

Detection of Anomalous Vehicle Loading

by

Troy McKay

B.A. SUNY Geneseo, 2003

M.S. SUNY Buffalo, 2006

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Chester F. Carlson Center for Imaging Science

College of Science

Rochester Institute of Technology

August 5, 2016

Signature of the Author _____

Accepted by _____
Coordinator, Ph.D. Degree Program Date

CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE
COLLEGE OF SCIENCE
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

Ph.D. DEGREE DISSERTATION

The Ph.D. Degree Dissertation of Troy McKay
has been examined and approved by the
dissertation committee as satisfactory for the
dissertation required for the
Ph.D. degree in Imaging Science

Dr. Carl Salvaggio, Dissertation Advisor

Dr. Edward Hensel, External Chair

Dr. Alfred Garrett, Affiliated Professor

Dr. David Messinger

Date

Detection of Anomalous Vehicle Loading

by

Troy McKay

Submitted to the
Chester F. Carlson Center for Imaging Science
in partial fulfillment of the requirements
for the Doctor of Philosophy Degree
at the Rochester Institute of Technology

Abstract

Determining the mass of a vehicle from ground based passive sensor data is important for many security and traffic safety reasons. A vehicle consists of multiple dependent and independent systems that each respond differently to changes in vehicle mass. In some cases, the responses of these vehicle systems can be measured remotely. If these remotely sensed system responses are correlated to the vehicle's mass, and the required vehicle parameters were known, it would be possible to calculate the mass of the vehicle as a function of these responses.

The research described here investigates multiple vehicle phenomenologies and their correlation to vehicle load. Brake temperature, engine acoustics, exhaust output, tire temperature, tire deformation, vehicle induced ground vibration, suspension response, and engine torque induced frame twist were all evaluated and assessed as potential methods of remotely measuring a vehicle's mass. Extensive field experiments were designed and carried out using multiple sensors of various types; including microphones, accelerometers, high-speed video cameras, high-resolution video cameras, LiDAR, and thermal imagers. These experiments were executed at multiple locations and employed passenger vehicles, and commercial trucks with loads ranging from zero to beyond the recommended load capacity of the vehicle. The results of these experiments were used to determine if the signature for each phenomenology could be accurately observed remotely, and if so, how well they correlated to vehicle mass. The suspension response and engine torque induced frame twist phenomenologies were found to have the best correlation to vehicle mass of the phenomenologies considered, with correlation values of 90.5 % and 97.7 %, respectively.

Physics-based models were built for both the suspension response, and the engine torque induced frame twist phenomenologies. These models detailed the relationship between each phenomenology and the mass of the vehicle. Full-scale field testing was done using improved remote detection methods, and the results were used to validate the physics-based models. The results of the full-scale field testing showed that both phenomenology could accurately calculate the mass of the vehicle remotely, given that certain vehicle parameters were accurately known. The engine torque induced frame twist phenomenology was able to find the mass of the test vehicle to within 10 % of the true mass. Using the suspension response phenomenology the mass was accurately predicted as a function of its location on the vehicle. For either phenomenology to be effective, certain vehicle parameters must be known accurately; specifically the spring constant and damping coefficients of the vehicle's suspension, the unloaded mass, the unloaded center of gravity, and the unloaded moment of inertia of the vehicle. The models were also used to propagate measurement and parameter uncertainty through the vehicle mass calculation to arrive at the uncertainty in the mass estimation. Finally, the results of both the phenomenologies were combined into a single vehicle mass estimate with a smaller uncertainty than the individual vehicle system estimations taken alone.

Acknowledgements

I would like to thank all the people that helped me with this work, without them I never would have finished. Dr. Carl Salvaggio has been the best advisor I could have hoped for. He constantly went above-and-beyond the responsibilities of an advisor and I will sincerely miss our weekly meetings. He provided guidance, encouragement, and support throughout my graduate career. Dr. Al Garrett always had great ideas to test out, and his enthusiasm helped drive me along the way. Dr. Ed Hensel's guidance was invaluable to the mechanical modeling work. Dr. David Messinger agreed to be on my committee and always made time to talk about any issues I had.

This work required several field tests, and I was very lucky to have a group of hardworking people to help out. Whether it was November in upstate New York or Summer in South Carolina I always had a great support team. I'd like to thank Jason Faulring for all of his help and expertise during this project, I couldn't have done it without him. Jason supported every field collect and worked tirelessly to ensure each was a success. Dr. Dave Coleman, Dr. Larry Koffman, and Dr. Si Young worked hard to facilitate the field collections at the Savannah River Site and provided on-site support. Several RIT students and faculty helped out at the various field collections and their contributions made everything go much smoother. Thank you to Don McKeown, Glenn Sweeny, Dr. Phil Salvaggio, and Steve Cavillia for your help with the sensors. Thank you to Colin Axel and Ryan Laclair for your help processing the data.

Finally, I would like to thank my family. My mom, dad, and brother were always there for me, and encouraged me when it seemed I would never finish. Thank you to Courtney who supported me throughout all the nights and weekends of coding and writing, and thank you to Katherine for sharing me with this work.

*This work is dedicated to Dr. David Meisel who taught me to love science and started me down
this path.*

Contents

1	Introduction	1
2	Theory	4
2.1	Physics	4
2.1.1	Linear Kinetic Energy	4
2.1.2	Rotational Kinetic Energy	4
2.1.3	Damped harmonic oscillator	5
2.1.4	Heat Transfer	8
2.1.5	Ideal Gas	10
2.2	Instruments and Sensors	11
2.2.1	Thermocouple	11
2.2.2	Microbolometer Infrared Camera System	12
2.2.3	LiDAR Sensing Theory	13
2.3	Internal Combustion Engine	13
2.3.1	Two-Stroke	14
2.3.2	Four-Stroke	14
2.4	Mathematics	15
2.4.1	Euclidean Distance	15
2.4.2	Mahalanobis Distance	16
2.4.3	Homogeneous Coordinate System	17
2.4.4	Hessian Matrix	18

