Error Mitigation in Roughness Measurements

Zhuosong Wang

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
In
Mechanical Engineering

John B. Ferris

Alfred L. Wicks

Saied Taheri

June 9, 2014
Blacksburg, VA

Keywords: roughness, IRI, terrain measurement

Error Mitigation in Roughness Measurements

Zhuosong Wang

Abstract

Road roughness is an important factor in determining the quality of a stretch of road. The International Roughness Index, a specific measure of road roughness, is widely used metric. However, in order to measure roughness, an accurate road profile must exist. To measure the roads, terrain profiling systems are commonly used. Modern systems based on laser scanners and inertial navigation systems (INS) are able to measure thousands of data points per seconds over a wide path.

However, because of the subsystems in the profiling systems, they are susceptible to errors that reduce the accuracy of the measurements. Thus, both major subsystems – the laser and the navigation system – must be accurate and synchronized for the road to be correctly scanned. The sensors' mounting was investigated to ensure that the vehicle motion is accurately captured and accounted for, demonstrated in the Vehicle Terrain Performance Lab's (VTPL) Ford Explorer profilometer. Next, INS errors were addressed. These may include drift in the inertial measurement unit or errors due to poor reception with the global navigation satellite system. The solution to these errors was demonstrated through the VTPL's HMMWV profilometer.

Acknowledgements

I would like to thank…

My advisor, Dr. John Ferris, for his guidance throughout my work. I am grateful for his supporting my academic interests and challenging me to do my best.

My committee members, Dr. Saied Taheri and Dr. Alfred Wicks, for their input to my thesis and project.

My research sponsors, the Center for Tire Research (fund number 500019), US Army TARDEC (417270), Chrysler (458081), and Virginia Tech (441703) for making my research possible.

All others who have helped with my research, including my friends and the other members of the Vehicle Terrain Performance Lab (VTPL).

My family, for their continual support. Their encouragement has helped bring me to where I am today.

Table of Contents

List of Figures

1. Introduction

This thesis focuses on characterizing and mitigating the effects of errors in measuring road roughness. Currently, the international roughness index (IRI) is widely used in the road measurement community as a roughness metric [1]. However, because of numerous sensors involved in road profiling, there are many instances where errors could be introduced. Regardless of the eventual application, it is necessary to ensure the accuracy of the data.

1.1 Motivation

Automobiles receive an input excitation from the roads that they travel on. Knowledge of this excitation can be determined in many ways. From simple, 2 dimensional devices to characterize roughness to modern, 3 dimensional road profiling systems, it is important that the data collected accurately reflects the actual road condition. Obviously, vehicle based measurement systems are also subjected to these input excitations, which may cause unwanted vibrations and movements that would distort the measurement. Thus, this movement must be documented taken into account to provide the most accurate data.

One method to track sensor movement is to use an inertial navigation system (INS). This generally utilizes an inertial measurement unit coupled with a satellite navigation system in order to maintain knowledge of the vehicle's position and attitude. By subtracting the vehicle's motion from the measurement data, the terrain profile can be generated without any unwanted environmental input. However, the additional sensors require a high degree of accuracy; if their measurements were incorrect, then the corrected data would still have errors.

Even with good profile data, they need to be effectively processed to be useful. As modern systems are capable of capturing over a million points per second, the raw data, e.g., a point cloud of road height, is not inherently useful. The proper manipulation of this data allows users to better see how it may represent the terrain. This presents a multitude of uses for the data. One such application is to determine road roughness. Transportation departments often use this to metric to track road quality. In contrast, terrain models can also provide automobile designers the ability to make informed decisions about necessary

changes to make in the chassis design. Regardless, the usefulness of the data hinges on both accurate data and good processing techniques.

1.2 Problem Statement

The IRI, a metric for road roughness, is widely used to determine road conditions. To calculate the IRI, a road profile must be provided. Modern methods of measuring the road profile can involve laser scanners coupled with an INS. This equipment requires a high degree of accuracy to synchronize with each other and provide good data. However, both the laser scanners and the INS may have errors that would negatively affect the quality of the collected data.

First, the sensors must be mounted in such a way that the INS can track the movement of the laser scanners. This ensures that any changes to the scanner due to vehicle motion are not erroneously added to the measured profile. This will be investigated in a case study regarding the VTPL's Ford Explorer profilometer, where a new mounting frame was developed for attaching two Selcom Roline 1130 lasers.

Next, the INS may have its own measurement errors, from either the GPS or IMU. A problem was detected with the VTPL's HMMWV profilometer in early 2013 - after collecting data for extended periods the terrain surface displayed discrete steps not found in the actual road. Due to the nature of the problem, there seemed to be an issue with the IMU. Because the IMU tracks the attitude of the HMMWV, any errors in the IMU measurements would affect the perceived location of the vehicle, which would ultimately cause the discrepancies between the actual road and measured data. This error was investigated

1.3 Thesis Statement and Scope of Work

Thesis Statement: The effects of road profiling measurement errors on roughness metrics can be identified and mitigated with proper hardware and software correction.

The objective of this work is to develop methods of ensuring hardware integrity in order to allow for accurate road profile measurements. Because a profiling system is made up of several different components taking different measurements, error in any one subsystem will contribute to the overall error. This work focuses on issues found in the VTPL's Ford Explorer and HMMWV profiling systems, including refining a calibration method for the INS. While the scope of the specific hardware work is limited, the

calibration method has a broader application of any profiling system using terrain measuring sensors and an INS.

1.4 Main Contributions

The main focus of this research was to characterize and minimize errors found in IRI measurements. This involved investigating the hardware and finding how to prevent the occurrence of common errors, as well as determining the effects of these errors on the roughness metric. Ultimately, the increased measurement accuracy should help result more accurately defined roughness in roads.

1. Method of identifying and rectifying errors from the process of measuring terrain profiles
    a. Movement in the sensor mounting frame
    b. Inaccuracies introduced by the INS
2. Demonstrated the effects of noisy data on IRI values

1.5 Thesis Outline

This thesis is organized in the following manner. The first chapter introduces the work and shows how it fits into the current profiling technologies. The second chapter contains a review of background information that is relevant to road roughness and road profiling. The third chapter addresses errors found in mounting the sensors, focusing on the Ford Explorer profiling system. Chapter four investigates INS based error, specifically IMU drift and INS alignment issues as found in the HMMWV profilometer. A summary of research from the thesis is given in chapter five. Finally, chapter six suggests some future work, including an improvement to the IRI.

2. Background

2.1 Road Roughness

Based on historical engineering methods, the international roughness index (IRI) has been the standard measure of road roughness for many years [1]. Developed by Sayers, the IRI is used to quantify the roughness of road profiles [2, 3]. It is defined as the average of the absolute suspension travel per unit of longitudinal distance, usually in equivalent units of mm/m and m/km (or in/mi) [4]. In other words, the $IRI = \frac{1}{L}\int_0^T |\dot{z}|dt$, where $\dot{z}$ is the velocity of suspension travel, L is the length of road, and T is the time it takes to travel from the beginning to the end of the road.

The IRI specifically targets passenger vehicles by using the golden quarter car model in its calculations. The golden quarter car model has the following parameters [5]:

Parameter	Value
Spring stiffness (k_s)	63.3 m/s²
Tire stiffness (k_t)	633 m/s²
Damping (c)	6 m/s
Unsprung mass (m_u)	0.15 kg
Sprung mass (m_s)	1 kg

These parameters are normally used in a way such that they are normally the ratio over the sprung mass; as such, the value of the sprung mass is 1. The quarter car is simulated to travel at 80 km/h. The setup of the quarter car model is shown in Figure 1, where z_u and z_s are the positions of the unsprung and sprung masses, z is the suspension travel, and x is the input from the road. Finally, a 250mm filter is applied to the road profile to simulate the contact size of the tire.

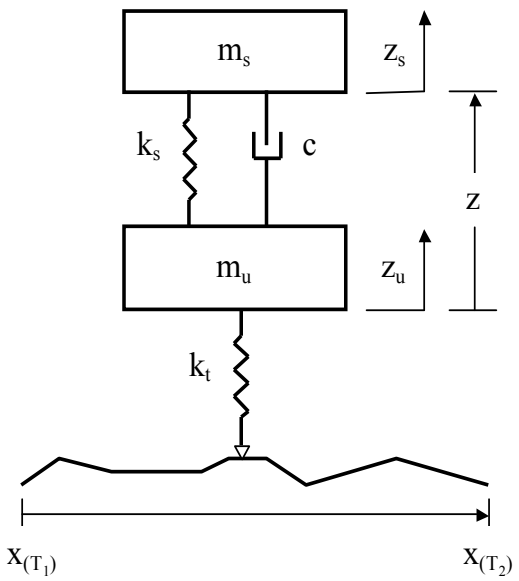

Figure 1: quarter car model

There are two main disadvantages that must be considered when using the IRI. First, the IRI is an average over the length of the road. This means that multiple different roads with very different profiles make produce identical IRI values (e.g., one large event in one otherwise smooth road versus another road that is more rough throughout its length) [6]. Second, the speed of 80 km/h at which the IRI is calculated is not necessarily applicable to all roads [1], especially for surface streets, where speeds are typically 50 km/h, and highways, where speeds exceed 110 km/h.

2.1.1 Sample Road Roughness

Throughout this work, three sample roads will be referred to for demonstration. These three roads comprise of an unused airport runway in Danville, VA, a brick road in Danville, VA, and a stretch of Main Street in Blacksburg, VA. Their roughness data came from 3 dimensional profiles taken by the VTPL; each road had 7 linear profiles, spaced evenly between 0.7 meters to either side of the center. The following table shows the average of the seven IRI values for each road, as well as their standard deviation. These IRI calculations used the ASTM (E1926-08) method, transcribed to Matlab.

Road name	IRI (m/km)	st. dev	Road type
smooth	1.88	0.136	Main St
medium	3.19	0.143	Old runway
rough	9.69	0.446	Brick path

The profiles of the roads are also plotted. Figure 2 shows each of the sever profiles of Main St. (smooth), Figure 3 shows the profiles of the airport road (medium), and Figure 4 shows the profiles of the brick road (rough).

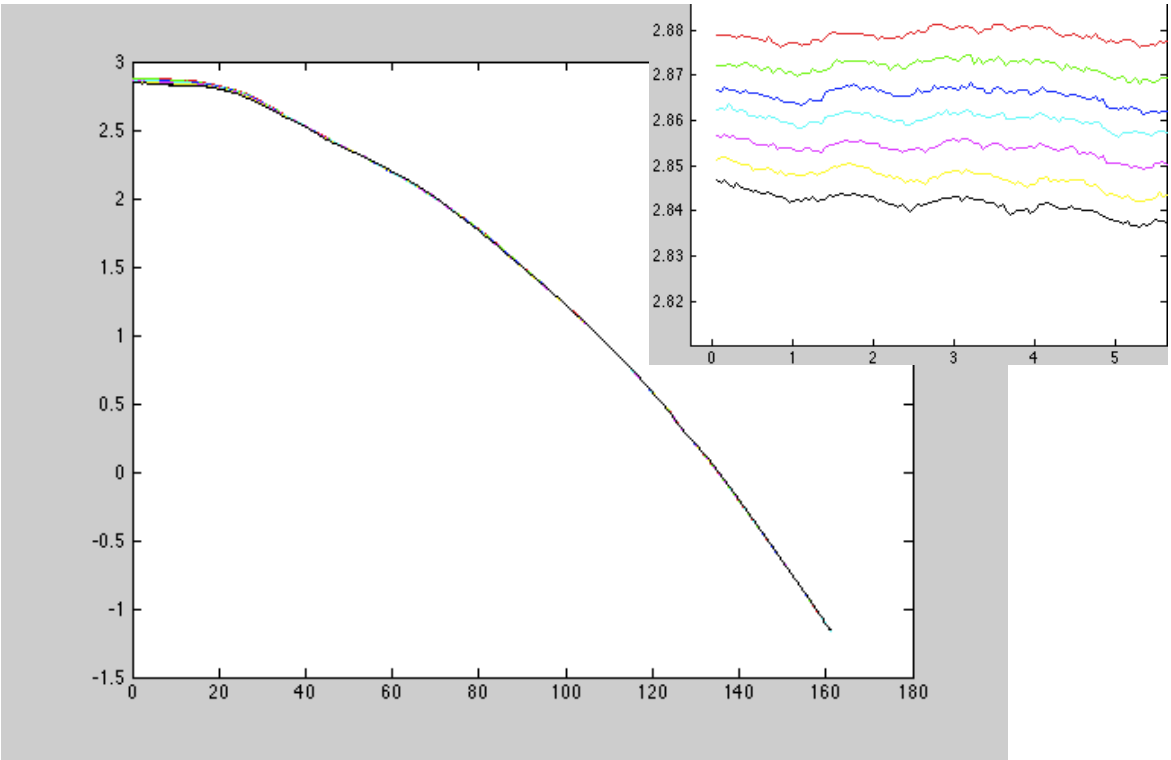

Figure 2: Main St. Road height (m) vs horizontal location (m). The overlaid road is zoomed in on the first 5 meters.

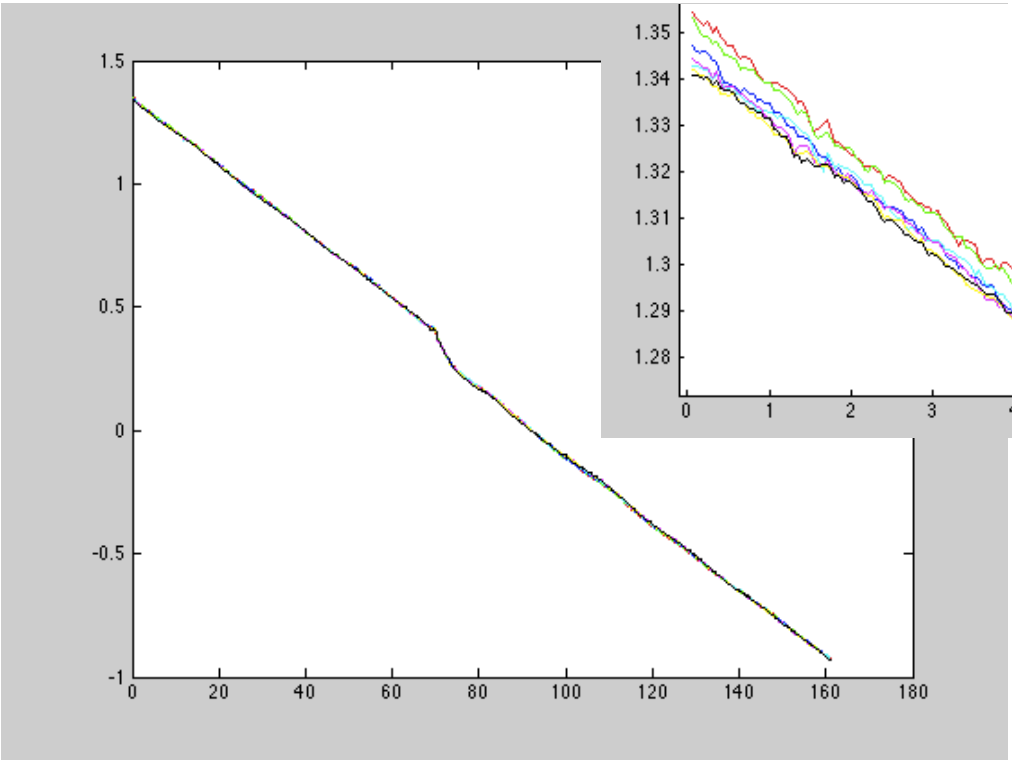

Figure 3: Airport runway. Road height (m) vs horizontal location (m). The overlaid road is zoomed in on the first 5 meters.

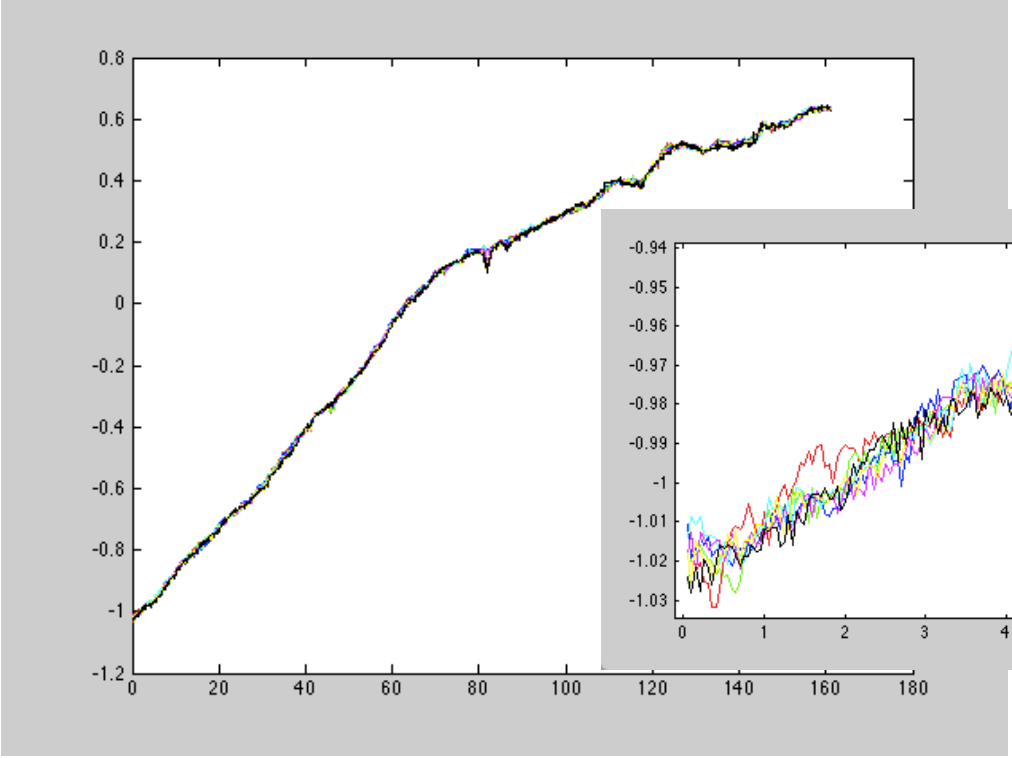

Figure 4: Brick path. Road height (m) vs horizontal location (m). The overlaid road is zoomed in on the first 5 meters.

2.2 Road Profiling

In order for the road profiling system to be able to effectively measure the road, it needs to not only measure the road surface, but also keep track of its location and orientation. While this is not absolutely necessary for a linear profile (assuming constant velocities), it is important for a three-dimensional road profile [7, 8]. Because the measurement of the road surface is taken with respect to a moving datum (the scanning laser is mounted to the vehicle, which moves), the measuring vehicle's location and attitude significantly affect the processing of the road surface data. A global navigation satellite system (GNSS) is generally used to keep track of location; this is complemented by an inertial navigation system (INS) that also tracks the vehicle's attitude data.

2.2.1 Laser scanners

Using lasers to measure distance is a popular method of non-contacting terrain measurement. As terrain measurement shares many characteristics with light radar (LIDAR) systems, there is an overlap in technology between the two fields [9]. Most terrain profiling lasers use one of two methods, time-of-flight or triangulation.

A time-of-flight system uses the speed of light to measure the distance of the object of interest to the sensor. Similar to how radar works, a pulse of light is emitted, which then bounces off the object and back into a sensor. By finding the time between the start and receiving of the pulse, the distance travelled can be calculated [10]. See Figure 5 for a diagram of this type of system.

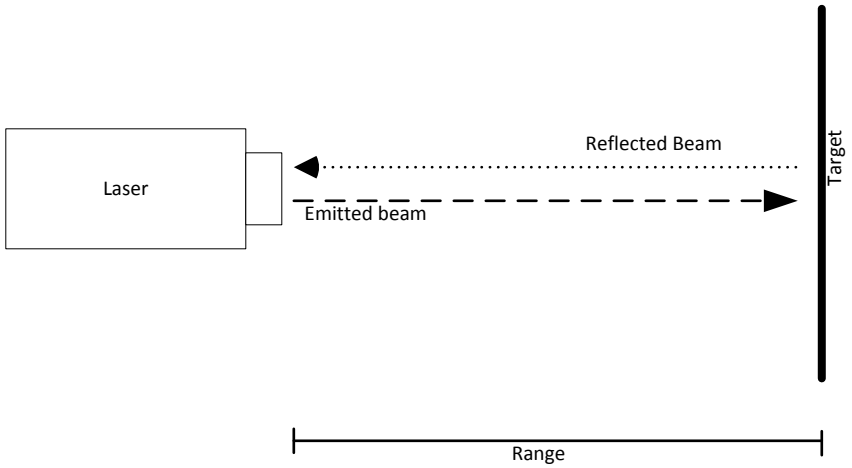

Figure 5: This demonstrates a time of flight laser that uses the total time of travel.

Another time-of-flight system uses a continuous laser beam with an embedded low-frequency pulse. By finding the difference in the phase between the returning signal and the output signal, a distance can be calculated [10]. This is shown in Figure 6.

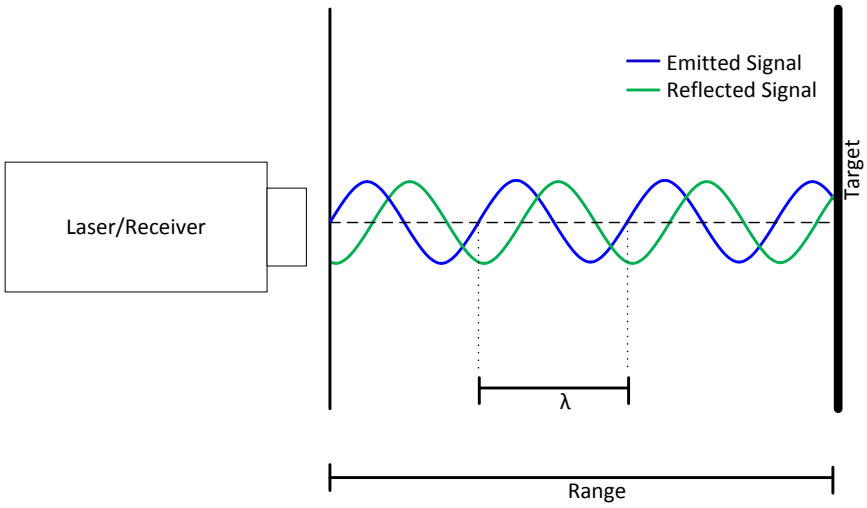

Figure 6: This represents a time of flight laser that measures the phase difference between the emitted and reflected signals.

Because the speed of light is so fast, these sensors require accurate internal clocks and mirrors to maintain their accuracy. Furthermore, the scanning speed is generally more limited. However, the time-of-flight systems are capable of measuring longer distances [11].

A triangulation laser system uses a laser and a separate camera to measure the distance. As the laser and camera are at known distances and angles to each other, the projection of the laser on the terrain will occur at different locations, depending on its height. From the angular difference that the camera sees, it can then calculate the unknown distance from the known values [12]. A diagram of a triangulation laser system is shown in Figure 7.

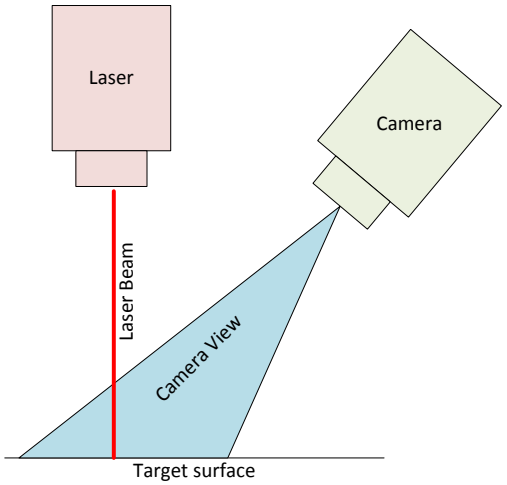

Figure 7: This shows the components of a triangulation laser scanner.

The triangulation systems are capable of sensing a wide line, with high precision. This results in a very high data collection rate. However, the maximum range of these systems is generally much smaller [12].

