

Creating Metric-topological Maps for Large-scale Monocular SLAM

Eduardo Fernández-Moral, Javier González-Jiménez and Vicente Arévalo

Departamento de Ingeniería de Sistemas y Automática, University of Málaga, Bulevard Louis Pasteur 35, Málaga, Spain

Keywords: Monocular SLAM, Metric-topological Map, Map Partitioning.

Abstract: In the last very few years, monocular SLAM approaches based on bundle adjustment are achieving amazing results in terms of accuracy, computational efficiency, and density of the map. When such solutions are applied on large scenarios it is crucial for the system scalability to maintain a map representation that permits efficient map optimization and augmentation. In order to cope with such large maps, we present an on-the-fly partitioning technique which allows abstraction from the metric map to operate more efficiently. The result is a metric-topological arrangement where the areas with highly-connected observations are grouped in submaps weakly interconnected to each other. This is accomplished by progressively cutting a graph representation of the map, where the nodes are keyframes and the arcs between them represent their shared observations. The experimental results indicate that the proposed approach improves the efficiency of monocular SLAM and provides a metric-topological world representation suitable for other robotic tasks.

1 INTRODUCTION

Monocular SLAM is an appealing way of solving the localization and mapping problem in mobile robotics because cameras are inexpensive, compact, easy to calibrate and consume low power. During the last years monocular SLAM has advanced notably with the use of parallel processing and efficient algorithms for data association and map optimization. It has made possible that current state-of-the-art approaches can operate accurately in some large scale scenarios, facilitating its application in a wide range of areas such as augmented reality, scene reconstruction and, particularly, mobile robotics.

The increasingly larger maps that are now possible with monocular SLAM are fundamental to cope with a wider range of real autonomous robotics applications. Such ability to operate in large scale brings the need of appropriate strategies for managing the map. Applying abstraction (as humans do) is an effective way of dealing with the huge amount of detail present in large metric maps. The result of such abstraction process is the so-called metric-topological map, consisting of a two-layer representation, one containing pure geometrical information and a second one containing higher level symbolic information (Thrun, 1998).

The benefit of a metric-topological arrangement is twofold: on the one hand, it offers a natural

integration with symbolic planning that permits a robot to reason about the world and to execute high level tasks (Galindo et al., 2005; Segvic et al., 2009). On the other hand, the efficiency and scalability of the SLAM process itself are improved by limiting the scope of localization and mapping to the region of the environment where the robot is operating. Also, loop closure and relocalisation can be more efficiently solved using topological information (Savelli and Kuipers, 2004; Angeli et al., 2009; Fernández-Moral et al., 2013).

In this work, we present an online submapping technique which creates a topological representation of the world from the metric map being built by a monocular SLAM technique. The key idea of our proposal is to cluster in the same submap those keyframes with higher observation overlap. This presents some important advantages over other approaches (as it will be explained latter on). The generated map consists of a topological structure composed of nodes representing local metric maps and arcs representing relative geometric transformations among the so-called submaps. In this paper, we will focus on the benefits of such a hybrid map for improving the efficiency and scalability of conventional (metric) monocular SLAM, concretely PTAM (Klein and Murray, 2007).

Next, we discuss some relevant related work and

explain in detail the advantages of our approach. We then describe our partitioning procedure and show how it is combined with the SLAM process (PTAM). The experiments and its results are presented next, and finally, we expose the conclusions of our work.

2 RELATED WORK

2.1 Construction of the Metric Map

Many solutions have been presented to build metric maps with monocular SLAM since Davison (Davison, 2003) presented the first real-time solution for the problem in 2003. Two main strategies have been applied since then: Bayesian filtering (following the work of Davison) and Bundle Adjustment (BA) on keyframes, as introduced in (Klein and Murray, 2007). The latter represents the base for the current state of the art since it allows handling denser maps and generally offers a better ratio accuracy/cost (Strasdat et al., 2010).

BA, traditionally used as an offline method for *Structure from Motion (SfM)*, is now widely used in visual SLAM thanks to the introduction of parallel processing and efficient algorithms which exploit the sparse structure of the problem. Its application to visual SLAM was inspired by real time visual odometry and tracking (Nistér et al., 2005), where the most recent camera poses were optimized to achieve accurate localization. In such line, PTAM selects keyframes and applies BA in a fixed size window, around the last keyframe incorporated, to obtain good metric maps and accurate localization. Then, once the local optimization is performed, a low priority global BA is run to improve the map consistency. This approach is extended in (Holmes et al., 2009) by combining it with relative bundle adjustment - RBA - (Sibley et al., 2009), allowing fixed-time, consistent exploration. An improvement of the latter to exploit the problem' sparse structure was recently presented by (Blanco et al., 2013).

