

# Global Alignment of a Multiple-Robot Photomosaic using Opto-Acoustic Constraints<sup>\*</sup>

Ricard Campos<sup>\*</sup> Nuno Gracias<sup>\*</sup> Albert Palomer<sup>\*</sup>  
Pere Ridao<sup>\*</sup>

*<sup>\*</sup> Computer Vision and Robotics Group  
University of Girona, Spain  
(e-mail: {rcampos, ngracias, apalomer, pere}@eia.udg.edu)*

---

**Abstract:** Distributing the mapping problem over several vehicles provides the ability to survey an area in fewer time than with a single one. However, the communications required to maintain the vehicles in formation underwater are highly complex. We present a global alignment method for multi-AUV trajectory optimization. We take advantage of the acoustic messages passed between vehicles as well as of the further constraints posed by optical matches between images in some of the vehicles in the formation, both within a single robot and from different ones. By merging all these constraints, we can obtain both an optical georeferenced mosaic of the surveyed area and a joint optimization of the poses of the vehicles in the formation at each time. We present and discuss an application of the method on a real dataset collected within the MORPH EU-FP7 project.

*Keywords:* Robot navigation, Robot vision, Underwater vehicles, Cooperative navigation, Relative localization, Optical mapping,

---

## 1. INTRODUCTION

The use of multi-robot systems for underwater mapping operations are gaining popularity in recent years due to the availability of acoustic modems and advances in control mechanisms dealing with these complex configurations. Obviously, having more than a single vehicle surveying an interest area reduces the mapping time. Given the locality of optical images, this is of special interest for large area mapping.

The work presented in this paper has been developed within the framework of the MORPH EU-FP7 project (Kalwa et al., 2012). This project proposes a novel concept of an underwater robotic system that emerges out of using different mobile robot modules with distinct and complementary resources. The robots will navigate at a very close range as a group, and will have the ability to change the formation to adapt to changes in the terrain. The project will provide efficient methods to map the underwater environment with great accuracy in situations that defy existing technology: namely on surveys over rugged terrain and structures with full 3D complexity, including walls with overhangs. Its applications cover a wide range of scientific and commercial areas such as monitoring of cold water coral reefs, oil and gas pipeline inspection, or harbor and dam protection. The idea that concerns more

this paper is that an underwater vehicle equipped with a multibeam sonar profiler (LSV), advances at the forefront of the formation, flying at a “safe” altitude from the sea-floor, while two other vehicles (camera vehicles CV1 and CV2) fly very close to the bottom, with a much higher risk of collision, acquiring images. The communication vehicle (GCV) is at the back with the main goal of keeping a good communication channel with the surface vehicle (SSV).

Whereas the MORPH concept envisions application to rugged terrain, the work of this paper addresses the particular case of surveys over approximately flat areas.

## 2. RELATED WORK

Underwater optical mapping has been an intensively studied topic in the underwater robotics community. Due to the common use of downward-looking cameras in the vehicles, a widely extended methodology is that of 2D mosaicing. Even if estimates for the robot pose are available at the time of image capturing, the accuracy of the navigation of the vehicle is orders of magnitude worse than the ground resolution of the images provided by the camera. In order to solve the misalignments between images in the mosaic plane, we can deal with the mosaicing problem as a pure image-based registration, as in Elibol et al. (2013). However, it would be desirable to allow the poses of the vehicles to be jointly optimized with the constraints provided by the matches between images during the map construction, as in Ferrer et al. (2007) and Pizarro et al. (2009). Moreover, posing the problem as a point-based bundle adjustment we can formulate the localization and mapping as a pose-graph optimization, such as the one

---

<sup>\*</sup> This work was partially funded through MINECO [grant number CTM2013-46718-R], the European Commissions Seventh Framework Programme as part of the project MORPH [grant number FP7-ICT-2011-7-288704], project Eurofleets2 [grant number FP7-INF-2012-312762] and the Generalitat de Catalunya through the AC-CIO/TecnioSpring program (TECSPR14-1-0050).

discussed in Kunz and Singh (2013). New trends such as the one proposed in Elibol et al. (2014) use strategies to optimally deal with the creation of mosaics of the same area as surveyed from different concurrent vehicles.

