Evaluating the Effect of Robot Group Size on Relative Localisation Precision

Frank E. Schneider and Dennis Wildermuth

Fraunhofer Institute for Communication,
Information Processing and Ergonomics (FKIE), Wachtberg, Germany
{frank.schneider,dennis.wildermuth}@fkie.fraunhofer.de

Abstract. Looking on co-operative position estimation in multi-robot systems, the question to what extend the number of robots has an influence on the quality of the resulting localisation is an important and interesting issue. This paper addresses this relation regarding a pure relative localisation approach based only on mutual observations between the robots. The intuitive expectation that more robots should improve the position estimation is motivated and the design of the experiments with special respect to possibly distorting parameters is discussed and reasoned in detail. An in-depth analysis of the collected data explains the only partial conformance of the experimental results with the expected outcome.

1 Introduction

It is an obvious fact that most of the actions performed by mobile robots require some type of localisation. While localisation itself is a field of ongoing research, it is also a vital component for the co-ordination of navigation and movement. This holds especially when dealing with multi-robot systems (MRS).

The problem of localization can be addressed by different approaches. Local localisation evaluates the robot's position and orientation through integration of information provided by miscellaneous encoders and inertial sensors. All these sensors are mounted on the robot itself, and no external information is used [1, 2]. However, due to inherent uncertainty and unbounded error growth this method is usually not simply extendable to MRS.

Global localisation is normally based on some kind of map and uses sensor information to localise the robot with respect to these maps. In recent years, the problem of global localisation as well as typical approaches like Simultaneous Localisation and Mapping (SLAM) was extended to multi-robot localisation [3, 4, 5]. In the related approach of absolute localisation the vehicle determines its position directly through an exterior reference system, usually a satellite-based positioning system, navigation beacons, or passive landmarks. Since at least satellite-based systems do not have the accuracy needed for most robotic tasks, absolute localisation is often combined with other localisation techniques [6].

Relative localisation uses sensor observations to localise the robot with respect to its environment – including other robots – without having an environment model. Most of the authors working on this topic use the mutual observations of the robots as

R. Groß et al. (Eds.): TAROS 2011, LNAI 6856, pp. 149–160, 2011.

means for improving global positioning [7, 3, 4]. A different approach – which is also applied in this paper – is to change the aim of relative localisation and to maintain only a relative positioning between the robots. Hence, the resulting reference co-ordinate system is not global in the sense that it has a fixed reference to world co-ordinates. It is just shared among the members of the MRS and can diverge from world co-ordinates over time [8, 9].

For all multi-robot localisation approaches an interesting topic is the evaluation of the results in terms of precision, stability, scalability, or environment dependency. Thereby, the question whether the number of robots sharing the common co-ordinate system has an influence on the precision of the resulting localisation is one core issue. So far, this question has been studied almost always in terms of global positioning [10, 11], comparing the robots' position estimates with their corresponding world co-ordinates. In [11], for example, an analytical upper bound for the global positioning uncertainty in Cooperative Simultaneous Localization and Mapping (C-SLAM) is obtained. In this paper we are going to examine the relationship between robot group size and localization accuracy with regard to "pure" relative localization.

The remainder of the text consists of four parts. The next chapter shortly introduces the employed relative localisation method. Afterwards, we give a detailed description of the experimental setup. Special emphasis is put on the goal of gathering a good data basis for the evaluation. Design decisions concerning possibly distorting parameters and preconditions are explained and justified. The final chapter presents a detailed analysis of the collected data and tries to explain the only partial conformance of the experimental results with the expected outcome.

2 The Relative Localisation Approach

As introduced in [9] an Extended Kalman Filter (EKF) is used to integrate measurements into a global state estimation. The state vector $x(k) \in \mathbb{R}^{3n}$ describing all important information in a robot group of n robots at time step k is denoted

$$x(k) = [p_1(k) \quad \dots \quad p_n(k)]^T \tag{1}$$

where $p_i(k) = [p_{i,x}(k) \quad p_{i,y}(k) \quad p_{i,\varphi}(k)]$ means position and orientation of the i^{th} robot. For the initial error covariance P_0 we took $\sigma_t = 3cm$ as translational and $\sigma_o = 3°$ as rotational error.

