

Event-Driven Loop Closure in Multi-Robot Mapping

Teresa A. Vidal-Calleja, Cyrille Berger and Simon Lacroix

Abstract—A large-scale mapping approach is combined with multiple robots events to achieve cooperative mapping. The mapping approach used is based on hierarchical SLAM – global level and local maps –, which is generalized for the multi-robot case. In particular, the consequences of multi-robot loop closing events (common landmarks detection and relative pose measurement between robots) are analyzed and managed at a global level. We present simulation results for each of these events using aerial and ground robots, and experimental results obtained with ground robots.

I. INTRODUCTION

Many robotic missions can be more efficiently and robustly achieved by a team of robots. Our interest is focused on air/ground robotics cooperative systems: within such systems, building and sharing environment representations is a key functionality to allow effective cooperation between the robots. The environment mapping problem is then more complex than in the single robot case [10], [2]. Notably, some issues arise because of communication constraints (it can not be taken for granted that all the acquired information can be centralized), and because robots might have heterogeneous perception means. The latter point naturally impacts the data association problem, but in this paper we focus on the distributed nature of the mapping.

Distributed mapping naturally leads to approaches where the robots build independent maps. Various “events” can occur during the exploration, which provide information on the spatial organization of these maps: *i.e.* relative robot position observations or map matching. The goal of this paper is to analyze these events, and to properly manage the mapping process among the robots when they occur. For that purpose, we propose an approach that generalizes the hierarchical SLAM approach introduced in [3] to a multi-robot context: each robot builds a set of independent local maps (or “submaps”) and maintains an adjacency graph between these maps, that constitutes the global level. The occurrence of the aforementioned events raises the following questions: How to manage the global level and the local maps between the robots? What information must be exchanged between the robots for that purpose?

Related work: Different scalable SLAM approaches in which a single vehicle builds multiple local maps have been proposed, mainly to reduce computational complexity and to delay linearization errors until the map merging [4], [5], [3],




Fig. 1. Loop closure events for multiple maps and multiple robots. Here, the relative initial position of both robots is known (*e.g.* georeferenced): a single occurrence of an event generates a cycle in the multi-robot global graph. A new local map is started when the event occurs.

[6]. When the maps are merged into a single one, on the basis of either common landmarks between local maps, or simply the sequential constraint, a fully correlated map of the environment is obtained. Successful fast implementations exploiting the topology of the representation to systematically join and fuse local map have been proposed for the single vehicle case, such as trees [8], or binary trees [6].

Work on map merging in a multi-robot context has been presented in [9]: the authors propose an algorithm to align maps built by different vehicles. The information provided by the robot-to-robot measurements (*rendez-vous*) is used to obtain the transformation between their coordinate frames. In [10], the authors use a maximum likelihood approach to merge maps with initial unknown correspondences. This work subdivides the multi-robot problem into three subproblems: incremental localization and mapping, loop closure and island merging. This is related to our approach, in the sense that we also consider this subdivision, although we rely on the *rendez-vous* and landmark correspondences events to trigger loop closures at the global graph level. The advantage of the global level is that it keeps consistency, via a batch optimization, on the global position of each robot without the need of communicating all the local maps.

Outline: The next section briefly describes the hierarchical SLAM approach and sketches its extension to multiple robots, considering two events among robots that trigger loop closures. Section III depicts these events and their impact on the overall map building strategy. Sections IV and V respectively present simulation results with aerial and ground robots, and experimental results with ground robots. Finally, Section VI concludes this paper.

T. Vidal-Calleja, C. Berger and S. Lacroix are with CNRS; LAAS; 7 Av. du Colonel Roche, F-31077 Toulouse, France and Université de Toulouse; UPS, INSA, INP, ISAE; LAAS; F-31077 Toulouse, France. {teresa.vidal,cyrille.berger,simon.lacroix}@laas.fr

Cyrille Berger is supported by Thales Optronics.

This work is partially supported by the DGA contract “PEA Action” – action.onera.fr.

II. APPROACH

A. Hierarchical SLAM

In the hierarchical SLAM approach [3], independent consecutive local maps are represented in their own reference frame (*lrf*), and the upper global level is a graph whose nodes correspond to the origin of the local maps, and whose edges are relative transformations between them. The nodes poses are expressed in the world reference frame (*wrf*).