2.2.2 GNSS

A global navigation satellite system utilizes a collection of satellites orbiting the Earth to determine a receiver's location (in latitude and longitude) and elevation. While there exist multiple systems, such as the European Union's Galileo and Russia's GLONASS, because the work in based in the United States, only the American based Global Positioning System (GPS) will be discussed.

Each GPS satellite continuously transmits a signal, which is read by a receiver on Earth. A diagram of the GPS satellites is found in Figure 8 [13]. By finding the length of time each signal traveled, the receiver is able to triangulate its exact position on Earth. Because radio signals travel at the speed of light, the satellites require precision timing. Thus, they carry multiple atomic clocks onboard to ensure the proper synchronization of time throughout the system. The receiver is also able to calculate the proper time, given one additional satellite (3 for position plus 1 for time) [14].

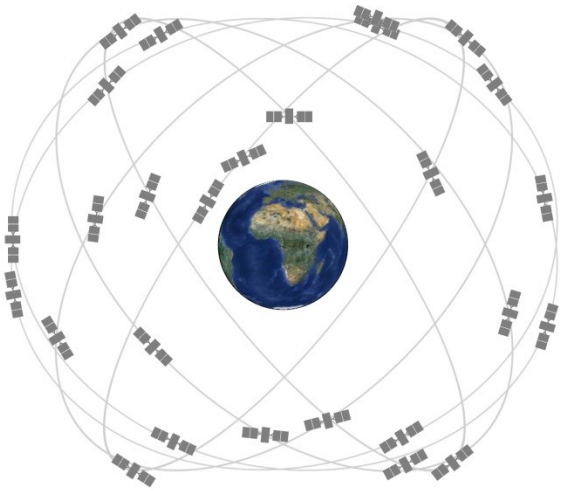

Figure 8: This figure represents the GPS satellites in orbit around the earth.

In 2000, the US government ended the former policy known as selective availability (SA), which greatly limited civilian usage of GPS by forcibly increasing the signal error [15]. A sample set of data collected by the National Oceanic and Atmospheric Administration shows that before selective availability was ended, 95% of the calculated positions in a day fell in a circle with a radius of 45 meters, while this decreased to 6.3 meters after its expiration, seen in Figure 9 [16]. Modern high-quality civilian GPS receivers can have accuracies of less than 5 meters [17].

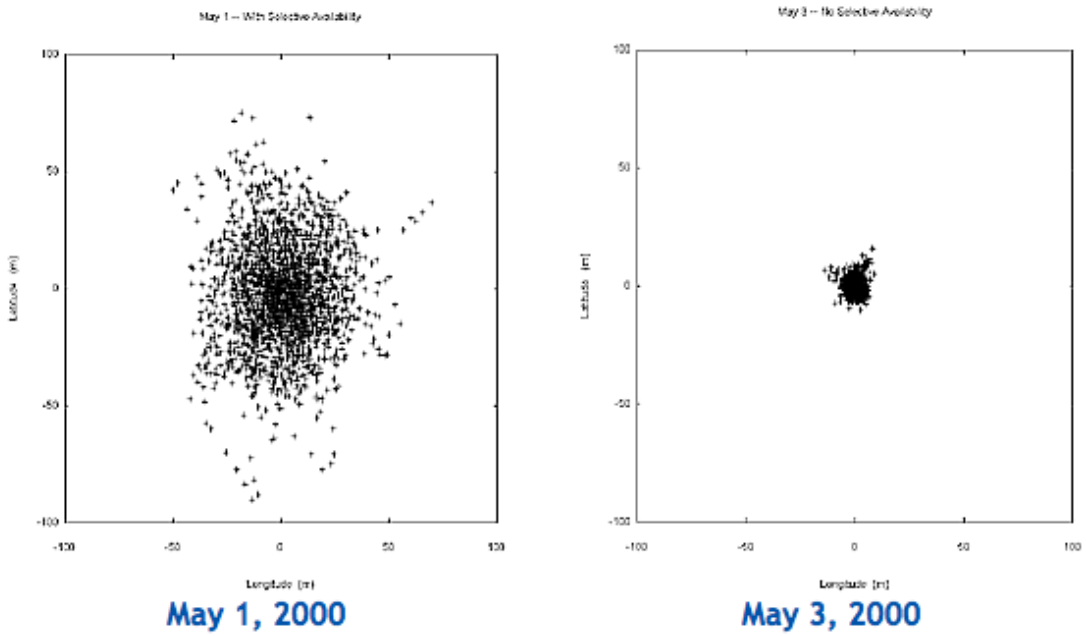

Figure 9: NOAA data on GPS positioning before (left) and after (right) the end of selective availability

Because of the limitations of SA, methods were developed to allow for more accurate positioning. One such method of increasing the accuracy of GPS is differential GPS (DGPS), which uses another signal from a known location. However, even today, because of the precision that is no longer limited by SA, DGPS allows for much better accuracy, down to less than 10 cm [18].

DGPS makes use of signals collected from one base station to improve the accuracy of nearby GPS receivers. Generally, the base station remains at a surveyed location, and its known absolute position can be compared to the calculated position from the GPS satellites or other corrections methods, allowing signal correction information to be generated. This allows for any compatible GPS receiver located nearby to utilize this correction information generated to eliminate error in the signal [19]. These errors may come from a variety of sources, including multipath interference (signals bouncing off nearby objects such that the receiver collects multiple signals), atmospheric interference (slowing of the signal because of different atmospheric conditions), or satellite geometry (the further the satellites are from each other, the better the triangulation solution) [20]. Figure 10 demonstrates one base station providing correction information to a remote receiver [21].

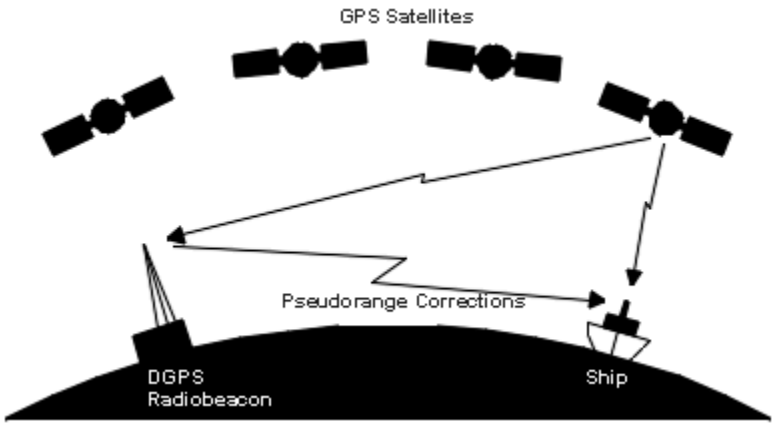

Figure 10: This shows how a base station sends correction information in a DGPS setup

One method of implementing DGPS is the installation of a system of base stations spread across a designated service area. These systems, such as the Federal Aviation Administration's (FAA) Wide Area Augmentation System (WAAS) or the NOAA's Continuously Operated Reference Stations (CORS), provide a multitude of base stations located around the United States [22]. With the assumption that any GPS receivers located

within that service area are receiving data from the same group of satellites and experience similar atmospheric conditions, the correction information sent by the base stations will increase the individual receivers' accuracies. Alternatively, a temporary base station can be set up such that a better solution can be post-processed. This is the setup that the VTPL profilometer uses.

2.2.3 IMU

The inertial navigational system uses an inertial measurement unit (IMU) to keep track of orientation. Commonly used in aircraft, this is used in the Profilometer because the vehicle and sensor movement must be accounted for in the laser measurements [23]. Using three orthogonal accelerometers (one in each x, y, and z orientation) to measure acceleration and three orthogonal gyroscopes to measure rotation, the IMU can keep track of any body motion in space. Due to the necessary precision of the IMU (because a small angular change can result in a large linear change with a long moment arm), it is also subject to errors. For example, a bias develops each time the IMU is powered on. The axes of the internal IMU sensors also may not be perfectly perpendicular to each other. However, these errors can generally be mitigated [24]. A rigorous calibration process can address issues present in the hardware, while alignment algorithms could correct errors encountered during usage.

3. Equipment Mounting Errors

3.1 Introduction

This chapter addresses errors that are caused by the sensors and how they are affixed to the vehicle. Lasers sensors used in road profiling measure the relative distance from the sensor to the road surface. However, if the vehicle moves vertically (e.g., due to a bump), the collected measurement will no longer reflect the actual road surface. Thus, an INS is used to track the movement of the sensors to cancel all movement. In order for this to be effective, there must a rigid connection between the INS and other sensors. If the sensors move relative to the INS, then it no longer correctly tracks the actual sensor movement. Figure 11 shows a road profiling system where the INS and sensor are connected by a small metal bar. While this may be sufficient during most situations, any large acceleration will very likely cause relative movement between the two sensors.

Figure 11: A road profiling system. The INS is in the red box, and a sensor is in the blue box. Note they are connected by a relatively thin piece of metal. Photo by author, 2014.

To ensure that the connection is sufficiently rigid to collect data, the equipment mounting frame design must be analyzed to ensure enough stiffness. Furthermore, a modal test should be performed and compared to the expected vehicle's excitation frequencies.

3.2 Ford Explorer Mounting Frame

A mounting frame was designed for the Roline lasers by Rainey [25]. The original configuration had the GPS/IMU mounted on a frame on top of the rear of the vehicle, as this was close to the location of the previous Phoenix Scientific laser. However, due to the limited width of the Roline lasers (about 10 cm), they were placed in front of the vehicle in the track of the wheels. However, this separation from the positioning system meant potentially greater inaccuracies, so a closer mounting location for the GPS/IMU was desired. Furthermore, a modular system would allow for the quick removal of the whole Profilometer system from the vehicle.

The original design, shown below in Figure 12, would allow for the GPS/IMU to be placed at the middle of the frame. The two Roline lasers would be placed on the same level, but in front of the track of the wheels. Finally, it has mounting spots for up to five additional Gocator lasers on the top front of the rack, which would scan a full lane width.

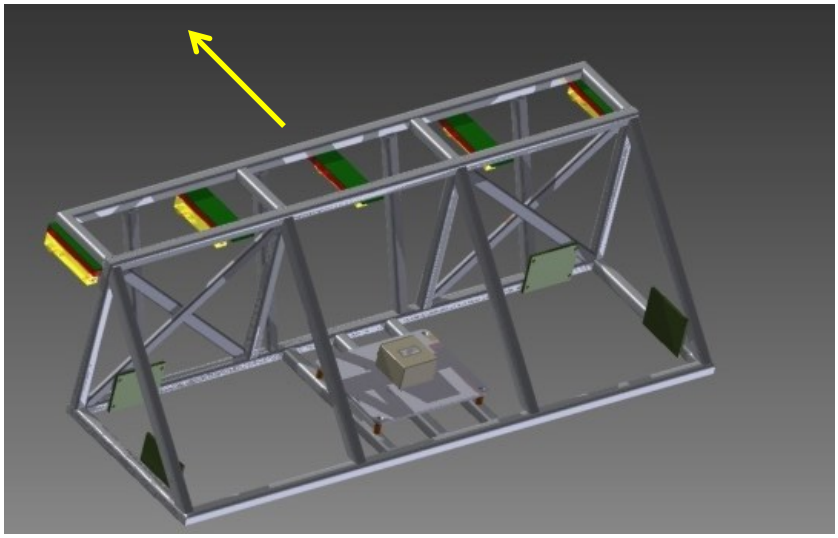

Figure 12: The original design for the mounting frame. This would be mounted in front of the vehicle, where forward is in the direction of the arrow.

However, due to a miscommunication with the machine shop, the design changed slightly between the design and the constructed frame. The diagonal supports in the front of the frame were placed all the way on the ends. Furthermore, because of the desire to keep the Roline lasers along the track of the wheels, the proposed mounting method was changed. Two supports were added on the bottom part of the frame to allow the Roline lasers to stay in line with the wheels. The new design is shown in Figure 13.

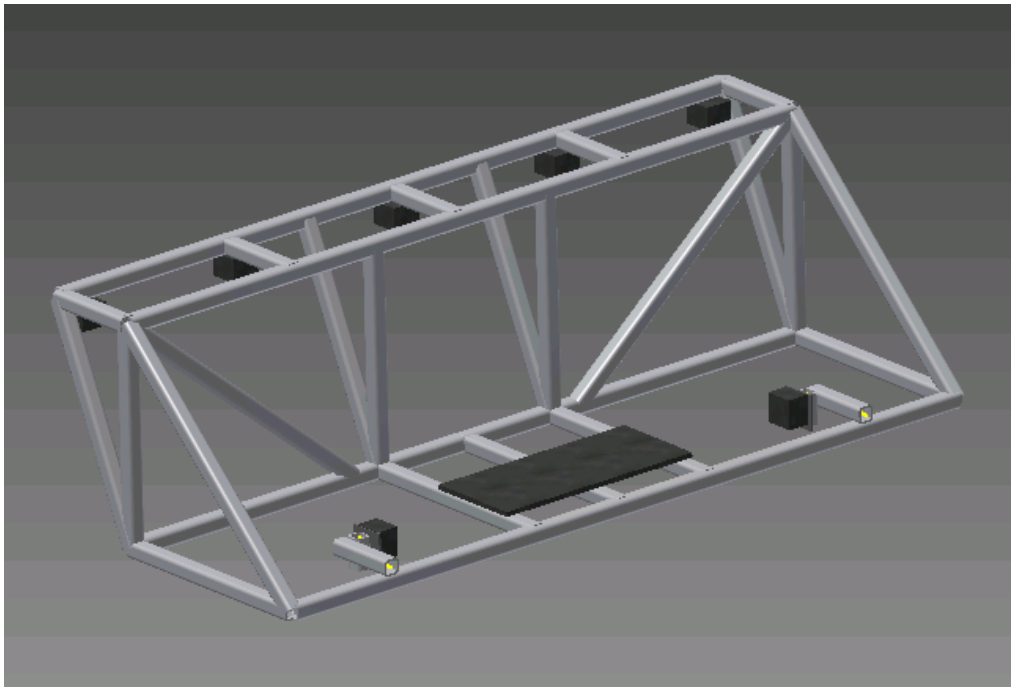

Figure 13: This is the actual design of the mounting frame.

3.2.1 Design Considerations

The assumption of rigid-body dynamics is important for the IMU to track the movement of the sensors. Because the IMU is able to sense changes in its own orientation, the sensors (GPS antenna and Roline lasers) must remain in the same relative locations despite any loading. This was the primary reason behind the design of the frame – it was overbuilt to minimize any deflections that occur during use. Due to the changes to the frame from Rainey's design, a new finite element analysis was performed. Below, Figure 14 and Figure 15 show the results of the finite element analysis. The arrows pointing to parts of the frame are the equivalent to the loading of the attached sensors, while the arrow floating near the center represents gravitational loading. Padden and Griffin measured 25 passenger vehicles and found that the maximum experienced vertical acceleration was less than 1 m/s^2 [26]. Uys, Els, and Thoresson further found that vertical acceleration while off-road did not exceed 8 m/s^2 [27]. However, a loading of 5 times normal gravity was simulated to include any possible circumstances, regardless of how unlikely. This was also the value used in Rainey's testing [25].

The maximum deflection was found to be 0.5 mm. This is very unlikely to occur in real world conditions, as it in improbable to achieve acceleration on the order of 5 times that of gravity. While there is not a specific limit for maximum allowable deflection, the

acceptable amount depends on the application. For example, measuring very smooth roads (requiring more detail) would require more precision than a very bumpy road, which may not need to be as precise (smaller percentage change compared to the surface). A rule of thumb is to ensure the expected deflection to be no more than 20% of the desired precision. For example, measuring a smooth road may require data on the order of 0.1 mm, which means deflection less than 0.02 mm. Using the 1 m/s² acceleration from Padden and Griffin would translate to a deflection of 0.01 mm, which is acceptable.

Additionally, it is necessary to ensure that the frame's natural frequencies of vibration are not close to the vehicle's natural frequencies of excitation (usually around 2 Hz and 15 Hz). These values were found to be 1.2 and 12 Hz by Smith [28]. This is to prevent any uncontrolled excitation of the frame while the vehicle is in motion. The first natural frequency was found to be 33 Hz, which is well above the expected excitation frequencies.

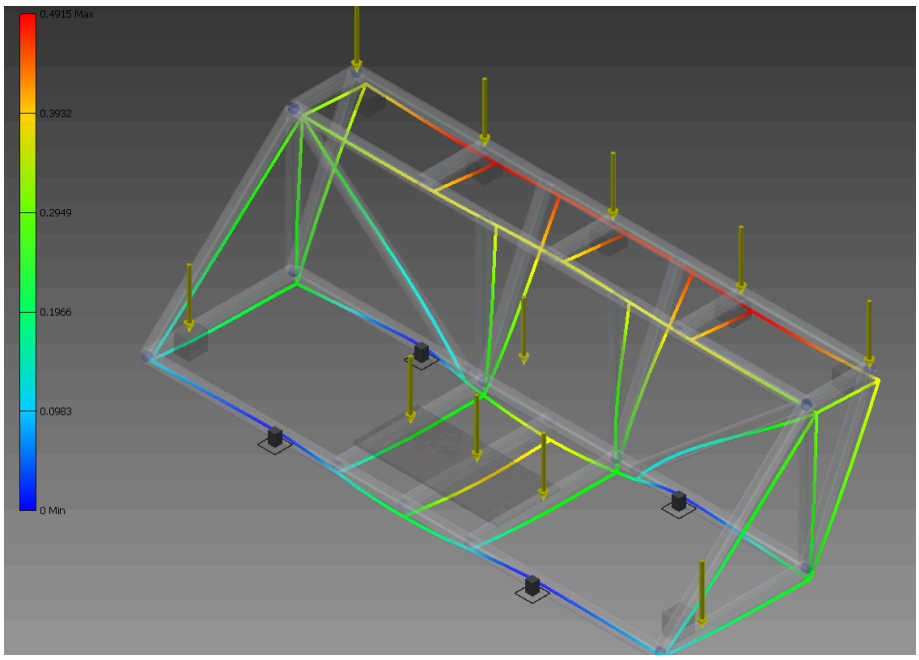

Figure 14: This shows the results of the frame under load.

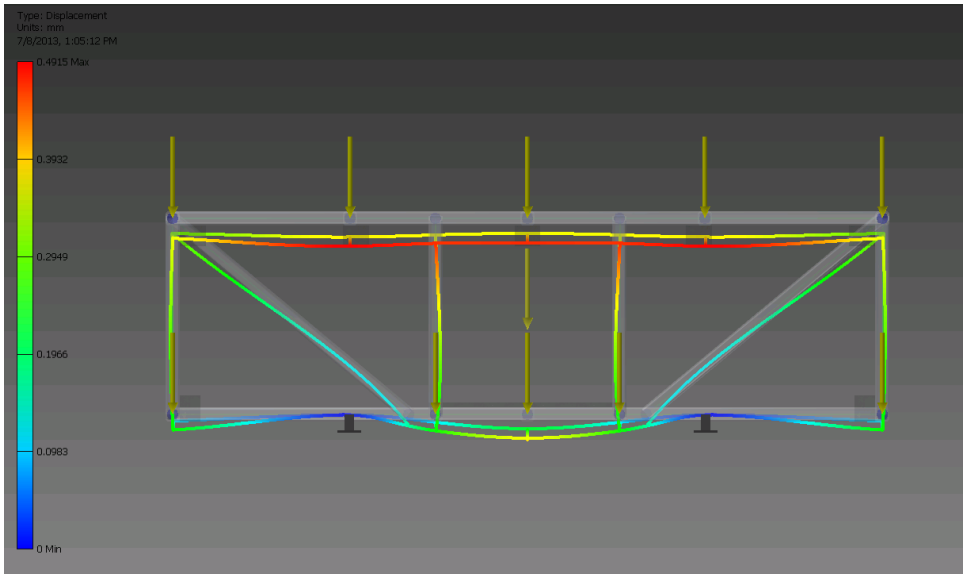

Figure 15: This provides another viewpoint of the frame under load.

Previous vertical acceleration data collected on the Ford Explorer show that the vehicle excitation occurs at 1.5, 6, 13.7, and 35.9 Hz, shown in Figure 16.

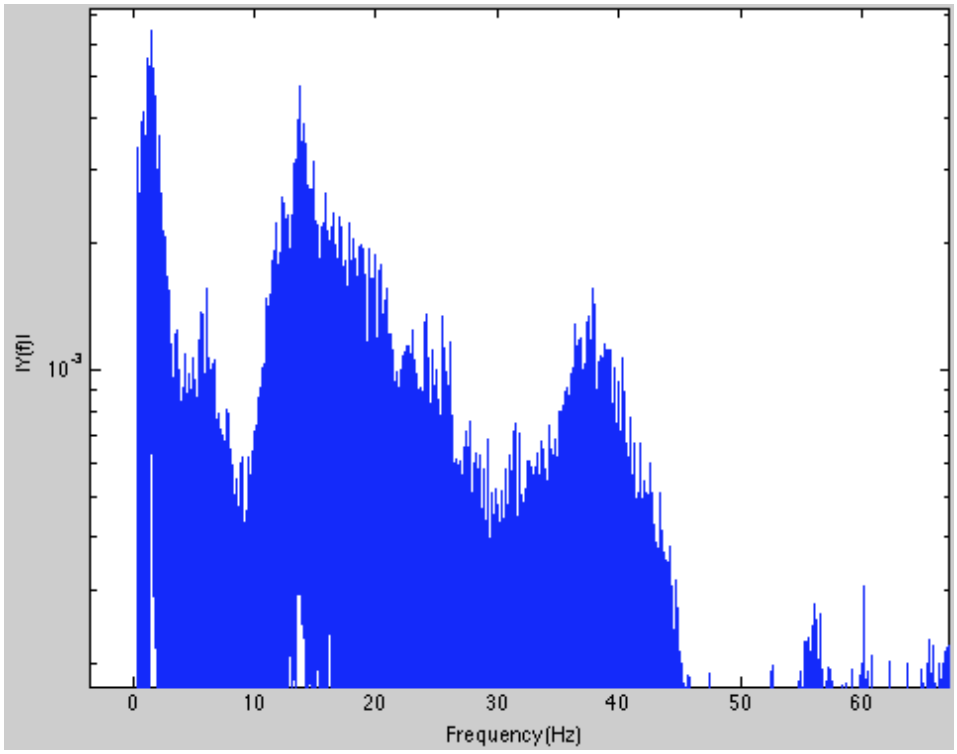

Figure 16: FFT of acceleration data. Note the peaks at 1.5, 6, 13.7, and 35.9 Hz.

When this data is integrated into displacement, the peak at 6 Hz drops out, shown in Figure 17.

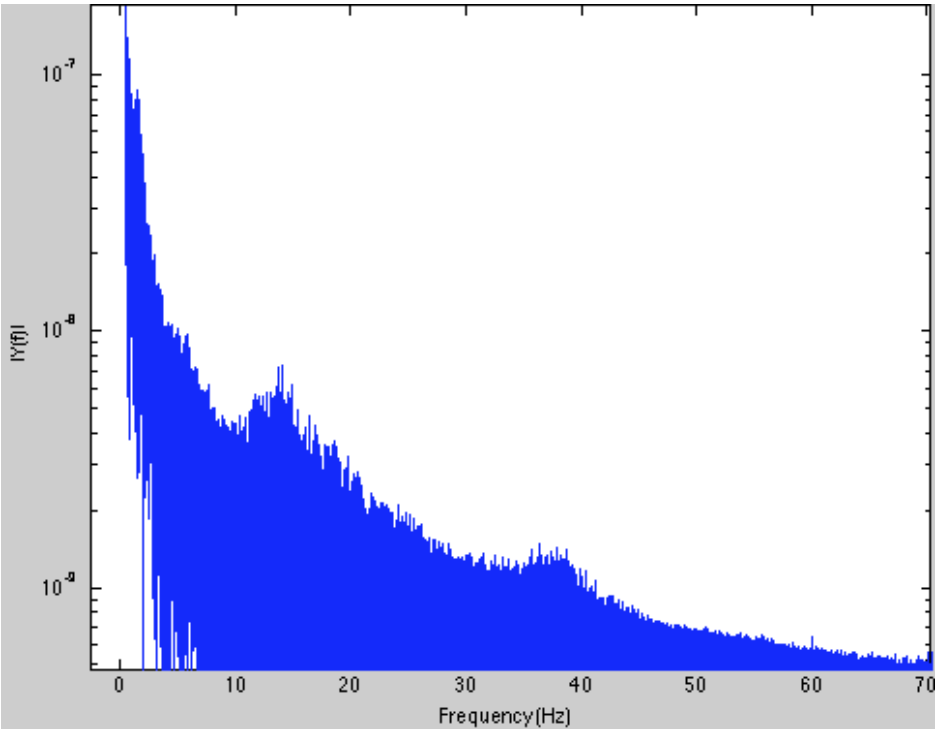

Figure 17: FFT of displacement, taken from accelerometer data.

