Enhancement and Field Test Evaluation of New Battery-Less Wireless Traffic Sensors

Final Report

Prepared by:

Sean Pruden
Krishna Vijayaraghavan
Rajesh Rajamani

Department of Mechanical Engineering
University of Minnesota

CTS 11-22

Technical Report Documentation Page

1. Report No.	2.	3. Recipients Accession No.	
CTS 11-22			
4. Title and Subtitle Enhancement and Field Test Evaluation of New Battery-Less Wireless Traffic Sensors		5. Report Date October 2011 6.	
7. Author(s) Sean Pruden, Krishna Vijayaraghavan and Rajesh Rajamani		8. Performing Organization Report No.	
9. Performing Organization Name and Address Department of Mechanical Engineering University of Minnesota 111 Church Street, SE Minneapolis, MN 55455		10. Project/Task/Work Unit No. CTS Project #2009028 11. Contract (C) or Grant (G) No.	
12. Sponsoring Organization Name and Address Intelligent Transportation Systems Institute Center for Transportation Studies University of Minnesota 200 Transportation and Safety Building 511 Washington Ave. SE Minneapolis, MN 55455		13. Type of Report and Period Covered Final Report 14. Sponsoring Agency Code	
15. Supplementary Notes http://www.its.umn.edu/Publications/ResearchReports/			

16. Abstract (Limit: 250 words)

This project focused on the enhancement of a previous battery-less wireless traffic flow sensor so as to enable it to provide weigh-in-motion (WIM) measurements and provide enhanced telemetry distance. The sensor consists of a 6-feet-long device which is embedded in a slot in the road flush with the pavement. As a vehicle travels over the sensor, vibrations are induced in the sensor. Using piezoelectric elements, energy is harvested from the vibrations and used to power the electronics in the sensor for signal measurements and wireless transmission.

The sensor's performance was evaluated by embedding it in a slot in concrete pavement and driving various vehicles of known weight over it at a number of different speeds on different days. The sensor was found to meet the specification of 500 feet telemetry distance. It was able to provide WIM measurements with an accuracy of better than ±15% in the absence of vehicle suspension vibrations.

However, much of the WIM data during the latter period of sensor testing was obtained in the presence of significant suspension vibrations. The project also evaluated the use of 4 consecutive WIM sensors in the road to remove the influence of suspension vibrations.

17. Document Analysis/Descriptors	18. Availability Statement		
Sensors, Weigh in motion scales, Energy transfer	No restrictions. Document available from: National Technical Information Services, Alexandria, Virginia 22312		
19. Security Class (this report)	20. Security Class (this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	110	

Enhancement and Field Test Evaluation of New Battery-Less Wireless Traffic Sensors

Final Report

Prepared by:

Sean Pruden
Krishna Vijayaraghavan
Rajesh Rajamani

Department of Mechanical Engineering
University of Minnesota

October 2011

Published by:

Intelligent Transportation Systems Institute
Center for Transportation Studies
University of Minnesota
200 Transportation and Safety Building
511 Washington Avenue S.E.
Minneapolis, Minnesota 55455

ACKNOWLEDGMENTS

The authors wish to acknowledge those who made this research possible. The study was funded by the Intelligent Transportation Systems (ITS) Institute, a program of the University of Minnesota's Center for Transportation Studies (CTS). Financial support was provided by the United States Department of Transportation's Research and Innovative Technologies Administration (RITA).

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

This project focused on the enhancement of a previous battery-less wireless traffic flow sensor so as to enable it to provide weigh-in-motion (WIM) measurements and provide enhanced telemetry distance. The sensor consists of a 6-feet-long, largely one-dimensional structure, which is embedded in a slot in the road flush with the pavement. As a vehicle travels over the sensor, vibrations are induced in the sensor. Using piezoelectric elements, energy is harvested from the vibrations and used to power the electronics in the sensor for signal measurements and wireless transmission.

The new sensor design utilized a 3-layer structure, consisting of a top frame, a WIM layer and an energy-harvesting layer. The energy-harvesting layer was designed to have a nonlinear structural response so that the sensor could harvest adequate energy from low vehicle loads, while at the same time not failing structurally for heavy vehicle loads. The WIM layer was designed to have an elastic structural response, so that voltage measurements on the WIM layer could provide a measure of vehicle weight.

The sensor's performance was evaluated by embedding it in a slot in concrete pavement and driving various vehicles of known weight over it at a number of different speeds on different days. The sensor was found to meet the specification of 500 feet telemetry distance. It was able to provide WIM measurements with an accuracy of better than ±15% in the absence of vehicle suspension vibrations.

However, much of the WIM data during the latter period of sensor testing was obtained in the presence of significant suspension vibrations. This is because the sensor could not be maintained flush with the road surface. Hence the passing of any vehicle over the sensor led to suspension vibrations.

The project evaluated the use of 4 consecutive WIM sensors in the road to remove the influence of suspension vibrations. Three different algorithms (average of the 4 sensors, least squares weight calculation and least square biased-sine fit) as well as a combination of these three algorithms were evaluated for static weight calculation in the presence of suspension vibrations.

Overall this system was found to be accurate to ±40% for smaller vehicles on the order of a few thousand pounds, and ±15% for heavy truck vehicles weighing 80,000 pounds, in the presence of suspension vibrations. The best results were obtained by using the lowest value of the least square error method and the average value method. Overall, the sensor is effective at measuring the weights of heavy vehicles and less so when measuring light vehicles. Despite this, possible improvements could be made through manufacturing changes, better installation practices and system characterization. In addition, the design could be made more robust to better handle high speeds, withstand the elements year round, and not be a source of perturbation for passing vehicles in the road itself.

I. INTRODUCTION

A. Project Objectives

The overarching goal of this project is the enhancement of a battery-less wireless traffic sensor to enable it to measure vehicle axle weights, as vehicles travel over the sensor embedded in the road.

In a previous project, a preliminary battery-less wireless sensor for measurement of traffic flow rates was developed by this research team. A photograph of the developed sensor is shown in Figure 1

Figure 1: Generation 1 battery-less wireless traffic sensor.

This sensor is 6 feet long and can be embedded in a slot in the road. As a vehicle travels over the sensor, it harvests energy from vibrations induced in the sensor beam. The energy is harvested using piezoelectric elements and is used to power electronics in the sensor which transmit a message to a roadside receiver. The roadside received can thus keep track of the number of vehicles that travel over the sensor, thus measuring traffic flow rate. Since the sensor requires no batteries or other power source and wirelessly transmits traffic flow rate measurement, it is called a battery-less wireless sensor.

In this project, the functionality of the battery-less wireless sensor is significantly enhanced:

First, the energy harvested from the sensor for each passing vehicle is increased by re-design of the sensor. The increased energy enables the sensor to communicate wirelessly over larger distances. With a transmission distance of 500 feet, the sensor can communicate with a receiver placed in one of the regular roadside cabinets on the highway. Further, the increased energy is useful to measure and transmit the vehicle axle weights, in addition to measurement of traffic flow rate.

Second, the sensor design is enhanced to enable it to measure the weight of each passing vehicle's axles. The new sensor design, associated electronics and the sensor performance in measuring axle weights are discussed in detail in this report.

Finally, the use of a series of such sensors in the road is considered, in order to enable compensation for vibrations in the suspension of each vehicle. Since suspension vibrations in a

truck can change the load experienced by the road by as much as a factor of two, it is important to be able to compensate for these vibrations so that the static load of the vehicle can be accurately calculated. Since the developed weigh-in-motion sensor is inexpensive, it is possible to use a series of consecutive sensors in the road, instead of just one embedded sensor. By analyzing the measurements from multiple sensors, algorithms to accurately calculate static vehicle axle weights in the presence of suspension vibrations are developed. A significant portion of the project report focuses on the weigh-in-motion measurements in the presence of suspension vibrations and algorithms to accurately obtain static weight.

To summarize, the key objectives of this project are:
a) To enhance the energy harvested by the sensor so as to enable increased telemetry distance and to enable supply of energy for additional weigh-in-motion related electronics.
b) To re-design sensor so as to enable weigh-in-motion measurements of vehicle axles.
c) To utilize a series of consecutive weigh-in-motion sensors in the road and develop algorithms to remove the influence of vehicle suspension vibrations for providing accurate static weights of moving vehicles.

B. WIM Sensors

Considering the huge cost of maintaining and rebuilding roads, transportation agencies have a great incentive to protect their roads. Many things are known to contribute to the rate of road deterioration. These include initial road construction quality, environmental conditions such as temperature and weather, traffic frequency, and traffic weight concentration. Among these, traffic weight is the most damaging factor to affect road surfaces. Since it is possible for trucking industries to overload trucks to increase profits, excessive loading on road pavements must be prevented through monitoring and enforcement.

Currently, many different weighing technologies are used to detect vehicle weight. First, weigh stations are present at exit ramps on some major roads. These require trucks to take an exit, pull off from the main road, and drive onto scales either embedded in the pavement or present on the pavement at the weigh station. These measurements are taken while the vehicle is at rest and as a result are very accurate. CAT scales are one of the most well known scale companies, and boast an accuracy of within 80 pounds [1]. A second weigh system category involves measuring vehicles without the need for them to slow down or take an exit. These are called weigh-in-motion (WIM) systems and are embedded in the main roadways themselves. Though these systems are inherently less accurate than static weight systems, they can be used to automate overweight vehicle monitoring and enforcement. In addition, these sensors can also be used to flag certain vehicles for further static measurement.

Current WIM technology is often wired to a roadside station which powers the device and stores measurement data. This roadside station may also interface with other technologies such as cameras, other data logging equipment, lights, or transmission devices in order to record vehicle tracking information together with axle weights. Even without the supplementary technology, however, current WIM sensors are very expensive to purchase and install. The table below shows the estimated cost of the three main WIM technologies.

Table 1: Current WIM Technology Information

	Piezoelectric	Bending Plate	Single Load Cell
Principle of operation	Load transferred through pavement to quartz sensing elements	Two steel platforms 2' x 6' ; Use of strain gauges	Two weighing platforms 6' x 3'2"; Single hydraulic load cell at center
Accuracy (95% confidence)	+/- 15% - 30%	+/- 10%	+/- 6%
Expected Life	4 years	6 years	12 years
Installed sensor cost per lane	$9,000 - $60,000	$15,000 - $85,000	$48,700 - $100,000
Roadside controller	$20,000	$20,000	$20,000
Annual life cycle cost	$4,750	$6,400	$8,300

Though WIM technology is in use today, it is too expensive to be as widely used as it needs to be to sufficiently monitor many roadways and prevent overloading. In fact, this monitoring is dwindling in recent years due to budget constraints [2]. Thus, in order to increase the use of these systems to reduce road damage, a WIM system is needed with comparable accuracy at a reduced cost.

C. New Sensor Motivation and Advantages

While the still measurement of a weigh station system is much more accurate than current weigh-in-motion systems, WIM sensors have an obvious advantage of being able to measure vehicle axle weights unobtrusively, since they are embedded in regular highway traffic lanes. With these sensors, vehicles do not have to exit the roadway, slowing down their travel, for their weight to be measured. As a result, WIM sensors ensure that all vehicles are measured for weight, instead of only a subset of them which are asked to be measured by exiting the road. As previously mentioned, a significant issue with WIM sensors is the very high cost associated with their hardware and installation, as shown in table 1. Highly accurate WIM sensors can cost as much as $100,000 per sensor per lane. In addition, a roadside controller can cost $20,000 and annual maintenance fees can add several thousand dollars per sensor.

Given the high cost and current technology of a WIM system, it would be most beneficial if a new inexpensive sensor were developed that was as accurate or more accurate than current WIM sensors. These sensors should be able to perform repeatably and reliably over life, and work for a large range of vehicle weights. As an added plus, the sensor would operate without the use of batteries or any power supply. Thus, it would be self powered. In addition, it would operate without the requirement of any wiring, significantly reducing installation time. The sensor would transmit its data wirelessly. If such a sensor could be produced, there would be a potentially large market for this product.

The development of such a sensor is the topic of this project. While the work done by the team does not result in a finished product as mentioned above, it does work towards that goal, and lays the groundwork for future developments which would make the sensor a reality. Thus, this report describes the development and testing of WIM technology which will supplement the creation of a battery-less wireless sensor.

II. FIRST GENERATION WIM SENSOR

As described earlier, a battery-less wireless traffic sensor was first developed by this research team and has been shown in Figure 1 of Chapter 1. Each sensor is 6 feet long, 1 inch wide, and embedded in the road such that one half of each axle (one of the two tires on each axle) passed over the sensor. This sensor utilized energy harvesting from the vibrations induced by each passing vehicle using piezoelectric elements on the sensor. This energy was then used to power a circuit and transmit a wireless signal corresponding to the passing of each vehicle axle.

A basic schematic of the electronics used to harvest energy from the piezoelectric elements is shown in Figure 2.

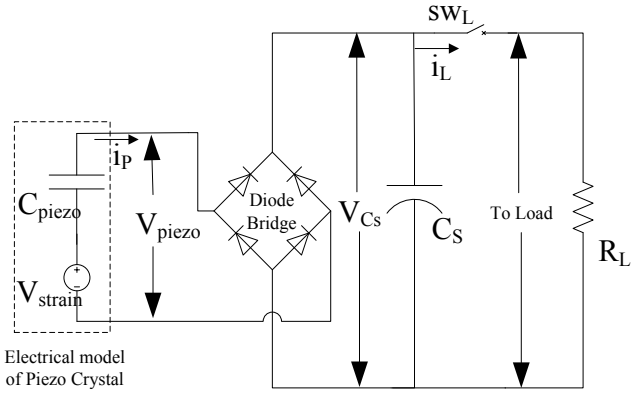

Figure 2: Basic schematic of energy-harvesting system.

The piezoelectric element is represented by a voltage source V_{strain} and a capacitor C_{piezo} in series with each other, as shown in Figure 2. The voltage from the piezo is rectified using a diode bridge and used to charge a storage capacitor C_s. A MOSFET switch controls the supply of power from the storage capacitor to the MSP430 microprocessor used for wireless transmission. It is important to supply power from the storage capacitor only after adequate voltage has been reached. Since the maximum voltage reached depends on the vehicle weight and is not known in advance, the optimal solution for maximal power transfer is to close the switch when the voltage across the storage capacitor reaches a maximum. The circuit used to implement the above schematic is implemented on a PCB and is shown in Figure 3.

A signature is assigned to each sensor in the form of a unique identification number. When a vehicle passes over the sensor, the circuit obtains energy from the mechanical vibrations. An encoder is used to encode this identification number to a sequence of '1's and '0's. A transmitter modulates an RF carrier wave to transmit this sequence of '1's and '0's to a receiver. Upon receiving this signal, the receiver demodulates the RF signal and recreates the series of '1's and '0's at its output. A decoder is used to convert this sequence back to the identification numbers number that was transmitted. The decoded signature is used to identify the sensor over which the automobile has passed.

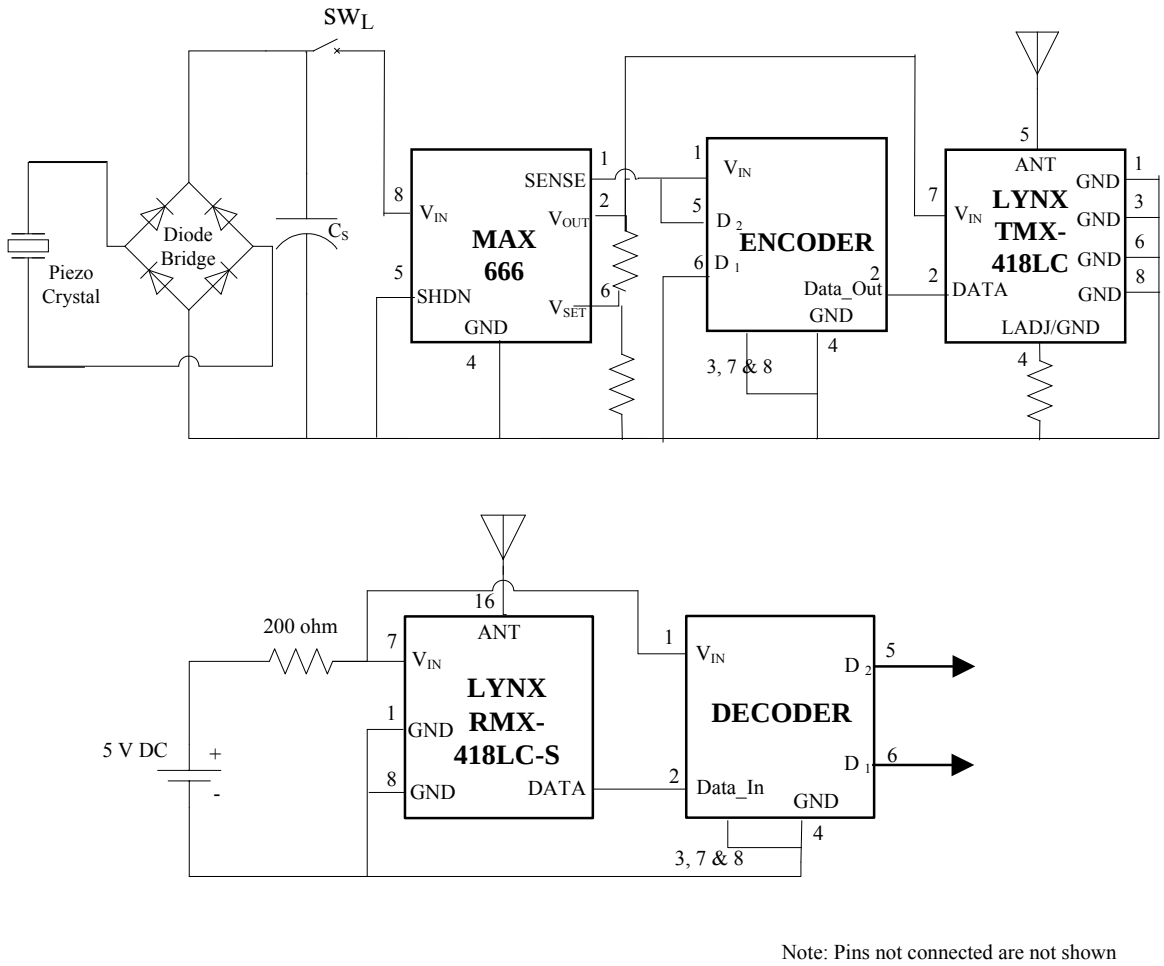

Figure 3: Circuit schematics for transmission and receiver circuits.

The performance of this traffic flow sensor proved that piezoelectric technology was a repeatable means to generate energy for signal transmission without the use of batteries or wires. Furthermore, this design proved that such a sensor could be built and installed at a relatively low cost. The sensor itself was built for an estimated $200 in material costs. The sensor performance was proven through installation and testing. It was secured in a slot in the road and multiple vehicles drove over it, with the sensor transmitting a signal every time a vehicle axle passed over it.

Following the traffic flow sensor, a similar sensor was developed with added functionality as a WIM sensor and with ability to harvest more energy for increased telemetry distance. Figure 4 shows a photograph of the original traffic sensor and the new sensor that incorporates weigh-in-motion measurement.

The new sensor has a 3-layer design. The sensor has a top beam (first layer) and three legs. A close-up photograph of each leg is shown in Figure 5. As seen in Figure 5, each leg has 2 layers, consisting of a weigh-in-motion layer and an energy-harvesting layer. The first layer in the

sensor will be referred to as the frame, the second as the WIM beams, and the third as the energy-harvesting beams.

First, note that the strain in the WIM beams and in the energy-harvesting beams is independent of the lateral location of the vehicle in the lane. This is illustrated in Figure 6. While the wheels can pass over the top layer (the frame) at any location along the frame, the total load acting on the frame is passed down to the WIM and the energy-harvesting beams. The total load on the three WIM beams is equal to the load acting on the frame. Further, the load acts at the center of each WIM beam. Thus the sum of strains on the three WIM beams is independent of the lateral location of the vehicle in its lane. Likewise, the sum of strains on the three energy-harvesting beams is also independent of the lateral location of the vehicle in the lane.

Figure 4: Generation 1 WIM sensor (left) and TFS (right).

Figure 5: Three layer structure of the first generation WIM design.

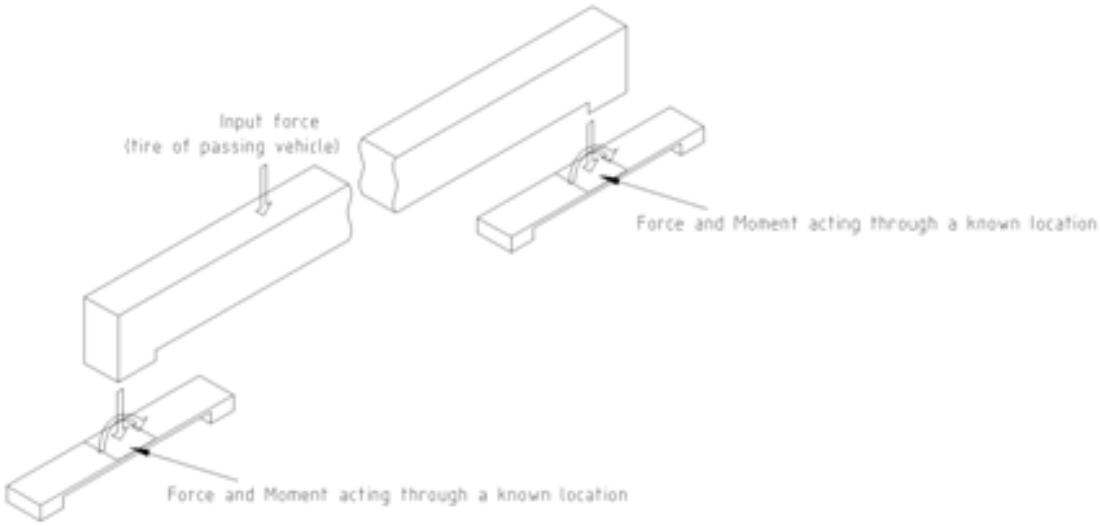

Figure 6: Voltage generation of Piezos independent of lateral load position.

Next, in order to increase the energy harvested from vibrations, the energy-harvesting beams were designed to have a nonlinear elastic response. Each energy-harvesting beam had thin plates that were significantly thinner than the beams used in the WIM layer. This enabled the beam to have a high strain for a given vehicle axle load. However, a mechanical stop prevented strain beyond a design limit. This allowed the energy-harvesting beam to have high strains for low loads (e.g. a motorcycle), but no significant additional strain for very heavy loads (e.g. a heavy truck). Since the voltage produced by a piezoelectric element is proportional to the strain on the element, the nonlinear elastic response ensured that the energy produced by vehicle loads was

adequate even for light vehicles, while at the same time ensuring that the beam did not fail for loads from heavy vehicles.

The WIM layer beams, on the other hand, were designed to have a linear elastic response. Each WIM beam was designed to be thicker and produce strains proportional to the vehicle axle load on the beam. Since the piezoelectric elements on this beam also produce a voltage proportional to the strain, the vehicle weight could be determined from measurements related to the voltage generated by the piezoelectric elements.

The traffic sensor function in the new sensor worked very effectively and was able to provide an enhanced telemetry distance of 500 feet.

Unlike the traffic flow sensor, the first generation WIM sensor failed to yield repeatable results. Rather than showing a reliable correlation between the sensor reading and vehicle weight, the output of the WIM sensor was not repeatable. This was a result of design problems in the sensor. While the sensor was able to measure static loads accurately, it was not able to measure dynamic loads from vehicle in a repeatable manner. This was analyzed as being due to resonances in the sensor that were excited by the vehicle loads. A final problem with the first generation sensor was that its width was too thin. The result was that as a vehicle traveled over the sensor, its full weight would not only be transferred to the sensor, but also to the road adjacent to the sensor. This made the sensor unable to fully weigh a given vehicle by design, and instead experience only a proportion of the weight. Since contact patches vary by vehicle and tire pressure, it would have been too difficult to estimate total axle weight using a sensor of a given width for multiple vehicles.

In order to develop an improved battery-less wireless WIM sensor, many improvements needed to be made to the design of the first generation WIM sensor. In order to gain a better quality sensor, the WIM portion of the sensor was worked on independently, since the performance of the energy-harvesting performance of the sensor had already been proven. It was important to refine the WIM response of the sensor, so that consistent readings were generated in which sensor voltage was proportional to vehicle weight.

III. SECOND GENERATION WIM SENSOR

A. Design Goals

Before designing the second generation WIM sensor many design goals were identified. These objectives were driven by the aforementioned motivations in addition to the shortcomings of the first generation WIM sensor:

- To fully weigh a vehicle, the full length of the contact patch of each tire had to be less than the width of the sensor.

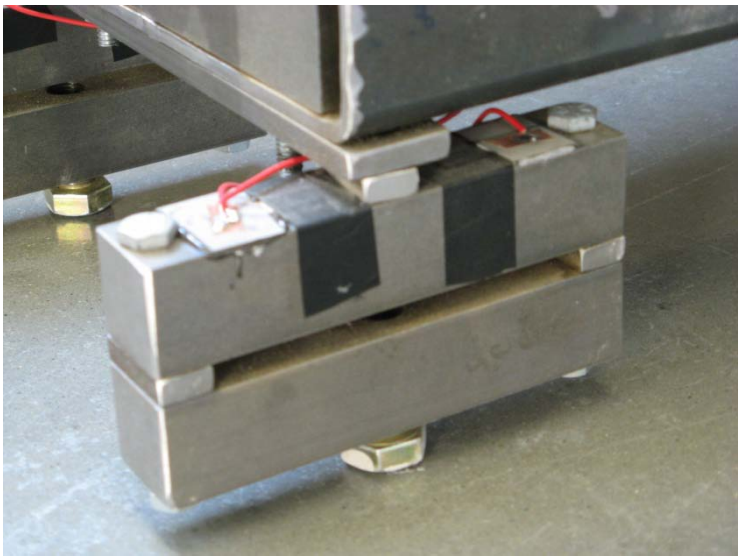

Figure 7: Fixed-fixed support beam with Piezos attached.

- The sensor must not be a source of disturbance in the road. This is to make driving over it a more pleasant experience, prevent damage to the sensor from snowplows or street sweepers, and to prevent fatigue of the pavement adjacent to the sensor from repeated side loading.

- The sensor must not show any resonance in the range of frequencies that the sensor might see.

- The sensor must be structurally robust enough to allow for the support of the dynamic weight of a semitrailer.

- The sensor system must data log the weight readings of each axle and be ready for weight measurement of the next axle at all reasonable highway speeds.

B. Load Transfer

In order for the sensor to read the weight of each vehicle axle at least half of the axle must exert its full weight on the sensor at some point in time. The half axle is assumed to represent the full

axle on condition of symmetry. This assumption allows for the sensor to be smaller and less expensive. Since a semi truck tire was thought to be the largest size contact patch which the sensor might see this contact patch was measured and found to be roughly 32 inches wide by 9 inches long. The sensor length was then designed to be 42 inches long by 10 inches wide to accommodate this contact patch.

In order to measure the half axle, the sensor surface also had to be able to withstand this weight without yielding or major deflection. Since the sensor was created using a similar structure as the previous WIM sensor without the third layer, this weight was distributed onto a single main beam with four supporting pillars (legs). The structure of a pillar is shown in the photographs in Figure 7. This was done to accommodate the increased width of the sensor while still forcing the load through discrete points. To be structurally sound, the worst case scenarios had to be considered for the main beam and pillars. For the main beam, this would be a load placed at the center of the beam. For the pillars, a single load placed on the corner of the frame would be a worse case as one pillar would have to bear all the vehicle weight. In addition to designing for this worse case so that no yielding would occur, it was also decided that the sensor should not deflect more than an eighth of an inch to prevent vehicle perturbation.

