2950 Niles Road, St. Joseph, MI 49085-9659, USA
269.429.0300 fax 269.429.3852 hq@asabe.org www.asabe.org

An ASABE Meeting Presentation

Paper Number: 12-1337075

Wireless Load Weight Monitoring Via a Mobile Device Based on Air Suspension Pressure

Alexander W. Layton, Research Assistant

Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave, West Lafayette IN, 47907, awlayton@purdue.edu

Andrew Balmos, Research Assistant

Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave, West Lafayette IN, 47907, abalmos@purdue.edu

Derek Hancock, Research Assistant

Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave, West Lafayette IN, 47907, dhancock@purdue.edu

Aaron Ault, Research Scientist

Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave, West Lafayette IN, 47907, ault@purdue.edu

James V. Krogmeier, Associate Professor

Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave, West Lafayette IN, 47907, jvk@purdue.edu

Dennis R. Buckmaster, Associate Professor

Agricultural and Biological Engineering, Purdue University, 225 S. University St, West Lafayette IN, 47907, dbuckmas@purdue.edu

**Written for presentation at the
2012 ASABE Annual International Meeting
Sponsored by ASABE
Hilton Anatole
Dallas, Texas
July 29 – August 1, 2012**

Abstract. *A system was designed and tested to measure pressure in the airbag suspension components of an over-the-road tractor trailer and use those signals to wirelessly display load weight. Three subcomponents (a HID -- human interface device and two DPUs -- digital processing units) were required and communicated using Zigbee wireless devices. Except for a region directly in front of the cab, the unit transmitted weights up to 45 m. Effect of number of calibration points on system accuracy was evaluated.*

Keywords. Air bag, load, suspension, truck, weight, wireless

Introduction

In order to optimize loading and comply with weight regulations, truck drivers need to be able to measure the weight of their vehicle, both total weight and individual axle weights. In many situations, e.g., field loading, fixed scales may be far away and portable scales may be impractical or too costly. A fast, real-time, cost-effective system for vehicle weighing is needed. There are current systems on the market, but some do not wirelessly transmit weight information (AirWeigh, 2011; Farmtronics, 2012; TruckWeight, 2011) and those that rely solely on air-spring pressure for mixed air-spring and mechanical suspension trucks are not well documented.

A system is proposed which weighs the vehicle by measuring air pressure in its air-spring suspension. It can be shown that there is a linear relationship between suspension air pressure and the weight carried. The coefficients in this linear weight-pressure model can be derived via standard analysis in the static case, if the geometry is known. As it is impractical to know the geometry in all cases and because it can vary with leveling valve setting, a calibration method and least squares fitting is suggested.

Note that the air-spring suspension system supports only the body and load, and not the axles or the suspension system itself. In addition, in many trucks the front steering axle is connected to a traditional spring suspension, which would require some other technology for measuring weight, e.g., load cells. It may be possible to compromise by calibrating around the missing front axle measurement using an affine linear weight-pressure model and least squares.

The purpose of this paper is to describe the design of such a system including the calibration procedure and to investigate the accuracy attainable subject to the compromise above.

Objectives

The specific objectives of this work were to:

1. Design a method and test a prototype system for estimating a semi-truck's weight.
2. Provide weight estimations wirelessly in real time.
3. Determine how to sufficiently calibrate the system to achieve estimations with 1% of full scale 80% of the time.

Methods

Air-Spring Suspension Systems – weight determination

Air-spring suspension assemblies use air pressurized rubber cylindrical structures, called airbags, to support the applied weight. A top metal plate transfers the load to the airbag (Figure 1). From Pascal’s principle, the force resulting from a pressure applied to a cylinder is proportional to that pressure. Approximating the airbag by a cylinder would imply that the air bag pressure times the metal top plate area equals the load carried.

A leveling valve (34 in Figure 1) is responsible for injecting air into or venting air out of the airbag as the load changes. With a reduction in vehicle load the height of the airbag increases and the leveling valve opens allowing air to vent, returning the airbag to its original height. On the other hand, if the vehicle load increases there is a reduction in airbag height that causes the leveling valve to engage, allowing air from the supply lines into the airbag, returning it to its original height. The leveling valve also serves to counteract changes in pressure caused by temperature.

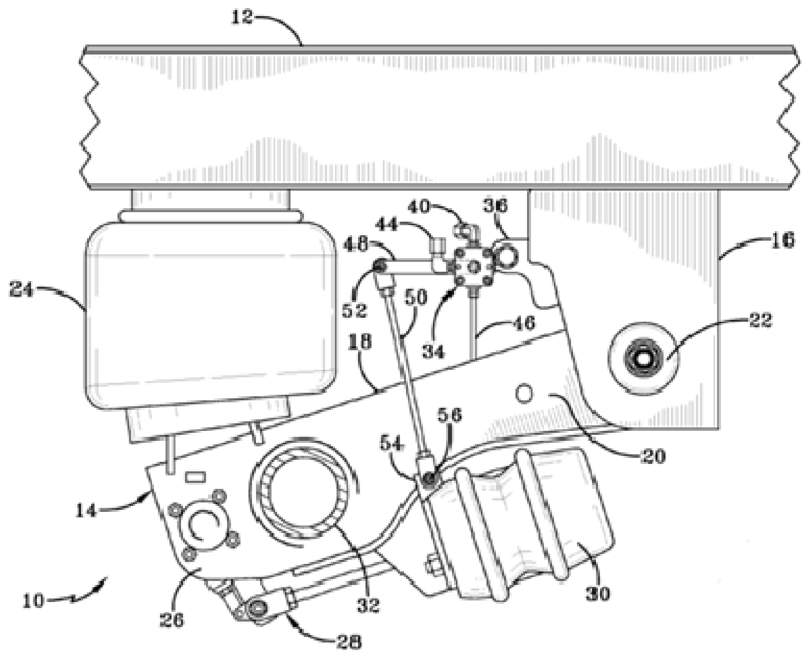

Figure 1: Mechanical Drawing of Leading Arm Air-Spring Assembly. 16 is the Frame Hanger, 32 is the Axle, and 24 is the Airbag. (Pierce and Cervantez, 2009)

The air-spring suspension assembly distributes the load weight between the airbag (24 in Figure 1) and the frame hanger (16 in Figure 1). From Figure 1 it is evident that the distribution of the load between the airbag and the frame hanger is a function of the height of the airbag (among other quantities). However, assuming that the leveling valve holds the airbag height constant, then the proportion will remain fixed.

