Acquisition of 2-D Ground Truth Data in Multirobot Experiments

Quentin Alexander
Computer Science
University of Alabama
Tuscaloosa, AL, 35487 United States
Email: qaalexander@crimson.ua.edu

Harold Hockett II
Computer Science
Tennessee State University
Nashville, TN, 37209 United States
Email: hhockett@mytsu.tnstate.edu

Aparna Veluchamy
Computer Science
University of Alabama
Tuscaloosa, AL, 35487 United States
Email: aveluchamy@cs.ua.edu

Dr. Monica Anderson
Computer Science
University of Alabama
Tuscaloosa, AL, 35487 United States
Email: anderson@cs.ua.edu

***Abstract*— This research examines three different approaches to acquisition for 2D ground truth data via an external localization system for subsequent analysis. We suggest criteria for localization system selection based on our experiences. There are many positioning system options available and we promote the use of localization methods that can provide real-time data and require little overhead in terms of processing and environmental setup. It is hoped that this study will assist other researchers in making better localization data collection decisions.**

I. Introduction

Although many robot algorithms can be tested in simulation, only experiments with real robots in real environments can quantify performance in the presence of unmodeled phenomena. In coverage algorithms specifically, performance is measured by the ability of a team of robots to attend to all reachable areas in the region of interest. Coverage is best calculated in terms of exteroceptive sensors. Proprioceptive sensors (odometry) can accumulate error over time and ceases to accurately represent the true position of the robot [1].

This research presents the criteria for selecting a localization system that can collect ground truth position for subsequent analysis. This work is proposed in the context of indoor experimentation with multiple mobile robots with or without digital communications ability. We examine three different approaches to illustrate the usefulness of these criteria. It is hoped that this study will assist other researchers in making better localization data collection decisions.

Criteria for development/selection of indoor localization systems:

- Accuracy in x, y, and theta dimensions and data frequency
- Cost per robot and scaling to experiment space size
- Integration into power systems
- Integration into data collectors (USB, Aircable/Bluetooth, Serial, etc)
- Ability to collect data remotely and/or online
- Complexity of environment installation
- Calibration to enable collection of data in absolute units (meters or feet)

II. Localization Approaches

A. Approach 1 - Camera-based system using a discretized field

Three cameras (see figure 1) are used to record the positions of robots at 1- second intervals (robots moved at .1 m/s). Using wireless networked cameras relieved the need for running data interface cables. The cameras (2 normal angle and 1 wide-angle) were attached to the walls of the room roughly 1.5 meters from the floor. Although the camera placement provided a full view of the test area, some issues arose because the optical axis of each camera was not perpendicular to the floor.

Fig 1 shows the grid of 1 meter squares that was placed on the floor. Due to the camera placement, cells closer to the camera are larger in the image than the cells further away. This affected the post-processing of images and increased the difficulty of tracking the position of the robot.

Position information was obtained from post-processing the images for the location of the robots relative to the marked cells in the images. Position resolution using this method was limited to about .5 meters in practice due to the small size of the robot in some positions in the images. Orientation was effectively unavailable. Two points on the robot were needed to calculate orientation and two points could not be reliably found due again to the small size of the robot in the images (+/- 30). Also differentiating between robots in the images was also difficult, even with special color markers on each one. This method provides very coarse position information at a low equipment cost ($500 per camera) but high manual labor and storage costs.

Post-processing can be partially automated using a Matlab system that allowed a user to view each image to click on the robot position. From the X and Y coordinates in the image, the robot position was calculated. This method provided position information at the sub 0.5 meter interval. However, based on the ability of the operator to select the exact position on the robot each time, the position estimate varied greatly. In fact, the position accuracy also varied with the size and position of the robot in the image. Automatic robot identification in images may not be able to perform better without more cameras and better camera placement. Each additional camera creates an additional set of images to be processed. We deemed the cost of additional cameras in terms of processing was prohibitive.

Camera-based localization approachs that have been implemented successfully by several researchers. Those interested in pursuing a camera positioning system can find additional resources in works such as [3, 4, 5].