2.4.5	Haar Transformation Matrix	18
2.4.6	RANdom SAMple Consensus	20
2.5	Imaging Science	21
2.5.1	Morphological Processing	21
2.5.2	Cross-Correlation	24
2.5.3	One-Dimensional Special Functions	27
2.5.4	Rectangular Function	27
2.5.5	Comb Function	28
3	Phenomenologies	30
3.1	Acceleration and Deceleration Rate	30
3.1.1	Measurement Techniques	31
3.2	Brake Temperature	41
3.2.1	Measurement Techniques	44
3.3	Engine Acoustics	46
3.3.1	Measurement Techniques	48
3.4	Engine Heat Generation	48
3.4.1	Measurement Techniques	50
3.5	Exhaust Plume	50
3.5.1	Measurement Techniques	52
3.6	Frame Twist	53
3.6.1	Measurement Technique	55
3.7	Suspension Response	55
3.7.1	Measurement Techniques	56
3.8	Tire Deformation	60
3.8.1	Measurement Techniques	61
3.9	Tire Temperature	62
3.9.1	Measurement Techniques	67
3.10	Vibration Response	67
3.10.1	Measurement Techniques	68

4	Collection Campaign	69
4.1	Wyoming County International Speedway Full-Scale Field Test	69
4.1.1	Test Vehicle	69
4.1.2	Test Site	71
4.1.3	Test Description	72
4.1.4	On-Board Sensor System	74
4.1.5	Off-Board Sensor System	83
4.1.6	Test Execution	85
4.2	Savannah River Site Full-Scale Field Test I	90
4.2.1	Test Vehicle	90
4.2.2	Test Site	90
4.2.3	Test Description	91
4.2.4	On-Board Sensor System	94
4.2.5	Off-Board Sensor System	97
4.2.6	Test Execution	100
4.3	Savannah River Site Full-Scale Field Test II	104
4.3.1	Test Vehicles	105
4.3.2	Test Site	105
4.3.3	Test Description	106
4.3.4	On-Board Sensor System	107
4.3.5	Off-Board Sensor System	107
4.3.6	Test Execution	108
4.4	Small-Scale Testing	111
4.4.1	Test Vehicle	111
4.4.2	Test Site	111
4.4.3	Test Plan	112
4.4.4	On-Board Sensor System	113
4.4.5	Off-Board Sensor System	113
4.4.6	Test Execution	114

5	Experimental Results and Modeling	118
5.1	Brake Temperature	119
5.1.1	Wyoming County International Speedway Full-Scale Field Test	119
5.1.2	Savannah River Site Full-Scale Field Test I	128
5.2	Engine Acoustics	136
5.2.1	Wyoming County International Speedway Full-Scale Field Test	136
5.2.2	Savannah River Site Full-Scale Field Test I	140
5.3	Exhaust Braking	144
5.3.1	Savannah River Site Full-Scale Field Test I	145
5.4	Tire Temperature	147
5.4.1	Savannah River Site Full-Scale Field Test I	147
5.5	Vibration Response	156
5.5.1	Savannah River Site Full-Scale Field Test I	156
5.6	Tire Deformation	162
5.6.1	Small-Scale Testing I	162
5.6.2	Small-Scale Test II	167
5.6.3	Wyoming County International Speedway Full-Scale Test	174
5.6.4	Savannah River Site Full-Scale Field Test I	176
5.6.5	Savannah River Site Full-Scale Field Test II	184
5.7	Suspension Response	192
5.7.1	Wyoming County International Speedway Full-Scale Field Test	192
5.7.2	Savannah River Site Full-Scale Field Test I	205
5.7.3	Physics-Based Model	216
5.7.4	Savannah River Site Full-Scale Field Test II	237
5.8	Frame Twist	248
5.8.1	Physics-Based Model	249
5.8.2	Savannah River Site Full-Scale Field Test I	253
5.8.3	Savannah River Site Full-Scale Field Test II	262
5.9	Combined Phenomenology and Error Analysis	272

5.9.1	Suspension Response Error Analysis	273
5.9.2	Frame Twist Error Analysis	283
5.9.3	Combined Phenomenology	286
6	Conclusions	289
6.1	Brake Temperature	289
6.2	Engine Acoustics	290
6.3	Engine Exhaust	290
6.4	Tire Temperature	291
6.5	Vehicle-Induced Ground Vibration	291
6.6	Tire Deformation	292
6.7	Suspension Response	292
6.8	Engine Torque-Induced Frame Twist	293
6.9	Combined Phenomenology	294
6.10	Major Accomplishments	295
6.11	Future Work	295
6.11.1	Tire Deformation	296
6.11.2	Suspension Response	296
6.11.3	Frame Twist	296
A	Suspension Response Model Derivation	298

List of Figures

2.1	Calculated emitted spectral radiance for a blackbody at a temperature of 800K . .	10
2.2	Structure of a typical bolometer	13
2.3	Diagram outlining the operation of a two-stroke internal combustion engine [1] .	14
2.4	Diagram outlining the operation of a typical four-stroke internal combustion engine [2]	15
2.5	A set of data, normally distributed about the origin with a standard deviation of 25 along the x-axis and 5 along the y-axis (blue). The point at (0,30) in green is a Euclidean distance of 30 and a Mahalanobis distance of 35.6 from the data set mean. The point at (30,0) in red is a Euclidean distance of 30 and a Mahalanobis distance of 1.5 from the data set mean.	17
2.6	Example of two binary images	21
2.7	Example of the union and intersection operator	22
2.8	Example of the union and intersection operators	23
2.9	Example of the closing and opening operators	24
2.10	An image and a template image	25
2.11	Cross-correlation of image I and template T	26
2.12	Normalized correlation coefficient for image I and the template T	27
2.13	The rectangle function $RECT\left(\frac{x-x_0}{b}\right)$	28
2.14	The comb function $COMB\left(\frac{x-x_0}{b}\right)$ where the vertical axis denotes the area of the Dirac delta function	29

