Spring June 7th, 2014

Road Profile Sensor: A Detection Method for Active Suspension Systems

Matthew Edel
Seattle Pacific University

Road Profile Sensor: A Detection Method for Active Suspension Systems

Matthew Edel, SPU

***Abstract*—Active suspension systems adjust the suspension components of an automobile to adapt to bumps or potholes that are encountered in the road as the vehicle is driving. These systems have the potential to improve safety, performance, and ride comfort in automobiles. An integral part of active suspension systems is a device to detect irregularities in the road. Current detection systems that are available lack either in precision, resolution, or speed. A senior design project, Dynamic Automatic Adjusting Suspension (DAAS), at Seattle Pacific University expressed a need for a high-performance road scanner that could be paired with their suspension system. The design would need to take into account problems with latency and resolution. I therefore began development of the Road Profile Sensor (RPS). The RPS was implemented through the design of a range finder that uses a linear photodiode array paired with a laser to measure distance. The distance from the sensor to the road changes as irregularities in the road are encountered, and this change in distance is measured by the RPS to determine the size of the irregularity. The proposed system runs on a soft processor core in an FPGA chip that is both a part of the DAAS system and communicates with the DAAS suspension controller. The RPS can be sampled 62 times per second and has height resolution of 4mm. With further development, the RPS has the potential to run at very high speeds in relatively low-power, low-cost FPGA's. This design will yield much greater resolution in road scanning, which will lead to better suspension control, and a generally more reliable active suspension system. These improvements in road irregularity detection are expected to improve isolation between the road and the chassis of the vehicle, thereby improving the vehicle's handling, versatility, and safety.**

I. Introduction

THE Road Profile Sensor (RPS) is a detection system for active suspension systems in automobiles. It scans a road for bumps in real time, and sends that information to an onboard suspension controller. The RPS is a subsystem in a senior design project at Seattle Pacific University by Everan Chaffee, Bartley Hallinan, Jamey Frykholm, and Matthew Edel. The senior design project, DAAS (Dynamic Automatic Adjusting Suspension), seeks to implement an active suspension system that would adjust the spring and damping rates of the suspension system in an automobile. DAAS adjusts the suspension for specific road conditions that it detects prior to impact. The system that detects these road conditions is the Road Profile Sensor (RPS). The development of the RPS is discussed at length in this document. I will begin by describing project DAAS and the role of the RPS within the greater project, and then I will discuss the development and testing of the RPS.

M. Edel is a student in the Department of Electrical and Computer Engineering, Seattle Pacific University, Seattle, WA, 98119 USA e-mail: edelm@spu.edu.

Manuscript received June 05, 2014

A. DAAS Problem Statement

Vehicles have suspension components that help to isolate the chassis of the vehicle from the road over which it travels. This is accomplished by placing a spring and damper system between the wheels and the chassis of the vehicle. These suspension components are chosen to optimize the automobile's handling on flat ground with minor irregularities, and do not handle large bumps well. Common automobiles drive very poorly in rough situations. A vehicle that is tuned to handle large irregularities in the road, such as an off-road vehicle, will handle rough roads much better, but will not handle as well at high speeds or in sharp corners where stiffer suspension improves performance.

If a vehicle were able to change the characteristics of its suspension system, it could provide for optimal handling in different scenarios. A vehicle that could adapt in anticipation of specific irregularities in the road, like a specific bump or pothole, would handle much better in typical city driving than either a vehicle tuned for speed or a vehicle tuned for the countryside. The goal of DAAS is to develop such a system. DAAS seeks to pair an automatically adjusting suspension system with a road detection system in order to achieve the best possible all-around suspension performance.

When an automobile passes over a bump in the road that it travels on, a sequence of events push the chassis upward. The wheels of the vehicle are displaced upward by the bump, and so they compress the suspension components. This increased spring compression pushes the chassis upward. If the wheels drop down again quickly, the chassis will see less vertical displacement. If the increase in road height is prolonged, the chassis will be pushed upward until the springs decompress to a state of equilibrium. In this, the suspension components are effectively acting as a long pass filter between the wheels and the chassis.

After the vehicle has passed over the bump, the chassis will oscillate on the suspension components according to their spring and damping rates. The spring and damping rates determine how much the chassis will oscillate for a given bump, and are tuned to minimize oscillations. Different size bumps have different ideal spring and damping rates, though, so a passive suspension system cannot be ideally damped for all bumps. On account of this, a vehicle is typically fitted with suspension components that are tuned to handle general road irregularities. Speed bumps, for example, are often larger than the bumps a vehicle is intended to handle, and so the vehicle must pass over them slowly, and will oscillate considerably after they have been passed. In this case, the suspension system is considered to be under damped for the bump it has passed over.

A suspension system is considered critically damped when the chassis does not oscillate and returns to its natural state as quickly as the spring will allow without oscillation. This is the ideal damping for a suspension system. Ideally, the vertical displacement should be minimized as well, which is a joint effort between the spring and damper system.

If an automobile could adjust its suspension damping rates and spring rates, it would be able to adjust its suspension system to be critically damped for whatever road irregularities it encounters. This would give the vehicle the ability to critically damp its suspension for any given bump, and then return back to a stiffer setting. By adjusting spring rate as well, the system could minimize the amount of displacement the chassis experiences, yielding a smooth ride. The effects of this are improved handling, greater ride comfort, and a safer ride.

The goal of DAAS is to detect irregularities in the road ahead of a vehicle and adjust the spring and damping rates of the vehicle's suspension components to minimize the vertical displacement of the chassis due to the irregularity, as well as eliminate oscillation after encountering the irregularity. Other systems have been developed to achieve this goal; they have succeeded in some areas, and failed in others. DAAS does not seek to rebuild one of these existing systems, but rather to develop a system that has potential to contribute to future suspension design. A survey of existing active suspension is in the following section.

B. Existing Actuation Systems

Adjusting suspension systems already exist and are currently used by many luxury and sport automobiles. Some systems adjust for ride height by lifting the chassis of the vehicle off of the suspension components with hydraulics or pneumatics (these are mostly found in luxury SUV's.) Such an adjustment would give vehicles greater clearance over the road when raised, and improved handling at higher speeds when lowered.

Air shocks are another type of adjustable suspension component. Rather than using traditional coils as the spring mechanism, they use a piston to compress air in a cylinder to form a spring. The spring rate of these can be adjusted by changing the amount of air that is compressed in the cylinder. The issue with air shocks is primarily that adding more air to the cylinders requires a lot of force, and so are not easily adjusted in real time.

Preload adjustment is another method that is currently used by Mercedes Benz in a feature they call Active Body Control (ABC) [1]. ABC effectively adjusts the preload in the suspension system by forcing the springs together some. This changes the engagement point of the spring; it effectively keeps the springs from responding to small forces on them while allowing normal response to larger forces. With preloading, the chassis of the vehicle effectively sees a very high spring rate for small road irregularities, while maintaining the system's normal spring rate for larger road irregularities.

Finally, some suspension systems today have the ability to adjust their damping rates. Porsche Active Suspension Management (PASM) is an example of such a system [2]. PASM uses extra passageways within the damper that are electronically regulated to alter the flow of hydraulic fluid. This method is fast and responsive and could be used alongside a spring rate adjustment method.

The closest an active system has come to adjusting spring rates is Mercedes' preload system. The other existing systems only adjust for damping or ride height. While these systems are effective to some extent, they do not match the potential performance improvements of a system that can actually change spring rate. This is the goal of DAAS.

An active suspension system must be paired with a detection system. Detection systems have been employed for decades, and with, the increasing capabilities of modern technology, they are expected to continue to improve in general performance. The existing systems are indeed quite powerful, though they have just begun to tap into their potential.

C. Existing Detection Systems

A few different road detection systems have been developed for use today. These systems are generally divided by whether they detect specific road irregularities or the general condition of the road. A system that detects the general condition of the road, a bumpy gravel road for example, would be very different from a system that detects each bump in that gravel road.

General detection systems typically measure the position of the suspension components or the speed of the vehicle. Speed based systems, for example, adjust suspension settings based on the speed of the vehicle. While this approach is effective, it too is a generalized approach that does not take specific road irregularities into account.

Systems have been developed that monitor the compression of the suspension system. When the springs compress more than normal, a response is triggered from the control system. These systems have been employed in general active systems and specific active systems. In both cases, though, the detection is reactive in nature rather than proactive. DAAS seeks to be proactive by detecting irregularities prior to impact.

Specific, proactive detection systems must employ a non-obtrusive detection method; they must be able to detect changes from the plane of the road before the chassis passes over that section of road. These systems typically use image processing or Doppler techniques. Mercedes is developing a system called Road Surface Scan (RSS) that looks at the road ahead of the vehicle with a camera mounted on the car's windshield. Mercedes system is able to measure bumps to within 1 cm of precision, and is paired with their Active Body Control system. This system is very similar to the detection system DAAS seeks to design.

The DAAS system is designed in light of these existing systems. DAAS does not seek to rebuild a concept that has been tried in the past, nor does it seek to implement a system that can be purchased from an automotive supplier. The DAAS system is both novel and effective, and is outlined in the following sections.

