a real-world demonstration of distributed, autonomous map making and vision-to-laser point-cloud registration from differing aerial views and mixed sensor data.

In this context, we employ a novel probabilistic data association method [3] that robustly aligns two maps that were generated by different sensors. Compared to [3], we see the following major contribution: Our data was recorded on different agents, sensor units and flight paths and hence represents a real-world scenario. In contrast, the data presented in [3] was recorded by the same agent and/or even with the same sensor which simplifies the registration process[1].

[1] Due to space constraints in this publication, only a subset of the data can be presented. However, the datasets can be requested from the authors.

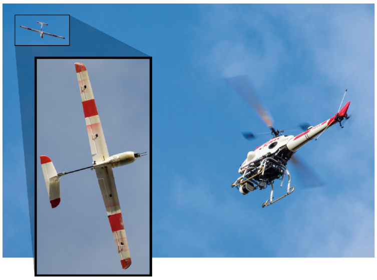

Fig. 1. Two UAV platforms during their cooperative scanning mission.

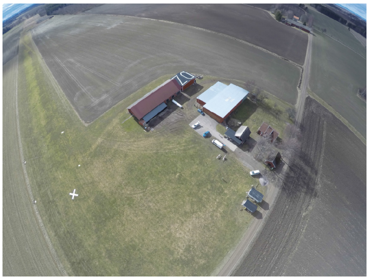

Fig. 2. Aerial image of test site near Motala, Sweden.

3 Technical Approach

This section opens with an overview of the delegation framework. For more details we refer to the companion paper [4]. The section proceeds with a description of the state estimation, point-cloud generation and concludes with a focus on the point-cloud registration methods.

3.1 Mission Process and Delegation Framework

A high-level depiction of the mission process is provided in Fig. 3. The delegation framework [5] which includes delegation modules from each of the participating agents, provides both a formal and software infrastructure for specifying and generating collaborative multi-agent plans to achieve complex goals such as multi-UAV 3D reconstruction of selected regions. Delegation is based on a recursive algorithm that sends requests of the following type, Delegate(*Agent1, Agent2, Task, Context*), where *agent1* makes an attempt to delegate *Task* to *Agent2*, given a specific *Context* specified as a set of constraints. Examples of constraints would be temporal constraints or restrictions on flight altitudes. Agents can be humans or robots. Tasks are represented using Task Specification Trees (TSTs). TSTs have both declarative and procedural descriptions. Internal nodes in a TST represent control modes such as sequence and concurrency, while leaf nodes represent domain dependent elementary tasks executable by different participating platforms. The delegation process, as illustrated in Fig. 4, itself begins with a goal request TST often provided by a human operator and if successful, results in an expanded TST where all constraints are satisfied. Sub-trees in the final TST are also appropriately allocated to those platforms with the proper capabilities. TSTs can be generated dynamically using automated planning techniques, or by using generic TST templates that can be instantiated appropriately. In this example (Fig. 4), a concurrent scanning plan is generated for one region where two separate sub-regions are covered by each of the UAVs involved, respectively. The scan-map task in the TST calls a region partitioning algorithm to determine appropriate sub-regions for platforms to scan based on their capabilities.

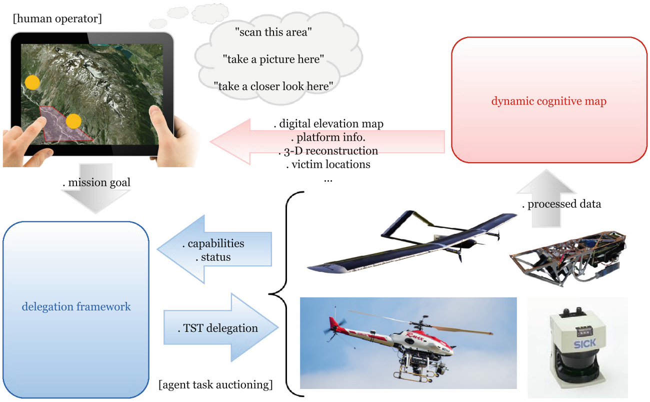

Fig. 3. Mission process: a human operator broadcasts a goal request for a data acquisition mission via its delegation module. Platforms with available capabilities reply and a delegation process ensues among each of the platforms' delegation modules. If successful, the net result is a joint plan to execute. Upon execution, raw/processed data can be stored locally or globally. During the mission or upon mission completion the human operator can access the results via specialized interfaces.