The peaks at roughly 1.5 Hz (ride frequency), 14 Hz (wheel hop frequency), and 37 Hz (chassis vibration) remain apparent as expected. Furthermore, it is important to note that there is no significant vibration after 50 Hz. This means that the IMU, which records data at 100 Hz, should not pick up any aliasing due to higher frequencies. Coupled with the IRI's frequency response, shown in Figure 18, this means that the frequencies of interest are less than 20 Hz [29].

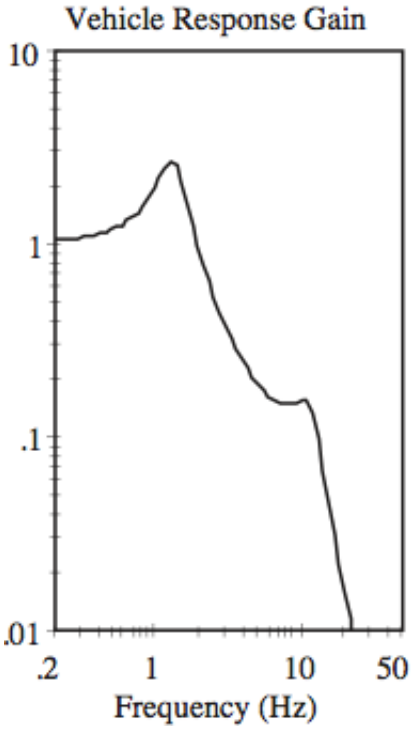

Figure 18: Frequency response of acceleration of the IRI

3.2.2 Manufacturing Validation

The inertial navigation system of the Profilometer relies on knowing the location offsets between the IMU and the other sensors (GPS antenna and scanning lasers). Because of minor tolerance issues that were likely encountered during the manufacturing process, it is very unlikely that the actual dimensions perfectly match those in the drawings. Thus, in order for the system to be able to accurately track the movement of the sensors, the actual locations of the sensors must be determined. However, due to the unique geometry of the system, it would be very difficult to use traditional methods to measure these offsets between the IMU and other sensors. Alternative methods of measuring the frame were considered, such as using three-dimensional scanners, but ultimately, it was decided that using a FARO Arm would provide the most accurate measurements. The FARO Arm is a 6 degree-of-freedom mechanical arm that is able to calculate distances from the encoders located in the joints. The base of the arm is clamped onto a fixed location, and as the user manually moves the measurement tool at the end of the arm, the locations are calculated from the measured angles (from encoders) and arm segment lengths. Because only the distances relative to the IMU are important, the arm

was clamped onto the frame (to prevent any movement) and used to measure the locations of the various sensors. This data was then used to create a new coordinate system with the origin at the center of the IMU, from which the final offsets were determined.

3.2.3 Modal Analysis

Because of the nature of the simulated modal analysis, it is very likely that the real results are considerably different. For example, the four brackets on the bottom of the frame simulate the attachment points of the frame to the vehicle. This was modeled as fixed constraint, while in practice, the system may more closely resemble free vibration. Thus, a modal response test was performed on the constructed frame to find the actual natural frequencies (and hopefully confirm the results of the finite element analysis). To simulate testing conditions, the frame and sensors were attached as they would be in the case of a normal test. Figure 19 shows this setup.

Figure 19: This shows the normal data taking setup of the equipment frame. The two yellow Roline lasers are seen on the sides in the wheel track. The black Pelican case in the middle houses the IMU and GPS receiver, while the circular GPS antenna is on the top of the frame.

Next, since the primary interest lies in the relative movement of the Roline lasers with respect to the IMU, the accelerometers were placed on these three sensors, seen below in Figure 20.

Figure 20: This shows the accelerometers attached to the Roline laser (left) and IMU (right).

These three signals and the output from an impact hammer were fed into a SigLab digital signal analyzer, which recorded the data. The system was excited from one of the two steel supports connected to the vehicle, as this is where the excitation would occur during regular use. Each set of data was an average of five test results to minimize any noise that may have affected the results. Multiple sets of data were collected, all showing similar results. The following plot, Figure 21, shows the frequency responses of the three accelerometers. The coherence is shown in Figure 22. The coherence is high except at the very low frequencies (less than 1 Hz) and the anti-resonance at 11 Hz.

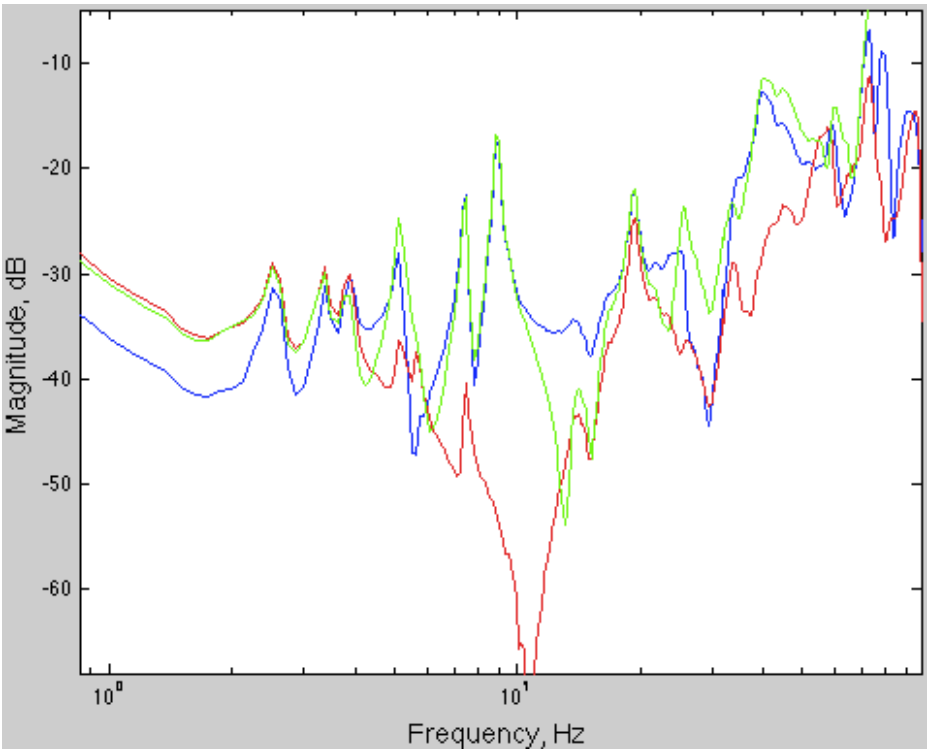

Figure 21: red – IMU; blue – passenger side laser; green – driver side laser. This plot shows the frequency response function of the response, with excitation on the driver's side. 0 dB is scaled to the acceleration from the impact hammer

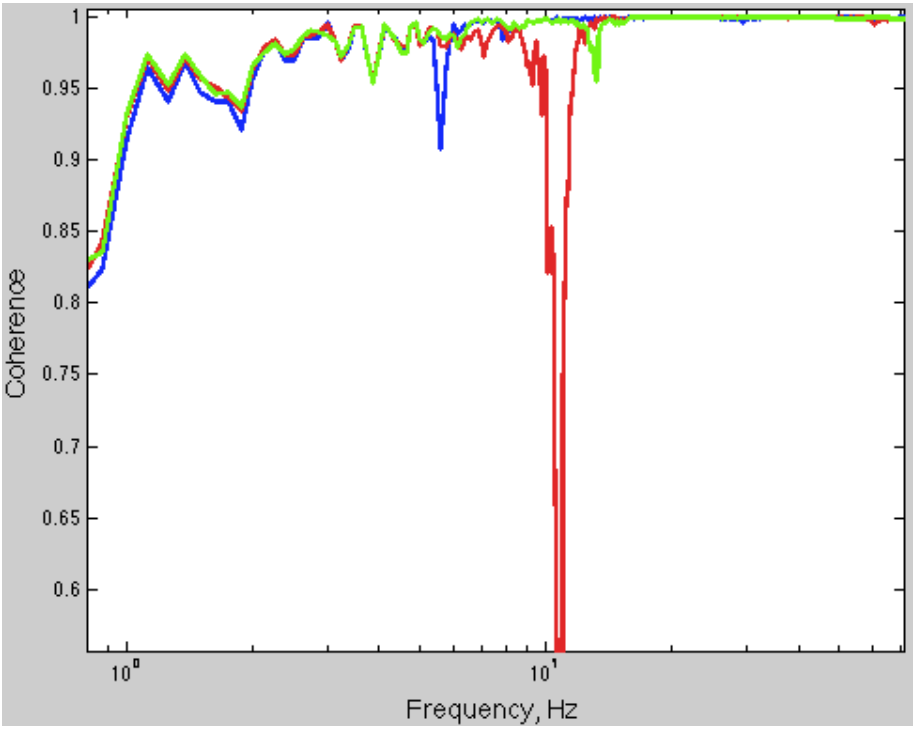

Figure 22: This shows the coherence of the test data over 5 trials; red – IMU; blue – passenger side laser; green – driver side laser

From the frequency response function, it can be seen that there is a very apparent peak of the whole system at 19.3 Hz. Figure 23 shows the phase of the sensors – all three are in phase, meaning that the three accelerometers are moving in the same orientation as each other. Because the phase of the three sensors are the same, even though there may be excitation close to this frequency, the IMU should not have any trouble accounting for the movement. Furthermore, the phase shift of all three sensors as they pass through this frequency suggests that this is a mode for the system.

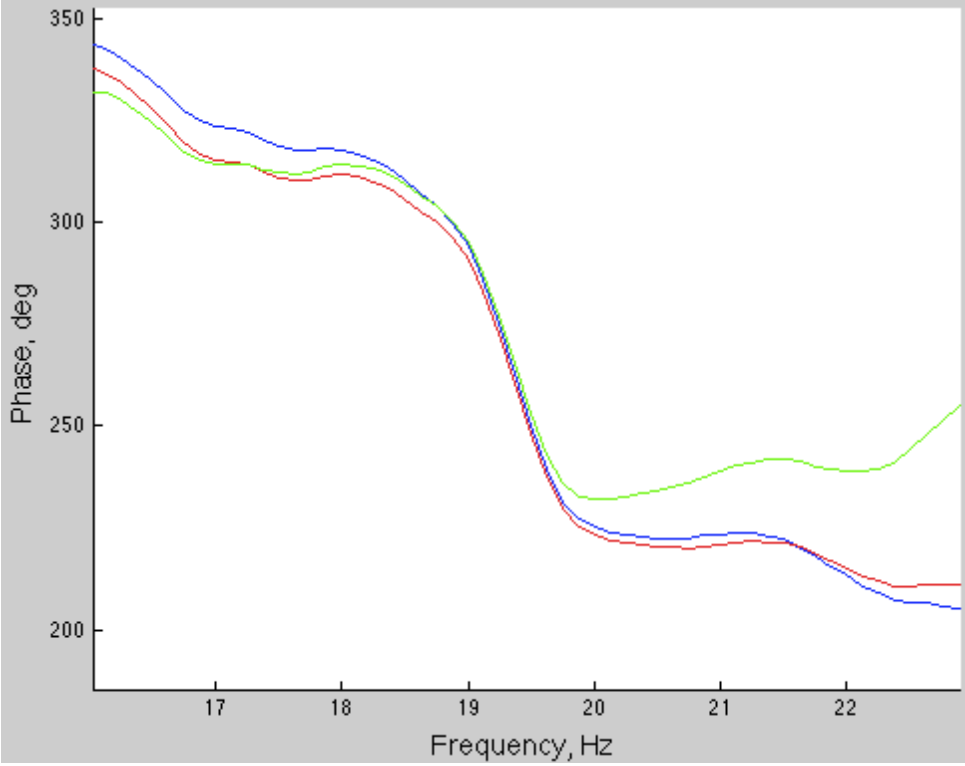

Figure 23: This shows the phase of the response centered around 20 Hz; red – IMU; blue – passenger side laser; green – driver side laser

There were also other areas of interest in the data before 20 Hz. At 2.5, 3.4, and 3.9 Hz, all three locations showed spikes in the frequency response, while at 5, 7.5, and 8.9 Hz, the two lasers showed excitations without the IMU. Figure 24 shows the phase of the response between 1 and 4.5 Hz. At the peaks 2.5 Hz, the IMU is mostly in phase with the two lasers, so its motion (translation) should be captured. At 3.4 and 3.9 Hz, there is a small phase difference between the two lasers, while the IMU is in the middle of the difference. This means that there is some roll and translation, which the IMU should also be able to accurately record.

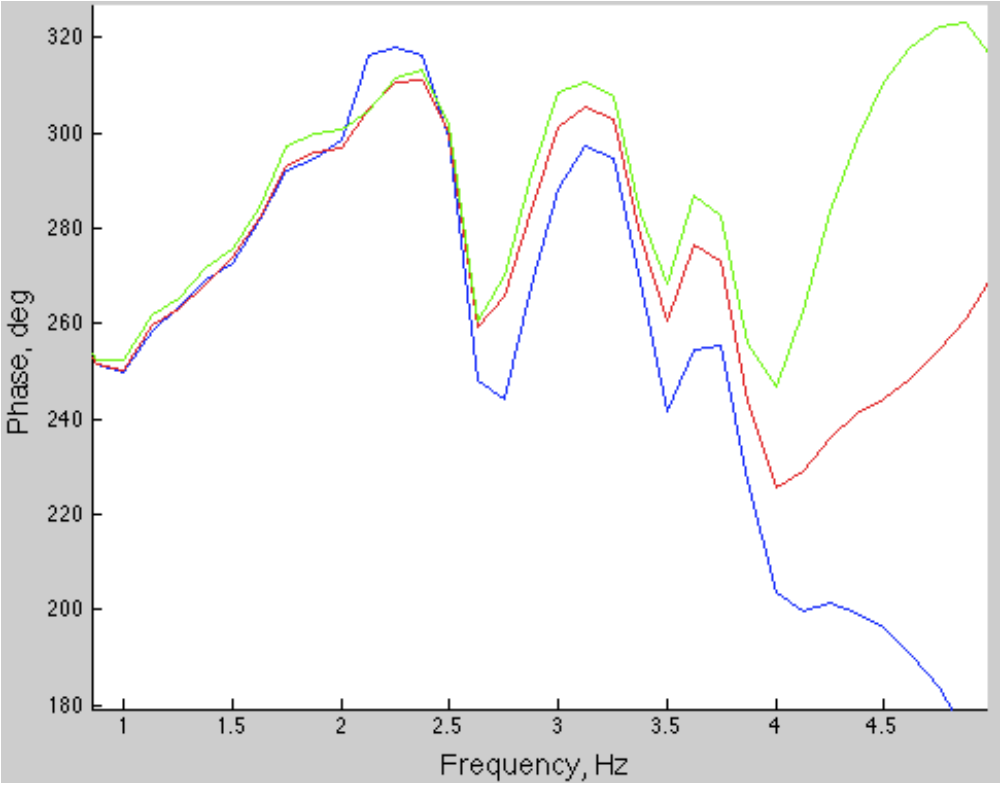

Figure 24: red – IMU; blue – passenger side laser; green – driver side laser; This plot shows the phase response from about 1 to 4.5 Hz.

Next, at 5, 7.5, and 8.9 Hz, only the lasers had peaks in the frequency response. Therefore, only the phase of the two lasers are shown in Figure 25. This plot shows that their phase is offset by 180 degrees, which suggests that the frame has a rolling motion. This should also be captured by the IMU.

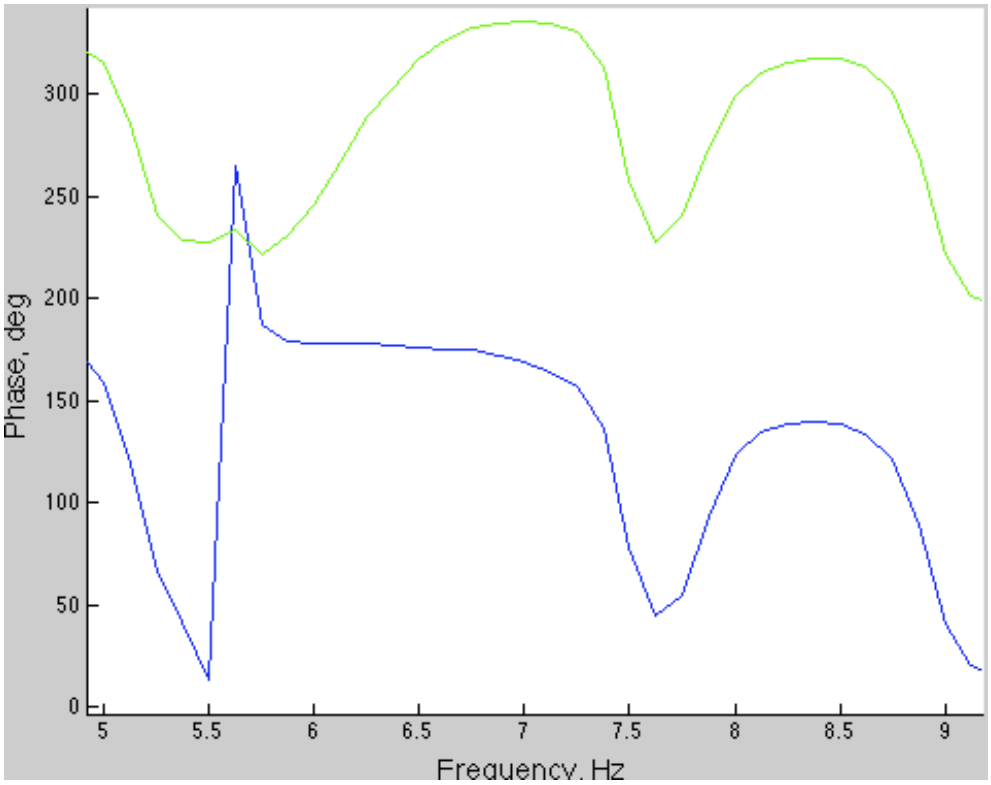

Figure 25: This shows the phase of the response of the two laser; green – driver's side laser; blue – passenger's side laser

3.3 Effects on Roughness

From the modal testing, it is clear that the INS should be able to keep track of all the expected motion from the equipment frame. The natural frequency of the frame is lower than the expected input excitation from the vehicle. Moreover, any excitation to either support at the natural frequency is in phase at the IMU and the two lasers, meaning that any motion should be correctly captured.

In the event that a frame was not correctly designed, where the IMU cannot accurately record the motion experienced by the other sensors, more errors will be present. To demonstrate, a sinusoidal movement with amplitude of 0.5 mm at 15 Hz (to simulate vibrations at the ride frequency) will be superimposed on the sample roads, simulating vehicle motion that is poorly accounted for. Figure 26 shows the profile of this sinusoidal error.

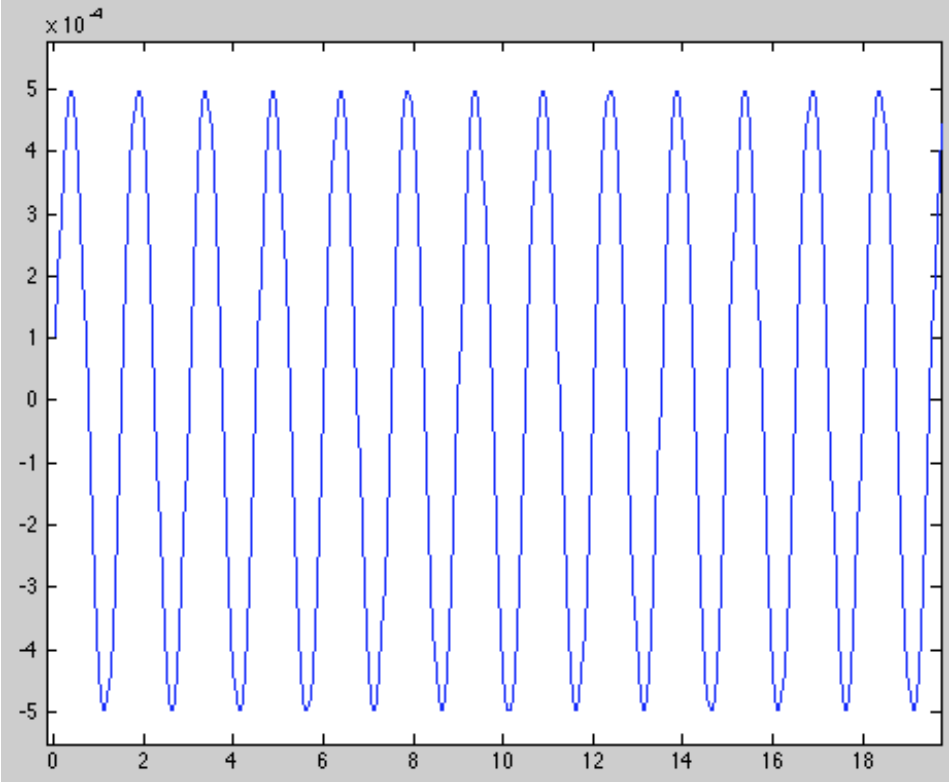

Figure 26: simulated sinusoidal error (zoomed in to 0-30 m). Height (m) vs distance along road (m)

The roughness values of the sum of the sample roads and error are shown below. It also calculates the confidence that the new IRI (containing the simulated error) is greater than the original IRI using a two sample, one tailed t-test. In addition, it also shows the cutoff amplitude required for a significant change in IRI (90% confidence of change).

	Original		Original + sinusoidal		Change	Confidence	Amplitude for
Road	IRI	st. dev	IRI	st. dev	in IRI	of change	90% confidence
smooth	1.88	0.136	2.11	0.119	0.23	99.2%	0.34 mm
medium	3.19	0.143	3.35	0.121	0.16	96.8%	0.42 mm
rough	9.69	0.446	9.73	0.412	0.04	56.6%	1.40 mm

The IRI values for all three roads increased, though not significantly for the brick road. Thus, these results show that this error has bigger effects on the smoother roads, while barely changing the roughness of the more irregular brick road. This is likely due to the to the existing road roughness overshadowing the simulated error conditions. The higher amplitudes needed to cause a significant change in the rough road reiterates this belief. Thus, especially in roads that need to be smooth (such as highways), this type of error would create significant undesired increases in the measured roughness. It is also

important to note that this error is not easily seen in the data. Because the error is both smooth and small in amplitude, it may not be very visibly apparent when inspected. Thus, it is important for the hardware design to be robust, as in the case of the Explorer's mounting frame, to prevent this type of error from occurring.

Finally, excitations at different frequencies were tested. The following table shows the results of various frequencies on the IRI given a 0.5 mm amplitude.

	10 Hz	12Hz	14Hz	16Hz	18Hz	20Hz
smooth	2.06	2.10	2.11	2.11	2.08	2.03
med	3.24	3.28	3.33	3.36	3.36	3.34
rough	9.71	9.72	9.73	9.73	9.73	9.72

The IRI actually increases as the excitation frequency approaches 14-16 Hz, while decreasing as the frequency goes further away. These results confirm the frequency response of the quarter car model, which has a peak at 15 Hz. This means that any frame excitation further from 15 Hz would likely have a smaller impact on the IRI.