3.1	Imagery taken using an uncalibrated camera system, highlighted in the image are the lines used to determine the vanishing point and the location of the vanishing point [3]	33
3.2	RIT DIRS laboratory camera calibration facility	36
3.3	Diagram of a typical disc brake system	43
3.4	Diagram of a typical drum brake system	44
3.5	Operation of a 4-stroke gasoline powered internal combustion engine [2]	46
3.6	A schematic diagram of the test setup used by Burgard and Bishop et al. [4].	53
3.7	A block diagram of the data processing steps used to track position from video using the custom video tracking algorithm	59
3.8	The region to track (left) and a edge image of the region to track (right)	59
3.9	An array of binary, edge images generated by varying the contrast stretch (mild contrast stretch on the left to severe contrast stretch on the right)	60
3.10	An example of the intensity image created from the normalized cross-correlation of a search area image and the template image (both binary, edge images)	60
3.11	Proposed configuration for one-dimensional scanning LiDAR system to measure tire contact area	62
3.12	Stress-strain plots for ideal elastic material (left) and a viscoelastic material (right). [5]	63
3.13	Predicted force function for a position on the tire tread for a loaded and unloaded tire traveling at the same velocity	65
4.1	The 1992 Ford F600 chosen for as a test vehicle	70
4.2	The Wyoming County International Speedway	71
4.3	Portable speed bumps acquired for the suspension response test	73
4.4	The location of the thermocouples installed on the rear brakes	75
4.5	The location of the thermocouples installed on the front brakes	75
4.6	Location of the thermocouple installed on the exhaust	76
4.7	Location of the thermocouples installed on the rear differential	76

4.8	Location of the thermal imager installed on the rear of the truck, directed at the rear brakes	77
4.9	Location of the thermal imagers installed in the front, driver side wheel-well of the truck, directed at the front brakes	77
4.10	Location of the triaxial accelerometers that were installed on the rear axle and the frame	78
4.11	Location of the microphone installed in the engine compartment directed at the valve train	79
4.12	Location of the microphone installed on the frame of the test vehicle directed at the exhaust	80
4.13	One of the six wireless tire pressure monitors that were used to monitor tire pressure	81
4.14	One of two LEDs installed on the bed of the truck which were used to synchronize the data collected by the off-board system with the data collected by the on-board system	82
4.15	The computer and data acquisition system installed in the cab of the truck	83
4.16	The WASP system setup to monitor the test vehicle	84
4.17	The test vehicle being loaded at the track	86
4.18	The test vehicle being weighed	87
4.19	A fiducial marker being added to the frame of the test vehicle	88
4.20	The suspension response test in progress for the test vehicle loaded with 3 concrete blocks	89
4.21	The tractor-trailer chosen for Savannah River Site field tests	90
4.22	The Savannah River Site testing location	91
4.23	Location of the thermal imager installed on the front driver's side of the tractor directed at the front passenger's side brakes and tire	95
4.24	Location of the thermal imager installed on the front passenger's side of the tractor directed at the rear driver's side brakes and tire	95
4.25	Location of the microphones installed on the test vehicle trailer directed at the tractor	96
4.26	The computer and data acquisition system installed on the test vehicle's trailer . .	97

4.27	The location of RIT's WASP sensor during testing	98
4.28	Location of the triaxial accelerometers that were mounted on the railroad tracks	99
4.29	The camera and laser trigger used to capture images of the test vehicle's tires while moving	100
4.30	The suspension response test in progress for the test vehicle loaded with 40 tons, the arrow marks the location of the speed bump	103
4.31	The vibration response test in progress for the test vehicle loaded with 40 tons, the circles mark the locations of the accelerometers	104
4.32	The Honda Accord used for this field test	105
4.33	The instrument configuration for the full-scale field test conducted at The Savannah River Site in December of 2012	110
4.34	The LiDAR system position and orientation used for the first small-scale experiment	114
4.35	The rear tires of the test vehicle positioned on foam-board prior to measuring contact area	115
4.36	Several volunteers loaded on the flatbed of the test vehicle	116
5.1	Gray-scale LWIR image of the front driver's side disk brake collected by the on-board imager and the tracking ROIs	120
5.2	Gray-scale LWIR image of the rear passenger's side drum brake collected by the on-board imager and the tracking ROIs	120
5.3	Comparison of on-board temperature and radiance measurements.	123
5.4	The characteristic values of the driver side front disc radiance signal in digital counts. Note that the final radiance was chosen to be 3 seconds after the vehicle was stopped.	124
5.5	The characteristic values of the passenger side rear drum radiance signal in digital counts. Note that the final radiance was chosen to be 3 seconds after the vehicle was stopped.	125
5.6	Gray-scale LWIR images of the tractor's brakes collected by the on-board imager	129
5.7	Gray-scale LWIR image of the test vehicle immediately upon coming to a complete stop collected with WASP	130

5.8	Location of the areas-of-interest used to track the location and radiance of the hottest regions of the test vehicle's brakes.	132
5.9	Location of the area-of-interest used in the WASP imagery to track the location and radiance of the hottest region of the test vehicle's brake.	133
5.10	The characteristic values of the passenger side front brake radiance signal in digital counts as measured by the on-board IR sensor.	134
5.11	The characteristic values of the passenger side front brake radiance signal in digital counts as measured by the off-board IR sensor (WASP)	135
5.12	Spectrogram from the on-board microphones for the test vehicle loaded with one block accelerating from a stop	137
5.13	A spectrogram from the on-board microphone located near the exhaust of the test vehicle, collected while idling. The top line illustrates how the fundamental frequency was calculated from the average distance between local maxima (harmonics) in the frequency domain. The bottom 9 lines show the location of the first nine harmonics.	138
5.14	Spectrogram from the on-board and off-board microphones for the test vehicle loaded with two blocks accelerating from a stop	140
5.15	Spectrogram from the on-board microphones for the test vehicle loaded with one block accelerating from a stop	142
5.16	Location of the fifth, sixth and seventh harmonic of the test vehicle. The harmonics took place at 29.2, 35.1 and 40.9 Hz, respectively, meaning that the first harmonic occurs at 5.9 Hz and the test vehicle's engine is operating at 700 RPM.	142
5.17	The calculated engine RPM for the test vehicle as a function of time	143
5.18	Spectrogram from the on-board and off-board microphones for the unloaded test vehicle accelerating from a stop	144
5.19	Spectrogram from the test vehicle decelerating to a stop with and without the exhaust brake	146
5.20	Gray-scale LWIR images collected by the on-board imagers and the corresponding tracking ROIs	149

