Long-Horizon Active SLAM system for multi-agent coordinated exploration

Marie Ossenkopf[1,*] Gastón Castro[2,*] Facundo Pessacg[2] Kurt Geihs[1] Pablo De Cristóforis[2]

***Abstract*— Exploring efficiently an unknown environment with several autonomous agents is a challenging task. In this work we propose an multi-agent Active SLAM method that is able to evaluate a long planning horizon of actions and perform exploration while maintaining estimation uncertainties bounded. Candidate actions are generated using a variant of the Rapidly exploring Random Tree approach (RRT*) followed by a joint entropy minimization to select a path. Entropy estimation is performed in two stages, a short horizon evaluation is carried using exhaustive filter updates while entropy in long horizons is approximated considering reductions on predicted loop closures between robot trajectories. We pursue a fully decentralized exploration approach to cope with typical uncertainties in multi-agent coordination. We performed simulations for decentralized exploration planning, which is both dynamically adapting to new situations as well as concerning long horizon plans.**

I. Introduction

Nowadays robots are taking over more and more autonomous tasks, and for this reason, they need to be able to navigate through unknown environments safely. Hence generating a representation of the environment becomes crucial. This includes the necessity of exploring new areas as well as improving the map quality of already visited places. Coordinating multiple agents in the same environment poses additional control challenges but its inherently parallel sensory and computational facilities allows for faster exploration than a single agent.

The autonomous exploration of an unknown environment can be roughly defined as an iterative procedure that consists in the selection of a new goal to explore and navigation towards this goal, which ends by fulfilling a defined condition (mission objective). Meanwhile, the usage of resources (e.g., the exploration time, the length of the trajectory) is optimized. The exploration strategy determines the next target pose in each exploration iteration concerning the current robot pose, the current knowledge about the environment (i.e., current map), and a selected optimization criterion.

Exploration methods have been developed for more than four decades now. For the longest time, the main focus was put on coverage, hence, visiting the largest part of the environment in the shortest amount of time [1]. A central question of exploration is where to place the robot in order to obtain new information about the environment. One of the pioneering work on this was the frontier-based approach proposed by Yamauchi [2], where the robot is guided to the boundary between mapped and unknown environment. Typically, several frontiers exist in a given map, which allows a simple extension of this approach to multi-robot exploration [3]. These methods do however only optimize coverage of the map, leaving the resulting quality out of focus.

Active Simultaneous Localization and Mapping (Active SLAM) covers the problem of choosing control actions that allows exploration while improving map quality and localization performance. Hence Active SLAM needs to anticipate the changes that will occur in the map regarding possible future sensor readings and observations. Using the notion of entropy, information-gain exploration methods have been proposed for cooperative multi-robot systems [4]. Most Active SLAM researches consist of defining a metric to be used as a measure of information gain and optimize this measure to find control policies that maximize the information gain [5], and thus, reducing uncertainties. One significant difference between existing approaches lies in the control actions evaluation horizon and the accuracy for this anticipations.

In this work we propose a decentralized multi-agent Active SLAM approach that is able to evaluate a long planning horizon of actions and perform exploration while minimizing a path and map entropy objective function. Main contributions of this research can be summarized as:

- A decentralized planning arbitration is employed promoting paths that result in better exploration coverage.
- Method generates complex inter-robots loops that maintain estimated uncertainties bounded.
- Short horizon entropy evaluation is carried using a complexity bounded information filter.
- An entropy reduction on predicted inter-robot loop closures is proposed in the long horizon evaluation.

II. Related Work

Atanasov et al. [6] propose a decentralized solution for the multi-sensor active information acquisition problem and show how it can be applied in the context of multi-robot Active SLAM. They build their planning method upon a graph-base SLAM [7], [8] estimation layer which exploits the information matrix sparsity of the SLAM problem. Following this sparse representation, they employ a Square-Root Information Filter (SRIF) [9] for future control actions evaluation where filter updates compute the evolution of pose and map uncertainties in the information form. They claim

*Authors contributed equally to this work.