4. INS Errors

Another large source of error is in the inertial navigation system. This uses both an inertial measurement unit (IMU) and the global navigation satellite system (GNSS, but this work will focus on the US-based global positioning system (GPS)) to keep track of the vehicle's attitude (orientation) and location. Because these are fairly complex systems, they are likely to have error introduced during their usage. This chapter will focus on the diagnosis and resolution of INS based error in the VTPL's HMMWV profilometer.

4.1 Introduction

In order to solve troubleshoot INS errors, the following methods were developed as a way to ensure quality data. The first three steps address the potential causes of error, while step 4 is a backup step that removes any remaining error.

1. Ensure that all relevant hardware is functioning properly.
2. Maintain correct input voltage
3. Manually confirm INS solution quality
4. Post-process data to remove remaining error

This process was used to troubleshoot the HMMWV profilometer that the VTPL designed and manufactured for the US Army TARDEC. This Profilometer consists of a road profiling laser and a vehicle positioning system mounted on a HMMWV. The Phoenix Scientific Inc. (PSI) road profiling laser uses a 48 volt power source (consisting of four 12 volt lead-acid automotive style batteries in series), while the Novatel global positioning system/inertial measurement unit (GPS/IMU) uses a 12 volt power source (consisting of two additional 12 volt batteries in parallel). In early 2013, a problem was reported with the profilometer, stating that after collecting data for extended periods (such as over a workday), the output data was no longer smooth. Road scans would begin to display discrete steps not found in the actual road. Due to the nature of the problem, there seemed to be an issue with the IMU. Because the IMU tracks the attitude of the HMMWV, any errors in the IMU measurements would affect the perceived location of the vehicle, which would ultimately cause the discrepancies between the actual road and measured data.

4.2 Drift

4.2.1 IMU Troubleshooting

Initially, the HMMWV IMU was replaced by another of the same model (from the Ford Explorer profilometer) to check if the problem persisted or if the IMU was faulty. Because the same errors occurred again, the possibility of a failing IMU was ruled out.

If the IMU is not provided with adequate voltage and current, it will produce erroneous results. This issue was demonstrated on VTPL data acquisition trips in the past and was verified by TARDEC. Specifically, a problem occurs when the supply voltage decreases below the minimum threshold of the positioning system, which is listed at 9 volts. Previous experience also indicates that there may be some degradation in the quality of the IMU results when the voltage drops below 10 volts. No issues had been detected when the voltage is maintained above 10 VDC.

The supply voltage to the IMU/GPS system is nominally 12 VDC, however, there are several factors that may lead to a drop in voltage from the nominal battery voltage to that actually supplied to the IMU/GPS. First, there is a voltage drop when the batteries is under load – up to a 2 volt difference was observed between battery voltages when isolated versus when the sensors are connected. Next, the wiring and connectors between the batteries and sensors will account for some additional voltage drop. The six batteries, located in the bed at the rear of the vehicle, feed into a power distribution board (located in the rear of the vehicle cabin). This power distribution board connects four batteries in series (to supply a nominal 48VDC to the laser) and two batteries in parallel (to supply a nominal 12VDC to the GPS/IMU and accelerometers). Next, the power is connected to the sensors, which are located on a frame in the front of the vehicle, requiring another set of connectors and several meters of wiring. These factors, combined with the decreased battery voltage after constant usage, were thought to have brought the voltage below the 10 VDC threshold.

4.2.1 Proposed Solution

To mitigate this issue, a DC/DC converter was installed into the power supply before the GPS and IMU. Located physically near the sensors (to avoid unnecessary voltage drops due to connectors and wiring), each converter produces a consistent input voltage of 15 volts (at up to 30 watts) to the GPS/IMU and amplifier. The GPS receiver

accepts a voltage from 7-15 volts (typical 3.7 watts), while the IMU specifies a 15 volt input (typical 8 watts). Furthermore, the amplifier for the accelerometers accepts 9-18 volts (typical 10 watts). Separate DC/DC converters were used to account for potentially higher current draw during startup. This now places the input voltage above the minimum threshold (found through experience to be 10VDC), allowing the equipment to receive a consistent 15 volt input even as the supplied battery voltage drops to 9 Volts. While the minimum input to DC/DC converters is 9 volts, a more consistent supply voltage should allow for better accuracy. A wiring diagram is shown in Figure 27 below.

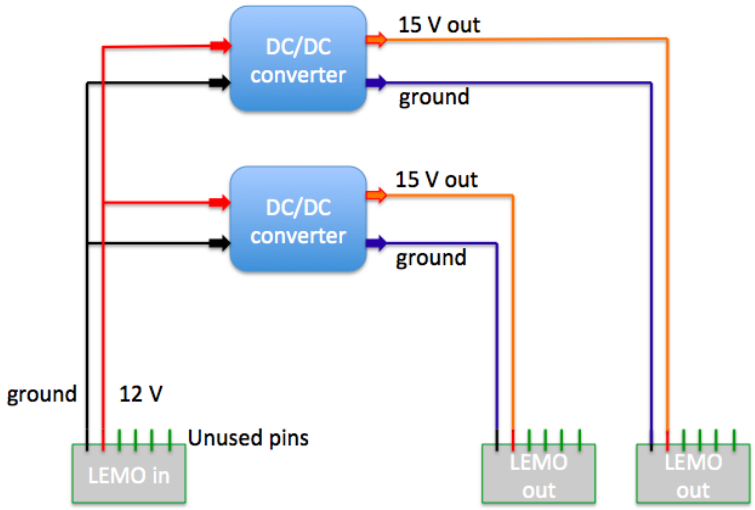

Figure 27: This wiring diagram shows the batteries' nominal 12V converted to 15V for the sensors. The LEMO connectors being used have 6 pins.

4.2.3 Verification

With the new power system in place, the Profilometer was tested. The HMMWV took data along the marked roads in Figure 28 in Blacksburg, VA.

Figure 28: Data was collected northbound along the red line and southbound along the blue line.

A total of 5 laps were driven, and data was also taken along the return trip to the Virginia Tech campus, as seen with the blue line in Figure 29, below. The red box in Figure 29 is roughly equivalent to what is shown above in Figure 28.

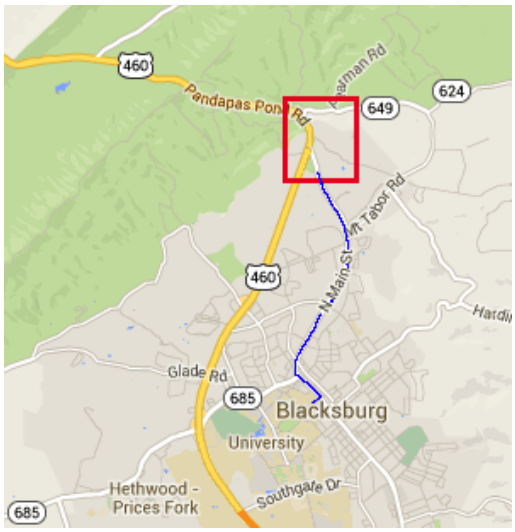

Figure 29: This shows where Figure 28 is located with respect to the university and the route back.

To ensure that the reported problem could be duplicated, the batteries were not charged before taking this data. Therefore, as the batteries started partially discharged, their voltages would more likely drop low enough without needing to take data for an entire day. The collected data began without any issue, although there was some noise in the middle of the road (possibly from some residual salt from a recent snowfall). Also, the scanner picked up the lane markers (the lines on the sides of the scan) as low points consistently throughout the data collection session, and can be seen below in Figure 30 and in most of the example figures.

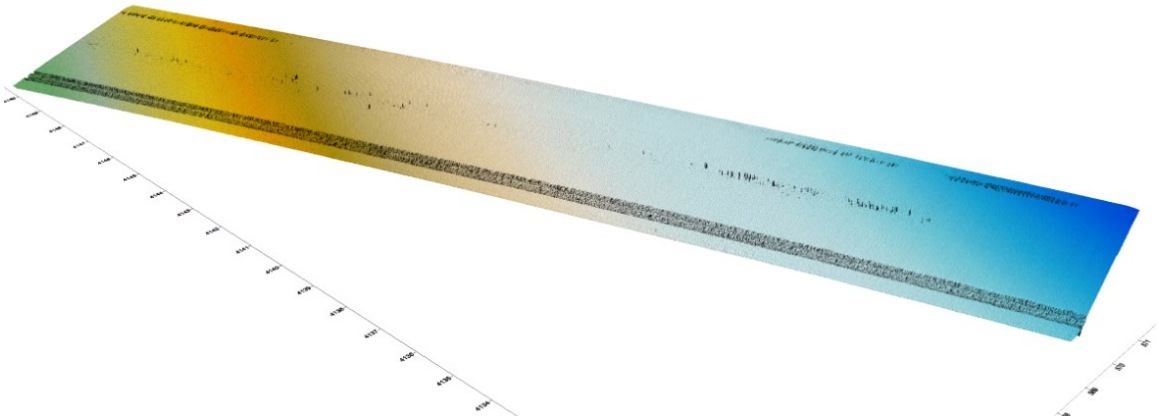

Figure 30: Some noise can be seen in the middle of the road, as well as lane markers on the edges. Units in meters

The noise in the middle of the road also increased at different points of the scan. There was some correlation of this noise between the scans, although they did not show the exact same characteristics. This noise was not identical most likely due to the small size of the salt residue remaining on the road. Between the scans, minor changes in the position and angle of the laser could affect how the data was interpreted.

4.2.4 Error Recurrence

Approximately halfway through the data collection session, errors similar to what were originally reported began to appear in the data. These errors were only in the data intermittently – a few frames would have the noise, but the rest of the data would be fine. The presence of the noise was particularly concerning because the Profilometer included the DC/DC converters meant to fix this issue. Thus, if the noise kept appearing, either the problem was not effectively fixed, or something else caused different errors. At first, the

errors were minor, appearing as small drops and rises at regularly spaced intervals, as seen in Figure 31 below. They appear at regular intervals corresponding to the GPS update rate: 2 Hz. It was hypothesized that this error is due to the interplay of the GPS updates and the IMU data.

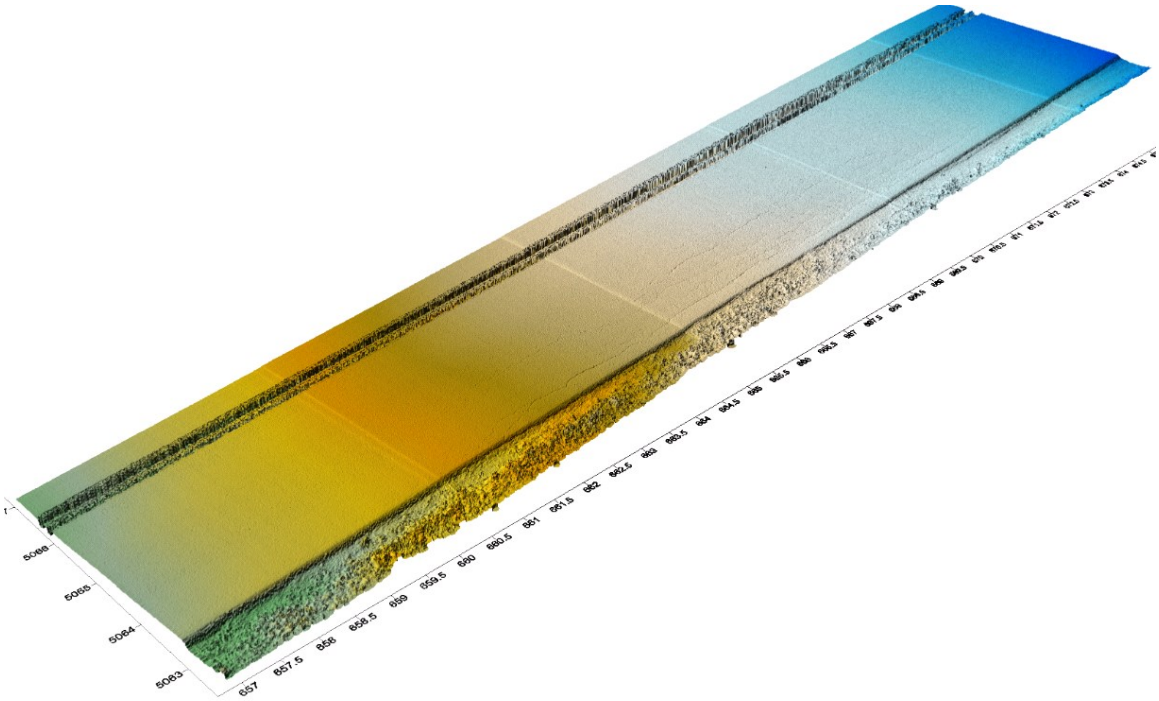

Figure 31: Small horizontal lines showing bumps began appearing in the road – this is similar to the originally reported errors. Units in meters

Later, more apparent steps were seen in the scans. As seen in Figure 32, the scanned data now shows distinct steps that are not present in the actual road. This was the very beginning of the last set of collected data, on the route back to campus. However, these large steps did not occur later in the same run.

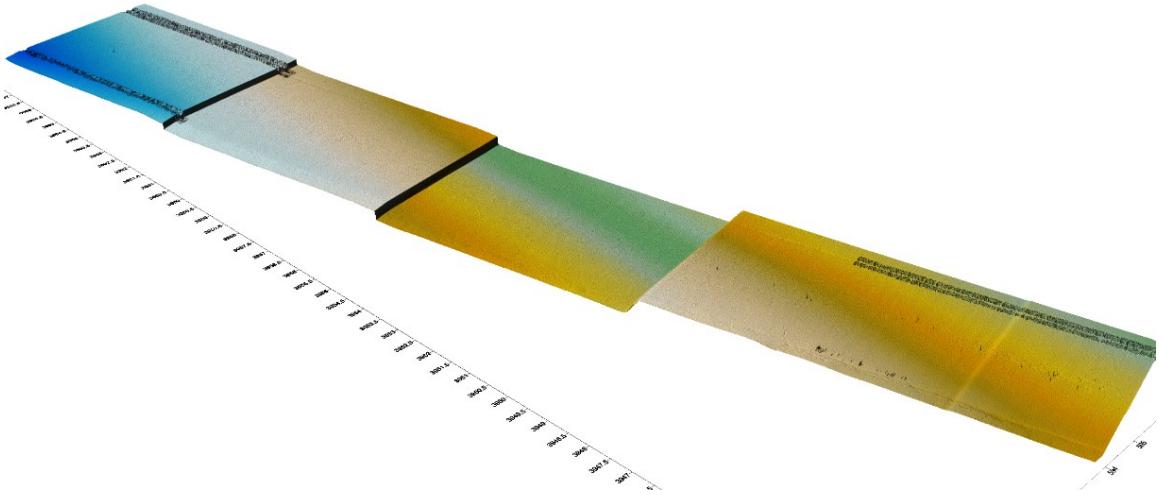

Figure 32: A section of the road scan showing considerable steps in the data at discrete intervals. Units in meters

Again, the appearance of these large steps occurred at regularly spaced intervals, but the appearance of steps was not continuous throughout the data set, but seemingly randomly distributed regions throughout the data. The previous smaller errors also appeared, but usually around these larger step errors. It should be noted again that the majority of the scanned data, even in these later scans, were smooth.

To eliminate the possibility of the laser or the accelerometers being the cause of this error, the GPS/IMU data of the noisy scans were compared to the data of the smooth scans. Figure 33 below shows typical plots of elevation versus sample number of the GPS/IMU data for profiles that were noisy (on the left) and for a typical profile that was correct (on the right). The errors in the GPS/IMU correlate to the errors in the profiles. The regular rate of the discrete steps in the data (corresponding to the 2Hz GPS refresh rate) clearly show that the cause of the steps is the GPS/IMU combination.

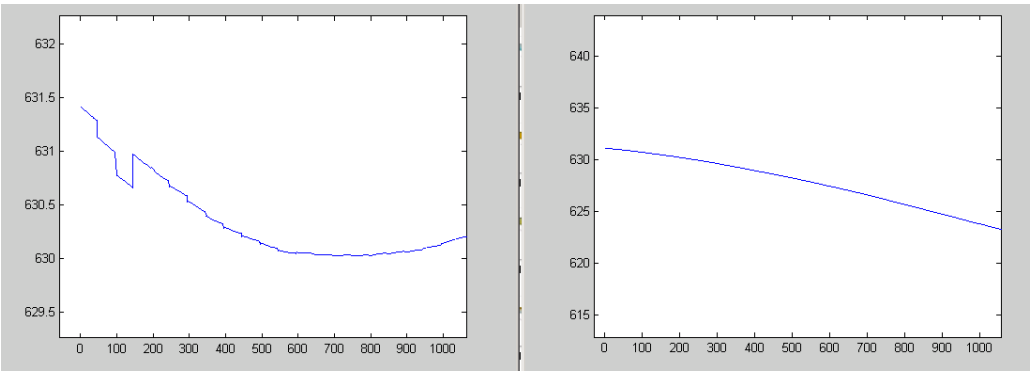

Figure 33: Elevation (m) vs sample number. The plot on the left shows the same data as Figure 32, while the plot on the right is from the first set (Figure 30). The steps in Figure 32 are apparent on the left plot.

4.3 INS Alignment

Since the issue resides with the GPS/IMU combination, the TARDEC IMU performs the same as the VTPL IMU, and the supply voltage to the IMU is now a consistent 15VDC (assuming that the batteries are producing at least 9VDC – the lower limit of the DC/DC converter), all potential causes were investigated. First, the provided voltages were checked. This data was taken with the new DC/DC converters in place, which was expected to solve the problem. At the beginning of the data collection session, the 12 volt supply measured 11.8 volts under no load. With the DC/DC converters and sensors connected, it read 10.1 volts. At the end of the session, the batteries were measured at 11.5 volts; meanwhile, under load of the DC/DC converters and sensors, they read 9.7 volts. This voltage is still above the minimum 9 volt requirement of the converter, and they still output 15 volts without issue. Thus, it is unlikely that the voltage was the cause of these problems.

4.3.1 GPS Reception

Next, the GPS and IMU output data were investigated. Figure 34 shows the recorded elevation, solely based upon the GPS data (no IMU), versus time. This data was collected along the route shown in Figure 28. The discontinuities in the plot correspond to missing GPS data (potentially from objects such as trees blocking satellites).

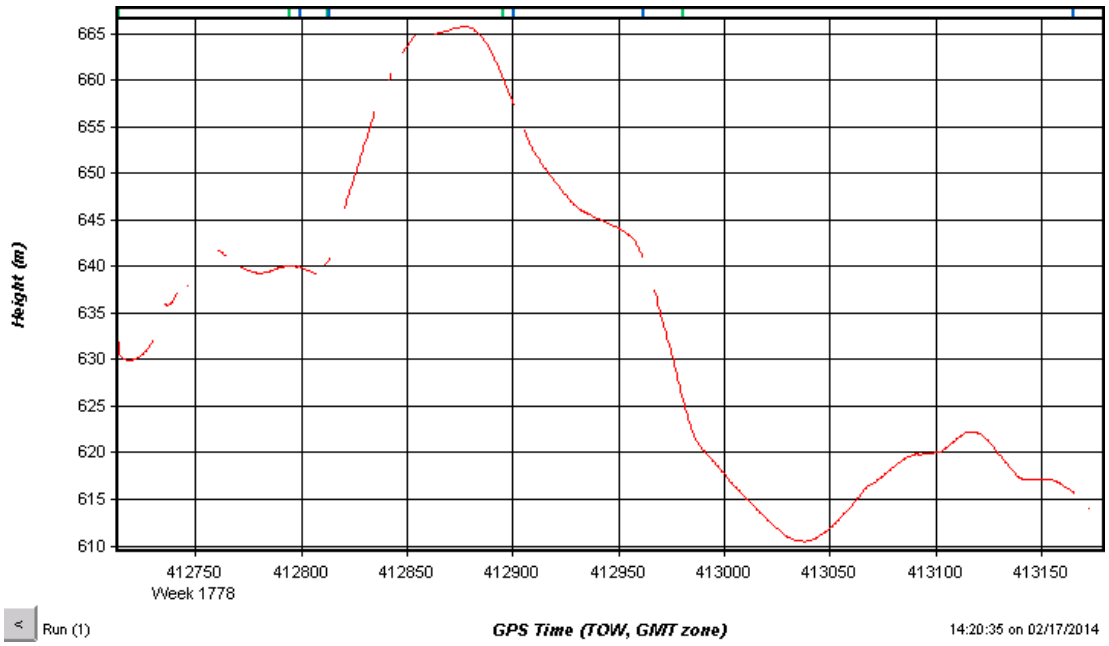

Figure 34: GPS only elevation profile of final run – this is a zoomed in section of a data collection run. elevation (m) vs. GPS time (s)

Although not shown on the plot, there is also some missing data immediately prior to starting this data collection run. Thus, it is hypothesized that poor GPS data may lead to noisy results due to drift in the IMU's estimate in location and orientation. Below, in Figure 35, Figure 34 is superimposed onto a plot of the number of shared satellites between the remote unit and base station. It is clear that where there are no shared satellites (most likely because the remote unit lost all reception), there is also no elevation data.

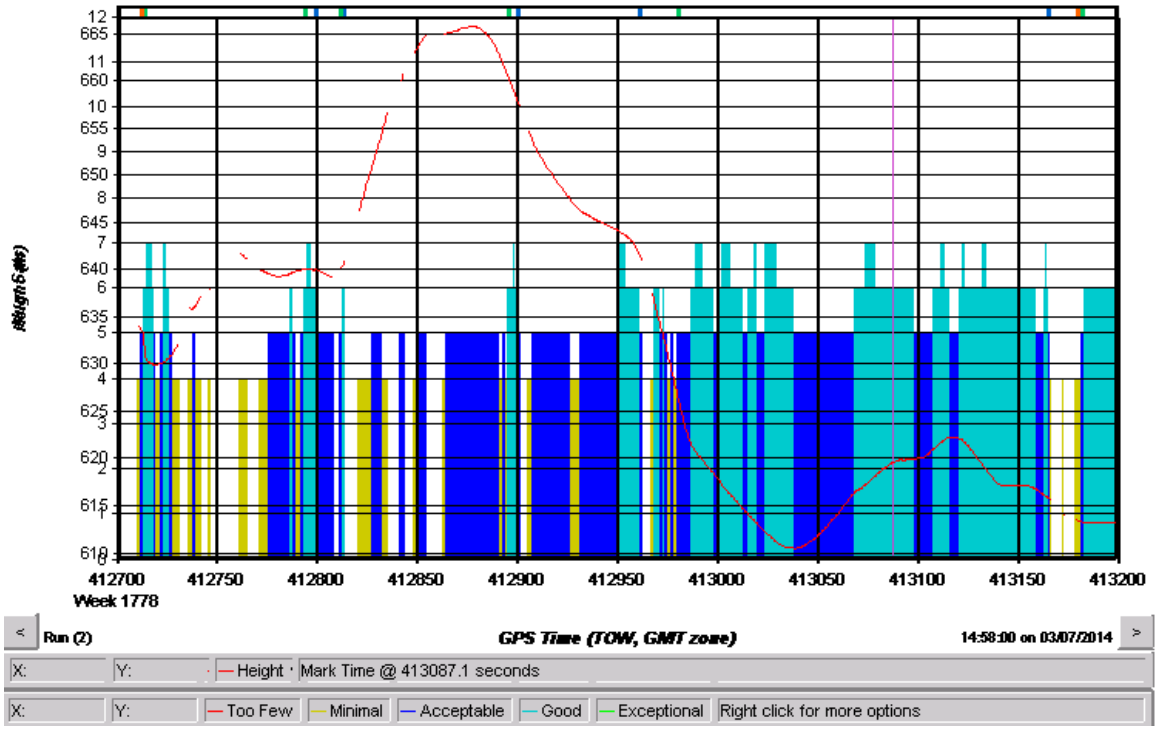

Figure 35: elevation data superimposed on number of shared satellites between remote unit and base station. Number of satellites connected and elevation (m) vs. GPS time (s)

Further investigation into the number of shared satellites shows that even the position solution with the IMU data is affected. Figure 36 shows the estimated standard deviation of the position (the horizontal positions – northing and easting – and vertical height) plotted on top of the shared satellites figure. This plot shows that where there is poor correlation between the GPS units, the deviation of the position data increases. This can be seen in the many small bumps in the position deviation that occur even when there correlation is poor for short timeframes (see approximate time = 412900). Therefore, it is hypothesized that due to the increased deviation of the position data, noise may be introduced when the GPS signal is reacquired and the position is fixed to its actual

position. Increased error in positioning during GPS outages was also seen by Yen, Ravani, and Lasky [30].