Figure 2 is based upon Figure 1 but it has been redrawn as a free body diagram to facilitate a discussion on determining the model of the air-spring suspension system when only the airbag pressure is known. The derivation limits itself to considering only one air-spring which is directly controlled by its own independent leveling valve. It will later be argued that a collection of air-spring assemblies can be viewed as one large air-spring system.

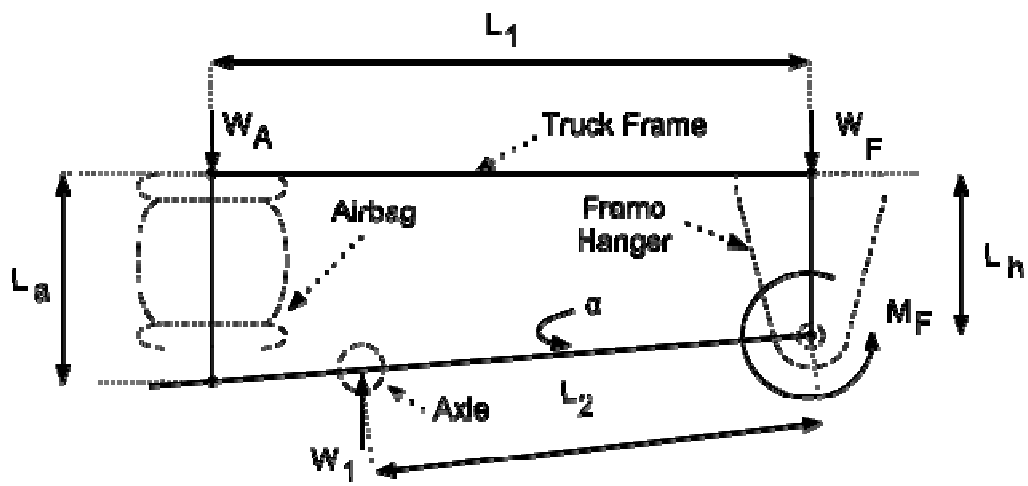

Figure 2: Free Body Diagram of Leading Arm Air-Spring Assembly.

Recall the previous assumption that the leveling valve holds constant the height of the airbag. With this assumption, the free body diagram may be treated as a rigid body, which allows the application of the standard statics analysis. In the figure, W_A is the weight on the airbag, W_F is the weight on the frame hanger, and W_1 is the weight on the axle. Setting the moment M_F equal to zero, it can be shown that the relationship between the axle weight and the airbag pressure is:

$$\frac{W_1}{W_A} = \frac{L_2}{L_1 \cos(\alpha)} \quad (1)$$

In other words, $W_1 \propto W_A$ and $L_2/L_1 \cos(\alpha)$ does not depend upon the load assuming the leveling valve holds the length L_a constant. This means that it is possible to calibrate and fit a model with least squares because the parameters of the model are fixed.

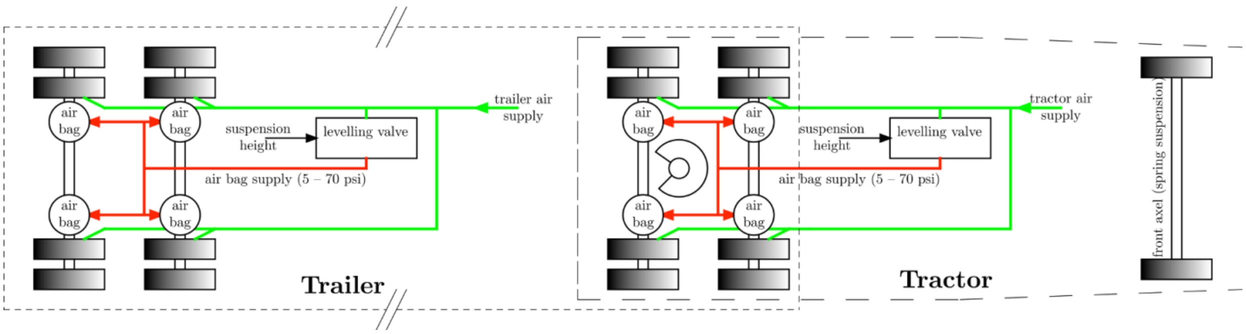

Figure 3: Schematic of a Standard Semi-Truck Air-Spring Suspension System

Figure 3 shows a simplified schematic of an air-spring suspension system. It features two airbags per axle and independent trailer and tractor pneumatic systems. A single leveling valve per system maintains the working pressure so that the changes in the airbags' applied force and height occur in unison. More complex situations exist, and the results to follow can be extended to cover them. For the simple model shown in Figure 3 it is assumed that the grouped air-spring assemblies for the trailer and for the tractor act as separate individual air-springs.

The weight of the entire vehicle can be expressed as the sum of the weight carried by each axle plus a static offset equal to the weight of the axles, tires and unmeasured weight carried by the front springs. Most of the static offset weight is tare weight, with the exception of any load weight transferred to the front axle.

