Leap-Frog Path Design for Multi-Robot Cooperative Localization

Stephen Tully, George Kantor, and Howie Choset

Carnegie Mellon University
{stully@ece, kantor@ri, choset@cs}.cmu.edu

Summary. We present a “leap-frog” path designed for a team of three robots performing *cooperative localization*. Two robots act as stationary measurement beacons while the third moves in a path that provides informative measurements. After completing the move, the roles of each robot are switched and the path is repeated. We demonstrate accurate localization using this path via a coverage experiment in which three robots successfully cover a 20m x 30m area. We report an approximate positional drift of 1.1m per robot over a travel distance of 140m. To our knowledge, this is one of the largest successful GPS-denied coverage experiments to date.

1 Introduction

Localization is critical for the navigational aspect of many robotic applications. Without accurate positioning, a mobile robot would get lost, wander away from its target workspace, and fail to complete its intended task. Additionally, there are many situations where an external positioning system, such as GPS, is unavailable to the robot, e.g. indoors, within dense vegetation, and underwater. To solve the localization problem, a team of robots can employ *cooperative localization* [1] to incorporate relative sensor measurements into a Kalman filter framework that estimates the pose of the robots.

It can be shown that the accuracy of such a filter is dependent upon the path the robots take. This is due to the fact that certain measurements are more informative than others, depending on the vantage point of the sensor. We believe a “leap-frog” path, as in [2, 3], is desirable because it temporarily grounds the increasing uncertainty of the system via stationary robots.

The contribution of this work is the introduction of a new leap-frog path designed to produce informative measurements for three robots performing cooperative localization. We also report a 20m x 30m large-scale GPS-denied coverage experiment with three robots (see Fig.1) that was only possible after the gain in positioning accuracy provided by this new leap-frog strategy.

Fig. 1. Three robots used for experimental evaluation of the proposed leap-frog localization strategy.

2 Related Work

The majority of recent work on cooperative localization ignores the path planning aspect of this topic and instead focuses on filtering. In [1], Roumeliotis et. al. define a Kalman filter framework, similar to our formulation in Sec. 3, that can be used for cooperative localization. In [4], they iterate on this work and present a distributed method for computing the Kalman filter. Finally, the work in [5, 6] studies the growth of uncertainty under different sensing modalities and with a varying number of robots.

Some recent work has addressed path planning and control for robots performing cooperative localization. Hidaka et. al. [7] present derivations to show that with any number of robots, the optimal formation for accurate localization is a “packed circles” configuration. For three robots, they claim that the optimal formation is an equilateral triangle. Trawny et. al. [8], on the other hand, perform optimization over possible multi-robot paths and demonstrate a performance improvement in simulation. Although the optimization is beneficial, we believe this method is susceptible to local minima.

Finally, some research has involved the investigation of leap-frog paths for cooperative localization. Navarro-Serment et. al. [2] use a group of small heterogeneous robots (Millibots) for localization and mapping. The authors use leap-frog paths to help maintain a better estimate over the occupancy grid map and their EKF localization. Kurazume [3] introduces leap-frog paths for cooperative localization as well, with a path that is designed to represent triangle chains of different configurations. Although these paths prove to be accurate solutions for localization, Kurazume does not analyze these paths from an information theoretic standpoint.

3 Cooperative Localization

To reduce motion error, a team of robots can employ *cooperative localization*, which improves the estimate of the state via relative sensor measurements between robots. This type of filtering method can be implemented with an extended Kalman filter, as in [1].

The state vector for the filter X_k is defined,

$$X_k = \begin{bmatrix} X_k^{0T} & X_k^{1T} & \dots & X_k^{N-1T} \end{bmatrix}^T \quad X_k^i = \begin{bmatrix} x_k^i & y_k^i & \theta_k^i \end{bmatrix}^T,$$

where N is the number of robots and X_k^i represents the state of the i -th robot at time step k . In this formulation, (x_k^i, y_k^i) represents the position of the i -th robot and θ_k^i represents that robot's heading.

