

Localizing Ground Penetrating RADAR: A Step Toward Robust Autonomous Ground Vehicle Localization

Matthew Cornick, Jeffrey Koechling, Byron Stanley, and Beijia Zhang

MIT Lincoln Laboratory, 244 Wood St. Lexington, Massachusetts 02420

Received 21 May 2014; accepted 27 April 2015

Autonomous ground vehicles navigating on road networks require robust and accurate localization over long-term operation and in a wide range of adverse weather and environmental conditions. GPS/INS (inertial navigation system) solutions, which are insufficient alone to maintain a vehicle within a lane, can fail because of significant radio frequency noise or jamming, tall buildings, trees, and other blockage or multipath scenarios. LIDAR and camera map-based vehicle localization can fail when optical features become obscured, such as with snow or dust, or with changes to gravel or dirt road surfaces. Localizing ground penetrating radar (LGPR) is a new mode of a priori map-based vehicle localization designed to complement existing approaches with a low sensitivity to failure modes of LIDAR, camera, and GPS/INS sensors due to its low-frequency RF energy, which couples deep into the ground. Most subsurface features detected are inherently stable over time. Significant research, discussed herein, remains to prove general utility. We have developed a novel low-profile ultra-low power LGPR system and demonstrated real-time operation underneath a passenger vehicle. A correlation maximizing optimization technique was developed to allow real-time localization at 126 Hz. Here we present the detailed design and results from highway testing, which uses a simple heuristic for fusing LGPR estimates with a GPS/INS system. Cross-track localization accuracies of 4.3 cm RMS relative to a “truth” RTK GPS/INS unit at speeds up to 100 km/h (60 mph) are demonstrated. These results, if generalizable, introduce a widely scalable real-time localization method with cross-track accuracy as good as or better than current localization methods. © 2015 Wiley Periodicals, Inc.

1. INTRODUCTION

Self-driving vehicles and driver-assist systems have been pursued on a worldwide basis. One main objective is to reduce the yearly vehicle accident fatalities (32 K US [National Highway Traffic Safety Administration, 2011] and an estimated 1.24 M worldwide [World Health Organization, 2013]). Effective GPS/INS (internal navigation system), LIDAR, and camera-based autonomous navigation techniques were developed and honed in the DARPA Grand and Urban Challenges, though operation was in a carefully staged and mapped environment (Buehler, Iagnemma, & Singh, 2009). However, GPS/INS approaches, even aided by wheel odometry, typically have real-time 2σ values well over 1 m (Kennedy & Rossi, 2008) which is insufficient to maintain a vehicle in a travel lane on most roads. Failure modes for GPS-dependent solutions include blockage and

multipath, such as arise in urban and heavily forested or mountainous environments.

LIDAR-based algorithms fused with GPS/INS, wheel odometry, and cameras offer a very successful approach to localization. One of the most successful approaches was the modification of LIDAR mapping algorithms to use surface intensity probabilistic maps (Levinson, Montemerlo, & Thrun, 2007; Levinson & Thrun, 2010). Notably, using this method, the GPS/INS solution was improved to approximately 9 cm in urban environments and demonstrated in traffic and during rainfall. The surface intensity probabilistic map approach breaks down when the LIDAR beam is significantly attenuated or blocked, such as occurs in snow, fog (Yamauchi, 2010), dust, or with dirt on the lens. In addition, changes to the road surface, such as would be expected on dirt roads or after repainting, may require updating of the map. Camera-based approaches continue to be refined, with active vehicle localization approaches such as topometric localization (Badino, Huber, & Kanade, 2012), FAB-MAP (Cummins & Newman, 2011), and linear mosaicing (Unnikrishnan & Kelly, 2002). These approaches also have similar limitations to the LIDAR systems because of their use of optics and their operation in dynamic environments.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited. This work is sponsored by the Department of the Army and 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 authors and are not necessarily endorsed by the United States Government.
Direct correspondence to: Byron Stanley, e-mail: stanley@ll.mit.edu

Self-driving vehicles must be robust to environmental conditions and related failures in order to be broadly useful and live up to their potential. Milford and Wyeth (Milford & Wyeth, 2012) sought robustness in camera-based localization by identifying sequences of matches rather than single feature matches. Nuske et al. (Nuske, Roberts, & Wyeth, 2009) used a multihypothesis particle filter to select among matches to three-dimensional (3D) edges in the environment. Brunner, Peynot, Vidal-Calleja, and Underwood (2013) augmented visual sensing with thermal sensing to do localization in the face of obscuration and darkness.

In this paper, we detail localizing ground penetrating radar (LGPR), a form of ground penetrating radar specifically designed to enable a priori map-based localization. LGPR offers complementary capabilities to traditional optics-based approaches to map-based localization, including the ability to transmit through air-based obscuration, such as fog and dust, and common surface obscuration, such as dirt and snow (Abe, Yamaguchi, & Sengoku, 1990; Hoekstra & Delaney, 1974). Hence, LGPR offers increased resistance to common failure modes of existing localization techniques, which potentially allows for significant improvements in robustness when fusing it with those existing approaches. In addition, LGPR senses a generally stable environment, discussed in Section 4.1, which complements the dynamics of surface environments. Many research challenges and risks, discussed in Section 4, remain to be addressed, including all-weather operation, data storage requirements, and array size reduction.

We developed an early stage version of the LGPR system, shown in Figure 2, which has been used to automate the steering of multiple armored vehicles and was tested in three US states prior to several months of operation in Afghanistan in 2013. This large-scale system, which operated at 7.5–15 km/h, was tested on several soil types, on and off road, and demonstrated over thousands of kilometers of operation.

While the rough concept of localizing GPR has been introduced in limited detail (Stanley, Cornick, & Koechling, 2013), here we present in sufficient detail as to allow peer verification, the first practical (small size, low power) design for an LGPR system on a commercial passenger vehicle, with the capability of achieving high-speed operation (at least 100 km/hr). In addition, we, for the first time, characterize the real-time localization accuracy, long-distance highway operation, and high-speed performance of an LGPR system.

Previous work has combined robotics with GPR. Herman (1997), Williams (2012), and Lever et al. (2013) all used autonomous systems to move GPR systems, enabling them to create subsurface maps. They did not, however, attempt to localize the robots using those maps.

We have integrated the LGPR sensor onto a 2000 Chevrolet Silverado truck (see Figure 3), which we also

equipped with two Oxford Technical Solutions (OxTS) RT3003 GPS/INS systems: one is a real-time kinematic (RTK) system solely used for truth reference while the other receives WAAS differential corrections and is loosely coupled with the LGPR system to provide the LGPR position estimates for both mapping and localization. The RTK truth reference unit is coupled with a local base station that allows local 2-cm accuracy location measurements. The RT3003 uses a MEMS-based IMU and dual GPS antennas to produce a 6-DOF (degrees of freedom) Kalman filter-based pose output with 1σ heading error of approximately 0.1 degrees. The LGPR array is mounted underneath the vehicle behind the front wheels. We use spacers underneath the chassis to fix the array height at 15 cm (6 in) above the ground, the initial design point. In general, the performance of GPR systems improves monotonically as the array is lowered to the ground surface, but this must be balanced with the need for ground clearance.

We first discuss designing an under-vehicle-mounted LGPR system, including key design parameters. We then move on to the algorithms supporting mapping and real-time localization before discussing the experimental results from high-speed highway testing. We conclude by discussing remaining concept risks that should be addressed in future research.

2. LOCALIZING GROUND PENETRATING RADAR SYSTEM

The localizing ground penetrating radar (LGPR) system, designed and developed by MIT Lincoln Laboratory, consists of both hardware and processing components. The hardware component is a uniquely designed type of ground penetrating radar that allows high cross-track resolution to accurately detect subsurface features and has specially designed uniform elements to allow comparison of measurements even if the array element overlaps a different portion of the map than during its original pass. The key parameters of the LGPR system are shown in Table I.

The LGPR operates at low frequencies to allow deep ground penetration and to reduce the amount of small clutter in the image. Details, including the unusually low-radiated power, are covered in Section 2.1.

The processing component is made up of the mapping and registration components. The unique aspect of the mapping and registration components is that they are streamlined so that they are all accomplished automatically in real-time at 126 Hz on a consumer-grade dual core processor. The key components of the current LGPR processing subsystem are shown in Table II.

The mapping and registration components are detailed in Section 2.2.




Figure 1. The LGPR array is shown mounted under the vehicle in this concept drawing. Radio frequency (RF) signals bounce off of underground features to localize a vehicle using a prior map of the subsurface.



Figure 2. An early-stage version of LGPR (out front of the armored vehicle). This system steered itself based on real-time localization from the LGPR system. Operation was shown over a range of soils and road types during thousands of kilometers of operation in the United States and Afghanistan.

2.1. Localizing GPR Sensor

In the area of subsurface sensing, ground penetrating radar (GPR) is one of the most versatile and prolific sensing modalities today. All soil and most road materials are semi-transparent to radio waves. GPR systems work by sending a pulse of electromagnetic radiation into the ground and measuring reflections that originate from scattering below the surface. Reflections occur at the interface between

objects of different electromagnetic properties; for example, the interface between soil and pipes, roots, or rocks. However, it is not these discrete objects but rather the inherent inhomogeneity in subterranean geology that often dominate GPR reflection profiles. This can be seen in Figure 4, in which soil layers and variations in moisture content cause extended reflections in the data. GPR data paints a fairly complete picture of the subsurface environment. Nearly



Figure 3. The low-profile LGPR system is mounted underneath the experimental vehicle. GPS/INS system antennas for the RTK truth and differential units are mounted on the outermost corners of the vehicle chassis.