5.21	Minimum, maximum, and median tire radiance for a fully loaded (40 ton) test vehicle.	150
5.22	Location of the area-of-interest used in the WASP imagery to track the location and calibrated radiance of the hottest region of the test vehicle's tire tread.	151
5.23	The characteristic values of the driver side front tire tread radiance signal in digital counts as measured by the off-board IR sensor (WASP).	152
5.24	Plots of the test vehicle's velocity, and the test vehicle's tire tread uncalibrated radiance as measured by the on-board sensors.	154
5.25	Air temperature measurements during the field testing (from Aiken, SC weather station)	155
5.26	Off-board video frame collected of the test vehicle traveling over the railroad tracks	157
5.27	Image frame of the tractor-trailer traveling over the railroad tracks. The distances L1, L2, L3 were manually measured (in pixels)	158
5.28	The acceleration magnitude experienced by the railroad track when the first wheel of the tractor-trailer crosses it	159
5.29	An example of a tire footprint outline, measured by spray painting around the perimeter of the tire	163
5.30	A plot of measured contact area of the inner and outer rear driver side tires versus added vehicle load	164
5.31	A three-dimensional point cloud aligned, rotated, scaled, and translated to best show the test vehicle's tire	165
5.32	Plot of the measured total footprint area of the rear driver side tires (measured from the footprint outlines), and the predicted total footprint area versus total vehicle load	166
5.33	Plot of each tire's footprint length and width as a function of total vehicle load (for the rear driver side two tires)	167
5.34	An example of a tire footprint outline, measured by spray painting around the perimeter of the tire	168
5.35	A three-dimensional point cloud aligned, rotated, scaled, and translated to best show the test vehicle's tire	170

5.36	Plot comparing the calculated test vehicle mass (derived from the footprint outlines) versus the actual test vehicle mass	171
5.37	Plot showing three adjacent horizontal lines of LiDAR data at the road/tire interface as the tire travels through the field-of-view	172
5.38	Plot showing three adjacent horizontal lines of LiDAR data at the road/tire interface as the tire travels through the field-of-view	173
5.39	Measured tire pressure for each of the test vehicle's tires as a function of time . .	175
5.40	Linear trends of the measured tire pressure for each of the test vehicle's tires as a function of time	176
5.41	Off-board image collected of the front passenger side test vehicle wheel while the test vehicle was traveling at a constant velocity	177
5.42	An example of the tire image after the high-pass filter was applied (left) and after the thresholding was applied (right)	178
5.43	An example of the tire image after the ten largest regions in the image were removed (left) and after the morphological closing was applied (right)	179
5.44	An example of the tire image after the flood fill operator was applied (left) and the most elliptical region (right)	179
5.45	The perimeter of the detected ellipse (blue), the major axis (green), and the minor axis (red)	180
5.46	The processed image, where the tire rim has been rectified	181
5.47	The processed image, with the measured rim to tire edge distances displayed . . .	183
5.48	Three dimensional point cloud rotated so the Z plane is on the road surface . . .	185
5.49	Visual representation of how the X axis is adjusted from raw unscaled units on the left to units of mm on the right by fine adjustments until the rim appears round . .	186
5.50	Comparison of undeformed horizontal slices (red & blue) and a deformed vertical slice (green)	186
5.51	Point cloud with intensity data overlaid	188
5.52	False color image with wheel/tire interface and road surface identified	189

5.53	Deformed tire thickness from wheel 6 (last trailer wheel) of the tractor-trailer tested during the full-scale field test at the Savannah River Site	190
5.54	The tire specifications printed on the sidewall	192
5.55	Off-board video collected of the test vehicle traveling over the speed bump; the fiducials added to the frame of the test vehicle are circled in red. Fiducials can also be seen on the wheels of the test vehicle.	193
5.56	Off-board video collected of the test vehicle traveling over the speed bump, the suspension response (the distance between the frame and the fiducial) is marked in red.	195
5.57	A block diagram of the processing steps used to convert the movie files into suspension response signals	195
5.58	Smoothed rear suspension displacement from static equilibrium as the vehicle's rear axle traveled over a speed bump, gray area represents the time the tire was on the speed bump	196
5.59	First derivative (velocity) of the rear suspension displacement as the vehicle's rear axle traveled over a speed bump, gray area represents the time the tire was on the speed bump	197
5.60	Second derivative (acceleration) of the rear suspension displacement as the vehicle's rear axle traveled over a speed bump, gray area represents the time the tire was on the speed bump	197
5.61	Points on the front suspension displacement plot where the supported mass ratios were calculated, green Xs denotes points where spring is in the equilibrium position, red Xs denote points where the first derivative will be zero	200
5.62	Points on the rear suspension displacement plot where the supported mass ratios were calculated, green Xs denotes points where spring is in the equilibrium position, red Xs denote points where the first derivative will be zero	200
5.63	Front suspension m/k_d calculations for three different vehicle loads	201
5.64	Rear suspension m/k_d calculations for three different vehicle loads	202
5.65	Front suspension m/k_s calculations for three different vehicle loads	202

5.66	Rear suspension m/k_s calculations for three different vehicle loads	203
5.67	Off-board video frame collected of the test vehicle traveling over the speed bump	205
5.68	A block diagram of the processing steps used to convert the movie files into suspension response signals	207
5.69	The locations on the test vehicle that were tracked throughout the video frames; A & B, C & D, and E & F each make up a suspension pair	207
5.70	Rear of the trailer suspension displacement from static equilibrium as the trailer's rearmost axle traveled over a speed bump, gray area represents the time the tire was on the speed bump	208
5.71	First derivative (velocity) of the rear of the trailer suspension displacement from static equilibrium as the trailer's rearmost axle traveled over a speed bump, gray area represents the time the tire was on the speed bump	209
5.72	Second derivative (acceleration) of the rear of the trailer suspension displacement from static equilibrium as the trailer's rearmost axle traveled over a speed bump, gray area represents the time the tire was on the speed bump	210
5.73	Location of the pair of points with values of 0.01 m/s on the velocity versus time plot for the rearmost trailer suspension loaded with 40 tons	212
5.74	Location of the pair of points with values of 0.008 m on the suspension position versus time plot for the rearmost trailer suspension loaded with 40 tons	212
5.75	Suspension m/k_s calculations versus total vehicle load for three different vehicle suspensions	213
5.76	Suspension m/k_d calculations versus total vehicle load for three different vehicle suspensions	214
5.77	Simplified suspension diagram for a standard two-axle vehicle	218
5.78	Example of the front axle traveling over a rise in the road	219
5.79	Example of the rear axle traveling over an indent in the road	219
5.80	Example of the completely unloaded vehicle traveling over flat road (both springs are in the uncompressed state)	220