The work of (Strasdat et al., 2011) is also related to RBA, they propose a double window optimization: a first window as in PTAM and a second one including the periphery of the first to improve consistency by optimizing a pose-graph. Despite the impressive results obtained, such unique map solution has intrinsic limitations for managing maps of real large environments. To avoid such a limitation, we propose a topological arrangement in local metric maps.

2.2 Dividing the Map

Map division has been addressed in a number of works. Some relevant examples are: the Atlas framework (Newman et al., 2003), where a new local map is started whenever localization performs poorly in the current local map, or the hierarchical SLAM presented in (Estrada et al., 2005), where sensed features are integrated into the current local map until a given number of them is reached. However, none of these provides a mathematically grounded solution based on the particular perception of the scene.

In (Eade and Drummond, 2007), the map is divided in nodes where the landmarks are represented in a local coordinate frame and, these landmarks are updated using an information filter. This method uses the common features between adjacent nodes to calculate their relative pose. A different approach called Tectonic-SAM (Ni and Dellaert, 2007) uses a "divide and conquer" approach with locally optimized submaps in a Smoothing and Mapping framework (SAM). This approach is improved in (Ni and Dellaert, 2010) to build a hierarchy of multiple-level submaps using nested dissection.

Other works employ "graph cut" to divide the map according to a measurable property of the map observations. On that mathematical sound basis, (Zivkovic et al., 2005) addresses the problem of automatic construction of a hierarchical map from images; (Blanco et al., 2008) generates metric-topological maps using a range scanner, and generalizes the approach for other sensors; and (Rogers and Christensen, 2009) splits the map within a Bayesian monocular SLAM framework to reduce the problem complexity.

Our method, which also relies on graph cut, differs from the above works in the way the graph is constructed, which is specifically tailored for BA-based monocular SLAM. Our approach resembles also the stereo-SLAM framework of (Lim et al., 2011) who divide the map keyframes into groups (called segments) according to their geodesic distances in the graph. On the contrary, our map partitioning is independent of the keyframe positions, and is only based on observations acquired from the scene. Concretely, the map is split where there are less shared observations, minimizing the loss of information and therefore, enforcing the coherency and consistency of the submaps.

3 MAP PARTITIONING

Splitting a map into locally metric consistent and globally coherent regions provides some relevant advantages for SLAM. Next, we explain the benefits of such map structure (subsection 3.1), and describe our proposal to obtain this metric-topological arrangement of the map (subsection 3.2).

3.1 SLAM Improvements through Hybrid Mapping

The advantages of applying a coherent map partition in monocular SLAM are diverse: a) all the metric data in each submap can be referred to a local coordinate system, what reduces error accumulation and numerical instability; b) localization can be achieved more efficiently since only those map points in the nearer regions are reprojected to estimate the camera position; c) this map structure permits to approximate the global BA by the individual optimization of the different submaps, thus reducing the computational cost of the optimization process. This last advantage is of special relevance due to the demanding nature of BA, whose complexity ranges from linear to cubic in the number of keyframes depending on the particular point-keyframe structure (Konolige, 2010). Next, we explain the details of this approximation for the global optimization.

Having a map of n landmarks obtained from observations at m keyframes, bundle adjustment can be expressed as

$$\min_{\mathbf{a}_j, \mathbf{b}_i} \sum_{i=1}^n \sum_{j=1}^m v_{ij} d(\mathbf{Q}(\mathbf{a}_j, \mathbf{b}_i), \mathbf{x}_{ij})^2 \quad (1)$$

where

- $d(\mathbf{x}, \mathbf{x}')$ denotes the Euclidean distance between the image points represented by vectors \mathbf{x} and \mathbf{x}' ,
- \mathbf{a}_j is the pose of camera at keyframe j and \mathbf{b}_i the position of landmark i ,
- $\mathbf{Q}(\mathbf{a}_j, \mathbf{b}_i)$ is the predicted projection of landmark i on the image associated to keyframe j ,
- \mathbf{x}_{ij} represents the observation of the i -th 3D landmark on the image of keyframe j and,
- v_{ij} stands for a binary variable that equals 1 if landmark i is visible in keyframe j and 0 otherwise.