In our case, since we are taking into account at the same time the creation of a map and the retrieval of a correct navigation, the problem falls closer to simultaneous localization and mapping (SLAM) methods. When more than a single vehicle is used, the algorithms are usually referred to as Cooperative Localization methods (Rekleitis et al., 2001). These works are normally concerned with both sequential and non-sequential odometry links between the poses. Different estimation algorithms have been proposed to solve this problem, such as the Extended Kalman Filter (Roumeliotis and Bekey, 2002) or the particle filter (Howard, 2006). Obviously, the underwater application of such algorithms has the burden of the communication channel. The common procedure is to communicate relative measurements between vehicles in the formation using acoustic modems. However, the quantity of data that can be passed is restricted, its latency is very large, and it is error prone (Paull et al., 2014). Most of the approaches found in the literature are targeted to real time application on-board the vehicles but, due to the mentioned restrictions, we want to unify all these methodologies in an off-line procedure to take advantage of all the information available.

### 3. OVERVIEW AND CONTRIBUTIONS

This paper presents a global alignment method for multi-vehicle trajectory estimation, with a twofold motivation. On one hand, we target the creation of a multi-vehicle 2D optical map of the area under exploration. On the other hand, we also need to optimize the vague initial guess for the navigation that we obtain from the raw messages passed in a MORPH formation, thus generating a better estimation for the pose of these vehicles throughout the mission.

The multi-AUV structure proposed in the MORPH project allows for real-time control of the trajectory of the vehicles using acoustic messages. The problem resides in the fact that, with the exception of the surface vehicle, the global trajectory of the AUVs is not stored. We may have an internal estimation process for the trajectory for each vehicle (e.g., based on the motion model of the robot, or by integrating the output of sensors), but this trajectory is not in agreement, nor has been corrected, with respect to the rest of the vehicles in the formation.

Consequently, we introduce an off-line global alignment optimization framework taking advantages of all the messages passed between vehicles in the MORPH concept to recover this trajectory. While acoustic measures are enough for the control architecture to maintain the formation in real time, they are noisy and not precise enough to use them to build an optical map. Thus, we add image-based constraints in the form of homographies, which are derived from the construction of a photomosaic from the two camera vehicles (CVs). This provides further constraints to the global positioning of the other vehicles in the formation. Note that the improved navigation provided by our off-line approach may be useful for further mapping

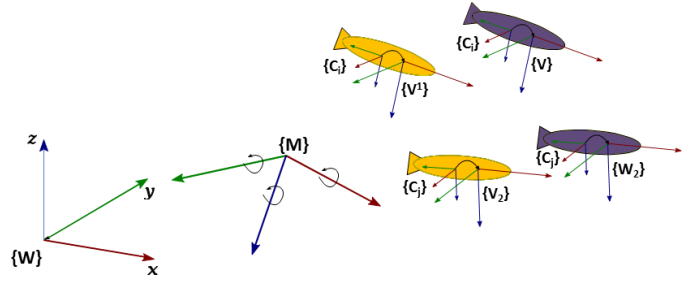


Fig. 1. Figure showing the different frames involved in the mosaic optimization: the world frame  $\{W\}$ , the mosaic frame  $\{M\}$ , and vehicle poses  $x_t^v$ , related with the camera frames  $c_t^v$  by a rigid transformation.

purposes, such as using it to map through the multibeam sensor carried by the vehicle at the forefront of the formation (LSV).

We start by introducing the creation of the optical mosaics, and the global alignment procedure involved, to then extend this least squares minimization with further terms coming from navigation and acoustic constraints.

### 4. OPTICAL MOSAICING

The two lower vehicles in the formation (CVx), are responsible for collecting imaging data in order to obtain an optical survey of the area. For these seabed surveys, and due to the downward-looking configuration of the cameras, we take the common assumption of the seabed being almost flat in order to create a photomosaic. We assume that the area under consideration does not contain enough relevant roughness to introduce parallax-related problems during the photomosaic construction.

