Range-Domain Integration of GPS and Laser-scanner Measurements for Outdoor Navigation

Mathieu Joerger and Boris Pervan
Illinois Institute of Technology

BIOGRAPHY

Mathieu Joerger is a PhD student in Mechanical and Aerospace Engineering at the Illinois Institute of Technology (IIT) in Chicago. In 2002, he obtained a Master in Mechatronics at the National Institute of Applied Sciences (INSA) in Strasbourg, France, and a Master of Science in Mechanical and Aerospace Engineering at IIT. He is currently working as a research assistant on the design of navigation and control algorithms for autonomous vehicles, and on GPS ranging augmentation systems.

Boris Pervan received a B.S. from the University of Notre Dame (1986), M.S. from the California Institute of Technology (1987), and Ph.D. from Stanford University (1996), all in Aerospace Engineering. From 1987 to 1990, he was a Systems Engineer at Hughes Space and Communications Group. Dr. Pervan was a Research Associate at Stanford from 1996 to 1998, serving as project leader for GPS Local Area Augmentation System (LAAS) research and development. He was the 1996 recipient of the RTCA William E. Jackson Award and the 1999 M. Barry Carlton Award from the IEEE Aerospace and Electronic Systems Society. Currently, Dr. Pervan is Associate Professor of Mechanical and Aerospace Engineering at IIT where he received the Outstanding Graduate Teaching Award (2002), and the University Excellence in Teaching Award (2005).

ABSTRACT

In this paper, an innovative navigation algorithm combining GPS and laser-scanner measurements is analyzed and tested in natural outdoor environments. Using carrier phase differential GPS, centimeter-level positioning of autonomous ground vehicles is achievable. However, GPS signals are easily attenuated or blocked, so their use is generally restricted to open-sky areas. In response, in this work we augment GPS with two-dimensional laser-scanner measurements.

The two sensors are integrated in the range domain for optimal navigation performance. The GPS/Laser range-domain integration is performed using a measurement-differencing extended Kalman filter. In addition, the use of laser measurements requires that we address the feature extraction process, which aims at selecting features in the environment that can be consistently identified, and the data association procedure, which establishes correspondences between these extracted measurements and a continuously updated map of landmarks. Using this algorithm, the vehicle's position is determined throughout GPS outages, without a-priori knowledge of the surrounding landmarks' locations.

Experimental testing in actual urban canyons in Chicago, where fewer than four GPS satellite signals are available, demonstrates that the performance of the range-domain integrated positioning and mapping algorithm exceeds that of a more traditional position-domain implementation.

I. INTRODUCTION

This paper describes the design, analysis and testing of a navigation system based on combined Global Positioning System (GPS) and laser-scanner measurements for use in natural outdoor environments. In order to fully exploit the complementary properties of the two sensors, their measurements are integrated in the range domain. The performance of the resulting positioning and mapping algorithm exceeds that of a simpler position-domain implementation, especially in places such as urban canyons, where fewer than four GPS satellite signals are available.

Carrier-Phase Differential GPS (CPDGPS) provides centimeter-level accuracy, which is critical for many precision applications such as land-mine detection [1], unmanned agricultural field plowing [2] and intelligent traffic systems [3]. However, robust CPDGPS navigation requires good sky visibility because GPS satellite ranging signals can be significantly attenuated or blocked by buildings, trees, and rugged terrain. In response, we augment GPS with two-dimensional laser-scanner measurements from surrounding obstacles, which are used as landmarks. Laser information is available when

GPS is not, and provides in addition, a means for obstacle detection. Because the laser-based positioning error increases with distance, as compared to inertial navigation which drifts with time [4], it constitutes an interesting alternative for bridging GPS-outages, especially in low-speed applications.

The range-domain integration of GPS and laser measurements is performed using a measurement-differencing extended Kalman filter (EKF). Laser measurements are processed according to a Simultaneous Localization and Mapping (SLAM) procedure [5], which determines the vehicle's position using previously unknown features in the environment.

As illustrated in figure 1, the hundreds of noisy raw laser measurements are not fed directly into the EKF. Two intermediary procedures are implemented: first, the feature extraction aims at selecting the few measurements describing actual landmarks; second, the data association consists in assigning these extracted observations to the corresponding landmark states in the EKF. These procedures have dramatic effects on the positioning performance and are particularly challenging in natural outdoor environment, where landmarks that can be consistently identified are rare.

In this paper, we analyze the three components of the algorithm: position estimation, feature extraction and data association. A first series of tests, conducted in a structured environment to ensure successful data association, demonstrates the performance of the position estimation process. Further testing in the streets of Chicago shows that failures to correctly associate measurements are a driving element in the final navigation error. In both cases, the range-domain integration demonstrates better performance than the position-domain implementation, not only for the position estimation, but also with regard to the data association.