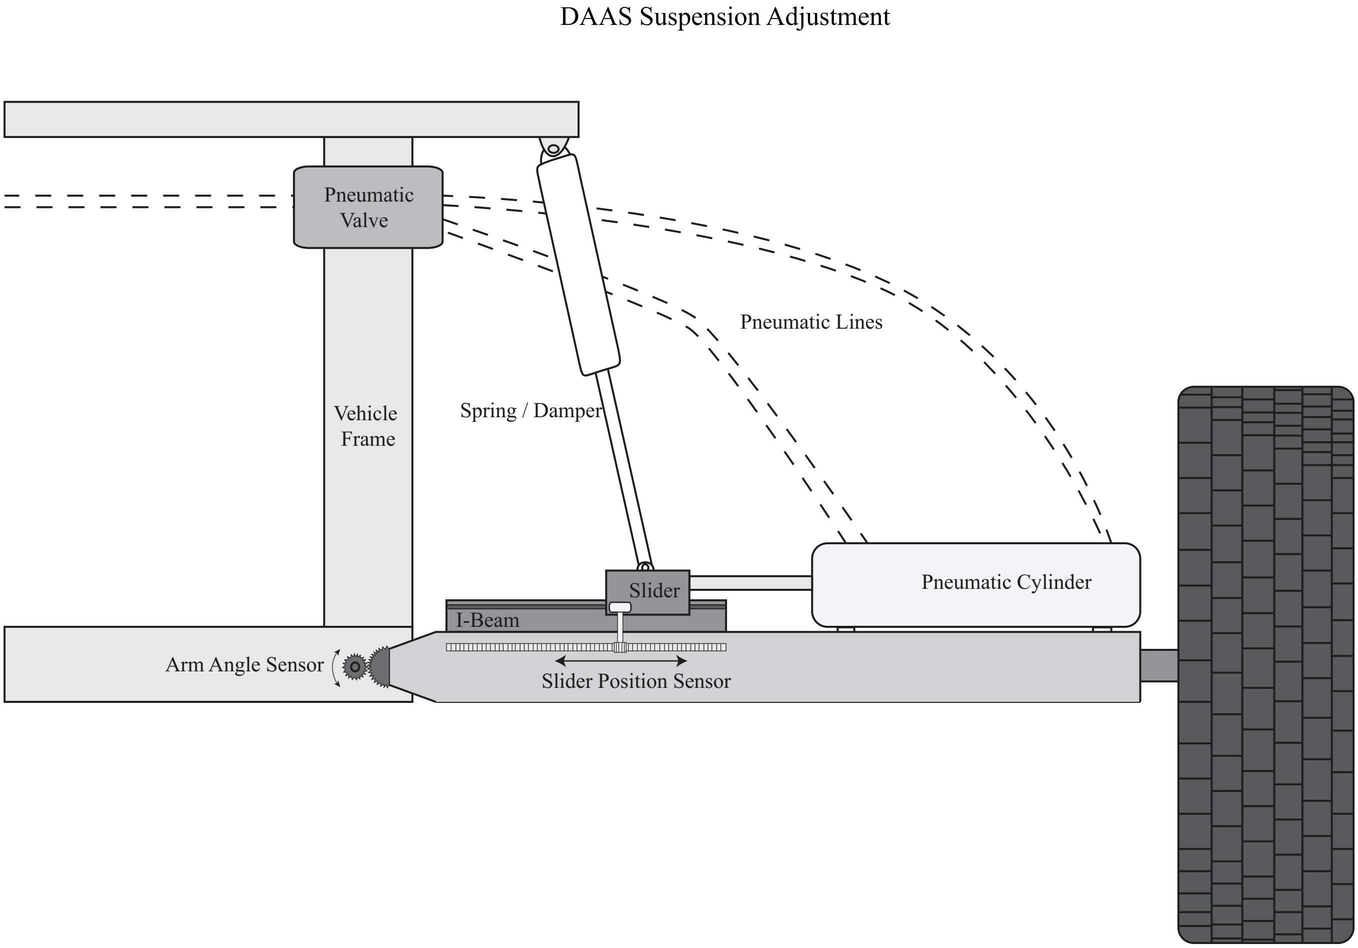

Fig. 1. DAAS Suspension System [3]

D. DAAS Suspension

The system that DAAS implements seeks to change the spring rate seen by the chassis as well as the damping rate, and to do so quickly enough to adjust proactively for each specific road condition. The mechanical implementation of the DAAS system uses standard spring and damper components, but changes their attachment point on the suspension arms to give the wheel more or less leverage on the suspension components (Figure 1). By moving the attachment point toward the vehicle, the leverage on the spring components increases, which compresses the springs to an equilibrium state.

The damping rate is not electrically controlled by the system, but the damping effectively decreases as the pivot point moves toward the chassis. This happens because the damper sees less linear displacement relative to the same angular displacement of the wheel.

To move the attachment point, DAAS connects the suspension components to a slider that moves along an I-beam on the suspension arm. The slider is pushed back and forth on the I-beam by pneumatic cylinders. The pneumatic cylinders are in turn controlled by pneumatic solenoid valves, which are driven by the electrical control system. This slider method is novel in nature. It has never been tried before in the context of the automobile. Very few systems, in fact, have been developed to change the effective spring rate seen by an automobile. The slider method is elegant in concept, though challenging in implementation; it nevertheless provides a spring rate adjustment method that DAAS seeks to implement.

The pneumatic cylinders, of course, do not actuate on their own. A computer must be used to take data from the RPS's, interpret it according to the position of the vehicle, and control the pneumatics accordingly. This is accomplished by the Electrical Control System.

E. DAAS Electrical Control System

The purpose of the Electrical Control System (ECS) is to take data from various sensors, calculate the ideal suspension settings, and control the pneumatic valves accordingly. The sensors used by the ECS return data pertaining to the speed of the vehicle, the position of the suspension components, and the profile of the road ahead of the vehicle. The ECS runs software drivers for each of these sensors in several processing cores instantiated in an economical field programmable gate array (FPGA).

The actuators in the ECS are the pneumatic hardware drivers. These drivers consist of switching transistors that supply ground to the pneumatic solenoid valves used in the mechanical implementation, and are driven by the GPIO port on the FPGA.

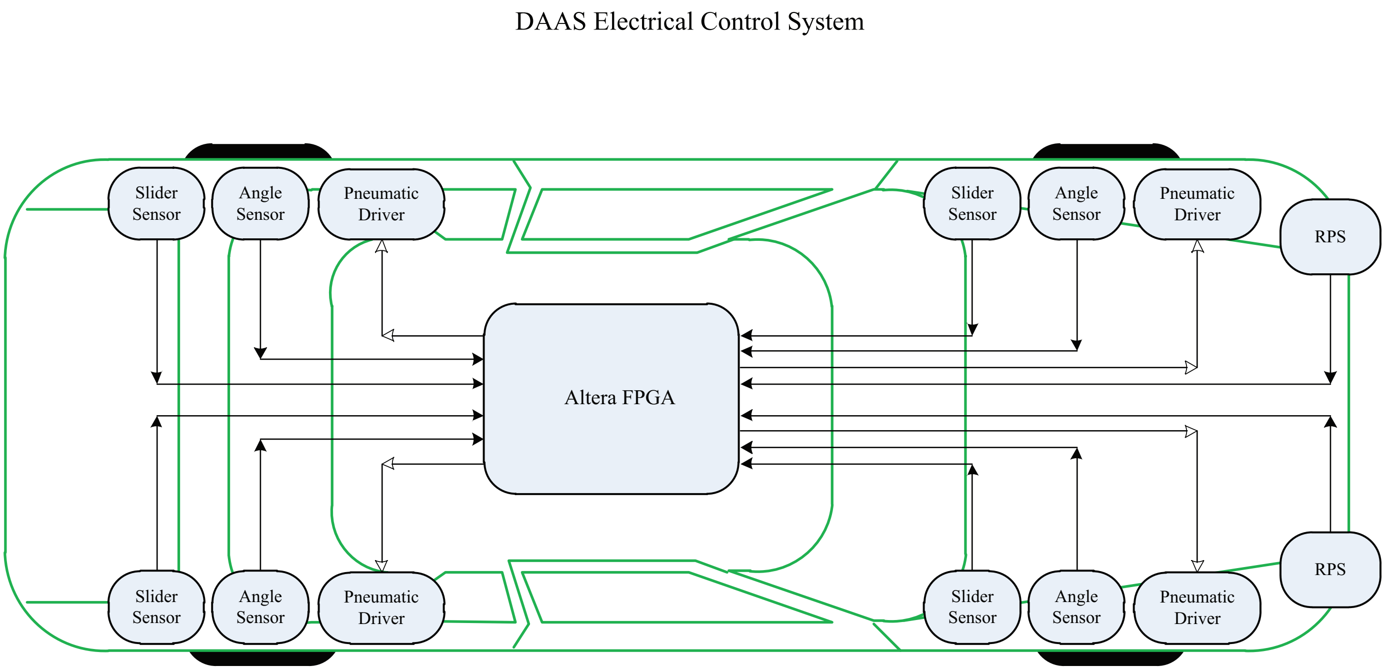

Fig. 2. DAAS Electrical Control System

The suspension sensors consist of rotary potentiometers that interface with the suspension arms (Figure 2). These sensors are used to measure both the position of the suspension slider and the angle that the suspension arm makes with the chassis. The slider position sensor is used by the pneumatic software drivers to improve accuracy in actuation, and the angle sensor is used to improve the algorithm that processes data from the RPS. The angle sensor also provides data that can be used to improve system performance. The final sensors in the ECS are the RPS's.

Each RPS has a driver that runs the communication port between the sensor and the FPGA. Each driver is run on its own processing core that is instantiated in the FPGA, and is implemented using general parallel input/output (GPIO) pins rather than through a traditional serial port. The data that is returned from the RPS is quickly processed in its driver before it is sent to main CPU. The main CPU calculates the best suspension settings from data it has taken from the suspension sensors and the RPS's and drives the pneumatic actuation system at precisely the right time, for precisely the right duration.

The RPS is a streamlined detection system that is capable of returning data to the CPU 60 times a second. This is faster than standard video cameras, and yet it runs on soft processors in an FPGA. The details of the requirements set forth for the design of the RPS are outlined in the following section.

F. DAAS Active Detection

DAAS uses an active road profile detection system. This is necessary to achieve the goal that DAAS has set forth of adapting to irregularities in the road before they are encountered. The following constraints were considered when selecting the detection system:

1) The detection system must detect changes in the plane of the road ahead of the vehicle in the vertical direction.
2) The detection system must look ahead of the wheels of the vehicle, but does not need to scan the center of the road, nor the peripherals.
3) The detection system must distinguish between changes on the left side of the vehicle from changes on the right side of the vehicle.
4) The detection system must sample the road at a minimum of 20 hertz.
5) The detection system must have a minimum bump height resolution of 1 cm.
6) The detection system must scan the road a minimum of 1 meter ahead of the vehicle.

The initial requirements developed by project DAAS were intentionally loose. The basic functionality was defined, though the implementation of that functionality was not determined.

In light of these basic requirements, a distance measurement system was chosen, (Other methods were considered, and are discussed at length in the documentation for project DAAS.) If the plane of the road could be known relative to the plane of the vehicle, a distance sensor could be used to determine the height of the road ahead of the vehicle. Effectively, a decrease in the distance to the road would correspond to a bump or rise in the road profile, while an increase would correspond to a dip in the road.

The chosen method to implement the distance measurement technique is the laser rangefinder. The laser rangefinder was chosen because it is unobtrusive by nature, has higher precision than sonar based methods, and requires less processing time than pure image processing techniques. There are two types of laser rangefinders, and they operate on two different principles: time of flight and triangulation. Time of flight measurement systems send out a laser pulse and measure the time it takes for the optoelectronic system to detect the reflection of the laser pulse. Triangulation methods, on the other hand, sense where the reflected laser pulse is received on an image sensor. Triangulation was chosen over time-of-flight because it is much more precise at the relatively shorter distances being measured by DAAS.

The development of this sensor is discussed at length in the following sections, including the design of the sensor, the testing of the sensor in different settings, the evaluation of the sensor in its intended setting, and concluding remarks about the effectiveness of the sensor as it functions within DAAS. More information on the DAAS project can be found in the documentation for the project. [4]

II. Road Profile Sensor

The Road Profile Sensor (RPS) is designed to meet specifications set forth by DAAS. These specifications are outlined in the introduction to this document. The RPS is a road sensor system for active suspension systems that uses laser triangulation to monitor the road ahead of the vehicle. The driving concept behind the RPS is discussed in the following sections.

A. Main Concept

The basic operating principle of the RPS is similar triangles. The RPS aligns optical components with a point on the road ahead of the vehicle to form a triangle. As the point on the road that the optical components are focused on changes elevation, the lengths of the sides of the triangle change as well in proportion to each other. By measuring the change in one side of the triangle with physical components, we can determine the change in location of the point on the road ahead of the vehicle.