We start by explaining, with the help of the depiction in Figure 1, all the frames and poses involved in the optimization. We parametrize a pose  $x_t^v$  of vehicle  $v$  at time  $t$  in the more general term using 6 degrees of freedom (DOF), with rotations represented as unit quaternions. All these vehicle poses are defined with respect to the frame  $\{W\}$ . Furthermore, we also have the frame  $\{M\}$ , whose X-Y plane frame comprises the 2D photomosaic frame where all the images should be projected. Note also that the camera pose/frame  $c_t^v$  is located at a given transformation with respect to the vehicle pose, which we consider to be known for each CV. Given the fact that just known rigid transformations separate  $x_t^v$  from  $c_t^v$  poses, we assume now to work directly with  $c_t^v$  and with respect to the  $\{M\}$  frame.

When dealing with the creation of photomosaics, the common assumption is that the vehicles have an altimeter that relates its pose with respect to the seabed (Ferrer et al., 2007). Nevertheless, in our case some of the vehicles do not contain an altimeter or are too far away from the seabed to compute a reliable altitude measure. Thus, we decided to define our poses with respect to depth measurements. This means that we do not have a reference for the base plane where to project the mosaic. Thus, the pose of  $\{M\}$  with respect to the world reference frame  $\{W\}$  is another variable to optimize. Note that this allows the mapping to happen on slanted terrains.

For simplicity, let us define the poses as separate  $3 \times 3$  rotation matrices  ${}^C_M R$ , defining the rotation from mosaic frame  $M$  to camera frame  $C$ , and translation matrices  ${}^M T = ({}^M t_x, {}^M t_y, {}^M t_z)$ , defining the translation with respect to frame  $M$ . Using this notation, and assuming the camera matrix  $K$  to be known (calibrated) for each camera, we can then define the classic pinhole projection model as follows:

$${}^i \begin{pmatrix} \phi u \\ \phi v \\ \phi \end{pmatrix} = K \cdot {}^C_M R \cdot \begin{pmatrix} 1 & 0 & 0 & -{}^M t_x \\ 0 & 1 & 0 & -{}^M t_y \\ 0 & 0 & 1 & -{}^M t_z \end{pmatrix} \cdot \begin{pmatrix} {}^M p_x \\ {}^M p_y \\ {}^M p_z \\ 1 \end{pmatrix}, \quad (1)$$

where  $({}^M p_x, {}^M p_y, {}^M p_z)$  is the 3D point referred in  $\{M\}$ , and  ${}^i(\phi u, \phi v, \phi)^T$  its 2D projection on the image frame  $i$ . From (1), we can define an image-to-ground plane mapping by projecting on the mosaic plane  $z = 0$ :

$${}^i \begin{pmatrix} \phi u \\ \phi v \\ \phi \end{pmatrix} = K \cdot {}^C_M R \cdot \begin{pmatrix} 1 & 0 & -{}^M t_x \\ 0 & 1 & -{}^M t_y \\ 0 & 0 & -{}^M t_z \end{pmatrix} \cdot \begin{pmatrix} {}^M p_x \\ {}^M p_y \\ 1 \end{pmatrix}, \quad (2)$$

Finally, we can define the absolute planar motion described in (2) as an absolute homography with respect to the mosaic frame  $M$ , representing a 3D point with respect to the plane  ${}^M z = 0$ :

$${}^i_M H = K \cdot {}^C_M R \cdot \begin{pmatrix} 1 & 0 & -{}^M t_x \\ 0 & 1 & -{}^M t_y \\ 0 & 0 & -{}^M t_z \end{pmatrix}. \quad (3)$$

Obviously, the relative planar motion between image frames  $\{i\}$  and  $\{j\}$  can be computed by composing the absolute ones:  ${}^i_j H = {}^i_M H \cdot {}^j_M H^{-1}$ .

With these components, we now have a way to determine if the poses in  $x_t^v$  are in agreement with homographies computed at image pairs using feature matching. The following section shows how we compute this minimization.

#### 4.1 Global Alignment

Localization and mapping problems such as the one introduced above are often modeled as a pose graph. This intuitive definition of the problem consists of a graph containing two types of nodes: poses and factors. Given a front-end definition of the graph, many available solvers can be used to deal with the non-linear minimization they describe. In our case, we implemented the optimization by extending the g<sup>2</sup>o framework (Kummerle et al., 2011).