The state process equation $f(X_k^i, u_k^i)$ is based on a unicycle model,

$$f(X_k^i, u_k^i) = \begin{bmatrix} x_k^i + v_k^i \cos \theta_k^i \Delta t \\ y_k^i + v_k^i \sin \theta_k^i \Delta t \\ \theta_k^i + \omega_k^i \Delta t \end{bmatrix} \quad u_k^i = \begin{bmatrix} v_k^i \\ \omega_k^i \end{bmatrix},$$

where u_k^i is a motion input for the i -th robot, which is composed of a translational velocity v_k^i and a rotational velocity ω_k^i .

The measurement equation for our state is a bearing-only measurement,

$$h_j^i(X_k) = \arctan\left(\frac{y_i - y_j}{x_i - x_j}\right) - \theta_j$$

which represents the relative bearing angle to the i -th robot as measured by the j -th robot. A typical sensor that provides bearing measurements in this form is a monocular camera.

The purpose of the extended Kalman filter (EKF) is to recursively estimate the state mean and covariance matrix with two stages: the prediction step, which produces the estimated mean and covariance, $\hat{X}_{k+1|k}$ and $P_{k+1|k}$ respectively, as well as the measurement update step, which produces an update to the estimated mean and covariance, $\hat{X}_{k|k}$ and $P_{k|k}$ respectively.

3.1 Prediction Step

The EKF prediction step is applied when processing the robot's internal velocities (usually from wheel encoders). The state mean and covariance matrix are computed as follows.

$$\hat{X}_{k+1|k} = \begin{bmatrix} f(\hat{X}_{k|k}^0, u_k^0)^T & f(\hat{X}_{k|k}^1, u_k^1)^T & \dots & f(\hat{X}_{k|k}^{N-1}, u_k^{N-1})^T \end{bmatrix}^T$$

$$P_{k+1|k} = F_k P_{k|k} F_k^T + W_k U_k W_k^T$$

where $\hat{X}_{k|k}$ and $P_{k|k}$ define the estimate from the previous time step, and F_k and W_k are the Jacobians of the state process equation, as defined below.

$$F_k = \begin{bmatrix} \frac{\partial f(\hat{X}_{k|k}^0, u_k^0)}{\partial X_k^0} & 0 & \dots & 0 \\ 0 & \frac{\partial f(\hat{X}_{k|k}^1, u_k^1)}{\partial X_k^1} & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & \frac{\partial f(\hat{X}_{k|k}^{N-1}, u_k^{N-1})}{\partial X_k^{N-1}} \end{bmatrix} \quad W_k = \begin{bmatrix} \frac{\partial f(\hat{X}_{k|k}^0, u_k^0)}{\partial u_k^0} & 0 & \dots & 0 \\ 0 & \frac{\partial f(\hat{X}_{k|k}^1, u_k^1)}{\partial u_k^1} & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & \frac{\partial f(\hat{X}_{k|k}^{N-1}, u_k^{N-1})}{\partial u_k^{N-1}} \end{bmatrix}$$

The matrix U_k is a covariance matrix with 2x2 matrices along its diagonal (all U_k^i for $0 \leq i < N-1$ as defined below).

$$U_k^i = \begin{bmatrix} \alpha |v_k^i| & 0 \\ 0 & \beta |v_k^i| + \gamma |\omega_k^i| \end{bmatrix} \quad (1)$$

Matrix U_k^i represents the covariance matrix for the additive white Gaussian noise that is expected to perturb robot i 's motion input u_k^i . Conventional implementations use a static covariance for this purpose, but we believe a velocity dependent noise model is more accurate. The model in Eq. 1 accounts for the fact that wheel slippage is more pronounced at higher speeds and that zero additive noise should be expected when the robots are stationary.

3.2 Measurement Update Step

To properly incorporate the information provided by the bearing sensors, we perform a correction to the predicted state estimate,

$$\begin{aligned} H_j^i &= \frac{\partial h_j^i(\hat{X}_{k|k-1})}{\partial X_k} \\ K &= P_{k|k-1} H^T (H P_{k|k-1} H^T + R)^{-1} \\ \hat{X}_{k|k} &= \hat{X}_{k|k-1} + K(z_k - h(\hat{X}_{k|k-1})) \\ P_{k|k} &= P_{k|k-1} - K H P_{k|k-1} \end{aligned} \quad (2)$$

where, for M bearing measurements, K is the Kalman gain, H is the measurement Jacobian, and R is an $M \times M$ matrix with diagonal elements σ_z^2 (the variance associated to a single measurement). The Jacobian H is constructed by appending together all row vectors H_j^i for each measurement between a robot i and another robot j . Likewise, h is constructed by appending together all h_j^i to form a column vector. z_k is the measurement vector to which h is associated.