For the pillars, they were made short to reduce the moment in the beam, but also long enough to allow for a sizeable piezo to be placed on the beam under a single stress state of tension or compression. The resulting geometry for the pillar was 4.75 inches in length, by 1 inch in height by 1 inch in width. The main beam result was a square tube 2 inches tall with side walls .25 inches thick. Last, reinforcement blocks of solid steel were placed at each end to increase the strength of the beam and increase the rigidity of the pillar to frame connections. See Appendix A for beam calculations.

C. Two Layer Design

In addition to strength, cost was also a consideration in the design of the sensor. Since the piezo electric sheets used to generate charge were expensive and relatively difficult to work with, the number of these sheets was minimized in the design. This was accomplished using two layers as previously mentioned. This configuration allowed for a voltage reading which was proportional to the weight of the vehicle but which did not depend on the lateral position of the vehicle in the lane (see Appendix B).

Figure 8: Current WIM sensor design.

If a single beam were used with piezos mounted to one location on its surface, the mechanical strain generated and resulting voltage produced would be dependent on the lateral position of the load. In order to create a single layer design which would be independent of lateral position, piezos would need to fully support the load of the vehicle and thus, span the entire length of the beam. Since this solution would be much more expensive than a two layer design and yield no benefit, a two layer design was used.

In the two layer design, the load is transferred from the top beam to the lower beams through a known central location at each of the lower beams (see Figure 5 and Figure 7). This well understood system of fixed-fixed supports may then be easily analyzed to determine the stresses in the beam. With this knowledge, the geometry was selected for the main frame and for each support as previously mentioned.

D. Stress Analysis

Calculating the stresses was also necessary to place the piezo electric pieces on the support beams. It allowed for proper placement at a location which was in tensile stress over the whole area. Failing to identify such a location could result in little or no voltage generation if the piezo were to be placed on an area which saw both tension and compression. With the piezos placed properly, the stresses induced in the piezo could be calculated to determine what the voltage output of the sensor should be for a given load. This calculation is shown in Appendix B.

E. Natural Frequency of Sensor

A final design consideration was for the natural frequencies of the beam. This was important because if the sensor had a resonant peak for the vibrational frequencies seen during vehicle loading, it could generate a greater voltage in the piezos for some operating conditions. Thus, in order to make the sensor output independent of velocity, the first bending modes of the sensors and associated resonant frequencies were calculated along with the translational mode of the beam.

These calculations came into account when designing the system and the structural design was checked for resonance as shown in Appendix C. The result was that the top beam showed resonance but the bottom beams did not. This resonance was calculated without the consideration of the stiffening supports at the end of the frame. To ensure the robustness of the design against resonance, testing was completed and no speed was found to show any sign of such behavior.

F. Result

The result of these constraints was that the frame was made of a large steel tube and supported by one pillar on each corner. These four pillars, or support beams, were placed to the very edges of the frame to make the sensor the most stable. The final tube dimensions measured 2" tall by 42" long by 10" wide with a .25" wall thickness. Each half of the support beams measured 1" by 1" by 4.75" in length. These pillars were symmetric in structure but only piezos were placed on the upper support beams to make them more accessible for replacement or maintenance. Two were used for redundancy in case one was to become separated from or damaged on the beam.

Bolts were used to support each pillar, and they could be rotated to raise or lower the structure. Without this adjustment it would be difficult to make the sensor flush with the road. This would result in excess dynamics which would hinder accurate measurement. The resulting design is shown in Figure 8.

IV. WIM MEASUREMENT ELECTRONICS

The electronics of the sensor consist of an energy producing element, an energy storage element and a device to measure and record the amount of energy stored. The piezo is the energy generating element and it is the most fundamental part of the sensor. It generates a charge in response to a strain. This charge can be quantified by equation 1. Here, Q is the total charge produced on the piezo surface per unit area, E is the young's modulus of the piezo, epsilon is the average strain which the piezo is under, and d_{31} is a constant relating the charge production of the piezo to the stress of the element.

$$Q = E\varepsilon d_{31} \quad (1)$$

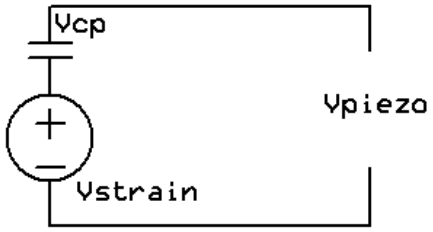

Figure 9: Model of a piezoelectric element.

The charge produced by the substrate of the piezo is stored on its surface and as a result, the piezo acts the same as a capacitor at low frequencies. For this reason an appropriate model for the piezo is shown in Figure 9. Here the piezo is shown as a voltage source in series with a capacitor. The open circuit voltage produced by the piezo is shown in equation 2, which is similar to equation one. Here, 't' is the thickness of the piezo, and 'g_{31}' is the open circuit electric field produced in relation to the applied stress. Note that the constant is different than in equation one and that voltage is not presented as being per unit area of the piezo.

$$V_{strain} = E\varepsilon g_{31} t \quad (2)$$

These piezoelectric elements were used to generate a charge by placing them at locations of strain on the lower beams of the sensor. In order to adhere the piezos to this surface, an epoxy was spread on the beam and one side of the element was pressed to it. To transfer the charge from the piezo to a circuit for use, copper adhesive tape was placed as electrodes on both sides of the piezo. When a strain was induced in the beam, because of how the piezo were bonded to it, this strain was also induced in the piezo and a charge generated.

These elements were wired in parallel such that the same load could be placed at any location on the sensor with the same total voltage resulting. This calculation is found in Appendix B. Once the charge is generated, it passes through a diode and is stored on a capacitor (see Figures 8 and 9 for schematic). This capacitor scales the voltage reading in relation to the capacitance of the piezos. This is advantageous as any data acquisition system or microcontroller used to measure the voltage across the storage capacitor will have a maximum voltage reading that it can

sample before saturation occurs. See Appendix D for the selection of the capacitor size and scale calculation.

With this hardware in place, the sensor will generate a voltage proportional to the load placed on the sensor. As the vehicle travels over the sensor the charge from the sensor will build until it becomes large enough to drive current through the diode. Once this occurs, charge begins to accumulate on the storage capacitor and continues building until the full weight of the vehicle is transferred to the sensor. As the vehicle moves over and off of the sensor the charge on the piezo will decrease until it goes to zero. The storage capacitor will remain charged, however; as the diode prevents the charge from leaking off. This charge can then be read by a microcontroller or data acquisition system.

For multiple axles traveling over the sensor, the aforementioned circuit would only be able to read the first axle or subsequent axles of a greater weight. This can be overcome by implementing a system to short the capacitor between axles readings. In order to do this before the next axle moves over the sensor, a microcontroller was used to monitor the capacitor voltage and short it when it reached a local maximum.

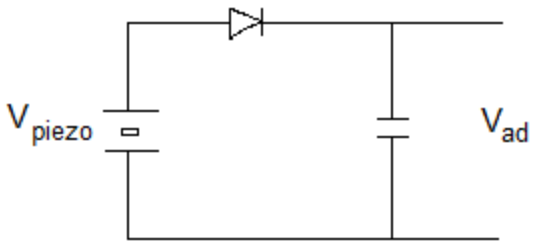

Figure 10: WIM sensor voltage storage circuit.

Figure 11: DIP configuration of the MSP430F2013.

The monitoring of the circuit and the shorting of the storage capacitor were both done by a Texas Instruments' MSP430F2013 (see Figure 10). In order for the chip to operate, C code was developed and uploaded to the chip controlling the processes of the microcontroller. The code was written to operate in an infinite loop and to constantly check three consecutive values of the voltage across the storage capacitor. If three consecutive values are found to be decreasing, it is assumed that the vehicle half axle has moved completely onto the sensor and is beginning to move off of it or is completely off. The C code tells the MSP430 to turn on one of its output pins. This pin is connected to the gate of two MOSFETs which short both the piezo and capacitor leads simultaneously.

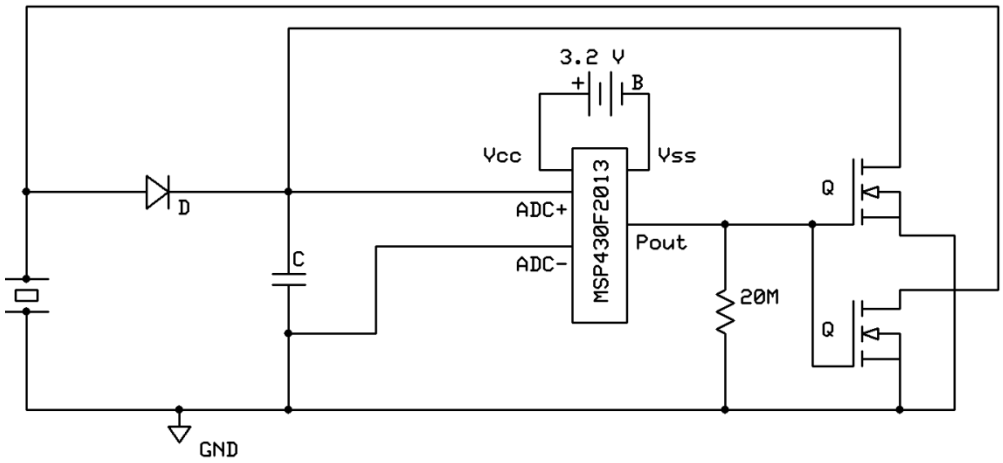

Figure 12: Sensor storage and measurement circuit.

In order to write this code there were a few details which needed to be addressed. First, there was found to be noise on the input signal. Without any set sampling rate, the MSP took consecutive readings at up to 200kHz. This would result in readings which were not representative of the overall trend of the data, but instead the local variation of the signal. In order to get around this, a timer was used to delay each sample so that an effective sampling rate of 100 Hz was used. Second, the registers of the timer and analogue to digital converter had to be set. This was difficult as many code registers needed to be established for the micro to perform as intended. It was often unclear which ones needed to be setup and which ones could be left to default values. Finally, the MSP430 needed a resistor on a certain port to be powered by a means other than a usb programmer. These issues were often not well documented and solving them took considerable time.

Once the program was working it was uploaded to the microcontroller. The microcontroller was then connected to the circuit as shown in Figure 12. This completed the test circuit.

V. EXPERIMENTAL SET UP

The most fundamental goal of a weigh-in-motion sensor is that it must measure the static weight of a vehicle of interest. In order to determine how effective the system was at reaching this goal, the testing of the sensor with real vehicles had to be completed. Four sensors were tested to find out how well they each worked given the variation inherent in manufacturing and installation. Four sensors were also used to determine if multiple sequential sensors in the road could be utilized to remove some of the error associated with the measurement of static weight using a dynamic vehicle.

Before any data could be taken, the sensors needed to be installed such that a vehicle could safely drive over them at high speeds. This required installing them into slots in the road so that they provided minimal perturbation to both vehicle and sensor during measurement. Many photos can be found in Appendix E which show this process. First, a large hole had to be made deep into the road surface for each sensor. Once made rebar was used to strengthen the new concrete base. Liquid concrete was then poured into the holes and wooden forms were installed before the concrete set to ensure that the sensors had a hole large enough to be set down into for testing. Once the forms were pressed into the wet concrete and the concrete allowed to harden, the forms were removed. The result was four slots in the road which were the correct shape as to fit each sensor. Figure 13 shows a single sensor placed into a slot and then adjusted to the correct height.

The geometry of each slot was determined by each sensor. For the spacing between sensors, it was desired to space them far enough apart to capture four points on one period of oscillation of vehicle suspension vibration for a reasonable velocity of travel (see section 7). Alternatively, the sensors could not be so close together as to result in a wall thickness separating the slots which would be structurally unsafe for heavy vehicles. The resulting wall dimension was chosen to be six inches.

Once the sensors were installed, data was taken on many different days. During the first testing day on 3-16-2011, data was taken at three speeds: a rolling speed, 15 mph, and 30 mph. At each speed for each axle, three runs were conducted. Three vehicles were tested for each of these conditions. The vehicles were a 1988 Toyota Corolla, a snowplow, and a Navistar semi truck (tractor trailer). The Corolla and the snowplow had two axles each, and the semi tractor trailer had 5 axles all together. In all, the axle half-weights (half is used because only one half the tires of each axle went over each sensor) ranged from 650 pounds to 9000 pounds.

Figure 13: WIM sensor installed in a slot at MnROAD.

The second day of testing yielded much more data at higher speeds. On 4-9-2011, the same vehicles were tested, but five runs were conducted at each of the following speeds: rolling speed, 10 mph, 20 mph, 30 mph, 40 mph, 50 mph, and 60 mph. Data was able to be gathered at 50 mph, though the semi made the sensor very unstable and it bounced as much as an inch out of the slot that it was installed into. For this reason not all five runs were conducted for the semi tractor trailer. For the Corolla only, 60 mph data was taken and the sensor was found to be quite stable. Another note of the testing for this day was that sensor 4 was found to have a bad connection and some poor results were unexpectedly gathered for a portion of the testing. Thus, some of the data for this sensor was thrown out.

The final testing day involved less runs, but acceleration data was gathered in conjunction with weight data. For the Corolla and semi tractor trailer, accelerometers were attached to the rear body and axle pair and drive body and axle pair respectively. The accelerometer data was then synched with the voltage output data from the sensors to see if the sensors themselves were a cause of increased vehicle vibration. Our data acquisition system for the accelerometer data was limited to sampling rates which aliased the accelerometer signal at frequencies of 50Hz and above. Thus, some of the data observed was not the correct frequency, as was found to be the case for the body vibrations measured of the semi truck.

VI. SENSOR PERFORMANCE

For actual testing, four sensors were used in order to obtain multiple readings for each axle as it traveled over the set of sensors. The voltage storage and monitoring methods used in a single sensor were replicated four times (one for each sensor) and the voltage stored across the capacitor for each sensor was recorded using multiple channels on a National Instruments Data Acquisition system (NI USB-6211). As a single vehicle axle traveled over each sensor, a channel of the data acquisition would record the voltage stored on the capacitor in time. This voltage would increase as each axle traveled onto and over the sensor where it would reach a maximum. It would then decrease very slowly in response to charge leakage through the internal resistance of the data acquisition until it was shorted completely by the microcontroller (after detecting a local maximum). Since the NI Daq recorded everything, the max value of the capacitor could be observed by graphing the data stored in this file using MATLAB. An example of such a file plotted is shown in Figure 14.

Once the data was plotted, each peak was manually recorded and documented to show for a given day, run, speed, axle, and sensor what the corresponding maximum voltage reading was which represented the vehicle's weight. This data was then graphed to show the relationship between the measured static weight of the test vehicle and the dynamic voltage reading as shown by the sensor.

An example of such a plot is shown in Figure 15. This data collected from sensor 2 in the Figure was shown for 9 different vehicle axle weights at low speeds where very little vehicle dynamics were present. The result was that the data appeared very linear and accurate from 5 to 10%. This behavior is a good representation of all sensors tested. It shows that the WIM sensor works and is accurate under conditions where there are no suspension vibrations.

Further processing was done using MATLAB to generate calibration results. First, all maximums were entered using MATLAB in a single multidimensional matrix holding all data from each run, for each axle, for each sensor, and at each speed. A program was then written to iterate through a range of possible solutions for a given estimator considering only the four data points from a single run. For loops chose a possible solution, found the error involved with the selection, and stored the information about the estimator if it was found to be favorable.

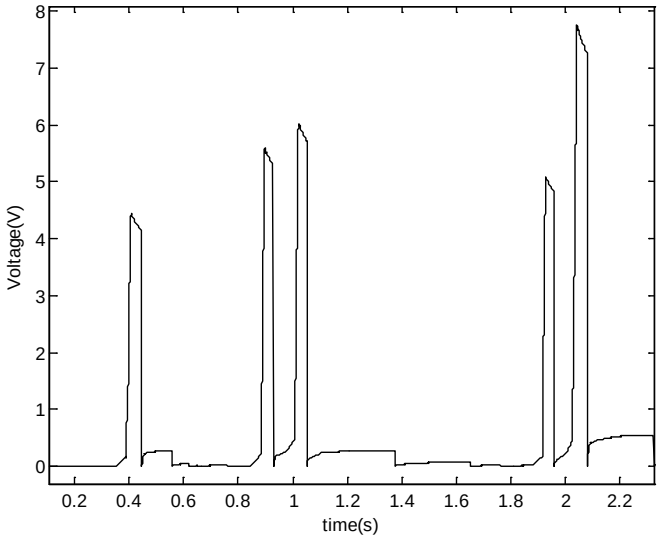

Figure 14: Typical semitrailer run over a single sensor.

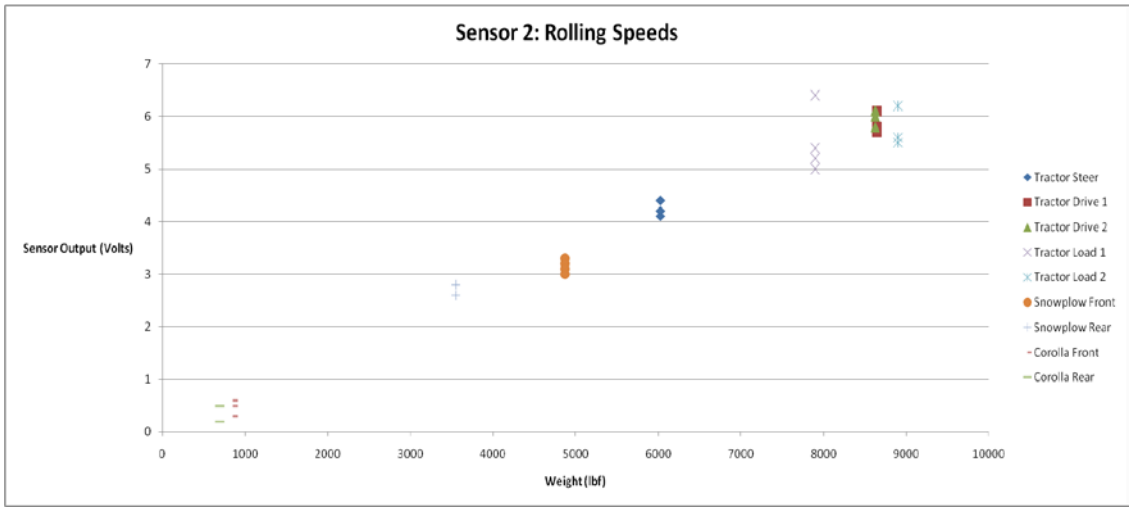

Figure 15: Low speed plot of voltage versus static weight for multiple test vehicles.

This methodology for finding a solution was attempted using many types of estimation methods:

- Taking the average of all four readings
- Finding the value which has the lowest square error between the estimator and the sensor values.
- Finding the best fit biased sine wave to the data points which reduces the error of the estimated phase, amplitude, and offset of the wave.

These methods discussed here will be presented in greater detail in the next section. Once produced these best fit estimators were written to a file and then entered into excel for plotting.

Excel was used to look at all the data, including estimators, to determine their relative effectiveness. All data was then plotted, with poor data noted and excluded.

VII. VEHICLE SUSPENSION VIBRATIONS

A. Theory

A major challenge to reading the static weight of a moving vehicle is that its horizontal motion is transformed into a vertical motion as it moves over any perturbation in the road surface. These road disturbances displace the tires of a vehicle and result in vertical oscillations of the vehicle axle and body masses. This displacement alters the force on the road surface as exerted by the vehicle itself. Since the desired output of the sensor is to capture static weight and not dynamic weight on the road, the oscillatory behavior must be understood and compensated for in order to gain the desired output for a WIM sensor.

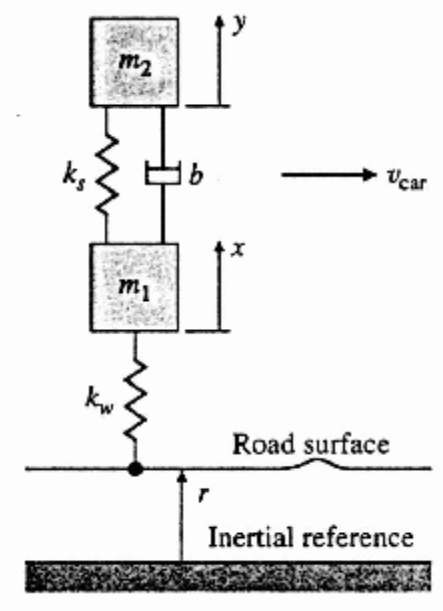

Figure 16: Quarter car model for vehicle tire and suspension.

Vehicle oscillation is commonly described by isolating a single tire and suspension combination called a quarter car model. This model looks at the vehicle suspension as shown in Figure 16 [3] using springs masses and dampers. Here m_2 corresponds to the mass of the vehicle body supported by the suspension element and m_1 to the vehicle half-axle. Note that the variables x and y are defined from the position of static equilibrium (gravity is excluded in vehicle accelerations). By analyzing this system the following equation can be found to describe the dynamic force exerted on the road by the vehicle.

$$F_r = -m_1\ddot{X} - m_2\ddot{Y} \quad (3)$$

The total force that includes static and dynamic components is given by:

$$F_r = -m_1\ddot{X} - m_2\ddot{Y} - (m_1 + m_2)g \quad (4)$$

Given the equation for the force on the road, we can quantify and better understand this oscillatory behavior. By looking at the equation, the road force is known if both the acceleration and static weight is known for both axle and body mass. Since any sensing element will experience the force exerted on the road, both these pieces of information will be needed to quantify road force in time.

B. Accelerometer Data

As previously mentioned, accelerometers were attached to the axle and body masses on two vehicles to obtain their accelerometer data while traveling over various perturbations. The vertical acceleration data of the axles and body was then graphed to understand the magnitude and frequency of these oscillations. The data shown in Figure 17 comes from one of the three main test vehicles, the Toyota Corolla. This was captured while driving at 20 mph over the WIM sensors when they were mounted 1/8 inch below the road surface. The vibration is shown to be on the order of 15 Hz for the axle and significantly greater than 50Hz for the body (the 100 Hz sample rate used was not able to capture the frequency). It should be noted that the vibration seen in the body is likely due to flexible body vibrations and does not correspond to the quarter car model. Also, it can be seen that the body vibrations are significantly smaller in magnitude compared to the axle vibrations.

Another set of data was taken many months earlier for a different perturbation for comparison. In Figure 17, the Corolla is shown driving over a large pothole. While this plot does not show the behavior of sustained oscillation of the suspension, it does show how the acceleration changes in time at a very high sampling frequency. Here a 5000Hz sampling frequency is used. There is shown to be more high frequency, low amplitude vibrations on the axle when compared to the body. This follows driver intuition and makes sense.

Before the road force could be plotted for this scenario, the Corolla weight needed to be measured. This was done at a CAT weigh station along with all other test vehicles. With this information Figure 19 could be plotted using the accelerometer data. Here it is observed that while the wave form could be modeled by a growing and decaying sine wave, there are many local high frequency variations in the data which deviate from the sinusoidal form. Thus, if the WIM sensor were to estimate static weight using a biased sine wave estimation method, the result would be less accurate because of these high frequency oscillations.

The semi truck was also used to gather accelerometer data for the axle and body. In Figure 20, the semi drive axle and body accelerations are shown for the same road perturbations used in Figure 17. Here the speed is 10mph rather than 20 mph. The resulting oscillatory behavior turned out to be around 15Hz for the axle and around 40 Hz for the body. The semi truck body data, as with the Corolla, showed signs of aliasing and rigid body motion, making the actual body behavior uncertain.

Figure 21 shows the semi drive axle traveling over a bump 1.5 inches high by 4 feet wide at a speed of 30 mph. In the plot only the axle acceleration is shown since the body acceleration is usually aliased and much smaller in amplitude. It is evident that the oscillations shown in the Figure are in the range of 10 to 20 Hz. Another observation is that the accelerations are larger than in Figure 17. This would correctly follow from the road perturbation being bigger.

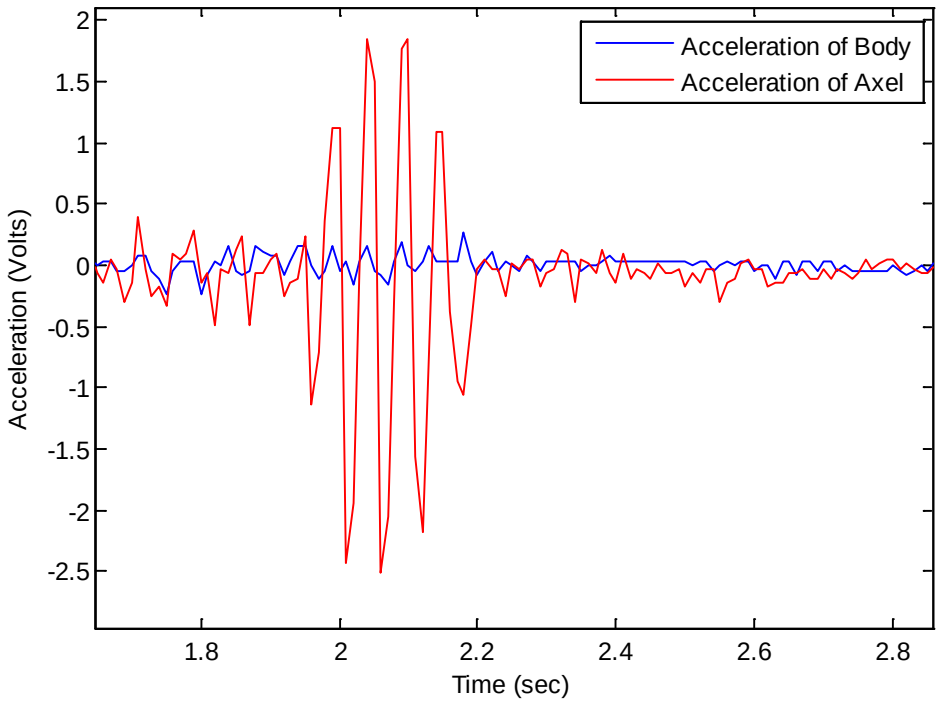

Figure 17: Accelerometer data for a 1988 Toyota Corolla while traveling over all four sensors.

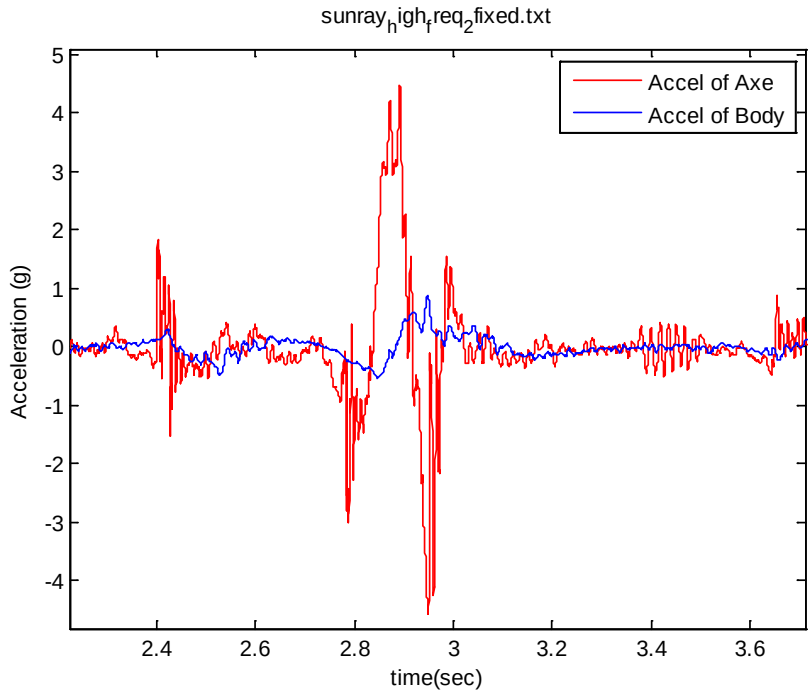

Figure 18: Accelerometer data for large pothole perturbation at slow speeds for 1988 Toyota.