It is important to notice that we are optimizing the  $x_t^v$  poses, as this is the connection with the larger definition of the optimization to be described in subsequent sections. Given our unknowns being poses  $x_t^v$  and another 6 DOF pose  $m$  referring to the mosaic frame  $\{M\}$ , the optimization we seek is the following:

$$X^*, M^* = \arg \min_{X, M} \sum_{v, w, t_a, t_b} \|e(x_{t_a}^v, x_{t_b}^w)\|_{\Sigma_{v, w, t_a, t_b}^h}^2, \quad (4)$$

where  $e(x_{t_a}^i, x_{t_b}^j)$  is the error measurement derived from the homographies. From each pair poses  $x_{t_a}^i, x_{t_b}^j$  having image matches, and following the set of rigid transformations and calculations described in the section above, we can

describe how well the image-computed homography agrees with the image-to-image planar homography derived from the poses:

$$\begin{aligned} {}^i e_k &= {}^i p_k - {}^i_j H \cdot {}^j p_k \\ {}^j e_k &= {}^j p_k - {}^j_i H \cdot {}^i p_k. \end{aligned} \quad (5)$$

This function is basically computing how far a 2D point  ${}^i p_k$  in one image falls from its computed match  ${}^j p_k$  through the homography  ${}^i_j H$  computed from the poses, with  $k = \{1..m\}$  being a determined number of correspondences per pair. As presented in the second line in (5), we do this bidirectionally. Thus, we have a 4-elements error vector for each correspondence. The size of the error vector, and its associated covariance matrix, depend on the number of matches taken into account, which we fix to four in the presented results.

Note that in (4), and in further terms to add in the minimization later on, we use the Mahalanobis distance  $\|e\|_{\Sigma}^2$ , where  $\Sigma$  represents the covariance matrix of the measure.

### 5. EXTENSION TO PRIORS AND ACOUSTIC MESSAGES

Once the mosaic has been built, we extend the proposed minimization by adding terms corresponding to the acoustic messages passed for controlling the MORPH formation, as well as some priors that may help in stabilizing the resulting trajectories. This further restricted optimization has the following terms:

$$X^*, M^* = \arg \min_X \{\mathbf{H} + \mathbf{A} + \mathbf{O} + \mathbf{U} + \mathbf{R}\}, \quad (6)$$

where  $\mathbf{H}$  is the term already introduced in section 4.1, and  $\mathbf{A}$ ,  $\mathbf{O}$ ,  $\mathbf{U}$ ,  $\mathbf{R}$  are the new terms restricting the trajectory, which will be introduced in the following sections.

#### 5.1 Absolute Priors

Absolute priors imposes some knowledge on the absolute measure for a pose vector. The minimization term has the following form:

$$\mathbf{A} = \sum_{v, t} \|x_t^v - a_t^v\|_{\Sigma_{v, t}^{a_t^v}}^2, \quad (7)$$

where  $a_t^v$  is a pose vector. This type of terms may be specially useful to define priors on the SSV vehicle, whose GPS readings and its known zero depth can be used to anchor the optimization. In fact, using different covariances, we can use absolute pose factors to define a prior uncertainty on the initialization poses, i.e., those corresponding to the trajectories we set at the beginning of the optimization. For example, we can impose a large reliability on the roll and pitch of the vehicles, as the robots used should be more or less stable on those angles. Additionally, we can impose reliability on the depth measures, while we let the other parameters of the poses  $x_t^v$  with a large uncertainty, meaning that they will be freed on the optimization, which will be the responsible of getting optimal values for them.

## 5.2 Relative Priors

Odometry (i.e., relative) constraints are used to keep a sequence on the trajectory:

$$\mathbf{O} = \sum_{v,t} \|x_{t-1}^v \oplus o_t^v - x_t^v\|_{\Sigma_{i,t}^o}^2, \quad (8)$$

where  $o_t^i$  is the relative pose measurement, and  $\oplus$  is the 3D version of the compounding operator as defined in Smith et al. (1987).