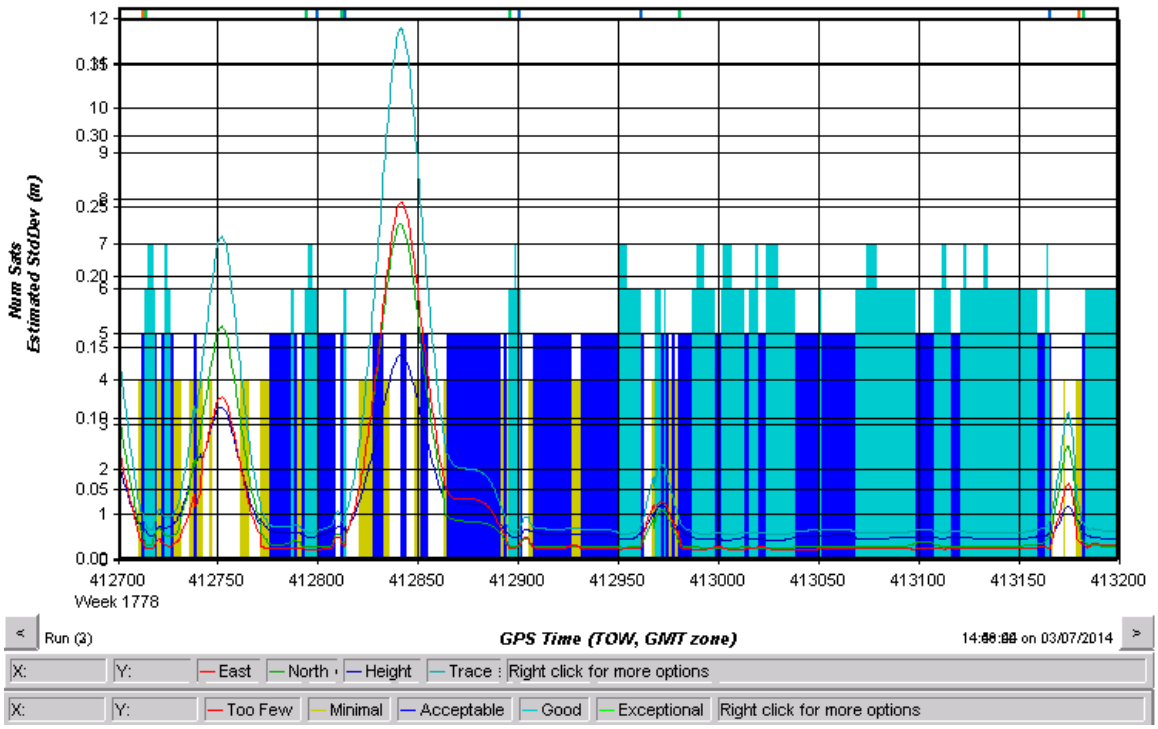

Figure 36: standard deviation of position plotted on shared GPS satellites between base and remote units. Number of satellites connected and estimated error (m) vs. GPS time (s)

4.3.2 Power Supply

Because no direct comparison data exists showing the difference of scan quality with and without the DC/DC converter, the Virginia Tech Drillfield was scanned with both the DC/DC converter enabled and with direct battery power (converter disabled). By measuring the same stretch of road with the different power supplies, it is possible to make a more accurate comparison between the scans. Unfortunately, while both scans were generally acceptable, they also both contained some noise. Figure 37 below shows a scan with the DC/DC converter and noise, while Figure 38 shows noise from a scan using direct battery power. To reiterate, these figures are not representative of the quality of the entire scans. Rather, they show some frames with noise – the majority of the scans have no such issues. Again, the noise was not consistent but only sporadically appeared in the data.

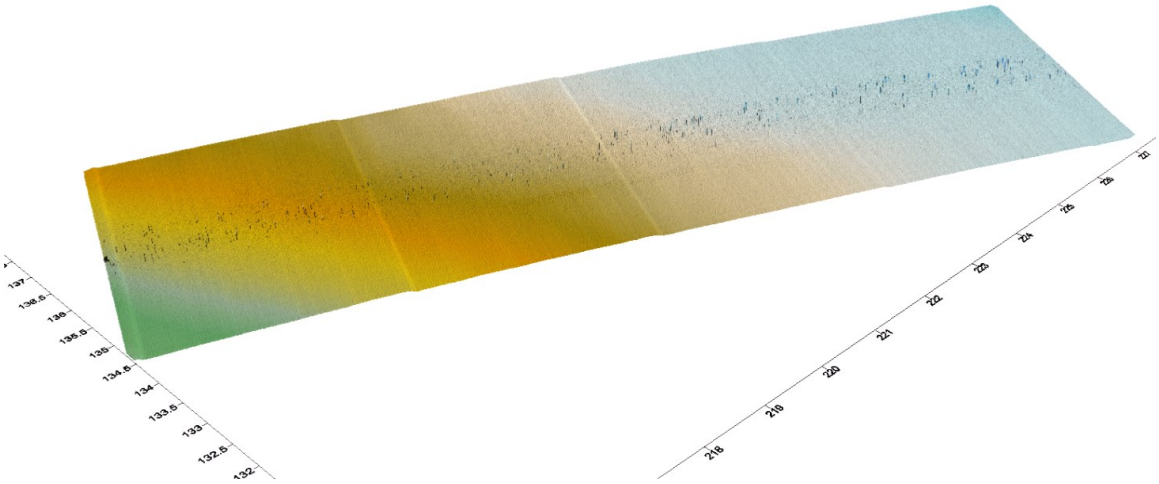

Figure 37: This shows some noise from a scan with the DC/DC converter. Units in meters

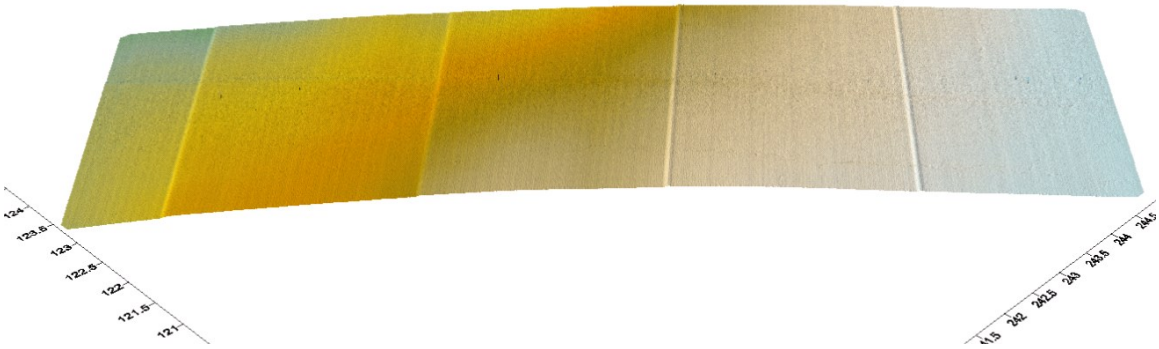

Figure 38: This is a scan without using the DC/DC converter. Noise in the data can also be seen in this plot. Units in meters

4.3.3 GPS Receiver

The inconsistent GPS reception and correlation with the base station was thought to be the major contributor to the noise found in the scans. To verify this, the remainder of the hardware in the GPS setup was checked. As the IMU was previously verified to function correctly (it was replaced by the army with another unit and gave similar data), the GPS receiver and antenna were the only local sources of error. First, the GPS receiver was replaced with another unit. The Drillfield was scanned again – if the receiver was the source of the error, then the new scans should have no noise. This test was also conducted with the DC/DC converter enabled and disabled, as the power supply may exacerbate the problem. Again, after the two scans with the new receiver were conducted, they showed no significant difference from data taken with the original receiver. A majority of the

frames of the scan had no issue, while a few had similar amounts of noise, as seen below in Figure 39.

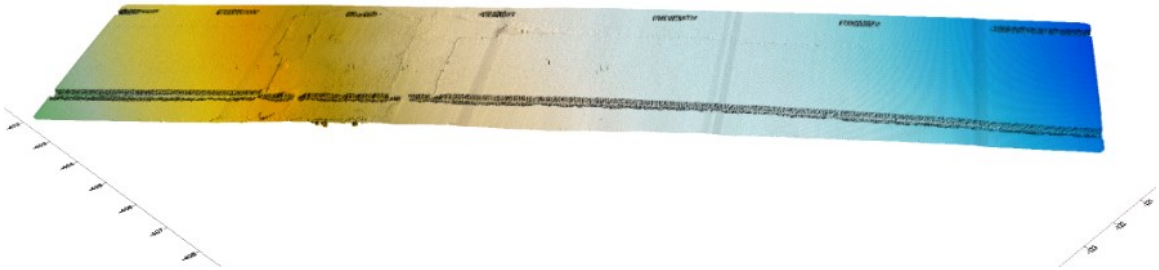

Figure 39: This is data taken with a different receiver. There are some cracks in the road, but there are also bumps due to a noisy GPS/IMU signal. Units in meters

Plots showing the position deviation and number of shared satellites for three configurations are presented in the following figures. Figure 40 shows the result of using the DC/DC converter and the original (TARDEC) receiver. Figure 41 shows the result of using direct battery power and the original (TARDEC) receiver. Figure 42 shows the result of using the DC/DC converter and a replacement (VTPL) receiver.

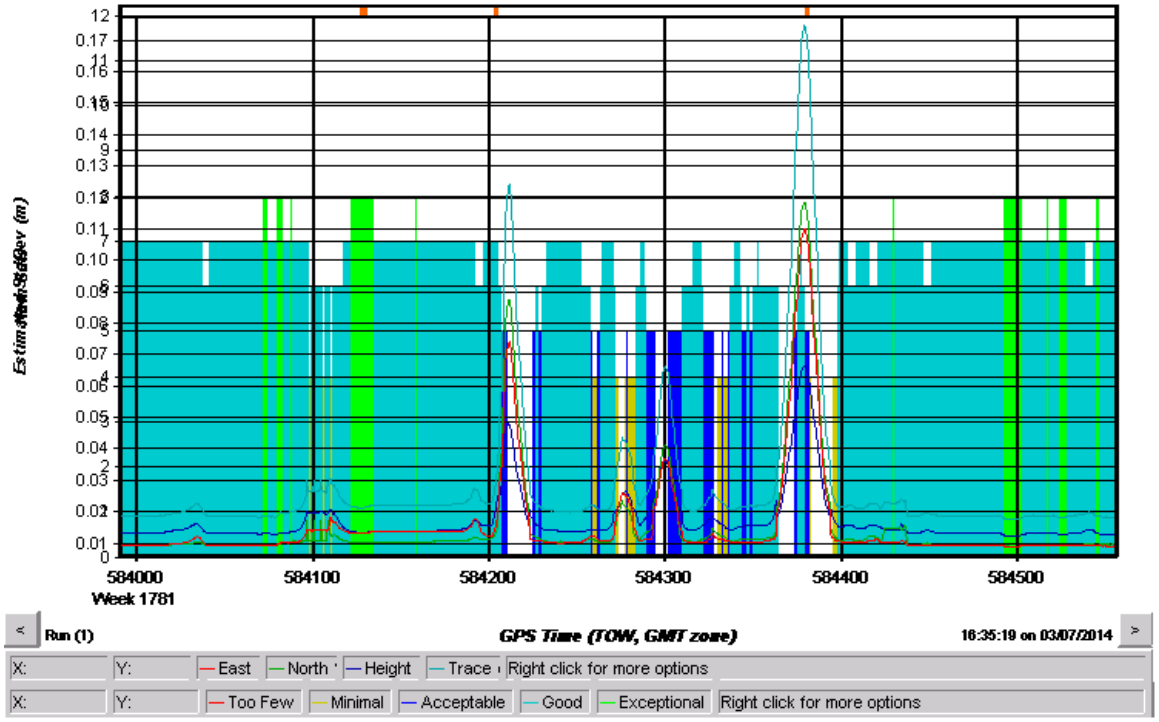

Figure 40: The position deviation overlaid on the number of shared satellites for the run with the DC/DC converter around the Drillfield (original GPS receiver). Number of satellites connected and estimated error (m) vs. GPS time (s)

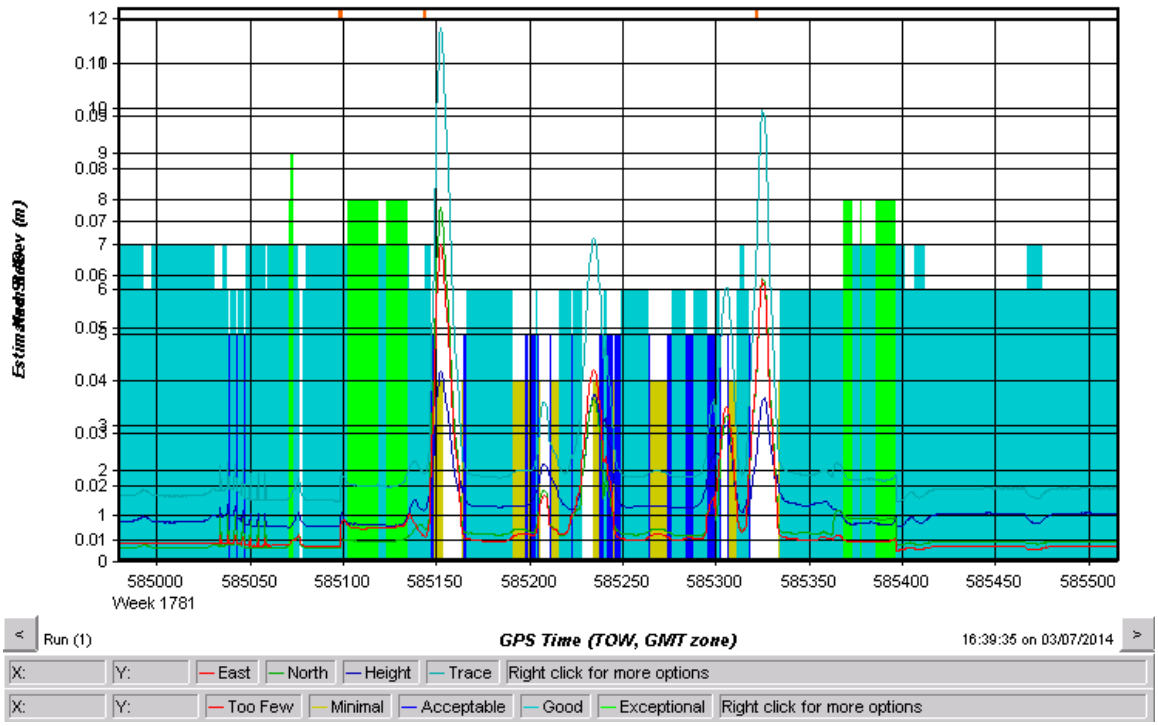

Figure 41: The position deviation overlaid on the number of satellites for the run using direct battery power around the Drillfield (original GPS receiver). Number of satellites connected and estimated error (m) vs. GPS time (s)

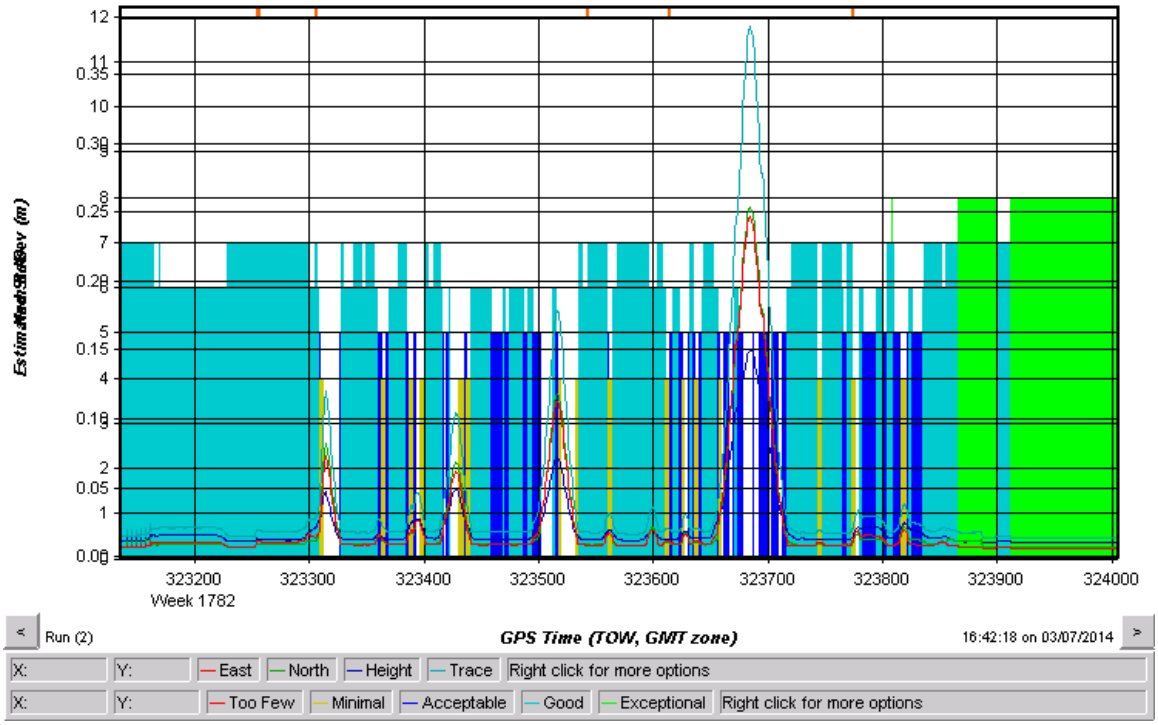

Figure 42: The error overlaid on the number of satellites, using the DC/DC converter with the VTPL GPS receiver. Number of satellites connected and estimated error (m) vs. GPS time (s)

4.3.4 Antenna

Finally, the antenna was also replaced with another unit to determine its status. Since the IMU and GPS receiver were ruled out, the focus of this test was to see differences in the quality of GPS data and satellite reception. As stated previously, the noise in the road scans correlated with areas of weak GPS signals. Figure 43 shows the signal before processing with the base station – note that there are areas of no data and also discontinuous points.

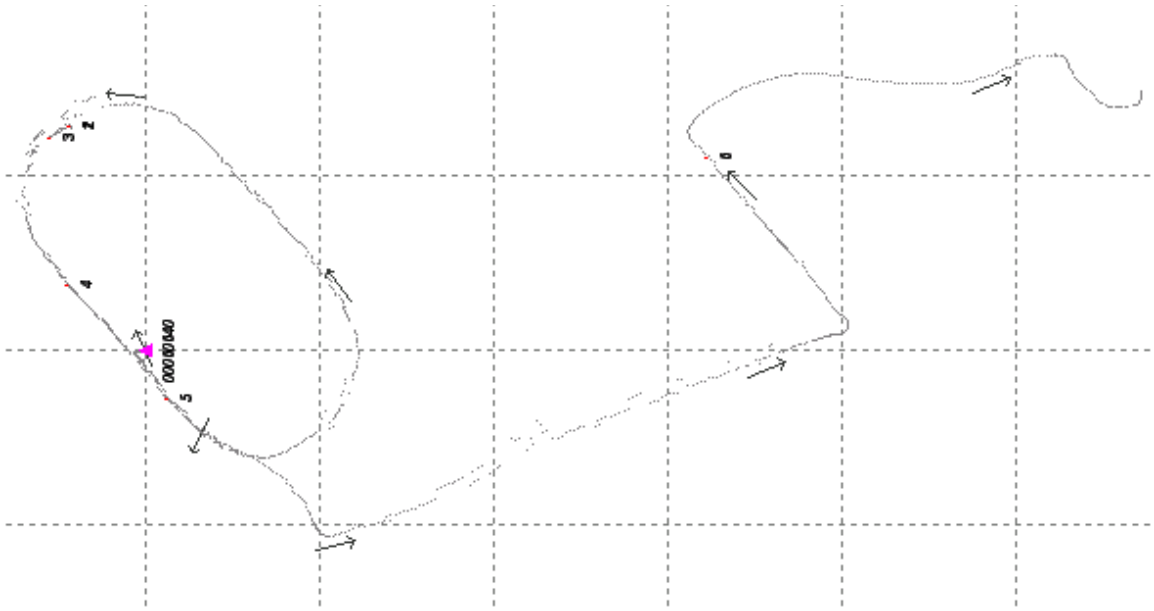

Figure 43: This shows the GPS data with the original antenna.

While a slightly different route was driven with the replaced antenna, it also shows discontinuities and holes in the data. This is shown below in Figure 44.

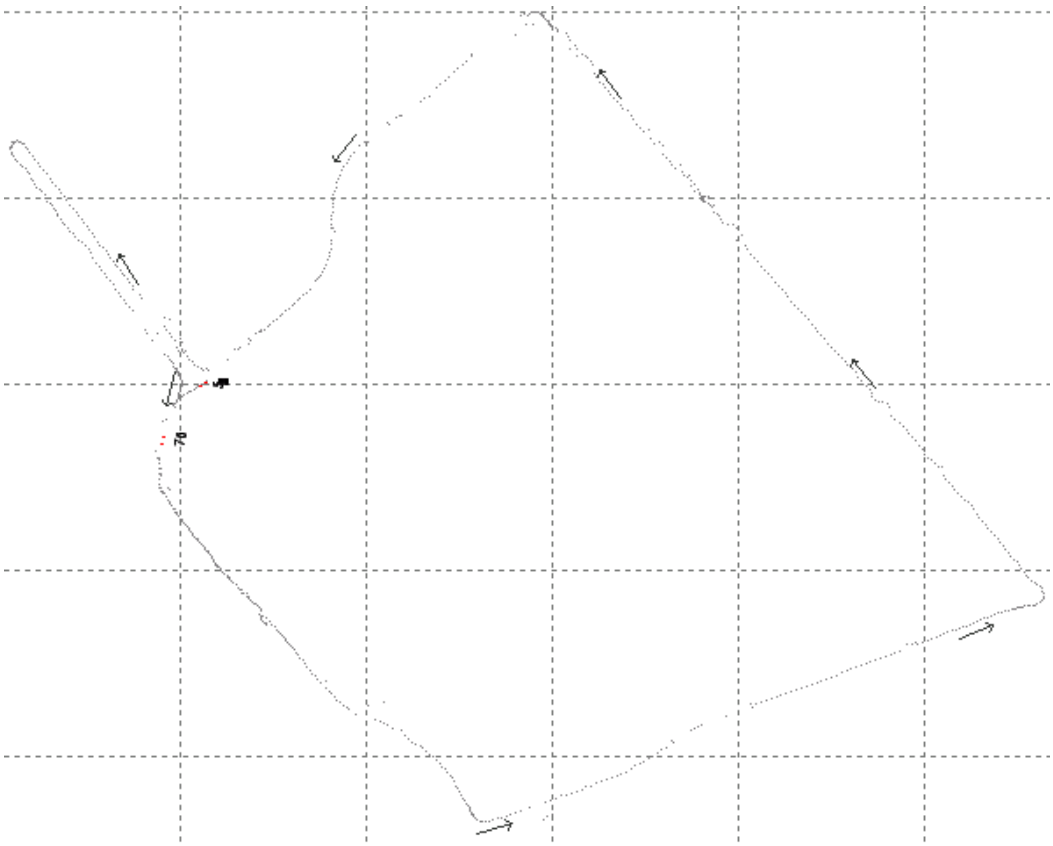

Figure 44: This shows the GPS data with the replaced antenna.

Furthermore, the actual connectivity to the satellites can be compared. Figure 45 shows the GPS satellite that the receiver (and original antenna) is connected to, while the red marks are bad data points.