To take this measurement, a laser and an optical detector are placed some vertical distance apart from each other. They each form one vertex of the triangle. A point some distance ahead of the laser-receiver pair is the third vertex of the triangle. The laser is focused on that point, and the optical receiver is focused to receive light from that point. This system makes up the RPS.

If the laser strikes a point closer to the RPS, the reflected laser light will be detected by a different pixel on the optical detector (figure 3). As the optical system detects this change, it is able to calculate the change in the location of the point that is reflecting the laser light.

If the RPS is mounted on a vehicle oriented toward the road, and the laser and optical receiver are both focused on a point on the road ahead of the vehicle, when a bump is encountered by the laser, the position of the laser point will change. This change corresponds to the bump height.

The RPS is effectively calculating the change in the length of one side of a triangle, the distance between the laser source and the pixel that receives the reflected laser light. Based on this change in distance, the change in the height of the road

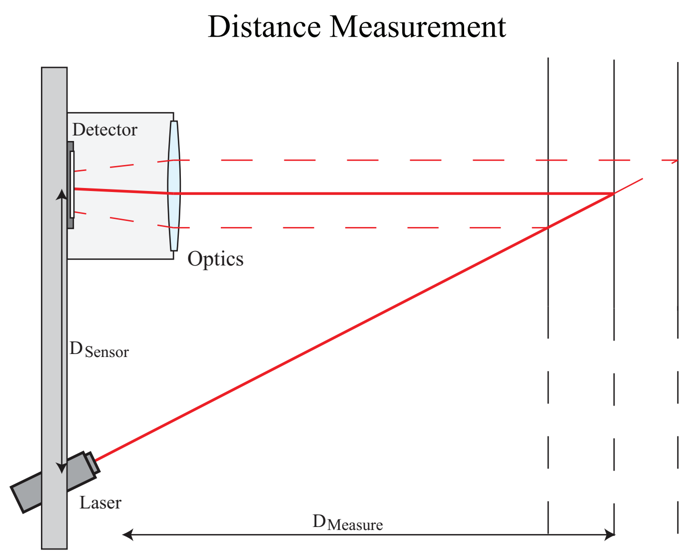

Fig. 3. Distance Sensor Geometry

relative to the RPS can be calculated. This is illustrated in figure 3.

Several factors are considered in setting the basic parameters for this design. They include the default location of the focus point of the laser and optical detector, the distance between the laser source and the optical detector, and the height at which the RPS is mounted above the ground. These parameters are determined through testing and are discussed later in the document. First, the major systems were designed, and components were chosen to begin testing before the design was finalized.

B. Major Components

The RPS is made of two major subsystems, a laser source and an optical detector. The selection of the laser and optical detector is outlined below.

1) Laser Diode: There are several factors that need to be considered when selecting a laser diode for the system. The major aspects that need to be considered are the power output of the laser and the wavelength of the laser.

The power output of the laser must be considered for both feasibility and for safety. As the power output of the laser increases, the feasibility of detecting the laser via the optical detector increases. The logic is simple; if more light is focused at a point on the ground, more of it will be reflected back to the optical detector. The effects of the increase in received light are:

- Increased operating speed: The sensor will not need to allow as much time for the light to be received because the light intensity has increased.
- Increased range: The distance that the laser will be able to project will increase with light intensity. The amount of light returned to the optical detector decreases with distance because of changes in angle, so to increase distance, a stronger laser must be used.
- Increased reliability: If the amount of laser light reflected by the road increases, but the amount of noise from other

light sources does not increase, the signal to noise ratio (SNR) will grow. As the SNR increases, the reliability with which the signal can be detected will increase as well.

The problem that arises is that as the power of the laser increases, the safety of the system decreases. Lasers are classified for safety in terms of their output power. [5] Low power lasers are typically only dangerous to an individual's eyesight. Prolonged exposure could cause burning of the retina. As the power output and classification of the laser increases, the risk also increases. Some of the laser classifications are outlined below:

- A class II laser has output power no greater than 1 milliwatt (mW). Class II can cause burning of the retina if exposed for several seconds. This would occur significantly after the blink reflex, which will prevent eye damage.
- A class IIIR laser has output power no greater than 5 mW. These lasers pose slight risk to cause eye damage before the blink reflex, though significant damage occurs after the blink reflex. They are considered to be potentially dangerous.
- A class IIIB laser have output power greater than 5mW and can cause permanent, immediate eye damage. These are dangerous, as even reflected light can cause damage if it enters the eye.

The RPS seeks to maximize laser output power, but at the same time to keep from developing a potentially dangerous product. The balance is achieved by using a 5mW laser at a high wavelength. This maximizes output power while not advancing to the IIIB classification. In addition to the lower IIIR classification, the laser selected is a high wavelength laser. The eye is less sensitive to high wavelength (Infrared) light, as it is in fact invisible to the human eye. This helps to minimize the risk from the class IIIR laser.

In addition to safety factors, the wavelength was chosen for other reasons as well.

The wavelength of the laser is important because materials reflect different wavelengths to varying extent, image sensors are more sensitive to different wavelengths, and ambient noise exists to varying extents at different wavelengths. The following three factors were considered in the selection of the laser wavelength:

- Reflectance: The material that will reflect the laser in operation is typically asphalt. Asphalt has much greater reflectance at higher wavelengths, increasing steadily up to 20 percent reflectance at 800nm, and increasing, though more gradually, from there.
- Sensitivity: Image sensors, and integrated circuits at large, are made of silicone and are somewhat more sensitive to infrared light than to lower wavelengths.
- Noise: The system is intended for use in the outdoors, where daylight fills the environment with light from varying wavelengths. The sun shines a very large amount of visible light, as well as some ultraviolet (UV) and infrared (IR) light [6]. Infrared light is less intense compared to visible light in sunlight, so noise in the infrared spectrum is lower than in other spectra. This being said, infrared light is still present in sunlight and poses a challenge to the system.

On account of these conditions, an infrared laser was chosen for the design. The specific wavelength is 780 nanometers (nm). This was chosen for the following reasons: there is very little ambient noise at IR wavelengths, asphalt reflects nearly as well at 780nm as it does at greater wavelengths, the image sensor operates very well at this wavelength, lasers at this wavelength are readily available, and the light is on the edge of the visible spectrum. The visibility of 780nm makes this laser much easier to work with in development than a laser outside of the visible spectrum.

A laser of higher wavelength was not used because 780 nm exhibits the benefits of high wavelength lasers while maintaining visibility. The system is much easier to calibrate when a laser is used in the visible spectrum. The benefits of high wavelength lasers in reflectance and sensitivity were not very significant relative to the 780 nm laser.

Laser Diodes need a lens to focus the beam, housing to mount the lens above the laser diode, and a driver to supply regulated current to the diode. The laser diode should be chosen with a housing, driver, and lens as a package. Such a configuration is chosen for this project because it is more cost and time effective than designing and assembling a similar system by hand.

2) Optical Detector: The optical detector is an integral part of the system, and it consists of two core components: an image sensor and a lens to focus the reflected laser point onto the image sensor. The selected image sensor is actually not a traditional rectangular pixel array, but rather a linear photodiode array (LPDA), and it is mounted behind a single biconvex lens.

The role of the image sensor is to detect a change in distance from the pixel that is illuminated by the reflected laser dot. The measured change only needs to be in a single dimension, as the only measured dimension is road height. In light of this, a linear pixel array can be used instead of a rectangular pixel array. A linear array accomplishes the same task as a rectangular array, but it does so with much greater simplicity. The specific details of the pixel array are:

- Pixel density: The pixel density directly affects the accuracy with which the reflected laser point can be resolved. As the pixel density increases, the pixel size decreases.
- Size of Array: The number of pixels in the array is important. As the number of pixels in a linear array increases, the precision of the measurement increases. The time required to read and process the data increases as well, which ultimately hurts the sample rate.
- Communication Protocol: The time required to sample the sensor and process the data needs to be minimized so as to maximize the amount of lead time that the mechanical system has to react. In order to minimize sample time, the array should be able to efficiently send data to the CPU.
- Pixel Type: There are various types of pixel technology. The charge coupled device and the complementary metal-oxide semiconductor (CMOS) are the two predominant

image detection technologies. They tend to be more sensitive than photodiode arrays (PDA), though they also have more noise and less range, making it easy to drive them into saturation. For the initial development of the RPS, the PDA will be easier to develop with because of these factors. At a later point in time, it may be good to use a CMOS array for increased sensitivity and faster operation [7].

A linear pixel array was chosen over a rectangular array for the aforementioned reasons. A photodiode array was chosen because of its simplicity, as well as availability. Linear arrays are much simpler than rectangular arrays, but they are also much less common, and finding one that was relatively easy to develop with was challenging. A linear array built on a different pixel technology could not be found.

The image sensor must have an optical system to focus the reflected laser dot onto the LPDA. The optical system must be designed to give the system a specific level of precision and range of possible values.

The factors of a lens or system of lenses that pertain to the basic requirements of the system are the focal length and the aperture. The focal length determines the range of distances that the system can measure. This also affects the precision with which the system can measure the road. If a fixed number of pixels are spread out by the lens over a greater range, the precision of those measurements decreases. All other factors held constant, the focal length is the determining factor on precision. The focal length should therefore be chosen such that the level of precision and the range of possible values are in balance with each other, and are in line with the design goals set forward by DAAS.

The chosen focal length is 95mm, as values between 80 and 100 provide a good balance between precision and range without becoming impractical to implement. The larger the focal length, the greater the precision of the RPS. As the focal length increases, however, the lens must be mounted further and further from the image sensor, which becomes impractical for focal lengths greater than 100 mm.

The lens that is used must have a large enough aperture given its focal length to focus all of the road heights that are in the range of the RPS. As the aperture increases, the ability of the lens to resolve light from different locations decreases, and the focus point becomes finer. At the same time, the light that is collected by the lens increases. A balance between these blurring and brightening effects was struck at 20mm. Further testing shows that the system does not need a greater aperture, but can function well at 20mm. With this lens diameter, the f-number for the lens system is 4.75.

Finally, a filter should be used with the lens system to help eliminate ambient noise. The laser has a specific wavelength that is in the infrared spectrum. An infrared longpass filter could be used to eliminate visible light and UV light while passing longer wavelengths. An alternative is the bandpass filter, which can be used to block all light except for the specific laser wavelength.

The RPS uses a simple colored glass longpass filter because it attenuates less at the laser bandwidth compared to a bandpass filter of similar quality. The majority of ambient noise due to the sun is ultraviolet or in the visible spectrum, so a longpass filter is sufficient even in sunny situations. The selected filter should have a passband limit that is at or below the laser wavelength.

The laser, image sensor, and optical components have been outlined here. In the next section, particular components are selected for use based on the preceding guidelines. The details of the design will be decided once preliminary testing has been conducted with the selected components.

C. Particular Components

The RPS major components are a laser, an LPDA, an optical system, and a mechanical mounting system. The constraints on each of these components are considered in this section, and particular components are selected.