Therefore the model for the total weight (W) is:

$$a_1 W_{Tractor} + a_2 W_{Trailer} + b = W \quad (2)$$

where $W_{Tractor}$ is the sum of the weight carried by the axles in the tractor's tandem group and $W_{Trailer}$ is the sum of the weight carried by the axles in the trailer's tandem group. As argued above the axle weights are proportional to the airbag weights, and the airbag weights are proportional to the airbag pressures. Therefore, Equation 2 could be rewritten in terms of airbag pressures. In fact, in the implementation described later it is the numeric outputs of the pressure sensors' analog to digital converters (ADC), which are recorded. Hence the model actually corresponds to a version of Equation 2 in terms of the ADC outputs.

With more than two calibration data sets (paired raw ADC data and known weights) the system of Equation 2 is over determined. Therefore the least squares method was used to find the coefficients (a_1, a_2, b) in the model.

System Design

To gather data and evaluate the proposed algorithm a system was designed which gave real time weight estimates to the user. The initial design only implemented a model which assumed a_1 and a_2 from Equation 2 are equal. This was done because the algorithm was originally implemented on a microcontroller with only fixed point arithmetic. However the system also displayed all raw data which the operator recorded for later processing.

The system consists of three parts, which communicate wirelessly. They are: two data processing units (DPUs) to measure the working pressure of their individual axle groups, and one human interface device (HID) which allows the user to interact with the system. To measure the pressure, the DPU air lines were spliced into the pressure lines of the suspension that feed the analog gauges currently installed on the truck. Both the HID and the DPUs are battery powered, and thus installation only requires connecting the airlines and securing the DPUs to the vehicle.

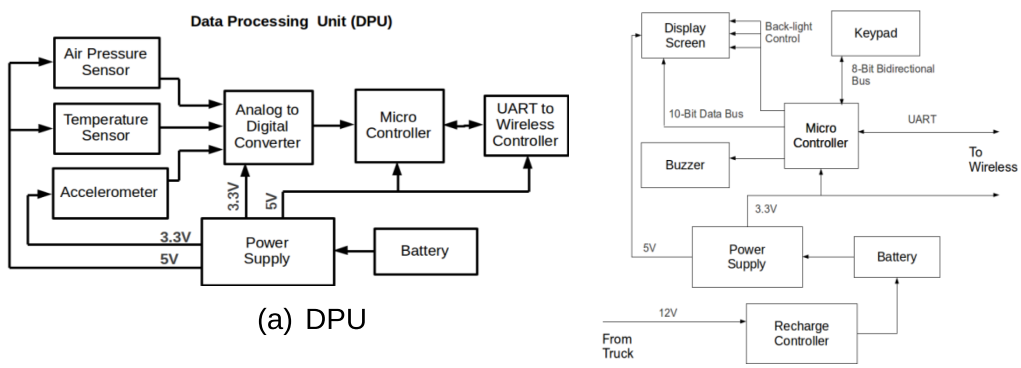

(a) DPU

(b) HID

Figure 4: Weighing System Block Diagrams

The DPU, shown schematically in Figure 4a, is the part of the system which takes necessary raw measurements. It is equipped with analog sensors for measuring pressure, temperature, and acceleration, along with some other internal components as described in Table 1. The acceleration measurements can be used to find the angle of the vehicle with respect to level. An analog-to-digital converter (ADC) periodically measures the various sensors and reports the values back to an on-board microcontroller. Every one half second, the microcontroller averages the values and instructs the Zigbee wireless module (Table 1) to broadcast those averages to the HID. The DPU requires both 3.3 and 5 Volt supplies which are efficiently generated via a switching regulator powered by two AA batteries in series.

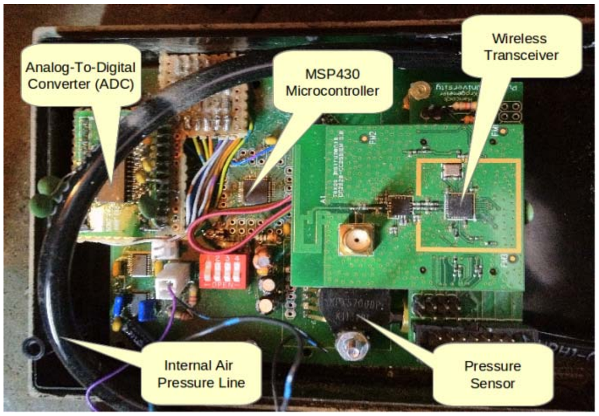

(a) DPU Interior

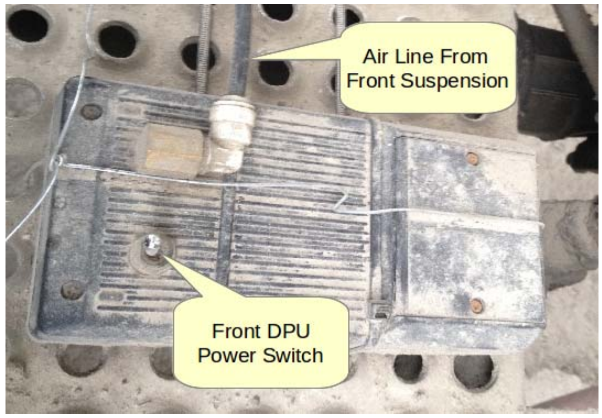

(b) DPU Exterior

Figure 5: Hardware of the Weighing System

Part Number	Manufacturer	Quantity	Description
MSP430F5310	Texas Instruments	6	Microcontroller used for all the processing in both the HID and the DPUs.
CC2520-CC2591EMK	Texas Instruments	3	Zigbee module used for wireless communications.
MPX5700	Honeywell	2	Pressure sensor in the DPUs.
TMP20	Texas Instruments	2	Temperature sensor in the DPUs.
LIS3L02AS4	STMicroelectronics	2	Accelerometer in the DPUs (for finding angle).
TPS60122	Texas Instruments	3	3.3 Volt regulator used in both HID and DPUs.
TPS60140	Texas Instruments	3	5 Volt regulator used in both HID and DPUs.
TLC3548	Texas Instruments	2	Analog to digital converter used in DPUs to sample all sensors.
MAX713	MAXIM	1	Battery recharge controller used in HID.