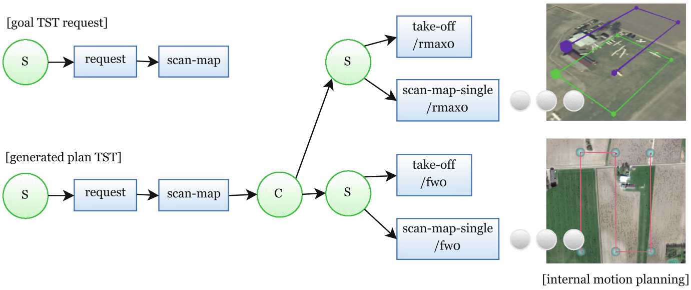

Fig. 4. Goal TST request from operator and generated plan TST involving both RMAX (/rmax0) and FW-UAV (/fw0). Internal nodes: (C) concurrent, (S) sequence.

The delegation process itself is quite complex and involves auctions, constraint solving and dynamic TST expansion. The scan-map task involves use of partitioning algorithms and the scan-map-single tasks involve internal path planning by the respective platforms. During the mission execution phase, each system executes its part of the mission TST relative to timing and other constraints.

3.2 State Estimation and Point-Cloud Generation

RW-UAV. The state estimation is used both for autonomous navigation and for point-cloud generation by incorporating laser scanner measurements in form of a direct georeferencing technique [6]. It is based on a Kalman filter algorithm which fuses inertial and GNSS position data. The deployed Kalman filter uses a linear state-space error dynamic model derived from a perturbation analysis of the equations of motion [7]. The Kalman filter produces state estimation at 50 Hz rate and performs the update step using GNSS measurement at 20 Hz rate.

FW-UAVs. The Pixhawk PX4 auto-pilot performs an indirect EKF-based state estimation as presented in [8]. Within the Kalman filter linear acceleration and angular rates measurements are used for propagation of the system state. Pressure, GPS velocity and position, as well as magnetometer measurements are used for the state update [8]. The estimated states involve the IMU's attitude and position in WGS84 coordinates[2]. The vision point-cloud was acquired with an optical camera using a classical photogrammetric approach[3].

3.3 Point-Cloud Registration

The alignment of point-clouds generated from two unmanned aerial vehicles with different sensors involves consideration of the following challenging aspects: Firstly, one point in the visual source point-cloud does in general not correspond to a point in the laser target cloud and vice versa. Secondly, the sensor noise models are different: Peaky for laser but more spread for visual points due to camera noise and triangulation errors. Thirdly, the laser point-cloud is in general denser than the vision point-cloud. Furthermore, the robots fly at high altitudes. Consequently, a dominant ground plane and few depth discontinuities are common for most datasets. Lastly, a rough initial alignment is given by global positioning systems such as GNSS or fused from the state estimator. To register a visual sparse point-cloud to a dense laser point-cloud the following registration algorithms are evaluated with respect to the problem specifications described above: Iterative Closest Point (ICP), Iterative Probabilistic Data Association (IPDA), Generalized Iterative Closest Point (GICP), and Normal Distribution Transform (NDT).

One iteration of the Probabilistic Data Association [3] approach consists of the following steps: For every point of the source cloud a kd-tree search is performed with maximal radius $r_{kd,}$ and maximal number of returned neighbours n_{kd} as shown in Fig. 5. For every source-target correspondence, the residuals and weights are calculated as illustrated in Figs. 6 and 7 respectively employing expectation-maximization (EM) in combination with e.g. a Gaussian or t-distribution. The red line indicates the evolution of the true correspondence residual and weight for 28 Levenberg-Marquardt optimization steps. Note that the data associations do not change during one iteration but only the residuals and weights update based on the iterative solution of the Levenberg-Marquardt algorithm. These steps can be performed iteratively to increase the area of convergence (IPDA). The advantages of this approach relevant to the problem specifications are the following: (1) a sensor model can be intuitively inserted in the EM-algorithm based on the expected noise, (2) a point of the sparse source cloud

[2] For more details about the state estimation framework we refer to [8].

[3] The vision point-clouds are generated using the commercial software pix4d.