1) LPDA: AMS TSL3301: The TSL3301 is a linear photodiode array that has 102 pixels that are spread across an 8mm length. It communicates to a central processor via a Universal Synchronous / Asynchronous Receiver / Transmitter (UART) that is capable of running at speeds up to 10 MHz. This device was chosen for the following features:

- Clock speeds: The device operates at clock frequencies up to 10 MHz. This is not as high as fast other communication protocols, but is relatively high speed and should be sufficient for the RPS.
- 102 pixels: The TSL3301 has enough pixels to provide good resolution and few enough for fast operation.
- Price: The TSL3301 is available for between 7 and 8 dollars, depending on the vendor, which is relatively inexpensive compared to other image sensors.
- Availability: LPDAs are far less common than rectangular pixel arrays, so the TSL3301 is somewhat unique.
- Light Sensitivity: The sensitivity of the TSL3301 increases with wavelength, and the chip is classified as highly sensitive. This should help in the detection of the reflected laser point.
- Efficient architecture: A single 8-bit command will output data, reset the pixels, and begin the next pixel integration cycle. On account of this, very little time is lost in the interface between the LPDA and the CPU.

The TSL3301 is a good chip for the RPS. It has built-in control over the photodiodes, and has a protocol that is easy to run in any environment on GPIO pins [8], or in a configurable UART port. It offers a good compromise between flexibility and usability, and has the appropriate resolution and cycle time to meet the requirements of DAAS. [9]

2) Optical Focusing System: The LPDA, like the image sensor in a camera, needs a lens or system of lenses to focus light onto it. The requirements for the optical system are simple: the focal length is calculated to be at least 80mm and at most 100mm, and the aperture should be close to 20mm. Based on these requirements, a simple biconvex lens was selected from Edmund Optics. The lens has a 95mm focal length and a 20.9mm diameter. It is di-convex, grade 1 lens, so it is not overly expensive. This lens was selected based on availability and price.

The lens is used to focus light onto the LDPA. It does not, however, discriminate between infrared light and visible light. That is accomplished by an optical filter. The requirements for the optical filter are that it pass 780 nm while stopping lower wavelength light. The selected filter to accomplish this is a 720 nm colored glass longpass filter from Edmund Optics. It was selected for economy, as colored glass filters are much cheaper than many other filter types, including most bandpass filters. The 720 nm filter has its passband at 750 nm, which is significantly less than 780 nm, so this filter should have no issue passing 780 nm light.

3) AIXIZ AH780-5-1230: The laser is specified to be a 5mW laser with a 780nm wavelength. Given these requirements, the Aixiz AH780-5-1230 was chosen. The selected laser module contains a laser diode, adjustable focus lens, and a driver circuit.

The only specific requirements for the laser are its wavelength and output power. Many laser diodes are available that meet these requirements; this specific laser was chosen for price and availability. The specifications for the selected laser module are:

- Output Power: 5 mW (Class IIIR)
- Input Current: 25 mA
- Input Voltage: 3.2 V
- Casing: 12mm x 30mm Can

The laser comes with an adjustable focus lens, and requires little configuration out of the box. A system, though, is needed to mount the laser and the LPDA a fixed distance apart from each other, and angled such that the laser will intersect the line of sight of the LPDA. The development of this system is briefly discussed later.

III. RPS Testing

Several tests were conducted to determine the feasibility and reliability of the RPS. Tests were conducted to determine the effect of angle, material, and surface conditions on the measured reflectance. Further tests were conducted to determine the effect of optical filters, focus, integration time, and surface conditions on received signal strength. These tests verify functionality of the RPS and provide information for the final design of the RPS.

A. Focus

If the optical system is not properly focused on the laser dot, it will not be able to resolve the dot to a single pixel on the LPDA, and the precision of the distance measurement will be affected.

To set the focus of the lens in the optical focusing system, the desired focal length was first calculated based on the lens equation. The lens equation relates the distance from a simple lens to an object and that object's image based on the focal length of the lens. The lens equation states that the inverse of the sum of the object distance and the image distance is proportional to the inverse of the focal length of the lens.

$$1/f = 1/d_o + 1/d_i$$

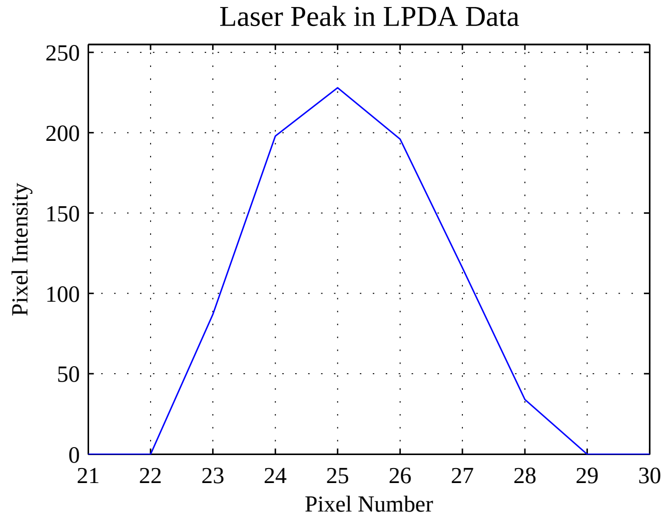

Fig. 4. Single laser peak in LPDA data.

In the RPS, a lens of focal length 95mm is used. The object is the laser point on the road ahead of the vehicle. While the road is flat, the object distance is the distance from the lens to the laser point on the road, which is set at 2.2 meters. The distance that the lens should be set away from the image sensor can be calculated from these values:

$$1/95 = 1/2200 + 1/i$$
$$1/i = 1/95 - 1/2200 = 421/41800$$
$$i = 41800/421 = 99.28mm$$

Based on this calculation, the lens was mounted 10 cm ahead of the LPDA in a screw-type adjustable housing. It was roughly focused on the laser point 2 meters ahead of the RPS. Data was collected while adjusting the lens mount 22.5 degrees at a time (1/16th of a turn). Because the laser dot is the object to be detected by the RPS, it was used as the object for the optical system to focus on. When the approximate focus was achieved, the mount was adjusted 4.5 degrees at a time until the precise focus was determined. The data from these tests is shown in figure 4.

The system was focused such that the laser point is resolved to a single pixel on the LPDA. The two adjacent pixels receive some light, but that amount is considerably less than the amount of light received by the primary pixel. A single laser peak as seen by the LPDA is shown in figure 5.

With the laser focused to a single pixel, that pixel receives the maximum light intensity possible, which enables the LPDA to function with a shorter integration time. As the integration time decreases, the frequency with which data can be sampled increases, enabling the RPS to return data to the CPU with increased frequency.

The goal of the sensor is to determine the location of the laser pulse. The driver for the RPS determines how much and how quickly the location of that laser pulse on the LPDA changes. The software must determine the location of the laser pulse on the LPDA each time the LPDA is sampled, and compare that location with the location taken from the

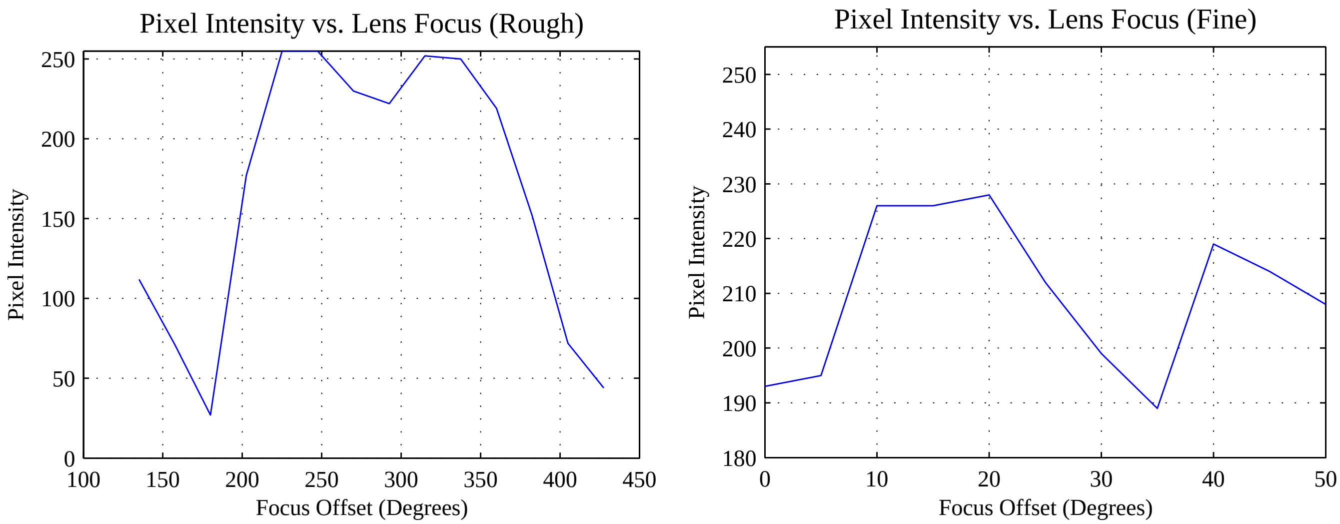

(a) (b)

Fig. 5. Focus Test Results: Figure A shows the pixel intensity as the focus is adjusted roughly. Figure B shows fine adjustment about the left peak in figure A.

previous trial. If the laser pulse is focused to a single pixel, the software architecture is greatly simplified.

If the peak were not focused to a single pixel, the software would have to determine where the center of the peak falls on the LPDA. This itself is not difficult, but testing has shown that a poorly focused laser beam often results in three independent peaks. The complexity in determining a LPDA location from three peaks is much greater compared to a single peak, so an ideal focus greatly improves RPS performance and simplifies the design. This test shows how sensitive the LPDA is to the focus of the lens. If the focus is off by a small amount, the system will not function as well. It is therefore important that a lens mounting system be devised that can position the lens at precisely the correct distance from the LPDA.

1) Material: The amount of laser light that is reflected back to the optical receiver depends in part on the material that the laser beam illuminates. The system is designed for asphalt, but various other surfaces have been used in testing. Those surfaces include both white and black paper, wood, cardboard, and asphalt. The goal of this experiment is to determine the received signal strength that corresponds to each of these surfaces. Such a comparison would determine how well the RPS will function on asphalt relative to the surfaces that have been used in testing.

In this test, the laser and LPDA were both focused on a surface 2 meters ahead of the RPS. Data was taken from the LPDA, and the section containing the laser peak was recorded. The surface was held perpendicular to the laser beam, and the integration time was held constant across all trials. Each trial consisted of 50 samples that were averaged together. The procedure was repeated with different surfaces: wood (pine, 2x4), white paper, black paper, gray cloth, and asphalt. The pixel intensity data was recorded. Figure 6 shows several pixel intensities surrounding the laser peak from each of the trials.