Table 1: Bill of Materials

As with the DPU, the HID requires both a 3.3 and 5 Volt supply generated using the same switching regulators and battery pack. However, unlike the DPU the HID has support for recharging its battery pack. . With the 16-key keypad, a 2-line, 16 character LCD display (Figure 6), and a buzzer, the HID is both the input and output device for the system; it includes a menu which puts the HID into different modes such as (i) weight display, (ii) enter alarm weight, and (iii) add a calibration point. The HID, shown schematically in Figure 4b, is the part of the system that uses the weight models and calibration points to estimate the system parameters and weight. It receives the average sensor measurements from the DPUs via the same Zigbee wireless module. When there is a new data set from both DPUs, the HID uses a microcontroller to estimate the weight. If the weight exceeds a pre-entered threshold, a buzzer sounds to inform the user that the truck is full.

Figure 6: Human Interface Device for Displaying Weights

and Identifying Calibration Data.

When entering a calibration point the user should make sure the truck is stationary, navigate to the “add a calibration point” mode, and then enter the current known truck weight. When the weight is entered the HID uses its on-board microcontroller to perform least squares regression with all calibration points to refine the current model coefficients.

Experiment Design

Working Prototype

To evaluate the system, the prototype was installed on a semi-truck owned by Ault Farms. The driver was asked to continue to operate the vehicle as normal while entering calibrations and collecting data.

Figure 7 shows the real air-spring in use on the test semi as well as the installation point for the front DPU. The air-springs are the same for all assemblies on the both the front and rear tandem axles. However, the arm length varies between the tractor and the trailer axles. The approximate dimensions of the air-springs were 10-inch in diameter and 10-inch cylinder height when loaded. Figure 7 shows the installation points for the rear DPU located in a panel at the very rear of the trailer. This figure also shows a close up example of the analog gauges that measure the same pressures as the DPUs.

Figure 7: Photos of an Installed Weighing System.

Data Collection

At various times the truck is weighed by scales approved for commercial use. Several different scales were used, and data was collected whenever a scale was visited. Figure 8 shows the test truck sitting on a commercial scale. When the semi-truck was weighed the driver recorded the following data:

- date
- location
- scale reading (assumed to be “the truth”)
- ambient temperature
- air-spring pressure readings from analog gauges
- estimated weight as computed by the system
- raw pressure and temperature ADC readings from both DPUs

Because this was a working truck, it was only convenient to record data points for full and empty weights. For the same reason, it was not feasible to record separate axle weights. A total of 39 data points were collected with 5 different scales involved.

Figure 8: Experiment Vehicle Being Calibrated on a Scale

Data Analysis

State and federal law requires that every vehicle traveling on public roads satisfy certain weight requirements (U.S.C. 2007). The national recommendation for weight limits of semi-trucks on federal highways is about: 36,287 kg (80000 lbs) for total weight, 15,422 kg (34000 lbs) per tandem axle, and 9,072 kg (20000 lbs) for single axles. Since these requirements were a main motivation in the design of the system, they were used as the main metric for comparison. Requirements like these are a main motivation for drivers to know vehicle weight, and so they are used as the main metric for comparison. The on-board scale system is considered "correct" as long as its weight estimate is within 1% of full weight requirements, which is 363 kg (800 lbs). As stated before, our goal was to see if such a scale could operate "correctly" at least 80% of the time.

The current dataset collected during the trials consists of 39 points, which is too small to accurately estimate the various probabilities of error. Some of the data did not include axle identification with the pressure data; that data was not used to determine a_1 , a_2 and b of Equation 2. However, once these coefficients were determined, proper identification of the remaining axle load data could be accomplished because the range of raw data for the front and rear axles were different enough. With the identification determined with confidence, all data was used for the analysis (all new and future data is properly labeled).

The trials are completely independent of each other, and so to increase the number of tests cases the data were interpreted as potentially occurring in every possible order. This allows the small set of data points to offer many possible sets of training calibrations to compare with recorded weight and pressure measurements. Unfortunately this interpretation provides far too many training possibilities to reasonably compute so 50,000 random permutations per calibration scheme were selected, with the requirement that each would have at least one low weight and one high weight point. Errors were then generated for each permutation by comparing the calibrated estimate to all other data points not used for training.

This process generated a set of errors grouped by the number of calibration points used during the training. Using various tools, such as cumulative distribution functions (CDFs) and probability of tolerance; the errors were visualized in the results section to help quantify how to calibrate an on-board weighing system. The following section considers if the data seems to fit the given models, the most general overall error of the system, comparison of errors at empty and fully loaded, and suggest a possible strategy to decide how to distribute given calibration points to improve accuracy. Using data recorded directly from the prototype system, how well the real-time implementation of the algorithm performed on actual hardware is presented.

Results

Once the data were collected and simulated, the estimation algorithms and their errors were analyzed. The first step was to see if either model fits the collected data; either the model where the tractor and trailer coefficients are assumed to be equal, or the model where they are estimated separately.

Prototype Performance

The cost of the prototype system components was approximately $800. The price would be lower in a production situation, as the parts would come at the much lower bulk prices. The batteries in the DPUs lasted approximately two to three months with typical use. Since the HID's batteries are rechargeable and it is frequently plugged in, its battery life is not known. Also, the prototype system has currently survived several months of active use. The environmentally sealed enclosures have not failed. When the parts of the system were opened, there was no dust, debris, etc. inside. A couple internal connectors shook loose; however, they were fixed in the field with just a screw driver.