5.81	Example of the loaded vehicle traveling over a flat road (both springs are in a compressed state)	221
5.82	The vertical position of the front axle (y_{A1}), and rear axle (y_{A2}) traveling over a simulated speed bump	224
5.83	The vertical position of the vehicle frame at the front suspension and rear suspension attachment points (y_{F1} and y_{F2}) as the vehicle travels over a simulated speed bump	225
5.84	The vertical position of the vehicle frame at the front suspension relative to the static equilibrium position (y_{M1}) as the vehicle travels over a simulated speed bump	230
5.85	The orientation of the vehicle frame relative to the static equilibrium orientation (θ_M) as the vehicle travels over a simulated speed bump	230
5.86	Heat-map showing the position error as a function of the mass of the load (kg) and the center of gravity of the load (m) (correct value denoted by red dot)	231
5.87	Heat-map showing the orientation error as a function of the mass of the load (kg) and the center of gravity of the load (m) (correct value denoted by red dot)	232
5.88	Heat-map showing the averaged error as a function of the mass of the load (kg) and the center of gravity of the load (m) (correct value denoted by red dot)	232
5.89	The vertical position of the vehicle frame at the front suspension relative to the static equilibrium position (y_{M1}) as the vehicle travels over a simulated speed bump with simulated measurement error (± 1 mm)	233
5.90	The angular orientation of the vehicle frame relative to the static equilibrium position (θ_M) as the vehicle travels over a simulated speed bump with simulated measurement error (± 1 mm)	234
5.91	Heat-map showing the position error as a function of the mass of the load (kg) and the center of gravity of the load (m) for data with simulated measurement error (correct value denoted by red dot)	235
5.92	Heat-map showing the orientation error as a function of the mass of the load (kg) and the center of gravity of the load (m) for data with simulated measurement error (correct value denoted by red dot)	236

5.93 Heat-map showing the averaged error as a function of the mass of the load (kg) and the center of gravity of the load (m) for data with simulated measurement error (correct value denoted by red dot)	237
5.94 Off-board video frame collected of the passenger vehicle traveling over the speed bump	238
5.95 Off-board video frame collected of the tractor-trailer traveling over the speed bump	238
5.96 Off-board video frame collected of the tractor-trailer traveling over the speed bump	239
5.97 Off-board video frame collected of the tractor-trailer traveling over the speed bump	239
5.98 Modeled data fit to the measured y-axis wheel position	241
5.99 Example data set for the front of the passenger car	242
5.100 Example data set for the rear of the passenger car	242
5.101 Example data set for the front of the tractor-trailer	243
5.102 Suspension m/k_d calculations versus total vehicle load for three different vehicle suspensions	245
5.103 Suspension m/k_d calculations versus total vehicle load for three different vehicle suspensions	247
5.104 Longitudinal engine imparting torque, τ_E , to the vehicle frame causing it to twist until the torque generated by the suspension forces (F_1 and F_2) are sufficient to match	250
5.105 The modeled position and velocity of the test vehicle versus time used for all simulated data	252
5.106 The modeled position and velocity of the test vehicle versus time used for all simulated data	253
5.107 Off-board video frame collected of the test vehicle accelerating from a stop . . .	254
5.108 A block diagram of the processing steps used to convert the movie files into frame twist signals	255
5.109 An image frame with the extracted grill region identified	255

5.110	The left image is the unrotated grill image, the middle plot shows the vertical variability values for a range of image rotations (-5° to 5°), and the right image shows the grill image rotated so the grill runs perpendicular to the horizontal image axis.	256
5.111	Left plot shows the grill intensity profile, middle plot shows the demeaned intensity profile with a Hamming window applied, and the right plot shows the power spectra of the middle plot	257
5.112	From left-to-right the raw pressure signal, the pressure signal with a Hamming window applied, and the power spectra (during idle)	258
5.113	From left-to-right; the angular frame twist, position, and engine speed of the unloaded test vehicle accelerating from a stop	259
5.114	The calculated mass to spring constant ratio for each test run versus the true test vehicle mass	260
5.115	The calculated mass versus the true test vehicle mass	261
5.116	Example frame from the Phantom Video camera	263
5.117	A block diagram of the processing steps used to convert the movie files into frame twist signals	264
5.118	The regions on the test vehicle tracked by the algorithm	264
5.119	The tracked points on the bumper used to calculate the vehicle to camera distance	265
5.120	From left-to-right the raw pressure signal, the pressure signal with a Hamming window applied, and the power spectra (during idle)	266
5.121	An example of the on-board video of the vehicle's instrument panel captured during each test run	267
5.122	From left-to-right; the angular frame twist, position, and engine speed of the unloaded test vehicle accelerating from a stop.	267
5.123	The calculated mass to spring constant ratio for each suitable test run versus the true test vehicle mass	269
5.124	The calculated mass for each suitable test run versus the true test vehicle mass	270

5.125	The simulated suspension response data used to characterize the effect of error on the mass estimation	274
5.126	Average error heat map for the simulated data calculated with the true parameter values	275
5.127	Average error heat map for the front spring constant k_{s1} with -70% error (left) and 70% error (right)	276
5.128	Average error heat map for the front spring constant k_{d1} with -28% error (left) and 28% error (right)	277
5.129	Average error heat map for the rear spring constant (k_{s2}) with -5% error (left) and 5% error (right)	278
5.130	Average error heat map for the rear damping coefficient (k_{d2}) with -5% error (left) and 5% error (right)	279
5.131	Average error heat map for the moment of inertia of the unloaded vehicle (J_0) with -3% error (left) and 3% error (right)	280
5.132	Diagram of the observation geometry meeting the minimum spatial resolution and field-of-view requirements for a passenger vehicle	282
5.133	Diagram of the observation geometry meeting the minimum spatial resolution and field-of-view requirements for a tractor-trailer	282
5.134	Diagram of the observation geometry meeting the minimum spatial resolution and field-of-view for a tractor-trailer. The left diagram shows the vehicle prior to acceleration and the right diagram shows the vehicle after traveling 6 m.	285
5.135	The simulated suspension response data used to characterize the effect of error on the mass estimation	286
5.136	The simulated suspension response data used to characterize the effect of error on the mass estimation with the frame twist error bars added	287
5.137	The simulated suspension response data used to characterize the effect of error on the mass estimation with regions excluded based on the frame twist analysis . . .	288