This test shows the reflectance of asphalt, and compares that reflectance to other materials. The maximum received signal

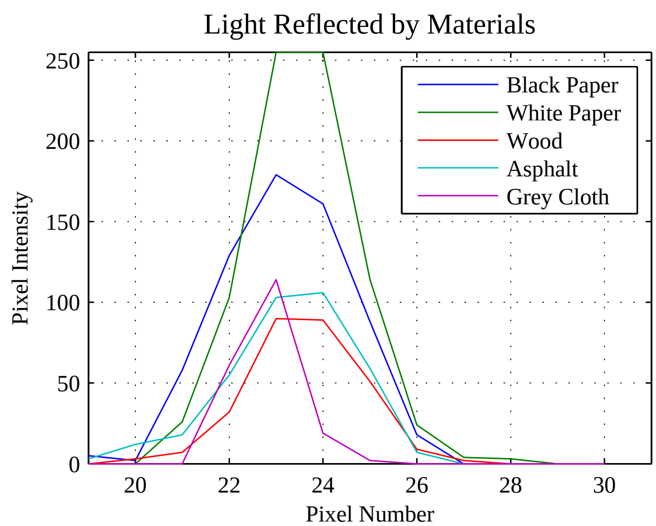

Fig. 6. RPS material reflectance testing results.

from the asphalt has a value of 106. This is high enough that the RPS should be able to detect the reflected laser light easily, especially with an increase in integration time.

The most significant comparisons are those between white paper, black paper, and asphalt. White paper has close to 100% reflectance, so it was used in many calibration phases in the design of the sensor. White paper, however, reflects so much laser light that the integration time cannot be decreased enough to keep from driving the LPDA output to its maximum, 255. Even in this test, the LPDA output goes to 255 with white paper.

Black paper was used in many tests because it was easy to work with given its smooth, planar surface, and it reflects much less light compared to white paper. On account of its lesser reflectance, it is a much better surface to use in testing. Black paper was used to determine various relationships, integration time vs. received signal strength, for example. Given that black paper was used in testing, it is important to know how its reflectance relates to that of the target surface, asphalt. It turns out that the black paper that was tested (a white sheet of paper with black laser-printer ink on it) reflects the laser much better than asphalt. The received pixel intensities are 179 and 106 respectively, making black paper 68 percent more reflective than asphalt in this test.

Wood was tested because some of the testing of the DAAS system was conducted with wooden speed bumps. This test reveals that wood reflects very similarly to asphalt in this configuration, with max pixel intensities of 90 and 106 respectively, a difference of 15 percent. Because wood has a very close reflectance (as measured by the RPS) to that of asphalt, it can be concluded that it is an appropriate material to use in the testing of the DAAS project.

This test ultimately, however, presents a very incomplete picture of the different surfaces that the RPS will encounter. Testing should be done to see how much the surface finish affects reflectance. A glossy black finish, for instance, might reflect differently than a matte black finish. A dense white sheet of paper may reflect differently than a thin white sheet of paper. This test does not take these factors into account. It rather seeks to relate different surfaces that were used in the development of the RPS for use in the DAAS project, and to determine how asphalt reflects infrared light compared to surfaces that the RPS has encountered and will encounter in testing.

2) Surface Conditions: Tests were conducted to determine the effect that water on the road has on system performance. The system is designed to be implemented on vehicles that presumably are operated in a variety of weather conditions, ranging from sunny and warm to rainy and cold. It is therefore important that the system maintain its level of performance in both wet and dry environments, as well as across a range of temperatures. This test focuses on the condition of the road ahead of the vehicle specifically in regards to water. The goal is to determine whether the presence of water significantly affects the amount of light that is reflected back to the RPS.

This test was set up with the laser and the LPDA focused on a tile floor 2 meters ahead of the RPS. Tests were conducted with a dry floor, a wet floor, and 3 mm of standing water on the floor. The entire LPDA was sampled, and the section that detected the laser peak was recorded. Figure 7 shows the values of the pixels that detected the laser pulse for the three tests.

The dry surface reflected more than the wet surface, with a peak pixel intensity of 189 compared to 182 (a difference of 3.7 percent). The surface that was covered with standing water reflected a bit less light than the wet surface, with a peak pixel intensity of 177 (a difference of 2.75 percent). These peak values, though, are relatively close together. It is difficult to determine whether or not the difference in pixel intensity is related to these conditions because the change in pixel intensity was so small.

This test shows that the RPS can indeed function in wet environments. Wet surfaces and standing water both reflect slightly less than dry surfaces, but the difference is small enough that the design of the RPS does not need to be altered to account for wet environments.

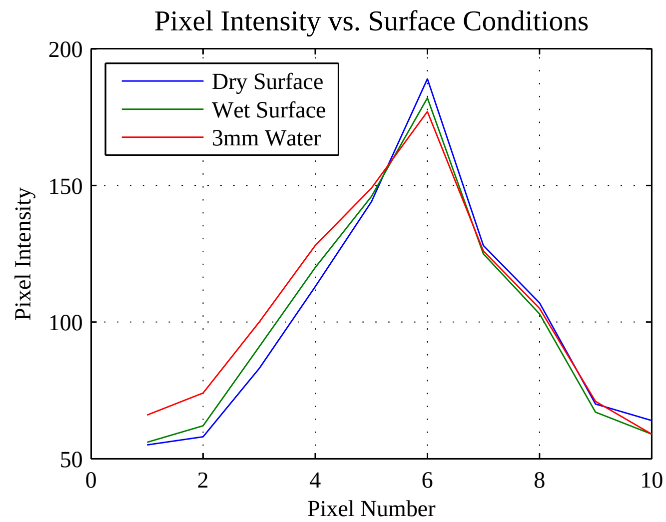

Fig. 7. RPS water testing results.

3) Optical Filters: Without an optical filter, there is no mechanism to keep the LPDA from detecting ambient light. The LPDA does not distinguish between wavelengths, so it would be likely that in high-noise environments, the reflected laser light could be indistinguishable from ambient light. This would especially be a problem when the system operates in sunlight, and has the potential to pose problems in other environments. A longpass filter is used to block ambient light that is of a lower wavelength than the laser wavelength. This test looks at the difference in signal to noise ratio with and with a longpass filter, a bandpass filter, and no filter.

Tests were conducted with each filter to determine the benefit that the filter yields. The laser and the LPDA were both focused on a black sheet of paper 2 meters in front of the RPS in a controlled lab environment, and data was taken from the LPDA for each trial.

The major difference between the longpass and bandpass filters is the level of attenuation at the laser wavelength. The longpass filter attenuates the passband very little, whereas the bandpass filter attenuates the passband significantly, requiring longer integration times. During indoor testing, the longpass filter had very little effect on the SNR compared to tests without a filter at all. This is because the laser point is bright in contrast to its surroundings. The integration time of the LPDA has therefore been decreased until only the laser pulse is detected, even in the absence of a filter. It therefore is clear that no filter is necessary for moderately-lit environments. In bright environments, like sunlight, however, a longpass filter is still necessary. A longpass filter is therefore included in the design of the RPS.

4) Angle: The laser beam in the RPS design makes an angle with the plane of the road. This angle must be calculated to maximize horizontal distance, as well as other factors. As this angle changes, the amount of light that is reflected back to the optical sensor changes. As a rise or fall is encountered

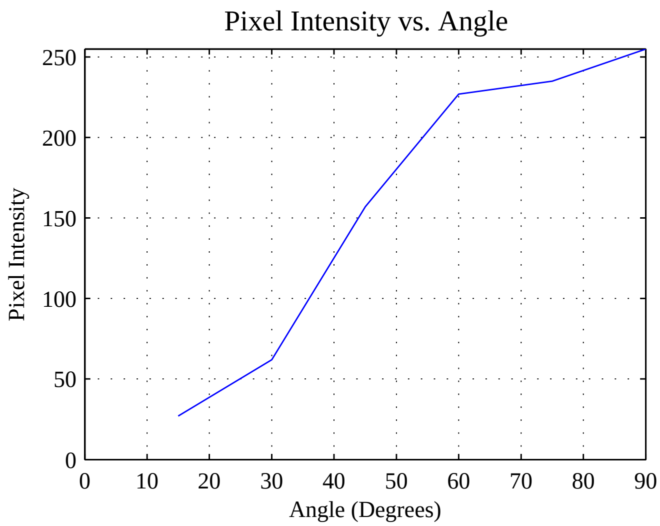

Fig. 8. RPS angle testing results.

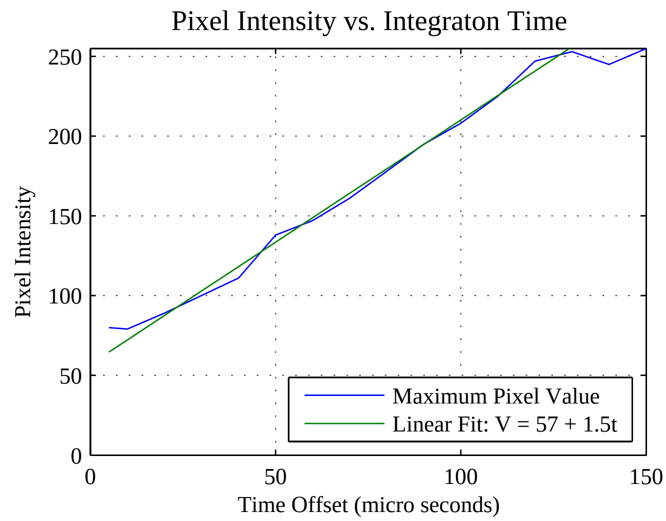

Fig. 9. RPS integration time results.

by the vehicle, the angle between the laser beam and the surface it illuminates changes. On account of these factors, it is important that the RPS be tested with a changing angle.

The test was set up with the LPDA and the laser focused on a flat piece of cardboard two meters ahead of the RPS. The cardboard was rotated such that the angle changed from 90 degrees down to 15 degrees relative to the laser while the distance was held constant. The LPDA was sampled for each trial, and the maximum value across the array, the peak value from the laser dot, was recorded. The maximum pixel intensity is shown in figure 8 with respect to the angle for each trial.

From this test we see that the received signal strength changes very little between perpendicular and 60 degrees, with pixel intensities changing from 255 to 227. After 60 degrees, the received signal strength decreases much more quickly, dropping down to 27 at 15 degrees. The data from this test shows that it is advantageous to maximize the angle between the laser and the road. Ideally, that angle would be close to 60 degrees.

The problem is that as the angle increases, either the sensor mounting height must increase or the lead distance (the distance from the front of the vehicle to the point on the road that the RPS samples) must decrease. As the lead distance decreases, the lead time (the time between bump detection and bump impact) decreases as well. The RPS must therefore compromise between greater reflected laser light and greater lead time for the DAAS system.

The smallest angle that was tested is 15 degrees, and the maximum pixel intensity at that angle is 27. Below 15 degrees, it became infeasible to continue sampling the RPS. Given the decrease in reflectance at very close angles, it is very difficult to detect downward slopes in the road. The greatest theoretical downward slope that can be detected will be almost parallel to the laser beam. This, of course, will be difficult to detect, as the angle will be approaching zero degrees. In the case that the downward slope increases to the point that it can no longer be reached by the laser beam, the LPDA will see a discontinuity in the road height, and the software can conclude that a dip in the road has been detected.