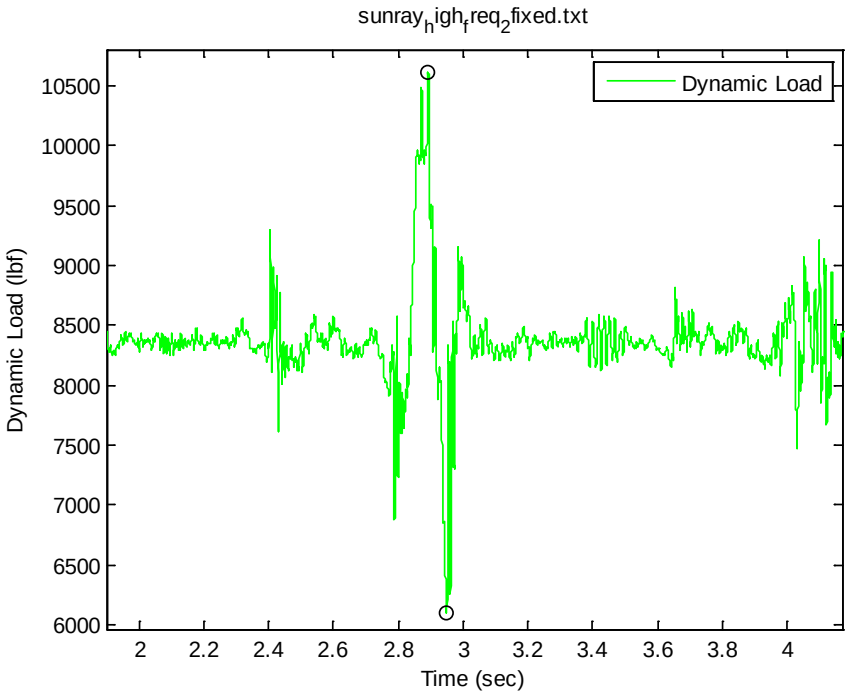

Figure 19: Road force plot for large pothole perturbation at slow speeds for 1988 Toyota.

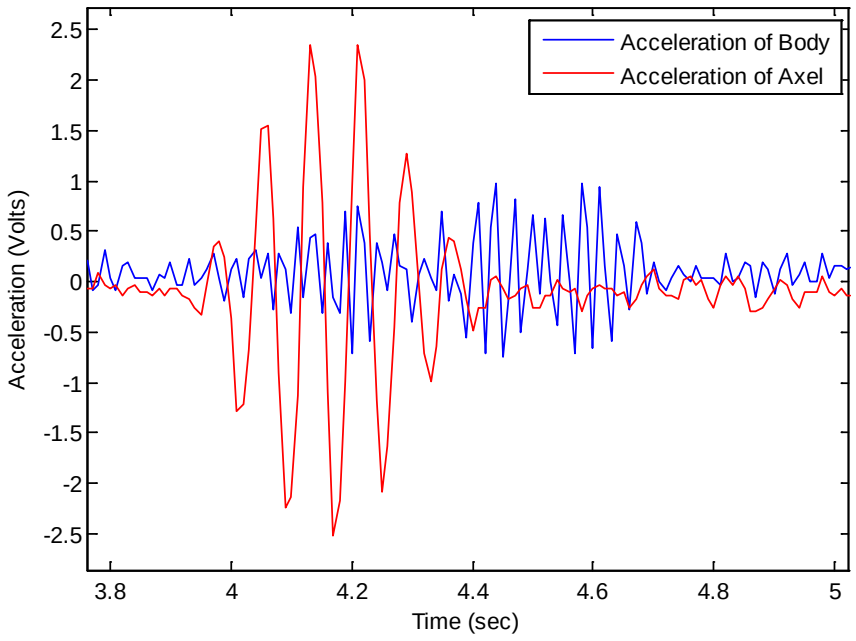

Figure 20: Accelerometer data collected for semi drive axle while moving over sensors.

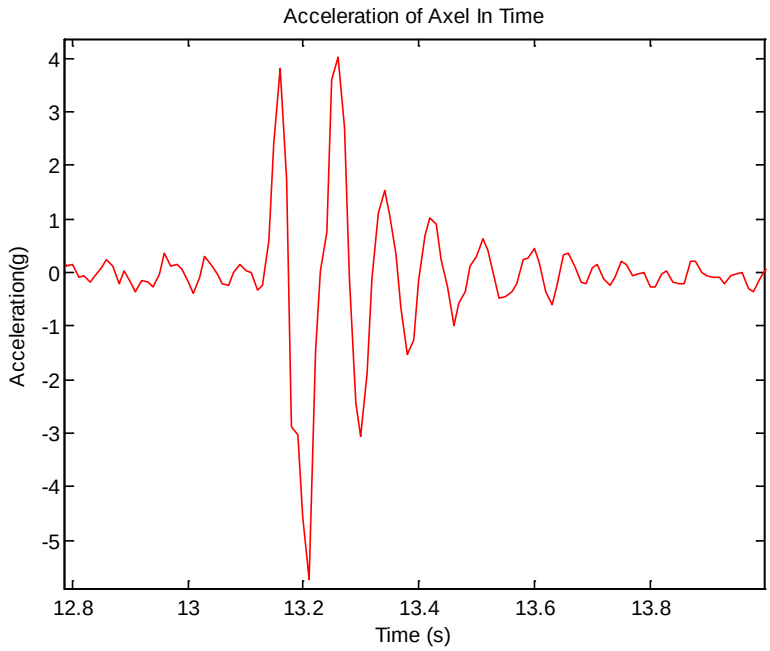

Figure 21: Semi drive axle accelerations for road perturbation 1.5 inches high at 30 mph.

From all the plots for both vehicles it can be assumed that vehicle axles oscillate around 15Hz. Since no data was able to be obtained experimentally which adequately showed significant body oscillations, it can only be known that the amplitude of these oscillations are much less than the axle. The data either did not show much variation in acceleration in time as to ascertain a frequency from or else it seemed to be too large in amplitude and too small in period to represent what is common for body oscillation. Thus, the oscillatory behavior of the body will be assumed which is commonly presented in literature on the order of 1Hz.[1]

Going forward it is clear that axle and body oscillations have specific frequencies which may be modeled accurately by sinusoidal behavior. These frequencies depend very little on perturbation size, or vehicle velocity. While these factors do not affect frequency of oscillation, they do affect amplitude and duration predominantly. Thus, these factors must be taken into account when analyzing the WIM data.

Table 2: Sensor Relation to Vehicle Oscillation for Spacing and Sensor Geometry

Velocity (mi/hr)	Time Between Sensor Centers(s)	Sensor Readings Per Period	Duration of 15 Hz Period on Sensor (%)
3	0.30	0.2	284%
10	0.09	0.7	85%
20	0.05	1.5	43%
30	0.03	2.2	28%
40	0.02	2.9	21%
50	0.02	3.7	17%
60	0.02	4.4	14%

[1] Rajamani, R. (2006). *Vehicle Dynamics and Control.* Verlag, New York: Springer.

C. Measurement Geometry

The next step is to understand how these dynamics will be captured by the WIM sensor to produce a single reading. This is important because the dynamic weight of the vehicle as measured by the sensor will change for different vehicle speeds. To understand this better the vehicle dynamics can be modeled using a sine wave with constant amplitude. A constant amplitude is acceptable here (rather than an increasing and decreasing one as shown by the data) because it is only needed to understand what percentage of the period of oscillation is captured by the sensor for a given vehicle speed.

From the data it was observed that both vehicle suspensions oscillate on the order of 15Hz. Assuming this frequency for all vehicles and the sensor geometry and spacing as previously described, table 2 shows how many readings would be taken in a single period at a given speed. Conversely, it is shown how much of the full period of oscillation will be captured by each sensor.

This table quantifies the trend showing that a faster vehicle velocity results in less of each oscillation being captured by a single sensor. Inversely, at faster speeds more consecutive sensor readings will be captured within a single period. Since it is the maximum value which the sensor sees in time that is important, the specific portion of oscillation observed should be understood and not just the percentage. To further understand this concept, plots were created to show which part of the oscillation each sensor would see. In Figures 22 through 24, a different color represents a different sensor. These plots were created for a 15 Hz oscillation about an 855 pound static load.

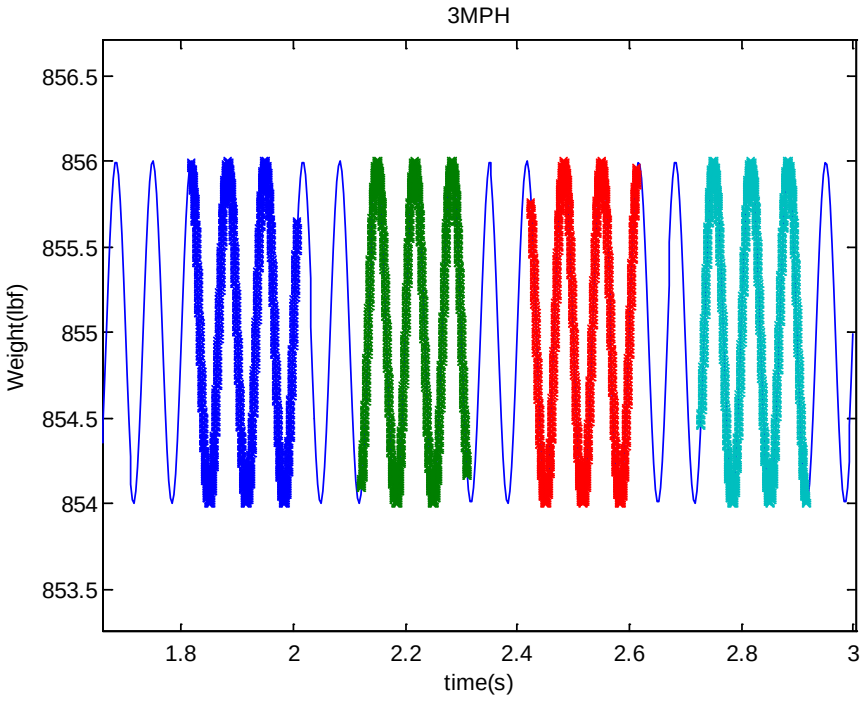

Figure 22: 3 mph contact between sensor and tire.

We can see that for slow speeds (3 mph) multiple periods are captured by each sensor. For intercity speeds (40 mph), we see that all sensors combined capture around one period. At highway speeds all 4 sensors would capture less than a full period.

Though these plots show the time during which the vehicle is in contact with the sensor, they are not representative of the weight which the sensors will see in time from the vehicle. As the vehicle enters onto the sensor, the tire and the sensor are in contact but the sensor reading of the weight is small since most of the weight is still supported by the road adjacent to the sensor.

Assuming that the contact patch of the vehicle is relatively uniform in pressure, the colored values in the previous plots will be scaled, with the beginning and end multiplied by zero and the middle region (when the contact patch is fully on the sensor) multiplied by one. The result is a more accurate representation of the theoretical weight which the sensors observe in time for the same velocities as shown in the previous Figures. In Figures 25 to 27 a black line is added to show this scaling factor (note that the amplitudes are adjusted).

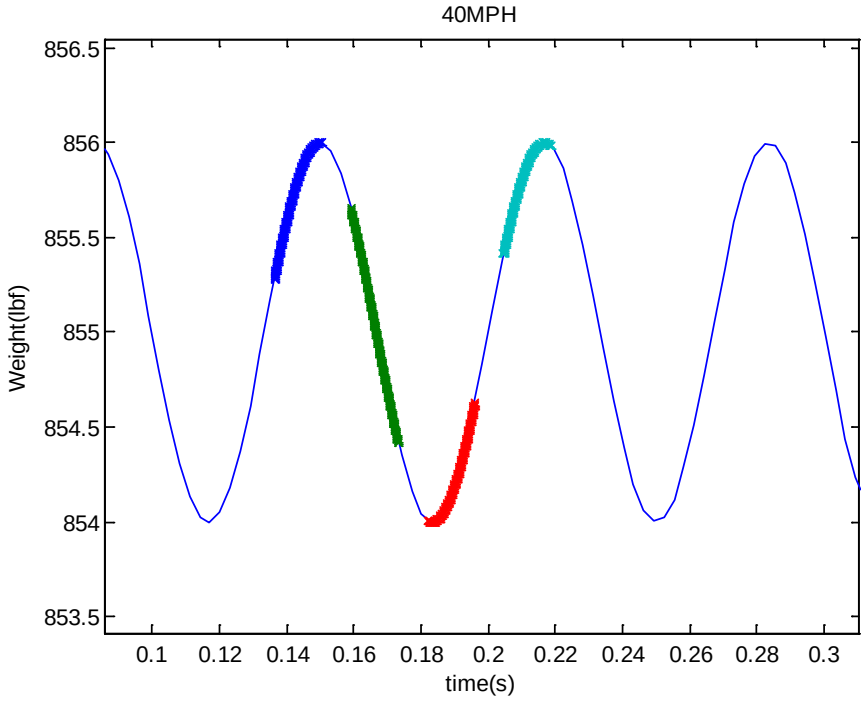

Figure 23: 40 mph contact between sensor and tire.

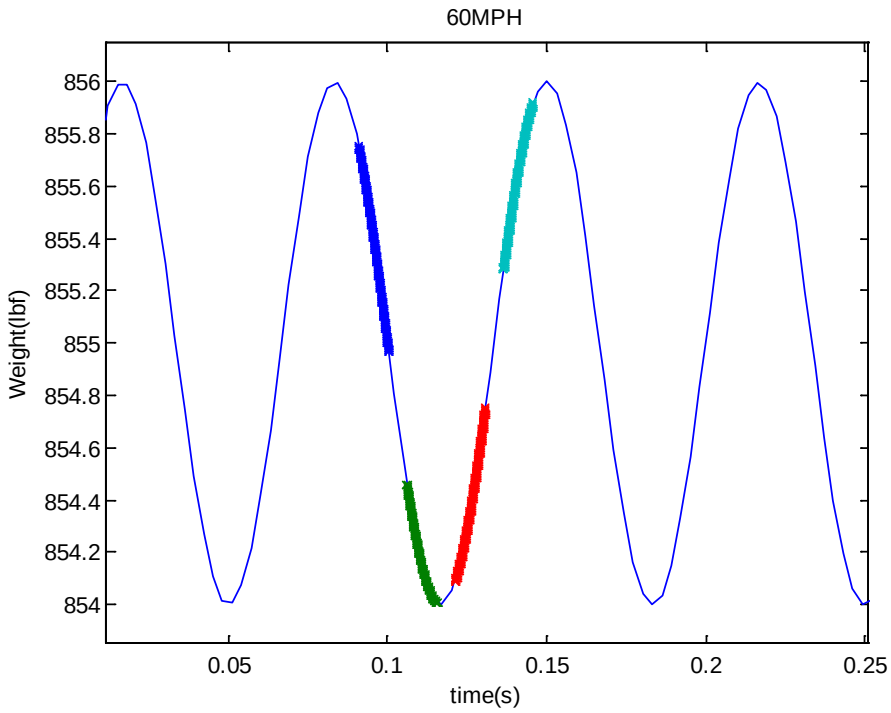

Figure 24: 60 mph contact between sensor and tire.

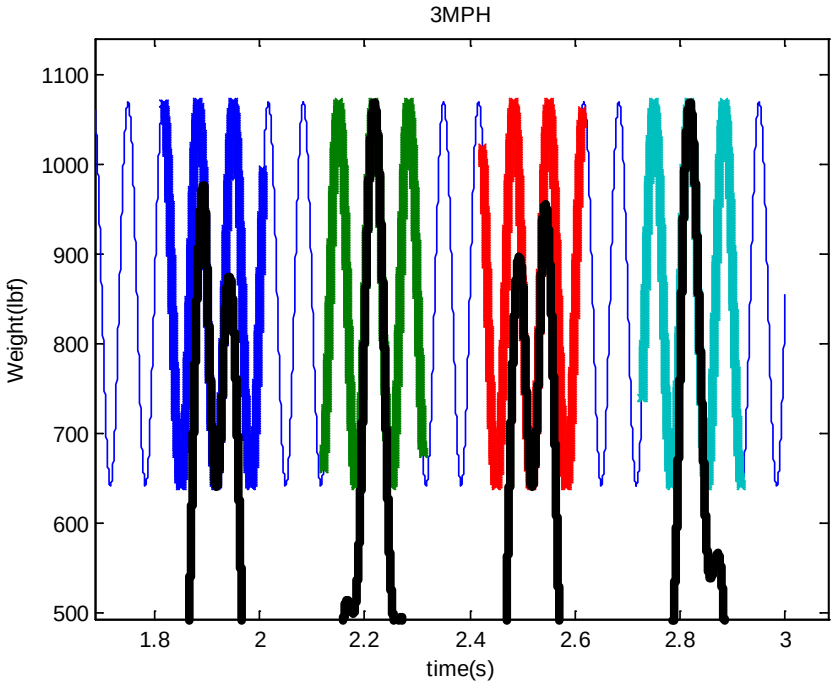

Figure 25: 3 mph theoretical load applied to sensor.

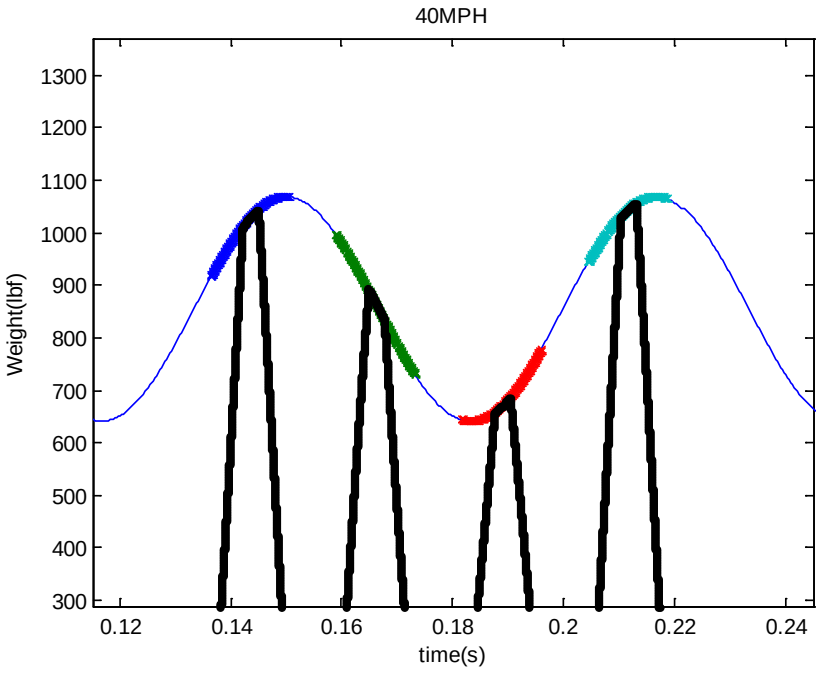

Figure 26: 40 mph theoretical load applied to sensor.

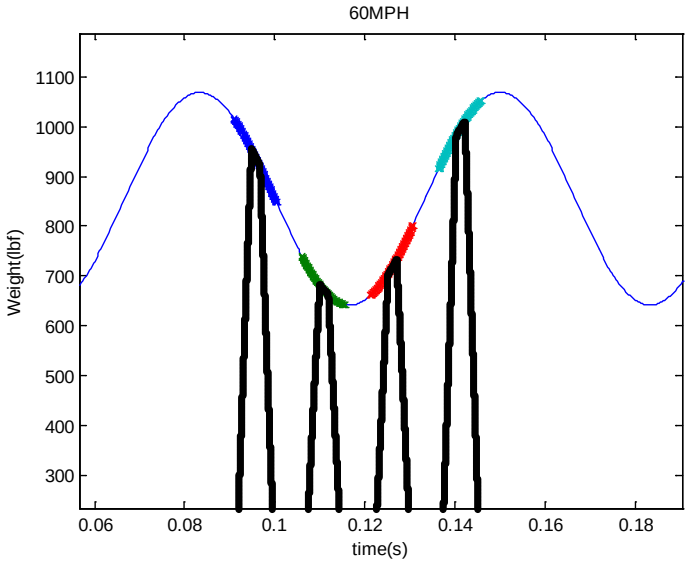

Figure 27: 60 mph theoretical load applied to sensor.

From these results it is evident that the actual observed weight is much different than the weight of the previous plots. More specifically, the amount of the vehicle dynamics each sensor sees is reduced. The updated plots show that a single sensor captures over one half of a period rather than three periods at 3 mph. At 40 and 60 mph, the same number of points is captured per period by the sensors but each sensor captures less of the vehicle dynamics.

D. Possible Sensing Algorithms

There are many algorithms which may be applied to this data to obtain static weight. Three are listed and implemented for actual sensor data. The first is to average the readings from all sensors. The averaging method is shown in equation 5 where Wi is each sensor reading.

$$Average = W_1 + W_2 + W_3 + W_4 \quad (5)$$

For lower speeds at high amplitude and sustained dynamics, the method would be very inaccurate as each sensor would estimate a value much too high. For higher speeds, however; the accuracy would improve. The higher the speed, the more accurate this method would become within the legal driving limits.

Another possible estimation method is a least squares error method. Using this method an estimator would be found that would minimize the square of the error between each reading and the estimated value. This method would be similar to averaging but for higher amplitude dynamics with a fast decay, the resulting estimator would favor outlying values while averaging would favor the grouped values. Thus, this method would be better if many points happened to be captured at the bottom of multiple periods and a single high point was also captured. Rather than estimate low, it would favor the outlier. This least squares estimator (LSE) is shown in equation 6.

$$LSE = \min\left(\mathrm{SE}(x)\right) \forall\, x \in \{0, \dots, \max(W_i)\} \quad (6)$$
$$\text{where } \mathrm{SE}(x) = (W_1 - x)^2 + (W_2 - x)^2 + (W_3 - x)^2 + (W_4 - x)^2$$

Finally static weight can be estimated using a biased sine approximation method. This would be most appropriate for very high speeds were multiple points all stay within one period and all are grouped closely together. This situation would not be suited well for the averaging method or least square error method. This sine approximation estimator (SAE) is shown in equation 7.

$$SAE = A\ for\ \min\left(\sum_{i=1}^{4} e_i^2\right) \forall\, A, B, and \emptyset \quad (7)$$
$$where\ A \in \left\{\frac{\max(W_i)}{2}, \dots, 2 * \max(W_1, W_2, W_3, W_4)\right\};\ B \in \left\{0, \dots, \frac{\max(W_i)}{2}\right\};$$
$$w = 15\,Hz;\ and\ \emptyset \in \{0, \dots, 2 * pi\}$$

and

$e_i = W_i - f(t_i)$ where t_i is equal to the time the sensor was crossed and

$f(t_i) = A + B * \sin(w * t_i + \emptyset)$

From the previous plots and proposed methods of estimation, it is clear that certain speeds are better estimated using certain methods over and above others. For slower speeds of up to 10-15 mph it is difficult to use any method as each sensor will only capture the maximum value. Though it would seem that this would be a cause of concern, it is not since any perturbation will most likely die out between the time the weigh-in-motion slab is entered onto and the time the final sensor is transitioned. If the sensor slab were to begin 2 feet before the first sensor at 10 Hz, the vehicle suspension would oscillate 4 times before the last sensor was crossed. From the data previously observed, such axle oscillations were never shown to produce high amplitudes for this number of oscillations. Thus, the last sensor could be read directly with a reasonable estimation of static weight assuming body vibrations were negligible. Note that this assumes each sensor is not a source of perturbation and that the slab which they are installed into is

perfectly flat. For actual sensor installation, these assumptions are good. For the experiments presented here this would not be true as sensor installation was relatively crude in comparison to industry standards.

Since this technology would most likely be used on highways, higher speeds are more likely to be relevant. For intermediate values of 30-40 as seen on higher congestion roads, we notice that multiple readings are taken within a single period. This would allow us the use of the sin approximation method, the averaging method, or the least square error method.

For very high speeds of 80, or 90 mph, an averaging or least squares method would not be as robust a method to use because readings are taken near one local part of a single period. This could cause a large amount of inaccuracy depending on the amplitude of the load as seen by the sensor. For this reason, the sin approximation method would be a good candidate for use. Given the velocities for which experiments were conducted, multiple methods may be used at most speeds. Thus, it will be attempted in this paper to use multiple methods at each velocity and discover which is best. In addition, estimator combinations will be also tried to see if a more accurate result can be found. Finally, it should be noted that all plots of weight in time as well as the application of certain estimation methods at different velocities assume that a 15 Hz signal will be observed by our system and that the points as recorded by each sensor will not deviate from this waveform to any great degree.

VIII. WIM MEASUREMENT RESULTS WITH SUSPENSION VIBRATIONS

A. Original Data

After extensive testing, many data points were obtained at all velocities. These voltage readings represent the maximum load which the contact patch of the vehicle was able to transfer to the sensor as it passed over the sensor. This weight was stored as a voltage which was proportional to the vehicle weight. For each run, a single data point was obtained. The original data is shown in Figures 28 and 29 where the known static weight of each half axle is plotted versus the voltage produced from the run. Both testing days are shown.

Figure 28: Sensor 1 data from 3-16-2011.

The accuracy of the readings from 3-16-2011 averaged about ±15% and the readings from 4-9 averaged about ±40-45% at the higher weights and around ±60% for lower weights. The 3-16 data appears more linear and a tighter spread where as data from 4-9 has a spread which is much larger and as a result the data appears to show a poor linear fit with a low accuracy.

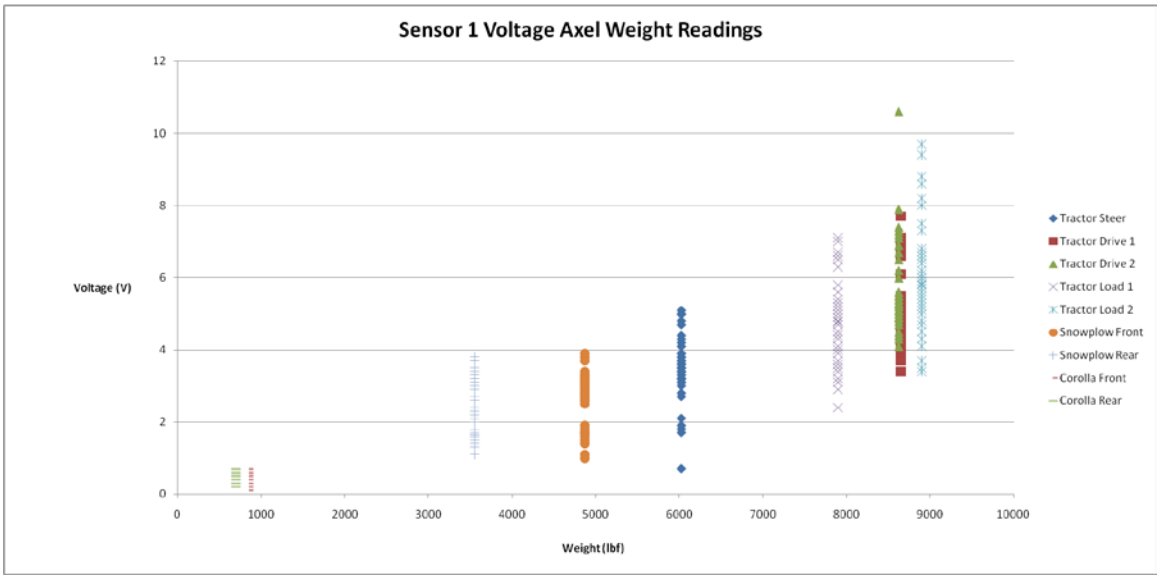

Figure 29: Sensor 1 data from 4-9-2011.