List of Tables

2.1	Values for k , p , and q when $N = 4$	20
4.1	The test vehicle loading conditions	86
4.2	The test vehicle loading conditions for the small-scale test	116
4.3	The test vehicle loading conditions	117
4.4	The test vehicle tire inflation pressures	117
5.1	Results from the Brake Test with transmission disengaged; where Peak is the peak radiance, Peak - Start is the difference between the peak radiance and the initial radiance, Final is the radiance 3 seconds after the vehicle stops, Final - Start is the difference between the radiance 3 seconds after the vehicle stops and the initial radiance, and Min is the minimum radiance	126
5.2	Results from the Brake Test with transmission engaged; where Peak is the peak radiance, Peak - Start is the difference between the peak radiance and the initial radiance, Final is the radiance 3 seconds after the vehicle stops, Final - Start is the difference between the radiance 3 seconds after the vehicle stops and the initial radiance, and Min is the minimum radiance	127
5.3	Results from the normalized correlation coefficient calculations that show the correlation between each signal parameter and the kinetic energy of the test vehicle .	127

5.4	Data calculated from each of the railroad track test runs; where Peak is the maximum total acceleration magnitude, AUC is the area under the curve of the total acceleration magnitude plot, and FFT Peak is the location of the peak frequency (frequency with the highest energy) in Hertz	160
5.5	The results of the data correlation analysis, the highlighted cells show substantial correlation to vehicle kinetic energy	161
5.6	Table of the internal tire pressure, vehicle load, and contact area for all of the test vehicle tires.	169
5.7	Correlation coefficient between vehicle weight and tire deformation for the tractor-trailer	191
5.8	Correlation coefficient between vehicle weight and tire deformation for the passenger vehicle measured	191
5.9	Normalized correlation coefficient for the calculated ratio m/k_d to vehicle mass .	203
5.10	Normalized correlation coefficient for the calculated ratio m/k_s to vehicle mass .	204
5.11	Normalized correlation coefficient for the calculated ratio m/k_s to vehicle mass .	215
5.12	Normalized correlation coefficient for the calculated ratio m/k_d to vehicle mass .	215
5.13	Table of passenger vehicle parameters	244
5.14	Table of tractor-trailer parameters	246
5.15	Recommended maximum error values for the suspension response estimated parameters	280

Chapter 1

Introduction

Estimating the total mass of a vehicle using only ground-based remote sensing techniques presents a substantial challenge. A vehicle is a complex machine consisting of multiple dependent and independent systems each with varying responses to changes in vehicle mass. Single-system physics-based or analytical modeling approaches towards vehicle mass estimation would require a comprehensive knowledge of the individual components making up the system. However, due to significant vehicle-to-vehicle variation, it would be very difficult to characterize the individual components to the level of precision necessary to make meaningful vehicle mass estimations. A multi-system physics-based or analytical modeling approach towards vehicle mass estimation would be less dependent upon precise characterization of individual system components. By combining multiple mass estimations, each with well understood, larger than desirable uncertainties it would be possible to generate a vehicle mass estimation with an uncertainty significantly less than any of the individual mass estimations. This method requires that each of the individual mass estimations have uncertainty values that are well understood and driven by uncertainty in the parameters of the individual components of the system. This requirement makes analytical modeling an unappealing choice, due to the vast amount of testing and data that would be required to fully characterize the effect of the parameters of the individual components of each system using empirical data. Physics-based modeling allows the effects of the parameters of the individual components of each system to be characterized mathematically. The work described here provides a

method for estimating a vehicle's total mass from multiple ground-based remote measurements of multiple vehicle systems and subsystems responses. The remote measurements are combined into a single vehicle mass estimate using physics-based models of the vehicle systems and subsystems and uncertainty calculations.

The approach described here uses multiple ground-based remote sensors to acquire vehicle system data. This data, along with estimated vehicle system parameters, is fed into physics-based system models that calculate vehicle mass. This results in multiple vehicle mass estimations and uncertainty values. These vehicle mass estimations and uncertainty values are combined to generate a single vehicle mass estimation and uncertainty.

Several full-scale field experiments were carried out to identify which vehicle systems were best suited to use for mass estimation. A well-suited vehicle system needs to have a response to vehicle mass that can be measured with a ground-based remote sensor system. The experiments were designed to provide remotely sensed measurements of the system response as well as truth data whenever possible. The remotely sensed measurements were made from ground-based sensors and the truth data measurements were made with on-board vehicle sensors. The experiments were conducted at multiple locations using multiple vehicle types. The vehicles were tested under load conditions that ranged from unloaded up to (and in some cases exceeding) the maximum recommended load capability of the vehicle.

The research described here resulted in multiple original contributions to science. While processing the video data collected during the first suspension response field test, a unique video tracking algorithm was developed to accurately track the position of a specific region on a moving vehicle as it traveled through the field-of-view of a video camera. The algorithm is robust to changing illumination conditions and slowly occurring changes in scale and rotation. During the research into the suspension response phenomenology, an algorithm was designed that fit the measured suspension response data from the front and rear suspensions of a vehicle to a unique system of second-order differential equations which described the response of a coupled damped-harmonic oscillator to an input signal (speed bump). The output of this algorithm is the mass and center of gravity of the load being carried by the vehicle, given that the suspension parameters, unloaded mass, unloaded center of gravity, and unloaded moment of inertia of the vehicle

are known. During the research into the frame twist phenomenology a method was developed to calculate the speed of an engine using Fourier analysis of acoustic data measured on or near the vehicle. The method exploited the strong harmonics present in the engine's acoustic signal, which allows it to be robust to background noise, echoes, and other issues associated with acoustic measurements. Finally, the work done here has defined a set of optimal collection requirements that would greatly help anyone in the accurate measurement of the frame twist or suspension response phenomenologies.

Chapter 2

Theory

The work presented here spans several seemingly unrelated areas of study. Topics ranging from thermodynamics, to morphological processing, to damped harmonic oscillators will be discussed at length in later chapters. For completeness, this chapter will provide fundamental information for the possibly unfamiliar topics that will be presented in later chapters.