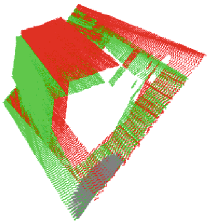

Fig. 5. Correspondences (grey) for one source point obtained by kd-tree.

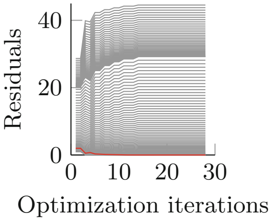

Fig. 6. Residuals for one source point after one iteration.

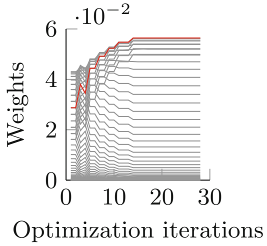

Fig. 7. Weights for one source point after one iteration.

can hold correspondences to many points in the target cloud. Iterative Closest Point (ICP) [9–12] is a widely used registration algorithm that has inspired many variants. For the evaluation we use the classic point-to-point approach implemented in the Point Cloud Library (PCL). It performs the following steps until convergence: (1) For every point in the source cloud find *the* closest point in the target cloud. (2) Estimate and apply the transformation **T** that best transforms the source cloud to the target cloud in the sense of a mean squared error. Naturally, ICP's assumption that one point in the source cloud has an exact correspondence in the target cloud is not fulfilled in the sparse-dense registration problem. Generalized ICP (GICP) [13] levers the classic ICP and point-to-plane ICP into a probabilistic framework. Applied to the aerial registration problem, GICP may profit from dominant ground planes due to the high flying altitudes. In the Normal Distribution Transform [14] the points of the cloud are represented in form of a probability distribution and hence no explicit point correspondences between source and target cloud are established. With regard to the sparse-to-dense point-cloud registration, NDT may fail if the visual cloud is too sparse. The parameter notations of the individual methods are presented in Table 1.

Table 1. Parameter notations for IPDA, ICP, GICP and NDT.

	Parameter	Description
IPDA	r_{kd}	Kd-tree radius
	n_{kd}	Kd-tree max. neighbours
	$iter_{max}$	Max. number of iterations
	ls	Leaf size of voxel grid
ICP	d_c	Max. correspondence distance
	ϵ_T	Transformation conv. criteria
	$iter_{max}$	Max. number of iterations
	ls	Leaf size of voxel grid
GICP	d_c	Max. correspondence distance
	ϵ_T	Transformation conv. criteria
	$iter_{max}$	Max. number of iterations
	ls	Leaf size of voxel grid
NDT	Δx	Step size
	Δr	Resolution
	ϵ_T	Transformation conv. criteria
	$iter_{max}$	Max. number of iterations
	ls	Leaf size of voxel grid

4 Platform Description

The platforms used for the experiments include a rotary-wing Yamaha RMAX and the two fixed-wing UAVs named Techpod and senseSoar.

4.1 RW-UAV

The Yamaha RMAX helicopter [15], shown in Figs. 1 and 3, has a rotor diameter of 3.1 m, a maximum take-off weight of 94 kg and a payload capability of about 30 kg. The platform is capable of fully autonomous navigation, including take-off and landing. The basic sensor suite used for autonomous navigation includes a fiber optic tri-axial gyro system and a tri-axial accelerometer system, a RTK GNSS positioning system and an infrared altimeter used for automatic landing. Onboard sensors used for mapping missions include color and thermal video cameras, as well as a class 1 SICK LMS511 PRO 2D laser scanner. The laser scanner's maximum range is 80 m with a maximum scanning FoV of 190°.

4.2 FW-UAVs

Techpod. The small unmanned research plane Techpod is shown in Fig. 1. It has a classic T-tail configuration, is equipped with one propeller, has a wingspan of 2.60 m and a nominal speed of around 12 m/s. The sensor and processing unit [1] as well as the PX4 auto-pilot are located inside the modified fuselage and allow autonomous mission execution such as GPS waypoint following.

SenseSoar. The highly versatile solar-UAV senseSoar was developed at the Autonomous Systems Lab for search & rescue missions and has a wingspan of 3.1 m. With its solar panels it is able to generate an electric power of around 140 W and has shown long-endurance capabilities. Likewise as Techpod, senseSoar is hand-launchable and carries the sensor pod inside the fuselage.