It is important to realize that the effectiveness of this cooperative localization filter is dependent upon the path of the robots. This can be seen in Eq. 2 where a positive definite matrix $K H P_{k|k-1}$ is subtracted from the predicted covariance $P_{k|k-1}$. Since H is dependent upon the state estimate $\hat{X}_{k|k-1}$, the reduction in uncertainty via subtracting $K H P_{k|k-1}$ will vary in amount depending on the configuration of the robots.

4 Leap-Frog Path Design

In Sec. 3, we discuss how the effectiveness of the cooperative localization filter is dependent upon the path of the robots. This suggests that by careful path planning, we can achieve better position accuracy during experiments.

As in [2, 3], we suggest the use of a “leap-frog” path for a team of robots, where at any given time, a subset of the robots temporarily act as stationary measurement beacons while the other robots are in motion.

A “leap-frog” path is intuitively beneficial for the Kalman filter because when robots are stationary, they will not gain any positioning noise, thus temporarily grounding the normally increasing uncertainty of the system. A moving robot can move around at will without concern for its added prediction noise because it can easily visit the nearby stationary robots to drive its position uncertainty down to their level via relative sensor measurements.

4.1 Three-Robot Path Design via Information Gain

The use of three robots for localization is a good fit for bearing-only measurements because the intersection of two bearing rays from two different robots will triangulate the location of a third robot, albeit with error due to noise. To investigate path design for a team of three robots, we consider the measurement update equation for the information filter, which is a dual to the Kalman filter and is commonly used in localization and mapping algorithms, such as [9]. The measurement update is as follows,

$$I_{k|k} = I_{k|k-1} + H^T R^{-1} H,$$

where the information matrix $I_{k|k} = P_{k|k}^{-1}$ is the inverse of the covariance matrix. In this work, we define the information gain $G(X_k)$ as a norm of the positive definite matrix that is added to the information matrix during a measurement update,

$$G(X_k) = \text{tr}(H^T R^{-1} H).$$

The information gain depends on the state X_k through the measurement Jacobian H . We argue that states producing a larger information gain will offer measurements that are more informative to the Kalman filter.

Fig. 2. Stationary robots 0 and 1 are a distance $d/2$ away from the x -axis. An analysis of the information gain is used to obtain the appropriate leap-frog path for robot 2.

To investigate the path optimization problem for three robots, we consider the situation in Fig. 2 where two robots (0 and 1) lie stationary on the y -axis an equal distance away from the x -axis (with a separation distance d). The third robot (robot 2) acts as the moving robot in this leap-frog strategy.

For any pose (x, y, θ) of robot 2 in Fig. 2, the filter will have the following information gain $G(X_k)$,

$$X_k = \begin{bmatrix} 0 & \frac{d}{2} & 0 & 0 & \frac{-d}{2} & 0 & x & y & \theta \end{bmatrix}^T$$

$$G(X_k) = \sigma_z^{-2} \left(6 + \frac{4}{d^2} + \frac{4}{(y - d/2)^2 + x^2} + \frac{4}{(y + d/2)^2 + x^2} \right)$$

To determine the optimal y value for any x position of robot 2, we can take the derivative of the information gain with respect to y , as follows,

$$\frac{\partial G(X_k)}{\partial y} = 8\sigma_z^2 \left(\frac{y - d/2}{((y - d/2)^2 + x^2)^2} + \frac{y + d/2}{((y + d/2)^2 + x^2)^2} \right)$$

By setting the derivative to zero, we can find the y that maximizes the information gain. The solution is $y = 0$, independent of the robot's x position. This implies that for a robot that is “leaping” past the two stationary robots along the direction of the x -axis, the optimal trajectory is for the robot to trace the x -axis itself, with position $y = 0$ throughout the path, and pass through the other two robots. This can be generalized for any position of robots 0 and 1 in the plane: the trajectory of robot 2 should move along the equidistant path between the two stationary robots to achieve maximum information gain.