The collected data showed that the weight is indeed proportional to the pressure in the air-spring suspension. Figure 9 demonstrates the linear relationship between the sum of the pressure sensors' digital outputs and the semi-truck's weight. This is how the current prototype estimates truck weight.

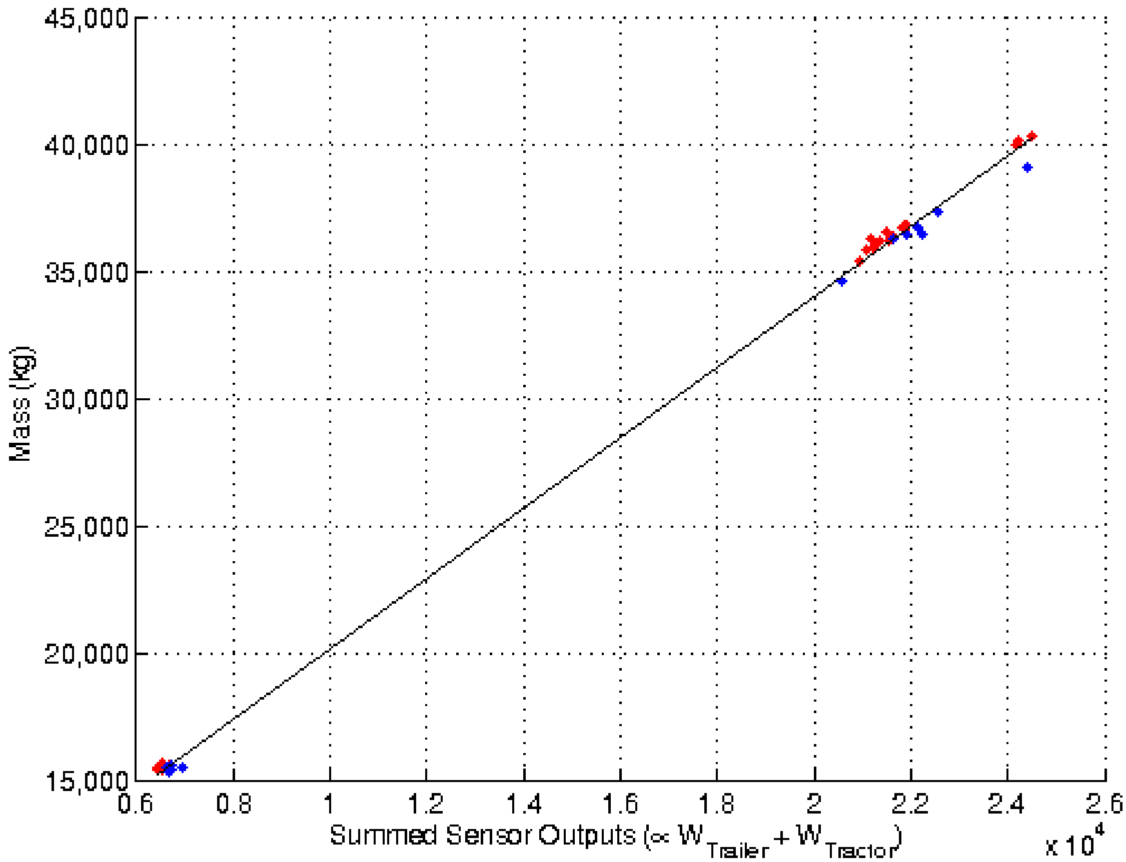

Figure 9: Scatter Plot of On-Vehicle Weighing System Data, Fit to a Line.

Similarly, Figure 10 shows the relationship between the weight and the two sensor outputs treating them as separate variables. This figure shows how well the separate coefficients model fits the data. In Figures 9 and 10 both methods look reliable, the estimate was always pretty close to the measured weight. When using the sensor outputs separately instead of summing them, the estimate stays even closer to the correct weight, which means a system using this method would be even more reliable.