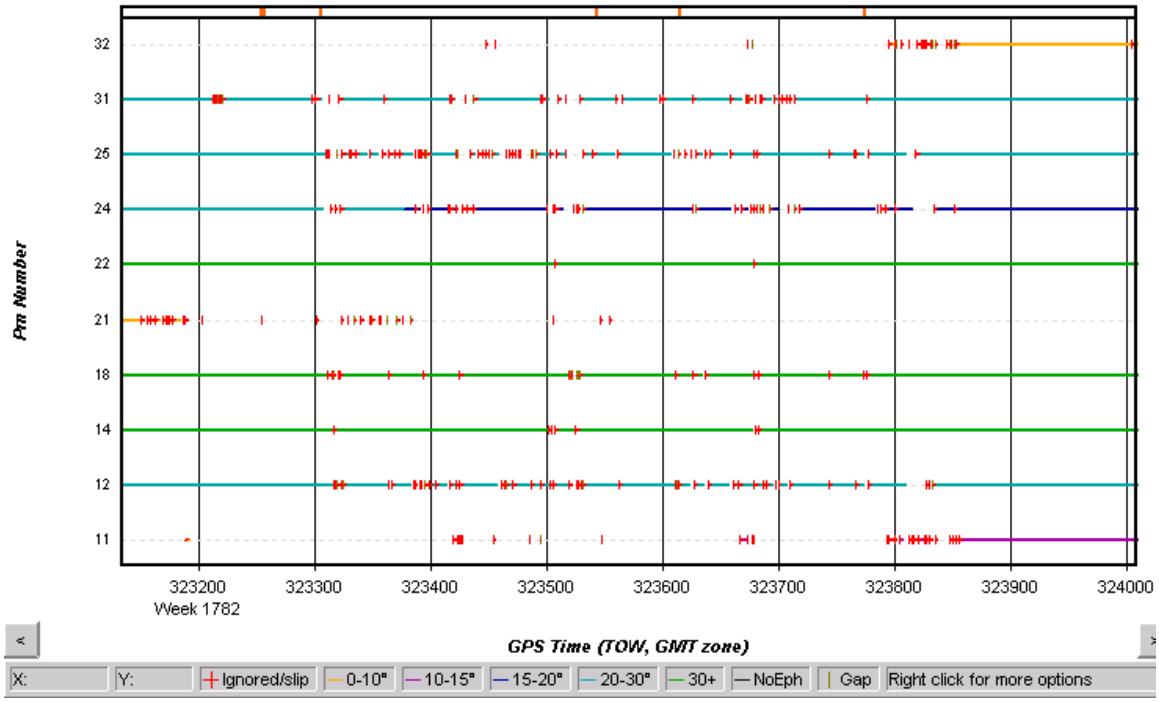

Figure 45: This shows the satellite connectivity with the original antenna.

This is compared to Figure 46, which shows the same plot but is taken with the replacement antenna. Note that they both exhibit points where there is poor reception.

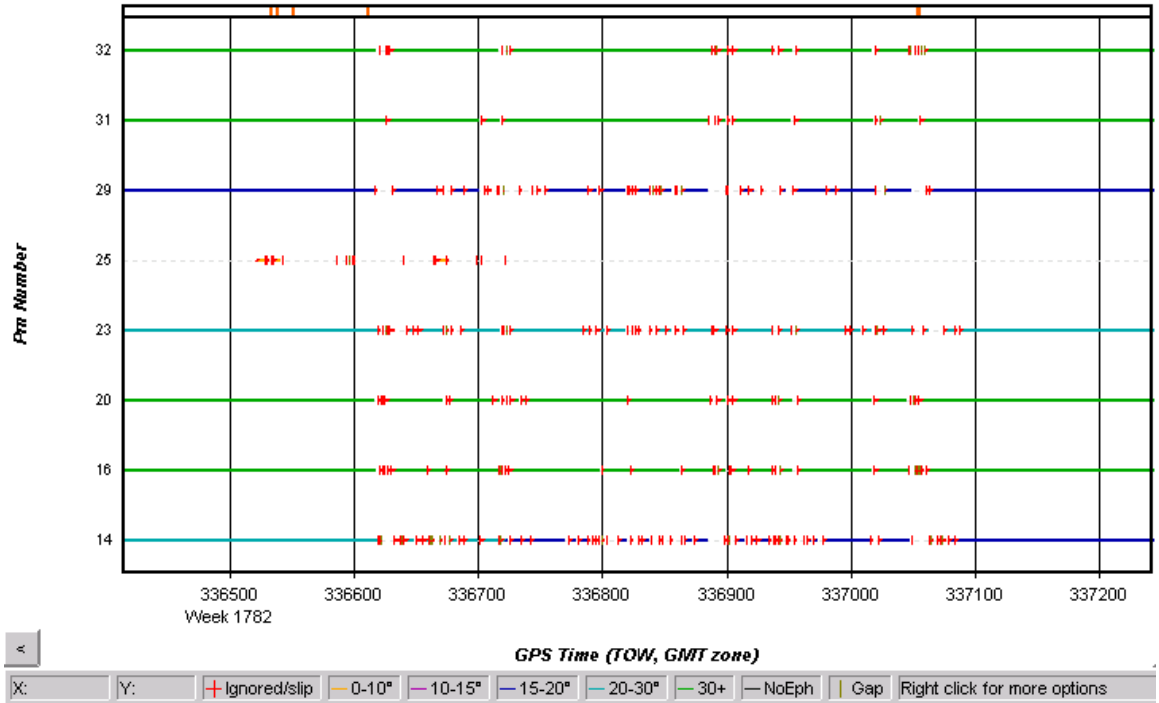

Figure 46: This shows the satellite connectivity with the replaced antenna.

4.3.5 Base Station

Additionally, to further ensure that this was a problem with the satellite reception, rather than error in the physical units, the HMMWV Profilometer was run concurrently with the Explorer Profilometer – this involved using two base stations and two remote units. Both base stations recorded similar data (especially the five satellites with the best coverage), as shown below in Figure 47.

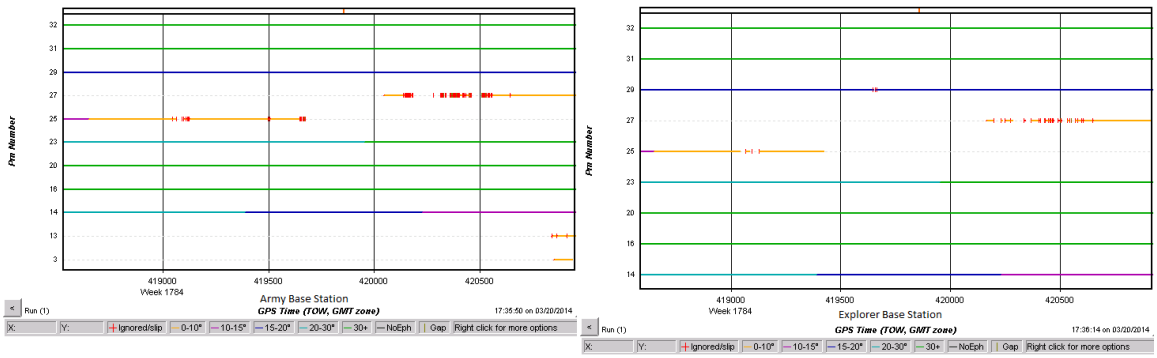

Figure 47: This shows the reception of the Army base station on the left and Explorer base station on the right. Note that there were not any signal gaps in the satellites with the best views (green lines) on either figure.

Moreover, the remote units similarly received similar data. Figure 48 below shows the two remote receivers' data overlaid on each other (with an offset to show the different source). Satellite connectivity was very similar, as were discontinuities in the data. Since both Profilometers recorded the same data, this is more likely attributable to a problem with reception, rather than a problem specific to the HMMWV system.

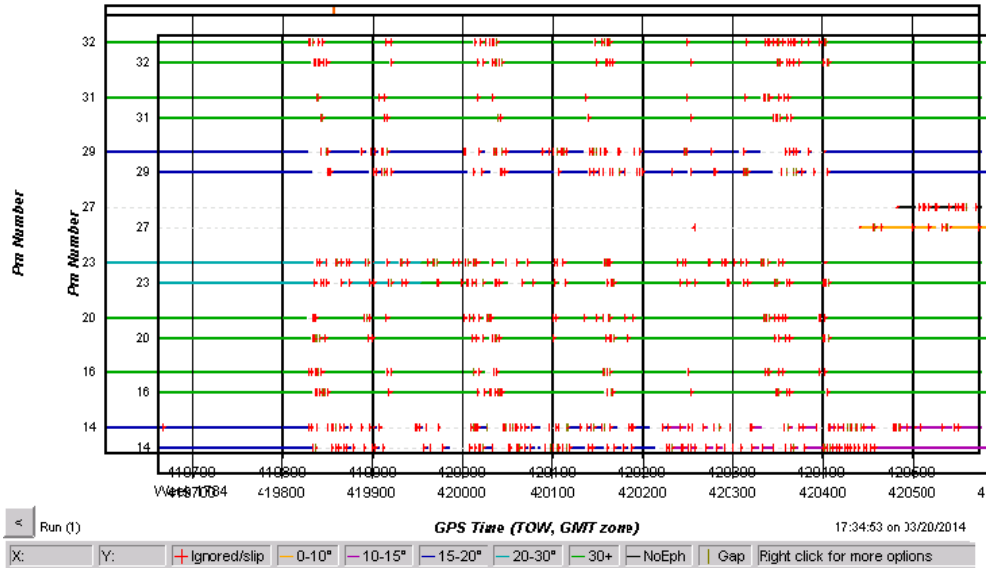

Figure 48: This shows the GPS satellite reception of the 2 remote units overlaid on top of one another.

The only consistent results demonstrated throughout all the testing is the correlation between the variation in profile position (the error in the final measurement of the terrain surface) and the lack of shared GPS satellites between the base station and the rover.

4.3.6 Proposed Solution

Because the INS consists of a GPS receiver and an IMU, these parts must be working together. The GPS must receive the proper signals from an outside source before it can begin to take data. Similarly, the IMU must also calibrate at the beginning of every power cycle in order to ensure that it is measuring its movement accurately. A more thorough initializing procedure should allow for better solutions, leading to less error.

The current instructions call for the GPS to be started, after which the IMU takes some time to perform a static alignment. The system then needs to be moved for a brief kinematic alignment (moving the vehicle, including turns), at which time, everything should be working comprehensively. However, after further investigation, it was found that this did not necessarily provide an adequate solution. Especially when the GPS signal dropped out, the IMU would often drift slightly off alignment. When the GPS signal gets reacquired, the measured travel path would have to be corrected to match the actual path, and this correction would often lead to some errors in the data. Godha and Cannon found similar GPS/IMU disagreements when taking data in urban areas [31]. A method was investigated to mitigate this error from occurring.

After the IMU performs its stationary alignment, a kinematic alignment is required so make sure that it is properly sensing motion. To do so, the driver must move the vehicle until the software recognized that the kinematic alignment section was completed. However, despite having this signal to proceed, the original alignment may have large confidence intervals. Obviously, it is more optimal if these were smaller. To ensure the best performance from the INS, the following procedures have should be used.

The GPS requires an absolute minimum of 4 satellites in view to calculate a position and time solution. To ensure a better solution, the satellites should optimally be apart from each other in the sky, therefore having large angles separating them from the receiver's point of view. A better solution can also be achieved by connecting to more GPS satellites. Because the setup of the GPS should be completed in an open area with good visibility of the sky, the GPS connectivity should not be a problem. Moreover, because

there is not much that the user can do to affect the GPS reception, ensuring that the program is finding a single, computed solution should be sufficient. This is shown in Figure 49, with the settings of interest circled in red.

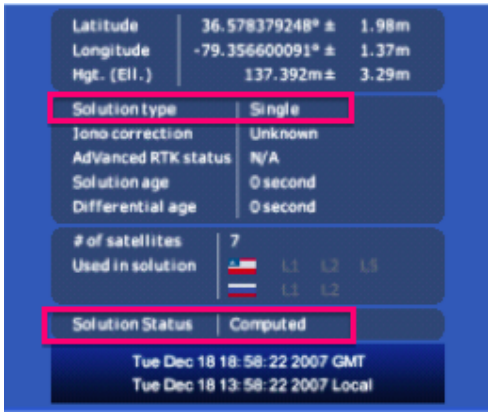

Figure 49: This shows the completed solution of the GPS.

The IMU setup procedure is slightly more complicated. After its static alignment period, while it has a solution, it still requires some movement to better determine alignment. Its pre-kinematic alignment error is shown in Figure 50, which is fairly high.

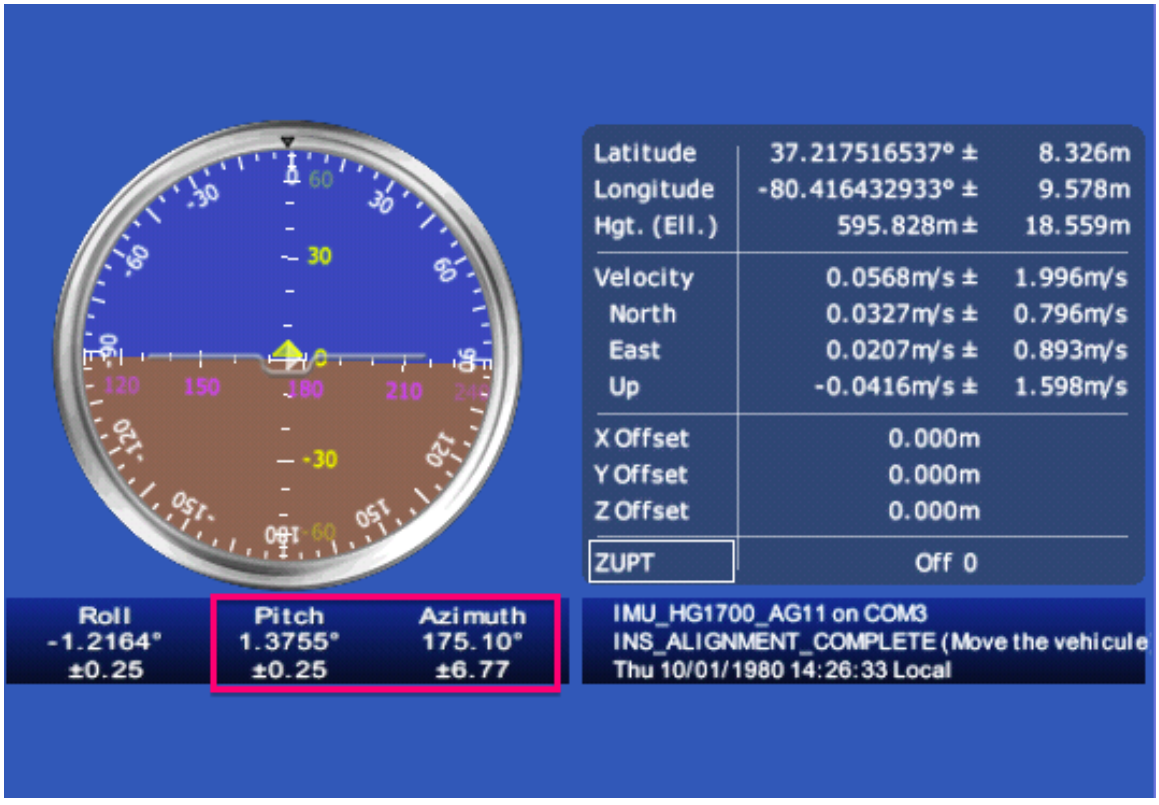

Figure 50: The IMU before any kinematic alignment. Note the high error values.

After the kinetic alignment, this error decreases but is still higher than preferred. It is believed that the higher error in the IMU causes a bad agreement between the GPS and IMU, which results in poor data quality. A sample set of these values are shown in Figure 51.

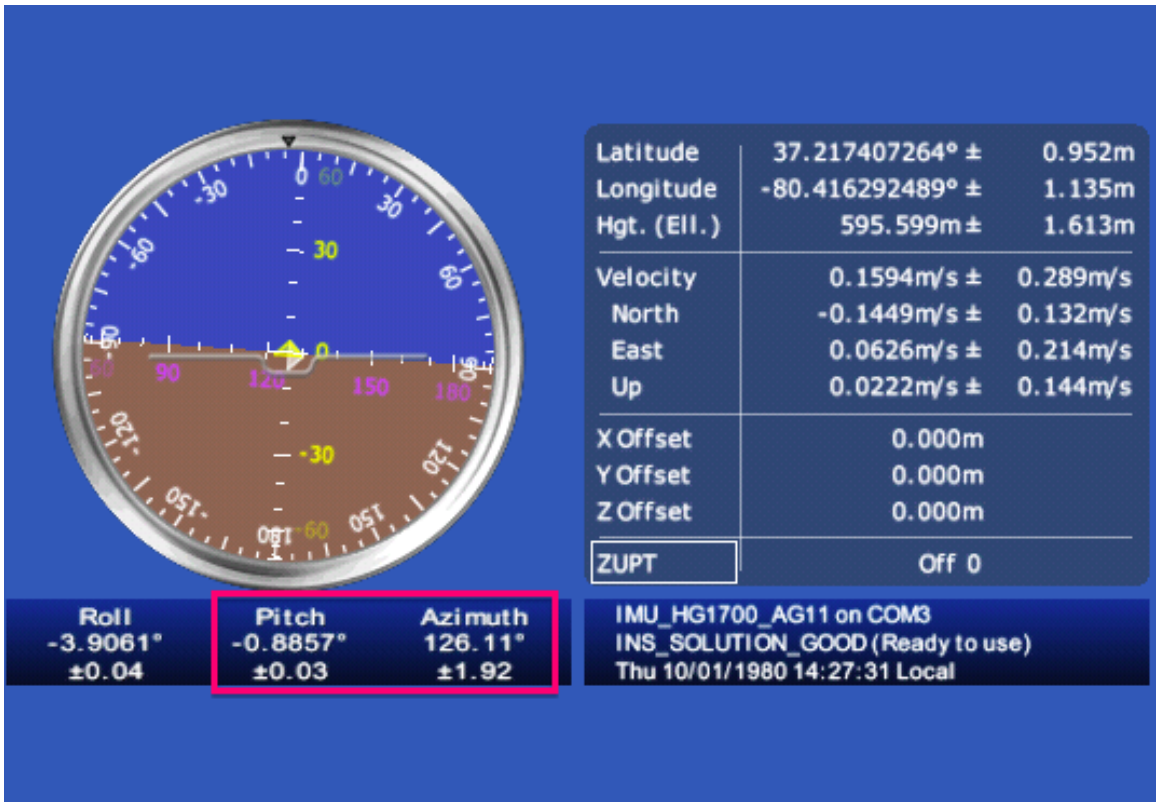

Figure 51: The IMU after the recommended kinematic alignment. The errors have decreased, but are still higher than optimal.

To overcome this issue, the vehicle should be given even more time to process a solution, even after the software claims that the IMU is ready. This should be performed in an open area with clear visibility of the sky (e.g., an open parking lot) to maintain GPS satellite reception, while introducing large changes in heading (figure 8s are recommended). The errors should decrease after approximately ten figure 8s to less than 0.2 for heading and 0.02 for pitch, as seen in Figure 52.

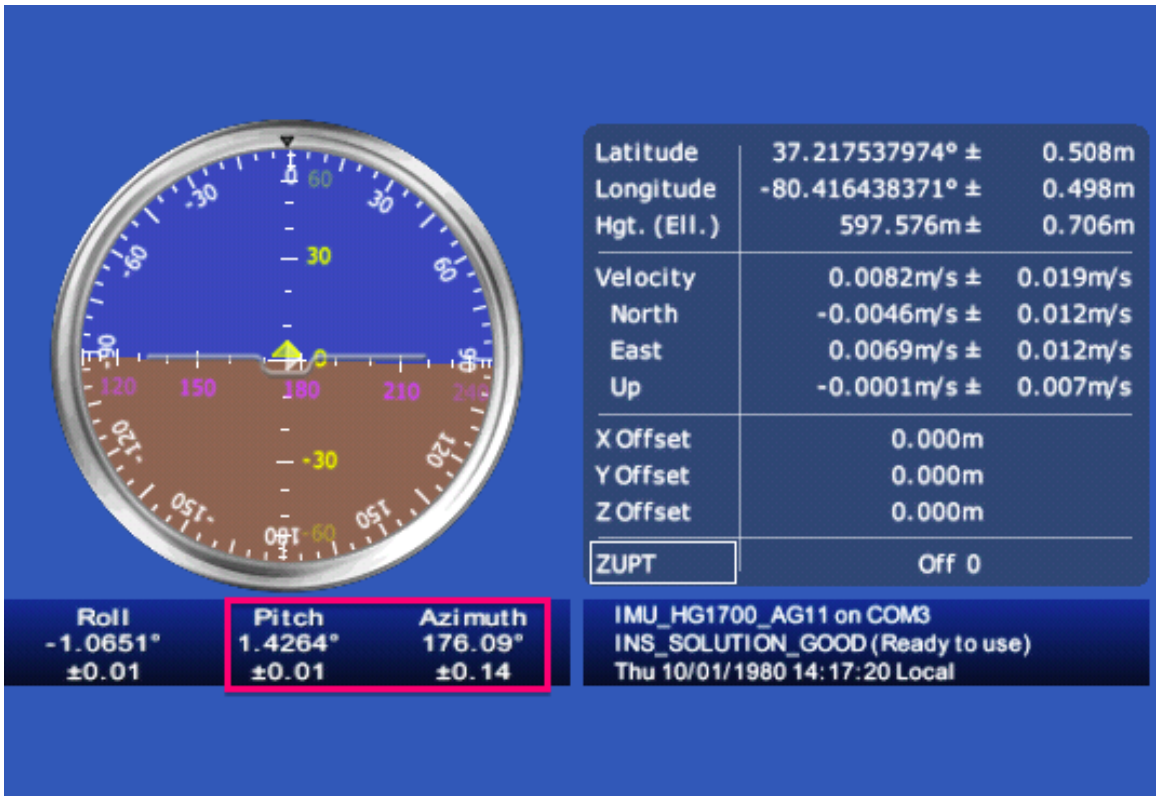

Figure 52: The IMU after the extra alignment (driving in figure 8s). Note the low error in pitch and azimuth.

With the added time for alignment, the INS should be able to more accurately record the path data of the vehicle. This means more consistent and accurate data while running through the remainder of the setup procedures. More importantly, when added to the normal profiling system setup procedures, it allows for better data to be taken in general, hopefully mitigating any negative effects that poor GPS reception may have.

4.3.7 Secondary Solution

Although it is preferable that the collected data has no unwanted noise or errors, it is likely that the previous steps will occasionally allow some errors through. To combat this, post processing of the data allows for the any remaining errors to be removed. Although not necessary in this case, software such as XYZ Tools, shown in Figure 53, allows the user to choose from a variety of filters to apply to the relevant sections of data. For example, the Median Despiker option looks at a moving average and decreases all points above a threshold to the median value. Furthermore, the SR Tool option contains several options that allow the user to smooth the data, depending on the type of error.

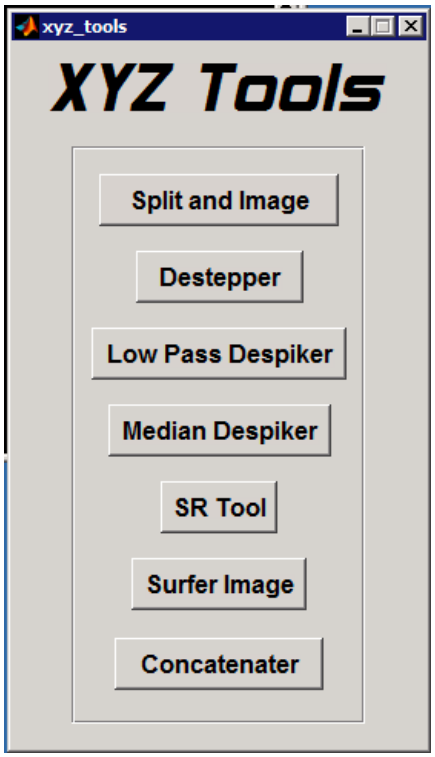

Figure 53: The user interface of XYZ Tools. This allows the user to edit processed terrain data.