In order to explain the EKF update step one has to consider the following situation. Each time one of the robots gathers a relative measurement of one of the others this information is deployed into the EKF in order to update the overall system state. For explanation of update step k, without loss of generality we will consider the case that the i^{th} robot has got a new measurement $z = [d \ \ \alpha]^T$ of the robot with index 1, thereby meaning d the distance between them and α the direction of a vector from the observer to robot no. 1. The information gained from the odometry sensors of the two involved robots is used for the prediction step of the EKF. Let

$$u_i(k+1) = \left[\Delta p_{i,x}(k+1) \quad \Delta p_{i,y}(k+1) \quad \Delta p_{i,\vartheta}(k+1)\right] \tag{2}$$

the i^{th} robot movement from time step k to $k+1$. Then the state prediction writes to

$$x^{-}\left(k+1\right)=\left[p_{1}\left(k\right)+u_{1}\left(k+1\right) \quad \ldots \quad p_{i}\left(k\right)+u_{i}\left(k+1\right) \quad \ldots\right]^{T}. \tag{3}$$

Accordingly, the projection of the error covariance is simply defined as

$$P^{-}\left(k+1\right)=P\left(k\right)+Q\left(k+1\right). \tag{4}$$

The process noise covariance $Q\ (k + 1)$ in our approach is not constant over time but depends on the distances the two involved robots travelled since their state had been updated last time. The further they moved the larger is the uncertainty on this movement because of possible odometry errors. This consideration led to the following definition:

$$Q\left(k+1\right)=\begin{pmatrix} Q_1 & & & \\ & \ddots & 0 & \\ & 0 & Q_i & \\ & & & \ddots \end{pmatrix} \tag{5}$$

a diagonal matrix with

$$Q_i(k+1)=\begin{pmatrix} \left(s_t\Delta p_{i,x}\left(k+1\right)+\hat{q}_t\right)^2 & 0 & 0 \\ 0 & \left(s_t\Delta p_{i,y}\left(k+1\right)+\hat{q}_t\right)^2 & 0 \\ 0 & 0 & \left(s_o\Delta p_{i,\vartheta}\left(k+1\right)+\hat{q}_o\right)^2 \end{pmatrix}. \tag{6}$$

and remaining diagonal elements $\hat{Q}=\begin{pmatrix} \hat{q}_t^{\,2} & 0 & 0 \\ 0 & \hat{q}_t^{\,2} & 0 \\ 0 & 0 & \hat{q}_o^{\,2} \end{pmatrix}\in\mathbb{R}^{3\times3}$, $\hat{q}_t$ and $\hat{q}_t$ very small positive constants. Thereby, s_t represents the mean error for translational movement of the robots and s_o the same for rotations. As long as all robots show the same error characteristic those values remain constant, otherwise they should depend on the actual robot they belong to. We assume constant values of $s_t = 5cm$ and $s_o = 5°$. The $\hat{q}_t$ and $\hat{q}_o$ represent the inherent uncertainty of the position information for every robot, which grows even if the odometry reports no movement at all.

After this, the EKF prediction step is finished. The new measurement $z\ (k + 1) = [d \ \ \alpha]^T$ of robot no. 1 obtained by the i^{th} observing robot is now used for the correction step of the EKF:

$$x\left(k+1\right)=x^{-}\left(k+1\right)+K\left(k+1\right)\left[z\left(k+1\right)-h\left(x^{-}\left(k+1\right),0\right)\right] \tag{7}$$

where the function $h(x, v)$ relates a state vector x, defined as in (1), to a measurement $z = [d \;\; \alpha]^T$, given the current measurement noise v. For $h(x, v)$ a standard mixed coordinate EKF approach [13] is used, which in this case writes to

$$h(x,v) = \begin{pmatrix} h_1(x,v) \\ h_2(x,v) \end{pmatrix} = \begin{pmatrix} \sqrt{\left(p_{1,x} - p_{i,x}\right)^2 + \left(p_{1,y} - p_{i,y}\right)^2} + v_d \\ \arctan\left(\dfrac{p_{1,y} - p_{i,y}}{p_{1,x} - p_{i,x}}\right) - p_{i,\vartheta} + v_\vartheta \end{pmatrix}. \tag{8}$$