2.1 Physics

2.1.1 Linear Kinetic Energy

Linear kinetic energy, as defined in classical mechanics describes the energy of a non-rotating object as a function of its mass and its linear velocity and has units of Joules. Linear kinetic energy of an object is expressed in terms of the mass of the object, m , and the velocity of the object, v .

$$KE_{Linear} = \frac{1}{2}m \cdot v^2 \quad (2.1)$$

2.1.2 Rotational Kinetic Energy

A rotating object has kinetic energy; however quantifying the amount of this energy using the equation for linear kinetic energy becomes problematic. If the object is taken as a whole, it is

not moving and therefore would have no linear kinetic energy. To calculate the amount of kinetic energy of the rotating object, one must think of the object as a system of infinitely small objects, each with an infinitely small mass and a measurable velocity. From this, one could calculate the linear kinetic energy of each object, sum them, and calculate the kinetic energy of the rotating system as seen in Equation 2.2.

$$KE_{Rotational} = \frac{1}{2} \int v^2 dm \quad (2.2)$$

The velocity, v , of the mass, dm , is going to vary based on its distance from the axis of rotation

$$v = \omega \cdot r \quad (2.3)$$

where, ω , is the angular velocity of the rigid system (radians per second), and, r , is the distance of the mass, dm , from the axis of rotation. Plugging Equation 2.3 into Equation 2.2 we get Equation 2.4

$$KE_{Rotational} = \frac{\omega^2}{2} \int r^2 dm \quad (2.4)$$

where the integral describes the mass distribution properties of the rigid system for a given axis of rotation. This is known as the moment of inertia I and has units of kilogram-meters squared (kg m^2).

$$I = \int r^2 dm \quad (2.5)$$

Combining everything, we get the simplified equation for rotational kinetic energy as a function of moment of inertia and angular velocity.

$$KE_{Rotational} = \frac{1}{2} I \omega^2 \quad (2.6)$$

2.1.3 Damped harmonic oscillator

A harmonic oscillator is a term for a system that experiences a restoring force when displaced from an equilibrium position. An example of a harmonic oscillator is a mass-spring system. Hooke's

law states that the force required to compress, or stretch a spring by a distance, x , is proportional to that distance. From that relationship, we can describe the forces acting on a simple harmonic oscillator

$$F = -k_s \cdot x(t) \quad (2.7)$$

where F is the force acting on the mass, k_s is the spring constant, which characterizes the force imparted by the spring per meter of displacement from the equilibrium position and has units of Newtons per meter (N/m), and x the displacement from equilibrium of the mass. In this document we will always consider the spring constant k_s as a positive term and the force imparted by the spring constant as opposite the spring displacement (note the negative sign in Equation 2.7). As the mass is moved away from the equilibrium position, a restoring force is applied by the spring back towards the equilibrium position. When the mass is released, it oscillates around the equilibrium position, and will continue oscillating until an external force acts on it. From Newton's 2nd Law, we can rewrite the above equation and describe the forces acting on a simple harmonic oscillator at any instant in time

$$\frac{d^2x}{dt^2} = -\frac{k_s}{m}x(t) \quad (2.8)$$

where m is the mass of the object. The general solution to this differential equation is a combination of a sine and cosine of the same frequency

$$x(t) = A \cos \sqrt{\frac{k_s}{m}}t + B \sin \sqrt{\frac{k_s}{m}}t \quad (2.9)$$

where the square root of the sum of the square of A and B is equal to the initial displacement of the system. This means that this system will oscillate at a frequency of $\sqrt{\frac{k_s}{m}}$, which we will call ω_0 . This term, ω_0 , is called the undamped angular frequency of the oscillator and has units of inverse seconds.

Now we will consider a damped harmonic oscillator. Damping, as it pertains to harmonic oscillators, refers to any force that reduces the oscillations. In the case of the mass-spring system, the damping force would be caused mainly by air resistance, or drag. The forces acting on a damped harmonic oscillator can be described with the following equation

$$\frac{d^2 x}{dt^2} + \frac{k_d}{m} \frac{dx}{dt} + \frac{k_s}{m} x = 0 \quad (2.10)$$

where k_d is the viscous damping coefficient which has units of Newtons-seconds per meter (N s/m). The viscous damping coefficient is linearly dependent on the velocity of the mass. In this work we will always consider the damping coefficient as a positive term and the force derived from the damping coefficient as opposite the first derivative of the spring displacement. If the spring constant, viscous damping coefficient, and mass are known, Equation 2.10 can be solved for $x(t)$. To solve this differential equation, we will assume $x(t) = e^{rt}$ and plug it into Equation 2.10.

$$r^2 e^{rt} + r \frac{k_d}{m} e^{rt} + \frac{k_s}{m} e^{rt} = 0 \quad (2.11)$$

Equation 2.10 and be simplified to Equation 2.12.

$$r^2 + r \frac{k_d}{m} + \frac{k_s}{m} = 0 \quad (2.12)$$

Solving Equation 2.12 gives Equation 2.13.

$$r = \frac{-k_d \pm \sqrt{k_d^2 - 4 \cdot m \cdot k_s}}{2 \cdot m} \quad (2.13)$$

Plugging Equation 2.13 back into our equation for $x(t)$, we are left with Equation 2.14.

$$x(t) = e^{t \frac{-k_d \pm \sqrt{k_d^2 - 4 \cdot m \cdot k_s}}{2 \cdot m}} \quad (2.14)$$

From Equation 2.14 we see there are three possible solutions for the damped harmonic oscillator; the underdamped, critically damped, and overdamped solutions. For the underdamped solution, $k_d^2 - 4 \cdot m \cdot k_s$ is less than zero. For the critically damped solution, $k_d^2 - 4 \cdot m \cdot k_s$ is equal to zero, and for the overdamped solution $k_d^2 - 4 \cdot m \cdot k_s$ is greater than zero. We will define ζ as a unitless term referred to as the damping ratio.

$$\zeta = \frac{k_d}{2\sqrt{m \cdot k_s}} \quad (2.15)$$

This simplifies Equation 2.14 to Equation 2.16.

$$x(t) = e^{-t\omega_0(\zeta \mp (\zeta^2 - 1))} \quad (2.16)$$

This expression describes the position of a damped harmonic oscillator as a function of time (this assumes an initial displacement of one).