4.4 Effects on Roughness

With the upgrades to the power supply, the INS should always receive an input voltage within the acceptable range for accurate measurements. Furthermore, with the updated initialization procedures, the INS should be able to correctly track the vehicle's movement after a proper alignment period. These changes should ensure the data's accuracy, resulting in good measures of roughness.

To simulate INS drift, error, in the form a random walk, will be added to the sample signals. The random walk is characterized by an autoregressive model, as found by Chemistruck [32]. It can be described as $d_i = \alpha\, d_{i-1} + \varepsilon$, where α is a coefficient ranging from 0 to 1 used to bring the drift closer to zero, and ε is the distribution of the error. In this case, α was set to 1 and ε was a normal function centered about 0 with a standard deviation of 0.15 mm. The whole function is rounded to the nearest millimeter. Figure 54 shows the profile of one instance of the generated drift error. A total 100 sample error sets were generated and averaged.

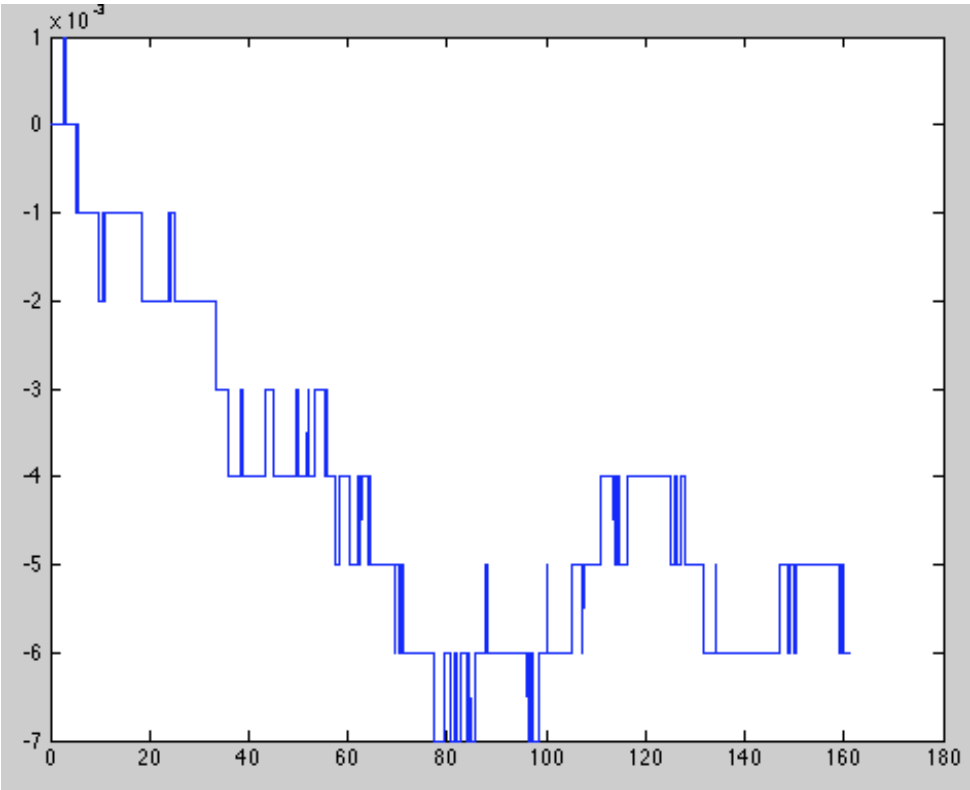

Figure 54: This is the profile of the simulated drift. height (m) vs distance along road (m)

After the simulated drift was added to each line of the three sample roads, their IRI values were calculated. This is shown below.

	Original		Original + sinusoidal		Change	Confidence	St. dev. for 90%
Road	IRI	st. dev	IRI	st. dev	in IRI	of change	confidence
smooth	1.88	0.136	1.99	0.124	0.11	90.2%	0.15 mm
medium	3.19	0.143	3.26	0.147	0.07	78.2%	0.25 mm
rough	9.69	0.446	9.71	0.446	0.02	53.0%	0.95 mm

This drift did not affect the roughness considerably. Part of the cause for this is that the minor 1 mm elevation changes gets smoothed by the IRI calculation's 250mm filter. Furthermore, there are not many changes within the total length of the road, which, when averaged, become an even smaller contribution to the IRI. However, it is worth noting that if the profile of the road is visually inspected (such as Figure 31), the lines will be visible. While they may not have a significant impact on the roughness, it will be apparent that the data is incorrect.

To simulate errors caused by bad INS measurements, a sawtooth pattern will be added to the road (as seen in Figure 59). This often occurs when the GPS signal is lost and

regained, but the software is attempting to correct the course as the INS has veered slightly off track. The simulated sawtooth, with a frequency of 2 Hz and amplitude of 5 mm (determined empirically), is shown in Figure 55. The offset of the signal does not affect the IRI calculation (as it only looks for relative changes), so the offset of the simulated sawtooth will not affect the IRI.

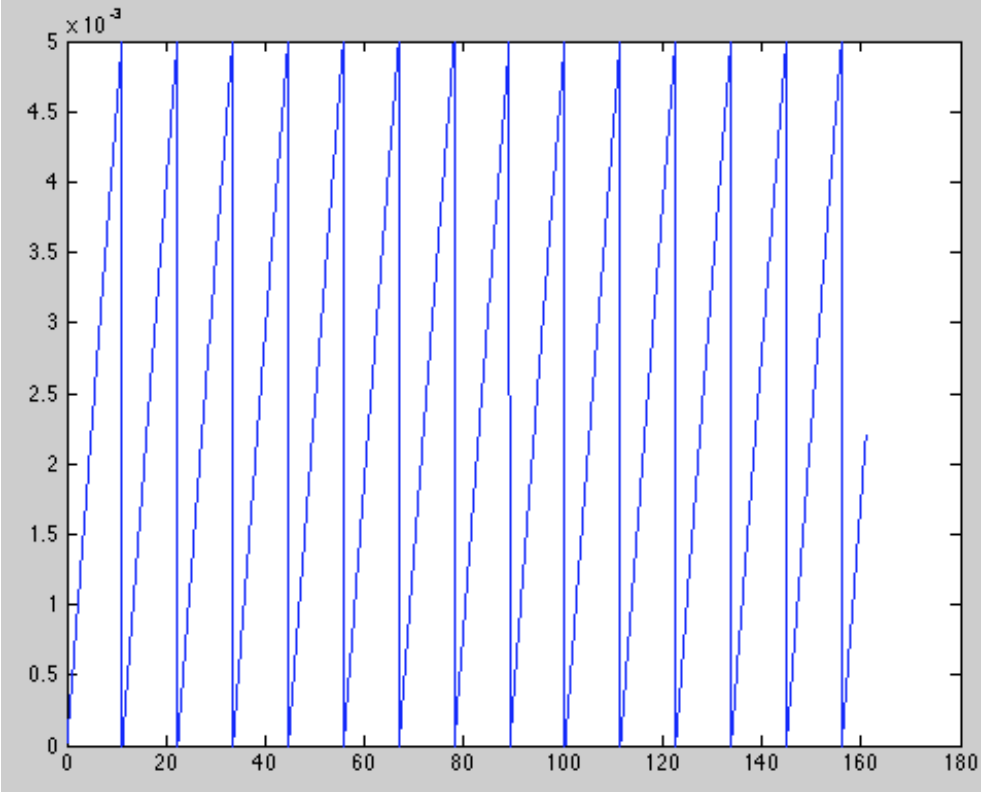

Figure 55: The simulated sawtooth error. It has an amplitude of 5 mm and a frequency of 2 Hz. height (m) vs distance along road (m)

The new IRI values (after error was added) are shown below.

	Original		Original + sinusoidal		Change	Confidence	Amplitude for
Road	IRI	st. dev	IRI	st. dev	in IRI	of change	90% confidence
smooth	1.88	0.136	2.45	0.125	0.57	99.9%	1.9 mm
medium	3.19	0.143	3.59	0.070	0.40	99.9%	2.6 mm
rough	9.69	0.446	9.75	0.416	0.06	59.8%	9.5 mm

The sawtooth error had a bigger effect on the IRI of the roads. Both the smooth and medium roads' IRI values increased significantly. However, the sawtooth error did not affect the brick path as much, and the smoother road was affected most of all. This is probably because the smoother roads have more possibilities to change at the peaks; the

rough roads may have rough values that slightly negate the effects of the increased roughness. Thus, it is reasonable to believe that the high amplitude necessary for 90% confidence for the rough road is because of the high variation already present in the road. Not only does the additional error have to change the filtered profile, the larger standard deviation also requires a bigger change to be significant. In addition, its 2 Hz frequency means that it has fewer than seven times the number of excitations as the sinusoidal error. This alone means that the amplitude must increase to account for the fewer instances of roughness. Nonetheless, it is apparent that this type of error would increase the IRI of most normal roads.

The frequency of the sawtooth error was also changed, with the results shown below.

	0.5Hz	1Hz	2Hz	3Hz	4Hz
smooth	1.92	2.06	2.45	2.90	3.38
medium	3.19	3.28	3.59	3.97	4.40
rough	9.67	9.68	9.75	9.88	10.07

This shows that the higher the frequency, the greater the change in IRI. This is logical because as the frequency increases, there are more sudden changes that are added to the road height in a given unit time.

Finally, another type of error that is commonly seen are steps in the road. These large, discrete steps occur are similar to the sawtooth error, but occur suddenly, rather than over time. This error is simulated using a length of some interval of the wavelength corresponding to 2 Hz (11.1 m), with a step size of 5 mm. The simulated profile is shown in Figure 56.

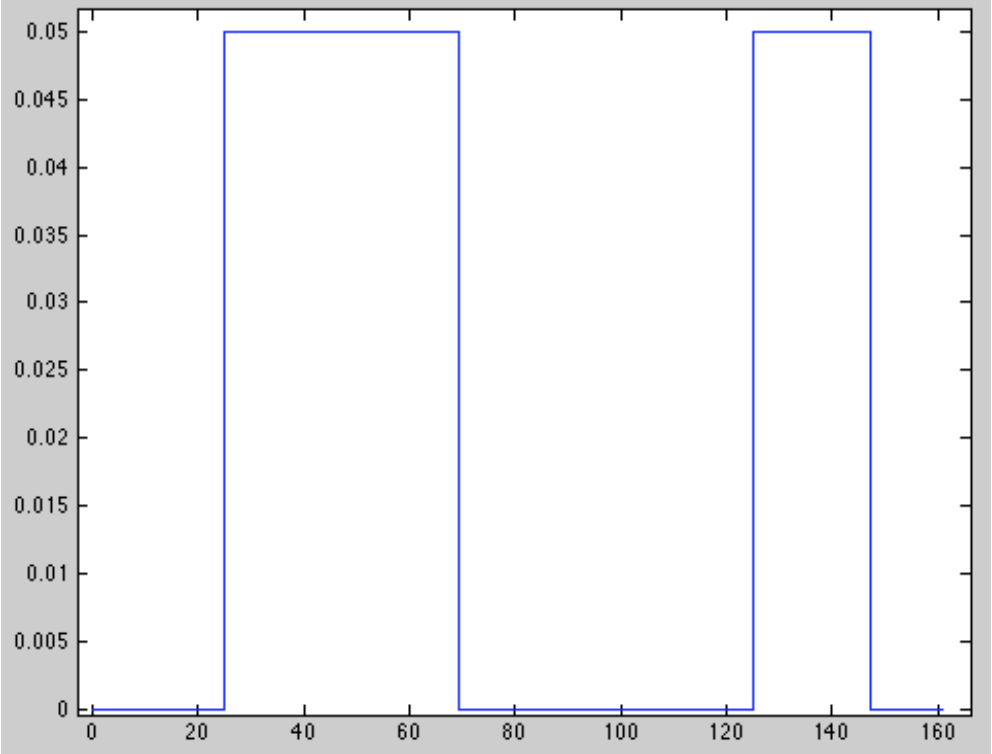

Figure 56: This simulates the step-type error. height (m) vs distance along road (m)

The resulting roughness values are shown below.

	Original		Original + sinusoidal		Change	Confidence	Amplitude for
Road	IRI	st. dev	IRI	st. dev	in IRI	of change	90% confidence
smooth	1.88	0.136	1.99	0.127	0.11	90.1%	5 mm
medium	3.19	0.143	3.32	0.126	0.13	92.7%	4.5 mm
rough	9.69	0.446	9.76	0.454	0.07	60.2%	14 mm

Again, the IRI of the smooth and medium roads were affected more than the rough road, although the significance of the change is still only around 90% confidence. This also requires high amplitudes to reach 90% confidence of change in the IRI values. Just as the sawtooth has many fewer excitations during the road, this step profile only adds four points of increased roughness. Thus, the amplitude must be much higher so that the IRI (which is an average value) increases enough to be significant. Because of the nature of the error, however, it is generally very distinguishable during a visual inspection of the data. Similarly, because of its simplicity, it is easily fixable in post processing. Regardless, it is still preferable that the error is not introduced at all, which requires a proper calibration and validation procedures.

Errors from the INS can affect the scanned road profile, leading to a variety of changes to the IRI. Any errors in the data must either be corrected, or fixed at the source. Thus, it is important to employ good practices when using an INS in order to get the most accurate results. By ensuring that the INS is receiving the correct input voltage and has enough time to align, most commonly occurring errors can be avoided.

5. Conclusions

The major goal of this work was to investigate errors commonly found in terrain profiling and characterize how they affect road roughness metrics. The major contributions of this work included:

1. Method of identifying and rectifying errors from the process of measuring terrain profiles
    a. Movement in the sensor mounting frame
    b. Inaccuracies introduced by the INS
2. Demonstrated the effects of noisy data on IRI values

Road profiling has many applications, from aiding vehicle design to the focus of this work – road roughness evaluation. Technologies have progressed from simple sensors that use a rod to determine the road height along a line to complex, 3 dimensional laser scanners, capable of scanning over a million points a second. However, with the new sensors comes a new set of challenges – guaranteeing the accuracy of the measurements. Measurement accuracy is important to ensure that the data correctly describes its physical counterpart. Obviously, this incorrect data may no longer be valuable to the user. Understanding the errors that may occur with the measurement system allows for the detection and prevention of these errors during data acquisition, leading to clean results. Simulation results show that the road roughness, especially of relative smooth roads, is quite susceptible to artificial increases due to noisy data.

First, the mounting system for a profiling system was investigated. As modern systems use an INS to track the position of the laser scanner, these sensors must be coupled together to ensure proper tracking. A good mounting frame behaves as a rigid body, permitting only negligible movement between the INS and the laser scanners, thus allowing the motion of the INS to correlate with that of the lasers. The mounting frame for the Ford Explorer mounting frame was analyzed, both during design and after completion. It was found to be rigid enough for the INS to track the laser movements given the expected input excitations from the Ford Explorer.

The inertial navigation system was also examined. The INS, composed of a GPS receiver and inertial measurement unit, tracks location and orientation. This allows for 3 dimensional profiles to be created. However, any errors from the INS are propagated to the data, which may cause bumps or other artifacts in the resulting data that do not

actually exist. Noisy data from the HMMWV profilometer pointed to INS issues, first indicating possible IMU drift. This was attempted to be solved by upgrading the power distribution system such that it output a more constant voltage, even as the batteries ran low. The error decreased but was not completely fixed. The next step was to ensure the proper alignment of the INS during operation. More thorough procedures were developed to give the INS to have enough time to align. This allows for better tracking, especially in cases where the GPS reception may be weak, thus solving the problem of the INS error.

6. Future Work

Two major drawbacks found during this work were the lack of an automated filter (XYZ Tool was mostly manual) and the loss of information in the IRI calculation. The following shows some preliminary work in developing a filter and an improving upon the IRI.

6.1 Automated Filter

A temporary software solution for this issue was beginning to be developed. A filter for the elevation was implemented (though this could also be applied to the position data if necessary). This filter identifies values beyond a certain threshold (based on deviation from the slope of the elevation immediately before the point) and attempts to correct it. Figure 60 show the filter in effect (red) compared to the original data (blue).

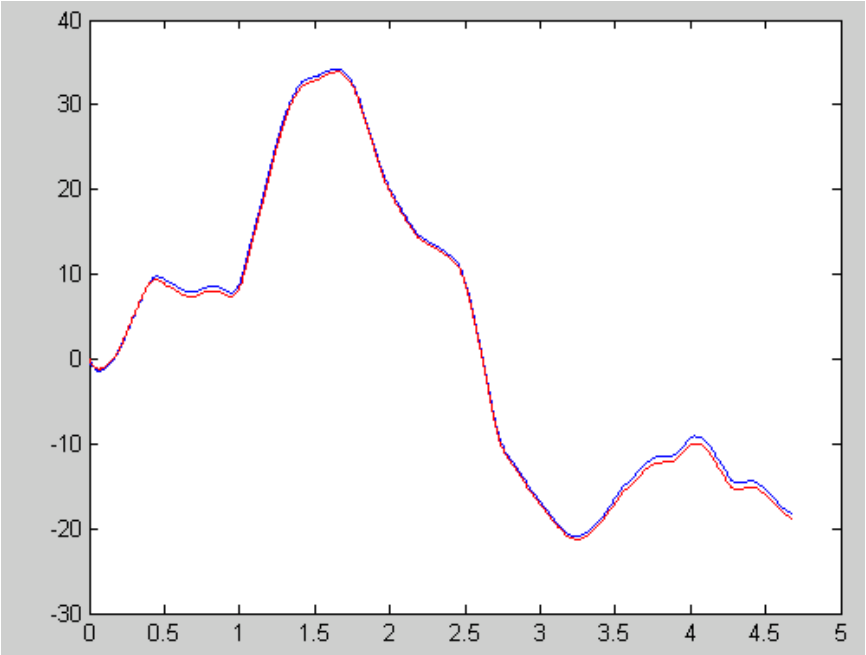

Figure 57: Elevation data of scan; original (blue) compared to filtered (red). Elevation (m) vs scan number (in ten thousands)

This is the overall elevation data for the last scan. Note that there is a small difference in elevation at the end of the data between the filtered and unfiltered data. This offset between the filtered and original data is a section that requires more work.

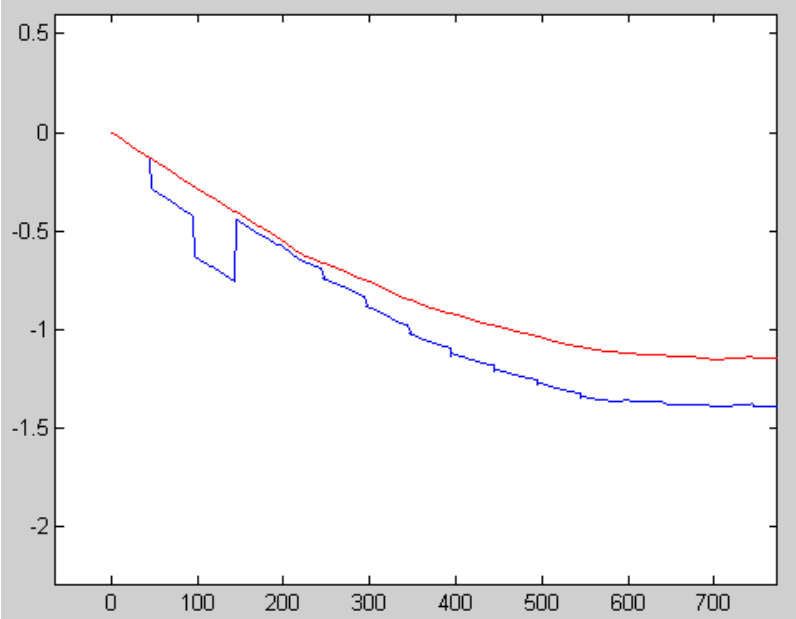

Figure 58: A zoomed in view of the beginning of the scan. A big difference can be seen between the original (blue) and filtered (red) data. elevation (m) vs scan number

Figure 58 shows the first part of the scanned road. Note that the steps have been smoothed away. However, the algorithm is not very effective when the noise is more rounded and less abrupt, as seen in Figure 59. Although this filter is not complete, the ability to develop a filter to address this issue shows promise. A completed filter, with adjustable parameters, should be able to be adjusted for any type of scanned terrain.

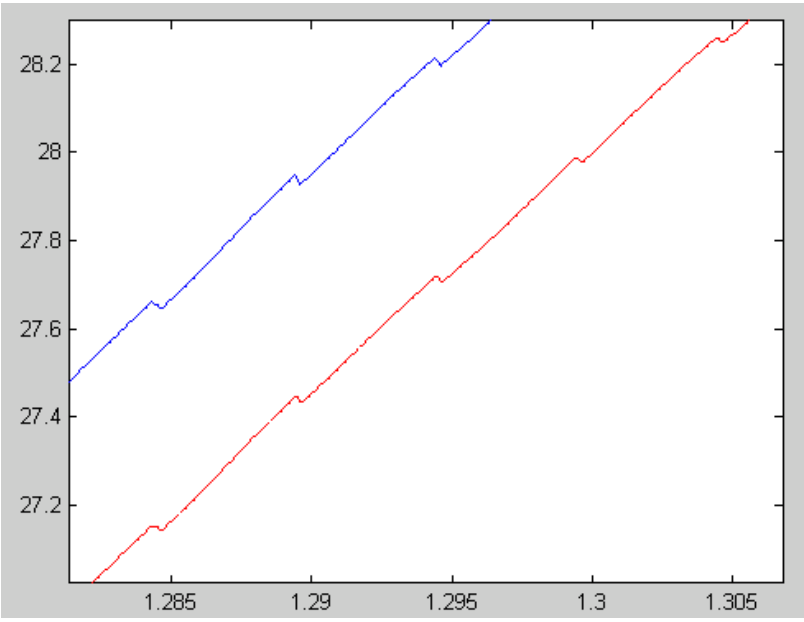

Figure 59: Small values of noise are where the filter does not work well. elevation (m) vs scan number (ten thousands)

6.2 Weaknesses of the IRI

The international roughness index (IRI) is a widespread measure of road roughness. However, one major drawback of the IRI is that it is only an average value. Therefore, roads with drastically different profiles may have the same IRI value. For example, both road profiles shown in Figure 60 have an IRI of 1.39 mm/m. However, they are clearly different. While the IRI may be a good measure for conventional roads, it clearly does not differentiate between different events in the road.