Fig. 3. This is the leap-frog path we use for simulations and experiments, which is based on an analysis of information gain for a team of three robots.

We introduce a new three robot leap-frog path in Fig. 3 to build off this result. To our knowledge, this is the only path for which the moving robot, at every time step, will trace the equidistant path between the stationary robots. The implementation of this path involves the trailing robot of an equilateral triangle configuration to pass through the stationary robots, establishing a new position and a new equilateral triangle configuration on the other side. The robots then switch roles and repeat the sequence.

4.2 Empirical Results

To test the generated path displayed in Fig. 3, we have performed a series of Monte Carlo simulations for three mobile robots performing cooperative localization. The robots (when moving) are instructed to drive at a constant 0.5 m/s with $\alpha = 0.006$, $\beta = 0.02$, and $\gamma = 0.003$ for the motion noise model in Eq. 1. Relative bearing measurements are obtained at 10 Hz and are assumed to have additive Gaussian noise with a standard deviation of 1 degree.

In Fig. 4, we compare the results of the Monte Carlo simulations for three different paths. Path (a) is a smoothed version of the leap-frog path designed in Sec. 4, path (b) is a trajectory obtained when the robots move in an equilateral triangle formation, and path (c) is the same as (b) but omits the measurements.

Fig. 4. Path (a) represents the leap-frog path presented in Sec. 4, path (b) represents the optimal formation for localization, and path (c) represents odometry only. The plot depicts the trace of the sample covariance matrix generated for a collection of 1000 Monte Carlo simulations.

Each path was simulated for 1000 different trials with randomly generated noise for measurements and motion. While the estimate of the state for each trial follows the intended path due to feedback control, the actual state for each trial is affected by the noise and drifts from the path. We measure the filter performance by observing the distribution of the robot state over all trials. A larger spread of data points implies worse tracking of the actual state. To quantify the performance, we compute the trace of the sample covariance matrix for the actual robot state computed over the 1000 simulated trials. Fig. 4 shows a graph of this metric for each of the path types. We note that our leap-frog path outperforms the optimal formation.

5 Experimental Evaluation

The motivating application for this work is GPS-denied autonomous coverage, for which accurate positioning is of critical importance. We apply a

smoothed version of the path depicted in Fig. 3 to a team of real robots performing coverage.

5.1 Experimental Setup

We use three robots for outdoor localization experiments, each of which is based on the Learning Applied to Ground Vehicles (LAGR) platform [10]. Each mobile robot has three on board computers, wheel encoders for odometry, and a set of four stereo cameras mounted above the chassis. We choose to treat the 4 stereo pairs as 8 individual bearing sensors in order to reduce the computational load. The filter is implemented according to Sec. 3 and is centralized (meaning that only one of the robots is running the Kalman filter at any given time). The robots measure bearing to each other by detecting large red spheres in the camera images with a circle Hough transform [11]. See Fig. 1 for a photograph of the robots in their experimental configurations.

5.2 Coverage Experiments

Fig. 5. An experiment on the football field at Gesling Stadium at Carnegie Mellon University. The three photos here are extracted from a video sequence used to record the ground truth position of the robots throughout the experiment.

The photos in Fig. 5 are from a video sequence recorded during one of our coverage experiments at Gesling stadium at Carnegie Mellon University. We are able to use this video sequence to post-process ground truth position data for each of the three robots. In order to do this, we compute a camera projection matrix based on known 3D points in the image (the markings on the football field). Then, after manually selecting a robot’s location in the image plane, we can infer its 3D position via its projection onto the plane.

The estimated path of the robots during the aforementioned coverage experiment is drawn in Fig. 6 (a). This estimated path follows an overall desired path that is composed of multiple smoothed versions of the leap-frog path displayed in Fig. 3 pasted together so as to sweep a region for coverage. The travel distance for each robot was approximately 140m during the experiment. The covered area was 20m x 30m. Ground truth points are shown in Fig. 6 for comparison and to help quantify the localization performance. Fig. 6 (b) shows the odometry-only estimate of the path. It is worth noting

Fig. 6. (a) shows the estimated trajectory of the robots during a coverage experiment along with ground truth points (yellow circles). The final ground truth position is displayed with yellow stars. (b) shows the same experiment with a filter that ignores the bearing measurements (dead reckoning only).

how erroneous this path is (most likely due to turning biases from unequally inflated tires) and how effective the filtering is in correcting the erroneous path to agree closely with the ground truth data.