[1]Distributed Systems Group, University of Kassel, Germany. {ossenkopf, geihs}@vs.uni-kassel.de

[2]University of Buenos Aires (UBA-CONICET-ICC), Argentina. {gcastro, fpessacg, pdecris}@dc.uba.ar

978-1-7281-3605-9/19/$31.00 ©2019 IEEE

to allow long horizon planning with reported experiments within a 12-step planning horizon. Exhaustive filter updates represent a substantial computing cost preventing them from evaluating a higher number of possible action combinations. Instead, the method uses a state machine and attractors [10] to encode long-term information-gain promises in the filter and, in this way, obtain the desired explorative behavior.

An essential element of optimal action search in a highly nonlinear domain as SLAM is action space discretization. It is necessary to create a finite set of candidates to search through. Atanasov et al. obtain their path candidates by combining a fixed number of possible control actions for every step of the planning horizon. This is one possibility of generating paths, which is aimed towards covering most of the action space in the close surrounding of the agent. However, there are better options of generating path candidates for long-distance planning. One of the well studied approaches is the RRT* algorithm [11]. It can quickly produce a tree of path options, that spread over the whole map, while iteratively optimizing a certain trait like distance travelled.

A lot of work in RRT related algorithms has been put into changing its inherent Voronoi bias, which favors a fast map coverage towards some goal oriented bias. E.g., in [12] there is a higher number of random nodes sampled inside a defined region between agent and goal to ensure short distance solutions. Yershova et al. [13] sample new tree nodes depending on the intersection of obstacles with the Voronoi graph to pass difficult obstacle constellations. These methods allow biasing the tree growth to produce more feasible solutions in shorter time. We will use this mechanism to produce path candidates that already support the exploration necessities.

Vallvé and Andrade-Cetto [14] explored the idea of using paths built by an RRT* [15] algorithm as candidates for single-robot Active SLAM instead of generating a path out of a fixed set of candidate actions. Thereby they exploit the favorable properties of RRT* like the Voronoi growing bias towards a fast exploration of the state space, ease of collision handling and node rewiring according to a given cost function. They introduce a joint approximation of path and map entropy as a way to evaluate RRT* nodes and maximize information gain by selecting actions paths that minimize entropy change divided by path distance. Entropy approximation relates three main behaviors: (i) An open loop path entropy estimation that models entropy growth over distance through noise propagation of the process model (ii) A closed loop path entropy estimation that introduces an entropy reduction when a loop closure is predicted between a new node configuration state and an already visited area (iii) An entropy exploration reward based on new explored space along the path. In all cases the method represents a quick entropy approximation for each RRT* node evaluated, allowing to consider a high number of control steps. Vallvé and Andrade-Cetto compare their approach in simulation against frontier based exploration and simpler heuristic Active SLAM approaches reporting benefits in terms of map coverage, estimates error and final overall uncertainties.

In this work action candidates are generated using a variant of RRT* and we use a combination of approaches previously introduced in [6] and [14] for entropy evaluation. Exploratory trajectories are generated with complex inter-robot loops exploiting the multi-agent aspect of the exploration without the need of artificial attractors. Tree nodes of a path are divided into two types: (i) short horizon planning nodes with actions that will have immediate effect and (ii) long horizon planning nodes with actions that will occur later on the path. Short horizon nodes will be exhaustively evaluated using filter updates while long horizon nodes will be evaluated using approximations proposed in [14] with extensions considering loop closures between different robots. Lastly we will describe a multi-agent coordinated scheme that allows decentralized planning arbitration promoting paths that result in better exploration coverage.

III. Decentralized Multi-Agent Exploration

The planning process is built over a graph-based SLAM estimation layer that maintains a probabilistic estimate of robots poses and discovered map landmarks positions. The problem is solved with a nonlinear least-squares approach performing iterative linearization and exploiting the inherent sparsity of the SLAM problem working over the Λ information matrix of the system. In this way, the problem is parametrized as $p(x) = \mathcal{N}(\mu, \Sigma)$ where $\Sigma = \Lambda^{-1}$. The associated linear system is solved using Cholesky factorization over the information matrix obtaining $\Lambda = C^{\intercal} C$ where C is upper triangular and known as the square root of the information matrix. Along with the sparse set of discovered landmarks we maintain also a grid map representation.