Local Level: The local level contains the feature-based locally referred stochastic map, built with the standard EKF-SLAM. A local map i is defined by:

$$\hat{\mathbf{x}}^{R_i} = \begin{bmatrix} \hat{\mathbf{r}}_n^{R_i} \\ \hat{\mathbf{M}}^{R_i} \end{bmatrix}, \quad \mathbf{P}^{R_i} = \begin{bmatrix} \mathbf{P}_{\mathbf{r}_n}^{R_i} & \mathbf{P}_{\mathbf{r}_n} \mathbf{P}_{\mathbf{M}}^{R_i} \\ (\mathbf{P}_{\mathbf{r}_n} \mathbf{P}_{\mathbf{M}}^{R_i})^\top & \mathbf{P}_{\mathbf{M}}^{R_i} \end{bmatrix}, \quad (1)$$

where $\hat{\mathbf{r}}_n^{R_i}$ is the pose of the n th robot and $\hat{\mathbf{M}}^{R_i} = [\hat{\mathbf{m}}_1^{R_i} \dots \hat{\mathbf{m}}_m^{R_i}]^\top$ is the map of m landmarks, both with respect to R_i frame. For one robot, local maps are obtained sequentially: the relative transformation between consecutive maps is given by the last robot pose in the *lrf*, i.e., $\hat{\mathbf{x}}_{R_i R_{i+1}} = \hat{\mathbf{r}}_n^{R_i}$.

Global Level: The last robot pose is marginalized out from the local map to become part of the global level. The global level is defined by the state $\hat{\mathbf{x}}_g$ of relative transformations between local maps and the associated covariance \mathbf{P}_g :

$$\hat{\mathbf{x}}_g = \begin{bmatrix} \hat{\mathbf{x}}_{R_1 R_2} \\ \vdots \\ \hat{\mathbf{x}}_{R_i R_j} \end{bmatrix}, \quad \mathbf{P}_g = \begin{bmatrix} \mathbf{P}_{R_1 R_2} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{P}_{R_i R_j} \end{bmatrix}. \quad (2)$$

The relative transformation $\hat{\mathbf{x}}_{R_i R_j}$ between the *lrf* of local maps R_i and R_j are the edges of the graph. The nodes (origins of local maps) are obtained by compounding the previous origin with the relative transformation¹. Note that the global level might be viewed as a sparse pose-SLAM [11], where local maps are hanging from past robot poses in *wrf*, but with no correlation before a loop closure.

Loop Closure: At the global level, a loop closure corresponds to a cycle in the graph, that appears for instance when a relative position estimate between non consecutive submaps is established by a map matching process. Such a cycle defines a constraint between a series of relative transformations:

$$\mathbf{h}(\hat{\mathbf{x}}) = \hat{\mathbf{x}}_{R_i R_{i+1}} \oplus \hat{\mathbf{x}}_{R_{i+1} R_{i+2}} \cdots \oplus \hat{\mathbf{x}}_{R_j R_i} = \mathbf{0}. \quad (3)$$

Given that $\mathbf{h}(\hat{\mathbf{x}})$ is not linear due to the angular terms, the enforcement of this constraint can be formulated as a nonlinear constrained optimization problem. A solution for instance could be based on the Iterative EKF as presented in [3]. As a consequence, the part of the state involved in the loop closure at global level becomes correlated, resulting in a non-sparse covariance matrix \mathbf{P}_g .

¹A detailed description of the compounding \oplus and inversion \ominus operations for 2D and 3D can be found in [12].

B. Multiple robots

A hierarchical SLAM approach in the multi-robot case seems a priori straightforward: each robot manages a set of submaps and a global graph of poses. But the interests of multi-robot mapping arise of course when the robots exchange position or mapping information, which allows to enhance the individual maps spatial consistency and to build up a *multi-robot pose graph*. The two following events provide such information (Figure 1):

- *rendez-vous* between robots, i.e. relative robot to robot pose observations,
- *map-matching* of independent submaps between two robots (or of landmarks belonging to independent submaps).

These two events create a link between the robot's global level. Whereas in a single robot case a loop closure only occurs when the robot revisits a previously mapped place, in a multi-robot case these events trigger loop closures: any cycle that appears in the overall graph defined by the concatenation of each robot graph (the *multi-robot graph*) is a loop closure. The compounding of all relative transformations that define a cycle is equal to zero as in equation 3, and a batch optimization over the transformations can be performed. Note that to obtain a cycle in the graph defined by the concatenation of two robots' global levels, at least two events between these robots are required².