## 5.3 USBL Measures

USBL messages are passed between the SSV and the GCV/LSV vehicles, and basically provide a relative translation constraint at a given time. This relative translation is in UTM, and with respect to another vehicle:

$$\mathbf{U} = \sum_{v,w,t} \|u(x_t^v, x_t^w) - u_t^{v,w}\|_{\Sigma_{v,w,t}^u}^2, \quad (9)$$

where  $u_t^{v,w}$  is the 3-element vector describing the measured relative displacement, and  $u(x_a, x_b)$  is the translation part of the compounded pose  $x_a \oplus x_b$ .

## 5.4 Range Measures

We also have simpler range measurements between all the vehicles in the formation, acquired through time-of-flight computation among messages passed between vehicles, and with the assumption of known speed of sound on water. They provide just a distance restriction between poses:

$$\mathbf{R} = \sum_{v,w,t} \|d(x_t^v, x_t^w) - r_{v,w,t}\|_{(\sigma_{v,w,t}^r)^2}^2, \quad (10)$$

$d(x_a, x_b)$  being the Euclidean distance between the two poses, and  $r_{v,w,t}$  being the measured scalar.

## 6. INITIALIZATION

Starting this minimization from raw data and a vague trajectory estimation is very unlikely for it to converge near the global optimum. Thus, we split the processing in two steps.

First, we build the multi-vehicle 2D optical map. We search for consecutive and non-consecutive image feature correspondences using SURF (Bay et al., 2008) for feature detection/description, and a RANSAC-based procedure (Fischler and Bolles, 1981) for the robust projective homography estimation. The trajectory is initialized using the in-vehicle pose estimation (an EKF filter integrating DVL, IMU and pressure sensors), and then optimized using just the restrictions in Section 4.1. For the initialization of the mosaic frame  $\{M\}$ , we take advantage of the CVx vehicles being close to the seafloor and carrying an altimeter. Thus, we can take the initial on-vehicle navigation estimation, and their depth plus altitude measurements to construct a set of candidate points being at the seafloor. From this point set, we fit a PCA plane whose normal serves to create an orthonormal vectors initializing  $\{M\}$ . Once we have the mosaics built for each camera vehicle,

we then search for image matches defining relative homographies between image frames from different vehicles, and perform a joint optimization, which finally renders the multi-vehicle mosaic.

Using the constrained estimation of poses obtained by the mosaicing strategy, we then add the acoustic measurements. The trajectory of the SSV is quite reliable, as this is the only vehicle in the formation keeping GPS fixes. Consequently, it is given a low variance in the optimization. For the GCV and LSV case, their initial translation with respect to the world is directly estimated by composing USBL readings and the SSV poses. Note however that these messages are not error-prone, and we may find outliers among them. Luckily, they are easy to detect, and we can filter them using a simple approach. This naive method consists in extrapolating the next pose in the trajectory given the last two already computed: if the new compounded USBL-based pose is within a given radial neighborhood from the predicted one, we accept it as valid, disregarding this USBL measure otherwise. Regarding the rotation part of the poses, we use the internal sensors of each vehicle.

Finally, we jointly optimize all the measurements. Note that the different restrictions are not synchronized in time. Thus, we augment the number of pose nodes in the graph for each new measure to include in the optimization. The initialization of these new poses is computed by interpolation.

## 7. RESULTS

We validate our method on a real dataset collected during the MORPH trials of September 2014 in the city of Horta, in the island of Faial (Azores, Portugal).

The vehicles involved in these experiments were 3 Medusa AUVs, playing the role of SSV, LSV and GCV in the MORPH formation, and Seacat AUV and Sparus II AUV, used as CV1 and CV2 respectively.

The survey performed consisted of a series of six loops (zero-shaped), whose central point is shifted incrementally in one direction parallel to the seafloor. The two CVs were required to keep a constant altitude of 2.5m with respect to the seabed, while the rest of the vehicles were keeping a constant depth. Since the CV vehicles are located at left/right positions with respect to the LSV vehicle in the formation, after a set of loops, the area surveyed by the two vehicles will overlap. Furthermore, we decided to start the trajectory with a straight transverse path in the middle of the surveyed area, so that more non-consecutive matches are found each time we cross these parts within the loops.