If we redefine (3) as $\left[\tilde{p}_{1,x} \quad \tilde{p}_{1,y} \quad \tilde{p}_{1,\vartheta} \quad \cdots \quad \tilde{p}_{i,x} \quad \tilde{p}_{i,y} \quad \tilde{p}_{i,\vartheta} \quad \ldots\right]^T = x^-(k+1)$ and define

$$P_1 = \sqrt{\left(\tilde{p}_{1,x} - \tilde{p}_{i,x}\right)^2 + \left(\tilde{p}_{1,y} - \tilde{p}_{i,y}\right)^2} = \left\| \begin{pmatrix} \tilde{p}_{1,x} \\ \tilde{p}_{1,y} \end{pmatrix} - \begin{pmatrix} \tilde{p}_{i,x} \\ \tilde{p}_{i,y} \end{pmatrix} \right\| \tag{9}$$

then for the Jacobian matrix $H(k + 1)$ for the remaining EKF step this leads to

$$H(k+1) = \begin{pmatrix} \dfrac{\tilde{p}_{1,x} - \tilde{p}_{i,x}}{P_1} & \dfrac{\tilde{p}_{1,y} - \tilde{p}_{i,y}}{P_1} & 0 & \cdots & 0 & \dfrac{\tilde{p}_{i,x} - \tilde{p}_{1,x}}{P_1} & \dfrac{\tilde{p}_{i,y} - \tilde{p}_{1,y}}{P_1} & 0 & 0 & \cdots \\ \dfrac{\tilde{p}_{i,y} - \tilde{p}_{1,y}}{P_1^2} & \dfrac{\tilde{p}_{1,x} - \tilde{p}_{i,x}}{P_1^2} & 0 & & 0 & \dfrac{\tilde{p}_{1,y} - \tilde{p}_{i,y}}{P_1^2} & \dfrac{\tilde{p}_{i,x} - \tilde{p}_{1,x}}{P_1^2} & -1 & 0 & \end{pmatrix}. \tag{10}$$

Note that $H_{1,3}(k + 1) = 0$ and $H_{2,3}(k + 1) = 0$, which means that no update on the orientation of the observed robot takes place during the filter step. This is, of course, due to the fact that in our setup the orientation information cannot be measured by the observer. But as long as all participating robots at least occasionally generate measurements on their own, orientation information for all robots is updated.

The Kalman gain $K(k + 1)$ is computed in the usual way with the measurement noise covariance matrix R considered as constant. Due to the characteristics of the laser device and the tracking algorithm for the experiments we used the values $r_t = 0.5cm$ and $r_o = 0.5°$. The filter step ends with the update of the error covariance matrix $P(k + 1)$. As a result $x(k + 1)$, as computed in (7) contains updates for the position of the first robot, which has been observed, and updates for position and orientation of the i^{th} robot, which in this case played the role of the observer. In other words, starting with initial position estimations the EKF maintains the relative positions for all robots during the ongoing run.

3 Description of the Experiments

Aim of this work is to examine the relationship between robot group size and localization accuracy. On one hand, since each robot measures distance and direction to every other robot in the group, the amount of mutual measurements rapidly increases. On the other hand, apart from the errors caused by the odometry, there are also a possibly growing number of measurement faults due to, for example, temporary occlusions while the robots work together.

The aggregated odometry error for a robot group grows linearly with the number of group members. The frequency of mutual occlusions depends on the kind of task the robots have to fulfil, but normally such occlusions have only a very limited duration, happening for example while two robots are passing each other. Therefore, one can expect that the increasing number of mutual position measurements outweighs the number of possible error sources. Thus, we expected that the overall precision of the relative localisation should increase with the size of the robot group.

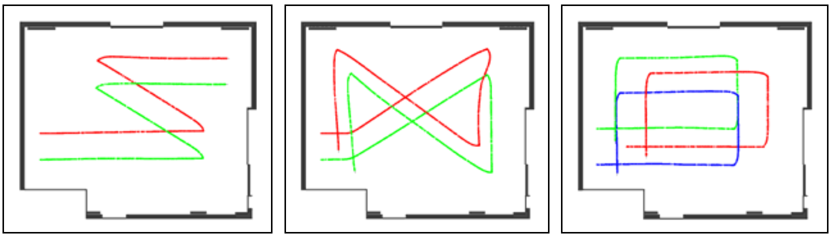

Fig. 1. The three different shapes on which the robots had to move in the large 15 x 18 m experimental hall

3.1 Design Decisions and Necessary Preconditions

The experiments were fully conducted in simulation. Apart from the obvious reason that a simulation with well-known parameters yields reproducible and comparable results, some other important causes lead to this decision. First, the simulation provides an exact ground-truth in form of absolute position information for each robot, far better than any means of positioning for "real" robots.