A. RRT* Tree Growing

Each robot performs an independent control actions search building his own RRT* tree. To ease the computational burden, on every iteration of our algorithm (see 1), only a pre-defined number of new nodes are sampled. We use a biased sampling to shape the RRT* tree growth towards regions of interest, while the rewiring minimizes the traveled distance. Inspired by [6], there are four different sampling types *a) Random, b) Explore, c) Improve Map, d) Improve Localization*, corresponding to sampling in: a) the complete configuration space, b) only over unexplored regions, c) around areas with highly uncertain landmarks, and lastly, d) around areas with well localized landmarks. According to the current situation, the chances for sampling from the four types differ. The higher the agent's localization uncertainty, the higher is the chance to sample from d). The same applies to the uncertainty of the worst localized landmark position and c), as well as to the distance to the closest unexplored region and b).

A novelty in comparison to [14] approach is that after every control action the RRT* tree is inspected for feasibility. The RRT* root node is redefined rewiring achievable branches and deleting those that are no longer valid (as we use a non-holonomic robot model not all of the old RRT* branches can be reached from the new position). It is critical that the trajectory execution is accurately tracked for practical real-time implementations of RRT*s preventing deviations from growing too large, which would make a complete re-initialization of the tree necessary [16].

B. Actions Evaluation

To choose the best path to take, it is necessary to evaluate the effect of a path on the derived map and resulting localization uncertainty. As we do not have prior information of where landmarks lie in the unexplored area, it is only possible to obtain an estimation of the true effect of a certain path on the underlying SLAM estimation layer. Every node of the tree represents a possible future state configuration and transitions between nodes defines actions. We adopt a similar notation to the one introduced in [14], where $x^i_{1:t}$ represents the trajectory realized by the robot i after executing a set of relative motions $u^i_{1:t}$, one grid maps m^i_t can be rendered for each robot based on the sensor measurements $z_{1:t}$ obtained. We will refer to $x_{1:t}$, $u_{1:t}$, $z_{1:t}$ and m_t as the combined trajectories, control actions, sensor measurements and grids of all robots. Every branch of the RRT tree i would define a path candidate a^i_t as a sequence of future control actions $u^i_{t+1:T}$, which would result in the sequence of node's configuration states $x^i_{t+1:T}$. Sensor measurements can be predicted along this path obtaining $z^i_{t+1:T}$ and a expected grid map m^i_T can be render.

Following [14], we choose the best action to follow as the one that minimize the entropy (H) change divided by the distance traveled

$$a^{i^*}_t = \underset{u^i_{t+1:T}}{\operatorname{argmin}} \frac{H(x^i_{1:T}, m^i_T | u^i_{1:T}, z^i_{1:T}) - H(x^i_{1:t}, m^i_t | u^i_{1:t}, z^i_{1:t})}{dist(u^i_{t+1:T})}. \tag{1}$$

Joint differential entropy will be approximated as proposed in [17] with

$$H(x, m|u, z) \approx H(x|u, z) + \alpha(p(x|u, z))H(m|\mu, z) \tag{2}$$

where $\alpha(p(x^i|u^i, z^i)) = det(\Sigma^i_{tt})^{-1}$. In this way, $H(m|\mu, z)$ is computed using the mean trajectory estimate μ and models a way to decrease entropy and introduce an exploration reward based on the amount of predicted space discovered by the path. This approximated map entropy is going to be relevant only when the robot localization is accurate as it will be weighted by determinant of the last estimated pose covariance $det(\Sigma^i_{tt})^{-1}$ as suggested in [14].

Planning horizon T is divided as $T = \{T_{sh}, T_{lh}\}$ where $[t+1 : T_{sh}]$ will cover those nodes that are within an defined area around the current robot pose and $[T_{sh} + 1 : T_{lh}]$ will encapsulate those that are beyond the short-horizon area. Proposed approximation (2) allow us to effectively divide evaluation treating path and map entropy estimation separately. Furthermore we divide path entropy estimation in two, a short horizon estimation will consider sparse landmarks of the SLAM estimation layer to produce a more accurate approximation while a long horizon estimation will model desirable long term behaviors.