B. Approach 2 Northstar Localization System (Evolution Robotics)

The Northstar Localization System is a commercially available system that allows each robot to record its position relative to an active IR image projected onto the ceiling. The cost for a development kit that instruments one robot within a 6mx6m area is $1795. Each additional robot requires a $1400 detector making this approach more expensive to scale to larger teams.

Unlike the previous approach, data is collected locally requiring the detector to be integrated into the robot power and data systems. Power integration was straight forward as it took a 3.3V connection, available on most robot bases. The detector has a serial interface but did not integrate well with either USB to serial cables or an Aircable (serial to Bluetooth converter), which limited flexibility regarding the use of our single available serial port. The detector comes with a C-based programming API, which should make integration into robot development environments easily. However, the raw data stream was not ASCII readable making the testing of the serial connection via a terminal impossible. A bigger issue was the absence of a calibration routine. Measurements were given in generic units since the ceiling height was not calculated and/or stored. Placement of the projector also proved difficult since the image was invisible to the naked eye and the presence of ceiling mounted light fixtures.

C. Approach 3 Stargazer (Hagisonic)

The Stargazer localization system uses an active detector (IR camera) on each robot with passive landmarks on the ceiling. The cost scales with the number of robots with each additional detector costing $1000. In this system, the camera moves and the object being identified (ceiling landmark) remains stationary. A key advantage to this approach is that image processing is handled by the device and position information can be reported in real time.

With the data being collected locally, integration into the existing lab framework is important. Data collection via the serial interface is straight forward with USB-to-serial connection. Data streams are ASCII readable, which allows for verification of the connection to be separate from the use of the programming API. In addition, the Stargazer includes a