At the time of writing this article, we decided to start validating our method with the initial part of this trajectory, that is, the transverse segment plus one and a half loop. The photomosaic, constructed from 1168 images (584 images for each CV) can be seen in the left-hand side of Figure 2. Moreover, on the right-hand side of the same figure we can observe the corrected trajectory with the detected consecutive, non-consecutive, and between-vehicles matches conforming the restrictions that guided the non-linear minimization proposed in Section 4.1. The shape of this last set of trajectories can be compared to

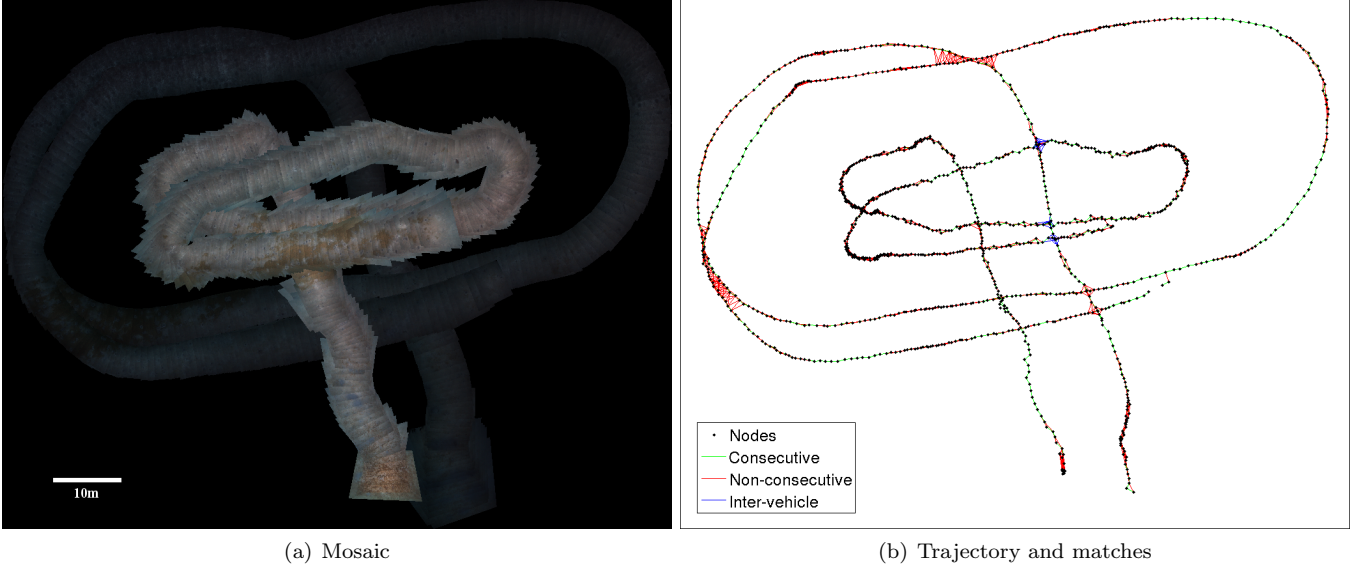


Fig. 2. Mosaic of the first loop of both camera vehicles in the dataset of Horta 2014. On the left figure, we find the constructed mosaic. Due to illumination changes, we can differentiate between the trajectories of CV1 (outer, darker loop) and CV2 (inner, brighter loop). On the right figure, we show the optimized trajectory, with consecutive matches marked in green, non-consecutive in red, and inter-vehicle in blue.