Table I. Key LGPR RADAR parameters.

Key Parameter	Value
Radar type	Stepped frequency continuous wave
Frequency range	100–400 MHz
Frequency spacing	51 tones spaced by 6 MHz
Array dimensions	152 cm × 61 cm × 7.6 cm (5.0 ft × 2.0 ft × 3.0 in)
Array offset from ground	15.24 cm (6 in.)
Number of antenna elements in array	12
Number of channels (number of pairs of elements)	11
Total radiated power	40 μ W continuous (one element at a time)
Leakage power (radiated above-ground)	4 μ W
Radar sweep rate (all 11 channels)	126 Hz
Depth of penetration	2 m–3 m (New England soil)
Radar range resolution	20 cm–30 cm

every discrete object and soil feature is captured, provided that it is not significantly smaller than a wavelength and that it has sufficient dielectric contrast with the surrounding soil. The premise of GPR localization is that these subsurface features, as represented in GPR data, are sufficiently

Table II. Key LGPR processing parameters.

Key Parameter	Value
Real-time localization rate	126 Hz
Type of RADAR Data Filter	High pass
Type of registration algorithm	Heuristic search, maximizing correlation over 5 DOF
Correlation threshold (integrated velocities used below this value rather than GPR solution)	0.9 (out of range –1 to 1)
Overlap threshold (integrated velocities used below this value rather than GPR solution)	2 elements (out of 11)

unique and static to permit their use as identifiers of the precise location where they were collected.

2.1.1. Introduction to GPR

A general guideline is that the maximum detection depth of a GPR will often be three to four skin depths, where skin depth is a measure of the depth to which the pulse can propagate before losing most of its energy (specifically $1/e$ lower in field values thus ~ 8 dB lower in power). Skin depth




Figure 4. Example cross section of GPR data collected by one transmit receive channel of the LGPR, showing subterranean features. A high-pass IIR spatial filter has been applied and thus vehicle reflections and the initial ground bounce at depth = 0 have been filtered out (see Section 2.1.2 for details of this process). Along-track distance is the distance in the direction of travel of the vehicle.

is determined by soil losses caused by Joule heating and dipole losses. High-conductivity soils, such as those with high moisture and salinity, have smaller skin depths (Jol, 2008). The range resolution of a GPR specifies the resolving power in depth and is approximately

$$\frac{c}{2 \cdot BW \cdot \sqrt{\epsilon_r}}, \quad (1)$$

where BW is the effective bandwidth, c is the vacuum speed of light, and ϵ_r is the soil's relative permeability (typically in the range of 5 to 15). The lateral (along-track and across-track) resolution achieved by a GPR system is dependent on the physical beam width of the antenna (in units of area), which increases with depth but in general is not better than the range resolution. Further information on GPR theory can be found in (Daniels, 2004).

Traditional GPR systems for road inspection are often centered in the 1 GHz to 3 GHz band (Saarenketo & Scullion, 2000), with nearly 100% bandwidth ($BW = 1 - 3$ GHz), which provides excellent resolution (2 cm – 5 cm) at the expense of penetration depth. One of our key findings is that such high resolutions can actually be detrimental to the task of localization, as it increases the fragility of the map correlation process, for several reasons. First, high-resolution features in GPR data become increasingly difficult to correlate when pass-to-pass offsets are present. For example, at 1 GHz, radar data decorrelates at distances around 2–3 cm, resulting in a very fine requirement for antenna element spacing. In addition, the phase differences resulting from vehicle motion (pitch and roll) are significant and can lead to difficulties in correlating coherent GPR data to the map, as each vehicle pass will have slightly different suspension trajectories. High frequencies also suffer from being too sensitive to small changes, such as trash on the road surface, and are more susceptible to phase errors due




Figure 5. Low-profile LGPR array, switch matrix, and processing chassis. Array as built is 152.4 cm × 61 cm × 7.6 cm (5' × 2' × 3").

to thermal drifts in the transmitter/receiver. For all of these reasons, the GPR data correlation process becomes easier at lower frequencies, which has the benefit of improving penetration depths as well. The only factor preventing one from lowering the frequency indefinitely is that the radar cross section of the most important subsurface geology tends to drop off steeply below 100 MHz as well, and the required antenna sizes grow rapidly below 100 MHz. For this reason, we have selected the frequency range 100 MHz to 400 MHz as best suited to the task of localization. This frequency range is capable of resolving large subsurface geology on the scale of 20 cm to 30 cm, while remaining robust to the sensitivities of high-frequency systems mentioned above. We note that the variation in range resolution is due to variation in ϵ_r , primarily driven by moisture content (dry soil will typically be in the range $\epsilon_r = 4$ to 6, whereas fully saturated soils are closer to $\epsilon_r = 25$ to 36). In the frequency range where our system operates (100 MHz to 400 MHz), soils have skin depths that range from $D = 10$ cm to 100 cm, depending on soil composition and moisture content (in New England soils, for example, $D \sim 100$ cm skin depths are common, leading to 2–3 m penetration depths).

2.1.2. System Components

It is important to note that our GPR design differs from traditional GPR systems to allow localization to be achieved. The LGPR consists of four basic functional components: a unique antenna array, a 2×12 switch matrix, a custom VHF stepped frequency continuous wave (SFCW) GPR, and one single-board computer (SBC). These components are shown in Figure 5 (the radar electronics and SBC are within the chassis shown). The switch matrix switches the individual transmit and receive channels of the radar to each of the array elements (Figure 6). Data sent to the SBC are processed using standard SFCW radar techniques to generate data as seen in Figure 4.

One key difference between the LGPR array and traditional GPR array designs is the spacing between the elements (12.7 cm), which is approximately one tenth of a center frequency wavelength. This resolution is finer than

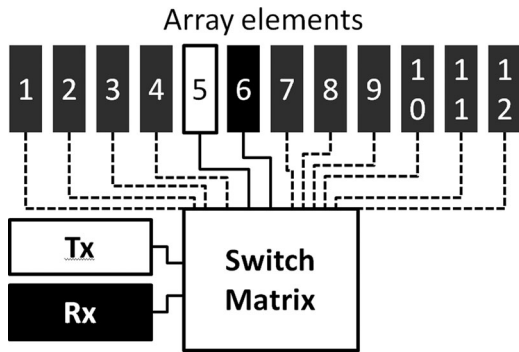


Figure 6. Layout of the GPR transmit (Tx) module, receive (Rx) module, switch matrix, and array antenna elements. Our GPR transmits on only one element at a time (example shown in white) while receiving on the adjacent element (example shown in black). The transmit/receive pair starts with 1/2 and sweeps to 11/12, creating 11 channels of data (5/6, the fifth channel, is shown in the example above).

typically seen in GPR arrays and is driven by a desire to allow for high-fidelity matching to baseline data. In addition, the elements and array cavity are designed so that every element has identical near-field (and thus far-field) patterns. This is required to allow path retraversal in which pass-to-pass offset or misalignment is present. This element similarity requirement is especially difficult to meet for our close element spacing (relative to wavelength), which, in traditional GPR arrays, ordinarily results in significant mutual coupling and array end effects.

An array with these characteristics, shown in Figure 2, has been developed previously at MIT Lincoln Laboratory (Fenn, Hurst, Pacheco, Cornick, & Parad, 2013) having favorable radiation characteristics for localization. That design was impractically large for consumer-vehicle applications and was demonstrated at speeds below 15 mph. After performing a study and simulating expected performance, we designed a miniature version, which consists of 12 dipole elements linearly aligned within a reflective rectangular metal cavity having dimensions 152.4 cm × 61 cm × 7.6 cm (Figure 5).

Other than this change in cavity size and number of elements, the design is identical to the larger array. This array is capable of mounting rigidly under the vehicle body. Care was taken in the array design to minimize the effects of the large metal vehicle chassis, which has the potential to alter the radiation patterns of the elements. The cavity, as designed, is able to reduce the impact of the vehicle chassis on radiation patterns of every element.

2.1.3. Signals

The transmit/receive data collection sequence, called a “sweep” pattern, begins by transmitting a pulse from the

first antenna element and receiving on the second. Once the first pulse is complete we transmit on the second and receive on the third and so on to the last pulse in which we transmit on the 11th and receive on the 12th. Hence, the 12-element array produces 11 channels of data. This sweep takes approximately 7.8 ms, so the vehicle advances several centimeters during the whole sweep (at highway speeds), which is sufficient to limit the effect of blurring the data due to motion. Each element “pulse” is formed by a sequence of 51 tones evenly spaced from 100 MHz to 400 MHz. The radiated power is extremely weak, with less than 40 microwatts (peak) radiated in total, and less than 4 microwatts (peak) leaked into the surrounding air (the majority is absorbed by the ground). This radiated peak power is significantly (1000x) less than traditional pulsed GPR designs and is made possible by the power efficiencies of the SFCW design approach (in which the radar transmits and receives continuously). This extremely low-power design alleviates spectrum conflict issues faced by traditional GPR designs.