Also, a third loop closing event is to be considered: absolute position observation, as provided by e.g. a GPS fix or an algorithm that matches landmarks with an initial georeferenced map. This establishes a link between the current robot pose and the georeferenced frame, triggering a loop closure for a single robot that has already been absolutely localized once (e.g. if it started from a georeferenced position).

Besides a low algorithmic complexity, the main advantage to exploit a hierarchical map structure in multi-robot mapping is the low communication bandwidth required among the robots: only the individual global graphs need to be exchanged to update the multi-robot graph.

C. Map Merging

Merging submaps into a global map is not necessary for the robots to operate, as is not required to maintain the consistency on the graph. Should one require a global map, the map fusion could be delayed until all possible loop closures are performed, e.g. at the end of the mapping process. An approach similar to the Divide & Conquer algorithm [6] is exploited for that purpose, on a node that centralizes all the information. We merge maps considering common map information in covariance form. However, fusion in Information form [2] has more interesting properties for decentralized systems: once the common landmarks for both maps are matched, the maps are transformed to the same reference and the information is just added.

²Unless they share the same *wrf*, in which case the *wrf* is a root node that connects the robot graphs: a single inter-robot event then defines a loop in the multi-robot graph (see Figure 1).

III. LOOP-CLOSING EVENTS

A. Setting up links

Rendez-vous: The relative transformation between two robots can be readily observed by a dedicated sensor (*e.g.* a robot camera that perceives a known pattern on another robot), or indirectly by matching two data sets acquired by both robots (*e.g.* a scan or feature matching process – as we do in our experiments, see section V). In both cases, the estimated transformation $\hat{\mathbf{x}}_{R_i, R_j}$ and the associated covariance generates a link in the multi-robot graph:

$$\hat{\mathbf{x}}_{R_i, R_j} = \mathbf{z}_{R_i}^{R_j}, \quad (4)$$

where $\mathbf{z}_{R_i}^{R_j}$ is directly the measurement.

Map-matching: Matches established between submaps (or landmarks within submaps) built by different robots set up a link between nodes of robot's global levels. Such matches must be searched between submaps that are likely to match, which can either be selected on the basis of their position estimate³, or by applying loop closure detection approaches – *e.g.* using image indexing techniques [13], such approaches having never been applied with multiple robots to our knowledge. Once submaps or landmark are matched, an estimation method yields the relative transformation between the submaps *lrf* with the associated covariance. For instance, for point landmarks, a statistic least-squares method can be applied:

$$\hat{\mathbf{x}}_{R_i, R_j} = \underset{\hat{\mathbf{x}}_{R_i, R_j}}{\operatorname{argmin}} \left(\sum_{m=1}^L |\hat{\mathbf{m}}_l^{R_i} - \hat{\mathbf{x}}_{R_i, R_j} \oplus \hat{\mathbf{m}}_l^{R_j}| \right), \quad (5)$$

where L is the number of matched landmarks.

Absolute localization: In an air/ground context, it is reasonable to assume that both kind of robots receive GPS fixes from time to time. The relative transformation provided by a GPS fix for vehicle i is simply $\hat{\mathbf{x}}_{R_i, G}$, where G is the georeferenced frame. Such information provide a link between a *lrf* and a global georeferenced frame, and generate a loop at the graph level for an individual robot.

B. Impact on the submaps

The hierarchical SLAM approach makes the assumption that the submaps are statistically independent. This is the case between submaps built by different robots, but it is an approximation for the submaps that are built by the same robot⁴. However, in the multi-robot case it can happen that two different events create a link on the same node, *i.e.* if a *map-matching* is established after a *rendez-vous*. To avoid this problem, new submaps are systematically started *after an event occurs*.

Similarly, to avoid counting information twice if one wants to merge all the submaps, after a *map-matching* event both submaps should be fused into a single submap. This requires that the submaps must be shared among the two robots, but on the one hand this is a pre-requisite to establish the matches, and on the other hand such events will occur when the robots are within communication range.

C. Enforcing constraints

In the centralized case, the loop constraint can be enforced whenever a cycle in the graph is found. If there are several cycles, the constraint is only applied to the largest one. For the distributed case, if a constraint is enforced by one robot and not the others, further global information exchanges will lead to incompatible graphs among robots: a loop constraint should only be enforced when it is known by all the robots. However constraints can be enforced locally, provided the graph state before the constraint application is memorized, which allows to backtrack to a globally consistent graph when new information is received from other robots.