The difference in data is most likely a result of the change in testing variables in the system which changed from the first to second testing day. The data from 3-16 has fewer points, taken at lower speeds, on a sensor which was newer and mounted flush with the road surface (within 1/8 of an inch). The later data from 4-9 shows more runs, at more speeds and at higher speeds. The sensor was also older and not installed to the same height as the previous test day. It was roughly ¼ inch from the road surface to the top surface of the sensor. This height difference between days effectively introduced a greater perturbation into the measurement system by the measurement system itself. At high velocities, the data points were higher than for data at lower speeds, effectively increasing the spread as faster runs were included to the set.

The age of the slot in the road also had a large effect on the height of the sensor as it was loaded and unloaded over time. The sensor is supported by four bolt heads which are about ¾ of an inch in diameter each. For heavier vehicle axles this pressure on the concrete could be as high as 5000 psi. When the feet of the sensor apply this much pressure to the same location of concrete cyclically, the feet work their way into the concrete base which supports them over time. Over the full duration of all testing, these holes increased an estimated 1/8 inch in depth.

Given the design of the sensor and its adjustability, the height reduction could be adjusted for. During actual testing, however, the height changes observed were subtle and went unnoticed until later in the testing. For consistency and because it was assumed that the height decrease would not have a large effect, they were not adjusted when discovered. In hindsight, this was found to be an erroneous assumption. Regardless, the bolts used to adjust the sensor height were too short to increase the sensor height any further. In the future, this could be corrected for by ordering longer bolts and adjusting the sensors back to the original height.

Over the course of both days of testing the height reduced 1/8 inch. The lowest point on the top surface of the sensor started at 1/8 inch below the surface and at a certain point during testing on 3-16-2011 no more adjustability was available to maintain the sensor height. On 4-9, this

problem continued as the bolts were not long enough to raise the sensor to a height which was maintained on 3-16-2011.

In addition to observing differences from day to day, it can be observed that a difference exists from sensor to sensor. There are many possible sources of sensor output variation which fall into two categories. The first deals with the installation of the sensor and the second with the construction of the sensors themselves.

For sensor construction, variation in support beam and piezo geometry will cause the most noticeable changes in sensor output. First, the geometry of the support beams is considered. Based on the print tolerances used to make the prototypes, the dimensions of length, width and height may vary up to .010 inches. The resulting changes of these parameters would mean a change in beam stress at the surface. This variation coupled with a possible variation in piezo placement of up to .1 inches could cause an estimated worst case variation of up to 30% in piezo output.

Another potential source of variation is in the piezo area. Since this area is important to the amount of scaling that occurs with the voltage output across the storage capacitor, this should be considered. With this variation estimated at 10%, the scaled output could also change by this amount. Overall, all sources of geometric variation could cause considerable variation in the output from sensor to sensor.

The second main source of variation lies in the position of the sensor top relative to the road surface. From sensor to sensor it is estimated that the variation could be as much as ¼ from the lowest point on one sensor to the highest point on the next sensor. This variation is made worse because although each sensor was leveled with the slot that it was placed into, these slots most likely varied 1/8 inch in height. This variation between sensors would also cause changes in the impact load from one sensor to the next, resulting in different sensor outputs.

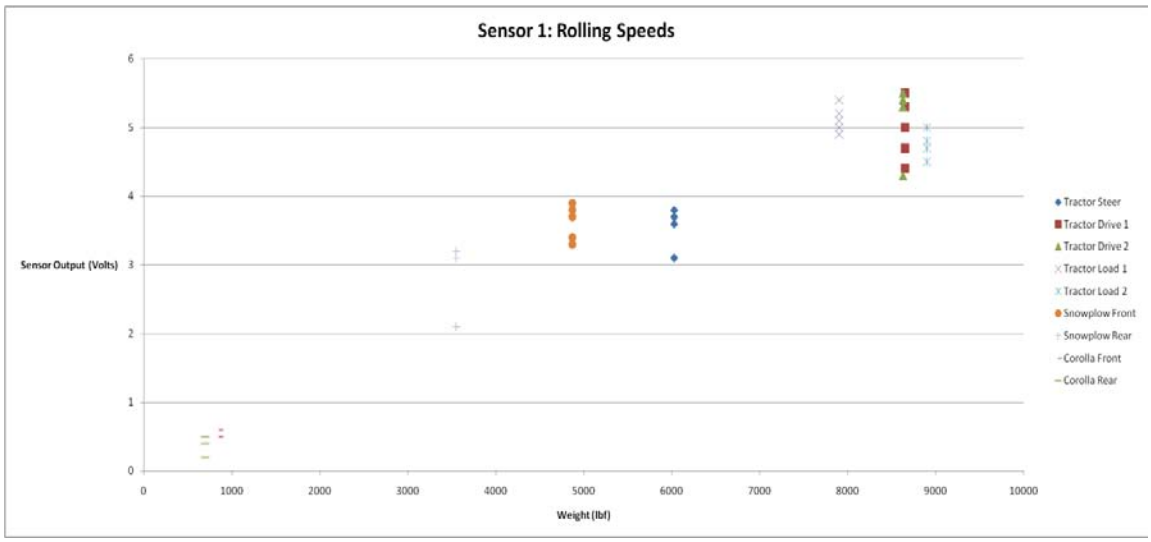

Figure 30: Sensor 1 data at 10 mph on 4-9-2011.

Both sources of variation could cause significant differences in sensor output from one to the next. An example of this is shown in Figures 15 and 30. These Figures were produced from the same set of testing conditions, but with different sensors.

The plots shown for the original data are linear and representative of the maximum weight observed by the sensor for the duration during which the vehicle moves over it. As discussed, the height difference between the sensor surface and the road surface is a significant source of perturbation which exacerbates the deviation of each reading from the static weight. This problem is made worse at higher velocities. It is also shown that the variation from sensor to sensor causes differences in the output for the same testing condition. Despite all this, it is desired to have a sensor output which shows the vehicle's static weight. If the current output is to be taken as a static weight reading, it is not desirable given that the accuracy can be as much ±45% as previously stated. To change the output of the system of four sensors to a static reading and not a dynamic one, different methods are explored here to convert the four readings obtained into a single reading of static weight.

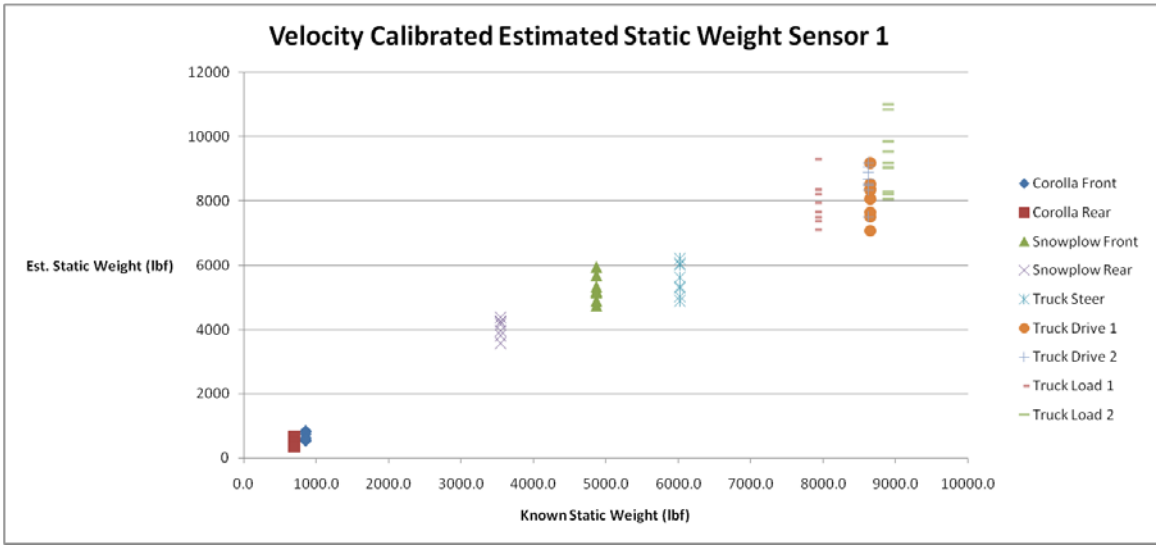

Figure 31: Speed calibrated data for sensor 1 on 3-16.

B. Sensor Output Correction

1) Speed Calibration and Velocity Dependence

It is known that for a given size perturbation in the road, a higher speed will result in vehicle dynamics with a higher amplitude of oscillation. In order to compensate for this, separate calibration lines were plotted for each sensor for each speed. Since the readings are most likely sensor specific and velocity specific, it is the intent of this method to compensate for these dependencies and reduce the spread in the data. The result should then be an output which more closely resembles the static weight of a vehicle.

In Figure 31, all voltage readings are plotted versus the known static weight for sensor one at 10 mph. It can be observed here that the data looks to have a spread which is significantly reduced from Figure 29, which shows all speeds.

Though less data points are taken here, the spread is very tight suggesting that for a given speed for a given sensor, the results are consistent. This shows that statements of the previous paragraph seem valid and the method feasible.

In order to calibrate the data, each sensor's data was plotted in subsets grouped by speed. A linear regression was performed on the data to obtain the best fit equation for the line. This

Table 3: Linear Regression Parameters

Multiple R	R Square	Adjusted R Square	Standard Error	Observations
0.976813219	0.954164064	0.953072732	0.417642588	44

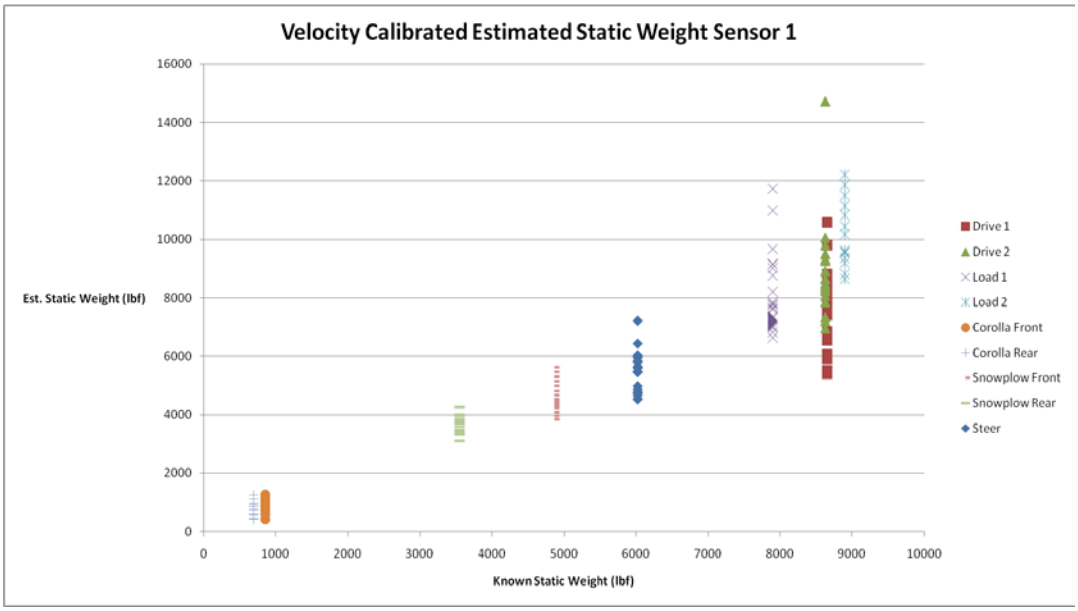

Figure 32: Speed calibrated data for sensor 1 on 4-9.

line was then used to obtain a weight value for each voltage value. The linear regression for Figure 15 is shown in table 3.The result of this calibration method is shown in Figure 32 for 4-9 and Figure 31 for 3-16. The result is that the accuracy for the sensor is improved from ±35% to around ±20% for the data on 4-9 at higher speeds. For 3-16 the calibration seems to improve the accuracy from ±15% to ±13%, if at all. The fact that calibrating the data in this manner helped the accuracy suggests that a velocity dependence exists, but that not all the spread shown on the original graph is because of speed dependencies.

Even though the data plotted for a given speed has a tighter spread than the all points plotted together, this trend reduced with increased speed. At 10 mph the spread of data is tighter than for 50 mph. This follows from the fact that at higher speeds the vehicle dynamics are larger in

amplitude but that the sensors do not necessarily capture the peak values (see Figures 25-27). If we look at the data for different velocities for sensor one, we see this behavior in the form of higher readings and increased spread. Appendix F shows sensor one calibration data plotted for variouos speeds as well as sensor two. It should be noted that sensors three and four also

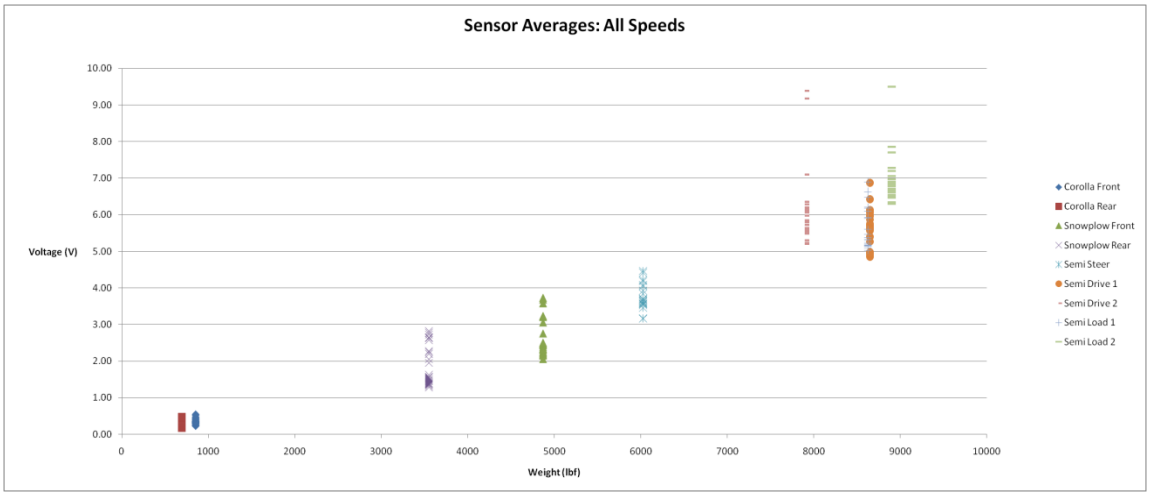

Figure 33: Averaging method applied to all data.

showed this behavior. Thus, the calibration breaks down at higher speeds and though the resulting data is more accurate, more needs to be attempted to improve the data.

2) Averaging

Averaging works well if all sensors used to capture the dynamic vehicle weight are able to capture points both below and above the static weight as the vehicle oscillates. These points should also be balanced in equal number below and above. This could be accomplished by large numbers of sensors but since this system only has four, this effectiveness of this algorithm will be limited. In addition, this method of estimation should only begin to work noticeably well at velocities in excess of 40 mph assuming the previous sustained oscillatory behavior. The larger the speed from which the data comes, the more accurate this method of estimation should be. Since the collected data only goes up to 50 mph due to sensor instability for heavy vehicles above this speed, it is possible that the benefit to averaging will be less apparent.

Appendix G shows the averages of all sensors for each speed. The result is that at lower speeds the averaging is not needed but once applied it does not make the data worse. For higher speeds, we see that the spread seems to improve or worsen depending on which sensor it is being compared to. Since sensor two has data which is very accurate, the averaging method seems to be worse, but for sensor 4, this data seems to be much better. As far as velocity dependence, the averaging method follows the trend such that higher speeds result in a greater spread.

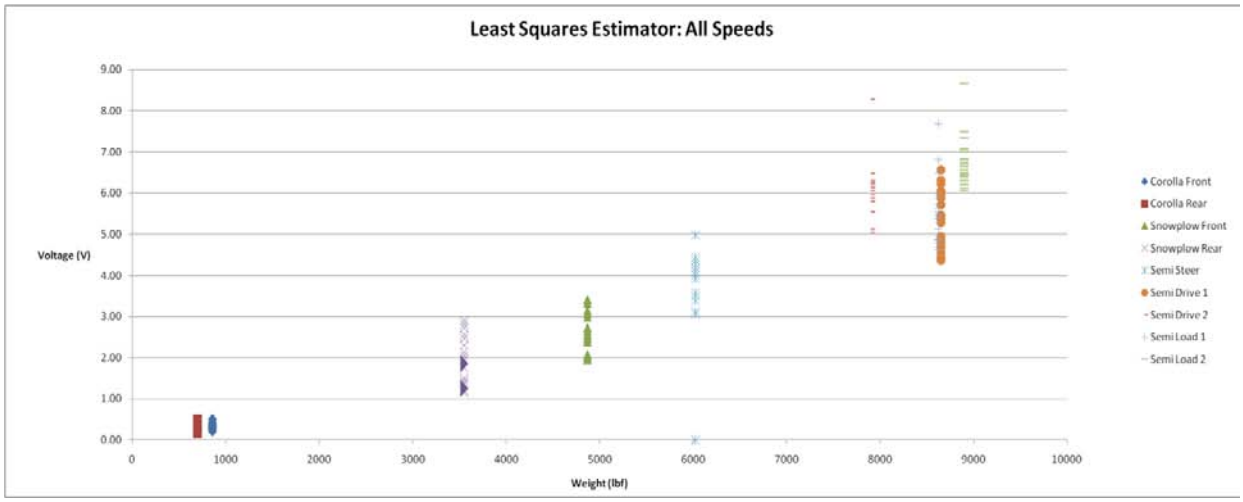

Figure 34: Least squares estimator results.

In Figure 33 all averaged data is shown together. Here the accuracy appears to be ±40% for higher weights (including outliers from 50 mph runs) and for lower weights around ±40-50%. Thus, this method does not help range of data much but it does help variance as the data seems tightly grouped.

3) Least Squares Estimator

Another method similar to averaging is the least squares error estimation method which will be referred to as the least squares method. This method iterates through the range of the output values of the sensors from the smallest to the largest value on any given run. The value which is found to have the lowest sum of the square errors between itself and each weight reading is the resulting estimator.

Unlike averaging, this method would favor outliers, especially for large data sets. For only four values, however this effect would not be obvious and the result would be similar to an average value. Regardless, this method is attempted here. The speed dependency of the estimator is shown in Appendix H. The result is very similar to averaging. For slow speeds the data is already good and a least squares estimator does not do much. For higher speeds the method is as relatively helpful as the single sensor which it is being compared to. The overall effect, though, is positive and can be seen in Figure 34.

The lower speeds have accuracies of ±50% and the higher speeds of ±20-30%. Thus, this method seems to be an improvement over averaging. Overall it is better than the raw data for larger vehicles.

4) Sine Approximation

The final method of static weight estimation involved fitting a sine wave to the data points from each run. It was previously shown that the vehicle dynamics follow a sin wave of a given frequency not only theoretically, but experimentally as well. Since this was found to be anywhere from 10- 20 Hz, a sine wave was fit to the data for 15Hz. For any set of four readings

from a given run, it was attempted to reduce the error between a sin wave (see equation 8) and the four sensor outputs.

Multiple iterations were tried until a best fit wave was found for the data points. In the equation, 'A' is the static weight of the vehicle, 'B' the amplitude of the sine wave, and Ø the phase shift. Phi was chosen to range from frequencies of 0-2*pi radians. Both 'B' and 'A' parameter ranges were chosen based on the readings observed from the four sensors. 'A' was ranged from half of the maximum weight reading for a given run to double this maximum reading. 'B' was chosen to range from 0 to half the maximum weight observed on the four sensors. This contingency method was chosen to maintain a good resolution for the iterated parameters. These ranges were picked based on the rule of thumb that peak oscillations would not exceed twice the value of static weight.

At higher speeds, 20-50 mph, the results were an improvement when compared to original data. A result is shown in Figure 35 for the snowplow test vehicle at 30 mph. At 10 mph, the result was not as good as shown in Figure 36 for the front snowplow axle. Overall however, the result was ±55% for lower weights and ±40% for the larger axles. This method was a minimal improvement over and above the raw data.

In light of the results, it should be noted that this sine approximation method assumes many things in order to be an effective method of estimation. First, the amplitude of the sine wave is a constant in the aforementioned equation. This is true if all data points are captured within one period or if repeating road vibrations in the pavement cause repeating and sustained

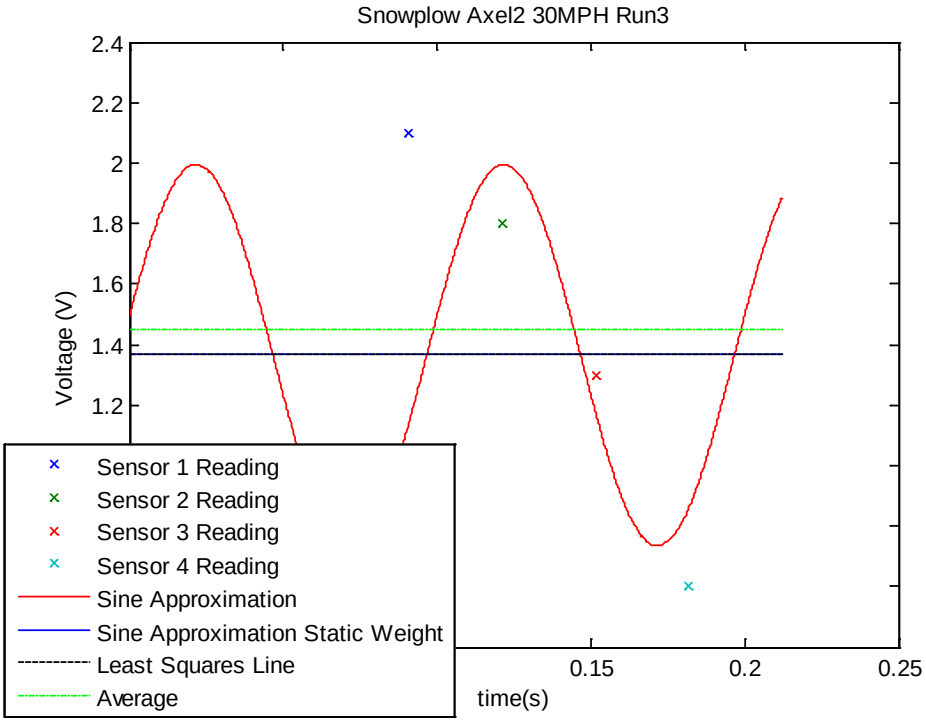

Figure 35: Sine wave fit to the sensor readings for the snowplow drive axle at 30 mph.

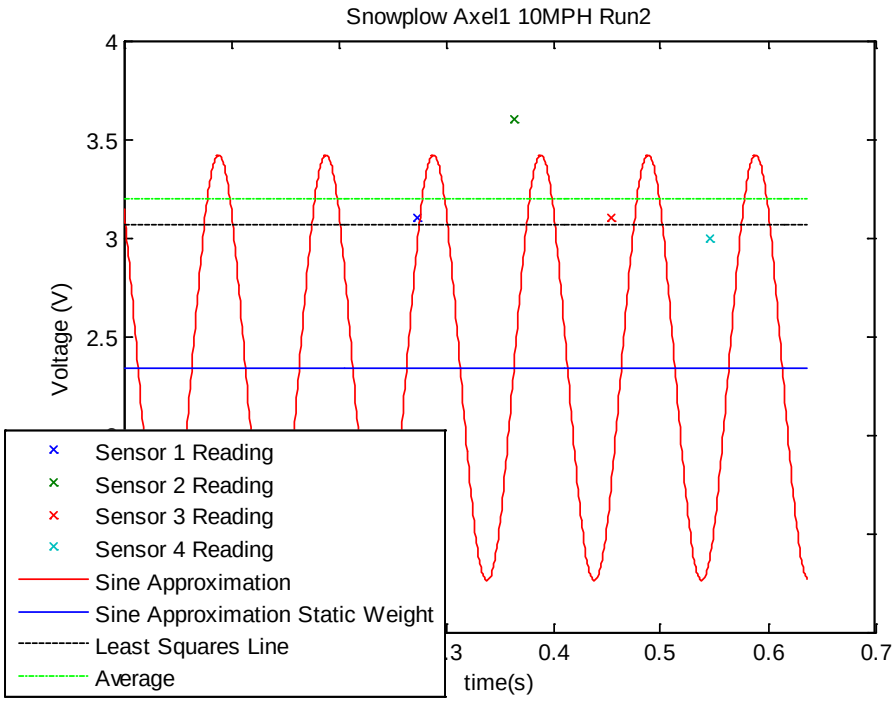

Figure 36: Sine wave fit to the sensor readings for the snowplow steer axle at 10 mph.

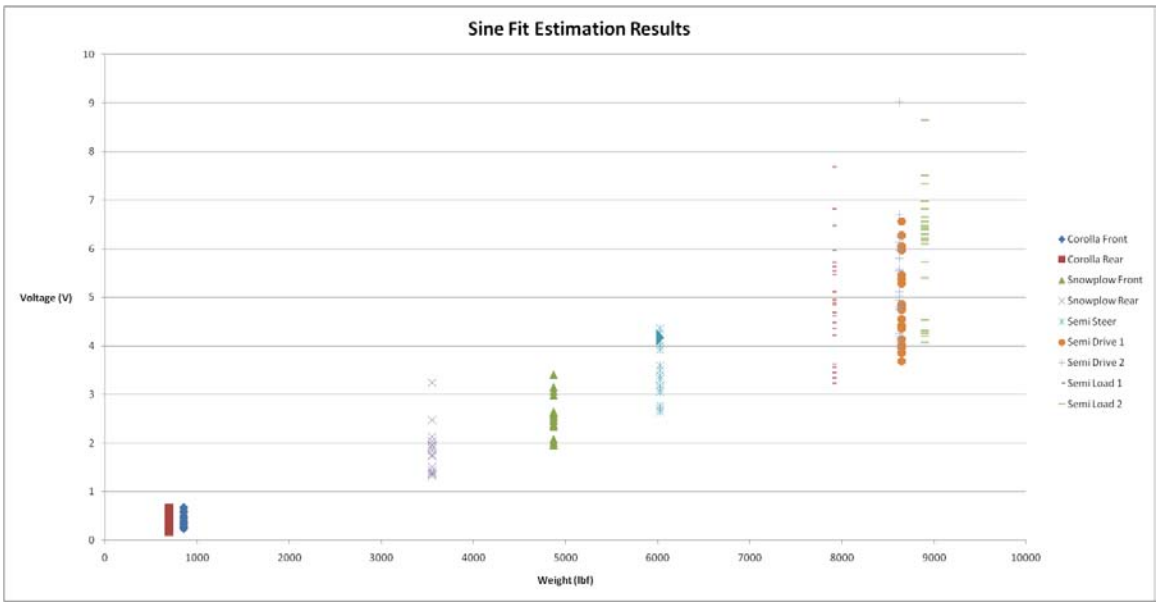

Figure 37: Sine approximation results for all vehicles at all speeds.

dynamics over the course of the measurement. For our measurement system and the road profile, this assumption is a reasonable one as shown in Figure 20.

Second, this method assumes that our experimental data does not deviate from the sine wave model. If we look at Figure 20 we see that our data fits a sinusoid beautifully. However, if we look at Figure 18 with a higher sample rate, we notice that the data does deviate from the fit.

The degree to which the noise affects the sensor reading would depend on the velocity of the vehicle. By looking at the aforementioned plots it is probable that this noise may or may not affect the sensor reading. In the future more accelerometer data should be taken at a higher sampling rate to better understand the vehicle vibrations. If large variations are found then for higher speeds the sinusoid model will be less accurate as the maximum reading of the sensor could be influenced.