Following a detailed explanation of each entropy estimation term, otherwise stated, we will omit i superscript referring to each possible robot.

*1) **Short-Horizon path entropy estimation**:* In short-term planning we are most interested in maintaining proper localization avoiding movements that leads to immediate counterproductive results. For this reason, each robot defines a set of variables of interest $s_t = \{x_t, l_1, \ldots, l_n\}$ where x_t is the latest estimated trajectory pose and $l_1, \ldots, l_n$ are close range landmarks up to a certain distance threshold. A sub matrix $\Lambda_t = [\Lambda]_{s_t}$ is taken using rows and columns corresponding to variables of interest s_t. This information matrix represents conditional probability distribution of said variable and allows us to work only with a fixed pre-defined number of variables. Applying Cholesky decomposition we have $\Lambda_t = C_t^\intercal C_t$ and, as described in [6], forward state configurations are used to predict future sensor readings that are integrated using a Square-Root Information Filter (SRIF) [9]. In this way, square-root information matrix progression is estimated through candidate nodes providing C_k with $k \in [t, \ldots, T_{sh}]$.

This process effectively discards global trajectory information relating old poses and landmarks allowing evaluation of information gain only in local terms relative to the last estimated pose and nearby landmarks. As the number of considered variables remains fixed covariances can be then computed in constant time by inversion of square-root matrices $\Sigma_k = C_k^{-1} C_k^{-\intercal}$ and pose marginals are used to approximate required path entropies as

$$H(x_{1:k}|u_{1:k}, z_{1:k}) \approx \frac{1}{2} \ln \left((2\pi e)^{\frac{d}{2}} det([\Sigma_k]_x) \right) \tag{3}$$

being d the dimension of the individual pose vector, $d = 3$ in our case.

*2) **Long-Horizon path entropy estimation**:* In long-term planning we are most interested in modeling proper rewards for robot paths that will probably produce loop closures, either with his own performed trajectory or with that of any other robot. For this, a candidate robot configuration x^i_k must fall inside the matching area of an estimated pose x^j_l belonging to the realized trajectory of robot j. An efficient iterative approximation introduced in [14] is extended to support the multi-robot case. In open loop, path entropy is averaged over all individual pose marginals

$$\begin{aligned} H(x_{1:k}|u_{1:k}, z_{1:k}) \approx &\frac{k-1}{k} H(x_{1:k-1}|u_{1:k-1}, z_{1:t-1}) \\ &+ \frac{1}{k} \ln \left((2\pi e)^{\frac{n}{2}} det(\Sigma_{kk}) \right) \end{aligned} \tag{4}$$

where $k \in [T_{sh} + 1, \ldots, T_{ls}]$ and Σ_{kk} is computed by noise propagation of the motion model.

To evaluate the effect of a predicted loop closure the path entropy reduction can be calculated using a loop closure sensor model and a predicted innovation covariance proposed in [17]. All pose marginal covariances Σ_{nn} change to new values Σ'_{nn} ($\forall n \in [1, k]$) if a loop closure is expected in a path and determinant ratio changes can be establish as $\rho_h = det(\Sigma'_{nn})/det(\Sigma_{nn})$. Taking only new covariances of loop nodes x^i_k and x^j_l, Σ'_{kk} and Σ'_{ll} the loop closure information gain can be linearly approximated through the path. In case that x^i_k and x^j_l belongs to the trajectory of the same robot ($i = j$), a "clean" loop is assumed where only marginal covariances of nodes enclosed in the loop will be affected. In case that the loop closure has been predicted between trajectories of different robots ($i \neq j$) the covariance with higher determinant is taken as reference and the approximation is done considering an improvement in the whole trajectory of the other robot up to the reference node.