IV. SIMULATION VALIDATION

We analyze simulation results for three settings; *rendez-vous*, map matching, and the full collaboration of three robots with several loop closures triggered by the latter events.

The world has 300 landmarks spread along a 10000 m² surface. Aerial and ground robots build local maps until any of the events occurs. We show robots' trajectories and consistency plots before the final global map merging. In each submap, the poses and their uncertainties are expressed in the associated *lrf*s: in order to plot the global positions, we compound the poses and the covariances to obtain them in the *wrf*. The robots move controlled by linear u and angular ω velocities on their own planes with the following odometry noise model: $\sigma_u = 0.01 \text{ m}/\sqrt{s}$, $\sigma_\omega = 1^\circ/\sqrt{s}$, $\sigma_{\omega u} = \sigma_{u\omega} = 0$. The uncertainty of each robot is considered for the 6 DOF $x, y, z, yaw, pitch, roll$. Robots make bearing only observations: the inverse-depth parameterization is used to map 3D point landmarks, with the parameters $\rho_{init} = 0.5 \text{ m}^{-1}$ and $\sigma_\rho = 0.5 \text{ m}^{-1}$ (see [15] for more information). The bearing-only observation model has a 0.2° standard deviation, and the initial parameters for the three experiments are shown in Table I.

Rendez-vous: For this setting we use two robots, \mathbf{r}_1 (aerial robot) and \mathbf{r}_2 (ground robot). We chose to start with a large uncertainty for \mathbf{r}_2 to see the impact of loop closures in the global localization – each local map is started with zero uncertainty, so this does not affect the local EKF-SLAM performance.

The robots meet in the middle of the environment where \mathbf{r}_1 detects \mathbf{r}_2 and computes the transformation between the two robots. This event triggers a loop closure because both robots initial positions are expressed in the same *wrf*. The standard deviation of the relative measure between robots is (0.02, 0.02, 0.02, 0.005, 0.005, 0.005) in (m, m, m, rad, rad, rad). The robots meet twice, the first time about 300s and then approximately at 440s as shown Figure 3. The first *rendez-vous* reduces the covariance on the position of \mathbf{r}_2 but unfortunately deteriorates its position (see Figure 3(b) at 300s). As one can see in the plot after the second *rendez-vous*, this problem is mostly fixed.

To evaluate the consistency of the approach, we performed a 50-runs Monte Carlo analysis of the normalized estimation error squared (NEES) of the current robot pose in *wrf*:

$$\mathbf{G}_k^{\mathbf{r}_n} = \mathbf{G}_{k-1}^{\mathbf{r}_n} \oplus \hat{\mathbf{x}}_{R_{k-2}R_{k-1}} \oplus \hat{\mathbf{r}}_n^{R_{k-1}}, \quad (6)$$

³Which implies that both robots share the same *wrf*.

⁴Conditional independent SLAM in [14] palliates this issue.

TABLE I
SIMULATION PARAMETERS FOR THE THREE SETTINGS.

Experiment	Robot	Real Initial Position ($x, y, w, \text{yaw}, \text{pitch}, \text{roll}$) (m,m,m,rad,rad,rad)	Slam Initial Position ($x, y, w, \text{yaw}, \text{pitch}, \text{roll}$) (m,m,m,rad,rad,rad)	Initial Standard deviation ($\sigma_x, \sigma_y, \sigma_w, \sigma_{\text{yaw}}, \sigma_{\text{pitch}}, \sigma_{\text{roll}}$) (m,m,m,rad,rad,rad)	Speed (u, ω) (m/s, rad/s)
<i>Rendez-vous</i>	\mathbf{r}_1	(0, -25, 8, 0, 0, 0)	(0, -25, 8, 0, 0, 0)	(0, 0, 0, 0, 0, 0)	(0.1, 0.01)
	\mathbf{r}_2	(0, 25, 0, 0, 0, 0)	(0.2, 25.5, -0.3, 0, 0, 0)	(1, 1, 1, 0.1, 0.1, 0.1)	(0.1, -0.01)
Common Landmark	\mathbf{r}_1	(0, 0, 8, 0, 0, 0)	(0, 0, 8, 0, 0, 0)	(0, 0, 0, 0, 0, 0)	(0.1, 0.007)
	\mathbf{r}_2	(0, 25, 0, 0, 0, 0)	(-0.1, 24.7, 0.1, 0, 0, 0)	(1, 1, 1, 0.1, 0.1, 0.1)	(0.1, 0.01)
Full Collaboration	\mathbf{r}_1	(25, 50, 0, 0, 0, 0)	(25, 50, 0, 0, 0, 0)	(0.1, 0.1, 0.1, 0.01, 0.01, 0.01)	(0.1, -0.01)
	\mathbf{r}_2	(-25, -50, 0, 0, 0, 0)	(-25, -50, 0, 0, 0, 0)	(0.1, 0.1, 0.1, 0.01, 0.01, 0.01)	(0.2, 0.01)
	\mathbf{r}_3	(0, -40, 8, 0, 0, 0)	(0, -40, 8, 0, 0, 0)	(0.1, 0.1, 0.1, 0.01, 0.01, 0.01)	(0.3, 0.008)