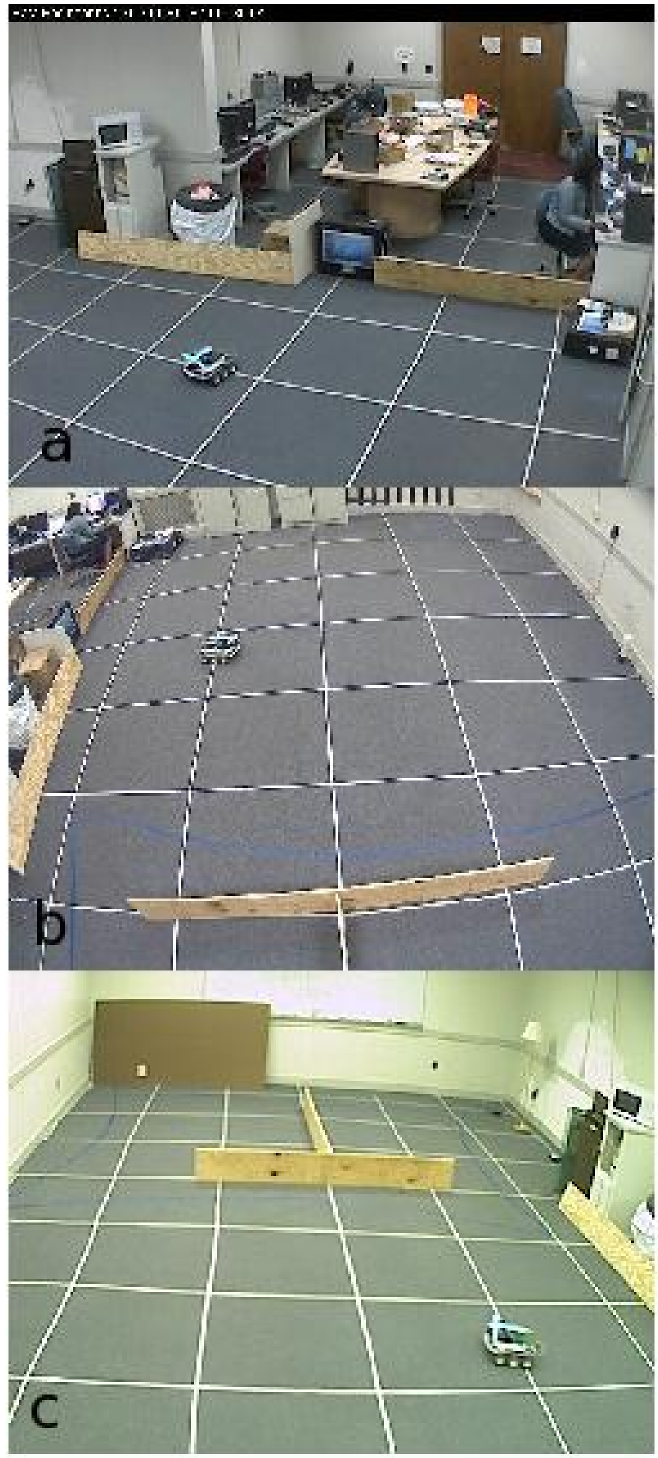

Fig. 1. Camera images obtained from coverage experiments in a 10x6meter lab. a) Normal angle view. b) Wide angle view. c) Normal angle view.

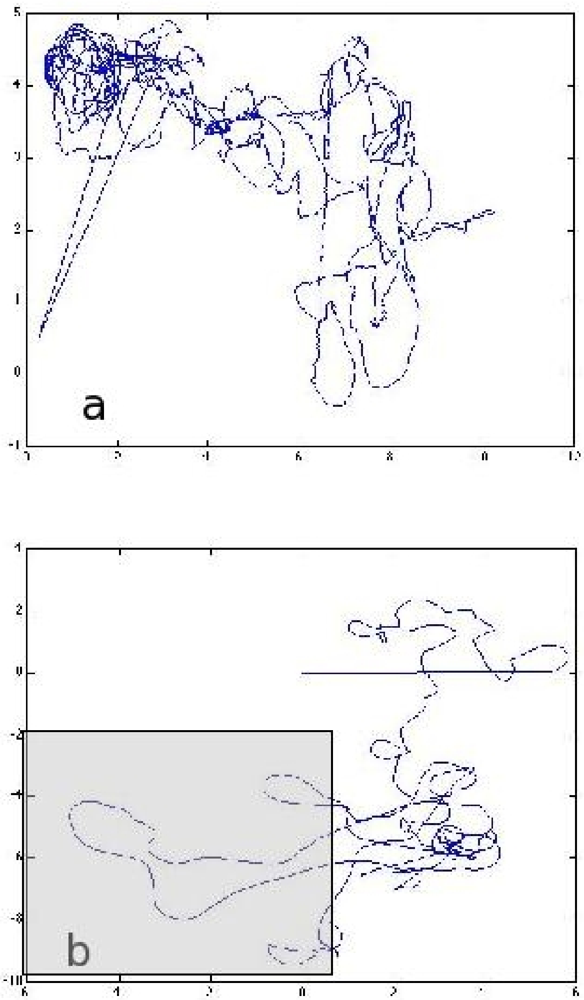

Fig. 2. Plots of x and y position data. a) Data obtained from the stargazer mapped to a global frame. Occasional misidentification of landmark IDs can create spurious data points. b) Data obtained from odometry. Gray shaded area does not exist in the real environment. Accumulated odometry error causes measurements to become more incorrect over time.

calibration step that calculates the distance to the ceiling based on the known structure of the ceiling landmarks. The result is an ASCII readable data stream in real units (cm and degrees). The power integration is much more complicated since both 5v and 12v power sources are required (two boards: processor and detector). We found the translational accuracy of the Stargazer relative to each marker is +/- 5cm with the heading accuracy to be +/- 3 degrees. However, each landmark only covers a 2m area so relative measurements between the landmarks are needed to give absolute position in a 10x6m lab. There is a global mode that did not work in our version.