A direct digital synthesizer (DDS) generates each tone and simultaneously generates a local oscillator (LO) signal at the sum of the tone frequency and the intermediate frequency (IF). The super heterodyne receiver multiplies the received radar signal with the local oscillator and filters the result to yield a signal at the intermediate frequency with the same amplitude and phase as the received radar signal. An analog-to-digital converter (ADC) samples the intermediate frequency signal four times per intermediate frequency period. Each tone lasts for six intermediate frequency periods. We skip the first and last intermediate frequency periods. For the remaining four intermediate frequency periods, we compute and average the complex Fourier parameter, or S-parameter, for that tone.

Because the environment’s ambient temperature may be significantly different from one pass to the next (i.e., the map is created on a hot summer day and then subsequently used through the winter), we have incorporated an automatic calibration channel. The purpose of this calibration path is to frequently adjust the signal to compensate for variations caused by component heating, thermal expansion, cable flexing, and so forth. We calibrate the system on every sweep (every 7.8 ms). To apply the automatic calibration, we divide the S-parameters (frequency domain I,Q) for each of the other element pulses, $S_{f,ant}$, by the S-parameters for the calibration channel pulse, $S_{f,cal}$.

$$\hat{S}_{f,ant} = \frac{S_{f,ant}}{S_{f,cal}} \quad (2)$$

Our experiments indicate that this self-calibration allows the system to operate in an ambient temperature range from -5°C to 50°C with only minor changes to the amplitude and phase response of the system (generally at the -20 dBc level or below). Temperature changes in the subsurface (such as freezing and thawing) may still cause

unwanted artifacts over temperature, but these effects have not been explored in this study.

In addition, a onetime “factory” array calibration is performed in a quiet anechoic chamber after the array is fabricated to measure all element direct-coupling responses. This measurement is then used to calibrate out, to first order, element differences. To apply the “factory” calibration, we first apply the automatic calibration to the factory calibration, $S_{f,ant,factorycal}$, using Eq. (3). We next divide the S-parameters for each element pulse by those from a “factory” calibration, and subtract 1 Eq. (4).

$$\hat{S}_{f,ant,factorycal} = \frac{S_{f,ant,factorycal}}{S_{f,cal}} \quad (3)$$

$$\hat{\hat{S}}_{f,ant} = \frac{\hat{S}_{f,ant}}{\hat{S}_{f,ant,factorycal}} - 1 \quad (4)$$

A time delay is next applied to the sequence to prepare the data for conversion to the time domain. Doing so keeps the negative part of signals, centered on zero, from appearing at the end of the reconstructed time-domain signal. To apply the time delay, we multiply the S-parameters of each element pulse by the frequency domain representation of an 8-ns delay.

The sequence is then windowed to reduce “ringing” in the reconstructed time-domain signal. To apply the window, we multiply the S-parameters of each element pulse by a Chebyshev window with 45-dB attenuation of the highest and lowest frequencies.

We next construct a time-domain sequence of 1,024 real, time-domain values by applying a complex-to-real discrete Fourier transform for each element pulse.

Finally, a high-pass IIR spatial filter is applied to all data to eliminate any residual internal reflections and suppress the ground bounce, which is uniform and not helpful for the matching process. The mean removal filter removes constant and very low-frequency components from the radar data. The filter is a first-order, linear, infinite impulse response, high-pass filter. The filter constant is chosen so that information decays with a half-life of 5 m and has a pass-band of 0.024 cycles/m. The running mean, M_i , the mean removed value, $C_{i,mr}$, and the parameter, β , are calculated as follows. The IIR filter distance, d_{iir} , is set at 5 m, while the current distance, d_i , is calculated from the difference in current and prior estimated positions.

$$M_i = \beta \cdot M_{i-1} + (1 - \beta) \cdot C_i \quad (5)$$

$$C_{i,mr} = C_i - M_i \quad (6)$$

$$\beta = 2^{-\left(\frac{d_i}{d_{iir}}\right)} \quad (7)$$

At the end of this process, a fixed range gain is applied to compensate for soil attenuation and geometric spreading of the electromagnetic wave (Daniels, 2004). This

final data product is one “sweep” of array data containing 11 channels and can be thought of as a scalar image (of signed 1-byte pixels) that is 369 pixels deep by 11 pixels across the array. This is the raw data output from the sensor, such as shown in Figure 7, which is sent to both the mapping process and the real-time localization process. We also maintain a database indexed by the geographic location of each sweep, as indicated by a GPS-aided inertial navigation system.

The LGPR concept depends on a sufficient density of subsurface features to be present for the registration process. As seen in Figure 7, the signal to noise ratio of the individual sweeps nears 0 dB near the 30-m-distance point. In this short region, there are very few subsurface features available for registration. We note that this short region has little effect on the localization process, as the system is able to drift through that short segment. Note that this drifting process is possible due to the temporary integration of the RT3003 GPS/INS solution mentioned in Section 1. It is interesting, however, to investigate the prevalence of these sparse feature regions. To this end, the average density of features was investigated with the LGPR in 3 states as well as in Afghanistan for the purposes of understanding the distribution of suitable subsurface features. Here we analyze the density of those features by computing the total signal strength (integrated over depth and over all array elements) of the IIR-filtered GPR data. In this filtered GPR data, low frequency and DC reflections such as vehicle scattering and the surface bounce are filtered out. With some simplifying assumptions (such as stable platform motion), the only remaining energy is from noise and true subsurface features having spatial structure size ranging from ~50 m (the approximate IIR filter cut-off) to approximately 0.2 m (the limit of radar resolution). In general, we find that the radar SNR of these features is sufficient for localization, even in highly attenuating soils (e.g., Arizona). Figure 8 shows the distribution of radar SNR for five test locations. Note that most locations have SNR values over 10 dB. Even in the worst-case scenario (fewest subsurface features—Afghanistan highway) less than 30% of the data is below our SNR ≥ 8 dB guideline, and the regions where the SNR drops below this value tend to be short (on the order of meters) so that an INS is more than capable of tracking motion between GPR locks.

2.2. Real-Time Localization

The LGPR system as implemented allows real-time creation of single-track maps with no offline processing, as well as real-time localization of the vehicle to a prior map. The radar operates in the way described above, both for mapping and for localization. However, when in localization mode, two extra processing steps, gridding and tracking, allow us to determine geographic location and orientation from the radar data. The process is depicted in Figure 9 and described below.




Figure 7. Views of a typical 50 m section of Massachusetts highway collected by the LGPR. The front view shows all 11 channels across the array from a single location, while the top view shows the total reflected energy (squared and integrated in depth). The side view is taken from channel #9. On the bottom, the SNR from each sweep is also shown.

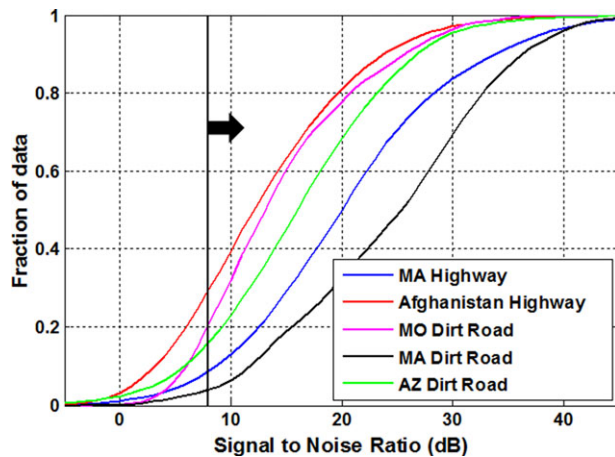


Figure 8. Example distributions of SNR for LGPR data collected from five sites: a clay-rich dirt road in Missouri (MO), a dirt road in a grassy Massachusetts (MA) field (sandy loam), a MA highway, an Arizona (AZ) dirt road (highly attenuating hard-packed clay), and a highway in Afghanistan. In our experience, LGPR data with SNR >8 dB are appropriate for achieving accurate localization.

We use a single standard consumer-grade dual-core processor running Linux for all processing. All map and localization code was written in C. Low-level firmware was executed on an FPGA. As this LGPR prototype used a fixed temporal sample rate (126 Hz), the uncompressed raw

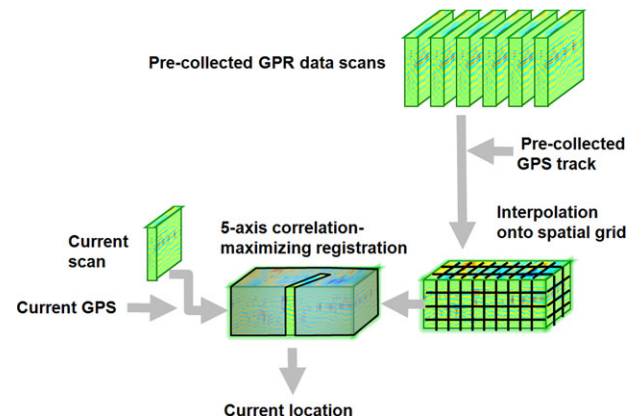


Figure 9. LGPR five-axis registration process. Data from the initial mapping collection is interpolated onto a regular grid and the current sweeps are registered in five dimensions to the grid.

map data sizes varied from ~32 MB/km for the 97 km/h collections to ~192 MB/km for the 16 km/h collections. Commercial zip packages have shown a factor of 20:1 compression is possible on raw LGPR data. It is likely, based on the results of the testing presented in Section 3, that lower spatial density data would be sufficient for localization at all speeds. Further discussion of data size can be found in Section 4.2. Similar to Levinson and Thrun (2010), map data were dynamically pulled to RAM in a dual buffer of local 50 m × 50 m grids to keep memory requirements constant.