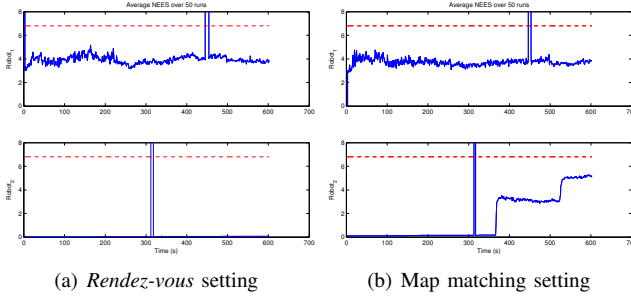


Fig. 2. Average NEES of 50 Monte Carlo runs of the robot pose (position and orientation) for the common landmark setting. The dashed line mark the single-sided 95% probability concentration region for the 6-dimensional state vector (the peaks are artifacts caused by a switch of the robot heading from $-\pi$ to π).

as explained in [16]. Figure 2(a) shows the average NEES for the global pose of both robots and the single-sided 95% region for the $N \times 6$ DOF during this setting. The average NEES shows good consistency for the full simulation, thanks to the use of sub-maps. For robot \mathbf{r}_2 , the average NEES is small compared with the consistency limits. This is because the robot starts with very large global uncertainty, while the estimation error is not so large.

Map Matching: For this experiment the two robots move clockwise in a circle. Robot \mathbf{r}_1 moves in the direction of the initial position of \mathbf{r}_2 . Once \mathbf{r}_1 reaches that position it is able to close a loop, as seen on 4(a). One of the main issues of this experiment is to compute the transformation between the two local maps which have common landmarks, because the least-square estimation is very sensitive to points whose localization is not accurate. The problem is even more difficult for bearing-only observations, when newly initialized points can be anywhere on the bearing line. The eigenvalues of each point's covariance matrix are computed, and points whose eigenvalues are below a threshold (e.g. 0.1 m^2) are selected to estimate the transformation.

The results of a single run experiment are presented on Figure 4. A first loop is closed at 380s, which increases the precision of the global localization of \mathbf{r}_2 . At 520s the two robots close another loop. Note that those two loops have mostly improved the global localization of \mathbf{r}_2 . When \mathbf{r}_2 reaches its original position, it can then close a loop with its first map, at around 610s: that map being connected to one of the map of \mathbf{r}_1 , a slight improvement on the localization of \mathbf{r}_1 can be seen through its covariance (Figure 4(b)).

Figure 2(b) shows the average NEES of the robots poses for this setting: after a loop closing the average NEES augments because the covariance is reduced significantly.

Collaboration between three robots: We show here the effects of multiple loop closures between three robots. Two ground robots \mathbf{r}_1 and \mathbf{r}_2 move along circles in different locations on the environment. They never meet, and their maps never overlap. The third robot \mathbf{r}_3 is an aerial robot that moves in a circular trajectory which extends from the area covered by \mathbf{r}_1 to the one covered by \mathbf{r}_2 .

The three robots start at a well known location, then at 240s, \mathbf{r}_1 and \mathbf{r}_3 have a *rendez-vous*, later at 610s and 700s the robot \mathbf{r}_3 detects a loop closure with two maps of \mathbf{r}_2 (the first part of the video shows a run of this experiment). The uncertainties are expressed in the *wrf* so that the effects of the loop closure can be seen. The consistency plots for a single run are shown in Figure 5. The final map and global graph before and after merging (at the end of the simulation) are shown in Figure 6.