Additionally, it was planned to compare group sizes of at least five robots, which means a great challenge with experimental robot systems and a lab environment of limited size. And the goal was not to record some impressive demonstration runs but to obtain data from a large number of runs for each scenario in order to establish a meaningful basis for further evaluation.

With different group sizes from two to five robots, at least three different driving scenarios and a minimum of 20 runs for each combination, this already results in a total of 240 experiments, not considering any possible technical troubles or problems in the design of the experiments. Realistically, this easily leads to hundreds of additional trials, consuming weeks of laboratory time. In simulation the original design regarding group size and driving scenarios was retained. Thus, in principle all results of these experiments could be validated using the real platforms and a large experimental hall. In simulation, now each of the scenarios described in the following section was conducted 60 times.

The three different paths on which the robots moved are pictured in figure 1. The left figure shows a group of two robots (red and green line) moving along a Z-like shape. The middle one presents the second scenario, a shape like a horizontal eight with corners. In the right figure a formation of three robots (red, green and blue line) drives on a rectangular path. These scenarios will be mostly referred to as “Z”, “eight”, and “rectangle” in this paper.

The three paths look similar in terms of driving distance, covered area, or overall turning angle. But there are also relevant characteristics in which the scenarios differ, mainly the kind of turns the vehicles have to follow. For many robot platforms – and for ours as well – any rotation induces a potentially large odometry error. Another frequent observation is the influence of the turning direction on the aggregated odometry error. During longer experiments it sometimes happens that the rotational error increases and decreases in turn. This seems to be because odometry errors caused by clockwise and counter-clockwise turns sometimes eliminate each other.

The multi-robot simulator was configured to use a realistic odometry error model, reflecting the foregoing considerations. Each robot had its own error characteristic with a normally distributed translational error ($\sigma = 3cm$ per 1m translation), and an orientation error spread normally around a small systematic fault ($\mu \in [1.5, 2.5]$, $\sigma = 0.3°$ per 90° turn). Especially the systematic fault causes the typical displaced odometry data. Figure 2 presents some characteristic example runs. In these examples always two robots have been recorded in order to simplify the illustrations.

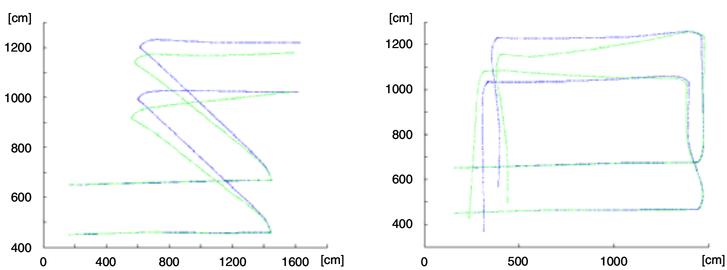

Fig. 2. Comparison of true positions (*blue path*) with information from odometry (*green path*)

The blue lines give the robots’ exact positions. The two figures correspond to the “Z” and the “rectangle” scenario described above. In contrast, the green lines show the faulty positions of the robots as they are recorded by the simulated odometry sensors. The left figure gives a good example for the before mentioned decreasing odometry error. The lower of the two robots nearly ends at its exact position, even though with faulty orientation information, and the odometry position of the upper robot also approaches to the real one. This is a quite recurrent behaviour in this first scenario, possibly because in this scenario the robots have to turn exactly the same angle of about 110° in both directions. In the eight and rectangle scenario the overall driving distance for each robot is between 30m and 40m, and the accumulated odometry error at the end is often more then 1m, sometimes even more than 2m. Like the right figure demonstrates, this deviance often points into completely opposite directions, due to the independent error model for each robot. As a result, looking only at the pure odometry data the distance (and orientation) error between any pair of robots easily reaches the magnitude of several metres during and especially at the end of a run. Deviances of this dimension mainly happen in the rectangle scenario. In this scenario the robots have to turn only counter-clockwise, three times about 90°.

