2013 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
AUTONOMOUS GROUND SYSTEMS (AGS) MINI-SYMPOSIUM
AUGUST 21-22, 2013 – TROY, MICHIGAN

Ground Penetrating Radar Based Localization

Byron Stanley
Matthew Cornick
Jeffrey Koechling
MIT Lincoln Laboratory
Lexington, MA

ABSTRACT

Localization refers to the process of estimating ones location (and often orientation) within an environment. Ground vehicle automation, which offers the potential for substantial safety and logistical benefits, requires accurate, robust localization. Current localization solutions, including GPS/INS, LIDAR, and image registration, are all inherently limited in adverse conditions. This paper presents a method of localization that is robust to most conditions that hinder existing techniques. MIT Lincoln Laboratory has developed a new class of ground penetrating radar (GPR) with a novel antenna array design that allows mapping of the subsurface domain for the purpose of localization. A vehicle driving through the mapped area uses a novel real-time correlation-based registration algorithm to estimate the location and orientation of the vehicle with respect to the subsurface map. A demonstration system has achieved localization accuracy of 2 cm. We also discuss tracking results for the first autonomous vehicle to use this technology and the potential for miniaturized general use systems.

INTRODUCTION

Among the motivations for automating vehicles is the expectation that by either removing human error (in the case of the 371,104 traffic deaths in the US in the previous decade [1]) or removing people from hazardous situations (in the case of many of the 8,148 coalition soldier deaths in OIF and OEF to date) a substantial number of lives could be saved.

Useful automated vehicles must know their location accurately enough to remain within a travel lane. Sensor systems for localization include GPS/INS units [2], 3D LIDAR scan matching, and image registration based scan matching. Even high performance GPS/INS units in open sky conditions are typically unable to localize well enough to keep a vehicle within a lane. Performance is worse when GPS reception is limited by jamming, multipath effects, or signal blockage from trees, buildings, or tunnels. Matching 3D LIDAR maps of the road surface distance, direction, and intensity [3] works well in fair conditions, but fails when the surfaces are obscured by obstacles like people or cars, or by atmospheric conditions like snow, heavy rain, ice, dust, fog, or smoke [4]. Image registration suffers from similar obscuration problems.

Given the limitations of existing sensors, it is natural to consider ways to augment them. Typical sensors make use of the environment by looking outwards or by trilateration using references in order to localize themselves. Downwards observation to date has been through road surface mapping and estimation. Each of these modes of sensing is affected by the complex and dynamic environment that they must handle in order to perform localization. Mapping and localization using the subsurface is a strong candidate given the relatively static nature of the ground, road, and other subsurface materials.

Ground Penetrating Radar (GPR) is widely used for geophysical investigation, civil engineering, archeology, forensics, humanitarian demining, mining, and space exploration (GPR profiles of the Martian subsurface have been collected). A key reason for its widespread use is its ability to detect objects of almost any material composition. In contrast, other

This work is sponsored by the Assistant Secretary of Defense for Research & Engineering under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government. Approved for Public Release.

subsurface sensing techniques such as pulse induction (metal detectors) and resistivity probes are best suited to detecting metal objects.

While soil and road materials are opaque to visible light, they are semi-transparent to radio waves. GPR systems send radio frequency electromagnetic radiation into the ground and measure reflections from scattering below the surface. A pulse is transmitted from a transmitting (Tx) antenna on or above the surface of the ground and propagates downward through the subsurface. When the pulse encounters an object, some of the energy reflects back up towards the surface, where it is detected by a receiving (Rx) antenna located near the transmitting antenna. Any interface between media of different electrical properties causes a reflection. As a result, spatial variations in soil properties or moisture content generate reflections visible to GPR, as do air voids, pipes, roots, and rocks.