This test shows how sensitive the RPS is to angles. From this test, we conclude that the angle between the road and the laser should be maximized, though that decision is based on many other factors as well. Finally, this test shows that the RPS will likely not be able to detect falling edges well at all, and will rather show them as a discontinuity in the data that it returns.

5) Integration Time: The LPDA functions by allowing light to illuminate the sensor for configurable amount of time. As light contacts the sensor, current is generated via photodiodes in the array. This current is integrated by a circuit for each photodiode, resulting in an accumulated charge specific to each photodiode. The amount of accumulated charge corresponds to the value that is returned by the LPDA. As either the time or the intensity of light increase, the amount of accumulated charge, and therefore the returned value, increases. This returned value will be referred to as pixel intensity.

This test determines the relationship between pixel intensity and the integration time. In the use of the RPS, various conditions and environments are expected that will affect the amount of light that is reflected back toward the optical sensor. The integration time must then be adjusted to compensate for the change in the amount of light that is reflected. In order to most effectively adjust the integration time, it is important to know the relationship between integration time and pixel intensity.

In this test, data was taken from the RPS with the laser and the LPDA focused on a black sheet of paper two meters away. The paper was aligned perpendicular to the laser. The integration time was adjusted by putting the processor to sleep for a configurable amount of time between the begin integration and the stop integration commands that are specific to the TSL3301. Data was collected from the entire array, and the maximum pixel intensity from the array was stored for each trial. The results of this experiment are shown in figure 9.

The relationship is close to linear in the middle section of the graph. Toward extreme values of integration time, the relationship becomes nonlinear. As the delay time approaches 150 microseconds, the maximum pixel intensity goes to 255, the maximum value of the LPDA, which presumably effects the linearity of the data at the upper end of the tested sleep times. As the delay time approaches zero, the integration time is effected by the code execution time. This is very brief, of course, but it could affect the linearity of the data at the lower end of the tested sleep times.

This test has demonstrated two main points: first, the LPDA behaves linearly in respect to integration time. Second, the LPDA is very sensitive to integration time. A single microsecond will affect the pixel intensity. This data, of course, is very dependent on several conditions: the material that is reflecting the laser, and the angle that the laser makes with the reflecting surface. These are two conditions that affect the relationship between time and value, and these parameters will change in implementation. The linear nature of the LPDA response is not expected to change, but the ratio between pixel intensity and time is expected to change with various factors that will be different during the integration of the RPS into DAAS.

IV. Final Design Considerations

The final design parameters were decided on after the testing phase and are described in this section. These parameters are the height that the RPS is mounted above the road, the distance ahead of the vehicle that the RPS scans, and the magnitude of the separation between the laser and the LPDA. These parameters determine the geometry of the triangle that is used in calculating road height. Based on these parameters, a final mounting system was designed to position the laser and the LPDA in an appropriate orientation to each other.

A. Height Above the Road

The height that the system is mounted above the road affects the angle with which the laser hits the road, and therefore the amount of light that is reflected back to the optical receiver.

The angle that the laser makes with the axis of the lens decreases as the height above the road increases if other parameters, like the distance of the laser point ahead of the vehicle, remain constant. This increases the distance over which the laser crosses paths with the field of view of the lens. As this distance increases, the range of heights that can be measured increases, but the precision with which those values can be measured decreases. In order to maintain resolution as the height increases, the separation of the laser and the LPDA must also increase.

A limitation imposed by DAAS is that the RPS must, to some extent/ be configured to work with the donor vehicle DAAS is implemented on. On account of this limitation, the maximum height is set to 1 meter. The height could be less than 1 meter, but this would decrease the angle between the laser and the road, which has been shown to have negative effect on the reflected signal strength in the testing section. On account of this, the mounting height should be maximized.

This being said, 1 meter of road height turns out to be a good height for the RPS in the context of DAAS. First of all, mounting the system at heights greater than 1 meter is difficult to achieve without disregarding practicality. Second, 1 meter of road height requires less separation between the laser and the LPDA. Finally, at 1 meter of road height, the lead distance can extend out to about 2 meters before becoming impractical, which will be discussed in the following section.

B. Distance Ahead of the Vehicle

The distance from the chassis to the laser point on the road affects many different factors. As the point moves further from the vehicle, the lead time increases, but it does so at the cost of angle. The issue is that the mechanical actuation system has a minimum lead time, and this corresponds to a minimum distance from the vehicle. Additional time would give the electrical control system more data to use in calculations, which would result in greater accuracy in control of the mechanical system. The goal, therefore, is to maximize the lead distance while maintaining functionality of the RPS. As the focus point moves further from the vehicle, the following issues arise:

- The angle that the laser makes with the road as well as the amount of light from the laser that is reflected back to the optical receiver decrease.
- The focal length of the lens in the chosen optics must increase proportionally to the distance in order to maintain precision. More complex lens systems are needed for greater focal lengths.

The largest problem is the decrease in reflected light back to the optical detector. On account of this, a distance of 2 meters has been chosen. Two meters gives DAAS .2 seconds to adapt at twenty miles per hour, which is longer than the minimum lead time for DAAS. With a height of one meter, as determined by the structure of the donor vehicle, a distance of two meters makes an angle of 22.5 degrees between the road and the laser. Much less than 22.5 degrees would result in very poor laser reflection, as determined by angle testing that has been conducted. On account of the decreasing angle and limited height, two meters was chosen as the lead distance for the RPS rather than 3 or 4 meters.

C. Distance Between Laser and Optical Receiver

The distance between the laser source and the LPDA determines the range of distance that can be measured by the RPS. As the distance increases, the resolution of the LPDA increases, and the range of values that can be measured decreases. It is important that the separation not be too small, lest the RPS be geared toward absurdly large bumps. The resolution shouldn't be so fine either such that the system can only detect small bumps.

On account of this, the separation has been chosen to be 25 cm. While greater separation would be beneficial in some ways, creating a rigid mounting for the laser and the LPDA such that they remain stationary relative to each other becomes more and more difficult as the distance increases. Also, if the separation were to exceed 25cm by much, the RPSs would no longer integrate well with DAAS simply because they would be too large to be mounted on the DAAS vehicle. It is therefore on account of many interwoven dependencies that the separation of the laser and the LPDA is set at 25 cm.

D. RPS Mounting System

The mounting system for the RPS was designed and implemented by mechanical engineering student Everan Chaffee from project DAAS. Chaffee designed the system under the direction of Matthew Edel. The mounting system is designed with a floating laser mount that is calibrated with set screws, and maintains calibration between the LPDA and the laser even in the violent environment of an automobile. The system is designed with the following constraints:

- LPDA should be mounted with its primary axis parallel to the laser beam.
- The angle that the laser makes with the LPDA should be adjustable.
- The LPDA and the laser should be mounted 25 cm apart.
- The lens should be mounted 99.8 mm above the LPDA.
- The longpass filter should be mounted above the lens.
- The system should be rigid.

With these guidelines in place, Chaffee developed a system that was machined out of a one inch square steel bar and 1.5 inch steel pipe that was on stock in SPU's machine shop. Chaffee milled a slot in the bar stock and mounted the laser in that slot with set screws for angle adjustment. Chaffee mounted 8cm of pipe perpendicular to the bar stock over the LPDA. The pipe would hold the lens system directly in front of the LPDA. The lens system was mounted on screw adjustable PVC conduit for fine adjustment, and the PVC conduit was mounted in the steel pipe. The system is illustrated in figure 10.

This design provided fine adjustment of the laser orientation, as well as fine adjustment of the lens focus. Once the laser is adjusted properly, it is held rigidly in relation to the LPDA. This is necessary to keep the laser from drifting out of the field of view of the LPDA. Other designs were considered that optimized adjustibility of the laser orientation and the LPDA orientation, but they were ultimately abandoned because they did not preserve rigidity as well.

With the RPS constructed with this set up, it is ready to be evaluated as a system, and ultimately integrated with project DAAS.

V. Performance Measurements

As a completed system, the RPS was evaluated in two ways: first the RPS was bench tested as a distance measurement device. The bench testing sought to determine the resolution of the RPS as a distance measurement device, as well as how long it takes the RPS to sample data. Next, the RPS was integrated with the DAAS project, and data was taken while a bump was simulated. These tests help determine how well the RPS functions as both a standalone system and as an integral part to an active suspension system. Some additional test data can be found in the report on the DAAS project.

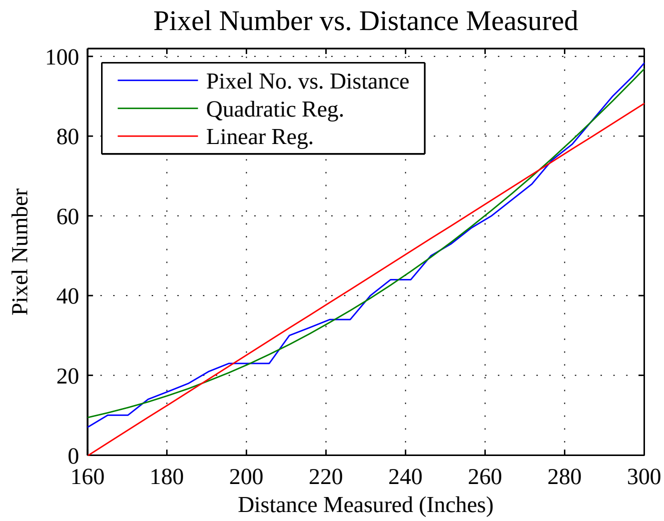

Fig. 11. RPS Final distance testing data.

A. Distance Measurement

The RPS was bench tested as a distance sensor to verify its basic functionality prior to mounting it to the DAAS vehicle. This test is valuable in determining whether the pixel number is linearly related to distance or not. The transfer function that is derived is never used, as the pixel number is the only data sent from the RPS to the DAAS system. In relation to this project, this test is a verification step to test functionality before the system is integrated with DAAS.

In this test, the RPS was positioned horizontally on a table. A white sheet of paper was placed vertically in front of the RPS such that the laser struck the paper. The paper was positioned close to the RPS such that the laser dot was at the extreme end of the field of view of the LPDA. The paper was then moved away from the RPS 5 cm at a time while data was being sampled and stored. This process was repeated until the reflected laser dot had crossed the entire LPDA and was therefore no longer in the field of view of the LPDA.

The data from this test was analyzed, and for each trial the number of the pixel that had the highest value was recorded. These pixel numbers are shown in figure 11 in relation to the distances they correspond to. Quadratic and linear curves were both fit to the data, as is shown in the figure.

The relationship between pixel number and distance is clearly nonlinear according to this test. The data does wander from the quadratic curve some, but it generally follows the same contour as the quadratic approximation. This test shows that the RPS does indeed function as a distance sensor, and it can therefore be implemented as a road sensor in the DAAS system.