Altogether, the average odometry error is smallest in the Z-scenario, largest in the rectangle scenario and somewhere in between for the eight. It is important to mention one more time that, although the presented data is simulation generated, it very well matches our experiences when collecting odometry data from the real robot systems.

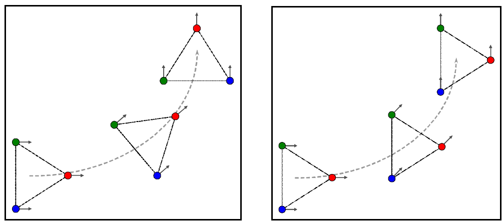

Fig. 3. Different modes of turning a formation: "normal" (*left*) – whole formation turns; "fixed" (*right*) – only the robots turn around

Apart from the impacts of the different paths it is worth discussing the mode in which the robot group had to move during the experiments. Of course, since all robots had to follow the same path, special co-ordination is necessary. Different approaches have been considered. One possibility, for example, is to assign exactly the same path to each robot and then let them start consecutively from the same starting position at a fixed time interval. The disadvantage is the nearly permanent mutual occlusion of robots which are moving along the same line of sight. Apart from the turning points, only direct neighbours can see each other. For relative localisation this is not the best choice for evaluating the overall precision of the approach.

Another idea was to simply give the robots the goal and turning points, leaving the remaining coordination to the collision avoidance. But the lack of a pre-defined and reproducible behaviour for each robot is a problem for the goal of precision measurement. Thus, it was decided to use fixed formations, which the robots had to maintain while following the pre-defined paths. "Fixed" in this context characterises a special mode of movement in formation. Whenever the formation reaches a turning point, instead of completely turning the shape of the formation only the robots themselves turn around (see figure 3). Different approaches to the field of formation navigation are addressed in detail, for example, in [12].

The actual goal formations for the varying group sizes can be seen in figure 4. They are straightforward from the 2-robot line over a triangle and rectangle shape leading to a pentagon for five robots. The minimal distances, as found between neighbouring robots, are 2m for the line shape, 2.25m for the rectangles and about 2.7m for the two other shapes. The largest distance of about 4.25m can be found in the pentagon formation.

As mentioned above the formation algorithm was configured to generate fixed formation shapes in order to obtain regular and straight paths for all robots. The simulated environment was a part of our indoor robotic lab, namely the large experimental hall of about 15 x 18 metres in size. This gives us the possibility to repeat at least some of the trials in a physical environment.

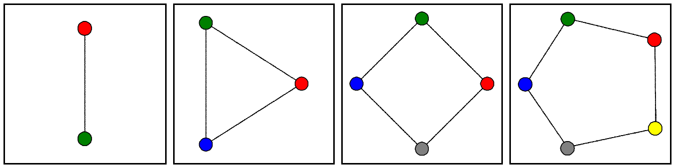

Fig. 4. Formation shapes used for the different robot group sizes

3.2 Collecting Data

Our robots are equipped with standard SICK laser range finders. According to the specification [14], the LMS200 devices deliver distance information with 1cm resolution and an error of less than 1.5cm. In our configuration the angular information has a resolution of 1° and a typical error of about 0.4°. The simulated laser scanners compute their sensor data with the same error characteristic.

In order to derive position information for the other robots from the raw laser readings, a straightforward stateless geometric algorithm was used. We achieve a mean error of about 0.5cm for the relative distance between two robots and less than 0.5° for the measured direction from one robot to another. This is a similar range as it can be achieved for sophisticated probabilistic trackers like Probabilistic Data Association Filtering (PDAF) or Probabilistic Multi Hypothesis Tracking (PMHT), both customized especially for extended targets [15]. Consequently, we omitted the implementation of any other tracking algorithm and continued to use the simple geometric method.

Figure 5 presents exemplary results for the mutual observation and measuring process in terms of quantity of position measurements. We counted the average number of relative position measurements which each robot gathered from all other group members during one experimental run. The y-axis gives this number and the x-axis splits the results for the three different scenarios and the different group sizes. One can expect the double number of measurements with three robots, triple with four and four times more measurements in the five robot group. Actually, in the "Z" scenario for two robots each one measures its counterpart's position at an average of 474 times during one experimental run. In a group of three robots each of them generates an average of 875 measurements, 1340 position values for four robots, and in the five robot group the mutual observation process delivers 1747 measurements per robots. Like figure 5 illustrates the situation is similar for the other scenarios.