5 Experimental Results

The datasets were collected at two locations: (1) in Motala, Sweden which includes a flight field with several houses and trees (Figs. 2 and 8). The resulting experiments are presented in Sects. 5.1 and 5.2; (2) at a mountainside in Isollaz, Italy as presented in Experiment III in Sect. 5.3.

5.1 Experiment I: Complementary Factor of Vision-Laser Point-Cloud Alignment

In a first experiment, the vision point-cloud generated with images recorded by a Sony ActionCam HDR-AS100V mounted on Techpod is aligned to the laser point-cloud generated with a SICK laser scanner onboard of the RMAX.

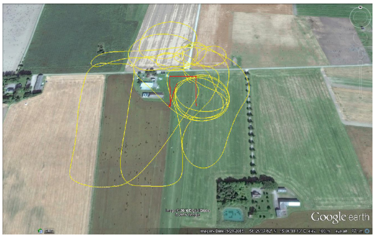

Fig. 8. Sample flight path of the FW-UAV (yellow) for Exp. I and II: The FW-UAV is loitering in-air until it receives the command for scanning the area from the delegation framework. Based on this request, the path-planner located on the ground station generates a scanning pattern which is transmitted to the FW-UAV via telemetry. After execution, the imagery is sent to the ground station via WiFi where the point-cloud is generated. The path of the RW-UAV is plotted in red. The nominal altitude of the FW- and RW-UAV is 100 m and 48 m.

Figures 9, 10 and 11 show a satellite image, colored vision and laser point-cloud of the region of interest. Note that the laser did not receive response for parts of the roof and barn due to a steep observation angle, relatively low altitude of the RMAX, and due to non-reflective surfaces. On the other hand, the laser point-cloud contains less measurement noise and a higher level of detail as can be best seen in Fig. 12 which e.g. depicts a wind vane in the top right corner not observed by the visible light camera. These observations underline the complementary factor of the two point-clouds which, when aligned, result in a more complete model of the environment. Figures 12 and 13 illustrate the initial misalignment from side and top view respectively. This georeferencing error is given in Table 2 and consists of a translational offset of several meters and a small rotation. The transformation was obtained by careful manual alignment of the point-clouds and used to evaluate point-cloud registration methods quantitatively. Note that due to the noisy character of the data, this manual alignment should not be considered perfect as slightly varying alignments seem still visually satisfying. Nevertheless, this method allows to reason about convergence and general trends. Figures 14 and 15 show the transformation error for IPDA and ICP plotted over the number of iterations. Both ICP and IPDA converge

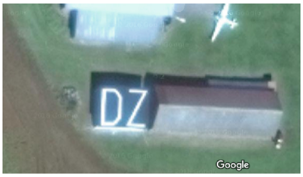

Fig. 9. Satellite image as reference.

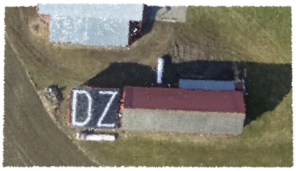

Fig. 10. Vision point-cloud.

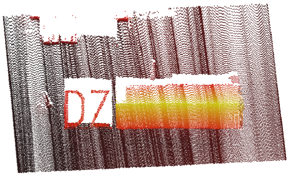

Fig. 11. Laser point-cloud colored by height

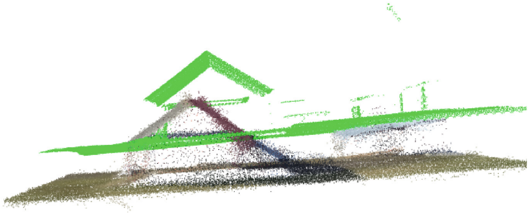

Fig. 12. Side-view: vision point-cloud colored by pixel intensity and laser point-cloud in green.

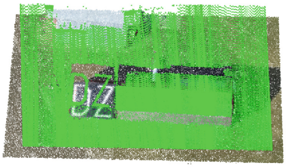

Fig. 13. Top-view: vision point-cloud colored by pixel intensity and laser point-cloud in green.