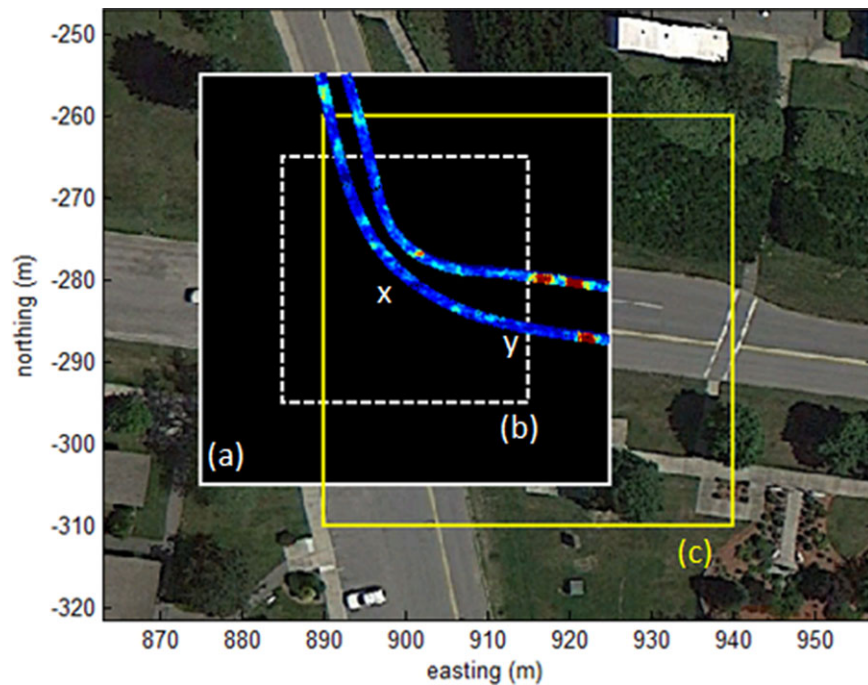


Figure 10. A top-down view of two LGPR data tracks is shown in blue, overlaid on a map. The local grid (a), which includes all data within a 50 m square area, is generated when the vehicle is at location x, and is used while the vehicle is within box (b). When the vehicle reaches the edge of box (b) at point y, a new local grid (c) is generated centered at y. Two passes appear in this grid since, while mapping the area, the vehicle traversed two different lanes on the outgoing and incoming journey. Grid points where no map data are present are colored black (the majority of the grid is empty).

2.2.1. Grid Map Preparation

At initialization, we prepare a local “grid,” a rectangular 3D data structure of scalar signed values representing the reflected radar signal from the subsurface environment centered on the vehicle. As the vehicle approaches an edge of the current grid, we replace it with a new grid centered on the vehicle. The gridding process is illustrated by Figure 10. Using a fixed-size local grid allows us to page in relevant data as needed without keeping huge data sets in memory. Furthermore, it permits fast calculation as data from a newly acquired sweep of the GPR can be quickly matched to the nearest data from the prior passes.

To create a grid, we query the database for previously collected radar sweeps that fall within the grid area. We interpolate the raw data from each of those sweeps onto grid points spaced 5 cm apart in the horizontal dimensions. This may appear too fine for the stated resolution of the radar (20–30 cm); however, it is often beneficial to overresolve this radar data. The extra information obtained by overresolving does diminish rapidly for finer grid spacing, but with a very high SNR (our SNR is typically >20 dB), this additional information is still useful for registration (useful because the small subresolution features are still above noise). Note that the accuracy of any correlation-based registration process

is driven largely by the SNR of the data, rather than by the resolution of the GPR or the grid spacing at which we sample that GPR data. Vertically, the samples in depth for each grid point retain their native density and span (369 points sampled over 60 ns). The data for each grid point are a weighted average of all raw data from prior passes that are within 12 cm. The weight for each raw data point is inversely proportional to its distance to the grid point. For this study, maps were created using single-pass data sets. Multipass data-set-based maps would require additional processing and registration steps to remove “ghosting” effects from relative misalignment of the data sets, an area we have not explored with GPR data.

2.2.2. 5-DOF Registration

As each sweep arrives from the radar, we search for the vehicle pose that best fits that sweep to the grid. The search region is five-dimensional: latitude, longitude, height, heading, and roll. We found that a simple two-axis (longitude, latitude) search domain produced inferior results, as variations in the vehicle suspension state between passes requires compensating for the subsequent vertical motion (height and roll) of the array, which can cause the individual elements to move as much as 5 cm up or down from their

height on the mapping pass. We used the correlation between the sweep data and the grid data as the heuristic fit criterion. The correlation is calculated as

$$\frac{\sum_{i,d} (A_{i,d} B_{i,d})}{\sqrt{\sum_{i,d} (A_{i,d})^2 \sum_{i,d} (B_{i,d})^2}}, \quad (8)$$

where the symbols A and B represent the current sweep and a candidate slice of the grid data, respectively, and the sums are over all channels i and 360 depth bins d . When correlation is $+1$, then the match is perfect, while a correlation of -1 indicates the worst-case correlation (in which one data set is 180 degrees out of phase). Correlation is a simple metric to quantify the matching of subsurface GPR signals and is independent of absolute signal strength. When the vehicle pose for optimal correlation is found, we regard the current sweep as being registered to previous grid data.

Particle swarm optimization (PSO) was chosen as the optimization technique to find the highest correlation 5-DOF array pose that matches GPR data to the grid. Particle swarm optimization is a well-known optimization technique (Kennedy & Eberhart, 1995) chosen due to its ability to search a very large space of candidate solutions. Furthermore, this technique makes no assumptions about the correlation space and is robust to local minima. This quality also allowed us to likewise make no assumptions about the GPR data as we expected a wide variation for different areas. Of the existing heuristic search strategies, the selection of PSO was driven by its relatively lower computational needs, ability to search large solution spaces, and ability to be scaled to the computational resources available (Shi & Eberhart, 1999). Moreover, it offers greater robustness in achieving near optimal solutions when compared to simpler techniques such as simplex or hill climbing.

The algorithm begins by creating “particles,” each of which occupies a randomly chosen point in the five-dimension search space associated with the 5 DOFs. During each iteration of the algorithm, each particle is evaluated for quality of fit (correlation). The algorithm keeps track of the best fit experienced by each particle, and the best fit overall across all particles. Each particle also has a “velocity” value associated with each search dimension. The velocity controls how the position is updated after each iteration of the algorithm.

Each particle then moves through the search space, its position \vec{P}_{i+1} updated from its previous position \vec{P}_i with the vector velocity \vec{v}_i :

$$\vec{P}_{i+1} = \vec{P}_i + \vec{v}_i \quad (9)$$

Next, the velocity of each particle is updated based on a combination of the best position achieved by that particle

and the best position of all the particles found thus far:

$$\vec{v}_{i+1} = \omega \vec{v}_i + C_\alpha r_\alpha (\vec{P}_{\text{Bestlocal}} - \vec{P}_i) + C_\beta r_\beta (\vec{P}_{\text{Bestglobal}} - \vec{P}_i) \quad (10)$$

ω is the inertial constant, C_α is the cognitive constant, and C_β is the social constant. r_α and r_β are random numbers chosen anew for each iteration. $P_{\text{Bestlocal}}$ is the best position that a particular particle was able to find in the current or past iterations. $P_{\text{Bestglobal}}$ is the best position that any of the particles was able to find. The behavior of the particles can be understood by examining each of the three components of the velocity update equation. A high C_α value steers the particle toward areas where it has found good performance. A high r_β value steers a particle toward the best position found by any particle. A high ω will cause particles to stay on their current trajectory, largely unchanged. A low ω value will allow the particle to be more strongly influenced by the other terms in the equation and thus will tend to veer toward other particles. In the implemented optimization algorithm, ω starts high and decreases with subsequent iterations. These values encourage an initially wide level of exploration by the particles, which subsequently converge onto a solution during the latter iterations.

As a further refinement, the size of the search region, the number of particles and the number of iterations evolve, depending on the correlation achieved. When correlation is high, the search region shrinks, allowing it to complete quickly. When correlation is low, the search region expands. We consider the tracker “locked” onto the grid if the correlation exceeds a threshold (0.9), at least two of the antenna elements overlap the mapped data, and the candidate location is physically possible based on the vehicle speed and previous location (it is within a delta of the velocity-extrapolated expectation). If the tracker is not locked, then we estimate our location using the velocity integrated position of the vehicle from GPS/INS inputs, and widen the search box for subsequent sweeps. For an example of how the particles search the space when first initialized with a large $5 \text{ m} \times 5 \text{ m}$ region, see Figure 11. The result of this registration process for one sweep is shown in Figure 12.

3. EXPERIMENTAL RESULTS

We demonstrate in this section that the LGPR can function at highway speeds while providing a cross-track accuracy of 4 cm RMS when using an Oxford Technical Solutions RT3003 RTK-based differential GPS/INS solution as a truth reference with 2 cm local accuracy. Cross-track accuracy is defined as the resulting error along the axis perpendicular to the vehicle heading. Along-track accuracy is defined as the resulting error along the direction of the vehicle heading. Our tests include a speed test in which the vehicle was driven in a large loop at speeds of 16 km/h to 97 km/h (10 mph to 60 mph), both for map creation and localization, as well as a long-distance test. RTK truth speed was used to




Figure 11. The locations of 200 particles over 100 iterations (each representing a candidate location for placing the current single array sweep into the map) are shown. A single high correlation peak is evident in the 5 m by 5 m region. (a) Initial estimate (from GPS/INS), (b) Highest correlation estimate. Note that the particle locations shown here in easting/northing represent just two of the five dimensional spaces in which they reside (height, heading, roll are not shown). Hence, the correlation space here appears to be multivalued. In this image, dark blue represents a correlation of 0, while dark red is 1.