Let's now consider that the map is divided into N submaps, each submap, say k , containing m^k keyframes and n^k landmarks, with $k = \{1, \dots, N\}$.

Then, (1) can be rewritten as

$$\min_{\mathbf{a}_j^l, \mathbf{b}_i^k} \sum_{k=1}^N \sum_{l=1}^N \left(\sum_{i=1}^{n^k} \sum_{j=1}^{m^l} v_{ij}^{kl} d(\mathbf{Q}(\mathbf{a}_j^l, \mathbf{b}_i^k), \mathbf{x}_{ij}^{kl})^2 \right) \quad (2)$$

where the combination of subscript i and superscript k refers to the i -th landmark of the k -th submap (e.g., \mathbf{b}_i^k), and similarly l over j refers to the j -th keyframe of the l -th submap (e.g., \mathbf{a}_j^l). Taking into account the observations shared between submaps, this expression can be written as

$$\min_{\mathbf{a}_j^l, \mathbf{b}_i^k} \sum_{k=1}^N \left\{ \underbrace{\sum_{\substack{l=1 \\ l \neq k}}^N \sum_{i=1}^{n^k} \sum_{j=1}^{m^l} v_{ij}^{kl} d(\mathbf{Q}(\mathbf{a}_j^l, \mathbf{b}_i^k), \mathbf{x}_{ij}^{kl})^2}_{A} + \underbrace{\sum_{i=1}^{n^k} \sum_{j=1}^{m^k} v_{ij}^{kk} d(\mathbf{Q}(\mathbf{a}_j^k, \mathbf{b}_i^k), \mathbf{x}_{ij}^{kk})^2}_{B} \right\} \quad (3)$$

where the term A stands for the reprojection error of those landmarks observed from keyframes of different submaps and the term B corresponds to the reprojection error of those landmarks observed from keyframes within the same submap. Both concepts are illustrated in figure 1.b. The first establishes the inter-connection between submaps which is represented by arcs connecting keyframes of different submaps (e.g. arc linking KF-2 and KF-11) and the second sets the intra-connection of the submap which includes the submaps' inner arcs (e.g. arc linking KF-1 and KF-2).

If we are able to divide the map in such a way that the different submaps have few common observations, and assuming that the reprojection errors are independent of the map division, then A becomes negligible with respect to B . Thus, the global optimization can be approximated by

$$\sum_{k=1}^N \left(\min_{\mathbf{a}_j^k, \mathbf{b}_i^k} \sum_{i=1}^{n^k} \sum_{j=1}^{m^k} v_{ij} d(\mathbf{Q}(\mathbf{a}_j^k, \mathbf{b}_i^k), \mathbf{x}_{ij})^2 \right) \quad (4)$$

This approximation is equivalent to optimize each submap independently, which leads to a significant reduction of computational burden. In fact, this approximation is equivalent to the original expression (1) when there are no connections between submaps.

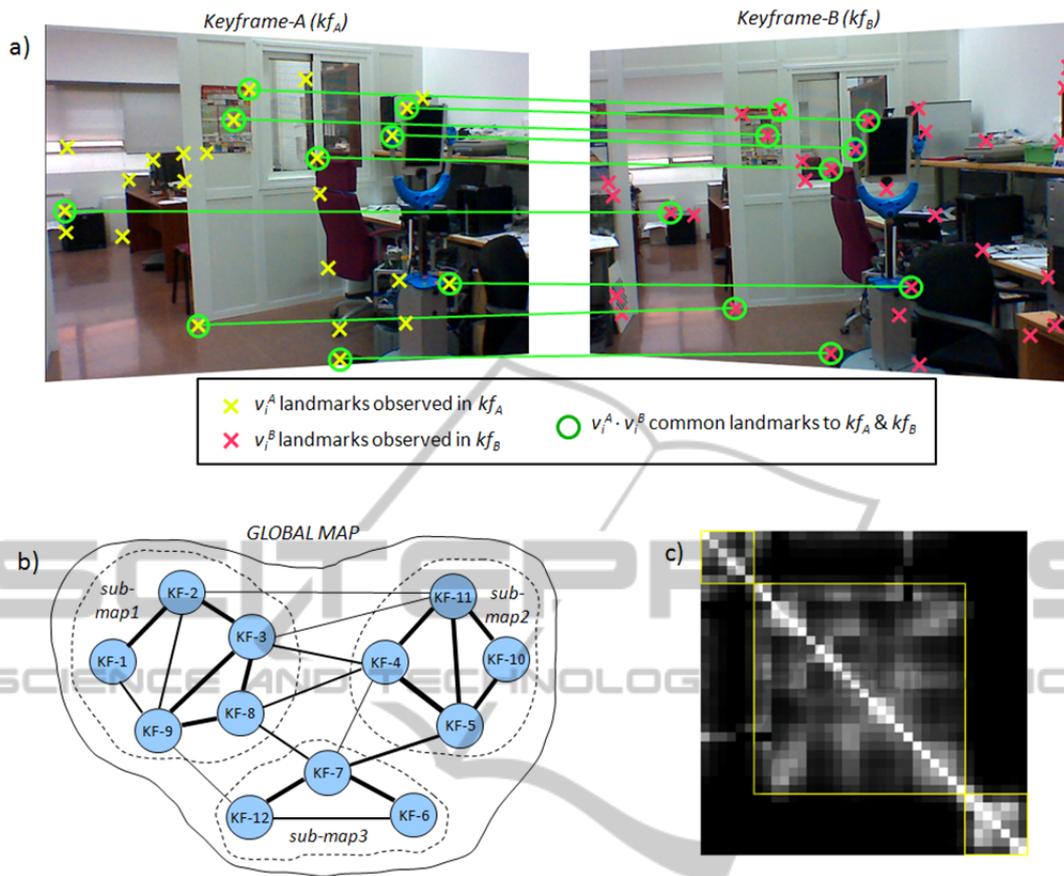


Figure 1: a) Common observations between two keyframes. This is used to calculate the *Sensed Space Overlap* (SSO) (see equation 5). b) Graph-representation of the map where each node represents a keyframe and the arcs are weighed with the SSO calculated between keyframes (thicker arcs mean higher SSO). c) Example of SSO matrix, in which the brightness of the element ij represents the SSO between the keyframes i and j .

3.2 Map Partitioning Method

The approach proposed here to divide the map into coherent regions consists in grouping together those keyframes that observe the same features from the environment. For that, we consider the map as a graph whose nodes represent keyframes and the weight of the arcs are a measure of the common observations between them. There are two critical issues in this partitioning approach: first, the computation of the arc weights; and second, the criterion adopted to perform the partition itself.

As for the first, the arc weights are assigned according to the Sensed-Space-Overlap (SSO), following our previous work (Blanco et al., 2006), particularized for landmark observations. This simple but effective measure represents the information shared by two keyframes. It is calculated with the relation between the number of common landmark observations and the total

number of landmarks observed in both keyframes (see figure 1.a). This is expressed as

$$SSO(kf_A, kf_B) = \frac{\sum v_i^A v_i^B}{\sum v_i^A + \sum v_i^B - \sum v_i^A v_i^B} \quad (5)$$

where v_i^A and v_i^B , similarly to the definitions of the previous section, are binary variables that equal 1 if landmark i is observed in the keyframes kf_A and kf_B , respectively.

Regarding the criterion for partitioning the graph, we follow previous works (Zivkovic et al., 2005; Blanco et al., 2008; Rogers and Christensen, 2009) that apply the minimum normalized-cut (min-Ncut), originally introduced in (Shi and Malik, 2000). The min-Ncut has the desirable property of generating balanced clusters of highly interconnected nodes, in our case clusters of keyframes that cover the same part of the environment. Figure 1 illustrates this concept: figure

1.a shows the common observations in a pair of keyframes whose arc weight is calculated with the SSO (see Eq. 5), and figure 1.b shows a map division into three submaps as produced by the min-Ncut procedure. Notice that the pairs of keyframes with higher SSO (thicker arcs) are grouped together. Figure 1.c shows the symmetrical SSO matrix corresponding to a different, larger map, where the keyframes are arranged according the min-Ncut to give rise to three groups of keyframes or submaps (matrix blocks).

It is important to notice that, in order to guarantee a scalable system when applying map partitioning to visual SLAM, the size of the submaps (i.e. number of keyframes) must be kept bounded. This requirement is not demonstrated mathematically here but it is intuitive to see that as the camera explores new parts of the scene, the new keyframes will have low SSO values (if any) with distant ones in the map. Therefore, the min-NCut will produce new partitions when the system explores unobserved regions of the environment. This can be more clearly understood with the following example: let's consider the case where there are features that are always observed (e.g. the horizon when travelling by train, or when zooming in the scene, or traversing a corridor with the camera pointing in the movement direction) as the new keyframes are selected, they will introduce new features and therefore will reduce the minimum normalized-cut, resulting in the eventual partition of the map. The last two examples represent another advantage of our partition method, which produces natural multi-scale maps when the camera zooms. This insight is supported by all the experiments we have carried out during this work.

4 COMBINATION OF MAP PARTITIONING AND PTAM

This section outlines the combination of our partition procedure and Parallel Tracking and Mapping (PTAM) (Klein and Murray, 2007). PTAM is a monocular SLAM algorithm which performs online BA on keyframes, separating the tracking and mapping stages in two different threads to permit efficient real-time execution. This technique requires an initial map before it starts working autonomously. Such initial map is acquired with a *Structure from Motion* procedure that involves user intervention to select two views with sufficient parallax. Once the

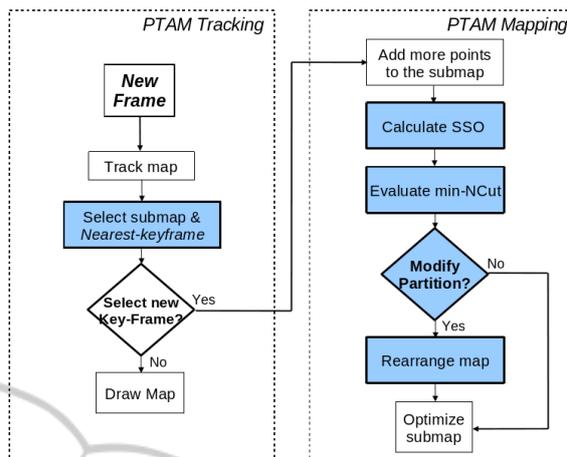


Figure 2: Tracking and mapping threads of PTAM. Blue boxes correspond to the embedded stages to perform the map partitioning.

initial map has been created, the system analyses the images retrieved by the camera to self-localize in the map, while the map is continuously optimized and augmented with new keyframes and landmarks. Such keyframes are selected according to some simple heuristics (see (Klein and Murray, 2007) for more details), and new landmarks are extracted through epipolar search between each new keyframe and its nearest keyframe in the map.

Figure 2 shows a scheme of the proposed partitioning method interacting with PTAM. Our submapping procedure takes action in both of PTAM threads. In the tracking thread, it selects the current submap and the nearest keyframe to the estimated pose after a new image is analyzed. In the mapping thread, after a new keyframe is selected and new landmarks are detected in it, the SSO is evaluated with respect to all the keyframes of the vicinity, which includes all the submaps directly connected to the current submap (see figure 3). The partitioning procedure comes into play after the SSO has been updated, then, the min-NCut is evaluated, and if it results in a different partition, the map is rearranged. This partitioning method is applied dynamically as the map enlarges and may create new submaps as well as merge existing submaps to maintain coherency by grouping keyframes with high overlap. The partitioning process, including SSO computation, min-NCut evaluation and map rearrangement depends on the number of keyframes and landmarks in the vicinity, taking up to 100 ms. in our experiments, which supposes a short time in comparison with the map optimization time.

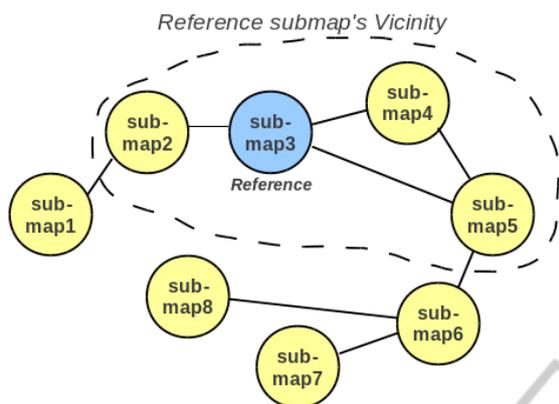


Figure 3: Topological representation of the concept of submap vicinity.

5 EXPERIMENTAL RESULTS

In this section we present some experiments which show the advantages, in terms of efficiency and scalability, of using the proposed metric-topological arrangement of the map instead of a single metric map. The experiments have been carried out using a Philips SPC640NC webcam, connected by USB to a linux-based laptop with an Intel Core2 Duo 2.4 GHz processor, 2Gb of memory and a nVidia GeForce-9400 graphics card.

A first experiment is aimed to illustrate the increase of efficiency in localization at frame rate. For that, we compare the time needed to project map points into the current frame with and without partitioning as the map grows. Both tests have been performed in the same environment, building maps composed of about 45000 points and 1000 keyframes, distributed in 52 submaps for the partitioning case. Figure 4 shows that the time with a unique map grows linearly with the number of map points, whereas with submapping, this time is bounded since only those points in submaps close to the camera are evaluated. This improvement in efficiency becomes more relevant when the map grows nonstop (note that this process is performed with each new frame captured by the camera).

The goal of a second experiment is to quantify the efficiency in the global optimization of the map with our submapping approximation. For that, we have run BA offline after every new keyframe is selected from a recorded video (that is, sequential SfM), measuring the times of each BA completion with and without partitioning. At the end of these tests, the maps created were composed of about

22000 points and 400 keyframes, distributed in 9 submaps for the partitioning case. In order to compare both alternatives in the same conditions, we have included the time of partition management in the BA time for the partitioning test. Figure 5 shows the optimization times vs. the number of keyframes of the whole map for both cases. As expected, for the case without partitioning, the computational cost follows an increasing polynomial trend with the number of keyframes. Conversely, when applying map partitioning, the computational burden is bounded since the BA is applied only on the current submap. For this case, we can observe some abrupt changes in the cost which are produced when the reference submap (the one where the system is localized) switches to a neighbor of different size. Figures 6.a and 6.b show the maps built with both alternatives (different colors represent different submaps in 6.b). We can verify visually their high similarity, and their good alignment, as a result of the continuous optimization previous to the map partition.

Additionally, we are interested in comparing the accuracy of the generated metric map. Due to the lack of a reliable metric to evaluate the map's quality, we have compared visually the different maps considering as ground truth the map obtained offline in the previous experiment (figure 6.a), which is the most accurate we can get. In the map obtained with PTAM (figure 6.c), we can appreciate some regions with depth errors and many outliers (e.g. landmarks detected behind physical walls). These inconsistencies are consequence of the premature interruption of global BA that happens when a new keyframe is selected, what leads to data association errors and the subsequent accuracy decrease with the map size. On the contrary, the map obtained with our approach (figure 6.d) presents no inconsistencies and considerably less outliers than the unique map solution (figure 6.c). This results from the higher efficiency of the submap local optimization, which optimizes regions with highly correlated observations to produce locally accurate submaps.

The results shown in this section have been supported in several tests performed under different conditions: exploring different rooms, re-visiting previous maps, traversing a corridor, zooming to get more detail of the scene, etc. The reader may refer to <http://mapir.isa.uma.es/eFernández/partition> for a video that illustrates the operation of our submapping approach with PTAM in different environments.

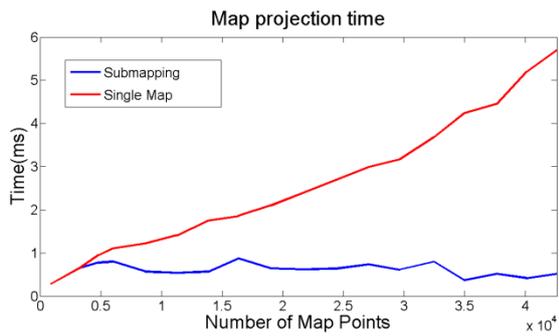


Figure 4: Map projection time for localization with and without map partitioning.

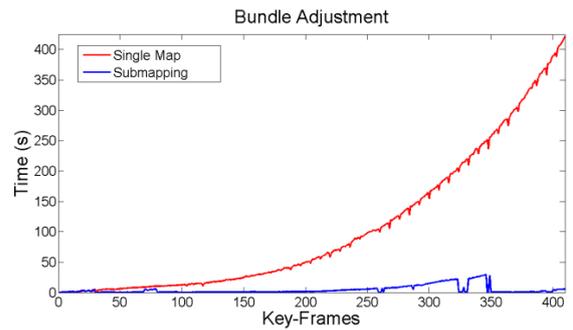


Figure 5: Bundle adjustment computation time (offline) with and without partitioning.

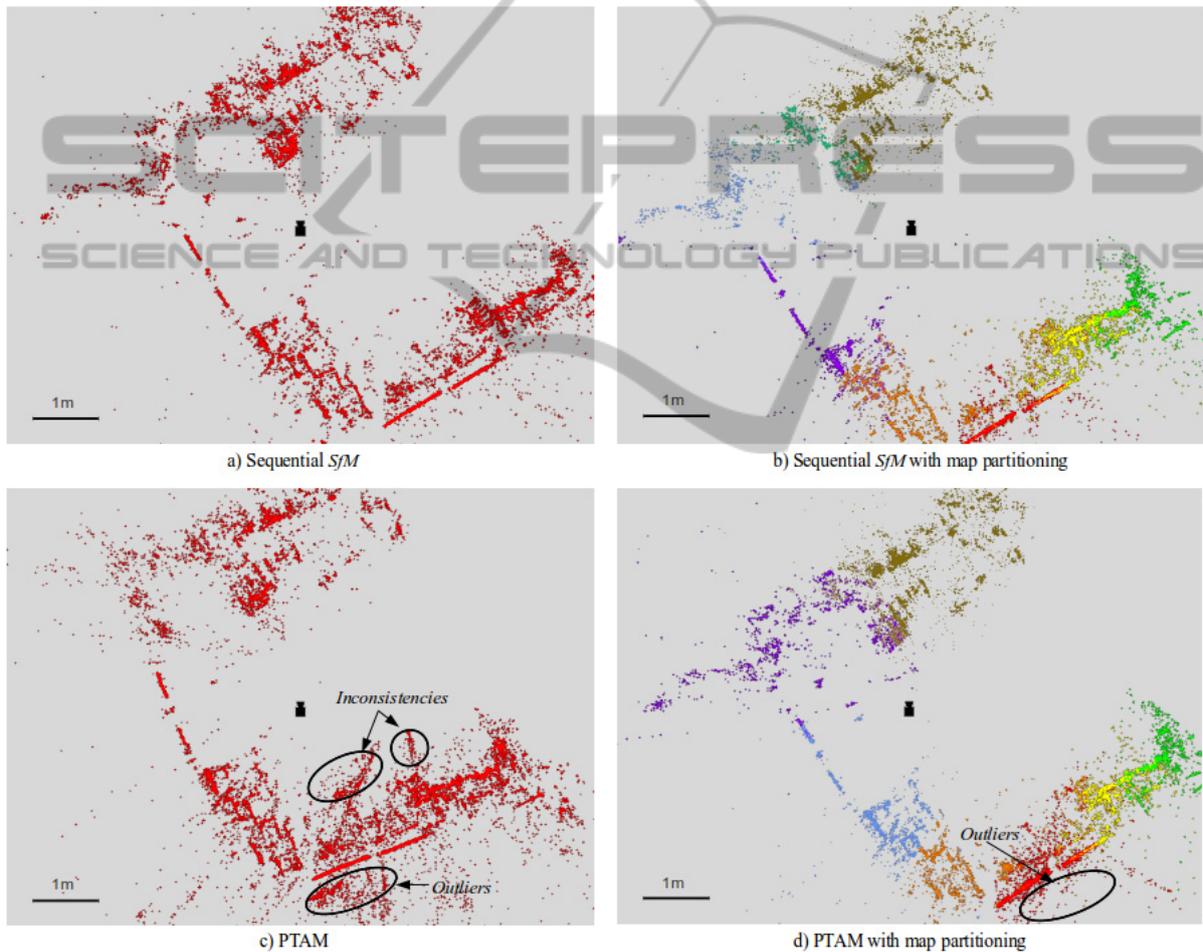


Figure 6: Top view of maps generated in our experiments. All the maps are composed of more than 400 keyframes and 22.000 landmarks. The different colors in b) and d) represent different sub-maps.

6 CONCLUSIONS

This article presents an online submapping method which transforms a metric map into a metric-

topological arrangement of it. This hybrid metric-topological structure improves the scalability of monocular SLAM in two aspects: first, the system rules out unnecessary metric information to perform more efficiently; second, it permits to use an

approximation of BA to reduce computational cost while maintaining map consistency. Besides, the topological arrangement of the map is useful for other tasks, as loop closure, global localization or navigation. Experiments have demonstrated the potential of our approach to obtain efficient map representation in large environments. Future work will focus on exploiting the topological structure of the map for tasks as loop closure and relocalisation.

ACKNOWLEDGEMENTS

This work has been supported by the project "TAROTH: New developments toward a robot at home", funded by the Spanish Government and the "European Regional Development Fund ERDF" under contract DPI2011-25483.