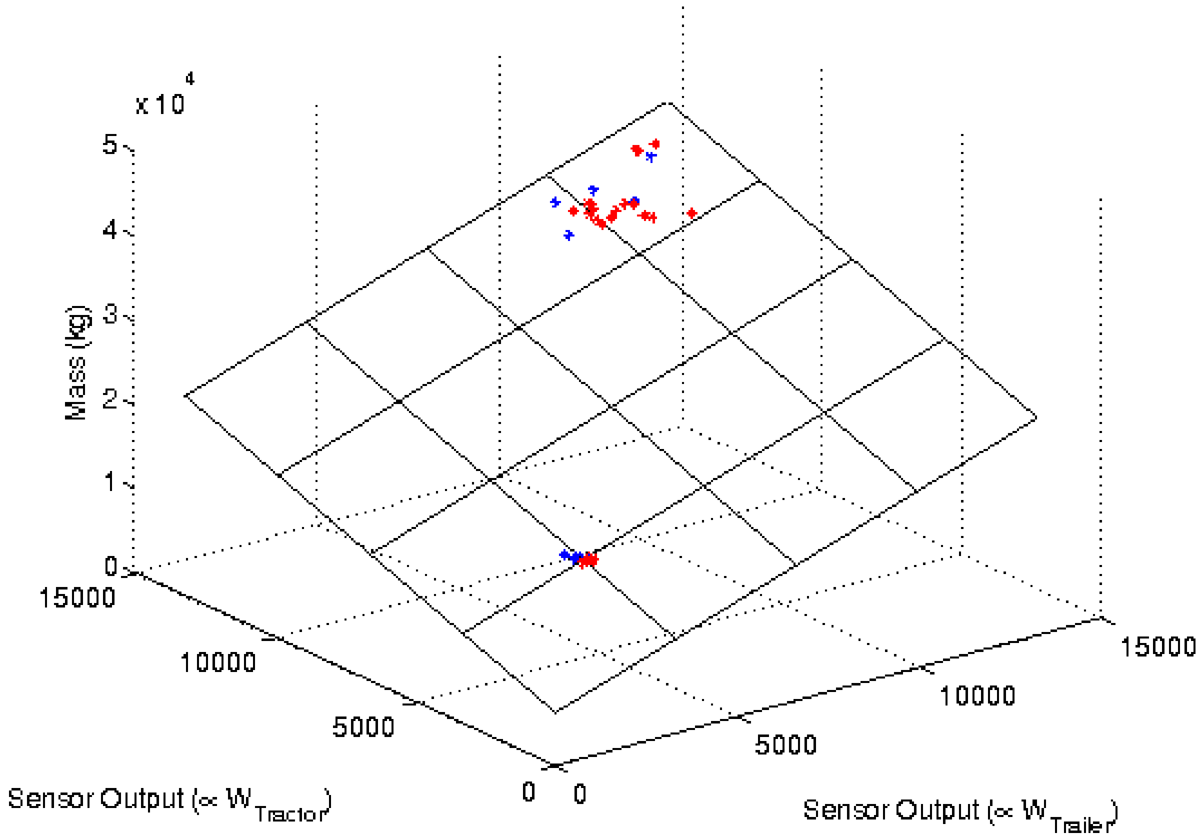

Figure 10: Scatter Plot of On-Vehicle Weighing System Data, Fit to a Plane.

Wireless Reporting

A field test of the wireless was performed on an actual semi-truck, to determine if the estimations could be reported reliably. When the receiver was in the cab, 95% or more of the data packets were successfully received. The reception was excellent when standing alongside the semi-truck, maintaining almost 100% success at up to 45 meters away. Performance dropped significantly directly in front of the cab of the semi-truck. The success rate dropped to 50% at around 12 meters in front of the cab. When the receiver was at either the front or the back or the semi-truck, it still received readings from both of the DPUs.

Calibration Strategy and Implementation

Having observed that the data was a close fit to the model, the next step was to try to evaluate just how close the fit was. The errors of the estimation needed to be measured and characterized. In particular how likely the estimate was to be within an acceptable range, and the performance with respect to the degree of calibration needed to be investigated.

To get an idea of the sorts of error, CDFs were a good place to start. Since all of the collected data points were either empty (the semi-truck had no load) or full (the semi-truck was nearly fully loaded), they were separated into two groups to see the errors in those two cases. Figure 11 shows the CDFs for the errors of those two groups. It was immediately evident that the estimates were better for the empty case than for the full case, and the full errors have a slight bias (non-zero mean). Of course the probability of no error is very small, so the next step in to see how likely the estimates are to be within an acceptable range of error, and how to calibrate it to be within that range.

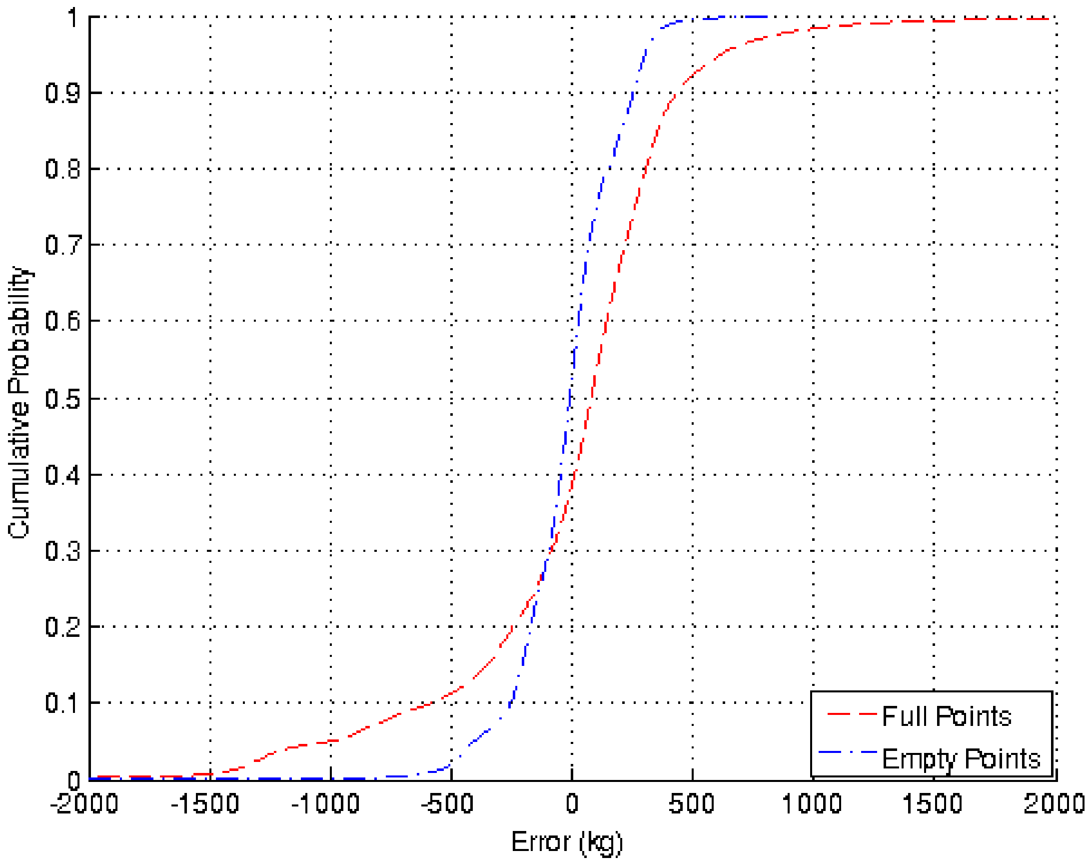

Figure 11: Weight Estimation Error CDFs

As previously mentioned, it was decided that estimates within 1% of the fully loaded weight (363 kg) were good enough to be considered “correct”. Using this threshold it can be determined how many calibrations are needed to maintain it at a given probability. The relation between the number of calibration points and the “correctness” rate is shown in Figure 12. Again the data points were separated by empty or full case.