Table 2. Initial misalignment transformation error and final translational and rotational offset for IPDA, ICP, GICP and NDT. The translation error e_{trans} and rotation error e_{rot} are computed as proposed in [12].

t_x	-1.05 m
t_y	3.52 m
t_z	6.46 m
φ	-0.0162 rad
θ	0.0152 rad
ψ	-0.0206 rad

	Parameters	e_{trans}	e_{rot}	iter.
IPDA	$r_{kd} : 5.0, n_{kd} : 50, iter_{max} : 1000, ls : 0$, student-t	0.5160	0.0129	34
ICP	$d_c : 10, \epsilon_T : 10^{-16}, iter_{max} : 500, ls : 0$	0.1992	0.0091	100
GICP	$d_c : 10, \epsilon_T : 10^{-16}, iter_{max} : 500, ls : 0$	4.075	0.0473	6
NDT	$\Delta x : 0.1, \Delta r : 0.1, \epsilon_T : 10^{-16}, iter_{max} : 1000, ls : 0$	1.8213	0.025	1000

to almost the same transformation. Furthermore, from the given plots it can be seen that the altitude offset converges first in very few iterations due to the dominant ground planes. The translational offset in x and y usually needs more iterations to converge. Figure 17 shows the aligned vision and laser point-cloud using IPDA. Figure 16(a)–(d) show the final alignments computed by the individual registration methods. The figures and Table 2 illustrate that all methods show convergence, however, small misalignment errors are visible for GICP and NDT.

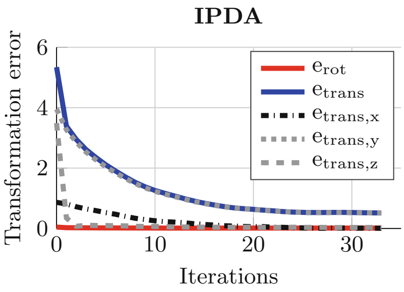

Fig. 14. IPDA: rotational (red) and translational (blue) misalignment error.

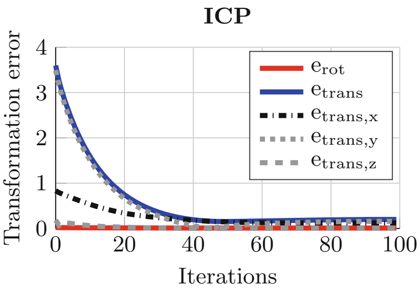

Fig. 15. ICP: rotational (red) and translational (blue) misalignment error.

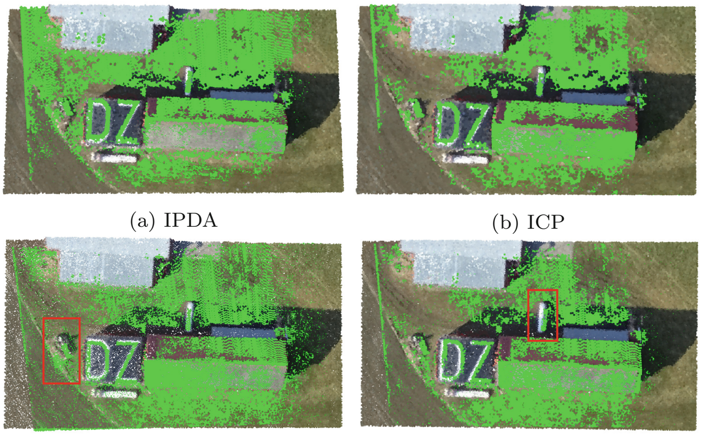

(a) IPDA

(b) ICP

(c) GICP

(d) NDT

Fig. 16. (a)–(d) illustrate that all registration methods converged, however, small misalignment errors are observable for GICP and NDT depicted by the red rectangles.

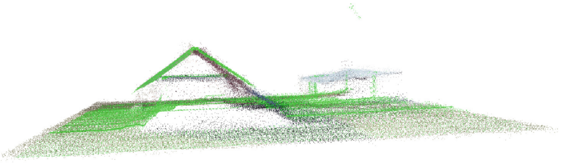

Fig. 17. Aligned vision and laser point-cloud (green) using IPDA. The experiment underlines the complementary factor of employing laser and vision to obtain a more complete model of the environment.

5.2 Exp. II: Changes in the Environment

This experiment evaluates if agents that possess only poor absolute position sensing capabilities can register to an a-priori obtained and well georeferenced point-cloud. This evaluation gives an idea of how well the different methods can deal with changes in the environment as well as about their region of convergence.