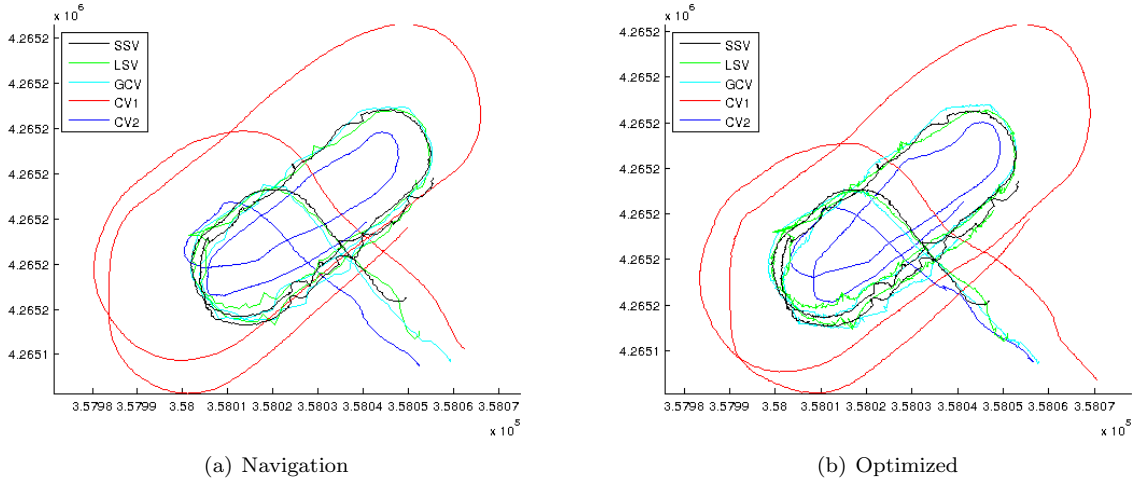


Fig. 3. Top view of the trajectories (UTM coordinates). On one hand, in (a) we find the initialization of the trajectories of each vehicle, as suggested in Section 6. On the other hand, in (b) we show the final optimized trajectories

that of Figure 3 (a), containing the raw navigation measures used to initialize the optimization. Due to changes in illumination conditions and camera configurations, we can clearly distinguish in Figure 2 (a) the part surveyed by CV1 (i.e., Seacat AUV, the darker outer loop) and CV2 (i.e., Sparus II AUV, the brighter inner loop). Note that, while no pre-processing was needed to create the presented mosaic, we should consider compensating for these differences in illumination as a further step after creating the mosaic. Furthermore, we can see in Figure 2 (b) that there is overlap between the trajectories of both CV vehicles thanks to the transverse transect, thus ensuring that the two initial rough trajectories are jointly optimized.

After the construction of this multi-vehicle map, we join the rest of the information to further constraint the trajectories of the non-camera vehicles using the acoustic measurements. In Figure 3 we can see the differences between

the non-optimized trajectories on the left (constructed as explained in Section 6) and the optimized ones on the right. Note how the in-vehicle trajectories of the CVs are moved to be consistent with the rest of the vehicles in the formation, thus improving the overall georeferencing of the system. Table 1 summarizes the numbers of nodes and factors in the optimized graph.

To conclude, Figure 4 shows a slanted version of the optimized trajectories in 3D, with a deliberately exaggerated Z. One can see the two CV trajectories being somehow slanted and intersecting with LSV/GCV trajectories, which kept constant depth. This suggests that the terrain, while being close to planar, had some slope and was far from parallel to the sea surface. Given the optimization of the mosaic frame  $\{M\}$  inside our framework, we are also able to measure this slope.

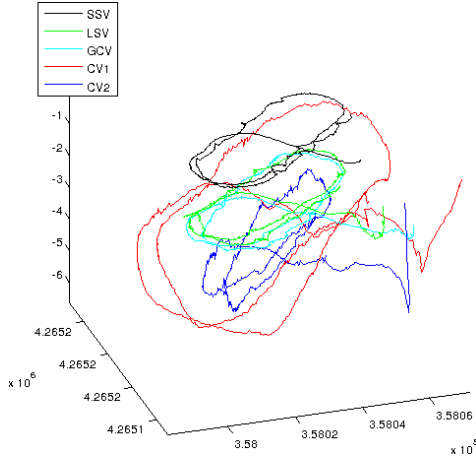


Fig. 4. Slanted view of the optimized trajectories (UTM coordinates), with exaggerated Z.

Table 1. Number of nodes and factors in the final optimized graph. Note that there is one additional node/factor in the total to account for the mosaic frame node and prior respectively.