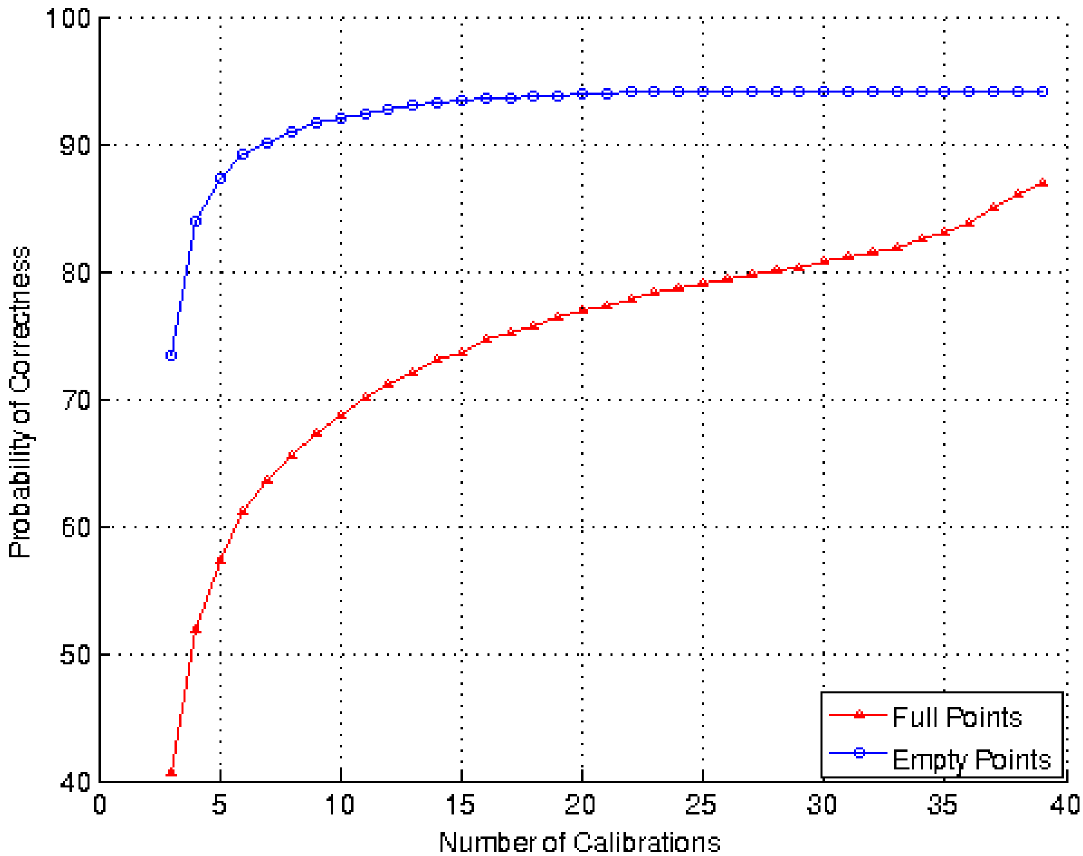

Figure 12: Relation Between Number of Calibrations and Correctness Rate

As expected, the estimates were more likely to be correct the more calibration points were used. The improvement has diminishing marginal returns, which makes sense because as the number of points increases the effect of a single point decreases. After around 20 calibrations the empty estimates show no noticeable increase in performance from further calibration. The full point's performance on the other hand does not level out at any of the calibration level our data allowed us to test.

From Figure 12 the point at which the both empty and full estimates meet the third objective (within 1% of full scale 80% of the time), is at 28 calibration points. This means a conservative level of calibration for meeting the objective would be 30 calibration points. A system using our method calibrated with this many points will work as well as desired.

Figure 13 shows the observed errors in weight estimation of the prototype system. It used the simplified method (summing the sensor outputs). The level of calibration for each data point is not known. The weight estimation errors from the prototype system vaguely resemble those seen in the theoretical analysis of the algorithm. If there were more data, the system's performance could be gauged more accurately and it may then more closely resemble the theoretical performance. The prototype's estimates were very close to achieving the third objective. At the desired frequency of correctness (80%), its errors were within 400 kg. If it were distinguishing between the two sensor outputs and had been calibrated with at least 30 points (as determined above), it would almost certainly achieve the objective.