1) Prior Stargazer-based experiments: Data collected from the Stargazer has been published as part of multirobot cooperation strategy research [2]. Odometry and Stargazer data is collected via Player/Stage drivers and saved during experiments with a time server verified time stamps. Not only can

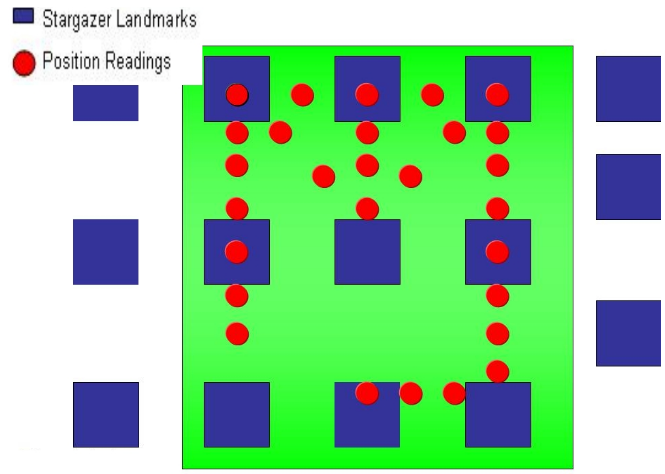

Fig. 3. Test environment used for the Stargazer, showing the location of the ceiling placards and the points at which the robot reported it's position. Note that some of the ceiling placards(blue) extend outside of the test environment.

robot coverage be calculated based on position (see Figure 2) and orientation (non-footprint coverage), but proximity between robots can also be calculated.

III. Localization System Testing

A. Test Setup

Tests were performed to evaluate the Stargazer, the Northstar, and the K-team Koala robot's internal odometry. The test environment is roughly 6x6 meters. In each test set, the relevant system reported its perceived values for x, y, and theta which were compared to the actual (true) values taken from ground measurements.

In the Stargazer tests, a K-team Koala robot equipped with that device traveled a predetermined path around the environment. Periodically the robot stopped moving and reported its position readings from both the Stargazer and the robot's internal odometry. Each time the robot reported its position, a mark was placed on the floor. The exact (true) position of the robot was determined by measuring the floor markings. Figure 3 shows the location of the Stargazer's ceiling placards and the general locations where the robot reported its position.

For the Northstar localization system tests, the projector was positioned such that its infrared 'spots' would be on the test area's ceiling. The detector was mounted on a small platform and reported position readings via a wired connection to a computer terminal. A program that is provided with the Northstar Development Kit was used for calibrating the device and retrieving position information. Position data was collected by moving the detector to different locations within the test environment. Figure 4 shows the test environment, along with the placement of the IR spots and the general location of the points where position measurements were taken.

Environment interference an restrictions can make it difficult to setup an ideal projector configuration. In order to produce

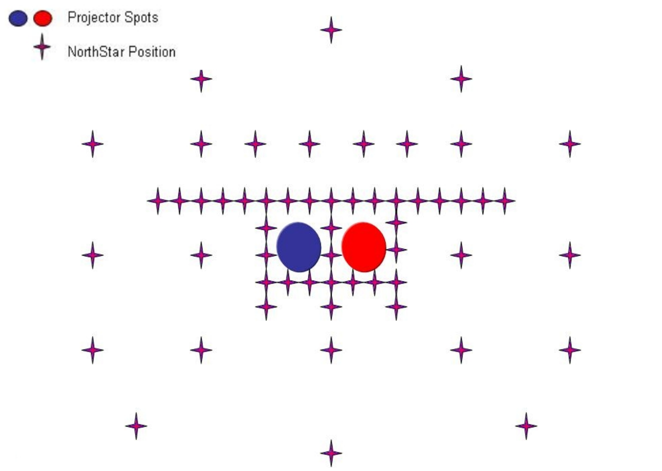

Fig. 4. Test environment for Northstar projector configuration A showing the location of the IR ceiling spots and points where the system reported position values.