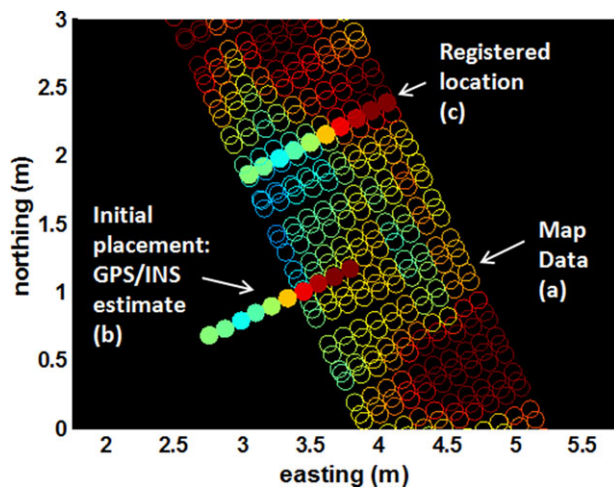


Figure 12. The image shows the initial GPS/INS estimate and the optimized registration of one sweep of radar data with filled circles for each channel, against a background of map data with empty circles for each channel. The LGPR produces, and the tracker processes, sweeps at 126 Hz to localize the vehicle relative to the map. Note that this figure shows only 2D (depth-integrated) energy (to enable visualization), while the process is actually performed on fully 3D data.

measure the actual speed of the vehicle when calculating the accuracy at each data point. We discuss below the results from each set of tests and how they improve on existing approaches.

3.1. Speed Test Results

We devised a set of tests for our LGPR system to quantify the impact of vehicle speed on baseline creation and real-time localization accuracy. Two of the key limitations of prior LGPR work, other than the inability to fit underneath a vehicle, were that the operations were limited to speeds below 24 km/h and that the real-time accuracy had not been characterized (Stanley et al., 2013). Table III shows the data collected during one day of testing. These tests were conducted with a manual driver. It should be noted that while it would be worthwhile exploring the results of high-speed baselines with low-speed tracking collections, those experiments were not attempted for this initial test.

On a clear autumn day, 6 baseline mapping collections and 21 separate one-lap 1.6 km localization tracking collections were manually driven on the test course. A baseline collection is a manual drive around the track with the LGPR in mapping mode. A tracking collection is a manual drive around the track where the LGPR system references a particular map, in this case taken at a particular speed, and uses it to localize itself in real-time. Speed was controlled within operator abilities, though the calculations were based on the truth measurement of the speed of the vehicle for each data point. To test performance of the LGPR versus baseline and localization vehicle speeds, each of the six 152 cm wide (the array width) baseline mapping collections was driven at a single nominal speed starting at 16 km/h in 16 km/h increments up to 97 km/hr. The 21 tracking collections were then driven at speeds and referenced baseline collections according to Table III.


The test course was a tree-lined asphalt four-lane highway with a grass divider. It contained two straight segments approximately 360 m each along with four significant curved segments with radii of 7 m, 22 m, 75 m, and 130 m as shown in Figure 13.

All speed test collections started on the southeast corner of the course and proceeded in a counterclockwise direction around the course and ended across the highway from the starting point. To determine the accuracy of the LGPR system, a separate truth GPS/INS system was setup to use a differential real-time kinematic (RTK) differential GPS solution. The RTK base station for this differential setup was placed mid-loop to provide approximately 2 cm relative localization accuracy in conjunction with the truth GPS/INS unit on the vehicle. By comparing the localization coordinates of the LGPR with this truth GPS, we were able to estimate localization accuracy to within 2 cm.

Accuracy calculations required aligning and comparing the baseline and tracking positions P (2D vectors of local

Table III. Tracking collections, related speeds, and baseline collection speeds taken during a single day of LGPR speed and real-time accuracy testing are shown.

	Tracking 16 km/h	Tracking 32 km/h	Tracking 48 km/h	Tracking 64 km/h	Tracking 80 km/h	Tracking 97 km/h
Baseline 16 km/h	1 run	1 run (removed)	1 run	1 run	1 run	1 run
Baseline 32 km/h		1 run	1 run	1 run	1 run	1 run
Baseline 48 km/h			1 run	1 run	1 run	1 run
Baseline 64 km/h				1 run	1 run	1 run
Baseline 80 km/h					1 run	1 run
Baseline 97 km/h						1 run

**Figure 13.** Highway speed test course with the tracking “truth” RTK data sets overlaid on an overhead image of the route.

easting and northing positions) for the truth RTK GPS/INS p^{RTK} , GPS/INS (WAAS) p^{GPS} , and LGPR p^{GPR} . Because the LGPR global location estimates are determined relative to the baseline map (which is created with a WAAS-aided GPS/INS), the LGPR location estimates may have global error due to the global error of the map coordinates. The accuracy of the LGPR approach is thus best measured with a relative metric. In particular, the location of the LGPR location estimate of the tracking collection relative to the GPS/INS track of the baseline can be compared with the actual difference between the tracking collection and the baseline from

the differential RTK truth GPS/INS. The complexity of this calculation could have been reduced by creating the map with RTK truth data rather than WAAS GPS/INS. Use of RTK systems, however, is limited in range and could have provided nongeneralizable performance improvements to the system. In addition, due to the configuration of the systems as built, this would have required offline processing of the maps after collection rather than the real-time WAAS-based map collection demonstrated here. To that end, the truth RTK GPS/INS position data from the baseline and tracking collections were aligned with one another using

interpolation of the two baseline nearest-neighbor positions to align with each tracking point. This alignment created a truth time, t_a , and a position mapping between the tracking and baseline collections. This then allows an estimate of the relative localization error described above. The vector error, E^{total} , was then calculated according to Eq. (11).

$$E^{total}(t_a) = [P_{Track}^{GPR}(t_a) - P_{Base}^{GPS}(t_a)] - [P_{Track}^{RTK}(t_a) - P_{Base}^{RTK}(t_a)] \quad (11)$$

The components of cross-track and along-track error were then calculated by splitting the error vector into body frame components using the RTK tracking heading as the along-track axis. The same approach, substituting the GPS/INS tracking measurements for the LGPR tracking measurements, was used to calculate the GPS/INS errors. In addition, to calculate speed-based results, actual RTK truth speed measurements were used in the error calculations rather than the nominal tracking collection speeds. Speed test results are summarized in Table IV.

During overall speed testing, as a function of distance, the differential GPS/INS (using WAAS) demonstrated cross-track errors of 35 cm RMS, whereas the LGPR demonstrated cross-track errors of 4.3 cm RMS. This overall RMS cross-track error is less than half of the documented 9 cm RMS cross-track error of earlier referenced LIDAR-based localization methods, noting that there may be significant differences in experimental conditions, equipment,

Table IV. Distance-based RMS tracking error performance for the LGPR and GPS/INS (WAAS) systems for all tracking collections.

	Cross-Track Error	Along-Track Error	Total Error
LGPR	4.3 cm RMS	11.9 cm RMS	12.7 cm RMS
GPS/INS (WAAS)	35.0 cm RMS	41.7 cm RMS	54.4 cm RMS

and execution. This high-accuracy result is encouraging for an early capability test of the LGPR technology.

The example speed collection in Figure 14 shows the along-track and cross-track error for the LGPR and GPS/INS (WAAS) solutions. The LGPR along-track error shows a negative bias that may indicate either a time misalignment in the error calculations or a time lag in the LGPR real-time solution. A typical speed collection lasted between 1 and 10 min, depending on vehicle velocity, which we varied from 16 km/h to 97 km/h (10 mph to 60 mph).

The summary data from 0 to 97 km/h (60mph) are shown in Figure 15. Interestingly, the low-speed region of the plot exhibits higher errors, which are a direct result of the northernmost 7-m-radius turn in the course, in which the driver was forced to slow down. This tight turn radius made it difficult for the driver to maintain overlap with our narrow 1.5 m baseline map. It is interesting to note that when the northernmost sharp corner above 718 m northing




Figure 14. Example speed measurement run (55 km/h mean speed). Differential GPS shows 40 cm or greater error, whereas the LGPR estimate closely tracks the RTK truth data. The larger LGPR error and transition period between 50 s and 60 s corresponds to the northern 7 m radius U-turn.

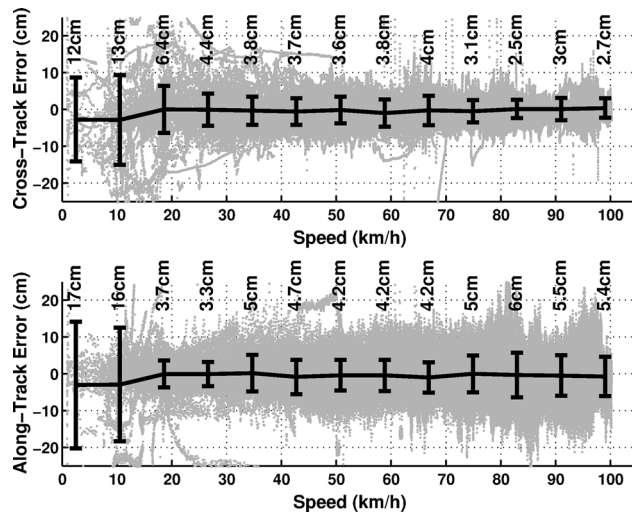


Figure 15. Time-delay corrected speed vs. error with the RMS error for each 8 km/h (5 mph) wide speed bin up to 97 km/h (60 mph). Larger deviations at low speeds were associated with the 7 m radius corner at the northernmost end of the course. Actual RTK truth speed measurements were used to bin all collected data as above.

is removed, the low-speed error is comparable with the errors at higher speeds.