In this test, the RPS measures distances spanning 140 cm, and it does so using a total of 94 pixels. The resolution is therefore 1.5 cm. When the RPS is mounted at 22.5 degrees relative to the road, the 140 cm of range become 35 cm of vertical measurement. The precision of that measurement becomes 3.75 mm. This data is validated in the next section when the system is integrated onto the DAAS vehicle.

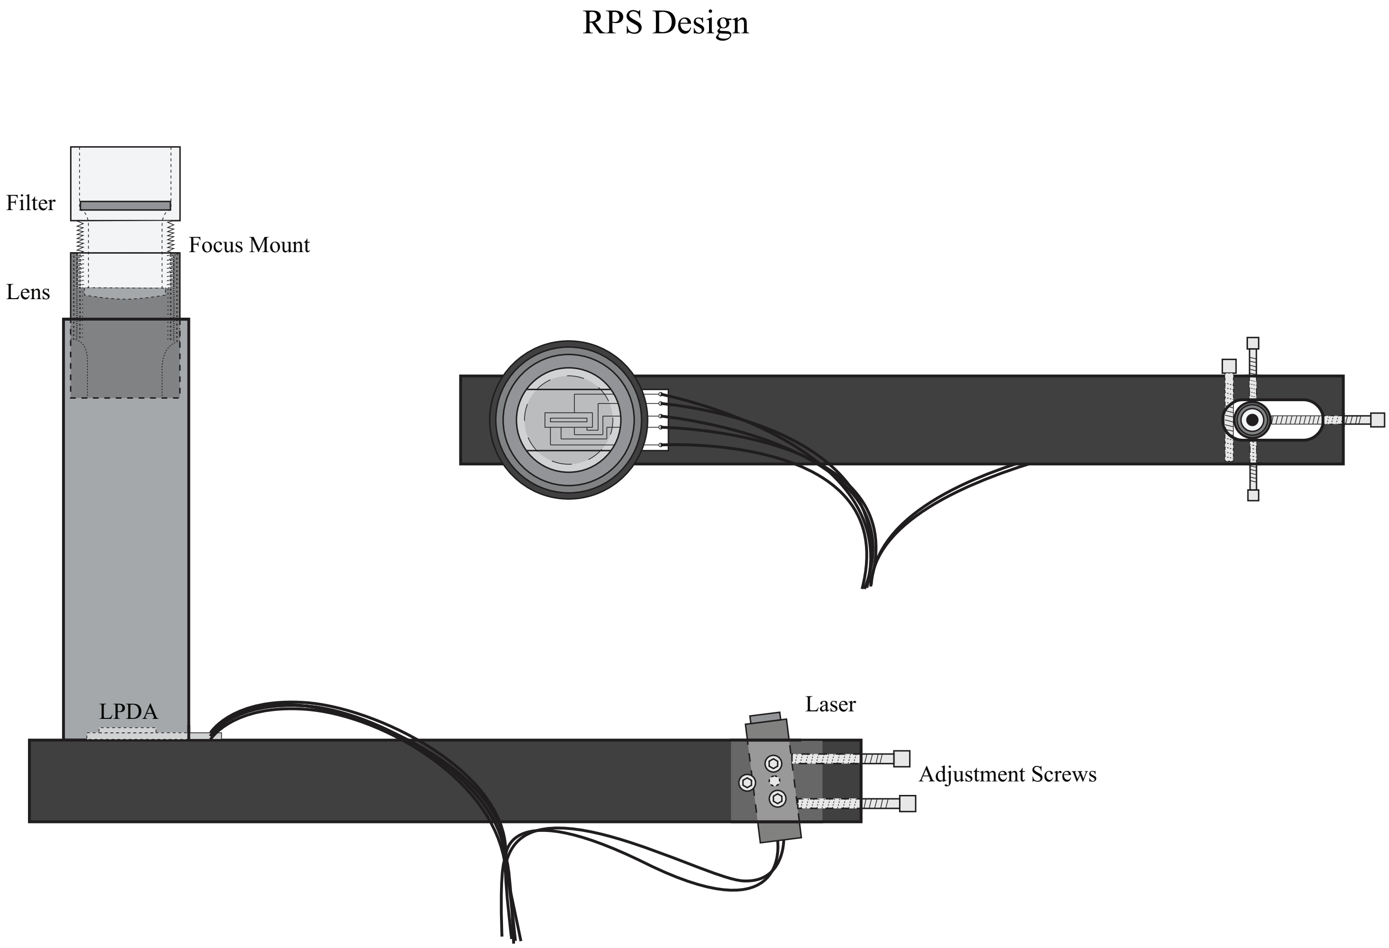

Fig. 10. RPS diagram.

B. Bump Detection

The RPS was finally integrated with DAAS. It was mounted at the front of the vehicle, slightly inside of the front wheel on the left side of the vehicle. Because the RPS was not in line with the front wheel, it was oriented toward the outside edge of the vehicle so that the height of the road would be measured at a point directly ahead of the front wheel. The sensor was oriented toward the ground 2 meters ahead of the vehicle, and approximately 2.3 meters ahead of the RPS due to the height at which the RPS was mounted (figure 13). The RPS was mounted with the LPDA positioned above the laser, as the received signal strength was found to be greater in this configuration.

With the RPS mounted on the vehicle, a bump was simulated by moving a two-by-four longitudinally toward the vehicle on the left side. Data was collected from the RPS each time the two-by-four was moved closer to the vehicle. The two-by-four was moved in quarter-inch increments, and required 22.25 cm to fully pass through the view of the RPS.

The data from this test was analyzed by comparing the number of the pixel with the greatest value and the location of the two-by-four relative to the vehicle. Figure 12 shows the pixel number as it corresponds to the two-by-four location as measured in centimeters from the vehicle.

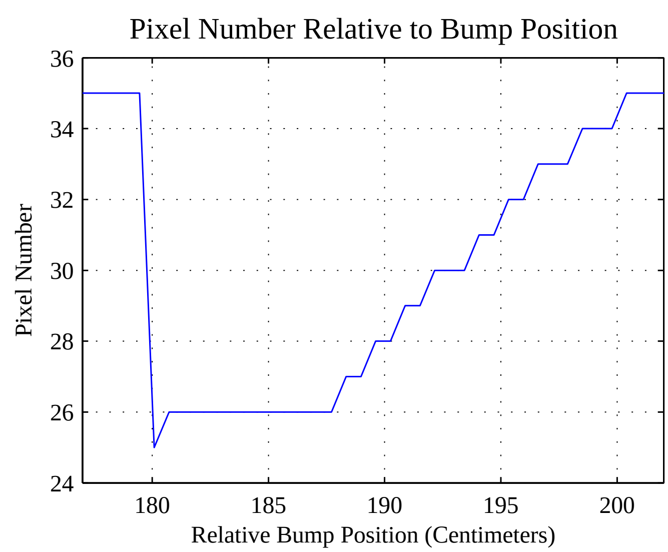

Fig. 12. Simulated bump data.

As the two-by-four moves into the path of the laser beam (from right to left in figure 12), the pixel number decreases. Once the laser beam begins to reflect off of the top side of the two-by-four, the pixel number stops changing. When the laser

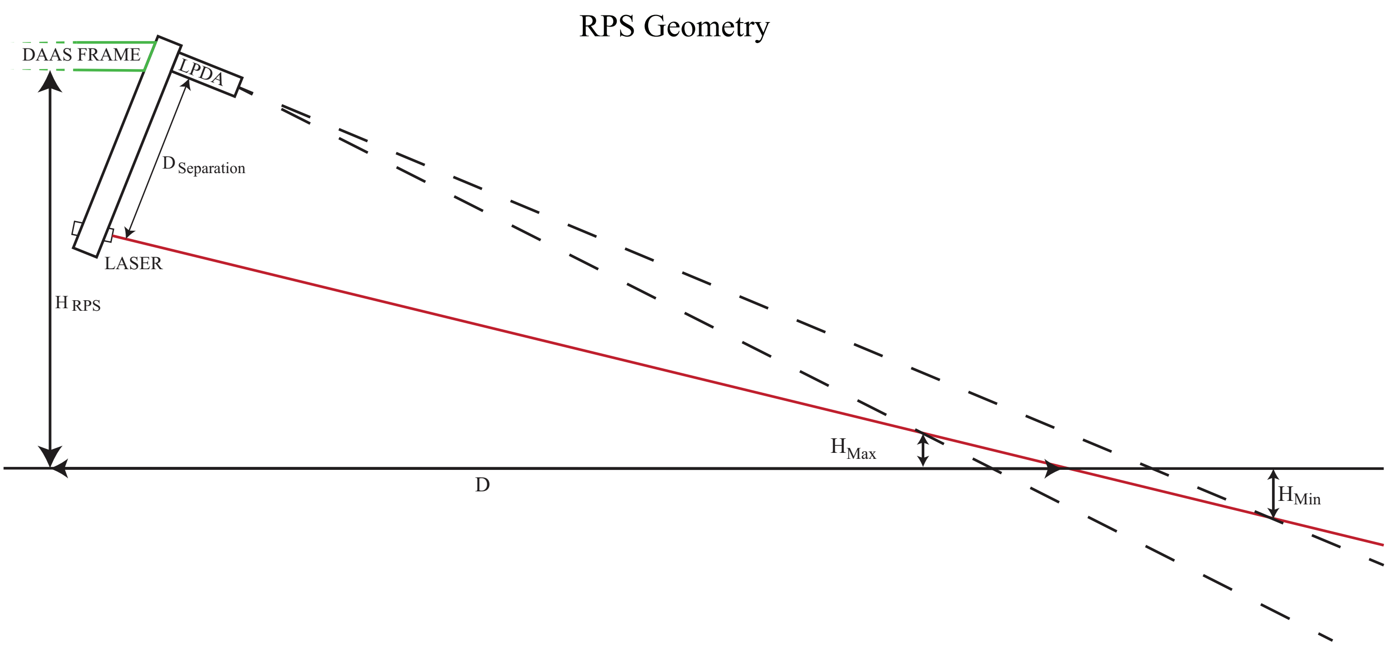

Fig. 13. RPS mounting relative to the road.

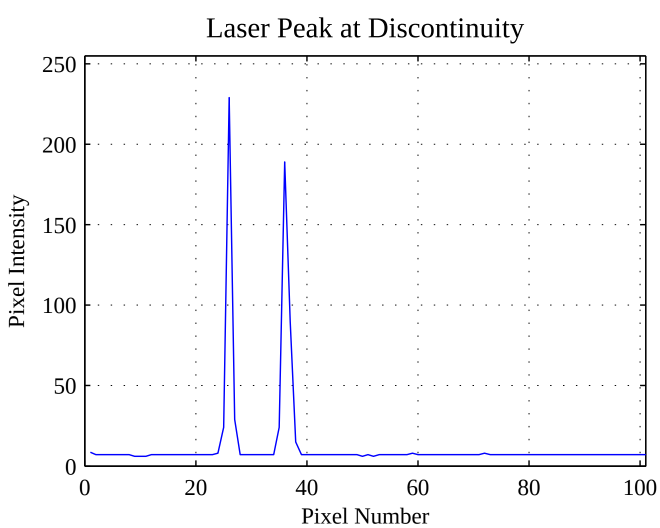

Fig. 14. Discontinuity in negative slopes.

beam falls off of the trailing edge of the two-by-four, the pixel number very quickly returns to its original value (the change in pixel number at position 63.5.)

