Using constraint propagation for cooperative UAV localization from vision and ranging

Ide-Flore Kenmogne, Vincent Drevelle, and Eric Marchand

Univ Rennes, Inria, CNRS, IRISA
`ide-flore.kenmogne@inria.fr`,
{`vincent.drevelle,eric.marchand`}`@irisa.fr`

Abstract. This paper addresses the problem of cooperative localization in a group of unmanned aerial vehicles (UAV) in a bounded error context. The UAVs are equipped with cameras to tracks landmarks, and a communication and ranging system to cooperate with their neighbours. Measurements are represented by intervals, and constraints are expressed on the robots poses (positions and orientations). Each robot first computes a pose domain using only its sensors measurements, by using set inversion via interval analysis [1]. Then, through position boxes exchange, positions are cooperatively refined by constraint propagation in the group. Results are presented with real robot data, and show position accuracy improvement thanks to cooperation.

Keywords: Intervals · Cooperative localization · Constraints propagation

1   Introduction

In this paper, we consider the problem of cooperative localization [2] in a group of N unmanned aerial vehicles (UAV). The robots are equipped with cameras, able to see landmarks of known positions. A communication and ranging system provides to each robot R_k a means of exchanging data and measuring distances with its neighbours and a base station B (Fig. 1). The goal for each robot is to compute a domain for its pose (position and orientation), assuming bounded error measurements.

The paper is organized as follows: we first present how each robot is able to independently compute a domain for its pose using constraints from camera


Fig. 1: Cooperative localization with camera and range measurements

measurements and distance to the base station. Then, in a second part, a co-operative localization method is introduced, in which neighbours positions and distances are introduced as additional constraints to tighten the pose domain of each robot, thanks to data exchange. We finally provide experimental results obtained with quadcopter mini-drones.

2 Vision-based pose computation

This section addresses pose estimation [3] from camera measurements. To compute the pose $\mathbf{r} = (x, y, z, \phi, \theta, \psi)$ of a UAV amount in estimating the transformation ${}^r\mathbf{T}_w$ between the world reference frame and a reference frame attached to the robot. We assume that the rigid transformation ${}^c\mathbf{T}_r$ between the camera and the robot frame is known from calibration [4], and the camera is calibrated.

For a known landmark of 3D coordinates ${}^w\mathbf{X}$ in the world reference frame, the normalized coordinates $\mathbf{x} = ({}^cx, {}^cy)$ of its projection in the camera frame are given by the pinhole model [5]:

$$\mathbf{x} = \mathbf{\Pi}\,{}^c\mathbf{T}_r\,{}^r\mathbf{T}_w(\mathbf{r})\,{}^w\mathbf{X} \tag{1}$$

where $\mathbf{\Pi}$ is the perspective projection operator.

For each visible landmark ${}^w\mathbf{X}_i$ $(i \in 1..m)$, we can derive the following constraints:

$$C_i : \begin{cases} ({}^cX_i, {}^cY_i, {}^cZ_i) = {}^c\mathbf{T}_r\,{}^r\mathbf{T}_w(\mathbf{r})\,{}^w\mathbf{X}_i \\ {}^cx_i = \frac{{}^cX_i}{{}^cZ_i}, \\ {}^cy_i = \frac{{}^cY_i}{{}^cZ_i}, \\ {}^cx_i \in [{}^cx_i], \ {}^cy_i \in [{}^cy_i] \qquad \text{bounded error camera measurement} \\ {}^cZ_i > 0 \qquad\qquad\quad \text{front looking camera} \end{cases} \tag{2}$$

We then define the image-based pose estimation problem as a constraint satisfaction problem (CSP) as:

$$\mathcal{H}_{\mathrm{img}} : \begin{pmatrix} \mathbf{r} \in [\mathbf{r}], \\ \{C_i, \ i \in 1...m\} \end{pmatrix} \tag{3}$$

This CSP is solved with Contractor Programming [6] and Set Inversion via Interval Analysis (SIVIA) [7]. The altitude $[z]$, pitch $[\theta]$ and roll $[\phi]$ components of the initial domain $[\mathbf{r}]$ are set thanks to onboard sensors (altimeter and inertial measurement unit). An outer approximation of the feasible domain for the pose $\mathbf{r}$ is obtained in the form of a subpaving. An example of solution set for one robot is shown in Fig. 2a (see [8]).