V. REAL-DATA EXPERIMENTS

We present preliminary results with two ground robots in an area of approximately 900 m^2 . The second part of the video shows this outdoors experiment. The robots explore the area independently, starting from known uncertain positions. Both robots have a stereovision bench, however we use a single camera to perceive the landmarks. The submaps are built using the inverse-depth parameterization as in the simulation settings. Each robot processes about 1000 images. Landmarks are Harris corners, tracked and matched with the algorithm depicted in [17] (see Figure 7(a) and 7(b)).

In our experiments, a *rendez-vous* is emulated using matches of interest points perceived by the two robots, the 3D coordinates of the points being obtained by stereovision: as a result, one has an estimate of the relative robot position. The effects of the *rendez-vous* event are shown in Figure 7. It also shows the image frames before the *rendez-vous* happens for both robots: new local maps are initiated afterwards.

The robot \mathbf{r}_1 starts with a small uncertainty in global (0.01, 0.01, 0.01, 0, 0, 0), while \mathbf{r}_2 starts with a huge uncertainty (15, 8, 0.5, 0.8, 0, 0) ($\text{m}^2, \text{m}^2, \text{rad}^2, \text{rad}^2, \text{rad}^2$) variances for the pose. The reason of this choice is to be able to locate both robots in the same *wrf*. This fact does not affect the consistency of the EKF-SLAM as mentioned before. The threshold for starting new local maps is 120 landmarks. The

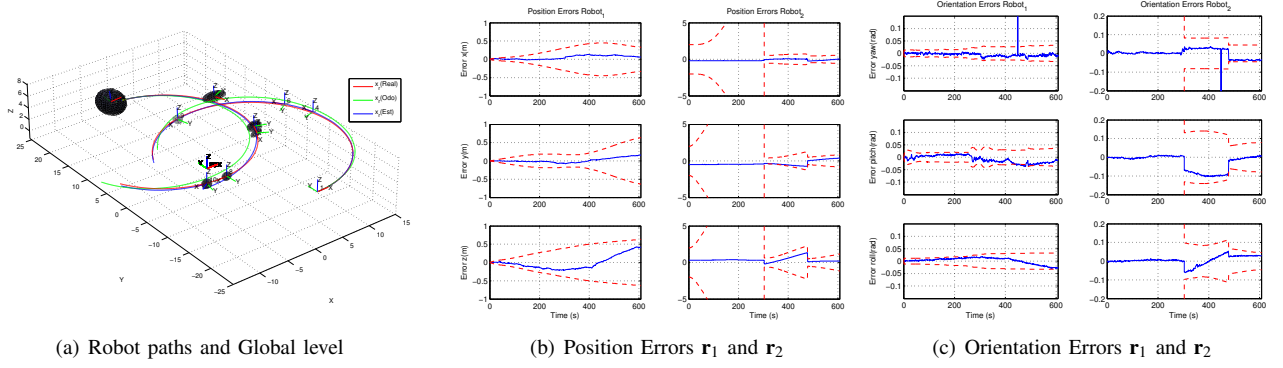


Fig. 3. *Rendez-vous* simulation results for one aerial and one ground robots. In a) the odometry is shown in green, real and estimated trajectories are shown in red and blue respectively. 3σ ellipsoids are plotted on the basis of each *lrf*. b) shows the global position errors for each robot and their global 3σ uncertainty bounds. c) shows the global orientation errors for each robot and their global 3σ uncertainty bounds (the jump on the yaw curve around $t = 450$ is due to a switch from $-\pi$ to π).