For this purpose, we align the vision point-cloud shown in Fig. 19 to the previously generated laser point-cloud given in Fig. 18. Several changes in the environment can be spotted, in particular, the vegetation, location of cars and of a small house. Furthermore, we generate a random large initial misalignment error between both point-clouds as shown in Table 3. Figures 20, 21 and 22 as well as Table 3 demonstrate that IPDA, in particular when employing t-distribution, results in the lowest final misalignment error, followed by GICP and ICP, whereas NDT diverged for this scenario.

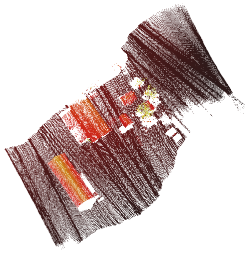

Fig. 18. Point-cloud generated by the laser scanner mounted on the RW-UAV. The point-cloud consists of 568’839 points and is colored by height.

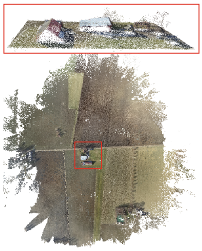

Fig. 19. Point-cloud generated by the RGB camera mounted on the FW-UAV. The house area shown in the top consists of 163’595 points.

Table 3. Initial misalignment transformation error and final translational and rotational offset for IPDA, ICP, GICP and NDT.

t_x	11.44 m
t_y	12.97 m
t_z	−8.32 m
φ	0.025 rad
θ	−0.0013 rad
ψ	−0.0411 rad

	Parameters	e_{trans}	e_{rot}	iter.
IPDA	$r_{kd}: 5.0, n_{kd}: 200, iter_{max}: 200, ls: 1.5,$ student-t	0.1152	0.011	74
	$r_{kd}: 5.0, n_{kd}: 200, iter_{max}: 200, ls: 1.5,$ Gaussian	3.3278	0.0243	80
ICP	$d_c: 10, \epsilon_T: 10^{-16}, iter_{max}: 500, ls: 1.5$	6.1891	0.0278	190
GICP	$d_c: 10, \epsilon_T: 10^{-16}, iter_{max}: 500, ls: 0.1$	2.6757	0.0201	33
NDT	$\Delta x: 0.1, \Delta r: 1.0, \epsilon_T: 10^{-16}, iter_{max}: 500, ls: 0$	22.2777	0.1126	500

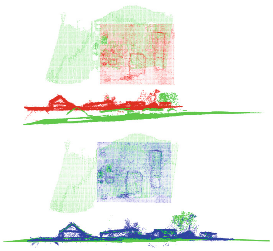

Fig. 20. Initial misalignment and final registration using IPDA with t-distribution.

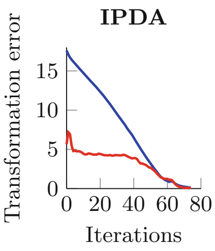

Fig. 21. IPDA: translation (blue) and rotation error (red; multiplied by $1e2$ for better visualization).

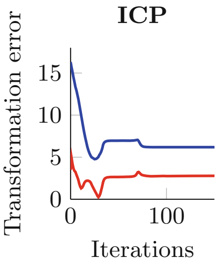

Fig. 22. ICP: translation (blue) and rotation error (red; multiplied by $1e2$ for better visualization).

5.3 Exp. III: Point-Cloud Sparsity

In Experiment III, we present a very challenging dataset consisting of a tree region with few man-made structures. The mission procedure is illustrated in Fig. 23: The FW-UAV, equipped with a Sony ActionCam and a grayscale Aptina MT9v034 camera, generates a rough initial point-cloud as soon as it receives the command from the delegation framework initiated by the human operator. Subsequently, the RW-UAV scans the region of interest with the aid of the SICK

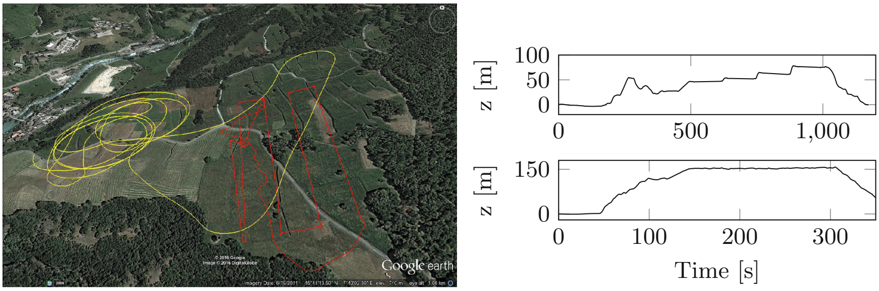

Fig. 23. The satellite image shows the flight path of the fixed-wing UAV in yellow and of the RMAX helicopter in red. The plots on the right show the altitude (top: RW-UAV, bottom: FW-UAV) in form of height above ground with respect to the individual starting positions.