Being $\gamma = \operatorname{argmax}_{n \in [k,l]} det(\Sigma_{nn})$, the loop closure information gain introduced is linearly approximated as

$$\Delta H(x^i_{1:k}|u^i_{1:k}, z^i_{1:k}) \approx \begin{cases} \frac{1}{k} \ln \prod_{n=1}^{k-l+1} \left(\rho_l + \frac{\rho_k - \rho_l}{k-l+1} n \right) & \text{if } i = j \\ \frac{1}{\gamma} \ln \prod_{n=1}^{\gamma} \frac{\rho_\gamma}{n} & \text{if } i \neq j \end{cases} \tag{5}$$

*3) **Map entropy estimation**:* We work over the occupancy grid map m_T with predefined cell size s. To each cell there is an associate classification probability p_c which values 0 when the cell is free and 1 if is occupied and, in case that we do not have information about a cell, p_c value is assumed as 0.5. The map entropy can be then calculated as a sum over all the cells in m_T

$$H(m_T|u_{1:t}, z_{1:t}) = -s^2 \sum_{c \in m_T} (p_c \ln\, p_c + (1-p_c)\ln(1-p_c)). \tag{6}$$

Note that if all map cells are known and classified, either free or occupied, total map entropy is 0. When $p_c = 0.5$ each cell increases entropy by a maximum amount $\lambda \simeq 0.7$.

Expected map entropy reduction after moving to a new state configuration is related with how many cells will change its classification probability from unknown ($p_c = 0.5$) to 0 or 1. As we want to reward long term trajectories we will speculate the amount of map grid that will be discover along a path by counting explorative nodes. A node is declared as explorative if his frustum is considered to reach a grid cell that has not been explored before. Furthermore, we will approximate classification probability change assuming that new discovered cells in the field of view will be completely classified as free or occupied. This will introduce an information gain of λ per predicted discovered cell for each explorative node. Using this approximation map entropy information gain is

$$H(m_k|u_{1:k}, z_{1:k}) - H(m_t|u_{1:t}, z_{1:t}) \approx -s^2 \sum_{n=t+1}^{k} d_{x_n} \lambda \tag{7}$$

where $k \in [t+1, \ldots, T]$ and d_{x_n} is the number of predicted newly discovered cells by x_n.

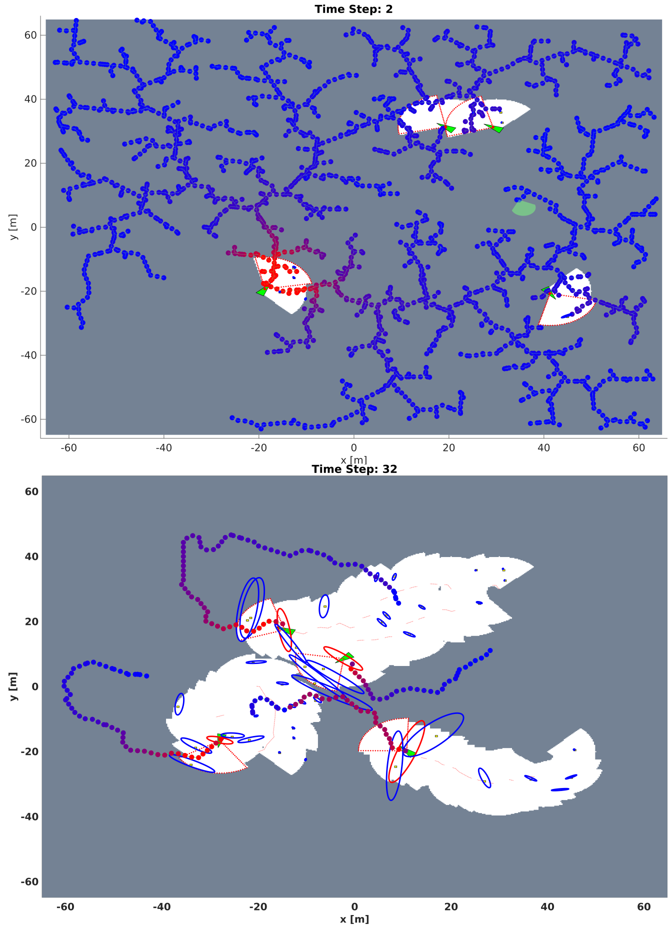

Fig. 1: Instances of the multi-agent active SLAM system simulation. Top image shows an early state with the complete RRT* tree of a robot. Bottom image shows only currently selected branches and can be seen that some robots seek for loop closures, in some cases, with trajectories of a different robot. Real robots positions are in green, red ellipses shows estimated covariances, dotted red areas exhibit robot's field of view. Landmarks positions are shown as yellow squares and his estimated covariances are represented with blue ellipses. RRT* branches are exhibit as connected dots colored representing entropy evaluation for each node. Blue nodes represent better joint path-map entropy evaluation.