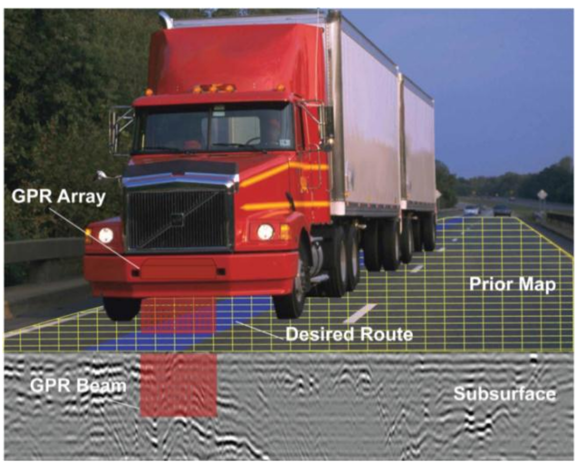

Figure 1: Conceptual visualization of localization using Ground Penetrating Radar.

The radar receiver collects a voltage profile of reflection strength at each instance of time. The delay between the transmitted pulse and the received reflection allows estimation of the depth of the feature from which the reflection originated. The received voltage profile from a single Tx pulse can be interpreted as a depth plot of objects in the subsurface environment. Using several Tx/Rx channels in a linear array allows one to construct a two dimensional image of the subsurface for a particular array location (along-array dimension and depth), and continuous pulses of an array while in motion allow one to construct a three-dimensional "tube" of data representing the subsurface reflections along a traversed path with a volume given by the travel distance, array width, and depth of penetration.

GPR data paints a fairly complete picture of the subsurface environment. Every discrete object and soil feature is captured, provided it is not significantly smaller than a wavelength and it has contrast with the surrounding soil. The premise of GPR localization is that these subsurface features, as represented in GPR data, are sufficiently unique and sufficiently static to permit their use as an 'identifier' of the precise location relative to previous collections. Our concept involves an initial collection of GPR data over a region. This first data collection forms a subsurface database referred to as the baseline. The baseline includes referenced detailed GPS or other global position data. On subsequent traversals of the same region, GPR data is matched against the baseline to determine the precise location relative to the baseline. Details of our registration process are discussed below. Registration results in a highest-likelihood location and orientation within a search volume dictated by GPS/INS uncertainty. Thus the pairing of GPR and GPS/INS allows a very precise global location estimate. It is important to note that the global accuracy of this approach is limited by the accuracy of the baseline map.

DESIGN

Radar Design

Our demonstration system is a Stepped Frequency Continuous Wave (SFCW) radar, the core components of which can be seen in Figure 2. A SFCW radar transmits a waveform consisting of a sequence of tones at discrete frequencies spanning a wide bandwidth. Only one frequency is transmitted at a time. The waveform begins at a low frequency and steps up to higher frequencies. When the maximum frequency is reached the waveform jumps back to the lowest frequency to begin the ramp again. There is no significant delay between ramps, so the duty factor of a SFCW radar is nearly 100%. The receiver is open at all times during the transmit pulse, and hence the strongest signal received is the direct coupling between the Tx and Rx antennas, which is compensated by a using a calibration pathway. Each tone is generated by a direct digital synthesizer (DDS) and fed to the Tx antenna element of the array. For each tone, the Rx antenna picks up the signal, and the super heterodyne receiver multiplies it by the output of an

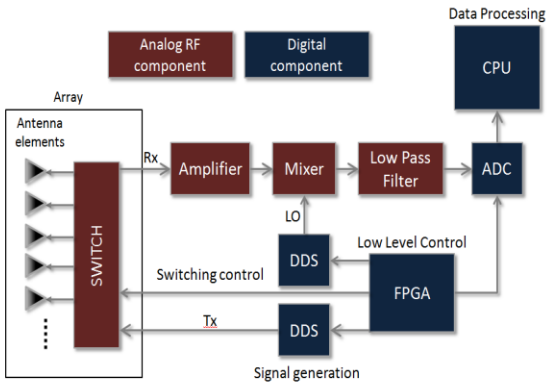

Figure 2: Block diagram of a stepped frequency continuous wave (SFCW) ground penetrating radar

onboard local oscillator to convert it to a lower intermediate frequency (IF). The signal at the IF is measured with an analog to digital converter (ADC) and sent to a CPU for processing, as seen in Figure 2. Our demonstration design uses frequencies in the 100 MHz to 400 MHz range. A modern dual-core Intel processor is used for all processing.