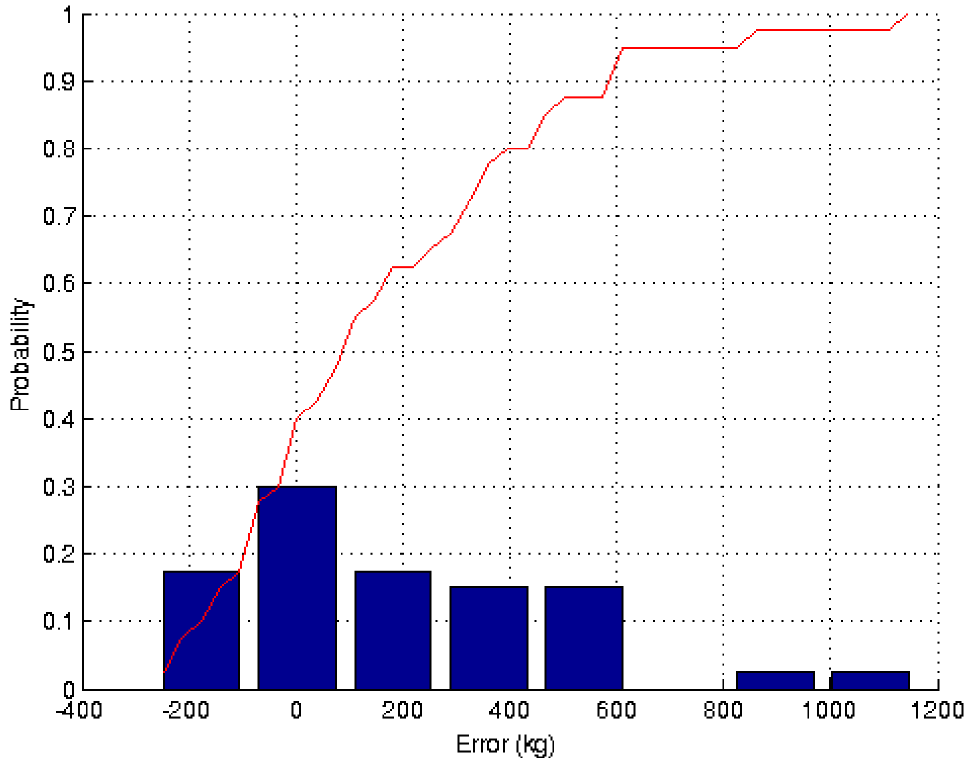

Figure 13: Empirical PMF and CDF of Errors in Estimates Reported by Prototype

Future Work

There are various approximations made in this work, including: all weight not support by an air spring is constant, the vehicle is sitting on flat level ground, and the vehicle is stationary. These assumptions limit the usage of the system and one would naturally like to eliminate them.

When the semi-truck is empty and only the tare weight of the vehicle is being measured the approximation that there is a constant weight being supported by the non-air spring suspension components is viewed as quite good. However, as the truck is loaded a small amount of the cargo weight is transferred to these components and quickly becomes a significant source of error. Modeling this effect, to account for it in the weight estimate, is challenging and is an unknown that requires some knowledge of how the cargo is distributed on the truck. Other small unpredictable weight variations also occur such as changes in fuel level and number of passengers.

It is possible to envision a scheme where the angle of the ground, with respect to the horizon,

can be accounted for in the weight estimates. The direction of W_1 in Figure 2 would change with this angle and would affect the summing of moments (Equation 1). However, a relatively clean measurement of the angle would be required. Accelerometer data was collected by the system and could be used to extend the model presented here.

Weight while stationary is important and useful the typical semi-truck driver. However, being able to measure weight while in motion, particular in bumpy terrain such as a farm field, would expand the usefulness of the system. To do this one can no longer make the assumption that the air-spring suspension assemblies are rigid bodies. A state-based model would need to be used to help average the pressure measurements of the dynamic in a logical way.

In actuality the system measures the load on the vehicle and not the whole vehicle. This would be useful for the weighing part of yield mapping, which usually uses load cells (Lee 2002). Such applications have no need of total vehicle weight, and are unaffected by the weight of the axles, cab, etc. However there are uses where total weight, and more precisely total axle weight, is needed by the user. It is theoretically possible to find the individual axle weights of the vehicle in the same way, given several full and empty axle weight calibrations. The original weight model was used to estimate the air-spring assembly pressure to weight coefficients and introduce a new weight model to estimate the axle weights. This new weight mode is as follows:

$$W_{axle} = m_{axle} P_{axle} + b_{axle} \quad (3)$$

In light of the previous realizations, a new prototype system needs to be developed. Work has begun on one which utilizes a smartphone as a replacement for the HID. This will allow for use of the more accurate model, given more processing power and floating point capability.

Conclusion

The analyses of this paper have shown that the inexpensive and effective method for estimating weight from air-spring suspension pressure presented here works. The prototype system is still in operation and successfully estimates the semi-truck's weight from its air-spring suspension. The estimates are delivered wirelessly in real time over a range of 15 meters, excluding a dead spot directly in front of the cab. The prototype nearly achieved estimation errors within 1% of full scale 80% of the time, and a method for calibrating it to achieve that goal was determined and verified using collected data. With the findings of this paper, the designed system could be updated to perform as desired.

Acknowledgements

Support from the Purdue School of Electrical and Computer Engineering was provided for this senior design project. Partial support for this work from the USDA NIFA grants program for the project titled "Improving Agricultural Management with Autogenic Mobile Technology" is acknowledged and appreciated. Cooperation of operators and managers at Ault Farms for device testing was also appreciated.

References

AirWeigh. 2011. Truck scales. http://www.air-weigh.com.au/allscales.html. Accessed 7/15/12.

Farmtronics. 2012. Loadmax Truck Scale Kit. http://www.farmtronics.com/proddetail.php?prod=E94720&cat=184&PHPSESSID=902bdf23cdd848ed0f87826fb3af6e40 Accessed 7/5/12.

Lee, W. S., Burks T. F., Schueller, J. K. 2002. Silage yield monitoring system.

Masser, L. D. 1961. Suspension for automotive vehicles. U.S. Patent No. 3,140,880.

Pierce, P. R. and J. C. Cervantez. 2009. Height control valve for vehicle axle/suspension system. U.S. Patent No. 7,621,537.

TruckWeight. 2011. Onboard Scales. http://www.truckweight.com/EN/New/products.html. Accessed 7/5/12.

U.S.C. 2007. 23 § 658.17