This approach allows to approximate map entropy reduction for nodes in unexplored areas predicting the amount of space that will be explored without considering actual feasibility of that map innovation. This results in long predicted paths towards unknown space and the method relies on the RRT* growing algorithm to check for feasibility and cut branches that will be occluded by newly discovered obstacles.

C. Anticipating Other Agent's Actions

The essential interference between two agents is the resulting map coverage of their paths. The lower the paths overlap the higher map coverage. There are several problems in anticipating the exact coverage a path will add to the map: (i) there is an uncertainty on localization and the agent may not take exactly the path it planned (ii) density of landmarks on the planned path may be unknown, hence, the amount of added information is unknown (iii) planned paths might be subject to re-planning and hence, won't be followed thoroughly.

We need a measure to evaluate the suitability of the path regarding the resulting coverage, which we then combine with the path entropy estimation to form the merit of a path candidate. To take the before mentioned uncertainties into account, we model path coverage as a probability distribution of visiting a particular partition of the map at a random time step along paths. We compute this probability distribution by assigning all path nodes into a set of bins, which represents a discretization of the map space. Given this representation, we can relate two paths by measuring the Kullback-Leibler divergence between them

$$x_{opt} = \arg\max_{x} D_{KL}(Y||X). \tag{8}$$

where X is the probability distribution resulting from the path candidate x, while Y is the probability distribution resulting from the chosen paths of all other agents combined, this gives us a measure for how likely is that a specific path overlaps with others or if it will add different information to the map. This measure, together with the entropy estimation decides both which agent is allowed to choose a new path candidate as well as which path will be chosen.

D. Decentralized Coordinated Planning

To reach true decentralization in the exploration process, we oriented on the idea of Desaraju and How [18] for token based multi agent coordination. They explore the setting of multiple agents planning goal oriented paths through an obstructed environment. A path chosen by one agent might make the path unfeasible for another one. To handle the problem that two agents might decide on the same path at the same time, they allow only one agent to change its path at a time. This is handled by a token that is passed to the agent with the highest potential in improving its path. This potential is called merit. The high number of path candidates given by the RRT does not allow to calculate the overlap estimation for every combination of paths. The calculation only stays feasible if every agent can assume the paths of others as fixed. Therefore we adopt the idea of only one agent being allowed to change its path at a time.

The decentralized coordinate planning is summarized in Algorithm 1

Algorithm 1 Decentralized Coordinated Planning

1) Perform next action of planned path
   - Apply control movement
   - Measure environment + SLAM Update
2) Update RRT and path candidates (**Section III-A**)
   - Prune now unfeasible branches
   - Grow up to 500 nodes of the tree
3) Evaluate path candidates (**Section III-B**)
   - Estimate entropy change for all branches
   - Choose the 20 highest performing branches
4) Measure path candidates distribution (**Section III-C**)
5) Calculate merit [18]
6) Agent with highest merit changes his current path

IV. Evaluation

We implemented our solution over the simulation framework introduced in [6]. The environment initializes as completely unexplored with predefined initial robots positions and sparse landmarks positions are randomly created. We performed experiments of the method without environment obstacles as that was not the principal scope of the research and RRT planning algorithms are already proven to efficiently handle this aspect. Fig. 1 shows two moments of a simulation. Left side figure exhibits all RRT candidate branches from one robot with low localization uncertainty resulting in $\alpha(p(x|u,z)) > 1$ promoting map exploration reward towards unexplored space (see equation 2). Right side figure shows only best candidate paths that minimizes entropy cost function (equation 1), colors go from red to blue representing nodes of higher to lower entropy evaluation. Colors are normalized only with respect to nodes of his own path, without relating with other robots. It can be seen that robots with higher localization uncertainty $\alpha(p(x|u,z)) < 1$ prioritize paths that will produce loop trajectories.