Array Design

A single Tx/Rx antenna pair can provide localization in the along-track direction. However, any lateral offset between the current data collection and the baseline erodes matching of current scans to the baseline. An offset as small as 30 cm could render single channel localization useless due to changes in the GPR data as the array moves off track. A linear array of multiple Tx/Rx pairs perpendicular to the direction of travel gives a full swath of recorded ground data, allowing for lateral localization, as well as providing a unique method for tolerating lateral offsets. As long as a sufficient fraction of the array overlaps the baseline, the GPR sweep matching can use the overlapping fraction of the total array.

It is important to note that the array design differs from traditional GPR systems so as to allow localization to be achieved. Two key modifications were fundamental to the design. The spacing between the elements is approximately one tenth of a center frequency wavelength, much closer than is typical for GPR systems. In addition, the elements and array are designed so that every element has identical field patterns to allow re-traversal using non-identical paths in which offset or misalignment is present. This element similarity requirement is especially difficult for our close element spacing, which results in significant mutual coupling and leads to array end effects. Further details on the array design are discussed in detail in [5].

Figure 3: Ground Penetrating Radar Array

Our demonstration system uses a 24-element array of dipole antenna elements. Each dipole element is constructed of resistively loaded patches printed on a PCB. The use of resistive loading is common in GPR antenna design, as it provides greater bandwidth (at the expense of efficiency). These elements are placed at 12.7 cm spacing within a metal box cavity having dimensions 3 m x 0.6 m x 0.3 m. This 12.7 cm spacing is approximately a tenth of a wavelength of our 250 MHz center frequency. This resolution is finer than typically seen in GPR arrays and is driven by a desire to allow for high fidelity matching to baseline data, especially in situations in which lateral offsets from the baseline are present. In the case of lateral offset, the elements from the current pass that overlap the baseline will be within 6.4 cm of some element from the baseline pass. The array includes a matrix switch that routes the Tx and Rx circuitry to a pair of adjacent antenna elements. There are 23 possible Tx/Rx pairs in the 24 element array.

Localization Algorithm Design

The goal of GPR localization is to correlate data from the baseline pass with current GPR data, and thus be able to infer the location of the array to high precision.

Baseline data is periodically fetched from a local data server for matching in real time when the vehicle is in motion. A 50 m x 50 m x 60 ns grid of baseline data is maintained in memory for matching. This grid is three-dimensional, with easting, northing, and radar range axes. When the vehicle nears the edge of the grid, the algorithm creates a new grid with baseline data from the server, centered on the expected vehicle position. In this way a local grid of baseline data is always maintained. The matching of a current data scan to the baseline grid is described below:

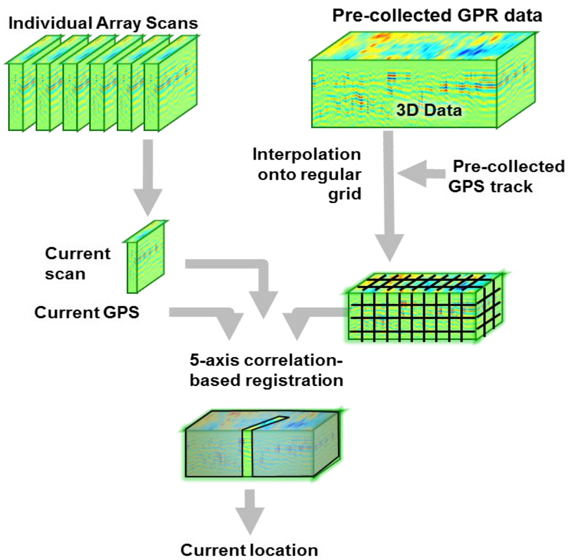

Figure 4: Processing chain of GPR localization system.

• A GPS/INS utilizing the OmniSTAR VBS service (0.5 m CEP) provides relatively accurate initial conditions for placing the sweep into the 3D grid of baseline data. The initial location of this sweep in the grid is determined by the latitude, longitude, heading, and roll of the array, as computed by the INS.

• A search region surrounding the initial location estimate is filled with a swarm of "particles" representing candidate array locations and orientations. Each particle represents the five variables easting, northing, height, roll, and heading. The particle poses are iterated so as to search the region to find the maximum correlation within the five-dimensional space. Each particle in the swarm exhibits inertia and experiences a force toward or away from other particles in the swarm based on their current correlations. In addition, each particle remembers the highest correlating location it has seen during a single registration, and experiences an additional force toward that point. Particle swarm optimization is well suited to optimization in cases where several local optima may be present, and is thus very robust [6].