The errors in the along-track direction increase with speed, which we believe to be a combination of a small time differential between the LGPR system clock and the DGPS system clock (an artifact of the way we measure truth), as well as increased error due to ground sample distance. To determine how much error was due to time misalignment versus ground sample distance, we plotted the along-track error versus speed and fitted a line to it. As increasing ground sample distance would be expected to grow the

Table V. Time-delay corrected distance-based RMS tracking error performance for the LGPR and GPS/INS (WAAS) systems for all tracking collections.

	Cross-Track Error	Along-Track Error	Total Error
LGPR	4.3 cm RMS	5.9 cm RMS	7.3 cm RMS
GPS/INS (WAAS)	35.0 cm RMS	41.7 cm RMS	54.4 cm RMS

along-track error, but would not be forward or reverse biased, the slope of the line indicates a 5.5 ms time differential between our LGPR system and the DGPS system used to measure truth, as shown in Figure 16.

The estimated resulting along-track error due to increasing ground sample distance is 5.9 cm RMS overall, as shown in Table V and 6.5 cm RMS in the high speed 89–97 km/h range.

LGPR cumulative cross-track localization error, shown in Figure 17, was less than 3.3 cm for 68.3% of travel distance and less than 7.4 cm for 95.5% of travel distance, which is well within the tolerance required to maintain a vehicle inside of a lane.

Cross-track error values shrink as the correlation approaches one, as plotted in Figure 18. This suggests that correlation is a useful metric to gauge tracking performance. It should be noted that the peak error data points were during a time when the correlation was zero, which implies that the LGPR solution was relying on the integrated GPS/INS solution alone during that time.

As the LGPR array overlap with the prior map decreases, localization accuracy decreases, as plotted in Figure 19. Good (≤ 6 cm) tracking performance is maintained until roughly 75 cm of offset (approximately 60 cm of overlap), above which error increases to just over 17 cm at 1.1 m. This would be expected for narrow 1.5-m-wide




Figure 16. Along-track distance based error vs. speed plot with least squares linear fit estimating time-related bias errors.




Figure 17. Cumulative distribution of cross-track and along-track error as a fraction of total distance traveled for all data collected. Percentages shown correspond to standard 1σ and 2σ values for a Gaussian distribution (the measured distribution shown here deviates slightly from Gaussian). Both time-delay corrected and nominal data are shown on the along-track error plot.

Table VI. Peak relative change values in key dimensions between track and baseline collections.

Metric	Max Value	Min Value
Height	0.2 cm	−0.2 cm
Roll	1.2 deg	−1.5 deg
Pitch	0.8 deg	−0.6 deg
Cross-track position	780 cm (130 cm trimmed)	−541 cm (−130 cm trimmed)
Cross-track error	70.4 cm	−120.9 cm
Along-track error	84.7 cm	−191.2 cm

array-based single-pass maps. As the overlap decreases, less data are made available to register and the chances of a poor fit increase. A wider array, such as could be mounted on a truck or bus, would also be expected to require approximately 60 cm of overlap for low error performance.

We note that the vehicle motion was relatively benign, as would be expected during paved highway operation. Table VI describes the peak relative motion of the vehicle during the tracking runs. Relative angles are the difference between the absolute orientation of the baseline map and the absolute orientation of the tracking collection. The cross-

track position is the difference between the RTK GPS/INS solution for the baseline map and the RTK GPS/INS solution for the tracking collection. The data were trimmed to only include data points at less than 1.3 m in position offset, where the array elements were actually overlapping with the map.

As this experiment was performed on a flat highway, large variations were not expected to be observed. The error did not appear to increase as a function of the limited changes in each of these dimensions. The cross-track and along-track peak error values result from a time, as noted earlier, where the LGPR had zero correlation during the sharp U-turn at the northern edge of the track. This implies that the peak cross-track (lateral) error was a function of the integrated GPS/INS solution during that time period.

We note that the summary data presented in this section have been modified in two ways: we have only included data where the vehicle has at least one element overlapping the map (130 cm offset or less), and we have eliminated data before the initial lock of the particle swarm optimization (less than 6 s into each collection in all cases). In addition, the 32 km/h nominal pass at 16 km/h baseline was removed due to what appeared to be unrelated operating system maintenance saturating the processing during the test. As the tracking algorithm drops data points when processing takes longer than the 8 ms period, such as when locking onto




Figure 18. Correlation (x axis) vs. Cross-track error (cm). Note that the registration process defaults to the integrated GPS/INS solution when the correlation is below 0.9 or less than two elements overlap.

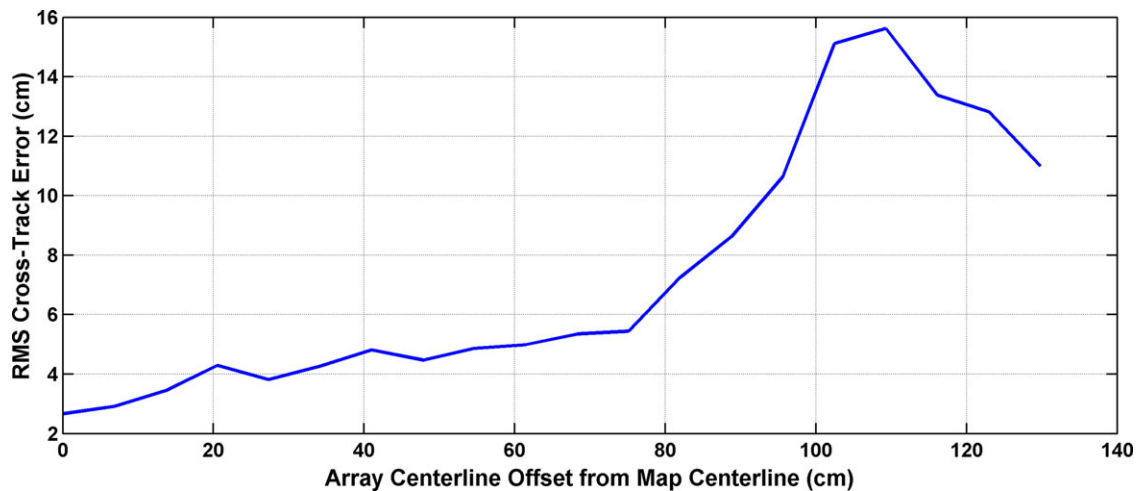


Figure 19. Cross-track error vs. cross-track offset of array. This plot shows that accurate localization is possible, even with only a limited overlap between the map and the array.

the initial position, the tracking rate is not always the same as the nominal loop rate. The tracking rate was measured at 126 Hz for approximately 68% of data points, 63 Hz for 27% of data points, under 56 Hz for 4% of data points, and measured at 111 Hz RMS overall.

3.2. Results for Improved Highway Test

We conducted highway testing to demonstrate the ability of the LGPR system to maintain lock over long distances

on improved roads (which include rebar) and bridges. Testing was conducted with a manual driver on the highway shown in Figure 20. Vehicle speeds varied based on traffic conditions, but most of the test was driven at approximately 97 km/h (60 mph). The total distance traveled was 77 km (48 miles). Because a differential RTK GPS solution is not feasible in this scenario (our differential GPS base station range is only a few miles) and we have not as of yet fused the system with other localization methods, such as LIDAR, we were unable to assess localization accuracy on




Figure 20. Overhead map of the improved highway testing loop (77 km), which included several bridges.

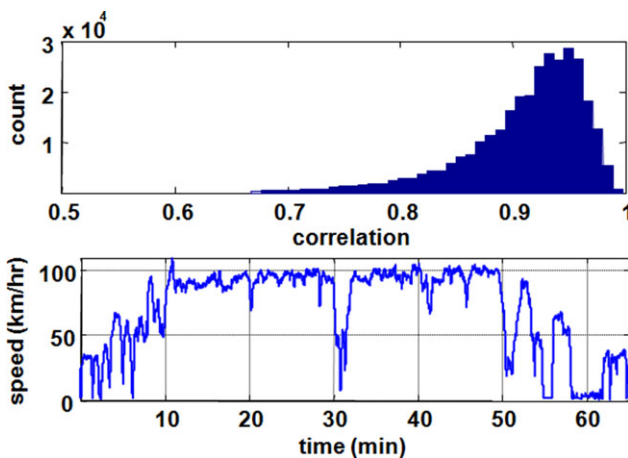


Figure 21. Correlation histogram and speed values for the entire localization test on improved highways.

this test. However, it is still a useful exercise to examine key measures of performance for this long-distance test. The data correlation, as defined in Eq. (8) over this collection was excellent, with a mean of 0.90 and a standard deviation of 0.07, as shown in Figure 21.