Nodes		Factors	
Vehicle	Num.	Type	Num.
SSV	6498	Optical	3519
LSV	1963	Ranges	2939
GCV	2350	USBL	1027
CV1	2146	Relative	15057
CV2	2105	Absolute	15062
Total	15063	Total	37605

## 8. CONCLUSIONS

We presented a global alignment off-line optimization procedure to compute the joint trajectory estimate of a set of vehicles surveying the seabed within the formation described in the MORPH EU-FP7 project. The system is generic enough to allow the augment/decrease of the number of vehicles to take into account, thus being usable in other multi-vehicle configurations.

We validated our method by testing it in a real environment. The results obtained show that the mosaicing pipeline renders smooth multi-vehicle mosaics, and also that the joint optimization with acoustic messages provides a coherent global alignment of all the trajectories of the vehicles in the same frame. Moreover, the deliberate abstraction in depth for the vehicle poses is not only useful to render the joint optimization with vehicles not having an altimeter, but also to compute the possible slope that the area under consideration may contain.

As future work, we want to study the inclusion of outlier-robust measures in the system. Range or USBL measures are likely to contain outliers, which the current framework could not deal with. Consequently, we want to study the possible integration of outlier-robust methodologies, such as those surveyed in Sünderhauf and Protzel (2013).

## REFERENCES

Bay, H., Ess, A., Tuytelaars, T., and Van Gool, L. (2008). Speeded-up robust features (SURF). *Computer Vision*

and *Image Understanding*, 110(3), 346–359.

- Elibol, A., Gracias, N., and Garcia, R. (2013). Fast topology estimation for image mosaicing using adaptive information thresholding. *Robotics and Autonomous Systems*, 61(2), 125 – 136.
- Elibol, A., Kim, J., Gracias, N., and Garcia, R. (2014). Efficient image mosaicing for multi-robot visual underwater mapping. *Pattern Recognition Letters*, 46(0), 20 – 26.
- Ferrer, J., Elibol, A., Delaunoy, O., Gracias, N., and Garcia, R. (2007). Large-area photo-mosaics using global alignment and navigation data. In *IEEE Oceans*, 1–9.
- Fischler, M.A. and Bolles, R.C. (1981). Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 381–395.
- Howard, A. (2006). Multi-robot simultaneous localization and mapping using particle filters. *The International Journal of Robotics Research*, 25(12), 1243–1256.
- Kalwa, J., Pascoal, A., Ridao, P., Birk, A., Eichhorn, M., and Brignone, L. (2012). The european R&D-project MORPH: Marine robotic systems of self-organizing, logically linked physical nodes. In *IFAC Workshop on Navigation, Guidance and Control of Underwater Vehicles (NGCUV)*.
- Kummerle, R., Grisetti, G., Strasdat, H., Konolige, K., and Burgard, W. (2011). G2o: A general framework for graph optimization. In *IEEE International Conference on Robotics and Automation (ICRA)*, 3607–3613.
- Kunz, C. and Singh, H. (2013). Map building fusing acoustic and visual information using autonomous underwater vehicles. *Journal of Field Robotics*, 30(5), 763–783.
- Paull, L., Seto, M., and Leonard, J. (2014). Decentralized cooperative trajectory estimation for autonomous underwater vehicles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 184–191.
- Pizarro, O., Eustice, R.M., and Singh, H. (2009). Large area 3-D reconstructions from underwater optical surveys. *IEEE Journal of Oceanic Engineering*, 34(2), 150–169.
- Rekleitis, I., Dudek, G., and Milios, E. (2001). Multi-robot collaboration for robust exploration. *Annals of Mathematics and Artificial Intelligence*, 31(1-4), 7–40.
- Roumeliotis, S. and Bekey, G.A. (2002). Distributed multirobot localization. *IEEE Transactions on Robotics and Automation*, 18(5), 781–795.
- Smith, R., Self, M., and Cheeseman, P. (1987). Estimating uncertain spatial relationships in robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, 850–850.
- Sünderhauf, N. and Protzel, P. (2013). Switchable constraints vs. max-mixture models vs. RRR – a comparison of three approaches to robust pose graph slam. In *IEEE International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany.