

Efficient Simultaneous Localisation and Mapping Using Local Submaps

Stefan B. Williams, Gamini Dissanayake and Hugh Durrant-Whyte

Australian Centre for Field Robotics, J04
The University of Sydney, 2006 Australia
e-mail: {stefanw, dissa, hugh}@acfr.usyd.edu.au

Abstract

Autonomous localisation and mapping requires a vehicle to start in an unknown location in an unknown environment and then to incrementally build a map of landmarks present in this environment while simultaneously using this map to compute absolute vehicle location. The theoretical basis of the solution to this problem, known as Simultaneous Localisation and Mapping (SLAM) problem, is now well understood. This paper presents a novel method for building maps of the environment that substantially reduces the computational complexity of the algorithm and improves the data association process. Rather than incorporating every observation directly into the global map of the environment, the Constrained Local Submap Filter (CLSF) described in this paper relies on creating an independent, local submap of the features in the immediate vicinity of the vehicle. This local submap is then periodically fused into the global map of the environment providing improvements in computational efficiency and data association.

1 Introduction

Simultaneous Localisation and Mapping (SLAM) is the process of concurrently building a feature based map of the environment and using this map to obtain estimates of the location of the vehicle. The theoretical basis of the solution to this problem is now well understood. In essence, the vehicle relies on its ability to extract useful navigation information from the data

returned by its sensors. The SLAM algorithm has recently seen a considerable amount of interest from the mobile robotics community as a tool to enable fully autonomous navigation [Castellanos *et al.*, 2000][Dissanayake *et al.*, 2000][Feder *et al.*, 1999][Leonard and Feder, 1999][Newman, 1999][Williams *et al.*, 2001].

This paper presents a novel approach to SLAM that exploits the manner in which observations are fused into the global map of the environment to manage the computational complexity of the algorithm and improve the data association process. Section 2 begins by summarising the key characteristics of the proposed SLAM algorithm, touching on the estimation process using this representation and highlighting the differences between it and the Absolute Map Filter (AMF) commonly used for SLAM. Section 3 presents results illustrating the application of the approach and highlighting the performance of the Constrained Local Submap Filter algorithm applied in simulation and to data collected using an Autonomous Underwater Vehicle. Finally, Section 4 summarises the paper and provides concluding remarks.

2 Constrained Local Submap Filter

The Constrained Local Submap Filter (CLSF) presents a novel scheme for addressing the computational complexity of the SLAM algorithm by allowing the update of the full covariance matrix to be scheduled at appropriate intervals. The method developed here maintains an independent, local submap estimate of the features in the immediate vicinity of the vehicle (see Figure 1). The observations are fused to create a local map of the environment referenced to a local frame of reference whose global position is known. At appropriate intervals, the information contained in the local map is




Figure 1: Local submap state estimation. The vehicle maintains a local map of the features around it. At appropriate intervals, the local map features are fused into the global feature map using appropriately formulated constraints. This approach to map building significantly improves the computational complexity of SLAM.

transferred into the global map. Subject to the usual linearising assumptions, the resulting map and vehicle estimates are identical to those obtained using the Absolute Map Filter (AMF) algorithm [Williams, 2001].

Figure 2 shows the basic steps in this approach to SLAM. At some time, a new frame of reference is defined at the current vehicle position. The estimate of the vehicle pose therefore represents an estimate of the pose of this frame of reference with respect to the global frame. The vehicle is now at the origin of this local reference frame with zero uncertainty at the instant the local frame is created. A new, local vehicle estimate is initialised relative to the new frame and the algorithm begins building a standard AMF SLAM map with respect to the local frame. The estimates in this frame of reference will be shown to be independent of the estimates in the global frame of reference, implying that only a small state covariance matrix must be updated with each observation.

When the decision is made to transfer information contained in the local map into the global map, indicated by the switch in Figure 2, the state vector will contain both local and global position estimates of some of the landmarks. A data association strategy is used to identify the landmarks that have duplicate estimates in the state vector and a constraint based projection is used to yield a single, consistent estimate

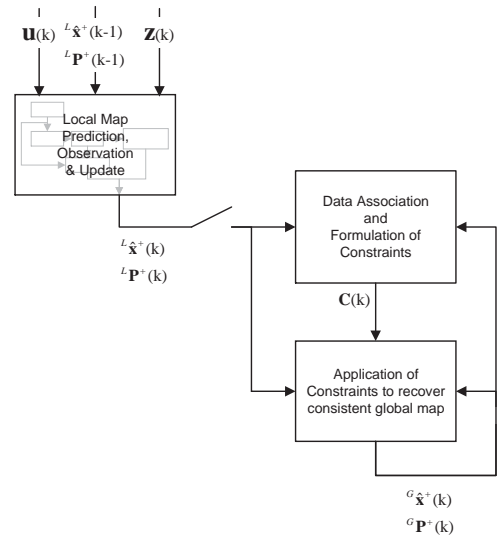


Figure 2: Scheduling of the application of constraints using the Constrained Local Submap Filter. The vehicle operates in a local frame of reference, building an independent map of the features around it. At appropriate intervals, indicated by the switch, the local map is transformed into the global frame of reference and the information fused into the global map using constraints to produce the updated global map estimate.

of these landmark states. This step effectively recovers all of the information available to the filter and allows the filter to generate the AMF state estimate despite the fact that the full covariance matrix has not been updated with each observation. Once the local map has been fused into the global map, a new local map is instantiated at the updated vehicle position and the process of building a local map begins again. This approach blends the ideas of applying constraints presented in the Geometric Projection Filter (GPF) [Newman, 1999] with the intuitive nature of the AMF algorithm.

2.1 The Estimation Process

For the CLSF formulation of the SLAM algorithm, the EKF is used to estimate the pose of the vehicle $\mathbf{x}_v(k)$ along with the positions of the n_f observed features $\mathbf{x}_i(k), i = 1 \dots n_f$. The major difference with this representation arises from the fact that only those feature estimates observed in the current local frame of reference must be updated during an observation. The remaining state estimates represent the feature estimates in the global map of the environment and are not updated until the information in the local map is fused into the global map. This leads to two distinct

independent state estimates while the vehicle is operating in the local submap.

$$\hat{\mathbf{x}}_{cls}^+(k) = \begin{bmatrix} {}^G\hat{\mathbf{x}}_L^+(k) \\ {}^G\hat{\mathbf{x}}_m^+(k) \\ {}^L\hat{\mathbf{x}}_v^+(k) \\ {}^L\hat{\mathbf{x}}_m^+(k) \end{bmatrix} \quad (1)$$

The notation ${}^L\hat{\mathbf{x}}_i^+(k)$ indicates a state estimate of feature $\mathbf{x}_i(k)$ in the local frame of reference. Estimates taken in the global frame of reference will be referred to using the notation ${}^G\hat{\mathbf{x}}_i^+(k)$. The special case of the estimate of the position of the local frame of reference in the global frame is referred to as ${}^G\hat{\mathbf{x}}_L^+(k)$.

The covariance matrix takes on the usual form and contains estimates of the vehicle state covariances and map feature covariances together with the appropriate cross-covariance terms. As shown in [Williams, 2001], the local map estimates are decorrelated from the global estimates and the covariance matrix therefore takes on a block diagonal structure with the global map representing one block and the local map the other. Only the local map estimates need to be updated during the estimation process as a result of this covariance structure.

$$\mathbf{P}^+(k) = \begin{bmatrix} {}^G\mathbf{P}^+(k) & 0 \\ 0 & {}^L\mathbf{P}^+(k) \end{bmatrix} \quad (2)$$

with

$$\begin{aligned} {}^G\mathbf{P}^+(k) &= \begin{bmatrix} {}^G\mathbf{P}_{LL}^+(k) & {}^G\mathbf{P}_{mL}^+(k) \\ {}^G\mathbf{P}_{mL}^{+T}(k) & {}^G\mathbf{P}_{mm}^+(k) \end{bmatrix} \\ {}^L\mathbf{P}^+(k) &= \begin{bmatrix} {}^L\mathbf{P}_{vv}^+(k) & {}^L\mathbf{P}_{vm}^+(k) \\ {}^L\mathbf{P}_{vm}^{+T}(k) & {}^L\mathbf{P}_{mm}^+(k) \end{bmatrix} \end{aligned}$$

where ${}^L\mathbf{P}_{vv}^+(k)$ represents the vehicle covariance in the local frame of reference, ${}^L\mathbf{P}_{mm}^+(k)$, represent the local landmark covariances, ${}^G\mathbf{P}_{mm}^+(k)$, represents the local landmark covariances and ${}^G\mathbf{P}_{LL}^+(k)$, represents the covariance of the estimate of the local frame of reference in the global frame.

2.2 Transforming to the Global Frame

The estimate of the pose of the local frame, ${}^G\hat{\mathbf{x}}_L^+(k)$, represents the relative transformation between the initial frame of reference, \mathcal{F}_G , and the current frame of reference, \mathcal{F}_L with associated covariance, ${}^G\mathbf{P}_{LL}^+(k)$. The transformation matrix $\mathbf{T}(k)$ transforms the local CLSF state estimates into the global frame of reference. This allows the vehicle and map estimates from local frame \mathcal{F}_L to be transformed to the global frame

of reference at any time using the estimated relationships between the frames of reference. The covariance estimate for the transformed states can be generated by also projecting through the transformation matrix

$${}^G\hat{\mathbf{x}}_{cls}^+(k) = \mathbf{T}(k)\hat{\mathbf{x}}_{cls}^+(k) \quad (3)$$

$${}^G\mathbf{P}_{cls}^+(k) = \nabla\mathbf{T}(k)\mathbf{P}_{cls}^+(k)\nabla\mathbf{T}^T(k) \quad (4)$$

2.3 Constraining the Independent Feature Estimates

As information available in the global map is not used while building local maps, some features present in the environment will have estimates in both the local and global state vectors. Once the local map features are transferred to the global frame, there will be multiple entries in the global state vector that correspond to a single feature in the environment. These multiple state estimates of the feature locations must now be fused together. This can be achieved by applying constraints that describe the relationship between these states.

The constraint operation can be considered a weighted projection of the estimates onto the space spanned by the constraints [Newman, 1999]. The weighting factors are functions of the variance of the prior estimates. Given an estimate of the position of landmark i in the global frame, ${}^G\hat{\mathbf{x}}_i^+(k)$, an estimate of the pose of the local frame with respect to the global frame, ${}^G\hat{\mathbf{x}}_L^+(k)$, and an estimate of the landmark position in the local frame, ${}^L\hat{\mathbf{x}}_i^+(k)$, the following constraint must hold

$${}^G\hat{\mathbf{x}}_i^+(k) - ({}^G\hat{\mathbf{x}}_L^+(k) \oplus {}^L\hat{\mathbf{x}}_i^+(k)) = 0. \quad (5)$$

where, following the compact notation given in [Smith and Cheeseman, 1986], \oplus denotes a compounding operator used to calculate the resultant relationship from addition between different frames of reference. Note that these transformations may involve rotations between the frames of reference and are therefore not equivalent to simple vector addition. An example of this relationship is shown in Figure 3. The two estimates of the state ${}^G\hat{\mathbf{x}}_i^+(k)$, indicated by the dashed and solid lines, must be equivalent since they both represent estimates of the same quantity. The application of constraints serves to recover the information contained in the two distinct estimates.

2.4 Computational Complexity

The computational savings that can be realised using this method arise during the update step of an observation. Assume that there are n_f features in the AMF




Figure 3: Multiple global feature position estimates in the submap encoding. The light dashed line shows the estimate of feature i generated through the frame \mathcal{F}_L while the estimate in the global frame is shown by the solid red line labelled ${}^G\mathbf{x}_i$. After constraints are applied, the updated local frame estimate and the estimate of the feature location are also shown.

state matrix. Furthermore, assume that there are n_G features in the global map and n_L features in the local map in which the vehicle is currently operating with $n_L \ll n_G$. Some of the states may be estimates in both the global and local maps such that $n_L + n_G \geq n_f$.

With each observation, the AMF requires that the full covariance update matrix be calculated. This step requires at best $O(n^2)$ operations to compute the matrix update. For the CLSF, however, the update requires only $O(n_L^2)$ operations - a considerable saving for each individual observation. The computationally intensive step in the filter of updating the full covariance matrix is deferred until the constraints are applied.

When the decision is made to apply the constraints to fuse the local estimates into the global map there will be some n_C common states between the global and local map estimates. The computationally intensive update of the global map is dependent on the number of common features between the maps. This will largely be a function of the rate with which new features are observed and thus will be somewhat application dependent.

The CLSF representation of the map presents a mechanism for generating consistent, high accuracy feature sets. Data association is simplified by maintaining an accurate local map of the features surround-

ing the vehicle. The local map of the environment generated using this approach is then fused periodically into the global map. This approach simplifies the data association problem in two significant ways. Firstly, when a new observation is received, it must only be matched against the limited number of features in the local submap. This can lead to significant computational savings given that the new estimate does not need to be compared against every estimate in the global map. Secondly, when the local map is fused into the global map it is necessary to establish the correspondence between feature sets in the local and global maps. As the distribution of features in the environment can now be used to establish correspondences, data association can be made more robust in comparison to the AMF where one observation needs to be matched against a map during data association.

3 Results

This section presents results of the application of the CLSF techniques to the SLAM problem in simulation and on an Autonomous Underwater Vehicle (AUV) operating off the Sydney coastline.

3.1 Simulation

This section presents simulation results of the application of the CLSF techniques to the SLAM problem. The landmarks are randomly distributed throughout the environment and the vehicle takes noisy range/bearing observation using an on-board sensor. The vehicle trajectory is approximately circular.

Figure 4 shows a comparison between the error in the vehicle estimate of the AMF versus the error of the CLSF along with the associated 2σ confidence bounds calculated using the resulting covariance matrix. In this instance, the constraint application is scheduled to happen at fixed intervals. As can be seen, the uncertainty of the CLSF vehicle estimate increases when the vehicle is operating relative to the local submap. When the constraints are applied, however, the uncertainties associated with the estimates of the two filters become identical.

Figure 5 compares the floating point operations required by each algorithm to build and maintain the map. It is clear that for the case of the AMF, the computational burden rises quadratically until all of the observable features have been incorporated into the map. The complexity of each update is then maintained. For the CLSF, on the other hand, each observation incurs only a small computational cost associated





Figure 4: The vehicle estimate errors with the 2σ covariance bounds. The AMF SLAM covariance estimates are shown together with the CLSF case. The global vehicle uncertainty grows for the CLSF case between applications of the constraints but the full covariance estimate is recovered when constraints are applied.

with the update of the local submap estimates. The application of the constraint, however, requires a computationally intensive update of whole the map. By proper management of the update, however, this approach can yield considerable computational savings when compared with the AMF approach, as can be seen in Figure 5 (c). The resulting maps were found to be practically identical.

3.2 Experiments with an AUV

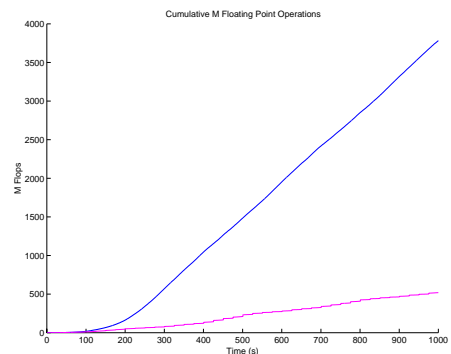
The experimental platform used for the work reported in this paper is a mid-size submersible robotic vehicle called *Oberon* designed and built at the Australian Centre for Field Robotics. The vehicle is equipped with two scanning low frequency terrain-aiding sonars and a colour CCD camera, together with bathymetric depth sensors, a fiber optic gyroscope and a magneto-inductive compass with integrated 2-axis tilt sensor.

The SLAM algorithms have been tested in a natural environment off the coast of Sydney, Australia. The submersible was deployed in a natural inlet with the sonar targets positioned in a straight line at intervals of 10m [Williams *et al.*, 2001]. The vehicle controls were set to maintain a constant heading and altitude during the run. Once the vehicle had reached the end of its tether (approximately 50m) it was turned around and returned along the line of targets. The slope of the inlet in which the vehicle was deployed meant that the



(a) MFlops

(b) Close-up of MFlops




(c) Sum of MFlops


Figure 5: The floating point operations required for the prediction and update stages of the filters. The CLSF (lighter line) has significantly less computational burden while operating in the local submap with a large burden imposed when constraints are applied. This can be significantly reduced by selective application of constraints.

depth of the vehicle varied between approximately 1m and 5m over the course of the run.

The data collected by the vehicle during these trials is presented here to illustrate some of the properties of the Constrained Local Submap Filter in a real world setting. In this instance, the outward journey is used as the original, global map. When the vehicle is turned around to return along the line of sonar targets, a new map is initialised and a new local map is generated. This local map is relative to the final position of the vehicle in the outward leg of the run. When the vehicle reaches the end of its journey, the local map is transformed to the global frame of reference, associations between the feature estimates in the two maps are established and the final map of the environment is generated using the constrained map estimates. Figure 6 shows the two maps generated by this approach.



(a) The first leg global map



(b) The second leg local map

Figure 6: Paths of the robot shown for (a) the outward leg, considered the global map, and (b) the return leg, considered the local map against the final map of the environment.

Finally, Figure 7 shows the resulting constrained map. It is plotted relative to the map generated by the AMF algorithm. As would be expected, there is a good correspondence between the two maps.

4 Summary

This paper has presented a novel approach to autonomous localisation and mapping that generates a local map of the features present in the immediate vicinity of an autonomous vehicle. This local map is then periodically fused into the global map to recover the full global map estimate. This approach to the SLAM problem allows the computationally intensive update of the global map covariance matrix to be scheduled at appropriate interval. It also allows a potentially large number of observations to be fused into the map in a single step, thus increasing the efficiency of the process.

The CLSF can also help to resolve ambiguity in data association. By deferring the association of observations to features in the map, a more informed choice of associations is possible. Asymmetry in the environment can help to resolve any ambiguity that might




Figure 7: The fused maps of the environment generated by applying consistency constraints to the estimated features. The map is plotted on top of the original map generated by the AMF. There is a very good match between the two maps, as should be expected.

arise from the association of a single range/bearing measurement.

This approach has been shown to perform very well in a simulated environment. Simulation allows the performance of the filter to be checked to verify that it is, in fact, generating consistent estimates of the map and vehicle states. It is clear from the results presented that the approach yields nearly identical results to the AMF despite the fact that the entire global map covariance matrix is not updated with each observation. The approach has also been demonstrated to perform well using data collected using an underwater vehicle equipped with scanning sonar.

References

- [Castellanos *et al.*, 2000] J.A. Castellanos, J.M.M. Montiel, J. Neira, and J.D. Tardos. Sensor influence in the performance of simultaneous mobile robot localization and map building. In P. Corke and J. Trevelyan, editors, *Experimental Robotics IV*, pages 287–296. Springer-Verlag, 2000.
- [Dissanayake *et al.*, 2000] M.W.M.G. Dissanayake, P. Newman, H.F. Durrant-Whyte, S. Clark, and M. Csorba. An experimental and theoretical investigation into simultaneous localisation and map building. *Experimental Robotics IV*, pages 265–274, 2000.
- [Feder *et al.*, 1999] H.J.S. Feder, J.J. Leonard, and C.M. Smith. Adaptive mobile robot navigation and mapping. *International Journal of Robotics Research, Special Issue on Field and Service Robotics*, 18(7):650–668, 1999.
- [Leonard and Feder, 1999] J.J. Leonard and H.J.S. Feder. A computationally efficient method for large-

- scale concurrent mapping and localization. In *Proc. Ninth International Symposium on Robotics Research*, pages 169–176. International Foundation of Robotics Research, 1999.
- [Newman, 1999] P. Newman. *On The Structure and Solution of the Simultaneous Localisation and Map Building Problem*. PhD thesis, University of Sydney, Australian Centre for Field Robotics, 1999.
- [Smith and Cheeseman, 1986] R. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *International Journal of Robotics Research*, 5(4):56–68, 1986.
- [Williams *et al.*, 2001] S.B. Williams, G. Dissanayake, and H.F. Durrant-Whyte. Towards terrain-aided navigation for underwater robotics. *Advanced Robotics*, 15(5):533–550, 2001.
- [Williams, 2001] S.B. Williams. *Efficient Solutions to Autonomous Mapping and Navigation Problems*. PhD thesis, University of Sydney, Australian Centre for Field Robotics, 2001.