In summary, the design of the experiments should be adequate to analyse the influence of the robot group size on the resulting precision of the relative localisation. Due to the way the robots moved, a predictable and comparable behaviour for all scenarios is achieved. The number and duration of mutual occlusions is limited and reproducible. There are linearly growing measurements per robot with growing group size. Although simulation based, the position measurements and the robots' odometry sensors have realistic accurateness and error characteristics. The aggregated odometry error for the robot group differs between the scenarios, being larger for the "rectangle" and smaller for the "Z", but in all cases it grows only linearly with the number of group members. Consequently, we expected the increasing number of mutual position measurements to outweigh the growing odometry error and, thus, the overall precision of the relative localisation to increase with the size of the group.

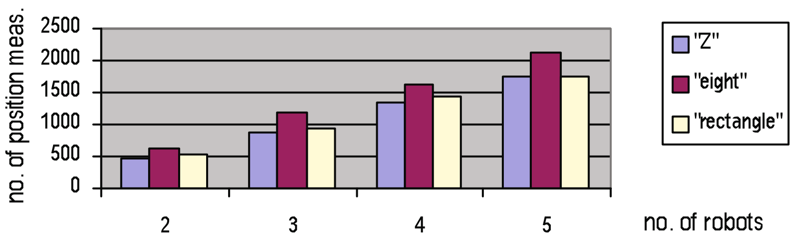

Fig. 5. Amount of position measurements (*y-axis*) for the different number of robots (*x-axis*)

4 Evaluation of the Results

As mentioned earlier a total of 60 runs have been conducted for each group size and each different driving scenario, leading to 720 simulation-based datasets in total. To evaluate the resulting relative localisation we defined a Mean Localisation Error (MLE). For each single localisation step the estimated robot positions were compared to the exact positions as delivered by the simulator. Since we are looking at relative localisation we could not simply compare absolute positions. Instead, for each pair of robots the difference between estimated and true distance was summed up. This deviation was then weighted by the number of robots to get a measure independent of the group size. Finally, an average over the entire dataset was computed, resulting in the MLE per run. One drawback of this approach is that orientation errors are not addressed. Hence, we tried some different approaches, including complex and mathematically demanding methods derived from metrics for robot groups as described in [16] and [17]. But since these approaches produced similar results we kept the computationally simple MLE as described above.

Looking at table 1 one can find the average MLE and the corresponding standard deviation calculated from the 60 runs for each scenario and group size. First of all, it is obvious that the resulting relative localization among the robots is very accurate. The mean error per time step for one robot consistently lies below 1.8cm. Considering only the tracking process with its ±0.5cm distance and ±0.5° angular error, then for typical robot distances between 2m and 4m the angular error delivers displacements

Table 1. Average MLE (in cm) and corresponding standard deviation per 60 runs for different group sizes and the three driving scenarios

"Z"	*no. of robots*	2	3	4	5
	mean MLE [cm]	1,542	1,552	1,364	1,441
	σ MLE	0,274	0,165	0,106	0,092
"eight"	*no. of robots*	2	3	4	5
	mean MLE [cm]	1,435	1,496	1,522	1,457
	σ MLE	0,078	0,160	0,132	0,085
"rectangle"	*no. of robots*	2	3	4	5
	mean MLE [cm]	1,695	1,781	1,501	1,541
	σ MLE	0,294	0,194	0,137	0,095

from 1.74cm in the near case and 3.48cm for the larger distances. Adding the perpendicular distance error this already leads to errors ranging from ±1.81cm up to ±3.52cm, thereby not even taking into account the errors coming from the odometry.

Unfortunately, the results from table 1 do not fully reflect the key expectation that the localisation precision increases with the number of robots. For the "eight" scenario there is no improvement at all, for the other scenarios there is at most a tendency of less than 0.2cm between the average MLE with two robots and with five robots. This might be due to the fact that even with only two robots the accuracy of the laser based mutual tracking outweighs the errors of the odometry sensors. At least the standard deviation of the MLE shows some improvement with larger group sizes. Again, for each combination of group size and scenario the table contains average values per 60 recorded example runs. Except for the "eight" scenario with 2 robots, the standard deviation shows less variation and therefore an improved stability in the relative localisation.