A final assumption is that a single perturbation is the cause of the oscillation of the vehicle and that the sensors merely capture this. This is not true in our system. The greater the care to mount each sensor in the road that there is, the better this sinusoid model will be. Since the majority of the data presented in this paper deals with sensors that do cause vehicle vibrations then the sin model would be less accurate. In light of this, Figure 20 shows a very sinusoidal form which would lead the author to believe that this fit would not be so bad after all as long as

Table 4: Estimator Average and Lowest Value Combination Rankings

Estimator Combination Evauator(all speeds)?											
Options	C1	C2	S2	S1	T1	T2	T3	T4	T5	Observed Linearity	Counts
A-SAL				A-SAL	A-SAL	A-SAL		A-SAL		A-SAL	4
A-SA		A-SA			A-SA	A-SA	A-SA			A-SA	3
A-AL						A-AL	A-AL		A-AL		3
A-SL					A-SL	A-SL	A-SL			A-SL	3
L-SAL										L-SAL	1
L-AL					L-AL	L-AL	L-AL	L-AL	L-AL		5
L-SL										**L-SL**	1
L-SA			L-SA	L-SA						L-SA	3

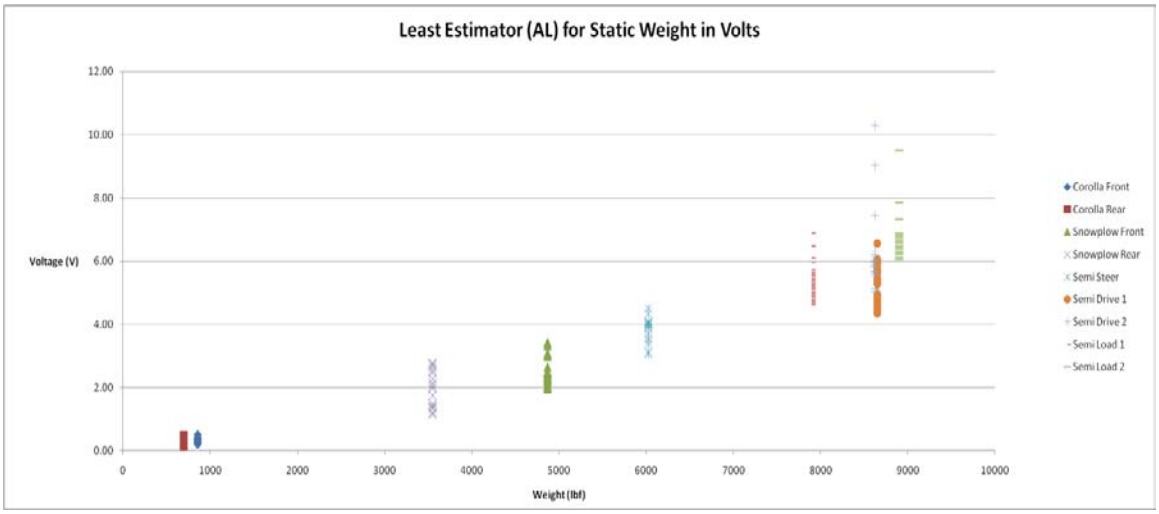

Figure 38: Lowest value used of least squares and averaging estimators.

the parameters are iterated to find a best fit. Thus, the sinusoid estimator is used here as shown in Figure 37.

5) Combining Estimators

In order to improve the accuracy of the measurement of static weight, multiple estimators were combined in multiple configurations. First, averages were taken of different estimators which

consisted of the sine approximation method (S), averaging method (A), and the Least Squares Method (L) as previously discussed. The resulting combination averages were SAL,SA,SL, and AL. In addition to averaging, a lowest value method was implemented where of the estimators considered in the combination, the lowest was taken from the set. For this method the same combinations were tried.

The best estimator combination was found using a ranking table where all combinations were equally ranked and the most useful combination was recorded. For each combination tried, it was noted how many times the method resulted in a reduction in the variation of the spread for an axle in relation to the other methods. In addition, it was considered how linear the estimation method was. Table 4 shows the results.

From the table, the best method is the lowest value taken when only considering the average and least squares estimators. The resulting plot can be found in Figure 38. Here it is observed that at lower weights the accuracy is ±40-50% and at higher weights the accuracy is ±20-30%. Thus the results are similar to the least square estimator.

C. Reflection on Results

1) Instability of High Speeds

Up to this point the best performance comes from the least squares method. It results in an increase of output accuracy from ±45% to ±25% at high weights and ±60% to ±50% at low weights. Thus, there is a large relative improvement in the sensor performance. Despite this finding, the accuracy of this system is worse than that of current WIM piezo systems. Thus, it would be ideal to improve the sensor performance to compete with these systems. In order to do this, the data will be evaluated more closely.

The first thing evident when looking at the estimation methods presented here is that often the data will be grouped closely and only a few outliers cause the overall accuracy to suffer. Without these values our system performance would improve greatly. As mentioned previously, higher weights often caused instabilities in the sensor system. This was especially true at 50 mph, where the semi truck repeatedly caused the sensor to leave the ground by about ½ inch or more after the vehicle passed over it. As a result of this, the measurement system behavior was quite different and unstable.

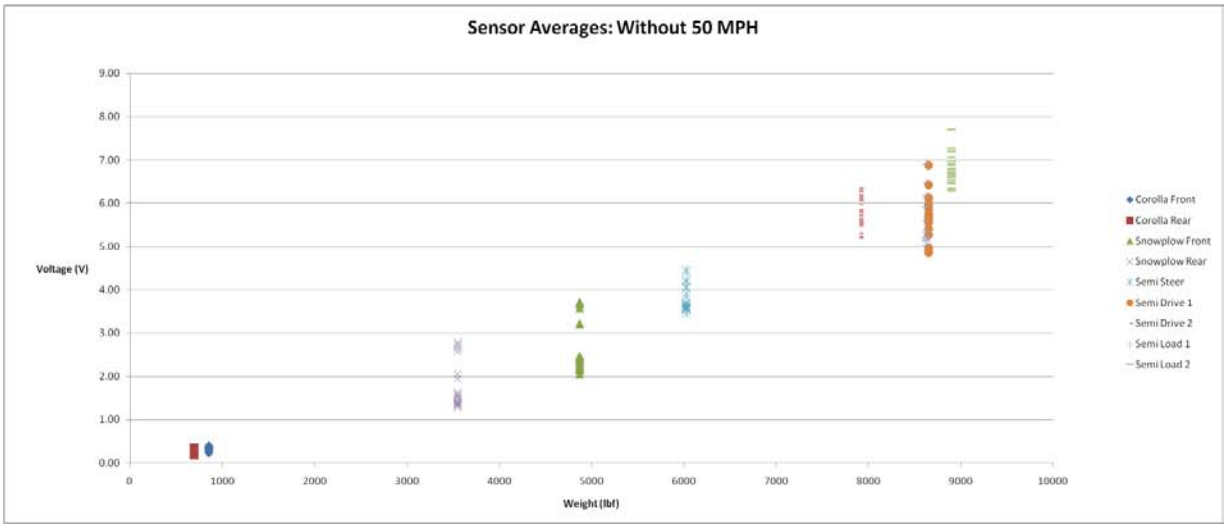

Figure 39: Sensor averages excluding 50 mph data.

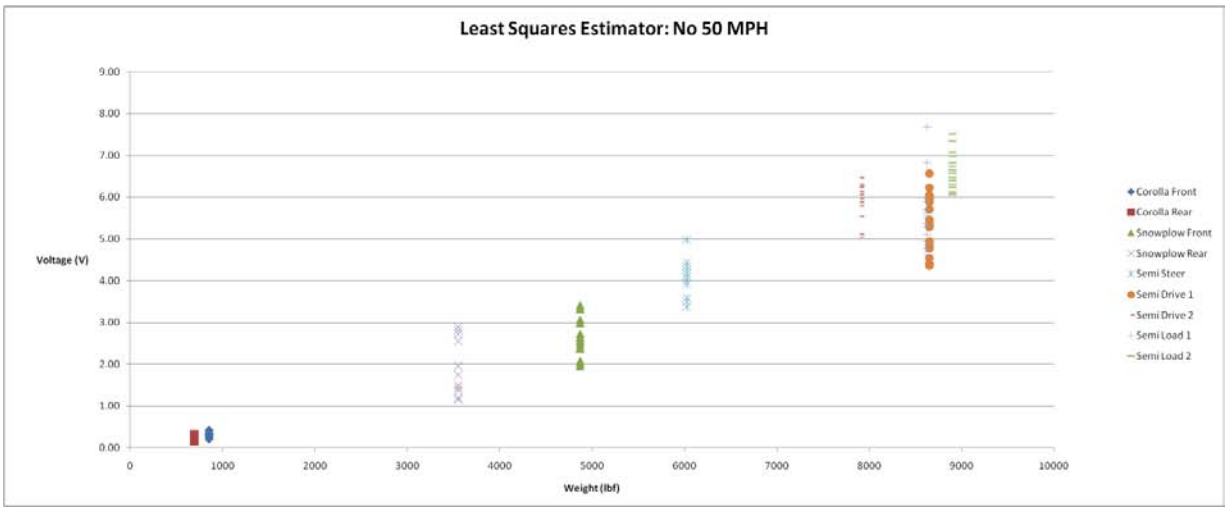

Figure 40: Sensor least squares estimator excluding 50 mph data.

When the data points were removed from the plots corresponding to these conditions the sensor performance dramatically increased. For the raw data there was little to no difference as shown in Appendix I. For the averages, however, a large improvement is shown (Figure 39). Low weights now show ±40% and high weights ±15%. For the least squares estimator improvements are also present with lower weights being ±50% and higher weights ±25%. The sin approximation method now is ±40% at lower weights and ±30% at higher weights. Finally, the lowest selection of the average and least squares estimators resulted in ±45% for low weights and ±15% for high weights.

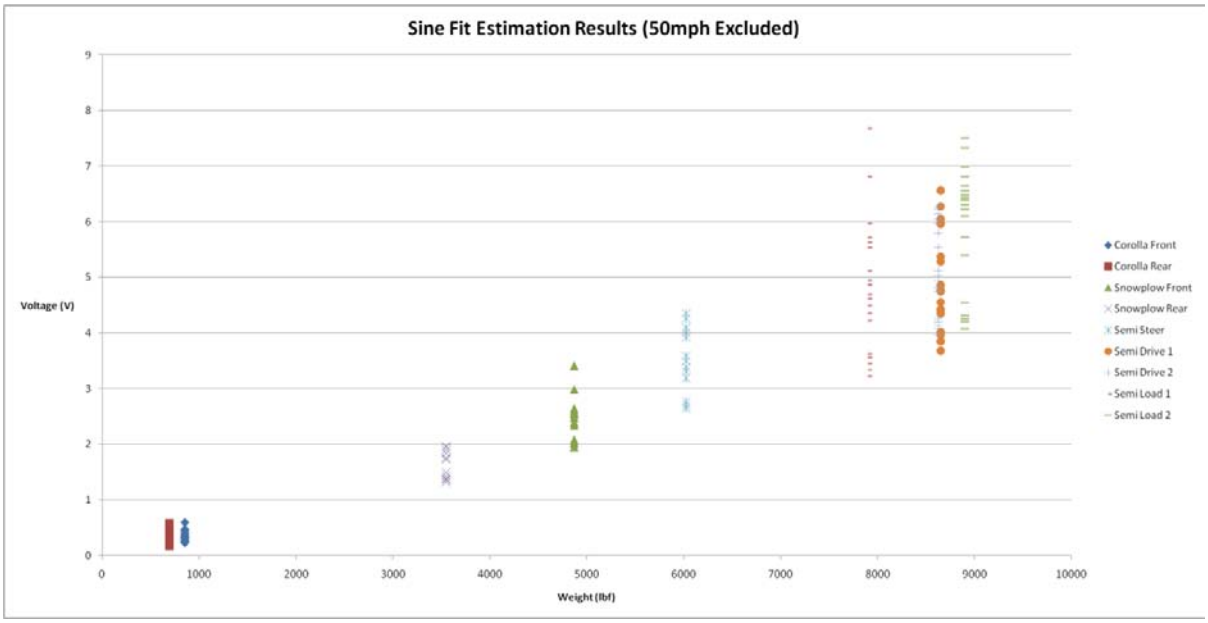

Figure 41: Sine wave estimator excluding 50 mph data.

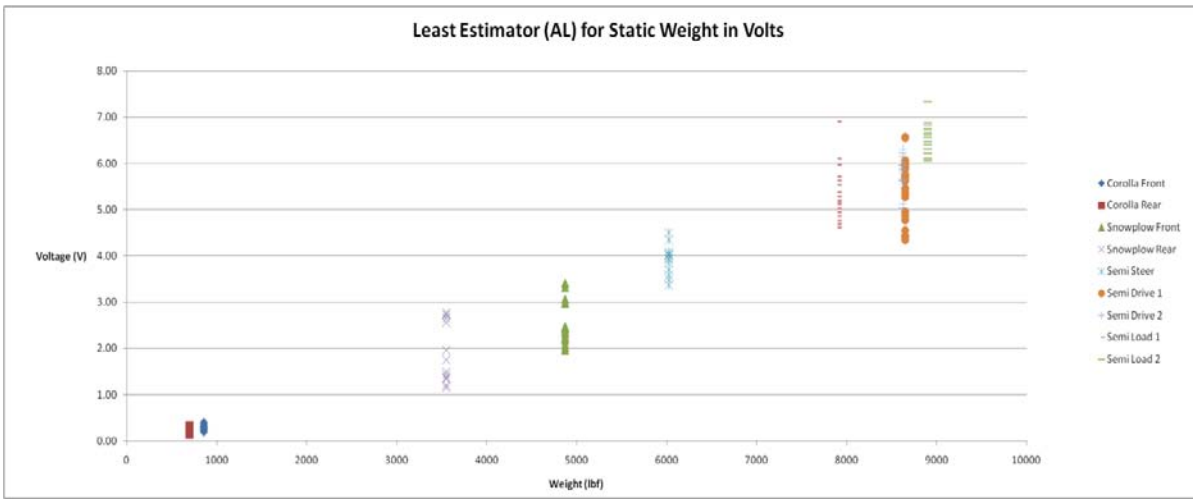

Figure 42: Lowest estimator for averaging and least squares methods excluding 50 mph data.

Observing the new results without 50 mph data shows a large improvement, with the best performance being shown by the straight average. This brings the raw data from ±65% to ±40% for low weights and from ±45% to ±15% for higher weights. Again, this improvement method is valid given the unstable system behavior shown for this speed only. Mechanical modifications to the system in terms of damping or simply securing the sensor to the concrete subbase would allow for the use of the system at 50 mph and above.

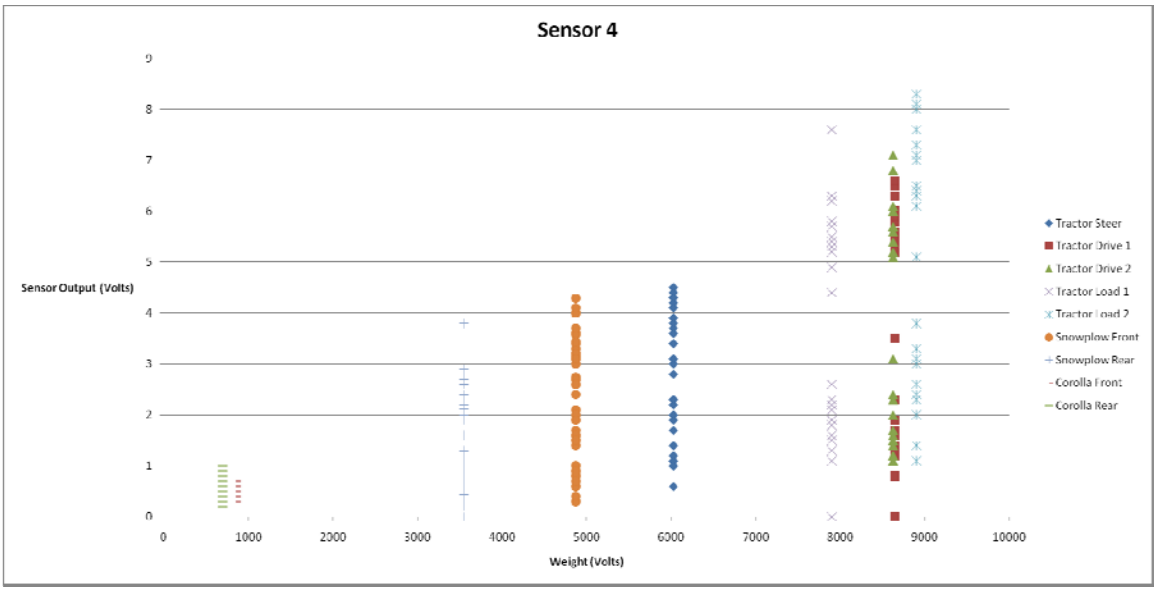

Figure 43: Sensor 4 raw data.

Overall, the results of the system are reasonable given that this system's primary use is to detect overloaded trucks. A potential customer would most likely not care about the information on smaller vehicles, and so the accuracy would be around ±15%. This is reasonable since other piezoelectric sensors have reported similar accuracies (see table 1).

2) Disregarding Sensor 4

Another observation that can be made while looking at the data is that sensor 4 was much less linear and accurate when compared to the other sensors. This data, as shown in Figure 43, has a much larger spread of data when compared to the other sensors involved (see Figure 29). For many runs, this data was disregarded all together because there was a bad connection and the output was scaled too small. After this condition was discovered and fixed, however, the data was used in the estimation of static weight.

Since sensor 4 data appeared so much worse it was removed from all estimation calculations to see if a better static weight could be gained. The result was not as good as anticipated. It turned out that since the average value of sensor 4 was lower than the other sensors, removing this data from the estimation methods improved accuracy for lower vehicle weights to around ±10-15% for the combined estimation method of least squares and averaging. However, for heavier weights where the vehicle spread was large, using sensor four to bring down the estimator was a plus. Without it, the spread increased around ±5-10 %(see Appendix J). For the averaging and least square estimation method individually, the outcome did not change considerably. Thus, removing sensor 4 data from the estimation calculations did not help with the accuracy overall.

3) Sensor Height

As observed during testing, the data from 4-9 was gathered from sensors which were mounted less flush with the pavement surface when compared to the data from 3-16. In order to see if the

mounting of the sensors was a cause of estimation accuracy, we will compare similar speeds from these data sets in the range of 5-30 mph for 3-16 and 10-30 mph for 4-9 for the same sensors. Appendix K shows all major presentations of the data for 3-16 and Appendix L for 4-9. This includes the raw data for sensor one as well as the velocity calibration, sin approximation, least squares, and averaging methods.

Observing the data from the first testing day, it is clear that the accuracy is much better than the second day. The velocity calibration method is the only high weight data type more accurate on 4-9 when compared to 3-16. With all other plots the accuracy is worse or the same. For the low weights the data from 4-19 is less accurate on all accounts. Though there are many variables involved in the testing it is reasonable to believe that the sensor height is the cause of this reduction in accuracy because it was the most prominent observed change. Thus, it is shown by the data that an 1/8" maladjustment is most likely the cause of the added vehicle dynamics, and the reduction in accuracy. For future testing it must be ensured that the sensor is not the cause of any external oscillation.

Table 5: Approximate Accuracy of Data from 3-16 for 5-30 mph

Data Type	Low Weights ± Accuracy	High Weight ± Accuracy
Original	20	20
Calibration	15	30
Sine Estimation	15	15
Averaging	30	20
Least Squares	20	15

Table 6: Approximate Accuracy of Data from 4-9 for 10-30 mph

Data Type	Low Weights ± Accuracy	High Weight ± Accuracy
Original	20	40
Calibration	20	20
Sine Estimation	15	40
Averaging	30	20
Least Squares	50	20

IX. FUTURE WORK

A. Design Changes

Since the work done on the sensors discussed here mainly dealt with determining the capability of the current sensor system, there is much more which could be done to improve it. First, the accuracy of the sensor could be reduced by identifying key areas for variation from one sensor to another. A study could be done to the current design's sensitivity to see which variables most affect the sensor output. Improvements could then be made to the most critical features such as locating piezo patches during installation, reducing the possibility of loosening bolts, or preventing the sensor components from moving during installation. The damping of the sensor could be reduced, both to improve the desired output for static weight, but also to anchor the sensor for higher speed runs to allow for heavier vehicles to be measured at all speeds rather than just up to 40 miles per hour. This solution might take the form of bolted studs holding the sensor into the concrete below, preventing it from jumping out of the slot. Also, the system could be characterized to better understand its susceptibility to vibration.

Another area for improvement is to calibrate each sensor before installation using known static weights. This would reveal the actual voltage versus force relationship for the system. This could provide more insight into what was actually happening with the dynamics of the vehicle as it passes over the sensor. For more testing, an accelerometer could be also added to the test vehicle as well as the sensor itself to observe what parts of the overall system were causing vibrations. Furthermore, readings could be taken at many locations on the sensor to ensure that no variation exists with lateral vehicle position as it drives over the sensor.

B. Installation Changes

Finally, more care could be taken to install the sensor into the road once all the aforementioned items have been improved upon. The height could be fine-tuned to be flush with the pavement and the pavement flush with the sensor. Also, the section of concrete under the feet could be reinforced with a metal plate to ensure that the sensor does not creep down into the pavement. Finally, a casing could be made to ensure that the sensor was able to weather the elements.
In all, the sensor could be improved by identifying key areas of variation in the design and removing them. After this has been completed the system could be characterized and calibrated so that it was more fully understood experimentally before testing. Following, the sensor could then be tested again to see if accuracy improvements resulted from these changes.
The aforementioned items all deal with gaining a more accurate static weight reading from a dynamic load. Once this has been improved as much as possible, previously proven technology could be added to make the design battery-less and wireless. Thus, there is much room for future work which has the possibility to greatly improve the design and its accuracy.

X. CONCLUSIONS

This project enhanced a previous battery-less wireless traffic flow sensor so that the sensor could provide weigh-in-motion (WIM) measurements and provide enhanced telemetry distance of 500 feet. The new sensor design utilized a 3-layer structure, consisting of a top frame, a WIM layer and an energy-harvesting layer. The energy-harvesting layer was designed to have a nonlinear structural response so that the sensor could harvest adequate energy from low vehicle loads, while at the same time not failing structurally for heavy vehicle loads. The WIM layer was designed to have an elastic structural response, so that voltage measurements on the WIM layer could provide a measure of vehicle weight.

The sensor's performance was evaluated by embedding it in a slot in concrete pavement and driving various vehicles over it at a number of different speeds. The sensor was found to meet the specification of 500 feet telemetry distance. It was able to provide weigh-in-motion measurements with an accuracy of better than 15% in the absence of vehicle suspension vibrations.

However, much of the WIM data obtained from the sensor after April 2011 was obtained in the presence of significant suspension vibrations. This is because the sensor could not be maintained flush with the road surface. The passing of the vehicle over the sensor lead to suspension vibrations.

The project evaluated the use of 4 consecutive sensors in the road to remove the influence of suspension vibrations. Three different algorithms (average of the 4 sensors, least squares weight calculation and least square biased sine fit) as well as a combination of these three algorithms were evaluated for static weight calculation in the presence of suspension vibrations.

Overall this system was found to be accurate to ±40% for smaller vehicles on the order of a few thousand pounds, and ±15% for heavy truck vehicles weighing 80,000 pounds, in the presence of suspension vibrations. Using the lowest value of the least square error method and average value method the best results were obtained. Overall, the sensor is effective at detecting the weights of heavy vehicles and less so when detecting light vehicles. Despite this, possible improvements could be made through manufacturing changes, better installation practices and system characterization. In addition, the design could be made more robust to better handle high speeds, withstand the elements year round, and not be a source of perturbation for passing vehicles in the road itself.

XI. COMPLETION OF PROJECT TASKS

The tasks in the project that summarize the major activities which were accomplished are as follows:

- Enhancement of sensor design so as to significantly improve telemetry distance and enable direct wireless communication to existing roadside cabinets.
- Experimental evaluation of enhanced sensor, measurement of maximum achieved telemetry distance and ability to communicate with a circuit board that can be placed in a roadside cabinet.
- Enhancement of sensor design to enable reliable determination of vehicle weight overloads (in addition to traffic flow rate and number of axles).
- Evaluation of weigh-in-motion ability of sensors and calibration by comparison with static vehicle weights.
- Evaluation of the use of a series of consecutive sensors for cancellation of the influence of dynamic vehicle vibrations. Development and implementation of algorithms for spatial filtering of data from consecutive sensors.
- Comparison of traffic flow rate and vehicle weight measurements provided by the new sensors with that from the current MnRoad sensors. Determination of accuracy and reliability of new sensors and statistical analysis of long-term test data.
- Preparation of comprehensive project report documenting technology developed, results obtained in the project and conclusions.

All of the above tasks have been completed. However, the comparisons in vehicle weight measurements for the new sensors were done with a range of vehicles which had been accurately measured at a weigh station, rather than with sensors embedded at MnRoad. Data was statistically compared using many tests with several vehicles over several days. However, the sensors were not permanently embedded into the road and hence variations with temperatures and other long-term trends were not analyzed. These effects will be analyzed in a new field study funded by the Minnesota Department of Transportation that will begin in Fall 2011.

REFERENCES

[1] Abelson, P. (2001, October). "Land Line." *Paul's Picks.* Retrieved July 12, 2011, from: http://www.landlinemag.com/Archives/2001/Oct2001/Features/pauls_picks.html.

[2] Gunderson, D. (2007, August 10). "Heavy trucks strain roads and bridges." *Minnesota Public Radio*. Retrieved July 12th, 2011, from: http://minnesota.publicradio.org/display/web/2007/08/10/bridgetruckweight/.

[3] (2006, September 2). "Automatic Control Laboratory." *System Modeling*. Retrieved July 12th, 2011, from: http://control.ee.ethz.ch/~ifa-fp/wiki/index.php?n=Main.Modelling.

[4] Rajamani, R. (2006). *Vehicle Dynamics and Control.* Verlag, New York: Springer.

[5] American Wood Council. (2005). *Beam Design Formulas with Shear and Moment Diagrams.* Washington, D.C.: American Forest and Paper Association, Inc.

[6] Beardmore, R. (2011, June 28). "Natural Frequencies to Traverse Vibrations." *RoyMech.* Retrieved July 12th, 2011, from: http://www.roymech.co.uk/Useful_Tables/Vibrations/Natural_Vibrations.html.

APPENDIX A: CALCULATIONS FOR SENSOR PILLAR AND FRAME GEOMETRY

The max load anticipated for the sensor is set by the heaviest vehicle used for testing. This was a Navistar tractor with an 80,000 pound configuration. The heaviest axle was 34,000 pounds. The axle was supported by four sets of tires. Thus, this number is divided by 4 to compensate for the max static weight that can be on the sensor. This amount is then multiplied by 2 to obtain the max dynamic weight which the vehicle will exert of the sensor. The result is below:

$$F_{max} = \frac{34{,}000 \text{ lbf}}{4} * 2 = 17{,}000 \text{ lbf}$$

Assuming a yield stress of 40 ksi and a safety factor of 2, the resulting design strength is to be no more than 20,000 psi.

First the frame geometry will be examined. The geometry of the frame is shown in Figure 44 below [5]. Thus, the calculation for applied bending stress in a pin-pin support with a distributed load is [5]:

$$\sigma = \frac{M * y}{I} = R_1 * \left(a + \frac{R_1}{2 * w}\right) * \frac{h * 12}{2 * b * h^3} = R_1 * \left(a + \frac{R_1}{2 * w}\right) * \frac{6}{b * h^2}$$

It should be noted that only normal stress is considered here as this would be the majority of the von-misses stress. Adding a shear stress into the calculation would gain no more than 10%. Thus, the safety factor will be sufficient to only account for shear stresses in these calculations.

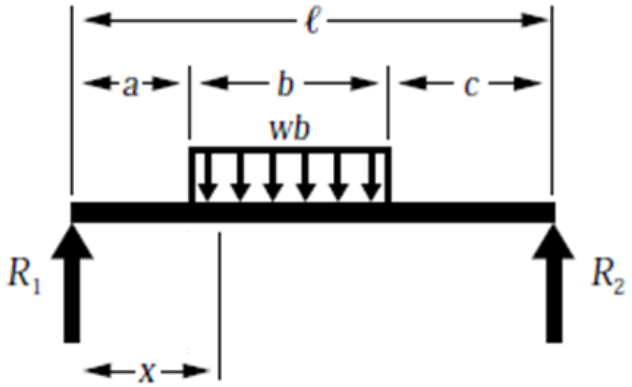

Figure 44: Diagram of a pin-pin support with distributed load.