Assuming a bounded error measurement $[d]$ of the distance d between the robot and the base station B is available, it can be used as an additional constraint on the robot position $\mathbf{p} = (x, y, z)$, to get a tighter pose estimate:

$$d = \|\mathbf{p} - \mathbf{b}\|_2, \quad d \in [d] \tag{4}$$

where $\mathbf{b}$ the known position of the base station. We obtain a tighter domain as presented in Fig. 2b.

a. using camera only   b. using camera and range to base

Fig. 2: Pose domain computation for a single drone (horizontal projection)

Dealing with camera tracking failures The subpaving computed with SIVIA can be empty. This situation corresponds to inconsistencies in the measurements, due to failure of the visual tracking of one or several landmarks (see [8] for more details). In this case, our approach put aside the image measurements and use only the distance to the base station B to compute $\mathbb{S}_{\mathbf{r}_k}^+$. The solution-set is in this case a ring centered on B.

3   Using range measurements for cooperative localization

Assuming each UAV is equipped with a communication system with ranging capabilities, each robot therefore measures range to its neighbours and cooperate with them by exchanging its position.

Sharing positions Lets consider one robot R_k ($k \in 1..N$) from the group. $\mathcal{N}(k)$ denotes the neighbours of R_k, i.e. the robots within communication range.

At each time step, R_k first computes an outer subpaving $\mathbb{S}_{\mathbf{r}_k}^+$, that contains all the feasible poses, considering the camera and base distance bounded-error measurements (as presented in Section 2).

Once the pose domain $\mathbb{S}_{\mathbf{r}_k}^+$ is computed, the robot computes the bounding box of its position domain $[\mathbf{p}_k] = \square\,\mathrm{proj}_{\mathbf{p}}\,\mathbb{S}_{\mathbf{r}_k}^+$, where $\square$ is the bounding box operator, and $\mathrm{proj}_{\mathbf{p}}$ is the projection on the position space. $[\mathbf{p}_k]$ is transmitted to all neighboring robots R_j, $j \in \mathcal{N}(k)$, and the distances $d_{k,j}$ between R_k and R_j are simultaneously measured.

Pose contraction At reception of information (position boxes $[\mathbf{p}_j]$ and bounded-error distances measurements $[d_{k,j}]$) from neighboring robots, each robot R_k tries to refine its actual pose domain, by propagating the new distance constraints between R_k and each of its neighbours. Recalling that $\mathbf{p}_k = (x_k, y_k, z_k)$ is the position of R_k, each robot R_k contracts a local CSP $\mathcal{H}_k$ defined as follows:

$$\mathcal{H}_k : \begin{pmatrix} \mathbf{p}_k \in \mathrm{proj}_{\mathbf{p}}(\mathbb{S}_{\mathbf{r}_k}^+), \\ \mathbf{p}_j \in [\mathbf{p}_j],\ j \in \mathcal{N}(k) \\ d_{k,j} \in [d_{k,j}],\ j \in \mathcal{N}(k) \\ d_{k,j} = \|\mathbf{p}_k - \mathbf{p}_j\|_2,\ j \in \mathcal{N}(k) \end{pmatrix} \tag{5}$$

We use interval constraint propagation to solve $\mathcal{H}_k$, in order to reduce the pose domain $\mathbb{S}^+_{\mathbf{r}_k}$.

Fixed-point The constraint network formed by the group of robots contains cycles spanning several robots. The contraction of the local CSPs $\mathcal{H}_k$ has to be propagated again through the network to improve the pose domain reduction of each robot. If after solving $\mathcal{H}_k$, the robot position bounding box $[\mathbf{p}_k]$ is reduced, then the robot R_k retransmits its updated $[\mathbf{p}_k]$ to its neighbourhood. This process is iterated until a fixed-point is reached (no more significant improvement of the robots positions bounding boxes).

4   Experimental results

The proposed method has been tested with data acquired on Parrot AR-Drone2 UAV, with 5 landmarks made with AprilTag markers (Fig. 3). Image measurement error bounds are set to ± 0.5 px and range measurement error is assumed to be within ± 5 cm.