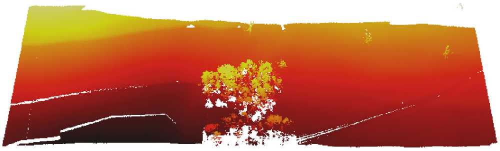

Fig. 24. One of the laser strips to be aligned to the vision point-cloud.

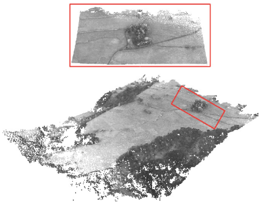

Fig. 25. Point-cloud generated by the grayscale camera mounted on the FW-UAV.

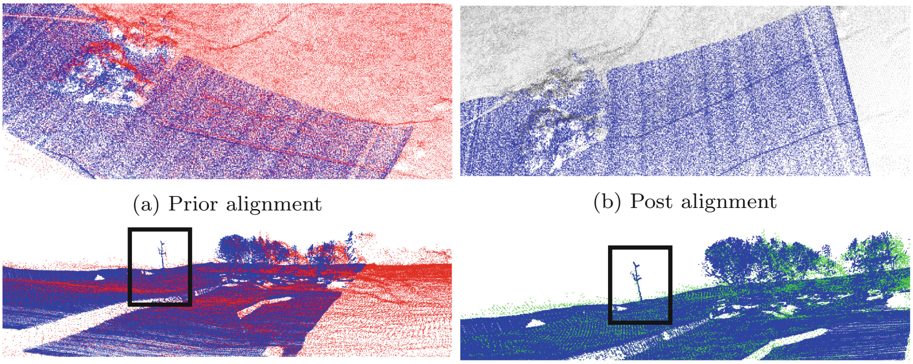

(a) Prior alignment

(b) Post alignment

(c) Prior alignment

(d) Post alignment

Fig. 26. The dense laser point-cloud is shown in blue. The prior misalignment is especially visible at the hill gradients. The tree region was only partially captured by the FW-UAV and could be densified by the laser scan.

laser for refinement. The generated laser and vision point-clouds are shown in Figs. 24 and 25 respectively. We deliberatively use the point-cloud generated by the low-resolution grayscale camera for point-cloud registration and employ IPDA to underline its performance when dealing with dense and sparse point-clouds. In contrast to Experiment I and II, the images are geo-registered by the onboard EKF instead of using the raw GPS measurements. As expected, the initial misalignment error is limited to $3.34, -2.51, -0.26$ m for the translational and $-0.016, 0.0082, 0.0089$ rad for the rotational offset. The initial misalignment and final registration are shown in Fig. 26.

6 Conclusion

In this paper, we presented an automated delegation framework that translates top-level commands of the human operator into low-level commands of the employed agents. We validated the framework based on realistic scenarios in two locations including more than 20 flights using a RW- and FW-UAV representing an arbitrary fleet of heterogenous agents. Furthermore, we chose the task of scanning a common area as *one* exemplary mission of the delegation framework. The point-clouds acquired during this scanning process are automatically registered and transferred back to the human operator and visualized in the dynamic cognitive map. Our experiments show the complementary factor of vision-laser point-cloud registration from aerial views and demonstrate the successful deployment of the Probabilistic Data Association algorithm. The final goal of this project will be to allow accurate path planning of unmanned ground vehicles (UGV) or smaller multicopter UAVs based on the aligned map and, for instance, to delegate them inside the buildings' interior. Future work will also include the integration of a previously presented human detection algorithm [16] into the delegation framework. The algorithm returns the UTM location of possible victims along with their detection uncertainties. Other agents may verify these possible human detections to decrease the false alarm rate.