Performance is quantified in Fig. 2. The two first plots show the Root Mean Square Error (RMSE) of each robot position and orientation which is shown to not exceed 0.6 meters and 8 degrees respectively. Robots pose entropy is effectively bounded through the exploration. The aggregated landmarks entropy exhibits a metrics that averages marginal entropies of all sparse landmarks, it shows an erratic behavior with some areas where entropy gets higher but an overall decrease tendency. This is due to our method not directly considering landmarks uncertainty as part of the entropy evaluation but while performing loop trajectories environment landmarks ultimately improves their estimation. Finally, the last plot represent the environment coverage that shows a quick coverage tendency. It worth mentioning that Atanasov et al. [6] method takes 700 time steps to fully cover the map while our takes just 400 steps.

V. Conclusions

We proposed a system to solve Active Pose SLAM problem, that enables long horizon planning for multiple agents. Thereby, we allow for effective multi-agent planning by giving the possibility to anticipate the effects of agent's actions. This system maintains the covariance pose estimation bounded for all the robots and at the same time improves exploration speed by rewarding long explorative trajectories with long term loop closures.

Acknowledgment

This work was supported by the German Academic Exchange Service (DAAD) and the Argentinian National Scien-

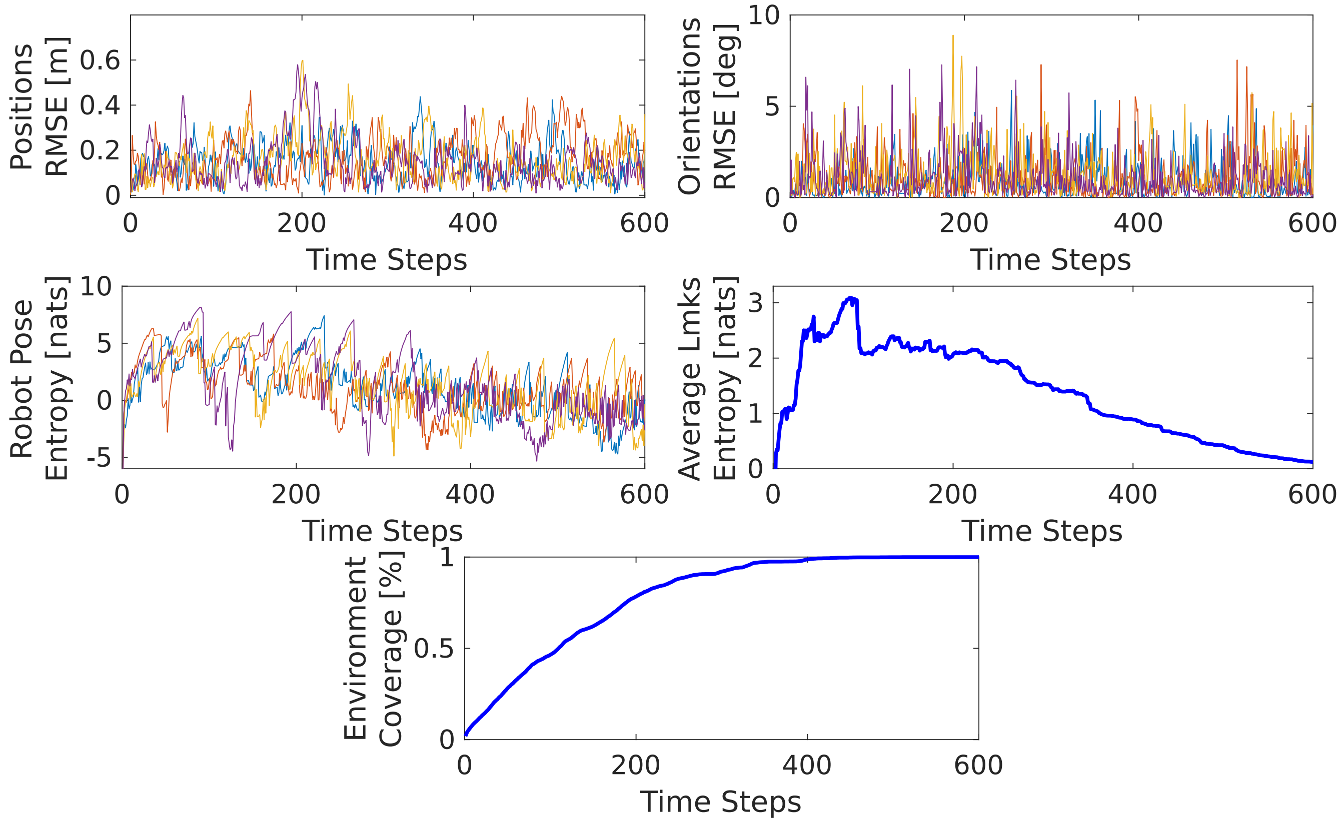

Fig. 2: Results of a simulation using 4 robots, one color for each robot is used for RMSE and entropy plots. 'Average Landmarks Entropy' is an aggregated metric averaging marginal entropies of every sparse landmark.

tific and Technical Research Council (CONICET) under the bilateral Project DA/16/08. We thank Nikolay A. Atanasov for happily sharing his simulation framework with us.

REFERENCES

[1] Sebastian Thrun et al. Robotic mapping: A survey. *Exploring artificial intelligence in the new millennium*, 1(1-35):1, 2002.

[2] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97.*, pages 146–151, July 1997.

[3] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376–386, June 2005.

[4] R. Rocha, J. Dias, and A. Carvalho. Cooperative multi-robot systems: A study of vision-based 3-d mapping using information theory. *Robotics and Autonomous Systems*, 53(3):282–311, 2005.