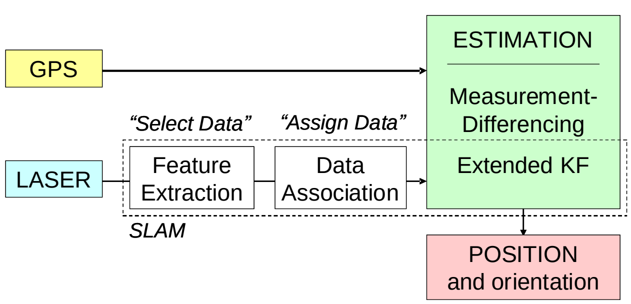

Figure 1: Block Diagram of the GPS/Laser Algorithm

MEASUREMENT-DIFFERENCING EXTENDED KALMAN FILTER

Determining the user's position based on GPS and laser-scanner information can be performed by simply combining the individual positioning outputs of each sensor. In this case, when less than four GPS satellites are available, the user's position estimate is based solely on laser measurements. In partially obstructed GPS environments such as urban canyons, there are often two or three satellites available, which with this position-domain approach are left unused. If in contrast we integrate the GPS and laser measurements in the range-domain, we can make use of these two or three satellite ranging signals. In the later sections of this work, we analyze and quantify the substantial contributions of these additional measurements for the final positioning solution. Note that one single satellite in view is useless because of the undetermined GPS receiver clock bias.

In this section, the challenges in the derivation of the range-domain algorithm are presented (for details on the mathematics of the filter, refer to [6]). A brief performance analysis illustrates the main characteristics of the estimation process.

A. Laser-based SLAM

The terms SLAM and Concurrent Mapping and Localization (CML) [7] refer to the problem of vehicle positioning using previously unknown features in the environment. Over the past two decades, several alternative solutions to the SLAM problem [8] have emerged. We selected an EKF approach which provides an incremental solution and can therefore be utilized in real time. The laser's angular and ranging measurements are expressed in terms of the vehicle and landmark position and orientation states in a Cartesian frame (e.g. north, east, down). The resulting non-linear measurement equations explain the need for an Extended Kalman Filter.

The complete solution to the laser-based SLAM problem includes the following tasks:

- feature extraction, in which we include noise rejection, segmentation and data selection; this procedure consists in extracting one measurement per landmark, out of the hundreds which constitute a laser scan
- data association, which aims at identifying from which landmark the extracted measurement is originating from, and
- vehicle and landmark position estimation, which is performed simultaneously using an EKF.

The first two tasks are addressed in detail in the next section.

The estimation process described in the third task can be summarized as follows: given an initial position estimate (here provided by GPS), the vehicle's absolute position and heading over time can be determined by keeping track of its relative location with respect to surrounding landmarks using laser measurements. Because the landmarks' locations are not known in advance, both vehicle and landmark position coordinates constitute the state vector to be estimated in the EKF.

B. GPS Algorithm for Range-domain Integration

In previous research, we designed a navigation and control system for an AGV [9], which we used for autonomous lawn-mowing applications [10]. Real-time cm-level positioning is achieved using differential code and carrier-phase measurements. Because of the time correlation in the measurement noise due to multipath reflections, our former GPS-based navigation system was composed of two separate processes [11]: (1) the integer estimation procedure was a Kalman Filter measurement update performed at infrequent intervals to whiten the measurement noise; (2) intermediary position estimation at each sample time was obtained using a weighted least squares solution. In the perspective of a unified GPS/laser integration, this dual-stage procedure is not an optimal solution because laser measurement updates need to be performed as frequently as possible.

In order to implement frequent GPS filtering updates, we choose to model the colored measurement noise as a first order Gauss Markov process. In this work, we use a measurement-differencing filter, which was first introduced in 1968 by Bryson and Henrikson as a way to model correlated measurement noise in a state space representation [12][13]. It is a very efficient alternative to state augmentation because the number of states remains unchanged and the measurement noise matrix is no longer singular. The core idea defining this filter is the elimination of the correlated measurement noise terms using two differenced consecutive measurements.

When compared to a more traditional state augmentation method, the measurement-differencing filter state-efficiency is well worth the cost of a few complications in the implementation (storage of the prior measurement vector and observation matrix, and initialization procedure). Indeed, when processing code and carrier phase measurements from a 12-channel dual-frequency receiver, the state-augmented filter requires 48 additional states (refer to [6] for more details). The resulting GPS algorithm has potential applications beyond this work since it performs the combined estimation of both position and cycle ambiguities at any update rate.

C. Analysis of the GPS/Laser Estimation Process

Differential code and carrier phase measurements as well as ranging and angular laser data are fed into a single measurement-differencing EKF to simultaneously estimate the vehicle three-dimensional position and its heading angle, the GPS receiver clock bias and cycle ambiguities, and the landmark locations, which constitute the states of the integrated system.

The following performance analysis assumes a 'forest scenario': the AGV starts in a GPS-available area (yellow in the bottom plots in figure 2) and roves across a GPS-unavailable zone (blue) using tree-trunks as landmarks (black circles). The essence of the GPS/Laser integration is illustrated in the transitional GPS-and-laser-available area (green). For this simulation, we assume that tree trunks are vertical cylinders, and the GPS satellite blockage due to the tree canopy is modeled using a horizontal plane on top of these cylinders. We make sure that no low-elevation satellite signal penetrates inside the forest because such signals would be affected by multipath reflections on tree trunks.