For a detailed analysis of the results we had a closer look at the changes of the Mean Localisation Error (MLE) over runtime. Figure 6 gives two typical examples, plotting the MLE in cm on the y-axis and the number of localisation steps from beginning to the end of the example run on the x-axis. Based on this data it was possible to check whether the error remains constant over time or whether there are peaks corresponding, for example, to turns or temporary occlusions of the robots.

The left part of figure 6 presents one dataset for the "eight" scenario and the two-robot formation, the right chart shows the same for five robots. One can see that, actually, the MLE fluctuates much less if more robots take part in the relative localisation. During a further inspection of the error peaks no general correlation to incidents like turnings or occlusions could be indentified. The variations seem to be arbitrarily spread over the whole runtime.

The only point which occasionally recurs is an increase of the MLE near the end of a run. In the left example in figure 6 the impression of a slightly higher error in the second half starting from time step 700 can be proven by numerical evaluation as well. The MLE is more than 0.5cm higher in the second half of the run. Figure 7 presents more examples for this effect. The left part is taken from a "rectangle" scenario with three robots, in the right part four robots move along the "Z" path. The "rectangle" example is comparable to the before described one. From time step 1700 the standard deviation of the MLE starts growing and from step 2300 the error increases by more then 1cm. The other example gives an even larger growth of the MLE. But in this case this does not go along with larger fluctuations of the localisation error, instead a constant increase of the error can be found. A possible explanation of this occasional effect might be an above-average odometry error aggregation which leads to very bad state predictions in the localisation EKF. But this is subject to ongoing research.

Another important fact currently under investigation is that only very few datasets result in a MLE below 1cm. Even during periods of precise localisation with very small deviation – like in the right example of figure 6 or in the first half of the "Z" run in figure 7 – the error only seldom drops below 1cm. We suspect a systematic fault somewhere in the design of the experiments or the analysis algorithms. But, as the section on the experiments should document, the whole preparation was done with great care. Thus, no such error could be found so far.

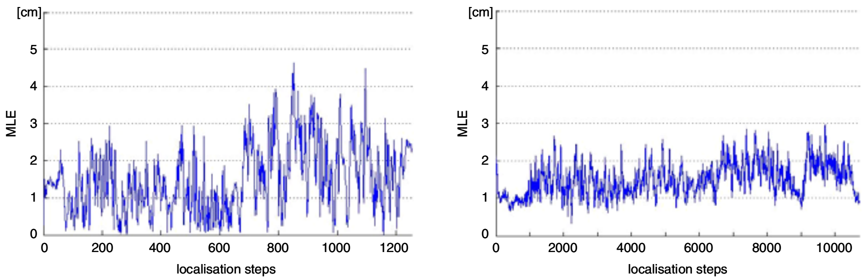

Fig. 6. Example plot of the MLE in cm (*y-axis*) for the "eight" scenario with two (*left*) and five robots (*right*); the *x-axis* runs through the localisation steps resp. the duration of the run. The error for two robots shows more variations, fitting well to the larger standard deviation.

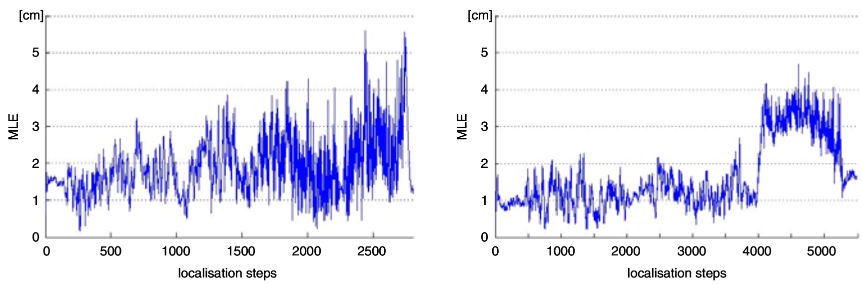

Fig. 7. MLE for the "rectangle" with three (*left*) and the "Z" with four robots (*right*), axes as above, showing larger variations or sections with constantly large errors near the end of a run.