Substituting the current beam parameters, the result is as follows:

$$\sigma = \frac{M * y}{I} = \frac{M * y}{I_o - I_i} = \frac{M * \frac{h_o}{2}}{\frac{b_o h_o^{\,3}}{12} - \frac{b_i h_i^{\,3}}{12}} = R_1 * \left(a + \frac{R_1}{2 * w}\right) * \frac{\frac{h_o}{2}}{\frac{b_o h_o^{\,3}}{12} - \frac{b_i h_i^{\,3}}{12}}$$

The variables used to solve the equation are as follows:

$b_o = 10\ in = outer\ width\ of\ sensor\ frame$
$b_i = 9.5\ in = inner\ width\ of\ sensor\ frame$
$h_o = 2\ in = outer\ height\ of\ sensor\ frame$
$h_i = 1.5\ in = inner\ height\ of\ sensor\ frame$

$$w = \frac{total\ load}{load\ distribution\ length} = \frac{17000\ lbf}{36\ in} = 472\ \frac{lbf}{in}$$

$$R_1 = \frac{17000\ lbf}{2} = 8500\ lbf$$

These values yield the following result:

$$\sigma = 27700 \frac{lbf}{in^2}$$

Thus, the worse case applied stress is slightly greater than the safety factor selected would allow for. Since the resulting safety factor was 1.5 and there were support blocks integrated into the design, the wall thickness of the rectangular tube considered here was deemed acceptable. To determine the deflection the system is treated at a point load at the center of the beam. This assumption yields the following results [5]:

$$\delta = \frac{F_{total} * L^3}{48 * E * I} = \frac{17000\ lbf * (46in)^3}{48 * 29 * 10^6 \frac{lbf}{in^2} * (\frac{10\ in * (2\ in)^3}{12} - \frac{9.5\ in * (1.5\ in)^3}{12})} = .22\ in$$

Since this assumption is very conservative, and the results are less than a tenth of an inch greater than our ideal design objective, the frame is assumed a viable geometry for both deflection and applied stress.

To calculate stress for the support pillars a fixed-fixed beam model is assumed. This is a reasonable assumption given that each end is bolted together and the length between supports is relatively small. Unlike the frame, only the applied stress will be calculated since the spacing between the upper and lower support beams is .25 inches and it is unlikely that such a beam would deflect such a large amount without permanently deforming.

The equation for an applied bending stress in a fixed-fixed with a point load is given by [5]:

$$\sigma = \frac{M * y}{I} = \frac{F_i * L * h * 12}{8 * 2 * b * h^3} = \frac{3 * F_{total} * L}{4 * b * h^2} = 24000 \frac{lbf}{in^2}$$

This calculation shows that the bending stress involved is just greater than a safety factor of two would allow. This is more than reasonable though since the full weight of the truck would never be supported by a single pillar alone on account of the geometry of the contact patch and the load distribution associated with it.

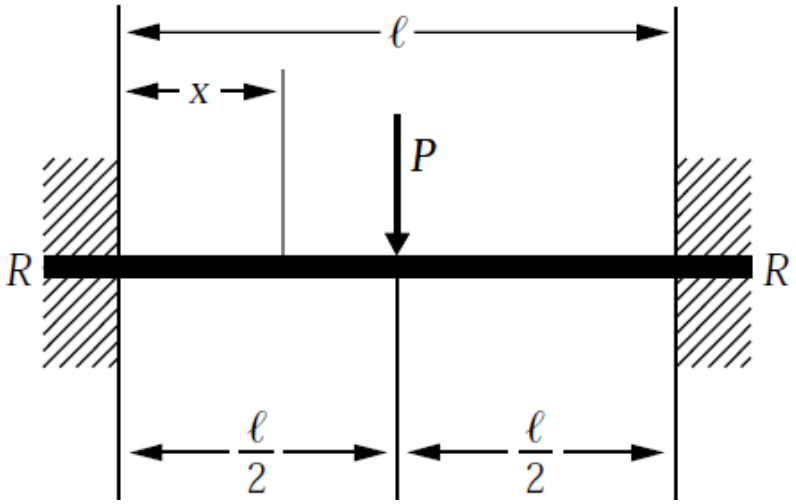

Figure 45: Diagram of a fixed-fixed support.

APPENDIX B: STRESS CALCULATIONS FOR LOAD POSITION INDEPENDENCE

Assuming the total load is on the sensor, it will be independent of position and given by:

$$F_T = F_1 + F_2 + F_3 + F_4$$

where F_i is the load supported by each pillar. F_i will be dependent on the position of F_{total}. The charge generated by each pillar is given by:

$$Q_i = C_i * V_i$$

And the resulting voltage is:

$$Vi = E * \varepsilon_i * g_{31} * t = \sigma * t * g_{31} = \frac{M_i * y * t * g_{31}}{I} = \frac{F_i}{8}(4 * x - l) * \frac{y * t * g_{31}}{I}$$

In the above equation, E is young's modulus of the beam, epsilon the strain at the location of the piezo, g_{31} a constant of the piezo, t the thickness of the piezo, M the average moment at the location of the piezo, y the vertical distance from the neutral axis to the piezo, x the distance from the left support to the center of the piezo, and l the length of the unsupported beam. The only variable to change with the location of the load is the force as previously mentioned. Since each support structure has the same geometry and piezo locations the previous equation can be reduced to:

$$V_i = F_i * k$$

where k is a constant. Hence,

$$Qi = C_i * F_i * k$$

Also, since all capacitances of the piezos on each beam are equal (C), the total voltage generated by the sensor is:

$$V_{total} = \frac{Q_{total}}{C_{total}} = \frac{(Q_1 + Q_2 + Q_3 + Q_4)}{(C_1 + C_2 + C_3 + C_4)} = \frac{k * (C_1 * F_1 + C_2 * F_2 + C_3 * F_3 + C_4 * F_4)}{4 * C}$$
$$= \frac{k * C(F_1 + F_2 + F_3 + F_4)}{4 * C}$$

$$V_{total} = \frac{k * F_{total}}{4} = K * F_{total}$$

Therefore, the total voltage output is independent of load position, and only dependant on the total load applied.

APPENDIX C: RESONANT FREQUENCY CALCULATIONS

In order to validate that no resonant peak exist in the range of frequencies that the sensor will see, many different natural frequencies are calculated for the current beam geometry. Before this however, the range of frequencies that the sensor could see must be identified. For this it is assumed that as the vehicle moves on and off the sensor, this represents half the period of a sin wave introduced to the sensor. Since the maximum speed the vehicle would travel would be 90 mph, and the sensor top is 10 inches in length, the range of operating frequencies are zero hertz to:

$$f_{observed} = \frac{90\ mi}{19\ in * hr} * \frac{5280\ ft}{1\ mi} * \frac{12\ in}{1\ ft} * \frac{1hr}{3600\ s} * \frac{\frac{1}{2} cycle}{frequency} = 42\ Hz$$

The first natural frequency to be considered in the beam is that of the first bending mode of the top wall of the frame. This can be modeled by a fixed-fixed support [6] as shown in Figure 46:

$$f_{resonant} = \frac{k}{2 * pi * L^2} * \sqrt{\frac{E * I}{m}} = \frac{k}{2 * pi * L^2} * \sqrt{\frac{E * b * h^3}{\rho * v * 12}}$$

$$= \frac{22.4}{2 * 3.1415 * (10in)^2} * \sqrt{\frac{30 * 10^6 psi * (.25in)^2}{\left(.284 \frac{lbf}{in^3}\right) * 10in * 12}} = 17Hz$$

The same calculation can be made for the support pillars as a fixed-fixed support (Figure 47):

$$f_r = \frac{k}{2 * pi * L^2} * \sqrt{\frac{E * I}{m}} = \frac{22.4}{2 * 3.1415 * (3.75in)^2} * \sqrt{\frac{30 * 10^6 psi * (1in)^2}{\left(.284 \frac{lbf}{in^3}\right) * 3.75\ in * 12}} = 388\ Hz$$

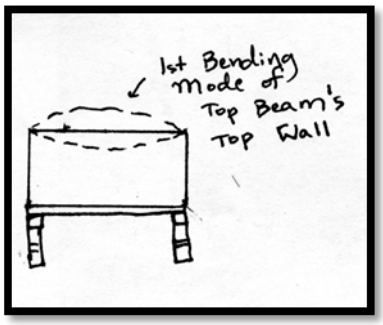

Figure 46: First bending mode of the top wall of the sensor frame.

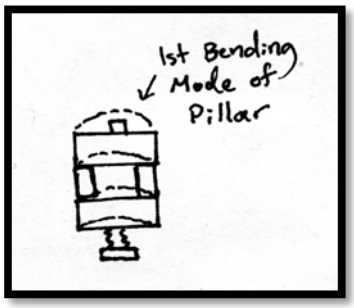

Figure 47: First bending mode of the pillar beams.

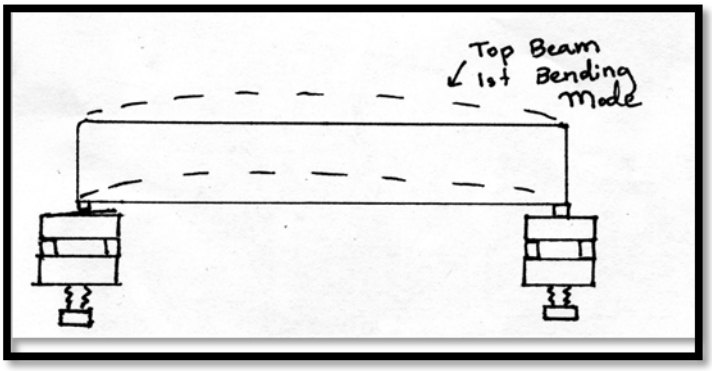

Figure 48: First .ending mode of the frame.

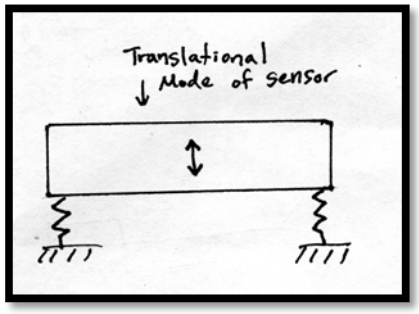

Figure 49: Sensor translation mode.

Another mode to be considered is the movement of the whole top beam in its first vibrational mode. This can be modeled as a pin pin[8] as shown in Figure 48:

$$f_r = \frac{k}{2 * pi * L^2} * \sqrt{\frac{E * (I_{out} - I_{in})}{m}}$$

$$= \frac{9.87}{2 * pi * (42\ in)^2} * \sqrt{\frac{30 * 10^6 psi * \left(\frac{10\ in * 2in^3}{12} - \frac{9.5\ in * 1.5in^3}{12}\right)}{\left(.284 \frac{lbf}{in^3}\right) * [(.5\ in * 2\ in * 42\ in) + (.5 * 9.5 * 42)]}} = 1.2\ Hz$$

The final frequency considered here is the translation mode of the beam where the frame is treated as a mass and the pillars are support spring (Figure 49). The following equation can be used to calculate this frequency:

$$f_r = \sqrt{\frac{k}{m}} = \sqrt{\frac{192 * E * I}{L^3 * m}} = \sqrt{\frac{192 * 30 * 10^6 psi * 1in * (1in)^3}{12 * (3.75\ in)^3 * \left(.284 \frac{lbf}{in^3}\right) * [(.5\ in * 2\ in * 42\ in) + (.5in * 9.5in * 42in)]}} = 364\ Hz$$

From these calculations we see that two vibrational modes have resonant frequencies which could potentially cause increased voltage readings from the piezo at 1 mph and 19 mph. At 1 mph, it is unlikely that this mode will have any effect on piezo output as the vehicle traveling over the sensor will not emulate a sin wave input at such a slow speed. At 19 mph, the input to the sensor will better represent a partial sine wave, however; it was uncertain whether this would affect the sensor output. After extensive testing, this speed showed no outputs which dramatically deviated from the outputs at other speeds. This is most likely due to the solid steel supports which added significant stiffness to the ends of the beams but which are not considered

in this equation. Thus, it is not likely that this resonant frequency or others would cause an issue with the sensor operation in the future.

APPENDIX D: CAPACITOR SELECTION FOR MICROCONTROLLER SATURATION AVOIDANCE

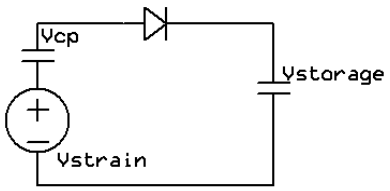

Figure 50: Piezo voltage storage circuit.

From the above Figure we see that:

$$V_{piezo} - V_f - V_{ad} = V_{strain} - V_{cp} - V_f - V_{storage} = 0$$

We can calculate the voltage generated by the piezo for any vehicle. It is done here for the heaviest weight observed during testing assuming an even distribution over all four pillars:

$$V_{strain} = E * \varepsilon * t * g_{31} = \sigma * t * g_{31} = \frac{M * y * t * g_{31}}{I} = \frac{-(F_{total}/4) * L * y * t * g_{31}}{16 * I}$$

$$= \frac{-\left(8500\,\frac{lbf}{4}\right) * 3.75in * 1in * .075\ in * -11.6 * 10^{-3}\,\frac{V * m}{N} * \frac{1N}{.225lbf} * \frac{1\ in}{.0254m} * 12}{16 * 1\ in * 1\ in^3 * 2} = 250\ V$$

Since the forward voltage drop for the typical diode is about 1 volt, this can be neglected from the equation resulting in:

$$V_{strain} - V_{cp} - V_{ad} = 0$$

This can then be manipulated in the following fashion to obtain the order of magnitude for the correct capacitor needed:

$$V_{strain} = V_{cp} + V_{storage}$$

Since both capacitors are in series, their charges are equal. Therefore:

$$Q_{cp} = C_{cp} * V_{cp} = Q_{storage} = C_{storage} * V_{storage}$$

Substituting into the voltage equation for Vstrain yields:

$$V_{strain} = \frac{C_{storage}}{C_{cp}} * V_{storage} + V_{storage} = (\frac{C_{storage}}{C_{cp}} + \frac{C_{cp}}{C_{cp}}) * V_{storage} = (\frac{C_{storage} + C_{cp}}{C_{cp}}) * V_{storage}$$

This may also be put in terms on Voltage across the storage capacitor:

$$V_{storage} = (\frac{C_{cp}}{C_{storage} + C_{cp}}) * V_{strain}$$

With the capacitance of the piezos calculated from the ratios of the applied charge to applied voltage:

$$C_{cp} = \frac{d_{31} * L * W}{g_{31} * T} = \frac{-190\ \mathrm{x}\ 10^{-12}\,\frac{\mathrm{m}}{\mathrm{V}} * 2 * 1\ \mathrm{in} * .75\ \mathrm{in}}{-11.6\ \mathrm{x}\ 10^{-3}\,\frac{\mathrm{V} * \mathrm{m}}{\mathrm{N}} .0075\ in} * \frac{.0254\ m}{1\ in} = 82{,}000\ pF$$

The final result is:

$$C_{storage} = \frac{C_{cp} * (V_{strain} - V_{storage})}{V_{storage}} = \frac{82 * 10^{-9} F * (250\ V - 10\ V)}{10\ V} = 2.0 * 10^{-6} F$$

Thus, the storage capacitor was chosen to be on the order of microfarads. The actual value was chosen experimentally given the variation in the system to be 10 microfarads.

APPENDIX E: SENSOR INSTALLATION

Figure 51: Holes were cut into the concrete and material removed from below the road.

First a section of useable road needed to be found and slots cut into the concrete. Once the concrete layer was breached, an extra 14" of depth was required to pour a concrete sub base capable of supporting the sensor and the heaviest of loads which it would see. Next, rebar was placed in the bottom of the hole on risers, to strengthen the concrete fill which would eventually be poured into the hole. Rebar was also loaded into holes in the sides of the slots to strengthen the section just under the sensor even more.

Once this was accomplished, wooden forms were made of the sensors and these were vibrated down into the freshly poured concrete. The concrete was allowed to dry and the forms were removed by chipping and burning the wood. Once this was done, the sensors could be placed into the slots which were the correct geometry and strong enough to support to sensor and any vehicle loading.

Figure 52: Rebar was placed in the hole before pouring to strengthen the new sub base.

Figure 53: Rebar was used to form a bond between the new and old forms.

Figure 54: Once the new sub base was poured, the sensors could be installed.

APPENDIX F: SPEED DEPENDANT DATA BEHAVIOR FOR SENSORS 1 AND 2

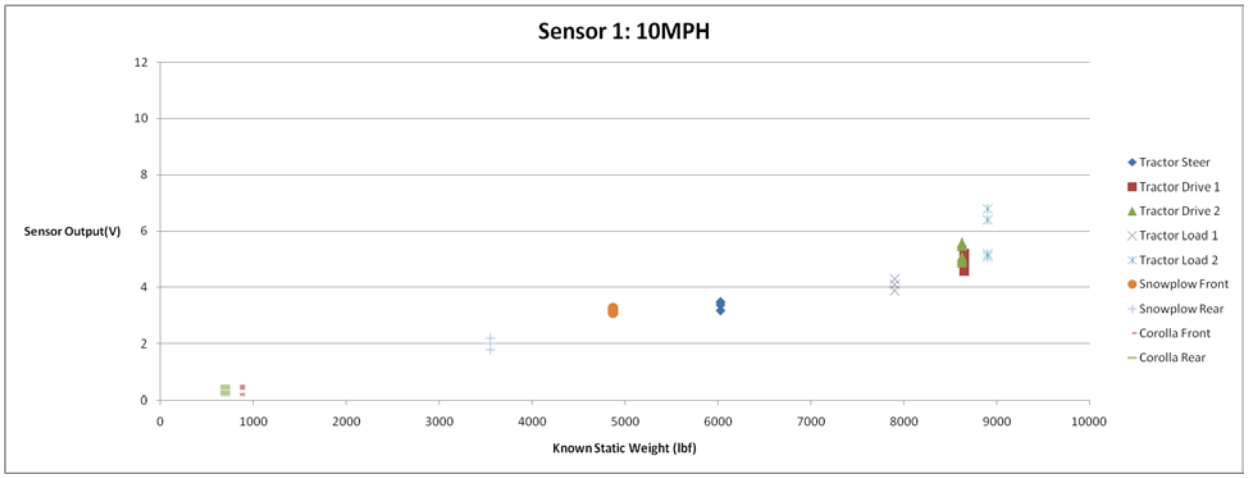

Figure 55: Static weight of vehicle vs. sensor 1 readings at 10 mph.

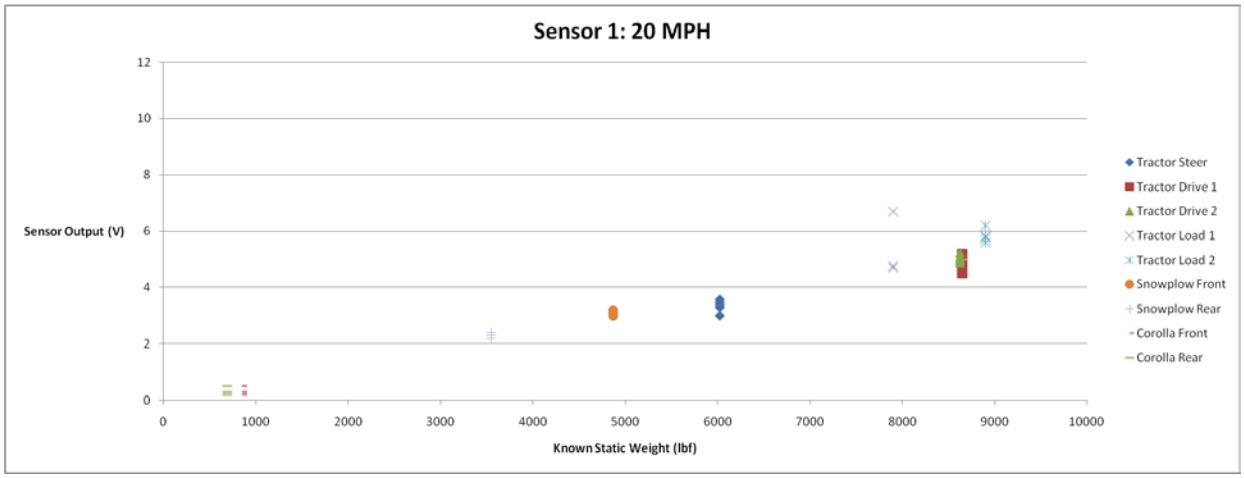

Figure 56: Static weight of vehicle vs. sensor 1 readings at 20 mph.

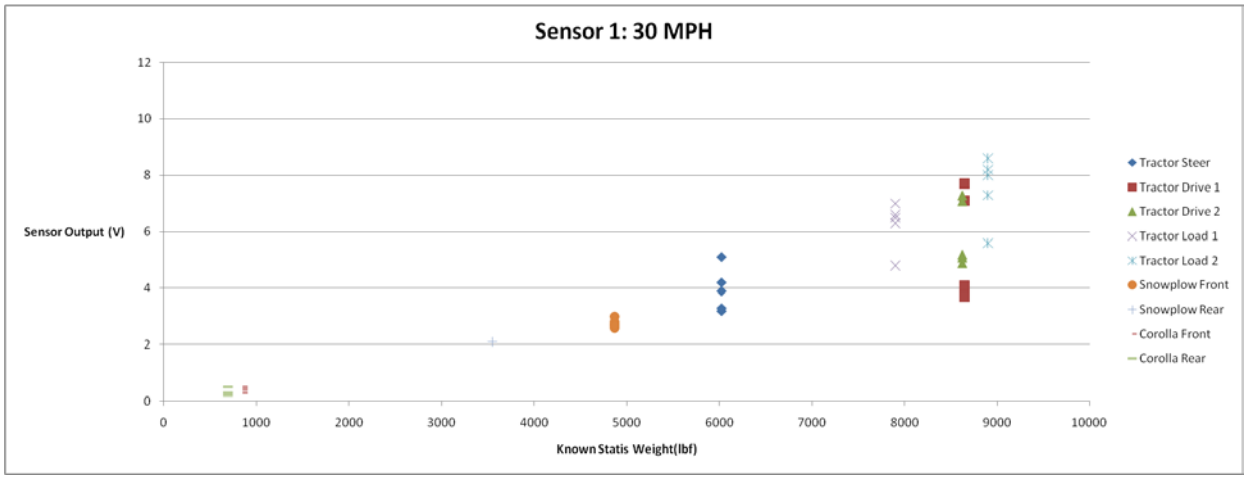

Figure 57: Static weight of vehicle vs. sensor 1 readings at 30 mph.

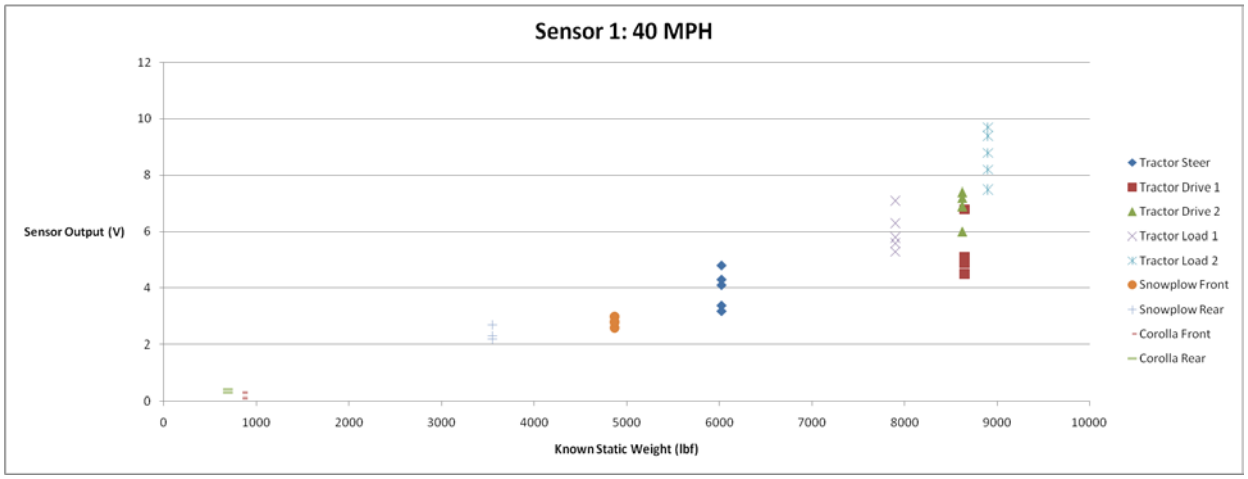

Figure 58: Static weight of vehicle vs. sensor 1 readings at 40 mph.

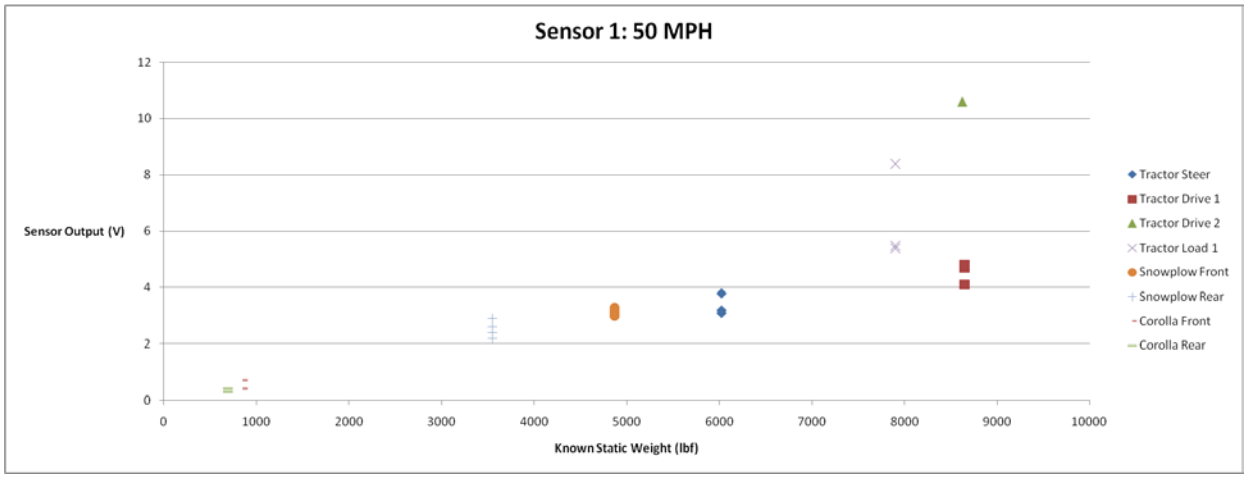

Figure 59: Static weight of vehicle vs. sensor 1 readings at 50 mph.