Figure 2 shows four successive snap-shots (a, b, c and d) of a direct simulation of the GPS/Laser algorithm while the vehicle roves across a forest. On the upper part, azimuth-elevation plots of the sky show the GPS satellite blockage due to the forest. On the lower part, the result of the estimation process is presented. The covariance ellipses represent the positioning error on the vehicle (red), and on the landmarks (blue).

In figure 2.a, the vehicle starts in a GPS available area. There are many satellite signals available, and the cycle ambiguities have been accurately estimated prior to this point, so that the uncertainty on the vehicle position does not exceed a few centimeters.

Then, in figure 2.b, the vehicle enters the green area. There are still more than four satellites available, so that the vehicle's position is accurately determined. A first landmark is within range of the laser scanner. Using GPS only (in the absence of a reliable dynamic model or heading sensor), the vehicle's heading is unknown; this is why the laser's angular measurement is of little use for the estimation of the landmark's position. The ranging measurement however determines the distance between the landmark and the vehicle. This explains the shape of the blue ellipse. Over time, as the system collects redundant observations from this landmark (we assume that the landmarks are stationary), and as the geometry between vehicle and landmarks changes due to the rover's motion, the landmark position estimate improves steadily, and so does the vehicle's heading estimate.

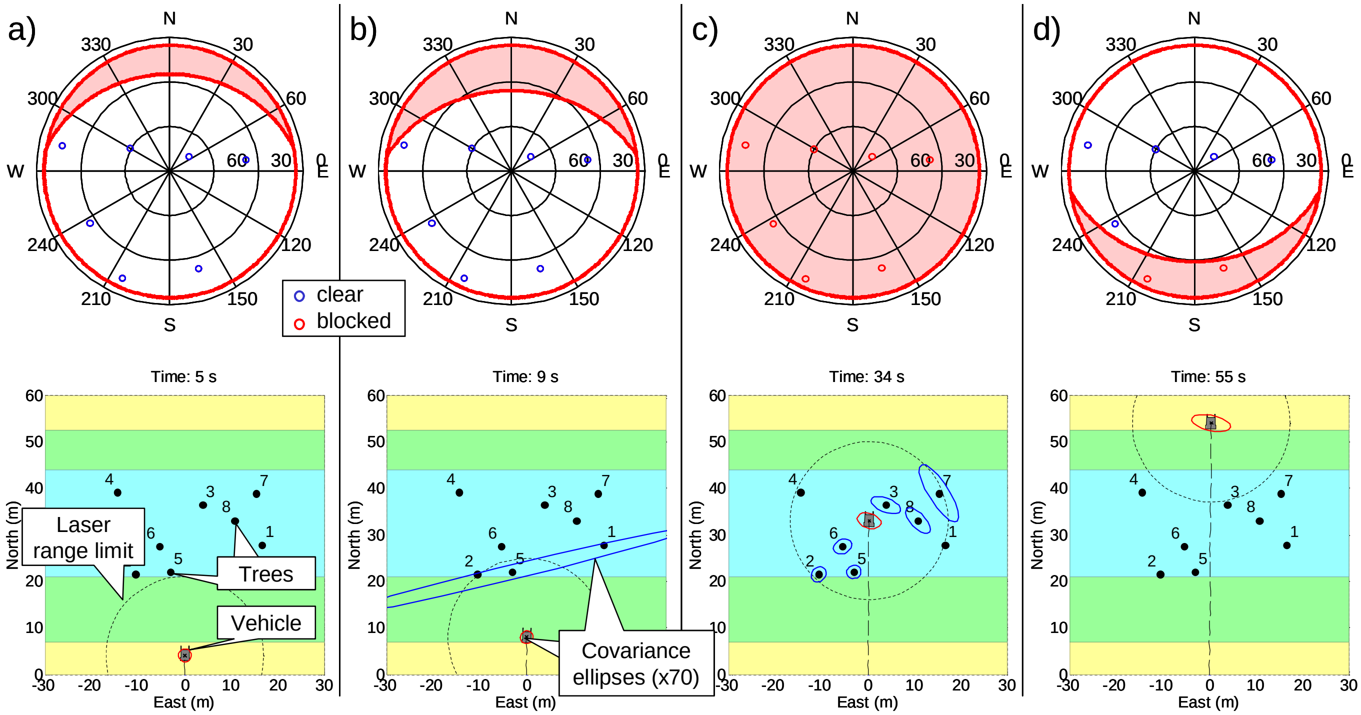

Figure 2: Direct Simulation of the GPS/Laser Algorithm

The next sub-figure (figure 2.c) shows the vehicle in the middle of the forest. In this case, there are no more satellites available, and the vehicle's position estimation relies on laser measurements only. Using laser-based SLAM, the rover's cross-track deviation drifts with distance as earlier landmarks get out of the scanner's range and new landmarks come in sight. Note that this drift also depends on the landmarks' geometry.

Finally, in figure 2.d, the vehicle is back into a GPS available area, and the cross-track deviation drift is stopped. Further filtering of GPS measurements over time will bring the cycle ambiguities and vehicle position estimates back to their initial precision, before entering the forest.

The experimental data in the section entitled 'Testing in a Structured Environment' validate the results of this simulation.

FEATURE EXTRACTION AND DATA ASSOCIATION

Practical implementation of laser-based SLAM requires additional processing before the measurements are sent to the EKF (figure 1). These intermediary procedures are discussed in this section.

A laser-scanner emits pulsed infrared laser beams that are reflected from surfaces of nearby objects and returned to the scanner. Signal time-to-return measurements are used to determine distances to the reflecting objects. The precision of the ranging measurement is affected by target surface properties (color, material reflectivity) and by the angle of incidence of the laser on the target surface. The pulsed laser beam is deflected with a rotating mirror to enable two-dimensional scanning [14].

As a result, a raw laser scan is made of hundreds of ranging measurements at regular angular intervals. Feeding hundreds of measurements directly into the EKF would be cumbersome; besides, not all of the data in a laser scan is useful. Therefore, we first get rid of noisy measurements and select the relevant information in a process called feature extraction. Next, the data association procedure aims at identifying the states in the EKF for which the extracted measurements provide observations.

Over the past two decades, extensive research has been dedicated to finding a resolution to these two problems [15], [16], which are especially challenging in natural environments. Elaborate algorithms have demonstrated promising results in specific environments. In this section, rather than describing the details of their implementation (see Appendices), we focus on the interactions between the feature extraction, data association and position estimation.

A. Feature Extraction

The goal of the feature extraction algorithm is to find features in the raw laser scan that can be repeatedly and consistently identified while the vehicle is moving and the laser's view-point is changing. With regard to the number of extracted measurements, the feature extraction routine should take the following tradeoff into account:

- on the one hand, more measurements generate better position estimates using an EKF,
- on the other hand, more extracted measurements increase the risk of failures in the data association process.

Because faults in the association have much more dramatic effects on the final position solution than the use of a few additional measurements in the EKF, we calibrate the feature extraction algorithm so that only the few easiest to identify landmarks are considered.

In the case of an urban canyon, the building's edges meet the above selection criteria: they are few within the range of the laser, sparsely distributed, and can be consistently extracted. The details of the implementation are given in Appendix I.

A laser scan taken on the site of one of the experiments (in an alley in Chicago), is presented in figure 3. The data is very noisy because of the wide variety of materials found in the street (wood, brick, glass, metal, vegetation), which all affect the laser beam differently. Also, walls, doors and fences often obstruct gaps between buildings so that the building's edges are no longer visible on the laser scan; therefore, we need to use edges of garage doors as landmarks. Such cluttered and inconsistent landmarks are likely to lead to confusions and failures in the data association process.

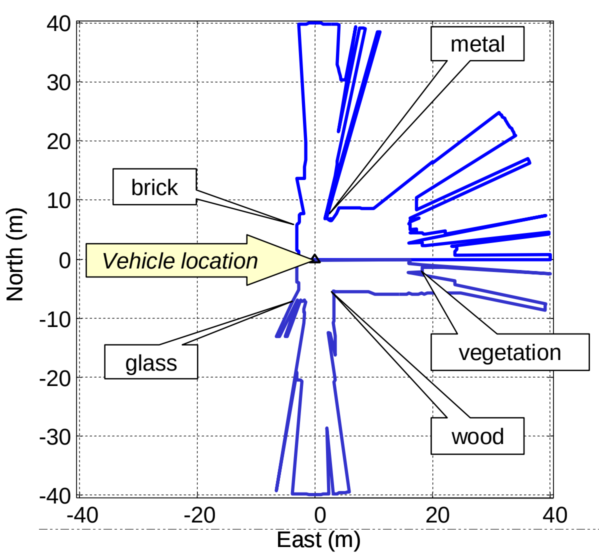

Figure 3: Raw Laser Scan Superimposed with a Satellite Picture of the Street

B. Data Association

The data association process establishes correspondences between consecutive sets of measurements and a continuously updated map of landmarks. More precisely, the current extracted measurements (resulting from the feature extraction) are matched with the predicted measurement estimates obtained from the EKF (which are based on all prior observations). Projections of the state estimates are evaluated using a linearized vehicle kinematic model.

We use a nearest neighbor approach based on the normalized innovation square to perform the association (see [17] and Appendix II for more details). This process is widely implemented in applications of this type, and more elaborate versions can be found in the literature [18], [19].

Although they are implemented separately, the data association and the estimation processes are deeply related. First, the influence of the estimation process on the data association is described in Appendix II. Then and more importantly, failures in the data association process, which we call miss-associations, can lead to the following outcomes:

- the measurement is not associated with its corresponding landmark, therefore we assume that it corresponds to a new landmark (probably nearby the former landmark)
- the measurement is associated with the wrong landmark.

In the first case, the consequence for the estimation process is that there will be fewer observations for this given landmark. The second case however can have catastrophic effects on the estimation process, as illustrated in figure 4 for the position-domain integration. In this example, because of erroneous vehicle position and orientation estimates, the system confused a landmark on the left of its trajectory for one on its right. In the following iterations, because the map of landmarks is built incrementally, the vehicle and landmark position errors accumulate and grow without bound.

Fortunately, other correctly associated measurements mitigate the effects of such miss-association. In addition, as we will demonstrate in the next sections, the GPS/Laser range-domain integration enables to utilize additional absolute ranging measurements which help recover from data association failures.

TESTING IN A STRUCTURED ENVIRONMENT

The first set of data was collected in a structured environment. Landmarks were located at locations sparse enough to ensure successful data association. Because the results presented here are free of miss-association, they describe specifically the estimation process.

Testing is carried out at the Illinois Institute of Technology. In order to obtain a full 360° laser scan, we assemble two 180° laser-scanners back-to-back. The GPS antenna is mounted on top of the front laser. The lever-arm distance between the two lasers is included in the model. The two lasers and the GPS antenna are mounted on an existing AGV platform equipped with a dual-frequency GPS receiver (figure 5). An embedded computer onboard the vehicle records all the measurements including the raw GPS data from the reference station transmitted via wireless spread-spectrum data-link. Synchronization and measurement projections on a common reference sample time are achieved using the computer's clock. Truth vehicle trajectory and landmark locations are obtained using a fixed CPDGPS solution.

A. *'Forest Scenario'*

Because there is actually no physical obstruction to the sky, satellite masking for the GPS/laser integration system is performed artificially using the same model as in the previous direct simulation (figure 5). Tree trunks are reproduced using cardboard columns and a garbage can.

Figure 6 confirms that smooth transitions between the open-sky and GPS-unavailable areas are achieved. The position error does not exceed 15 cm in spite of 35 meters of GPS outage; the vehicle was roving along a straight line at an average speed of 0.8m/s.

The result of the range-domain integration differs only slightly from the position domain implementation. This is because in this scenario, the system transitions quickly from an open-sky area to having all satellite signals blocked. Greater differences will emerge in the case of an urban canyon scenario.

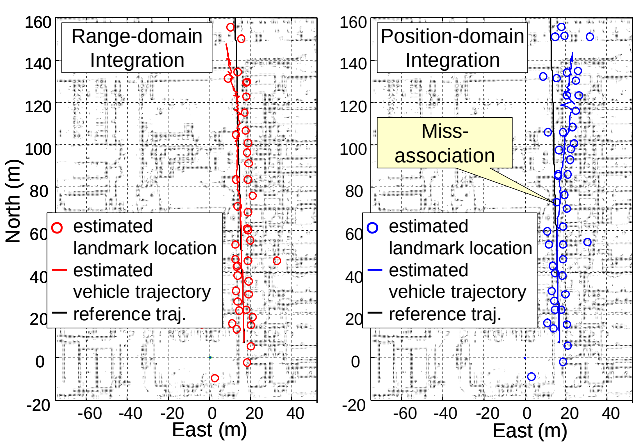

Figure 4: Consequence of a Miss-association in the Position-domain Approach

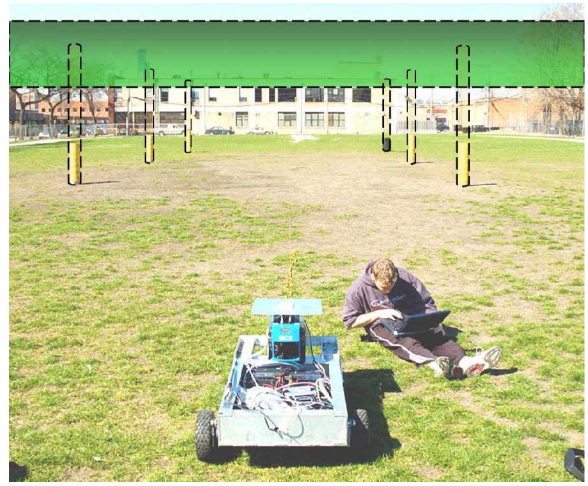

Figure 5: Experimental Setup and Artificial Satellite Blockage Model

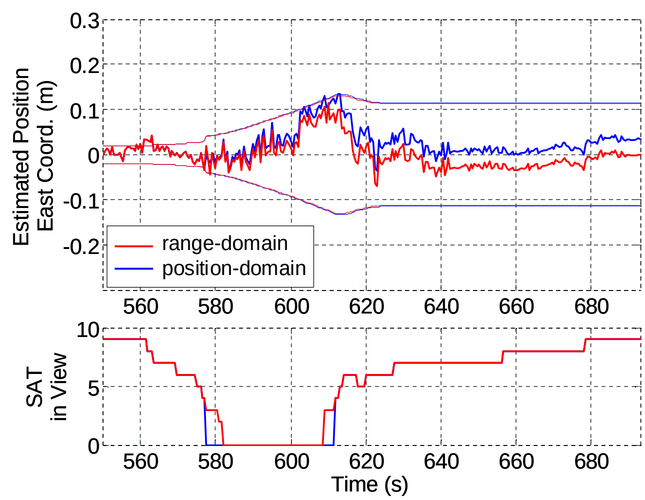

Figure 6: Experimental Result for the Forest Scenario

The result of the range-domain integration differs only slightly from the position domain implementation. This is because in this scenario, the system transitions quickly from an open-sky area to having all satellite signals blocked. Greater differences will emerge in the case of an urban canyon scenario.

B. *Miss-Association-Free 'Urban Canyon Scenario'*

GPS solutions are rarely available in urban canyons or forest roads because of the severe sky-blockage caused by bordering buildings and trees. The distinctive advantage of the range-domain integration is best illustrated here since the estimation process makes use of GPS signals that alone would be too few to generate a position fix. This sub-section aims at quantifying the navigation improvement brought by these few additional GPS signals as compared to position-domain integration.

For this miss-association-free case, we use the same set of data as in the preceding sub-section, but instead of artificially performing the satellite masking corresponding to a forest, we use a blockage model representative of an urban canyon (figure 7).

The results shown in figure 8 demonstrate as expected that as soon as there are less than four satellites in view, the range-domain integration is much superior to the position-domain implementation. In spite of 20 meters of GPS outage, the position error does not exceed 10cm.

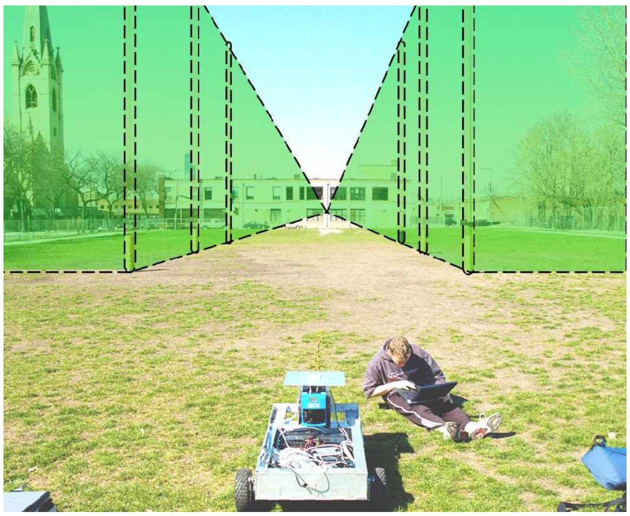

Figure 7: Satellite Blockage Model for the Urban Canyon Scenario

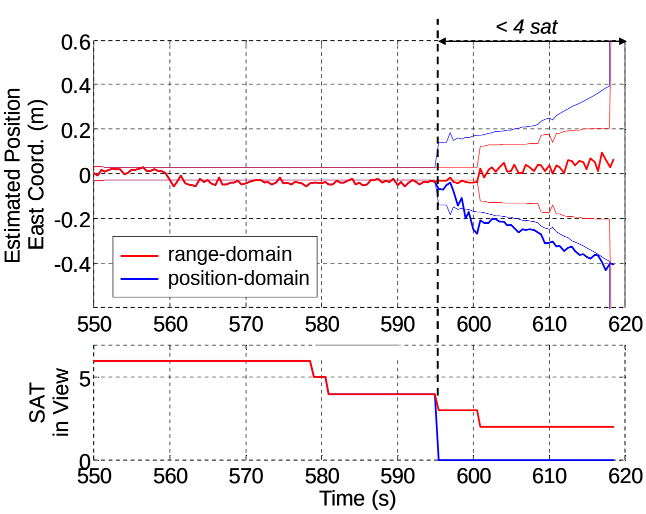

Figure 8: Experimental Result for the Miss-Association-Free Urban Canyon Scenario

TESTING IN URBAN CANYONS

Experiments in a natural environment, in the streets of Chicago, serve two main purposes: (1) they provide a measure of the performance of the algorithm in its current state when implemented in a realistic mission; (2) using the comparison with a position-domain approach, the testing results quantify the improvement brought by two additional GPS signals when miss-association of laser measurements are occurring.

In these experiments, the lasers and GPS antennas are mounted on a car, which is driven into an alley or a street bordered by buildings (figure 9).

For the set of data presented here, the GPS signals were actually blocked during the experiment, so that we could not use the fixed CPDGPS position solution to generate the truth trajectory over the entire experiment. Instead, we use a vehicle kinematic model to reproduce a reference trajectory. Additional testing in another canyon, in which a GPS-based truth trajectory was available, has been recently carried out. The results are consistent with the data presented here.

In figure 4, the estimated landmark locations are superimposed with satellite images of the alley. The estimated trajectory of the vehicle remains within the narrow alley, and the landmarks often match buildings' edges (also remember from the 'Feature Extraction' section that we had to use edges of garage doors as well).

In figure 10.a, the cross-track position error does not match the covariance envelope, which is an indication of miss-associations. As we know from the 'Data Association' section, miss-associations can generate great increases in the position error. Still, the absolute positioning error does not exceed 1.5m over 70 meters of GPS outage for an average vehicle speed of 1.5 m/s.

Figure 9: Experimental Setup for the Testing in the Streets of Chicago

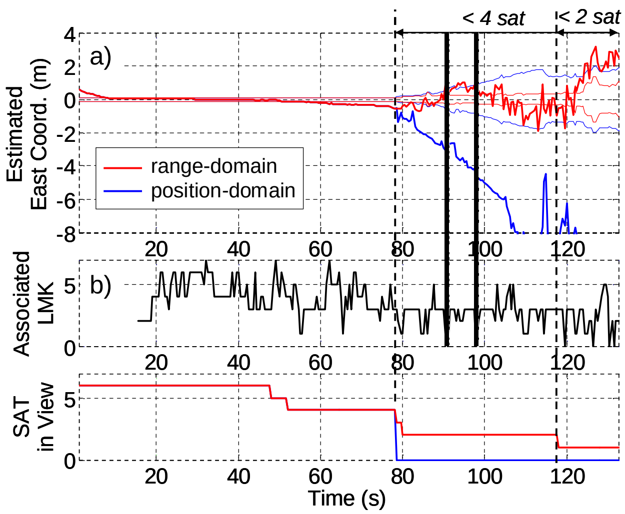

Figure 10: Experimental Results for the Test Conducted in an Alley in Chicago

To further explain these results, we notice that the number of associated landmarks (figure 10.b) sometimes drops to zero. In these cases, the only means to estimate the vehicle's position is the simple linearized kinematic model included in the algorithm. Note that we inflated the process noise on this vehicle model to make sure that it was not a decisive element of the estimation process.

Using the position-domain approach, the cross-track deviation quickly grows without bound (figures 4 and 10.a). Indeed, in this case, miss-associations led to a catastrophic navigation error. In contrast, for the range-domain integration, the two absolute GPS ranging measurements available were enough to recover from the miss-associations.

CONCLUSION

Laser-augmented CPDGPS greatly increases the availability of meter-level position solutions in outdoor environments. In partially obstructed GPS areas, the proposed range-domain integration exploits satellite signals that are too few to generate a position fix, by utilizing additional laser measurements.

Laser-based positioning in natural environments is difficult because it depends on the outcome of the challenging data association process. Using experimental data collected in both a structured environment and an actual urban canyon, we demonstrate that the use of two additional absolute satellite ranging measurements not only improved the estimation process, but also helped to recover from miss-associations.

AKNOWLEDGMENT

The authors would like to thank Bartosz Kempny and Samer Khanafseh for their valuable help in the acquisition of experimental data.

APPENDIX I: IMPLEMENTATION OF THE FEATURE EXTRACTION

This section presents an example procedure that we used to extract measurement to building's edges from raw laser scans. As illustrated in figure 3, raw laser data in natural environment is very noisy. Ideally, we would like to observe clean contours for each of the buildings, but fences, posts, garbage cans and other objects pollute the data. In addition, laser beams reflected at a low angle of incidence have very inconsistent returns.

Edges of buildings that point toward the laser-scanner are easy to identify because they correspond to local minima in the ranging measurements. We choose not to consider edges that point away from the scanner; they coincide with discontinuities in the ranging measurements and can therefore easily be confused with noise. The problem with this choice is that features targeted by the front laser are never the same as the ones extracted from the back laser scan.

The first stage of the procedure is to reject noisy measurements, while preserving the contour of the buildings. We segment the data based on a maximum threshold distance between consecutive measurements. If a segment contains less than a predefined number of measurements, it is considered as noise and the measurements are given average values based on the previous and following segments.

Next, in order to find the dominating local minima in the ranging measurements, we apply an averaging filter to the data. The size of the averaging kernel is tuned depending on the environment, so that the desired average number of landmarks is extracted.

Finally, once the angular values for the dominating local minima are found, we use the corresponding ranging measurement in the noise-free data as extracted measurements.

APPENDIX II: IMPLEMENTATION OF THE DATA ASSOCIATION

The data association aims at matching the current extracted measurements with the landmarks that are being estimated in the EKF. For each measurement j and for each previously estimated landmark i, we compute the Mahalanobis distance (or normalized innovation square):

$$\rho^{i,j} = \left(\tilde{\mathbf{z}}^{\mathbf{i,j}}\right)^{\mathbf{T}} \cdot \left(\bar{\mathbf{Z}}^{\mathbf{i}}\right)^{\mathbf{-1}} \cdot \tilde{\mathbf{z}}^{\mathbf{i,j}}$$

where $\tilde{\mathbf{z}}^{\mathbf{i,j}} = \mathbf{z}^{\mathbf{j}} - \bar{\mathbf{z}}^{\mathbf{i}}$.

$\mathbf{z}^{\mathbf{j}}$ are the elements (ranging and angular data) of the measurement vector corresponding to measurement j,

$\bar{\mathbf{z}}^{\mathbf{i}}$ are the elements of $\bar{\mathbf{z}}$ corresponding to landmark i,

$\bar{\mathbf{z}} = \mathbf{h}_{\mathbf{LAS}}(\bar{\mathbf{x}})$ is the best projected estimate of the laser measurement vector obtained from

$\mathbf{h}_{\mathbf{LAS}}$, the non-linear laser observation matrix and

$\bar{\mathbf{x}}$, the predicted state estimate vector

The matrix $\bar{\mathbf{Z}}^{\mathbf{i}}$ is made of the elements of $\bar{\mathbf{Z}}$ corresponding to landmark i, with:

$$\bar{\mathbf{Z}} = \mathbf{H}\bar{\mathbf{P}}\mathbf{H}^{\mathbf{T}} + \mathbf{V}$$

where: $\mathbf{H}$ is the observation matrix

$\bar{\mathbf{P}}$ is the covariance matrix of $\bar{\mathbf{x}}$

$\mathbf{V}$ is the measurement noise matrix

The first stage of the association is a validation stage. The Mahalanobis distance $\rho^{i,j}$ follows a chi-square probability distribution, with two degrees of freedom. Therefore, the association is validated only if $\rho^{i,j}$ is lower than a predefined threshold corresponding to a desired confidence level.