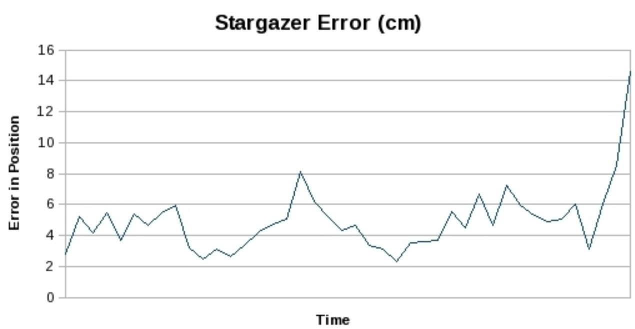

Fig. 5. Stargazer position error showing the distance between the robot's true position and the report generated by the Stargazer.

the most accurate and consistent position values it may be necessary to try multiple projector configurations and IR spot placements. In our experiments we tried several different projector configurations and we conducted detailed tests on the two configurations that showed the least amount of environmental interference. In Configuration A, the projector aims its IR spots near the center of the test environment's ceiling. The IR spots in Configuration B were directed closer to the front edge of the ceiling. Changing between configurations involved both moving the projector mount and adjusting the angle of projection. The Northstar detector was calibrated separately for each configuration.

The Northstar system sets its origin to an area centered between the two IR spots. The calibration program use that same origin and reports the data in units specified during calibration. Theta is available in degrees before the system is calibrated.

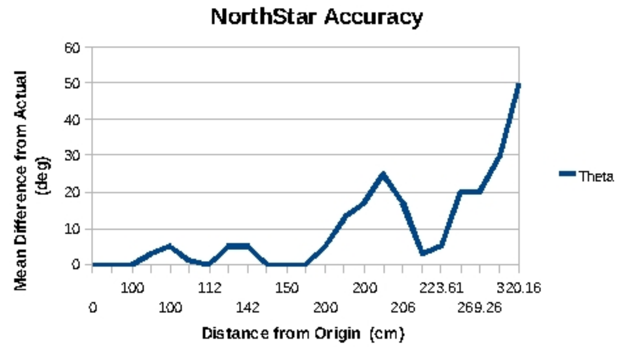

Fig. 6. Graph of the absolute difference between the actual theta and the value reported by the Northstar.

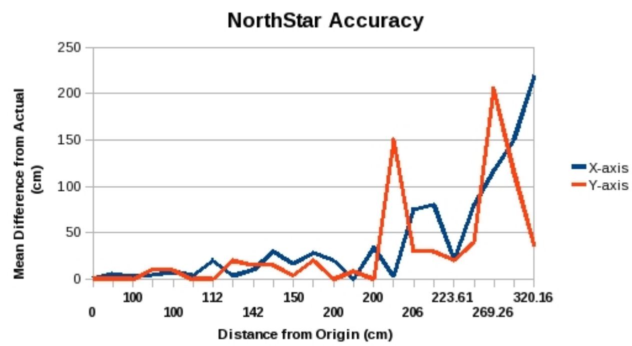

Fig. 7. Graph of the absolute difference between the true x and y axis and the values reported by the Northstar.

IV. Test Results and Analysis

A. Stargazer

The error in position readings reported by the Stargazer are illustrated in Figure 5. A graph shows the distance between the robots actual position and the values perceived by the Stargazer. The results for the Stargazer fell in line with the data mentioned earlier. The majority of position estimates were within a few centimeters of the robots actual position.

There are a few large spikes in the error graph. Analysis of the test data showed that these larger deviations occurred when the robot overlooked a nearby ceiling marker and localized itself using a more distance marker. Although environmental interference is most likely, the definite cause for this phenomenon has yet to be determined.

B. Northstar

Results from the tests on the Northstar are shown in Figures 6 and 7. Multiple values were recorded for each position and the figures provide the absolute difference between the mean and the actual (true value).

The Northstar system performed similarly for the two projector configurations. Within one half meter of the origin, the Northstar is very accurate. Readings at this distance are correct to within 1 centimeter and 2 degrees. Inaccuracies of up to 10cm and 3 degrees are seen at the one meter distance.