[5] Ayoung Kim and Ryan M Eustice. Active visual SLAM for robotic area coverage: Theory and experiment. *The International Journal of Robotics Research*, 34(4-5):457–475, 2015.

[6] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas. Decentralized active information acquisition: Theory and application to multi-robot SLAM. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4775–4782, May 2015.

[7] Frank Dellaert and Michael Kaess. Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *The International Journal of Robotics Research*, 25(12):1181–1203, 2006.

[8] Michael Kaess, Hordur Johannsson, Richard Roberts, Viorela Ila, John Leonard, and Frank Dellaert. iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree. *Intl. J. of Robotics Research, IJRR*, 31(2):217–236, February 2012.

[9] Gerald J Bierman. Sequential square root filtering and smoothing of discrete linear systems. *Automatica*, 10(2):147–158, 1974.

[10] C. Leung, S. Huang, and G. Dissanayake. Active SLAM using Model Predictive Control and Attractor based exploration. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5031, Oct 2006.

[11] Iram Noreen, Amna Khan, and Zulfiqar Habib. Optimal path planning using RRT* based approaches: A survey and future directions. *Int. J. Adv. Comput. Sci. Appl*, 7(11):97–107, 2016.

[12] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot. Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2997–3004, Sep. 2014.

[13] Anna Yershova, Léonard Jaillet, Thierry Siméon, and Steven M. LaValle. Dynamic-Domain RRTs: Efficient Exploration by Controlling the Sampling Domain. *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3856–3861, 2005.

[14] J. Vallvé and J. Andrade-Cetto. Active pose SLAM with RRT*. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2167–2173, May 2015.

[15] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011.

[16] Brandon D Luders, Sertac Karaman, Emilio Frazzoli, and Jonathan P How. Bounds on tracking error using closed-loop rapidly-exploring random trees. In *American Control Conference (ACC), 2010*, pages 5406–5412. IEEE, 2010.

[17] Joan Vallvé and Juan Andrade-Cetto. Potential information fields for mobile robot exploration. *Robotics and Autonomous Systems*, 69:68–79, 2015.

[18] Vishnu R Desaraju and Jonathan P How. Decentralized path planning for multi-agent teams with complex constraints. *Autonomous Robots*, 32(4):385–403, 2012.