The laser makes an angle of 22.5 degrees with the road. Because of this relatively small angle, the laser beam is stretched out in a line where it contacts the ground. In the trial with the two-by-four, the laser beam actually is split in half when it falls off of the trailing edge of the two-by-four such that a discontinuity can be seen in a single frame from the RPS. Figure 13 shows the pixel intensities on the LPDA as the laser is falling off of the edge of the two-by-four.

In this figure, the laser peak is clearly split into two distinct peaks. The higher peak is contacting the two-by-four, while the lower peak is contacting the ground behind the two-by-four. The difference between the location of the two peaks is the height of the two-by-four as it is measured by the RPS. This frame shows how the RPS detects potholes in the road; it does so by observing a discontinuity in the data. That discontinuity may not occur in a single sample, it may occur across two or more samples, but effectively, the RPS will see the data drop off suddenly when a downward slope is encountered that is greater than 22.5 degrees. This data confirms that the RPS will be able to detect potholes in addition to bumps.

The precision with which the RPS detects road irregularities is determined by the change in pixel number for a given bump height. In this test, the RPS was given a bump with a height of 3.8 cm, and the pixel number changed by 9. From this data we conclude that the bump height resolution of the RPS as it is integrated into DAAS is 4mm.

At twenty miles per hour and fifty frames per second, the RPS will be able to generate a graph with 17.3 cm spread between data points. The RPS detected the two-by-four over the course of about 22 cm, so it would likely still function to some extent at 20 miles per hour, but it certainly would not be able to detect the two-by-four with more than a single frame of resolution. At ten miles per hour, this distance decreases to 8.6 centimeters, which is more reasonable, though it still isn't ideal. Both of these numbers will struggle to detect small bumps very accurately, though they will theoretically still be able to detect a two-by-four sized bump.

With the current software driver, data can be sampled from the RPS every 16 ms, or 62 times per second. This would decrease the space between samples from 17.3 to 14.2 cm if the vehicle is traveling at 20 miles per hour. This is a significant improvement over the 50HZ cycle rate used by DAAS, though it could still be improved upon with further software driver development. The sample rate meets the specifications set forth by DAAS.

VI. Recommendations and Concluding Remarks

The Road Profile Sensor is a device that scans the road ahead of a vehicle, and is intended for use with active suspension systems. The RPS was developed under constraints set forth by a senior design project at SPU, Dynamic Automatic Adjusting Suspension, that sought to implement an active suspension system. The RPS met all of the specifications laid out for it by project DAAS, exceeding most of them greatly. In its final form, the RPS scans the road two meters ahead of the vehicle, it detects road height with a precision of 4 mm, and the maximum sample rate is 62 Hz.

The RPS in its current state has much potential for development. It is currently limited by the speed of the GPIO port on the Altera Cyclone IV development board, and so can only be sampled every 16 ms. The LPDA used in the design of the RPS is specified be able to operate with a 10 MHz clock, which would give the RPS a theoretical sample rate of 10 kHz. The limitation with the current implementation is that the communication port is bound to the GPIO pins. It is feasible to configure an SPI port on the Cyclone IV to work with the TSL3301 at a much higher clock frequency; it just hasn't been done yet. Such high sample rates from a chip like this would give it a major advantage over other road scanning devices currently being developed. The RPS would be able to model bumps at freeway speeds with the same resolution that is currently only attainable in the laboratory.

The RPS relies on very little processing power. The current driver takes the data from the serial port and puts it into an array of 102 bytes. The array is then processed in a loop that determines which byte has the greatest value. The number of the byte with the greatest value is returned to the CPU. This process has the potential to execute much faster while still consuming very few resources. Once the serial port is removed from the CPU, the communications will be able to run much faster. The CPU then would simply have to read data from the serial port instead of running the port it itself. The CPU would read a byte at a time and send each byte directly to a pipeline in the FPGA. The FPGA pipeline would then determine whether the laser peak changed pixels.

An implementation along these lines would not only run on few resources; it would run faster and leave the CPU cores open for other calculations. The system could potentially run on a smaller chip than the Cyclone IV, and so would be economical to manufacture while still maintaining a higher level of performance.

The sole remaining issue with the hardware is that the calibration of the laser and the LPDA is somewhat unstable. The LPDA detects a portion of the road that is 35 centimeters high but only 1 centimeter wide. If the laser source shifts at all from side to side within its mount, the laser point will likely drift out of the field of view of the LPDA, and the LPDA will not be able to detect it. If the laser shifts, then, the RPS's will stop functioning altogether.

Possible solutions have been considered to mediate this problem. A laser could be used that is focused to a line rather than a point, and that line could be oriented transversely with respect to the LPDA. This is likely the best solution; the problem is that a 5mW laser focused across a line will reflect much less light to the LPDA than a 5mW laser focused to a point. Another solution is to use a rectangular pixel array. The problem with a rectangular array is that it will require much more processing than the LPDA. The LPDA was chosen for its simplicity and elegance, and a rectangular array abandons that simplicity. The final solution is to mount the laser and the LPDA in such a rigid fashion that they simply cannot fall out of calibration.

This last approach has been tried with the RPS. It is not clear at this point how susceptible the current RPS design is to losing calibration. The problem could have been solved, but the issue cannot be considered resolved until further testing is done with the current RPS unit. The current design is thought to be much more rigid than previous designs, but the extent to which it will maintain calibration is not yet known.

The RPS in its current state is greatly limited by the driver that processes data from it. The RPS itself is capable of very fast, precise operation. The driver and the communication port should be redesigned for faster performance. This being said, the RPS has already met the specifications that were set for it. The RPS is somewhat prone to lose calibration between the laser and the LPDA, but this issue is thought to have been resolved. Further testing is expected to confirm this.

The RPS is currently being integrated with the DAAS project, and the DAAS team will begin testing with it in the coming weeks. The RPS is considered at this point to be a stable subsystem in the DAAS project.

M.Edel
Mat 23, 2014

References

[1] J. Meiners. (2010) Driven: Mercedes-benz magic body control. [Online]. Available: http://www.caranddriver.com/news/mercedes-benz-news-magic-body-control-system-driven

[2] J. Barnett. (2013) Porsche active suspension management explained. [Online]. Available: http://www.total911.com/porsche-active-suspension-management-explained/

[3] M. Edel, "Daas: by team etc, cornell cup," Poster, 2014.

[4] ——, "Dynamic automatic adjusting suspension," 2014, design project at Seattle Pacific University.

[5] R. L. Industries. (2000) Laser standards and classifications. [Online]. Available: http://www.rli.com/resources/articles/classification.aspx

[6] D. A. J. Marsh. (2011) Solar radiation. [Online]. Available: http://naturalfrequency.com/wiki/solar-radiation

[7] H. Choi. (2003) Advantages of photodiode array. [Online]. Available: www.hwe.oita-u.ac.$jp/kiki/ronnbunn/paper_choi.pdf$

[8] P. Pilgrim. (2001) Intelligent opto sensor designer's notebook: Using the tsl3301 with a microcontroller: Controlling the tsl3301 in the absence of a continuous clock. [Online]. Available: https://www.ams.com/eng/content/download/251600/993363/version/1

[9] A. M. S. AG. (2011) Tsl3301cl 102 1 linear optical sensor array with analog-to-digital converter. [Online]. Available: http://www.ams.com/eng/content/download/250205/975793/142526

Appendix
Faith Integration Statement

The development of the RPS is a very technically focused project. In this, it is easy to see how the project contributed to my development as a student of engineering. It is somewhat less clear how the project ties into my faith. The connection is there, but it is much more subtle than the projects role in my technical education. The connection, though, is strong; it is in fact the driving force behind the project.

The integration of my faith with my honors project is less causal, yet more driving. Perhaps it is just in my personality; perhaps it is faith. The distinction is not clear to me. My perspective, though, is that I am called by God to use my abilities to the limits of their potential. I am very skilled at engineering, and for me to not apply myself to the best of my abilities is wrong. American distance runner Steve Prefontaine put it simply: ”To give anything less than your best is to sacrifice the gift.” It is for this reason that I feel called to push myself academically. The honors project was an opportunity for me to put my skills to the test. It was an opportunity to develop those skills under mentorship from two very inspirational professors. Were I to squander that opportunity, or even walk away from the commitment that I made to the University Scholars program, I would feel that I had missed God's primary calling on me as a student.

I am not trying to say that I feel academically accountable to God. I do feel that I am called to push myself, though. A verse in the Bible states, ”Whatever you do, do your work heartily, as for the Lord rather than for men” (Corinthians 2:23). If I am to take the Bible seriously as a Christian, this verse has the power to terrify me. It would be simple to approach the honors project as an assignment to be completed for a grade, but to me that approach does not seem appropriate for God. If God is perfectly good, should I not aspire to make my honors project as flawless as I can? Shall I not also push myself to achieve something that is great, rather than something that is simple? The theme here is effort rather than perfection. At the end of the day, I am willing to quit so long as I feel I have registered my best attempt.

I understand that as humans we are not perfect, nor are we called to be perfect. We are called to live under the grace of God. It is on account of this that I can be proud of my work. I am proud of it because I feel that it represents the best of my work given the circumstances surrounding its development. The lamb that Able offered to the Lord in the beginning of Genesis was not flawless, but it was the best that he had to offer. In the same way, my project must be a work that I am proud of; this is clear to me.

In the beginning of the Bible, God tells Adam and Eve, ”fill the earth, and subdue it” (Genesis 1:28). As an engineer, my modern interpretation of the word subdue in this context is to design the automobile. I am no theologian, and I do think that we need to keep from destroying the planet that we live on. I do think, though, that God intends for us to excel in science, and to ultimately use the material world to pursue the greater good of humanity. Technology has the potential to save lives and ease suffering. If we stop pursuing science and its application in engineering, I think that we are in violation of God's intention for us. I am gifted in engineering. I therefore think that it is God's calling for me to pursue excellence in engineering. For me to walk away from engineering would be to walk away from the gifts that I have. On account of this, I think that I am not only called to push myself to the limits of my potential, but I am also called to push that potential further. I ultimately am coming to realize that God has given me gifts, and for me to walk away from these gifts is absurd. For me to simply use them to my benefit isn't the answer either. It is rather my calling to develop those gifts as I journey through this world and to use them to honor God. This, I think, is my highest calling.

ACKNOWLEDGMENT

The author would like to thank...

Advisors Dr. John Lindberg, Dr. Kevin Bolding & Thomas Edel, Team ETC, first mate John Metcalf, and fiancé Dani.

Matthew Edel is an undergraduate student at Seattle Pacific University studying Electrical Engineering and Computer Engineering.