• After 20 to 200 iterations the best fit particle is chosen. This provides an accurate estimate of the current array location and orientation.

• The search region is updated and either expanded or shrunk to reflect the new uncertainty in location based on the goodness of fit of the last estimate, as well as INS drift. The number of particles and iterations grows as the volume of the search region grows. The relationship between the number of particles, number of iterations, and the search volume is determined empirically.

Autonomous Vehicle Design

Drive-by-wire operation of a truck was achieved using a Kairos Autonomi appliqué kit for low level servo operation of the steering, throttle, and brake systems, though typical testing used the steering system only. All midlevel and high level design and implementation of the MIT Lincoln Laboratory autopilot control architecture was based on the Robot Operating System (ROS) [7] for communications and partially derived from the MIT DARPA Urban Challenge vehicle architecture [8]. The algorithms and nodes were written and implemented by MIT Lincoln Laboratory.

We implemented a steering control method based

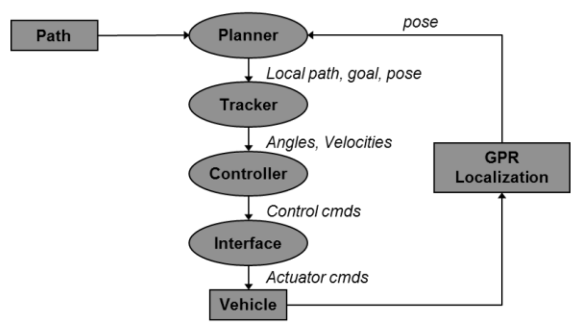

Figure 5: Autopilot navigation stack

on a novel derivation of the pure pursuit algorithm, which will be detailed in future publications along with detailed description of the autopilot architecture. The navigation stack is shown in Figure 5. Real-time corrected latitude and longitude data was streamed at approximately 67 Hz from the registration algorithm to the autopilot planning system. The planner created a local route at a 1 Hz rate. The autopilot system used the pure pursuit algorithm to calculate the appropriate steering angles to track the desired trajectory. The tracker calculated the pure pursuit algorithm at 40 Hz to track to the local route. The controller calculated the steering positions required to hold the pure pursuit arc. The interface controller then translated the commands into servo commands for the drive-by-wire system. Loss of correlation from the GPR localization system was handled through integration of the INS velocities and eventual default to the GPS/INS solution.

DATA AND DISCUSSION

Testing was conducted at several isolated locations, in primarily good weather over dirt, gravel, and asphalt roads. A typical test track, as shown in the figure below, was a few kilometers in length.

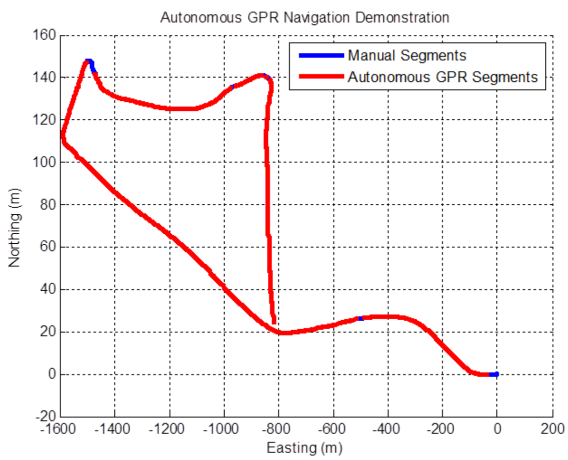

Figure 6: Autonomous GPR localization demonstration overhead view.

In some short segments, the vehicle was taken out of autopilot out of caution for a few areas where there was limited ground clearance for the array. As of the last review, the system had logged over 1279 km while using GPR localization, over 465 km of which were in autonomous operation. Operation was typically performed at between 2.2 m/s and 4.5 m/s (5 to 10 mph) due to vehicular rather than sensor limitations.