REFERENCES

- Angeli, A., Doncieux, S., Meyer, J.-A., Filliat, D., 2009. "Visual Topological SLAM and Global Localization", In *IEEE International Conference on Robotics and Automation*.
- Blanco, J. L., Fernández-Madrigal, J. A., González, J., 2008. Toward a unified bayesian approach to hybrid metric-topological SLAM. *IEEE Transactions on Robotics and Automation*, 24(2):259-270.
- Blanco, J. L., González, J., Fernández-Madrigal, J. A., 2006. "Consistent observation grouping for generating metric-topological maps that improves robot localization". In *IEEE International Conference on Robotics and Automation*, pp. 818–823.
- Blanco, J. L., González-Jiménez, J., Fernández-Madrigal, J. A., 2013. "Sparsen Relative Bundle Adjustment (SRBA): constant-time maintenance and local optimization of arbitrarily large maps", In *IEEE International Conference on Robotics and Automation*.
- Davison, A. J., 2003. "Real-time simultaneous localisation and mapping with a single camera," In *Proceedings of the International Conference on Computer Vision*.
- Eade, E., Drummond, T., 2007. "Monocular slam as a graph of coalesced observations". In *Proceedings of the International Conference on Computer Vision*.
- Estrada, C., Neira, J., Tardos, J., 2005. "Hierarchical slam: Real-time accurate mapping of large environments". *IEEE Transactions on Robotics*, vol. 21, no. 4, p. 588-596.
- Fernández-Moral, E., Mayol-Cuevas, W., Arévalo, V., González-Jiménez, J., 2013. "Fast place recognition with plane-based maps", In *Proceedings of the IEEE International Conference on Robotics and Automation*.
- Galindo, C., Saffiotti, A., Coradeschi, S., Buschka, P., Fernández-Madrigal, J. A., J. González, 2005. "Multi-hierarchical semantic maps for mobile robotics," In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2278–2283.
- Holmes, S. A., Sibley, G., Klein, G., Murray, D. W., 2009. "A relative frame representation for fixed-time bundle adjustment in monocular SFM". In *Proceedings IEEE International Conference on Robotics and Automation*.
- Klein, G., Murray, D. W., 2007. "Parallel tracking and mapping for small AR workspaces". In *Proceedings of the International Symposium on Mixed and Augmented Reality*.
- Konolige, K., 2010. "Sparse sparse bundle adjustment". In *Proceedings of the British Machine Vision Conference*.
- Lim, J., Frahm, J. M., Pollefeys, M., 2011. Online environment mapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Newman, P., Leonard, J., Soika, M., Feiten, W., Teller, S., 2003. "An atlas framework for scalable mapping". In *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, pp. 1899–1906.
- Ni, K., Dellaert, F., 2010. "Multi-level submap based slam using nested dissection". In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Ni, K., Steedly, D., Dellaert, F., 2007. "Tectonic SAM: Exact, out-of-core, submap-based SLAM," in *IEEE International Conference on Robotics and Automation*.
- Nistér, D., Naroditsky, O., Bergen, J.R., 2005. "Visual odometry". In *Proc. IEEE International Conference on Computer Vision and Pattern Recognition*, pages 652–659.
- Rogers, J. G., Christensen, H. I., 2009. "Normalized graph cuts for visual slam". In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Savelli, F., Kuipers, B., 2004. "Loop-Closing and Planarity in Topological Map-Building," In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, pp. 1511–1517.
- Segvic, S., Remazeilles, A., Diosi, A. and Chaumette, F., 2009. "A mapping and localization framework for scalable appearance-based navigation". *Computer Vision and Image Understanding* 113(2): 172-187.
- Shi, J., Malik, J., 2000. "Normalized cuts and image segmentation". *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, p. 888-905.
- Sibley, G., Mei, C., Reid, I., Newman, P., 2009. "Adaptive relative bundle adjustment". In *Robotics Science and Systems*.
- Strasdat, H., Davison, A. J., Montiel, J. M. M., Konolige, K., 2011. "Double Window Optimisation for Constant Time Visual SLAM". In *IEEE International Conference on Computer Vision*.

- Strasdat, H., Montiel, J. M. M., Davison, A. J., 2010. "Real-time monocular slam: Why filter?". In *IEEE International Conference on Robotics and Automation*.
- Thrun, S., 1998. "Learning Metric-Topological Maps for Indoor Mobile Robot Navigation," *Artificial Intelligence*, v.99, no.1, pp. 21-71.
- Zivkovic, Z., Bakker, B., Krose, B., 2005. "Hierarchical map building using visual landmarks and geometric constraints". In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, p. 2480-2485.