The LGPR system was able to track the vehicle location on over 99.9% of the 324,000 location estimates produced on the journey (excluding cases in which the vehicle did not overlap the map). The system experienced eight short periods of tracking loss during the localization test. Most periods of tracking loss lasted less than 1 s, while the longest period lasted 4 s. The causes of the brief losses of tracking are under study, as sufficient subsurface structure appears to be present. During this period, the system defaults to GPS (with Omnistar or WAAS, as available). This




Figure 22. Example loss of LGPR lock and the corresponding loss of correlation. The LGPR is able to relock onto position after just a few seconds.




Figure 23. Overlay of GPR data on map. The GPR data is 3D, so only the depth-integrated intensity is shown here. Substantial subsurface features can be seen in the data. Global alignment for this data set used a WAAS GPS/INS solution, so some alignment offsets are expected.

underscores that additional localization sensors would be required to achieve the high reliability ultimately required by self-driving vehicle applications.

A typical example of loss of tracking is shown in Figure 22. Notice that as the fraction of the array over the map was reduced, the correlation of the values drops to zero. As discussed earlier, the correlation provides a good indication of certainty in the LGPR location estimate and could be used to estimate covariance values for fusing multiple sensors in a Kalman filter.

Typical data are shown in top view and overlaid onto a map in Figure 23.

It is of interest to note that the LGPR data can potentially be used to evaluate the structure of the road and map underlying or surface features, such as pipes or potholes. This may, in itself, provide a use case for creating and maintaining a large-scale map of the roadways around the world.

4. DISCUSSION OF REMAINING CONCEPT RISKS

The authors acknowledge that this paper covers early work in the exploration of localization using ground penetrating

radar. There remain significant challenges and risks that should be further characterized. We discuss a number of those risks below in limited detail to provide context for future work.

4.1. Subsurface Map Stability

A fundamental requirement for the LGPR system is that subsurface features must be sufficiently stable so that high correlations can be made on future passes. The highway and subsurface move based on seasonal and water based variations, which offer the potential to reduce the effectiveness of LGPR based maps. The problem can be broken down into two main questions. First, do the subsurface features in the relevant locations change significantly over time? Second, do those changes prevent localization?

While subsurface feature localization is new, the concept of precise leveling for geodetic surveying has been studied in the United States since 1877 (Floyd, 1978). The primary direction of ground motion is in the vertical direction, barring significant events such as earthquakes. Primary near-surface sources of motion include direct impact or compaction from surface-based sources such as vehicles, motion due to frost heave, and sinking or swelling of soil and rock due to moisture content (Floyd, 1978). Roadbeds are also compacted during construction to mitigate vehicle-related compaction effects. The depth of frost penetration varies depending on the region as shown in Figure 24. The continental United States has peak frost depth at just over 100 in. in the worst-case areas, which, as a comparison, is less than the typical 3 m (118 in) peak depth of penetration of the LGPR.

The actual risk of frost heave, however, is substantially reduced by standard road building and drainage techniques. On top of the compacted roadbed itself, courses of gravel and other materials are used to improve drainage. The size of the particulate matter within the road subsurface is typically controlled to reduce susceptibility to frost heave (Moultan, 1980). The risk of swelling due to moisture is also substantially reduced by standard road drainage techniques. Even without human modification, the expected vertical rise of natural soils due to swelling is limited to peaks of a few inches in cases of extreme swell conditions (Christopher, Schwartz, & Boudreau, 2006), although changes in subsurface character may also occur. In short, there are regular, primarily vertical changes to the subsurface and surface of roads and highways, though it remains to be shown whether the changes are of any significance to the LGPR over time.

In regards to whether or not the changes in subsurface or surface affect localization, the answer is unclear in the first and known in the other. Surface changes have negligible effects on the LGPR correlation due to the low weighting of the measured ground bounce from the surface. Subsurface changes would need to be substantial enough to

significantly horizontally shift the location of the highest correlation or make the correlation sufficiently low over a long enough area that the measurements from the IMU (or INS with other sensors) are unable to be integrated to accurately smooth the location estimate. Frost heave on roads, for instance, typically results in multiple local disruptions in the road surface that may have little effect on the overall solution of the system. It remains, however, to be shown whether subsurface changes underneath roadways over long time periods will affect localization to prior LGPR maps.

4.2. Data Requirements

The LGPR saves data at a rate of 535 KB/s (2 GB/h). That rate can be reduced to 87 KB/s or (311 MB/h), by not oversampling the data in depth and by eliminating insertion of redundant INS data in the database. For reference, the popular Velodyne HDL-64E 3D LIDAR creates data at about 4 MB/s.

How much map data LGPR requires depends on the size of the geographic area in which the vehicle operates. When the vehicle follows a known route for which it has previously recorded data, it swaps in map data at about the ratio of the speed at which it tracks the route and the speed at which it collected the map data. For example, when tracking at 27 m/s over data recorded at 2.2 m/s, LGPR swaps in map data at a rate of about 6.5 MB/s (23 GB/h). Again, efficient data storage could reduce the rate to 1.1 MB/s (3.8 GB/h).

In the example above, the LGPR tracks a “ribbon” of data about 1.5 m wide. If it deviates more than a meter to either side of the nominal path, so little of each scan will overlap the map data that it is likely to get lost. More map data to each side would allow localization if the vehicle wanders. Meanwhile the map data are more closely spaced in the direction of travel than necessary, 1.7 cm compared to 12.7 cm spacing from side to side.

The LGPR actually projects the map data onto a rectangular grid. With radar resolution of about 12.5 cm in horizontal position and 25 cm in depth, data sampled at 5 cm in horizontal position and 10 cm in depth captures the features to which the radar is sensitive. The resulting map data density is 15 Kb/m². Providing a swath of map data 3 m wide for a vehicle traveling 40 m/s requires 1.8 MB/s (6.4 GB/h) carried onboard or communicated while traveling.

The United States has more than 215,000 lane miles (346,000 km) in its interstate highways and 8.5 million lane miles (13.7 million km) of paved road (Federal Highway Administration, 2010). The average interstate lane is 3.6 m (12 ft) wide (Federal Highway Administration, 2007), which is conservative for the average road lane width. Simple calculations give an area of 1,250 km² of interstate map data and 49,300 km² of road map data, which would consume 19 TB and a daunting 740 TB, respectively. Driveways, parking lots, and unpaved roads might double that. Carrying




Figure 24. Map of United States peak frost penetration depth (Moulton, 1980).

enough LGPR data for the entire United States requires advances in storage technology, but carrying enough data to operate in a small geographic area does not.

A fourth-generation (4G) mobile broadband data connection with 12.5 MB/s capacity would more than suffice to supply LGPR map data for a vehicle moving at highway speed. A third-generation (3G) connection with 625 kB/s capacity would fall short for highway speed but suffice for city driving. Engineering a system to dynamically download map data, with enough onboard storage to carry through gaps in the cell phone network, seems achievable.

4.3. All-Weather Operation

All-weather operation, particularly in snow and rain, is one of the key potential applications offered by LGPR. It remains to show, however, that the LGPR system can operate reliably in either condition. The LGPR system has been used without issue during and after rain of up to $\frac{1}{4}$ in; however, no direct study has been performed to further characterize performance. As noted in Hoekstra and Delaney (1974), saturation of the soil up to 30% can attenuate signals in the 100 MHz to 400 MHz range at up to 10 dB/m. This will reduce the signal to noise ratio of the system, though there is sufficient margin to overcome most of this effect. Variations in moisture content also cause GPR data to contract or expand in the depth dimension, as moisture has an effect on wave propagation velocity. As a result, the registration process may be impacted when soil moisture content is variable between the baseline and tracking passes. This is an active area of research, so the impact on the LGPR remains unknown, with the exception of a single known good performance immediately (<24 hr) after $\frac{1}{4}$ in of rain.

Propagation of 600 MHz signals in snow has been characterized, and the associated loss at these frequencies is approximately 1–2 dB over the 6 in offset of the array from the surface (Abe et al., 1990). Losses of similar magnitude or less would be expected for the 100 MHz to 400 MHz range. The practical risk associated with operation in snow is the salt used to melt the snow on roads, which attenuates RF signals more significantly. Although the LGPR has successfully measured subsurface characteristics through snow and road salt mixes, localization performance remains to be studied and characterized for LGPR systems.

4.4. Vehicle Chassis Reflections

Vehicle chassis reflections have the potential to cause variability and may be interpreted as additional subsurface reflections. When signals are emitted from the antenna, a fraction of those waves reflect off of the surface and bounce off of the vehicle chassis. These waves may then bounce off of the ground and into the receiving antenna. The strength of the reflections is a function of the power emitted, the characteristics of the vehicle chassis, and how well the signal couples with the ground. Limited simulations discussed earlier indicate that the strength of these signals is insufficient in the current design to prevent localization. Further study is warranted to understand differences between vehicle chassis, especially as the distance from the ground increases and the distance to the chassis decreases.

4.5. Antenna Polarization

The prototype antenna design yields polarized measurements. As a result, the map data depends on the

orientation (heading) of the vehicle. In cases where map collection paths and tracking paths are far off parallel (as might occur in an open field or parking lot), alternate antenna designs and algorithm adaptations may be required.

4.6. Miniaturization

The small LGPR array design presented and tested in this paper is 3 in \times 2 ft \times 5 ft. While a significant advancement from the original 1 ft \times 2 ft \times 10 ft arrays, the arrays will need to be further reduced in height in order to be broadly applicable to consumer vehicles. The 3 in height was chosen as approximately 1/16 of the wavelength of the 250 MHz center frequency of the antenna, though it is not clear that this offers significant benefits over other similar heights. The primary concern is to reflect waves in a way that does not create large ambiguity in timing of reflections and will not significantly destructively interfere with the emitted waves. As the height of the array shrinks, additional concerns include the proximity of the top surface ground plane to the antenna itself.