TABLE I

SUMMARY OF LOCALIZATION SYSTEMS

Criteria	**Sparse camera-based systems**	**Northstar Localization System**	**Stargazer Localization System**
Accuracy in x, y, and theta dimensions and data frequency	Not accurate, especially in theta	Accurate but may be sensitive to lighting	Accurate relative to each landmark
Cost per robot and scaling to experiment space size	Scales with area size (# of cameras) and amount of post-processing	Scales with number of robots and area size	Scales with number of robots
Integration into power systems	Trivial	Simple	Involved
Integration into data collectors (USB, Bluetooth, Serial, etc)	Trivial via network (separate subnet needed)	Complicated	Involved
Ability to collect data remotely and/or online	Remote access via network; position calculations are offline	Collection remotely is complicated due to issues with integration into standard wireless interfaces	Collection remotely and online somewhat simplified by integration into existing technologies
Complexity of environment installation	Camera placement is key to results	Complex placement of projector	Trivial to place landmarks on ceiling
Calibration enables data in real units (meters, feet, etc)	A lot of calibration required	No calibration available; must compensate in drivers	Calibration available; completed one-time for the environment

Around 2 meters from the origin, the detector reports readings that deviate up to 35 centimeters and 20 degrees from actual. Position values taken at over 2 meters from the origin are almost entirely unreliable.

The accuracy of position estimates from the Northstar depends primarily on the distance between the detector and the origin (area between the IR spots). As the detector moves along one axis, the values for the orientation and the other axis will eventually degrade. This effect is visible at distances above 1.5 meters. However, changes in theta do not show a similar effect on the coordinate values when the detector is far from the origin.

Variance in position readings remained low throughout much of the testing. With the detector located three meters or more from the origin, position readings became highly erratic. Values for both axis would fluctuate as much as 120 centimeters from the mean. At this distance, the detector also had trouble determining the sign (positive/negative) for the range of values being reported for a given position. It was also noted that prolonged use of the detector seemed to cause values to vary somewhat for the same position regardless of distance to the origin.

One interesting observation is that although the position readings became unreliable, the detector did not loose sight of the IR spots.

The Nortstar's configuration program has a feature that shows an image of how the detector perceives the location of the IR spots. Throughout testing, the detector maintained its view of the IR spots. When moving the detector around the environment, the general location and orientation of the IR spots was also maintained. At distances where position values become unreliable, this image feature shows the spots wavering around but their general location (as perceived by the detector) remains correct. It follows that the Northstar system can still provide some useful information, in terms of heading and so forth, even as its position coordinates degrade.

V. CONCLUSION

In conclusion, gathering accurate ground truth data is important and there are a number of good options available to mobile robot research groups. Key concerns when choosing a system include, cost, data processing, and the complexities of device setup and maintenance. Our experience has shown that systems that operate similarly to the Northstar or Stargazer maintain a fair balance between all of the important positioning system tradeoffs.

ACKNOWLEDGMENT

This work was supported in part by the following NSF grants: IIS-0846976 and CCF- 0829827.

REFERENCES

[1] J. Borenstein and Liqiang Feng, Correction of systematic odometry errors in mobile robots, Intelligent Robots and Systems 95. 'Human Robot Interaction and Cooperative Robots', Proceedings. 1995 IEEE/RSJ International Conference on, 1995, pp. 569-574 vol.3.

[2] M. Anderson and N. Papanikolopoulos, Improving Multirobot, Cooperative Search via Local Target Queues, IROS 2007.

[3] A. Georgiev and P. Allen, Localization methods for a mobile robot in urban environments, Robotics, IEEE Transactions on, vol. 20, 2004, pp. 851-864.

[4] W. Dixon, D. Dawson, E. Zergeroglu, and A. Behal, Adaptive tracking control of a wheeled mobile robot via an uncalibrated camera system, Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on, vol. 31, 2001, pp. 341-352.

[5] J. Wolf, W. Burgard, and H. Burkhardt, Robust vision-based localization by combining an image-retrieval system with Monte Carlo localization, Robotics, IEEE Transactions on, vol. 21, 2005, pp. 208-216.