Localization Accuracy

Accuracy of the system, using GPR localization, was demonstrated using an Oxford Solutions RTK GPS ground station, two Oxford Solutions RT3000 GPS units, and the new GPR system. The DGPS system used OmniSTAR differential corrections to provide realistic performance in a clear road situation (as shown below).

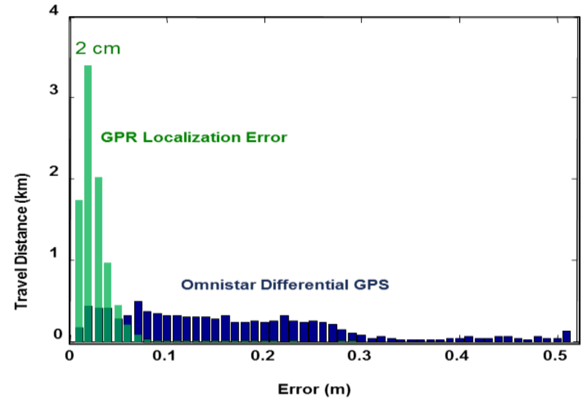

Figure 7: Position error histogram based on RTK data for truth, OmniSTAR differential corrections for DGPS raw position error, and GPR based localization.

An open, approximately 9 km course was driven in good conditions with the GPR system and the GPS units. As this calibrated test was prior to vehicle automation and real-time operation, the GPR data was collected, then post-processed using the localization registration algorithm. The results were then binned for the above histogram.

The 0.02 m RMS accuracy of the GPR localization registration algorithm was a substantial improvement over the approximately 0.2 m RMS performance of the DGPS system. The performance of the DGPS would be expected to degrade with a lower tier of differential GPS service, or in areas of GPS denial, obscuration, jamming, or multipath.

Autonomous Performance

Autonomous operation involved loading up the prior data and route, then engaging the autonomous pure pursuit tracking algorithm. Operation in autonomous mode typically tracked to 0.2 m RMS or less compared to the desired GPS baseline route. This is primarily a measure of the tracking performance of the vehicle and pure pursuit algorithm, but includes all large or small jumps in tracking position due to, for instance, loss of lock on position. The tracking performance of the system to the baseline for the route shown earlier in Figure 6 is shown in Figure 8 above. The tracking error is 0.12 m RMS for this route, which shows that many times it tracks well below the 0.2 m RMS measurement (which includes starting offset and post-route offset in the overall error calculation).

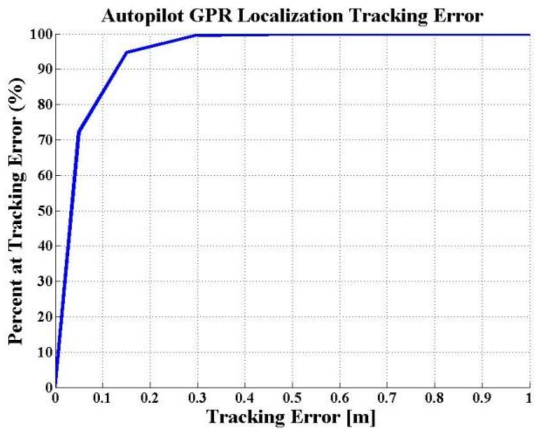

Figure 8: Autonomous vehicle GPR localization based route tracking performance.

We have shown that real-time navigation is possible using the GPR localization method and novel GPR design. Additional performance increases could be derived by improving the system steering model, steering state measurement errors, and further tuning the pure pursuit parameters.

Limitations and Unknowns

Most of the testing conducted by the demonstration system occurred on asphalt, dirt, and gravel roads in good weather. Hence, there are many unknowns that pertain to the use of GPR localization in less ideal environments. In particular, we performed limited testing on roads containing rebar, a common addition to concrete roadways, bridges, and tunnels. From the limited tests that we have conducted over bridges, correlation has been high and localization accurate. In addition, we did not perform significant testing of the effects of precipitation. We had two opportunities to test the effects of rainfall totaling approximately a quarter inch. In both cases the baseline was taken over dry ground, while the matching pass was taken after a storm and thus the soil was damp and puddles had formed. In both cases only minor effects on localization were observed. Still, it is expected that larger rain-fall totals may adversely affect GPR localization, and further testing of the demonstration system is required in these conditions.