Subpavings obtained with 4 robots in cooperative localization are presented in Fig. 4. Figure 4b clearly shows how cooperative localization reduces the feasible pose domain when some robots cannot clearly see the landmarks, by propagating position information of the neighbours.

Table 1 shows how making more robots cooperate in the fleet improves localization, first by reducing the width of the computed position domain (Table 1a), and also by improving the precision when using the center of the domain as a point estimate (Table 1b). With 4 UAVs, the average horizontal position error is less than 5 cm for all the drones.

Table 1: Horizontal position results, for 1 to 4 robots in the group

	R_1	R_2	R_3	R_4
1 UAV	1.29			
2 UAVs	0.86	0.56		
3 UAVs	0.80	0.52	2.98	
4 UAVs	0.33	0.37	0.78	0.61

Horizontal position domain width (m)

	R_1	R_2	R_3	R_4
1 UAV	0.226			
2 UAVs	0.159	0.088		
3 UAVs	0.141	0.071	1.220	
4 UAVs	0.040	0.046	0.248	0.192

Average horizontal position error (m)

Fig. 3: Onboard cameras views at $t = 8$ s. Landmarks are boxes with a printed pattern.

a. $t = 0$ s, the 4 robots cameras track landmarks

b. $t = 6.5$ s, the cameras of R_1 and R_2 are unable to track landmarks

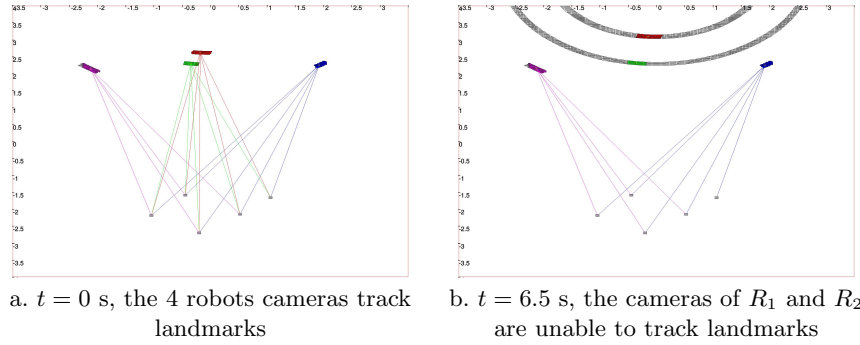Fig. 4: Position domains of 4 robots. *Black outline: subpavings before communication. Colored domains: subpavings after cooperative localization*

5   Conclusion

In this paper, we proposed a method to solve cooperative localization in a group of UAVs. Computations rely on interval constraint propagation, assuming bounded error image and distance measurements. Each UAV first independently compute a pose domain from its camera measurements, and then exchange position information with its neighbours to further reduce the position domains of the robots in the group. The method has been applied to real data, and enables to improve the positioning precision of the AUVs thanks to position information propagation. The experiments show that increasing the number of robots in the group provides additional constraints on position, and yields smaller uncertainty of the computed robots poses.

References

1. R.E. Moore. *Interval analysis*. Prentice Hall, 1966.
2. S. Roumeliotis, G. Bekey, et al. Distributed multirobot localization. *IEEE Trans. on Robotics and Automation*, 18(5):781–795, 2002.
3. E. Marchand, H. Uchiyama, and F. Spindler. Pose estimation for augmented reality: a hands-on survey. *IEEE Trans. on Visualization and Computer Graphics*, 22(12):2633–2651, 2016.
4. R. Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal on Robotics and Automation*, 3(4):323–344, August 1987.
5. Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2 edition, 2004.
6. G. Chabert and L. Jaulin. Contractor programming. *Artificial Intelligence*, 173(11):1079 – 1100, 2009.
7. J. Jaulin and E. Walter. Set inversion via interval analysis for nonlinear bounded-error estimation. *Automatica*, 29(4):1053 – 1064, 1993.
8. I. F. Kenmogne, V. Drevelle, and E. Marchand. Image-based uav localization using interval methods. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 5285–5291, Sept 2017.