Acknowledgment. The research leading to these results has received funding from the European Commission's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°285417 (ICARUS) and n°600958 (SHERPA). Furthermore, the authors want to thank Gabriel Agamennoni and Simone Fontana for providing an initial implementation of the probabilistic data association algorithm.

References

1. Oettershagen, P., Stastny, T., Mantel, T., Melzer, A., Rudin, K., Gohl, P., Agamennoni, G., Alexis, K., Siegwart, R.: Long-Endurance Sensing and Mapping Using a Hand-Launchable Solar-Powered UAV. In: Wettergreen, D.S., Barfoot, T.D. (eds.) Field and Service Robotics. STAR, vol. 113, pp. 441–454. Springer, Heidelberg (2016). doi:10.1007/978-3-319-27702-8_29

2. Conte, G., Rudol, P., Doherty, P.: Evaluation of a light-weight lidar and a photogrammetric system for unmanned airborne mapping applications. PFG Photogrammetrie, Fernerkundung, Geoinformation **2014**, 287–298 (2014)
3. Agamennoni, G., Fontana, S., Siegwart, R.Y., Sorrenti, D.G.: Point clouds registration with probabilistic data association. In: Proceedings of the International Conference on Intelligent Robots and Systems. IEEE (2016)
4. Doherty, P., Kvarnström, J., Rudol, P., Wzorek, M., Conte, G., Berger, C., Hinzmann, T., Stastny, T.: A Collaborative Framework for 3D Mapping Using Unmanned Aerial Vehicles. In: Baldoni, M., Chopra, A.K., Son, T.C., Hirayama, K., Torroni, P. (eds.) PRIMA 2016. LNCS (LNAI), vol. 9862, pp. 110–130. Springer, Heidelberg (2016). doi:10.1007/978-3-319-44832-9_7
5. Doherty, P., Heintz, F., Kvarnström, J.: High-level mission specification and planning for collaborative unmanned aircraft systems using delegation. Unmanned Syst. **1**(1), 75–119 (2013)
6. Skaloud, J.: Optimizing georeferencing of airborne survey systems by INS/DGPS. Ph.D. thesis, University of Calgary, Alberta (1999)
7. Britting, K.R.: Inertial Navigation Systems Analysis. Wiley, New York (1971)
8. Leutenegger, S., Melzer, A., Alexis, K., Siegwart, R.: Robust state estimation for small unmanned airplanes. In: 2014 IEEE Conference on Control Applications (CCA), pp. 1003–1010, October 2014
9. Besl, P.J., McKay, N.D.: A method for registration of 3-D shapes. IEEE Trans. Pattern Anal. Mach. Intell. **14**, 239–256 (1992)
10. Chen, Y., Medioni, G.: Object modelling by registration of multiple range images. Image Vis. Comput. **10**, 145–155 (1992)
11. Zhang, Z.: Iterative point matching for registration of free-form curves and surfaces. Int. J. Comput. Vis. **13**, 119–152 (1994)
12. Pomerleau, F., Colas, F., Siegwart, R., Magnenat, S.: Comparing ICP variants on real-world data sets. Auton. Robots **34**, 133–148 (2013)
13. Segal, A., Haehnel, D., Thrun, S.: Generalized-ICP. In: Proceedings of Robotics: Science and Systems, Seattle, USA, June 2009
14. Biber, P., Strasser, W.: The normal distributions transform: a new approach to laser scan matching. In: Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS 2003), vol.3, pp. 2743-2748 (2003). doi:10.1109/IROS.2003.1249285
15. Doherty, P., Kvarnström, J., Wzorek, W., Rudol, P., Heintz, F., Conte, G.: HDRC3 - a distributed hybrid deliberative/reactive architecture for unmanned aircraft systems. In: Valavanis, K.P., Vachtsevanos, G.J. (eds.) Handbook of Unmanned Aerial Vehicles, pp. 849–952. Springer, Heidelberg (2014)
16. Vempati, A.S., Agamennoni, G., Stastny, T., Siegwart, R.: Victim detection from a fixed-wing UAV: experimental results. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Pavlidis, I., Feris, R., McGraw, T., Elendt, M., Kopper, R., Ragan, E., Ye, Z., Weber, G. (eds.) ISVC 2015. LNCS, vol. 9474, pp. 432–443. Springer, Heidelberg (2015). doi:10.1007/978-3-319-27857-5_39