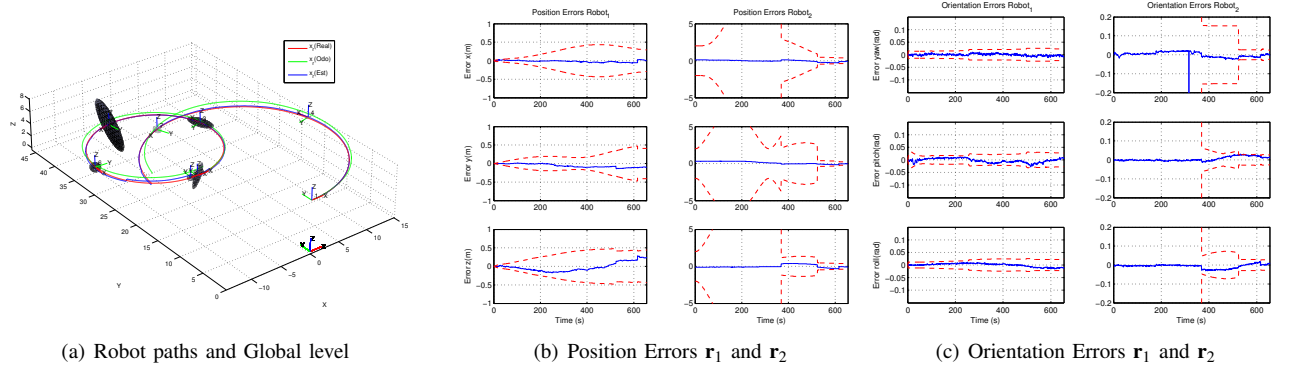


Fig. 4. Landmark matching simulation results for one aerial and one ground robots. See caption of Figure 3. .

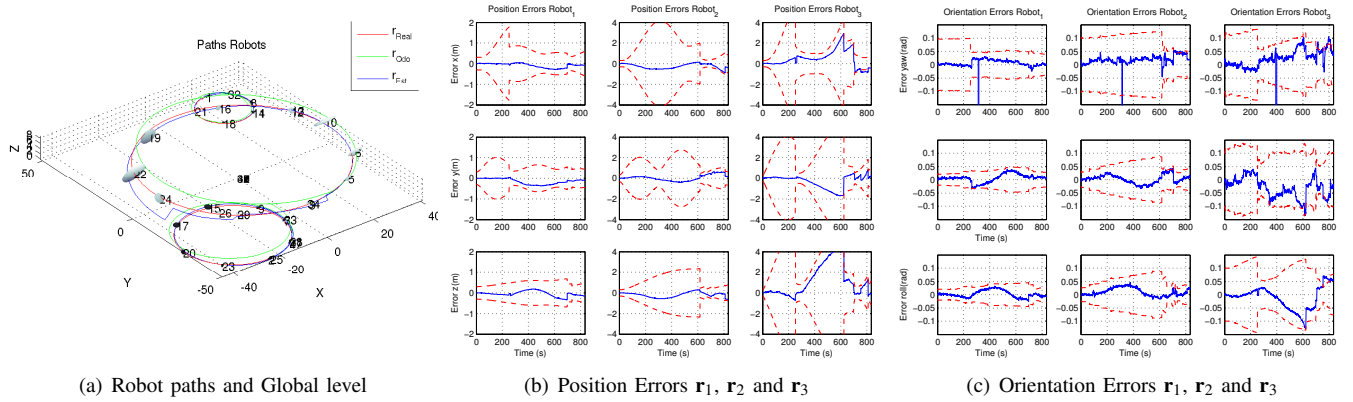


Fig. 5. Simulation results for the 3 robots exploration; one aerial and two ground robots. See caption of Figure 3.

final global level and the robot's trajectory of this experiment are presented in Figure 8.

VI. CONCLUSIONS

We explored the use of a multiple local maps technique for multi-robots. We exploit the way hierarchical SLAM approach solves the loop closure and the way map merging is done in Divide & Conquer SLAM. We proposed to trigger the loop closures using multi/single-robot events such as finding information correspondences between unconnected local maps, the *rendez-vous* between two robots or GPS fixes.

We analyzed the consequences of the appearance of cycles in the global multi-robot graph that these events can cause. Two events might be sufficient to close a loop between two robots, without requiring absolute position knowledge (GPS observations). Simulation results and a real-data experiment validate the approach.

Our approach is distributed, and the graph level is the sole information that must be exchanged between the robots. A central node is only required to merge maps, which is not required for the robots to build or navigate within local maps.




Fig. 6. Final map resulting from the 3 robots exploration. The figure shows the 3D global map, local map origins and the final robots location in *wrf*.