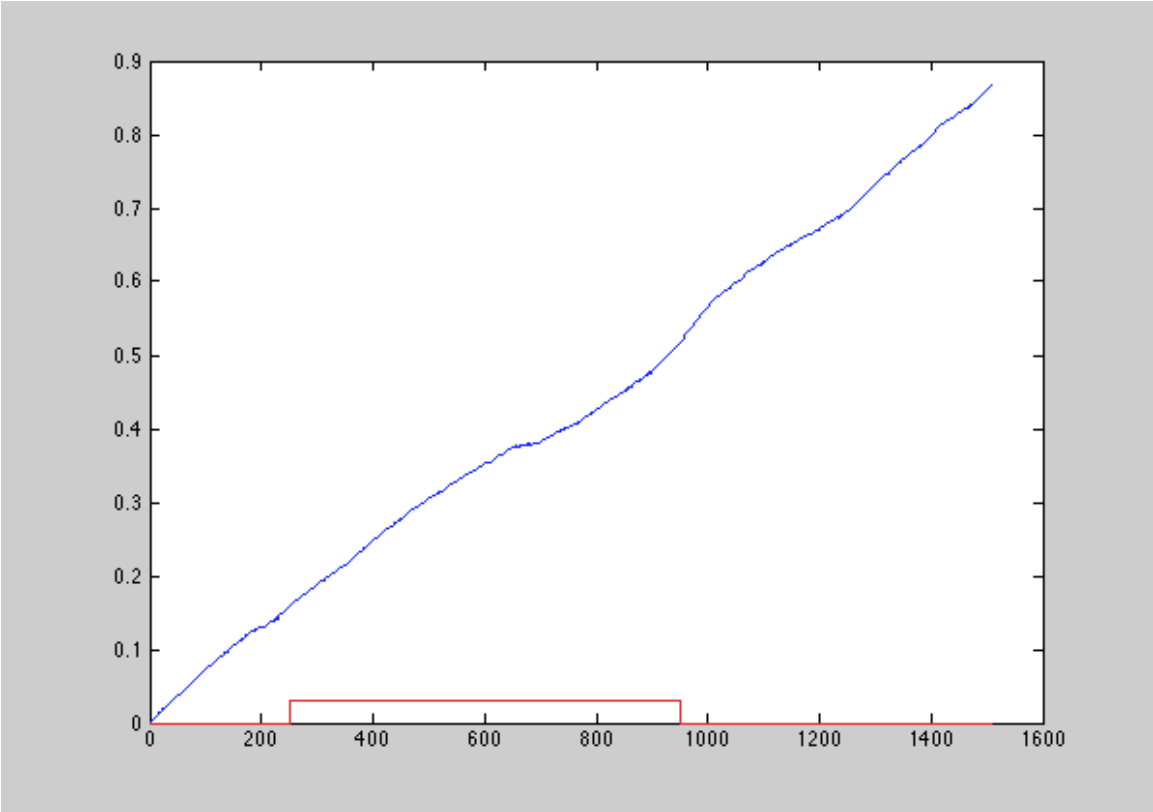

Figure 60: Two road profiles. height (m) vs data point

6.3 Developing the CRI

The idea of the continuous roughness index (CRI) is to provide the same information as the IRI, but to also show the contribution of each continuous section of road. With this information, the areas of high roughness can be pinpointed. This means that the cause of the roughness, whether it is unevenness throughout the length or a single, discrete event, can be seen. By definition, the sum of the CRI will equal the IRI for that particular section of road, so it will remain backwards compatible with existing applications of the IRI.

To begin, the IRI must first be examined. The IRI is based off the quarter car model, which is shown in Figure 61.

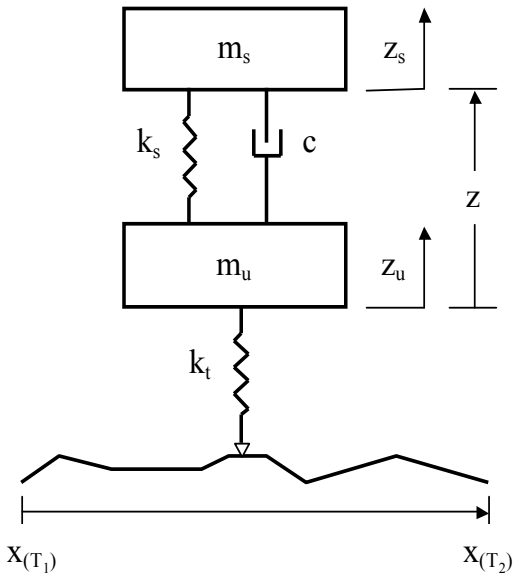

Figure 61: quarter car model. The unsprung mass (m_u) and tire stiffness (k_t) represent the tire, while the sprung mass (m_s) represents a quarter of the body. The properties of the spring are represented by the stiffness (k_s) and damping (c) coefficients, with z being the suspension travel.

The IRI is defined to be

$$IRI = \frac{1}{L}\int_0^T |\dot{z}|dt$$

where L is the length of road, T is the time it takes to travel from the beginning to the end of the road (at the defined 80 km/h), and $\dot{z}$ is the velocity of suspension travel. As the calculations must deal with discrete profile data, the discrete form of the IRI equation is

$$IRI = \frac{1}{L}\sum_{j=1}^{N}|\Delta z_j|$$

where there are N data points, and Δz_j is the change is suspension position at a given index.

The continuous roughness index aims to attribute the suspension travel to its cause. However, because most road profiles do not consist of only one excitation event, the contributions to the suspension travel must be calculated. The contribution fraction, f_{ij}, is defined as the fraction of the response at index j due to excitation from index i. Therefore, the CRI due to excitation from index i is defined to be the sum of the absolute

values of the contribution fraction multiplied by the suspension travel from *i* to the end of the given road profile, all divided by the total length of the profile. It can be written as such:

$$CRI_i = \frac{1}{L}\sum_{j=i}^{N}\left|f_{ij}\,\Delta z_j\right|$$

At a single index value for response, the sum of the contribution fractions of all the points of excitation equals zero. Furthermore, all contribution fractions are between 0 and 1. Finally, for all points of the excitation being after the response, the contribution for that point must be zero.

6.3.1 Finding the CRI

A preliminary program was written in Matlab to calculate the CRI. To ensure accurate calculations of the CRI, this calculation method uses the same suspension travel as determined by the ASTM standard. Taking a sum of the elements of the CRI matrix will give the IRI, while separately, the matrix shows the individual contributions of the road to the overall roughness.

The output CRI matrix can then be used for multiple applications. The most obvious is to find events of high roughness, which may lead to ride discomfort or vehicle damage. By applying a threshold to the CRI matrix values, locations of interest to the user may be easily distinguished. Similarly, a graphical approach also allows the user to see locations on the road profile that contribute significantly to the overall roughness. The sampling width of the CRI can also be changed. For example, if the provided road profile data has a very high sampling rate, the data points would be very close in real space. Users may not need resolution down to the nearest centimeter; instead, by changing the sample size, each user can tune the size to whatever best fits the application. If working on road repairs, having 1-3 meter steps would be sufficient for the user to find any issues on the real road.

One minor issue with the calculation and assigning the source of suspension travel is the 250mm moving average filter. Because this flattens some events out, single events may be shown on the CRI as having discrete, multiple contributions due to the filter. However, this issue can be fixed by increasing the sampling width to be greater than the 250mm filter width. This would allow for the multiple contributions to be summed

together, negating the issue of the filter separating excitations. Moreover, this is only an issue with certain road conditions (areas of road with high slope) and will not manifest itself in everything. Finally, the problem can be easily skipped by temporarily disabling the filter altogether.

6.3.2 Applying the CRI

Recall the example given in section 6.2 Weaknesses of the IRI, shown in Figure 60. The two road profiles both have IRI values of 1.39 mm/m. However, from the shape of the profiles, it is clear that they are different; one has some roughness as it progresses, while the other has a step in the middle. Thus, the IRI is not able to completely characterize the roughness of these two roads. On the other hand, the CRI can show the differences of each fairly easily.

Figure 62 shows the CRI overlaid with the road profile. Because of the nature of the road, the roughness only occurs when the road is going up and going down from the step. Thus, this figure shows that the main contributions of roughness to this road are at 25 and 95 meters, where the step begins and ends.

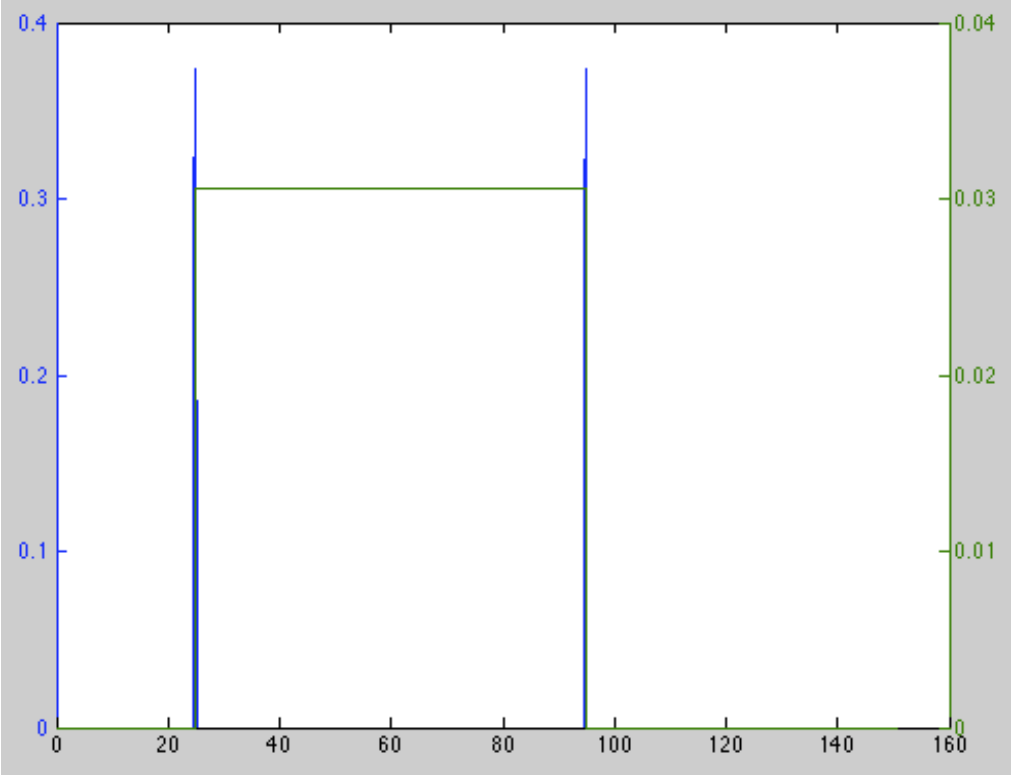

Figure 62: sample road profile 1: CRI (mm/m, blue) and road profile (m, green) vs distance along road (m)

Furthermore, by zooming in to one of these areas (at 25m in this case), the two excitations locations become more apparent. This is shown in Figure 63, where there are two separate points of excitation. This is because the 250mm moving average filter smooths the step up over the 0.1m spacing of the road profile over 0.3m, thus slightly changing the profile that the program uses when it calculates suspension travel. Finally, a sum of the CRI values equal 1.39, the same as the IRI, as expected.

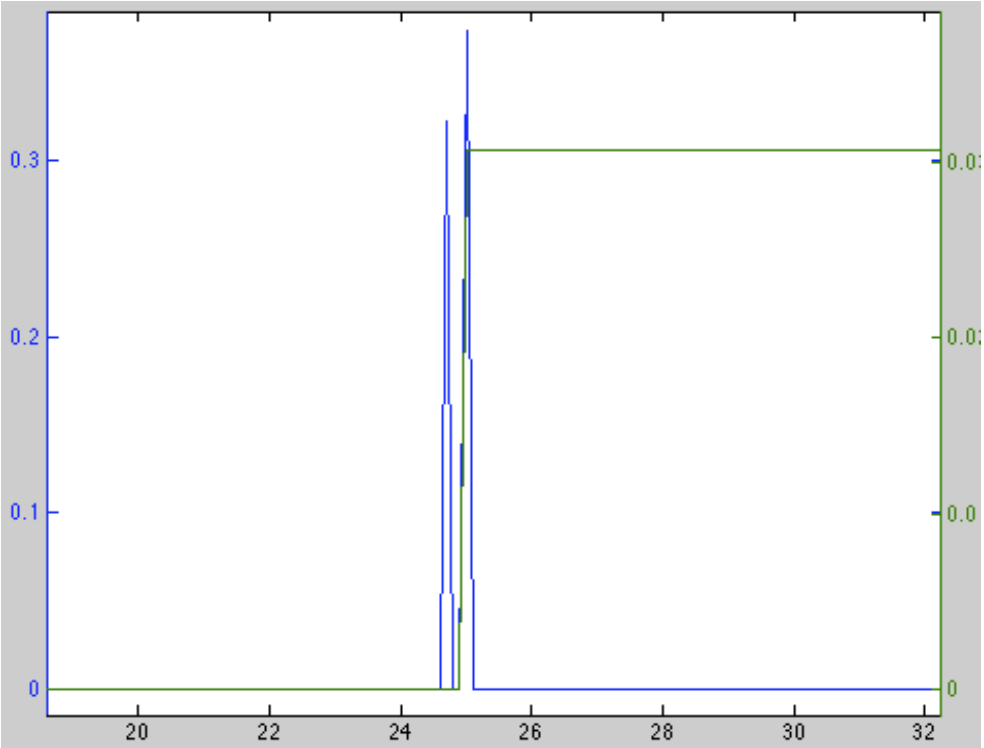

Figure 63: zoomed in at a distance of 25 m; CRI (mm/m, blue) and road profile (m, green) vs distance along road (m)

Conversely, Figure 64 shows the other road profile overlaid with its CRI. There are a number of differences it has compared to the first sample road profile, despite having the same IRI. First, the CRI values for the points are much smaller; whereas the first sample has spikes around 0.35 mm/m, the second sample's maximum CRI is only around 0.005 mm/m. Moreover, this is the peak value, occurring at only about four locations. The CRI at most other areas of the profile are lower than 0.002 mm/m. Looking at the plot of the profile, large changes in the slope can be seen at the positions where the CRI is higher, which is reasonable, as this would intuitively cause more suspension travel.

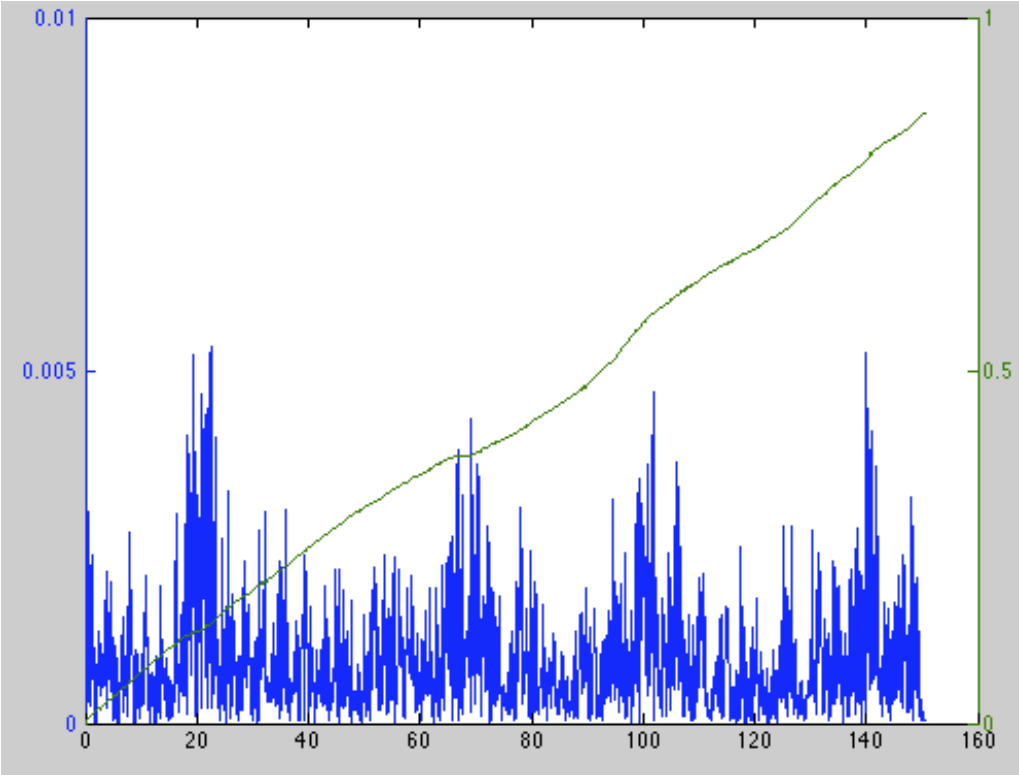

Figure 64: sample road profile 2: CRI (mm/m, blue) and road profile (m, green) vs distance along road (m)

Because of the high noise from the CRI, this would be a good candidate to apply a larger sample width. Using a sampling width of 15 data points (or 1.4 meters), Figure 65 shows the output CRI plot. This sums 15 consecutive CRI values at a time, which makes the out the areas of lower roughness more smooth. However, it does not lose any data and maintains the four areas of higher roughness. Adjusting the sampling width will change the smoothing effect that is applied.

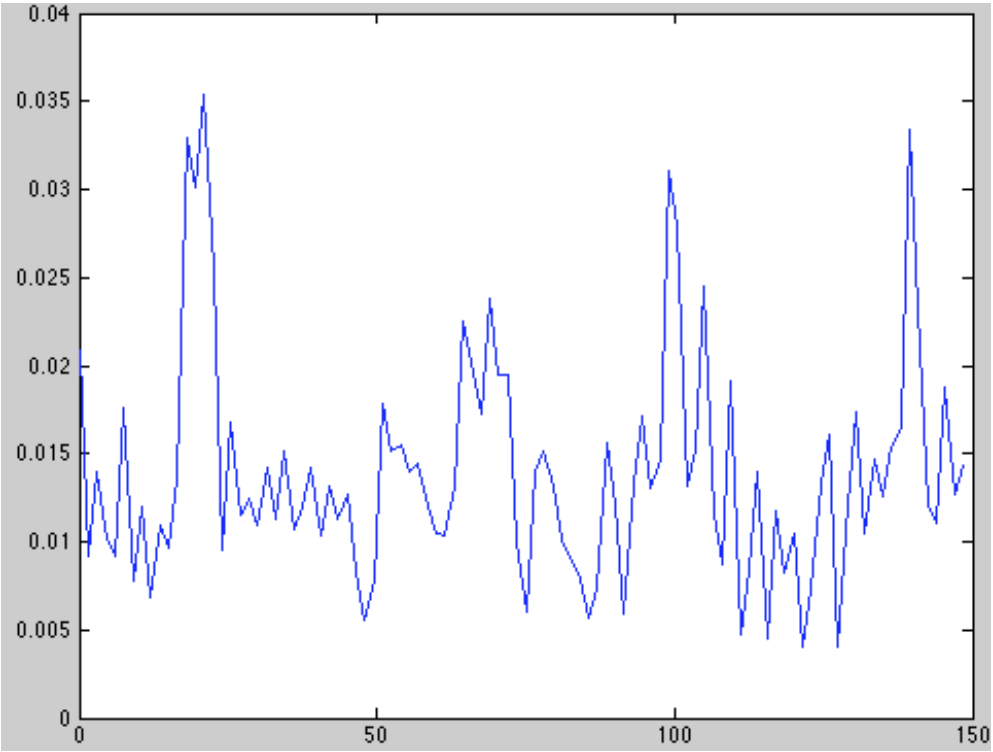

Figure 65: CRI with a sampling width of 15 data points. CRI (mm/m) vs distance (m)

6.4 CRI Conclusions

The CRI is a useful tool that expands upon the functionality of the traditional IRI. By being able to isolate the specific locations of increased roughness, users can more easily pinpoint areas of interest on real roads. Other advantages include more user-specific settings, which may differ depending on the application. Furthermore, by setting threshold values for the CRI, either as absolutes or as percentiles, the CRI can be used as a method of event detection along the road profile. Further work is required to fully develop the CRI concept.

References

1. Sayers, M. W., Gillespie, T. D., & Paterson W. D. 1986. *Guidelines for Conducting and Calibrating Road Roughness Measurements.* Washington, DC: The World Bank.
2. Sayers, M. W. 1989. Two Quarter-Car Models for Defining Road Roughness: IRI and HRI. *Transportation Research Record*, 1215, 165-172.
3. Sayers, M. W. & Karamihas, S. M. 1996. Estimation of Rideability by Analyzing Longitudinal Road Profile. *Transportation Research Record*, 1536. 110-116.
4. Sayers, M. W. 1995. On the Calculation of International Roughness Index from Longitudinal Road Profile. *Transportation Research Record*, 1501, 1-12.
5. ASTM. 2008. *Standard Practice for Computing International Roughness Index of Roads from Longitudinal Profile Measurements* (E1926-08).
6. Kropac, O. & Mucka, P. 2005. Be Careful when Using the International Roughness Index as an Indicator of Road Unevenness. *Journal of Sound and Vibration*, 287 (4-5), 989-1003.
7. Dembski, N., Rizzoni, G., and Soliman, A. 2006. Development of a Terrain Severity Measurement System Utilizing Optical Lasers. *Defense and Security Symposium*. 622805-622805. International Society for Optics and Photonics.
8. Dembski, N., Rizzoni, G., and Soliman, A. 2006. The Development of a Terrain Severity Measurement System. *ASME 2006 International Mechanical Engineering Congress and Exposition*. 217-226. American Society of Mechanical Engineers.
9. Leedy, B. M., Putney, J. S., Bauman, C., Cacciola, S., Webster, J. M., & Reinholtz, C. F. 2007. Virginia Tech's twin contenders: A comparative study of reactive and deliberative navigation. *The 2005 DARPA Grand Challenge*, 155-182.
10. Shan, J. & Toth, C.K. (Eds.). 2008. *Topographic laser ranging and scanning: principles and processing*. Boca Raton, FL: CRC Press.
11. Morin, K.W. 2002. *Calibration of Airborne Laser Scanners*. (Master's thesis). University of Calgary. UCGE Report 20179.
12. Clarke, T. A., Grattan, K. T., & Lindsey, N. E. 1991. Laser-based triangulation techniques in optical inspection of industrial structures. *Optical Testing and Metrology III: Recent Advances in Industrial Optical Inspection*, 474.
13. NOAA. 2014. Space Segment. *Gps.*gov. http://www.gps.gov/systems/gps/space/
14. Clynch, J. R. 2003. *The Global Positioning System*. http://www.oc.nps.edu/oc2902w/gps/gpsoview.htm
15. Adrados, C., Girard, I., Gendner, J. P., & Janeau, G. 2002. Global positioning system (GPS) location accuracy improvement due to selective availability removal. *Comptes Rendus Biologies*, 325(2), 165-170.
16. NOAA. 2014. Selective Ability. *Gps.gov*. http://www.gps.gov/systems/gps/modernization/sa/
17. Wing, M. G., Eklund, A., & Kellogg, L. D. 2005. Consumer-grade global positioning system (GPS) accuracy and reliability. *Journal of Forestry*, 103(4), 169-173.

18. Cannon, M.E. 1990. High-Accuracy GPS Semi-Kinematic Positioning: Modeling and Results. *Navigation*, 37(l), 53-64.
19. Morgan-Owen, G. J., & Johnston, G. T. (1995). Differential GPS positioning. Electronics & Communication Engineering Journal, 7(1), 11-21.
20. Walsh, D., Capaccio, S., Lowe, D., Daly, P., Shardlow, P., & Johnston, G. 1997. Real time differential GPS and GLONASS vehicle positioning in urban areas. *Space communications*, 14(4), 203-217.
21. Differential Global Positioning System. Retrieved from https://www.amsa.gov.au/navigation/services/dgps/
22. NOAA. 2014. Augmentation Systems. *Gps.gov*. http://www.gps.gov/systems/augmentations/
23. Groves, P.D. 2008. *Principals of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Boston: Artech House.
24. Titterton, D.H. & Weston J.L. 2004. *Strapdown Inertial Navigation Technology*. Reston, VA: American Institute of Aeronautics and Astronautics.
25. Rainey, C. S. 2013. Error Estimations in the Design of a Terrain Measurement System. (Master's thesis). Virginia Tech.
26. Padden, G. S. & Griffin, M. J. 2002. Evaluation of whole-body vibration in vehicles. *Journal of Sound and Vibration*, 253(1), 195-213.
27. Uys, P. E., Els, P. S., & Thoresson, M. 2006. Suspension settings for optimal ride comfort of off-road vehicles travelling on roads with different roughness and speeds. *Journal of Terramechanics* 44(2007), 163-175.
28. Smith, H. 2009. Improving the Quality of Terrain Measurement. (Master's Thesis). Virginia Tech.
29. Sayers, M. W., & Karamihas, S. M. 1998. *The little book of profiling*. Ann Arbor: Transportation Research Institute, University of Michigan.
30. Yen, K. S., Ravani, B., & Lasky, T. A. 2011. *LiDar for Data Efficiency*. Washington State Department of Transportation.
31. Godha, S. & Cannon, M. E. 2007. GPS/MEMS INS integrated system for navigation in urban areas. *GPS Solutions*, 11(3), 193-203.
32. Chemistruck, H. M. 2010. A Galerkin Approach to Define Measured Terrain Surfaces with Analytic Basis Vectors to Produce a Compact Representation. (Doctoral Thesis). Virginia Tech.