The demonstration system detailed above requires further improvements in the following areas before it is considered viable for widespread use:

- Size
- Weight
- Power (FCC requirements)
- Under vehicle mounting
- High speed operation
- GPS dependence
- Cost

Miniaturized Design

Simulation and dataset extrapolation indicate that it is feasible to design a miniaturized version of the system that addresses the above weaknesses. Through optimization of the power, signal to noise ratio, sampling rate, vehicle speed, and materials, we were able to form a rough design that would allow operation underneath a typical passenger vehicle.

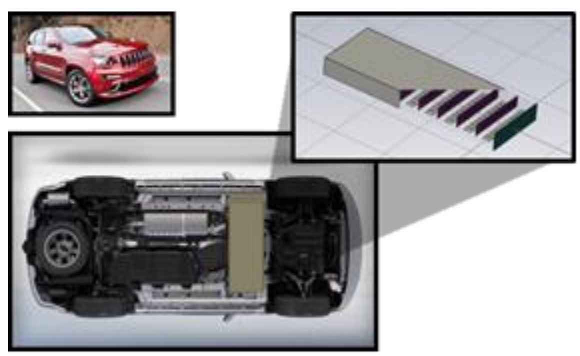

Figure 9: Miniaturized GPR Design

Further engineering will be required to detail, build, test, and optimize the miniaturized design. While not required for operation, additional research associated with modification of the road surface (such as adding reflective material or patterns to the asphalt or other layers during repair or replacement) could potentially lead to performance improvements and size reductions.

CONCLUSIONS

This paper shows that GPR localization is a viable method of localization for vehicles or systems in the ground domain. The localization performance is substantially improved over DGPS (OmniSTAR satellite service) and, with further verification required, will likely perform well in conditions that GPS/INS, LIDAR, and image registration systems would fail. Fusion of the systems with a tightly coupled or loosely coupled system could lead to substantially safer and more robust localization.

Additionally, we have demonstrated the first autonomous vehicle to use GPR localization, as well as real-time operation of the GPR localization algorithm. The localization method and similar GPR hardware could potentially apply to other areas in the industrial, commercial, and military sectors. There is reason to believe that this technology could be used for not only autonomous vehicles like patrol systems, but underground operation (such as mining), farming, infrastructure and bridge inspection, surveying, indoor operation and even including cases using the side walls, ceilings, or other surfaces as a reference rather than the ground.

Given the potential implications for vehicle localization, even in GPS-denied environments, we believe that substantial further study of this area and its potentials is warranted.

REFERENCES

[1] U.S. Census Bureau, " Statistical Abstract of the United States:," 2012. [Online]. Available: http://www.census.gov/compendia/statab/2012/tables/12s1103.pdf.

[2] E. Nebot, S. Sukkarieh and H. Durrant-Whyte, "Inertial Navigation aided with GPS information," in *Fourth Annual Conference on Mechatronics and Machine Vision in Practice*, Toowoomba, Qld., 1997.

[3] J. Levinson and S. Thrun, "Robust vehicle Localization in Urban Environments Using Probabilistic Maps," in *IEEE International Conference on Robotics and Automation*, Anchorage, Alaska, 2010.

[4] B. Yamauchi, "All-Weather Perception for Man-Portable Robots Using Ultra-Wideband Radar," in *IEEE International Conference on Robotics and Automation*, Anchorage, 2010.

[5] A. Fenn, P. Hurst, J.Pacheco, M. Cornick and L. Parad, "Ultrawideband Cavity-Backed Resistively Loaded Planar Dipole Array for Ground Penetrating Radar," in *IEEE International Symposium on Phased Array Systems and Technology*, Waltham, 2013.

[6] G. Venter and J. Sobieszczanski-Sobieski, "Particle Swarm Optimization," *AIAA Journal*, vol. 41, no. 8, pp. 1583-1589, 2003.

[7] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler and A. Ng, "ROS: an open-source Robot Operating System," in *Proc. Open-Source Software workshop of the International Conference on Robotics and Automation (ICRA)*, Kobe, 2009.

[8] J. Leonard, J. How and S. Tellar, "A Perception Driven Autonomous Urban Vehicle. Journal of Field Robotics," *Journal of Field Robotics*, pp. 727-774, 2008.