5 Summary

This paper addresses the effect of robot group size on relative localisation precision regarding a pure relative localisation approach based only on mutual observations between the robots. The intuitive expectation that more robots should improve the position estimation is motivated and the design of the experiments with special respect to possibly distorting parameters is discussed and reasoned in detail. A detailed analysis of the collected data shows to what extend the experimental results conform to the expected outcome. Possible explanations for the rather small improvements with larger group sizes are considered.

References

1. Barshan, B., Durrant-Whyte, H.F.: Inertial navigation systems for mobile robots. IEEE Transactions on Robotics and Automation 11(3), 328–342 (1995)
2. Borenstein, J., Feng, L.: Measurement and Correction of Systematic Odometry Errors in Mobile Robots. IEEE Transactions on Robotics and Automation 12(6) (1996)
3. Fox, D., Burgard, W., Kruppa, H., Thrun, S.: A Probabilistic Approach to Collaborative Multi-Robot Localization. Autonomous Robots 8(3), 325–344 (1996)

4. Rekleitis, I.M., Dudek, G., Milios, E.E.: Probabilistic Cooperative Localization and Mapping in Practice. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Taipei, pp. 1907–1912 (2003)
5. Madhavan, R., Fregene, K., Parker, L.E.: Distributed Cooperative Outdoor Multirobot Localization and Mapping. Autonomous Robots 17(1), 23–39 (2004)
6. Moutinho, A., Azinheira, J.R.: Comparison and fusion of odometry and GPS navigation for an outdoor mobile robot. In: Proceedings of the 11th IEEE International Conference on Advanced Robotics (ICAR), Coimbra (2003)
7. Kurazume, R., Hirose, S.: An Experimental Study of a Cooperative Positioning System. Autonomous Robots 8(1), 43–52 (2000)
8. Howard, A., Matarić, M.J., Sukhatme, G.S.: Cooperative Relative Localization for Mobile Robot Teams: An Egocentric Approach. In: Schultz, A.C., Parker, L.E., Schneider, F.E. (eds.) Multi-Robot Systems: From Swarms to Intelligent Automata, pp. 65–76. Kluwer, Dordrecht (2003)
9. Schneider, F.E., Wildermuth, D.: An Application of Relative Localisation for Multi-Robot Navigation. In: Proceedings of the IASTED International Conference on Automation, Control, and Applications (ACIT-ACA), Novosibirsk (2005)
10. Roumeliotis, S.I., Rekleitis, I.M.: Propagation of Uncertainty in Cooperative Multirobot Localization: Analysis and Experimental Results. Autonomous Robots 17(1), 41–54 (2004)
11. Mourikis, A.I., Roumeliotis, S.I.: Performance Analysis of Multirobot Cooperative Localization. IEEE Transactions on Robotics 22(4), 666–681 (2006)
12. Schneider, F.E., Wildermuth, D.: Directed and Non-Directed Potential Field Approaches to Formation Navigation. In: Proceedings of the 2nd IASTED International Conference on Automation, Control, and Applications (ACIT-ACA 2005), Novosibirsk (2005)
13. Bar-Shalom, Y., Li, X.-R.: Estimation and Tracking: Principles, Techniques and Software. Artech House, Boston (1993)
14. Technische Beschreibung Lasermesssysteme, doc. 8008969/06-2003, SICK AG, Division Auto Ident, Germany (2003)
15. Kräußling, A., Schneider, F.E., Wildermuth, D., Göb, R.: Zur Verfolgung ausgedehnter Ziele – eine Übersicht über ausgewählte Algorithmen und ein Vergleich deren Güte. FKIE-Forschungsbericht, vol. (94). Fraunhofer FKIE, Wachtberg (2005)
16. Schneider, F.E., Wildermuth, D., Kräußling, A.: Discussion of Exemplary Metrics for Multi-Robot Systems for Formation Navigation. International Journal of Advanced Robotic Systems 2(4), 345–353 (2005)
17. Navarro, I., Matía, F.: A Proposal of a Set of Metrics for Collective Movement of Robots. In: Proceedings of the Workshop on Good Experimental Methodology in Robotics, Robotics: Science and Systems Conference (RSS), Seattle (2009)