4.7. Cost

Cost is a driver for consumer use of self-driving vehicle technologies. There has been little work to date on determining the eventual cost of a mass-produced LGPR system. The LGPR design, however, has no moving parts, uses components that are readily available commercially, and can be prototyped for less than the typical cost of some commercial LIDAR systems. The antenna itself is made of simple printed circuit boards. The most expensive component in the system described in this paper is a processing chassis designed for military conditions; the military chassis would not be required for consumer use. Cost is already within bounds for industrial and military use, even produced as prototype systems. The authors are not aware of a technical reason the LGPR system could not be produced inexpensively in high volume.

5. CONCLUSIONS

Accurate, reliable, and robust localization is critical for operation and acceptance of autonomous vehicles. The vehicles must be capable of robust operation not only over an extended duration but also in adverse conditions, such as snow, fog, dust, and GPS-denied areas, as well as in the face of sensor damage or failure. Current approaches, including LIDAR, GPS/INS, and camera-based techniques work well in benign conditions for limited duration but are susceptible to failure or insufficient performance in many common weather and environmental conditions.

This paper details a new, complementary mode of map-based localization that is inherently insensitive to many of the failure modes of traditional and current localization approaches. We have demonstrated accuracy levels better than

or equal to other large-scale approaches using a relatively simple low-profile LGPR system. We have shown that a low-profile system can operate underneath a vehicle at highway speeds on existing road networks.

One of the fundamental limitations in our approach, similar to LIDAR-based localization, is the map generation and associated data requirements. Unlike LIDAR mapping, LGPR requires at least one pass to cover most of each lane width, though complete coverage is not necessarily required as partial overlap is generally sufficient for vehicle localization. There will be no holes in the map due to other vehicles blocking the sensor view, and underground geology is relatively stationary. In addition, LGPR data maps have alternative uses as a detailed map of underground infrastructure (pipes under roads, rebar in bridges, etc.).

The miniature LGPR array is small but would need to be reduced further from its current 7.5 cm height to fit under some passenger vehicles. We are currently studying the possibility of reducing this height further.

There are many fundamental areas of research associated with LGPR that remain to be explored. First, LGPR should be fused with current and traditional localization approaches in order to provide a robust overall system. The current heuristic approach is dependent on a GPS solution and would likely lose track in extended GPS-denied conditions. Second, it remains to be shown that LGPR functions over the wide range of soils and weather conditions, such as water-saturated soil and significant surface water accumulation in particular. While RF waves at 100–400 MHz penetrate snow with almost negligible losses, as can be inferred from (Abe et al., 1990), road salt and surface or subsurface water accumulation can attenuate the signal. Third, vehicle chassis radar reflections must be characterized across a wide space of vehicle types to ensure cross-vehicle compatibility. Last, of interest would be creation of large-scale maps using multipass mapping datasets.

ACKNOWLEDGMENTS

The authors would like to thank Ed LeFave, Tom Kelley, and the members of the RF Systems Test Facility for building and helping test the LGPR system. Special thanks to Jim Armao and the Lincoln Laboratory motor pool for allowing us to use the vehicle for testing. In addition, the authors would like to thank Birol Bekirov, Alan Fenn, Chip Coldwell, Jeremy Furtek, Peter Hurst, and others who made this approach possible.

REFERENCES

- Abe, T., Yamaguchi, Y., & Sengoku, M. (1990). Experimental study of microwave transmission in snowpack. *IEEE Transactions on Geoscience and Remote Sensing*, 28(5), 915. Retrieved from <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=00058981>

- Badino, H., Huber, D., & Kanade, T. (2012). Real-time topometric localization. International Conference on Robotics and Automation, St. Paul, MN. Retrieved from https://www.ricmu.edu/pub_files/2012/5/badino_icra12.pdf
- Brunner, C., Peynot, T., Vidal-Calleja, T., & Underwood, J. (2013). Selective combination of visual and thermal imaging for resilient localization in adverse conditions: Day and night, smoke and fire. *Journal of Field Robotics*, 30(4), 641–666.
- Buehler, M., Iagnemma, K., & Singh, S. (2009). *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*. Berlin: Springer.
- Christopher, B. R., Schwartz, C., & Boudreau, R. (2006). *Geotechnical Aspects of Pavements (FHWA NHI-05-037)*, Washington, D.C.: US DOT Federal Highway Administration.
- Cummins, M., & Newman, P. (2011). Appearance-only SLAM at large scale with FAB-MAP 2.0. *The International Journal of Robotics Research*, 30(9), 1100–1123. Retrieved from <http://ijr.sagepub.com/content/30/9/1100>
- Daniels, D. J. (2004). *Ground penetrating radar, volume 1*, Institution of Engineering and Technology. London, United Kingdom.
- Federal Highway Administration. (2007). *Mitigation strategies for design exceptions*, Washington, D.C.: US Department of Transportation. Retrieved from http://safety.fhwa.dot.gov/geometric/pubs/mitigationstrategies/chapter3/3_lanewidth.htm
- Federal Highway Administration. (2010). *2010 status of the nation's highways, bridges, and transit: Conditions and performance*, Washington, D.C.: US Department of Transportation. Retrieved from <http://www.fhwa.dot.gov/policy/2010cpr/chap2.htm>
- Fenn, A. J., Hurst, P. T., Pacheco, J., Cornick, M., & Parad, L. I. (2013). Ultrawideband cavity-backed resistively loaded planar dipole array for ground penetrating radar. *IEEE International Symposium on Phased Array*, Waltham, MA.
- Floyd, L. R. (1978). *Geodetic bench marks*, Rockville, MD: US National Oceanic and Atmospheric Administration.
- Herman, H. (1997). *Robotic subsurface mapping using ground penetrating radar (CMU-RI-TR-97-19)*, Pittsburgh, PA: Carnegie Mellon University.
- Hoekstra, P., & Delaney, A. (1974). Dielectric properties of soils at UHF and microwave frequencies. *Journal of Geophysical Research*, 79(11), 1699–1708. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1029/JB079i011p01699/full>
- Jol, H. (2008). *Ground penetrating radar theory and applications*, Amsterdam: Elsevier Science.
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *IEEE International Conference on Neural Networks IV*, Perth, WA. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=488968&tag=1
- Kennedy, S., & Rossi, J. (2008). Performance of a deeply coupled commercial grade GPS/INS system from KVH and NovAtel Inc. *IEEE/ION Position, Location and Navigation Symposium*, Monterey, CA.
- Lever, J. H., Delaney, A. J., Ray, L. E., Trautmann, E., Barna, L. A., & Burzynski, A. M. (2013). Autonomous GPR surveys using the Polar Rover Yeti. *Journal of Field Robotics*, 30(2), 194–215.
- Levinson, J., & Thrun, S. (2010). Robust vehicle localization in urban environments using probabilistic maps. *International Conference on Robotics and Automation*, Anchorage, AK.
- Levinson, J., Montemerlo, M., & Thrun, S. (2007). Map-based precision vehicle localization in urban environments. *Proceedings of Robotics: Science and Systems*, Atlanta, GA, USA.
- Milford, M., & Wyeth, G. (2012). Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. *IEEE International Conference on Robotics and Automation (ICRA)*. St. Paul, MN, USA. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6224623&tag=1
- Moulton, L. K. (1980). *Highway subdrainage design (FHWA-TS-80-224)*, Washington, D.C.: US DOT Federal Highway Administration.
- National Highway Traffic Safety Administration. (2011). *NCSA Data Resource Fatality Analysis Reporting System*. Retrieved 2013, from <http://www.fars.nhtsa.dot.gov/Main/index.aspx>
- Nuske, S., Roberts, J., & Wyeth, G. (2009). Robust outdoor visual localization using a three-dimensional-edge map. *Journal of Field Robotics*, 26(9), 728–756.
- Saarenketo, T., & Scullion, T. (2000). Road evaluation with ground penetrating radar. *Journal of Applied Geophysics*, 119–138.
- Shi, Y., & Eberhart, R. C. (1999). Empirical study of particle swarm optimization. *Proceedings of the 1999 Congress on Evolutionary Computation*. IEEE. Washington D.C. <http://staff.washington.edu/paymana/swarm/shi99-ccc.pdf>
- Stanley, B., Cornick, M., & Koechling, J. (2013). Ground penetrating radar based localization. *Ground Vehicle Systems Engineering and Technology Symposium*, Troy, MI.
- Unnikrishnan, R., & Kelly, A. (2002). Mosaicing large cyclic environments for visual navigation in autonomous vehicles. *IEEE International Conference on Robotics and Automation (ICRA)*, Washington, D.C.
- Williams, R. (2012). An autonomous robotic platform for ground penetrating radar surveys. *Geoscience and Remote Sensing Symposium (IGARSS)*, Munich, Germany, 22–27 July 2012.
- World Health Organization. (2013). *Global status report on road safety 2013: Supporting a decade of action*, Geneva: World Health Organization. Retrieved 2013, from http://www.who.int/iris/bitstream/10665/78256/1/9789241564564_eng.pdf
- Yamauchi, B. (2010). Fusing ultra-wideband radar and LIDAR for small UGV navigation. *Proceedings of SPIE 7692, Unmanned Systems Technology XII*, Orlando, FL, 6–9 April 2010.