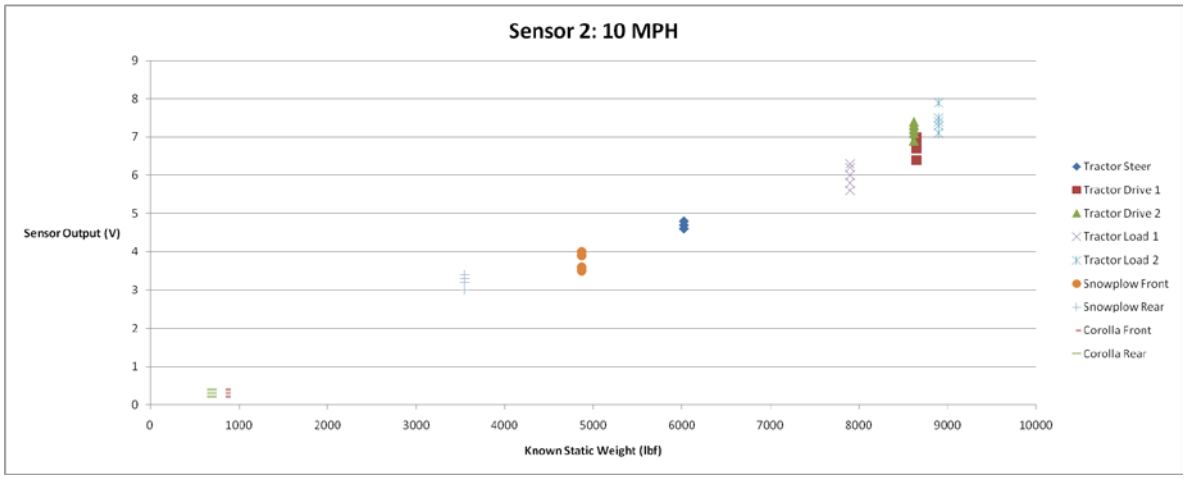

Figure 60: Static weight of vehicle vs. sensor 2 readings at 10 mph.

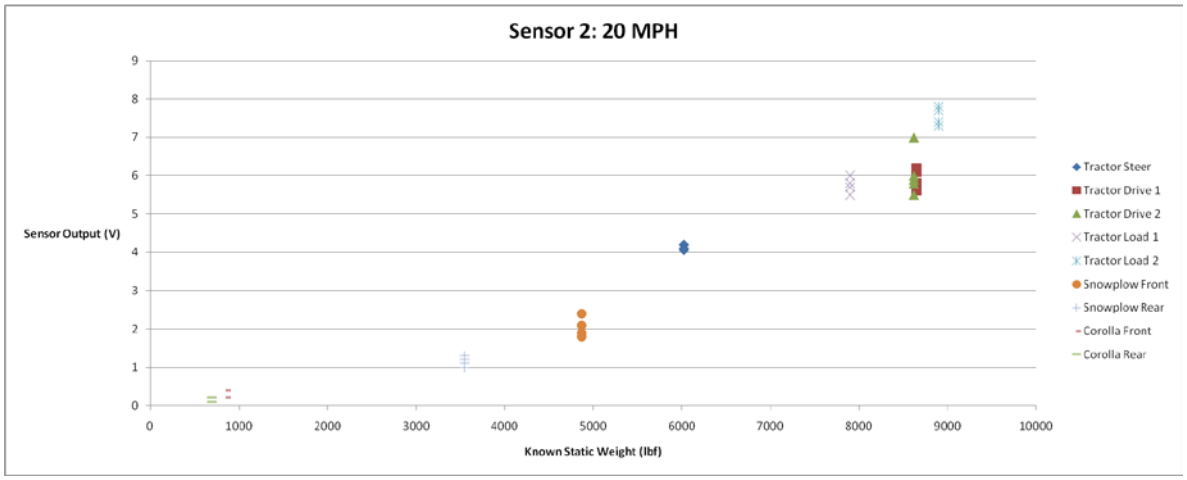

Figure 61: Static weight of vehicle vs. sensor 2 readings at 20 mph.

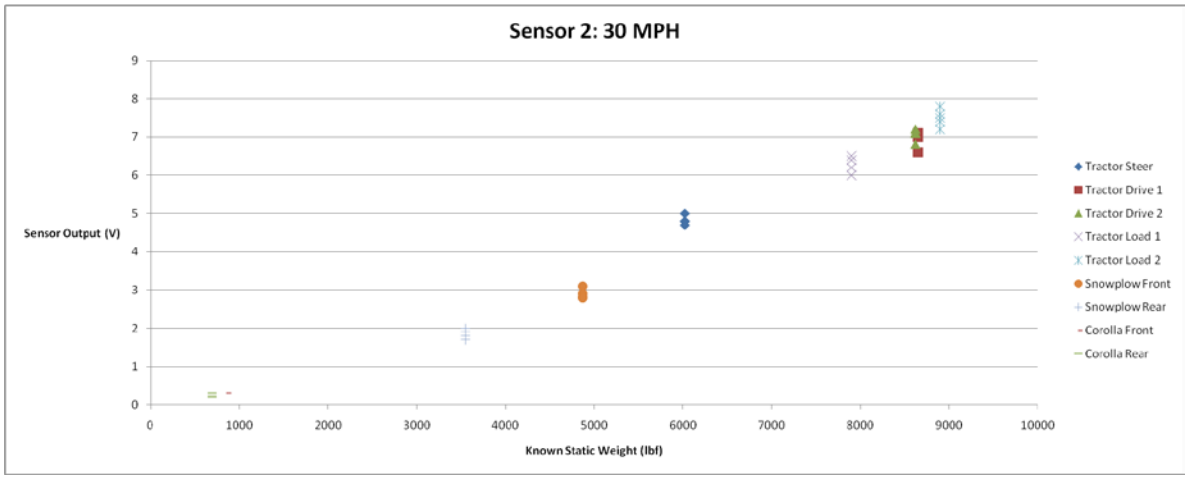

Figure 62: Static weight of vehicle vs. sensor 2 readings at 30 mph.

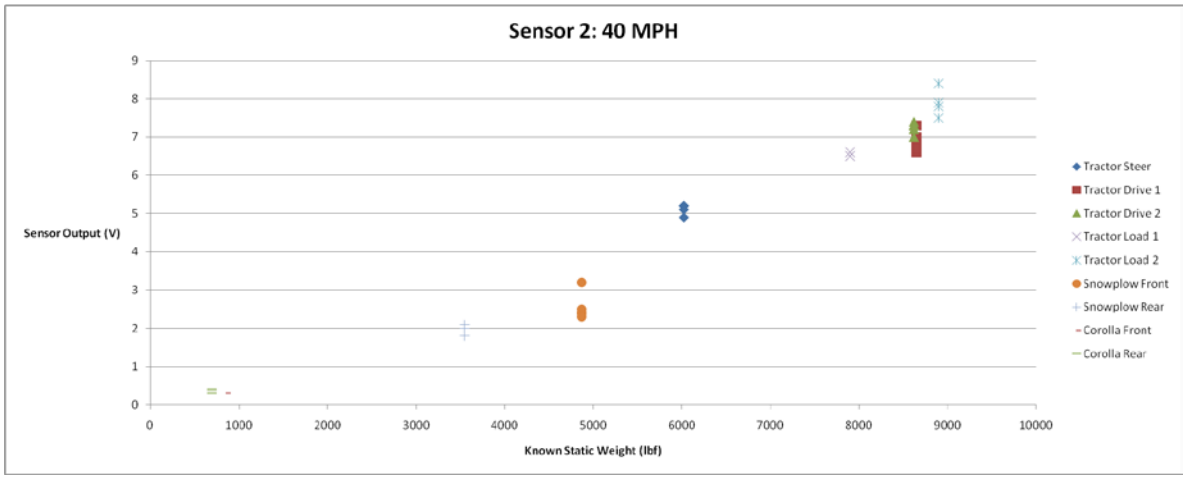

Figure 63: Static weight of vehicle vs. sensor 2 readings at 40 mph.

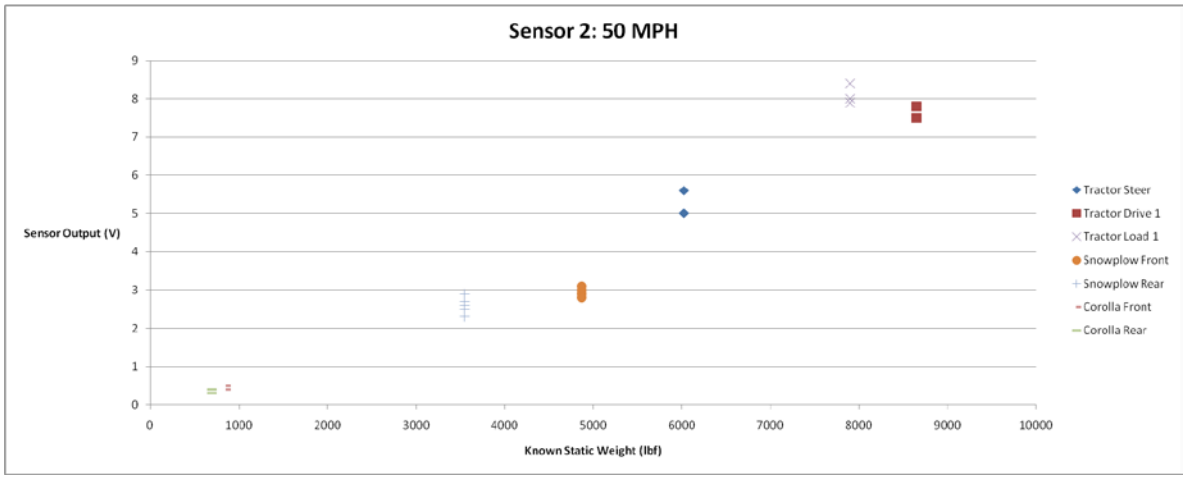

Figure 64: Static weight of vehicle vs. sensor 2 readings at 50 mph.

APPENDIX G: SPEED DEPENDENCY OF AVERAGE ESTIMATION

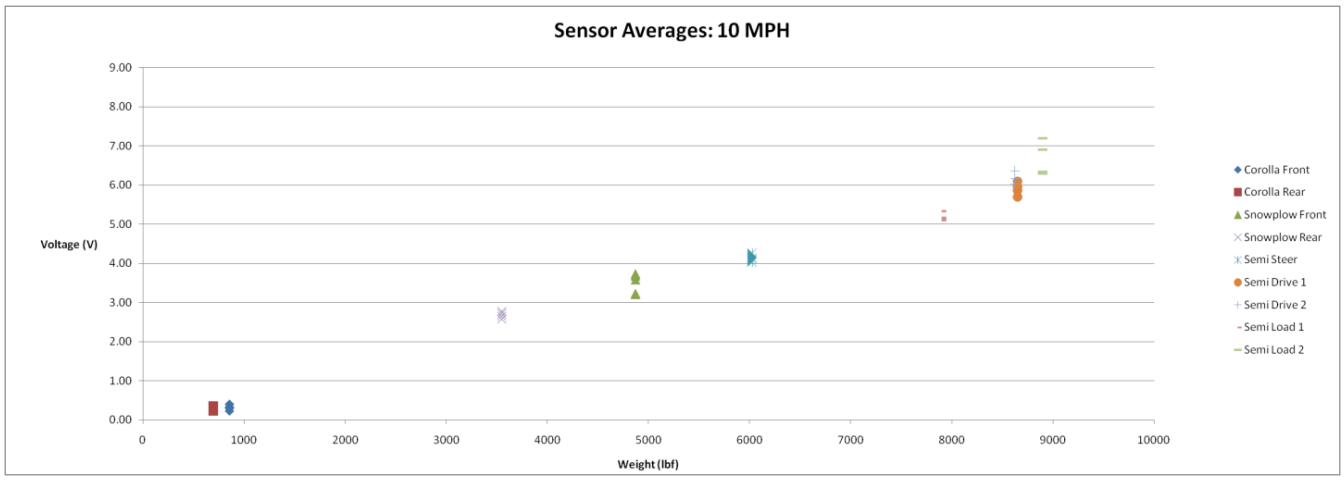

Figure 65: Static weight of vehicle vs. sensor average readings at 10 mph.

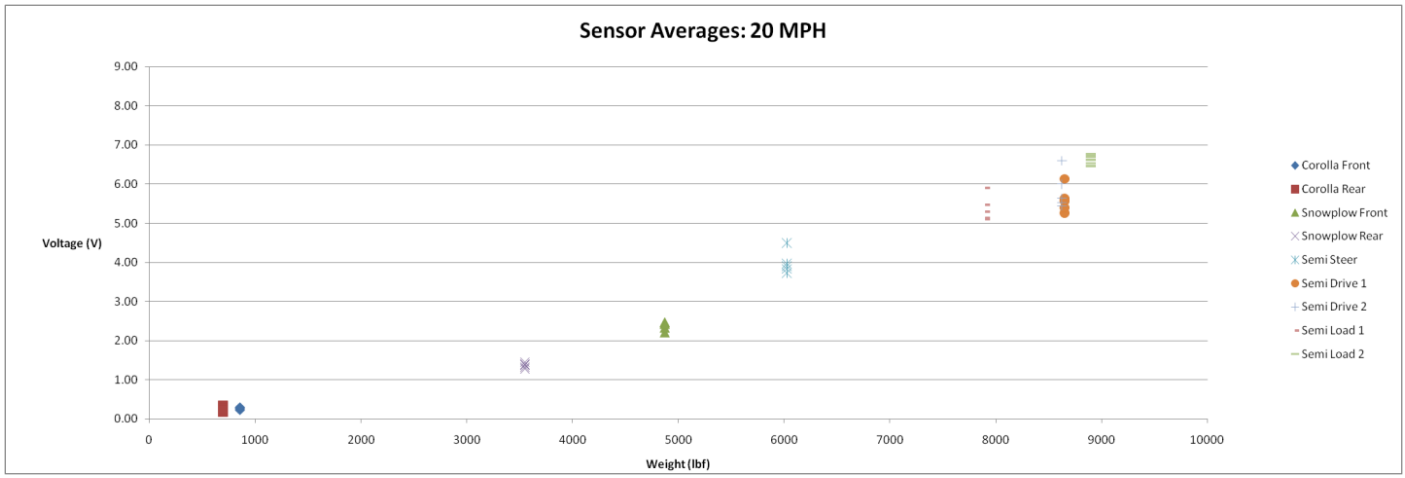

Figure 66: Static weight of vehicle vs. sensor average readings at 20 mph.

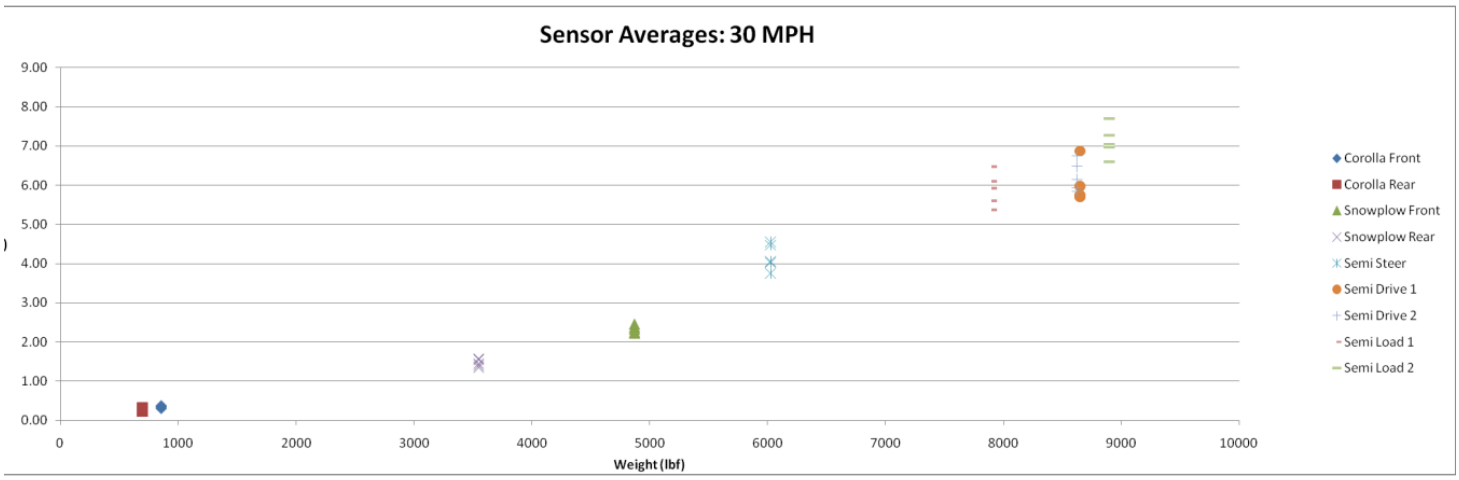

Figure 67: Static weight of vehicle vs. sensor average readings at 30 mph.

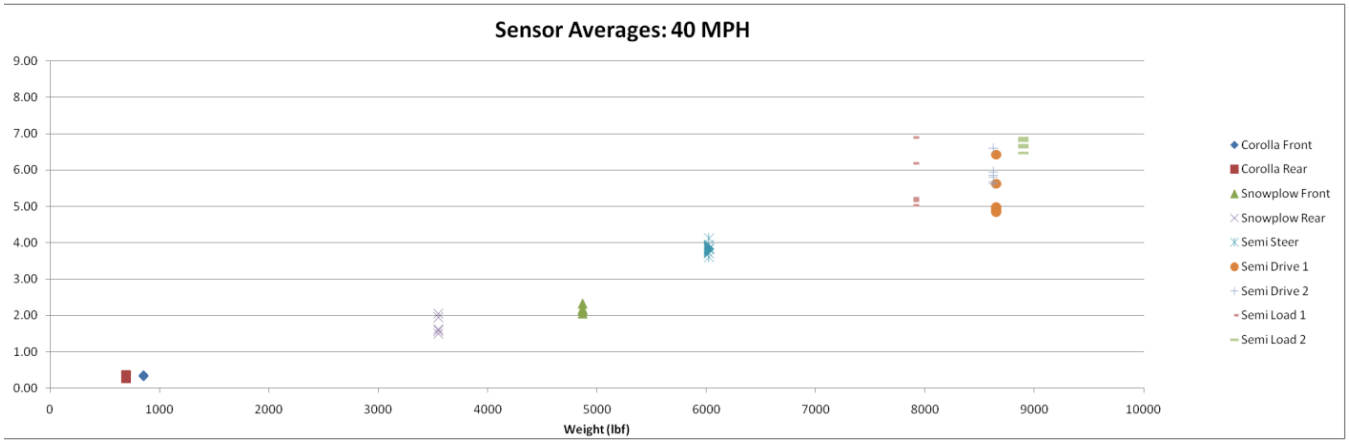

Figure 68: Static weight of vehicle vs. sensor average readings at 40 mph.

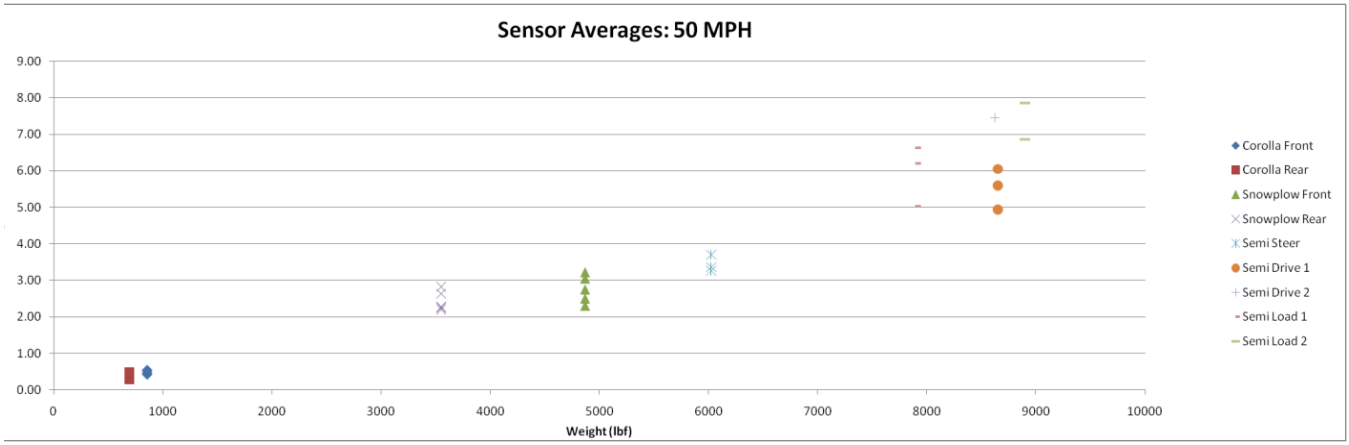

Figure 69: Static weight of vehicle vs. sensor average readings at 50 mph.

APPENDIX H: VELOCITY SPECIFIC OUTPUT OF THE LEAST SQUARES METHOD

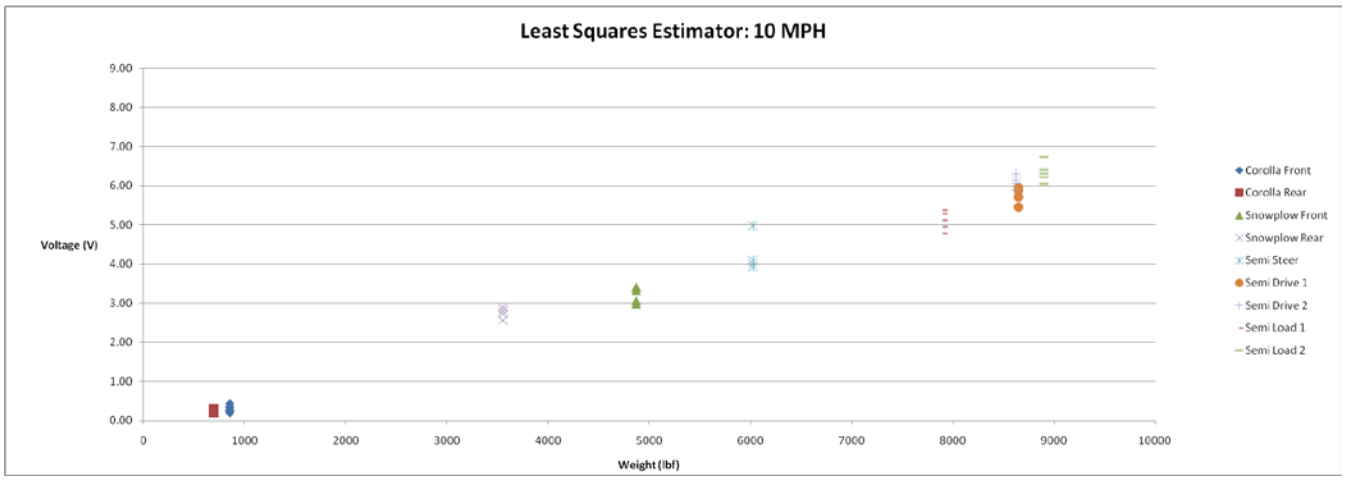

Figure 70: Static weight of vehicle vs. sensor least squares estimator at 10 mph.

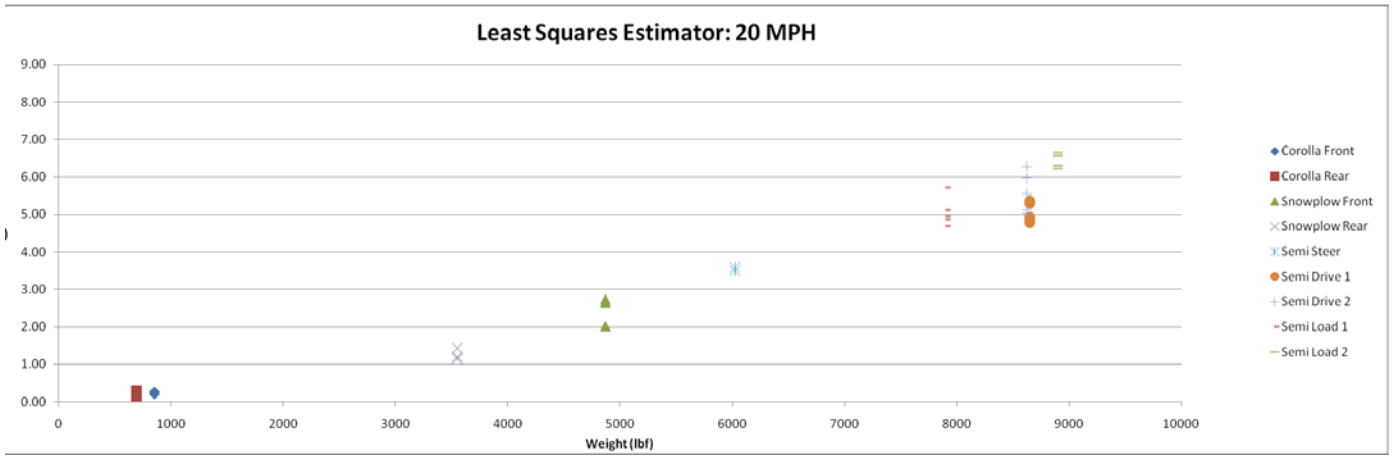

Figure 71: Static weight of vehicle vs. sensor least squares estimator at 20 mph.

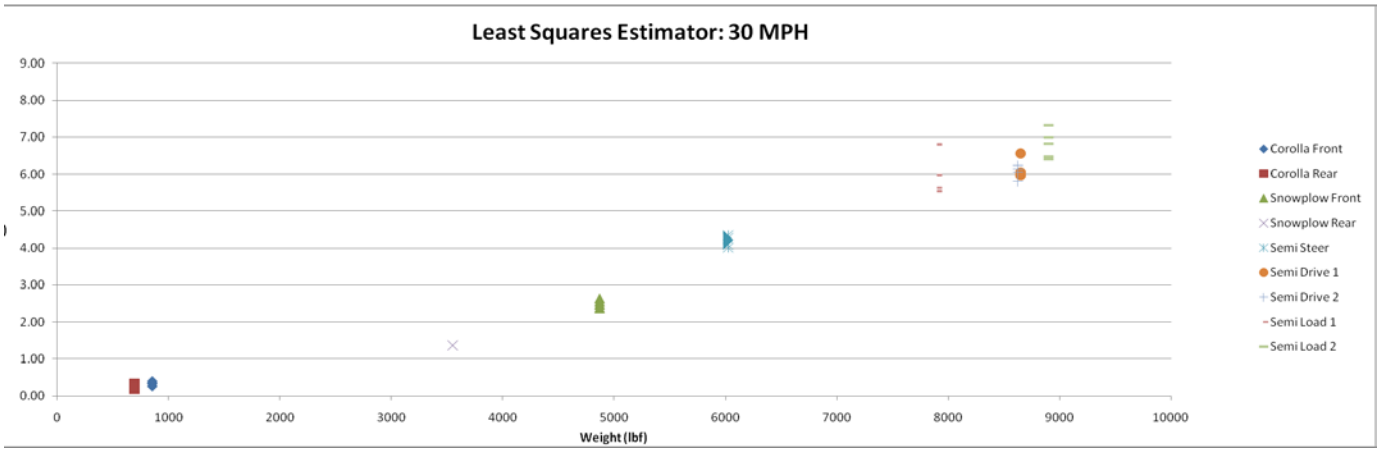

Figure 72: Static weight of vehicle vs. sensor least squares estimator at 30 mph.

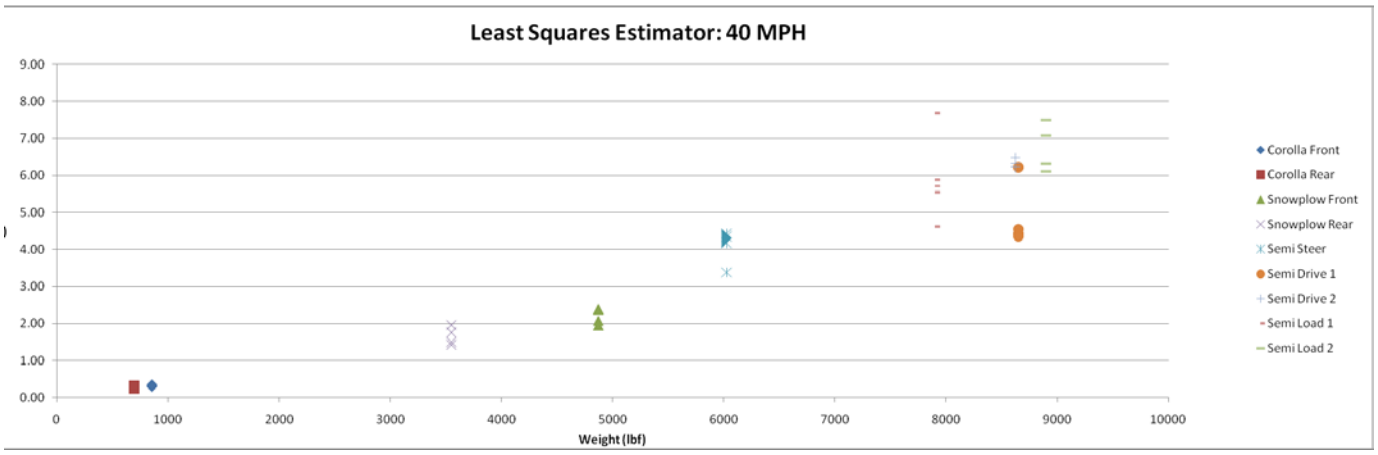

Figure 73: Static weight of vehicle vs. sensor least squares estimator at 40 mph.