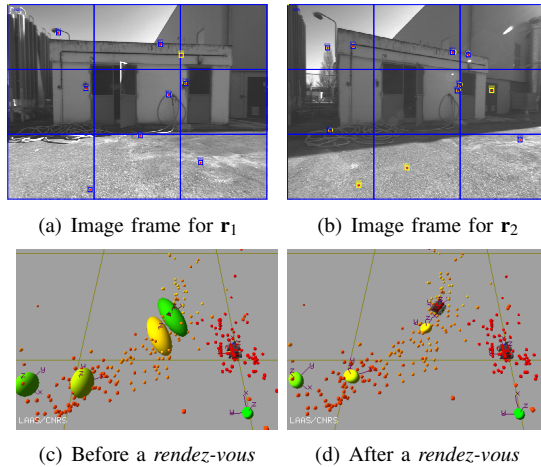


Fig. 7. Top: Image frames from both robots before *rendez-vous*. Yellow squares represents a new interest point declared as a landmark, blue squares the currently tracked landmarks and yellow ellipses the uncertainty in the image view. The blue line grid is used to reduce the number of image features used as landmarks in each frame. Bottom: *Rendez-vous* effect in the global map, with the global level (large ellipsoids), the 3D points map without uncertainty and the final robots location in *wrf*.




Fig. 8. Real-data outdoors experiment with 2 robots. The figure shows the final global level and the robot's trajectories in blue (r_1) and in red (r_2) in *wrf*. The robot r_2 jumps after the first *rendez-vous* is established and previously mapped local maps are re-localized at global level.

REFERENCES

- [1] A. Howard, G. Sukhatme, and M. Mataric, "Multirobot simultaneous localization and mapping using manifold representations," *Proceedings of the IEEE - Special Issue on Multi-robot Systems*, vol. 94, no. 7, pp. 1360–1369, July 2006.
- [2] S. Thrun and Y. Liu, "Multi-robot SLAM with sparse extended information filters," in *Proc. 11th Int. Sym. Robot. Res.*, Siena, Oct. 2003.
- [3] C. Estrada, J. Neira, and J. Tardós, "Hierarchical SLAM: Real-time accurate mapping of large environments," *IEEE Trans. Robot.*, vol. 21, no. 4, pp. 588–596, Aug. 2005.
- [4] S. B. Williams, H. Durrant-Whyte, and G. Dissanayake, "Constrained initialization of the simultaneous localization and mapping algorithm," *Int. J. Robot. Res.*, vol. 22, no. 7, pp. 541–564, Jul. 2003.
- [5] J. D. Tardós, J. Neira, P. M. Newman, and J. J. Leonard, "Robust mapping and localization in indoor environments using sonar data," *Int. J. Robot. Res.*, vol. 21, no. 4, pp. 311–330, 2002.
- [6] J. D. T. Lina María Paz, Patric Jensfelt and J. Neira, "EKF slam updates in $O(n)$ with divide and conquer slam," in *IEEE Int. Conf. on Robotics and Automation*, Rome, Italy, April 2007.
- [7] U. Frese, "Treemap: An $O(\log n)$ algorithm for indoor simultaneous localization and mapping," *Auton. Robot.*, vol. 21, no. 2, Sep. 2006.
- [8] X. Zhou and S. Roumeliotis, "Multi-robot slam with unknown initial correspondence: The robot rendezvous case," in *IEEE/RSJ IROS, Beijing*, 2006, pp. 1785–1792.
- [9] R. C. Smith and P. Cheeseman, "On the representation and estimation of spatial uncertainty," *Int. J. Robot. Res.*, vol. 5, no. 4, 1986.
- [10] R. Eustice, H. Singh, and J. Leonard, "Exactly sparse delayed-state filters for view-based SLAM," *IEEE Trans. Robot.*, vol. 22, no. 6, pp. 1100–1114, Dec. 2006.
- [11] M. Cummins and P. Newman, "FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance," *The International Journal of Robotics Research*, vol. 27, no. 6, pp. 647–665, 2008.
- [12] P. Piniés and J. D. Tardós, "Scalable slam building conditionally independent local maps," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, San Diego, CA, Oct 29–Nov 2 2007.
- [13] J. Civera, A. J. Davison, and J. M. M. Montiel, "Dimensionless Monocular SLAM," in *Iberian Conference of Pattern Recognition and Image Analysis*, June 2007.
- [14] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the EKF-SLAM algorithm," in *IEEE/RSJ IROS, Beijing*, Oct. 2006, pp. 3562–3568.
- [15] T. Lemaire, C. Berger, I.-K. Jung, and S. Lacroix, "Vision-based slam: Stereo and monocular approaches," *IJCV - IJRR*, 2006.