If one measurement is validated for more than one landmark, it is associated with its nearest neighbor in terms of Mahalanobis distance. If multiple measurements are validated for a single landmark, we make the conservative choice of rejecting all of these observations. The measurements which have not been validated are interpreted as corresponding to new landmarks, which are given a new identification number.

Note from the expression of $\rho^{i,j}$ that the better the estimate of a landmark, the more robust its association with incoming measurements; indeed for better landmark state estimates, the validation threshold with respect to the physical distance between estimated landmark and current measurement ($\tilde{\mathbf{z}}^{\mathbf{i,j}}$) is tightened.

REFERENCES

[1] K. Bossert, "Reclaiming Kaho'olawe", *Geospatial Solutions*, Vol. 14, No. 9, September 2004.

[2] M. O'Connor, "Carrier-Phase Differential GPS for Automatic Control of Land Vehicles," *PhD Dissertation*, Stanford University, 1997

[3] J. A. Farrell, T. Han-Shue, and Y. Yunchun, "Carrier Phase GPS-Aided INS-Based Vehicle Lateral Control," *Journal of Dynamic Systems, Measurement, and Control*, Transactions of the ASME, Vol. 125, pp. 339-353, 2003.

[4] Greenspan, R., "GPS and Inertial Integration," *The Global Positioning System: Theory and Applications*, B. W. Parkinson, J. J. Spilker, Jr., P. Axelrad, P. Enge, editors, AIAA Progress in Aeronautics and Astronautics Volume 163-4, Washington, DC, 1996

[5] J. Leonard, and H. Durrant-Whyte, "Directed Sonar Sensing for Mobile Robot Navigation," *Kluwer Academic*, 1992.

[6] M. Joerger, and B. Pervan, "Autonomous Ground Vehicle Navigation Using Integrated GPS and Laser-scanner Measurements," *Proc. ION/IEEE PLANS*, April 2006.

[7] I. Tena Ruiz, Y. Petillot, D. M. Lane, and C. Salson, "Feature Extraction and Data Association for AUV Concurrent Mapping and Localisation," *Proc. IEEE-ICRA*, 2001

[8] S. Thrun, "Robotic Mapping: A Survey," *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel, February 2002.

[9] M. Joerger, J. Christ, R. Duncan, and B. Pervan, "Integrated Design of an AGV for Improved GPS-based Path-Following Performance," *International Journal of Vehicle Design*, Vol. 42, Nos. 3/4, 2006

[10] M. Joerger, J. Christ, R. Duncan, and B. Pervan, "Design of an AGV for Improved CDGPS-based Control Performance," *Proc. ION GNSS*, September 2004.

[11] D. Lawrence, "Aircraft Landing Using GPS: Development and Evaluation of a Real Time System for Kinematic Position using the Global Positioning System", *PhD Dissertation*, Stanford University, 1996

[12] A. E. Bryson and L. J. Henrikson, "Estimation Using Sampled Data Containing Sequentially Correlated Noise," *AIAA J. Spacecraft*, Vol. 5. No. 6, June 1968.

[13] A. E. Bryson, *Applied Linear Optimal Control*, Cambridge University Press, 2002.

[14] C. Ye, and J. Borenstein, "Characterization of a 2-D Laser Scanner for Mobile Robot Obstacle Negotiation," *Proc. IEEE-ICRA*, May 2002.

[15] R. Madhavan, H. Durrant-Whyte, and G. Dissanayake, "Natural Landmark-based Autonomous Navigation using Curvature Scale Space," *Proc. IEEE-ICRA*, 2002.

[16] F. Tang, M. Adams, J. Ibanez-Guzman and W. S. Wijesoma, "Pose Invariant, Robust Feature extraction From Range Data With a Modified Scale Space Approach," *Proc. IEEE-ICRA*, 2004.

[17] Y, Bar-Shalom, and T. E. Fortmann, "Tracking and Data Association," *Mathematics in Science and Engineering*, Vol. 179, Academic Press, 1988.

[18] G. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte and M. Csorba, "A Solution to the Simultaneous Localization and Map Building (SLAM) Problem," *IEEE Trans. On Robotics and automation*, Vol. 17, No. 3, June 2001.

[19] D. Maksarov, and H. Durrant-Whyte, "Mobile Vehicle Navigation in Unknown environments: a

multiple hypothesis approach," *IEEE Proc. Control Theory Appl.*, Vol. 142, No. 4, July 1995.