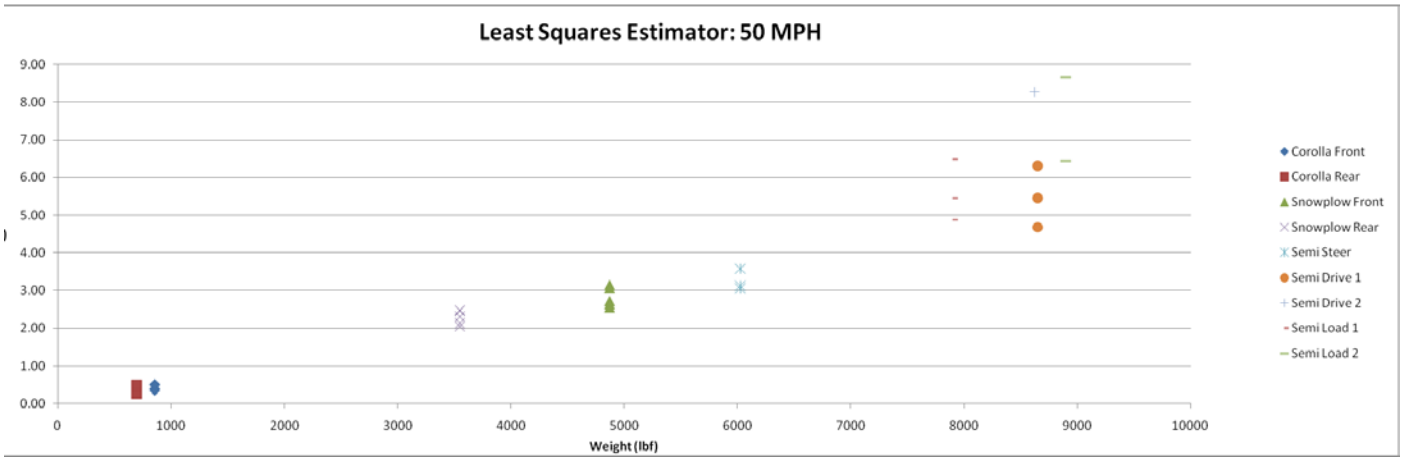

Figure 74: Static weight of vehicle vs. sensor least squares estimator at 50 mph.

APPENDIX I: PREVIOUS DATA PROCESSING METHODS WITH 50 MPH EXCLUDED

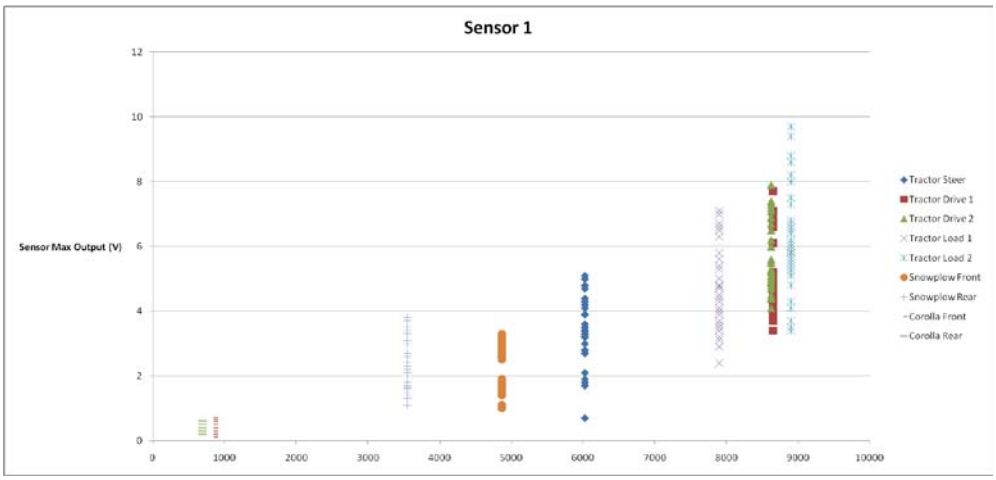

Figure 75: Static weight of vehicles vs. all sensor 1 readings (excluding 50 mph).

APPENDIX J: SENSOR 4 DATA EXCLUDED

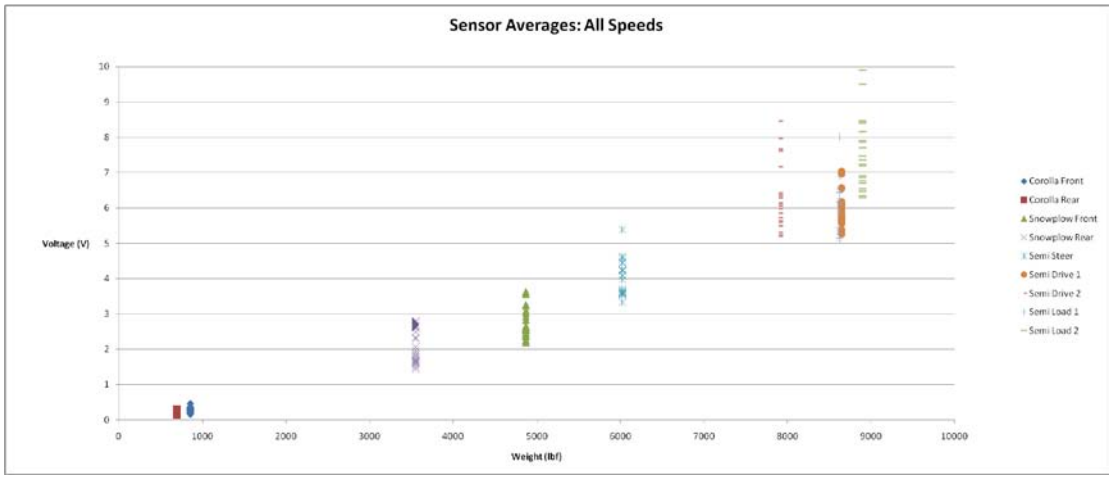

Figure 76: Static weight of vehicles vs. sensor 1, 2, and 3 average estimator.

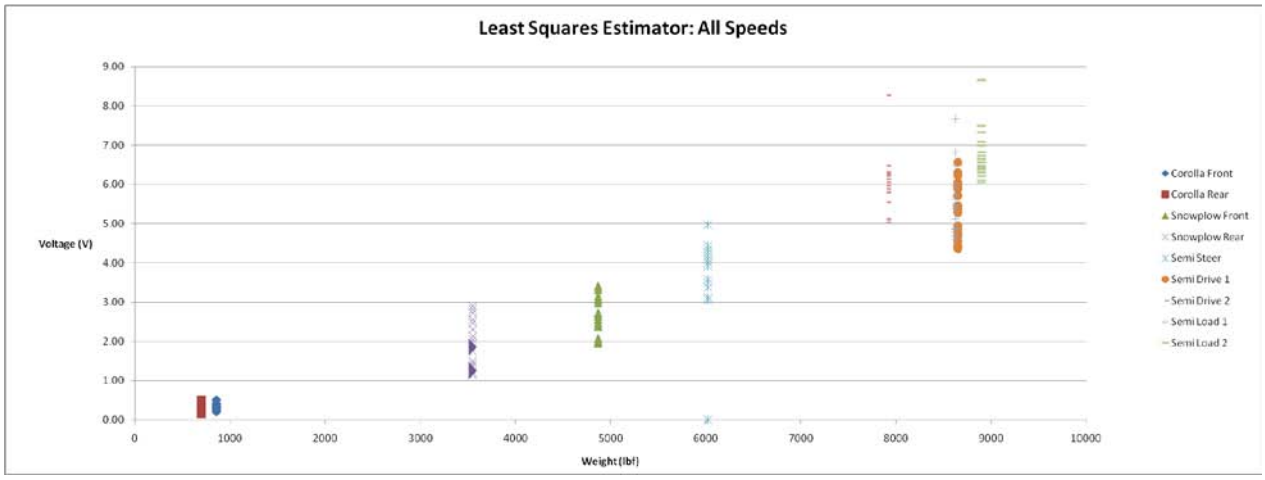

Figure 77: Static weight of vehicles vs. sensor 1, 2, and 3 least square estimator.

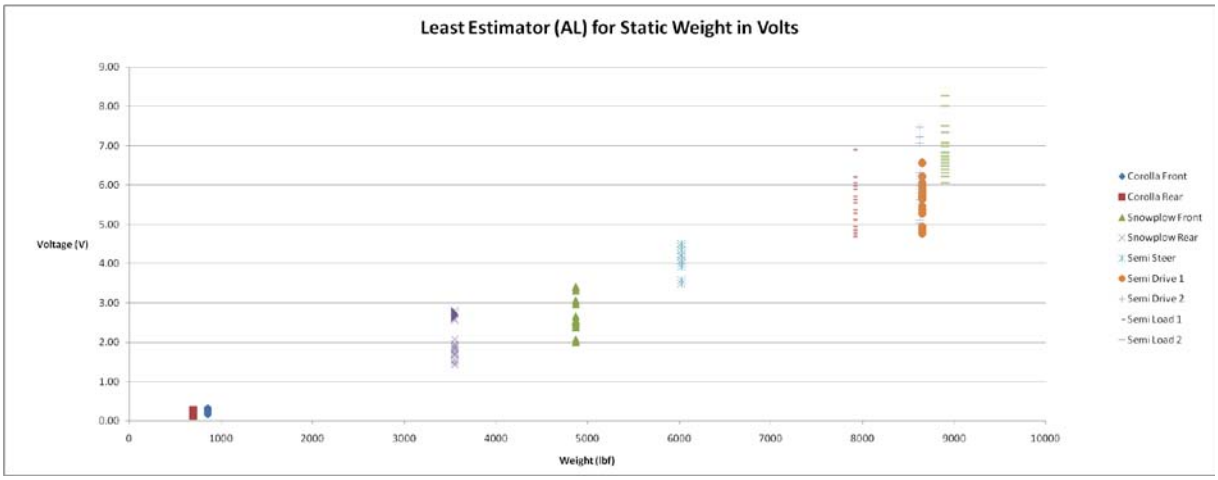

Figure 78: Static weight of vehicles vs. lowest of average and least squares estimators for sensor 1, 2, and 3.

APPENDIX K: 10-30 MPH DATA FOR 3-16

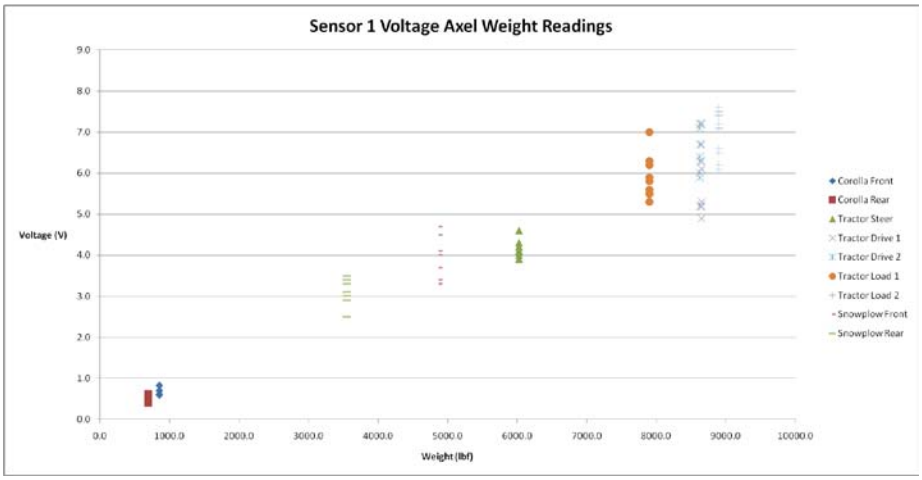

Figure 79: Static weight vs. sensor 1 output for 5-30 mph on 3-16.

Sensor 1 Raw Data

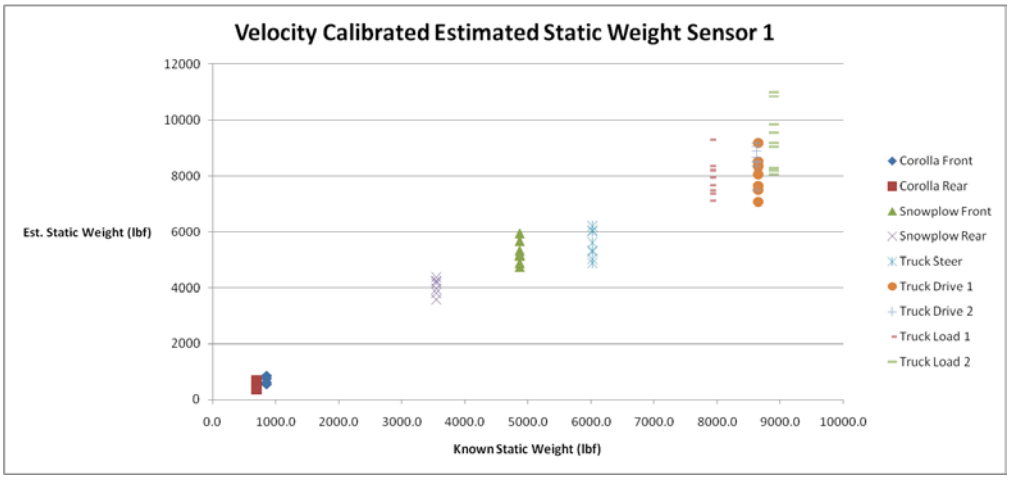

Figure 80: Static weight vs. sensor 1 calibrated output for 5-30 mph on 3-16.

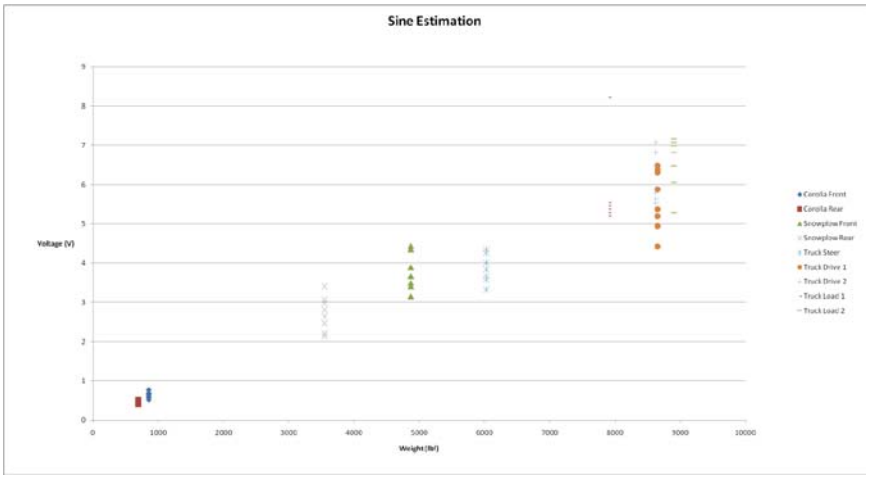

Figure 81: Static weight vs. sine estimator for 5-30 mph on 3-16.

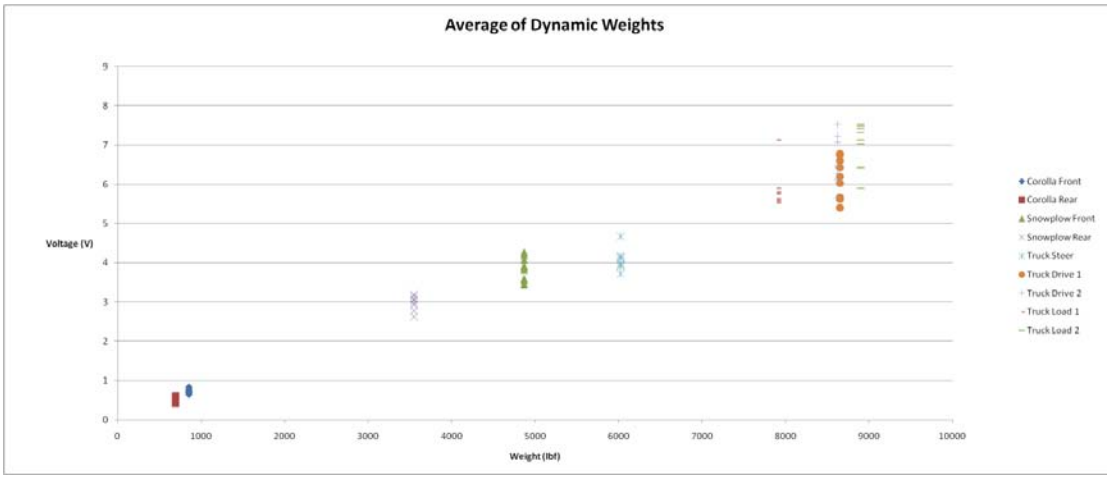

Figure 82: Static weight vs. sensor average estimator for 5-30 mph on 3-16.

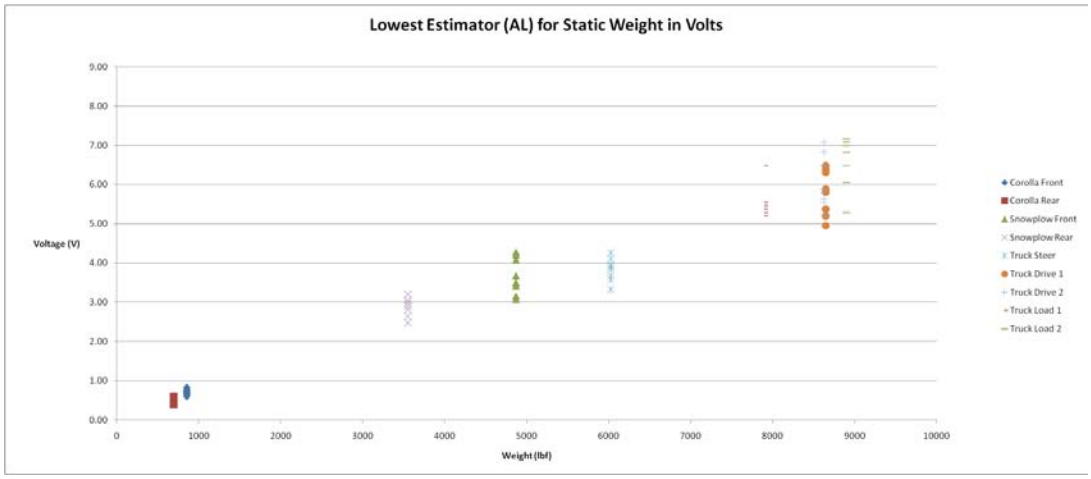

Figure 83: Static weight vs. sensor lowest of average and least squares estimators for 5-30 mph on 3-16.

APPENDIX L: 10-30 MPH DATA FOR 4-9

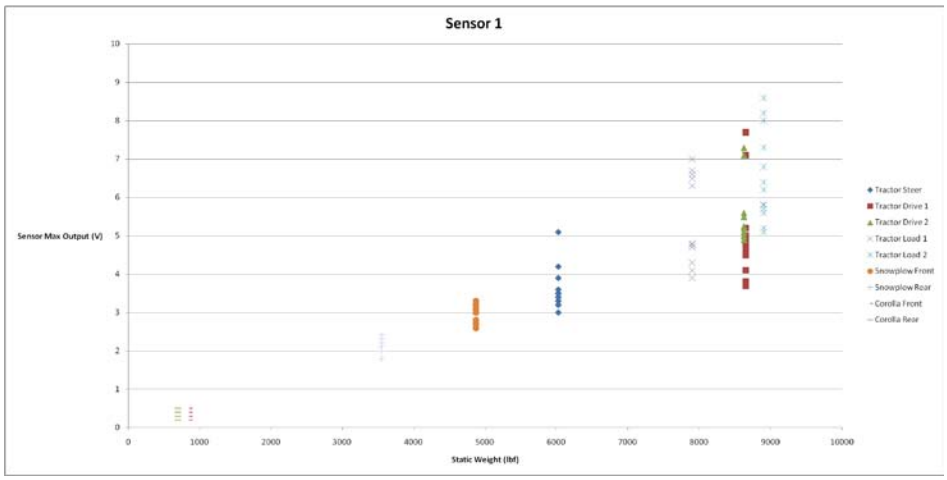

Figure 84: Static weight vs. sensor 1 data for 10, 20 and 30 mph on 4-9.

Original Data for Sensor 1

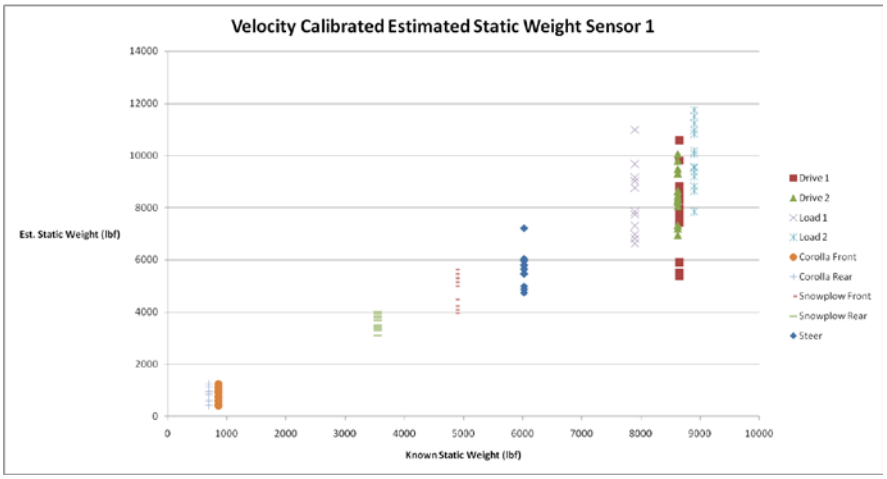

Figure 85: Static weight vs. sensor 1 calibrated data for 10, 20 and 30 mph on 4-9.

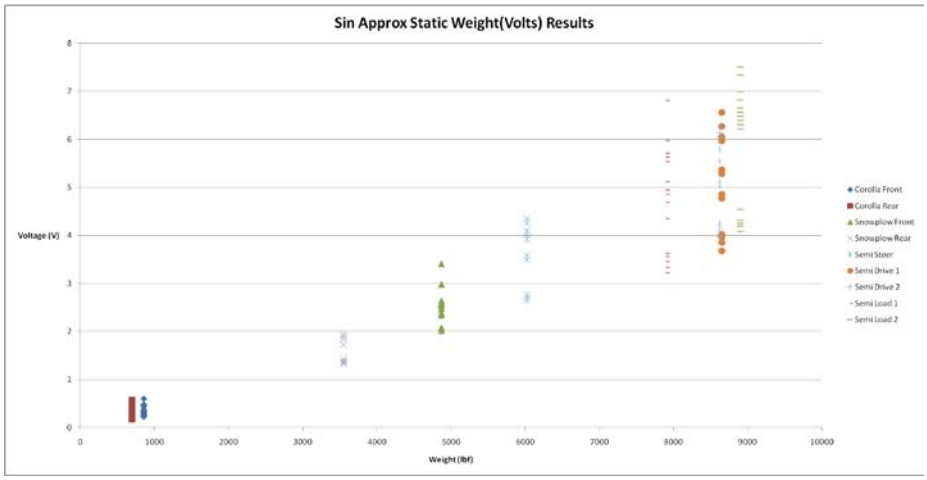

Figure 86: Static weight vs. sine estimator for 10, 20 and 30 mph on 4-9.

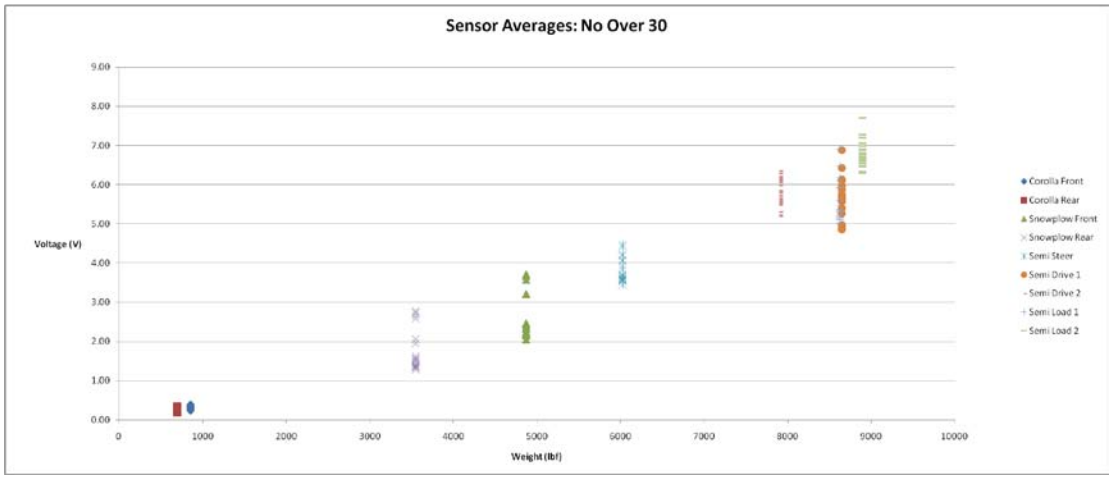

Figure 87: Static weight vs. sensor average estimator for 10, 20 and 30 mph on 4-9.

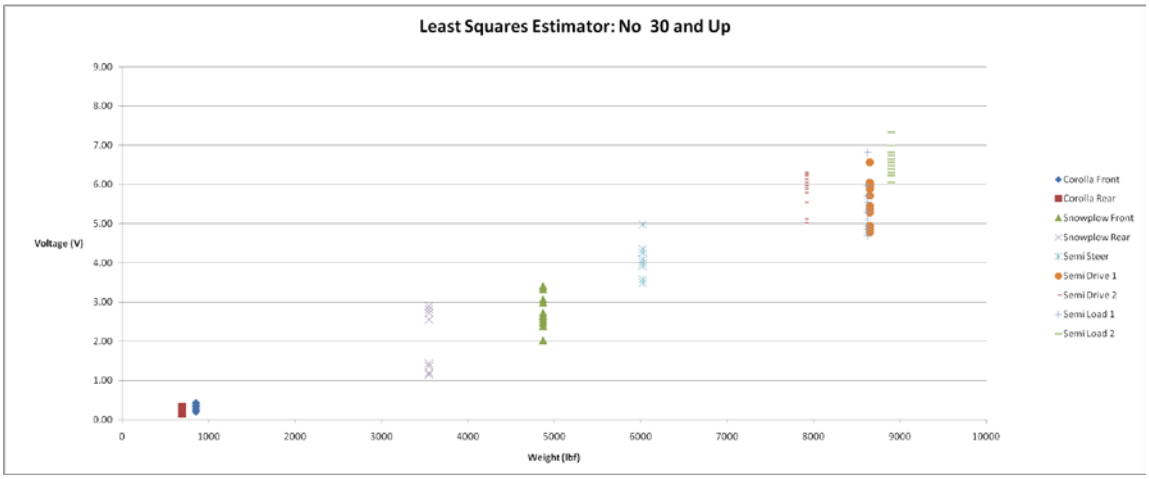

Figure 88: Static weight vs. sensor least squares estimator for 10, 20 and 30 mph on 4-9.