The true final position of the three robots is also depicted in Fig. 6. The approximated error between the filtered estimate and the measured final ground truth pose is: 1.09 meters for robot 0 (the red trajectory in Fig. 6), 1.01 meters for robot 1 (the green trajectory in Fig. 6) and 1.15 meters for robot 2 (the blue trajectory in Fig. 6). We note that the accuracy of these ground truth measurements is subject to possible user error when manually selecting the image points that correspond to the robots in the video sequence.

The localization accuracy for this experiment is quite remarkable for this type of outdoor robot. The presence of wheel slippage coupled with a difficult terrain can cause severe drift in the odometry estimate over a path this long. Additionally, the measurements that we acquire with vision can be fairly noisy compared to more expensive laser range finders. But when an informative path, such as the one we present in Sec. 4, is used, the accuracy improves significantly, as shown in our experiment.

6 Conclusion

This work presents a leap-frog path designed to aid localization for a team of three robots. The path is designed such that the moving robot travels along a path that adds maximal information to the filter. The resulting path outperforms the optimal formation-based path. The experiment that we describe is, to our knowledge, one of the largest outdoor GPS-denied coverage results, successful in part because of precise localization.

Although we believe the absolute optimal path (in terms of localization accuracy) for a team of three robots would involve a leap-frog motion strategy, the path we introduce in this paper is most likely not optimal. Precisely defining the optimal path is still an open problem, which may require running an exhaustive simulation to optimize over all possible combinations of motion inputs: a task that would be computationally infeasible.

Also, this paper has focused on developing paths for a team of three robots. We believe that a three robot team is a good fit for applications that require accurate positioning, in part because three robots can provide proper triangulation. That said, it is always beneficial to add additional information to the Kalman filter, and a way to do this would be to add additional robots.

References

1. S. Roumeliotis and G. Bekey, "Collective Localization: A distributed Kalman filter approach to localization of groups of mobile robots," in *Proceedings of the International Conference on Robotics and Automation*, April 2000.
2. L. E. Navarro-Serment, C. J. J. Paredis, and P. K. Khosla, "A beacon system for the localization of distributed robotic teams," in *Proceedings of the International Conference on Field and Service Robotics*, August 1999.
3. R. Kurazume and S. Nagata, "Cooperative positioning with multiple robots," in *Proceedings of the International Conference on Robotics and Automation*, May 1994, pp. 1250–1257.
4. S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 781–795, October 2002.
5. I. M. Rekleitis, G. Dudek, and E. Milios, "Multi-robot cooperative localization: A study of trade-offs between efficiency and accuracy," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.
6. S. I. Roumeliotis and I. M. Rekleitis, "Propagation of uncertainty in cooperative multirobot localization: Analysis and experimental results," *Autonomous Robots*, vol. 17, no. 1, pp. 41–45, July 2004.
7. Y. S. Hidaka, A. I. Mourikis, and S. I. Roumeliotis, "Optimal formations for cooperative localization of mobile robots," in *Proceedings of the International Conference on Robotics and Automation*, 2005, pp. 4137–4142.
8. N. Trawny and T. Barfoot, "Optimized motion strategies for cooperative localization of mobile robots," in *Proceedings of the International Conference on Robotics and Automation*, April 2004.
9. S. Thrun, Y. Liu, D. Koller, A. Ng, Z. Ghahramani, and H. Durrant-Whyte, "Simultaneous localization and mapping with sparse extended information filters," *Intl. Journal of Robotics Research*, vol. 23, no. 7-8, pp. 693–716, 2004.
10. L. Jackel, E. Krotkov, M. Perschbacher, J. Pippine, and C. Sullivan, "The darpa lagr program: Goals, challenges, methodology, and phase i results," *Journal of Field Robotics*, vol. 23, pp. 945–973, 2006.
11. E. R. Davies, "A modified hough scheme for general circle location," *Pattern Recognition Letters*, vol. 7, pp. 37–43, January 1988.