2.1.4 Heat Transfer

Heat transfer refers to the movement of enthalpy, or heat energy, within a system. There are three mechanisms by which enthalpy can move within a system; conduction, convection, and radiation.

Conduction

Heat transfer by conduction refers to the flow of enthalpy from a region of higher temperature to a region of lower temperature through adjacent molecule interaction. For example, when one end of a metal rod is heated, the enthalpy travels from the heated end towards the unheated end by interaction of adjacent molecules, not through migration of the molecules themselves through the rod.

Convection

Heat transfer by convection refers to the flow of enthalpy by the movement of molecules in a fluid. For example, in a pot of water being heated from the bottom, the enthalpy is passed from the heated pot to the water mainly by conduction, but the heated water molecules rise due to a decrease in density transferring the enthalpy upward where they come into direct contact with cooler water molecules and transfer enthalpy by conduction.

Radiation

Heat transfer by radiation refers to the flow of energy from matter in the form of electromagnetic radiation. All matter has some thermal energy, which causes their atomic particles to interact with each other. This interaction of charged atomic particles causes them to move from one discrete energy level to another. When a charged particle moves from a higher discrete energy level to a

lower discrete energy level, a photon is emitted with energy equal to the difference between the levels. Since the photon does not require a medium in which to travel (unlike conduction and convection), thermal radiation can occur in a vacuum.

Blackbody Radiation and Emissivity

All objects radiate electromagnetic (EM) energy; the amount and type of electromagnetic energy emitted by an object depends on the object's surface temperature and emissivity. Emissivity, $\epsilon(\lambda)$, is the ratio of the emitted spectral radiance, $L(T, \lambda)$, from an object at temperature T to the emitted spectral radiance of a perfect blackbody at the same temperature $L_{BB}(T, \lambda)$ [6]. Emissivity also describes how well an object absorbs EM energy, an emissivity of 1 means that an object absorbs all of the EM radiation incident upon it. Emitted spectral radiance is defined as the power per unit area radiated by a surface into a steradian of solid angle as a function of wavelength and has units of Watts per meters squared steradian micron ($\text{W}/(\text{m}^2 \text{sr } \mu\text{m})$). A perfect blackbody is an ideal material with an emitted spectral radiance that can be calculated using the following equation [6]

$$L_{BB}(T, \lambda) = 2hc^2\lambda^{-5} \left(e^{\frac{hc}{\lambda kT}} - 1 \right) \quad (2.17)$$

where h is Planck's constant (6.626×10^{-34} J s), c is the speed of light (2.998×10^8 m/s), λ is the wavelength of the EM energy being emitted (μm), k is Boltzmann's constant (1.38×10^{-23} J/K), and T is the temperature of the blackbody (K). An example of the calculated emitted spectral radiance for a blackbody at 800 K can be seen in Figure 2.1. Since the emitted spectral radiance for a blackbody can be calculated for any given temperature, it is possible to calculate temperature of a blackbody from a calibrated measurement of the spectral radiance leaving it. Real objects, however, do not behave as blackbodies. Real objects have emissivity values that are less than 1 (and greater than 0), this means that some of the EM energy incident upon it is not absorbed. If the EM energy incident upon an object is not absorbed, it must either be reflected or transmitted as the following equation states

$$\epsilon(\lambda) + \rho(\lambda) + \tau(\lambda) = 1 \quad (2.18)$$

where $\rho(\lambda)$ is the spectral reflectance of the object, and $\tau(\lambda)$ is the spectral transmission of the object. This means that for an opaque object ($\tau(\lambda) = 0$) the measured spectral radiance will be a combination of self-emitted radiance and reflected radiance. In the subsequent section, we will review a common method of separating the self-emitted radiance from the reflected radiance to calculate the surface temperature of the object.

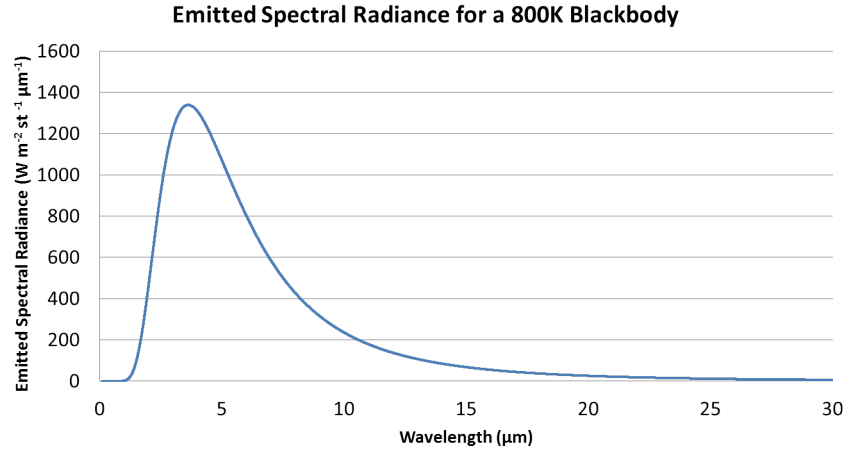


Figure 2.1: Calculated emitted spectral radiance for a blackbody at a temperature of 800K

2.1.5 Ideal Gas

An ideal gas exhibits a linear relationship between pressure, temperature, volume, and amount [7]. This means that at constant temperature, the volume of a fixed amount of an ideal gas is inversely proportional to pressure [7]. This is known as Boyle's law and is expressed as

$$V = \frac{1}{P} \quad (2.19)$$

where V is volume and P is pressure which has units of Pascals (Pa). For a fixed amount of ideal gas at constant pressure, the volume and temperature are proportional. This is known as Charles's law and can be expressed as

$$V \propto T \quad (2.20)$$

where T is temperature in units of Kelvin (K). For an ideal gas at constant pressure and temperature, the volume is proportional to the amount of the gas. This is known as Avogadro's law and can be expressed as

$$V \propto n \quad (2.21)$$

where n is the amount of the gas in units of moles (mol). Finally, all of the relationships can be combined into a single expression known as the ideal gas law which relates the four parameters; temperature, volume, pressure, and amount

$$R = \frac{PV}{nT} \quad (2.22)$$

where R is known as the universal gas constant, and has a value of $8.314 \text{ Pa m}^3/(\text{K mol})$.

Although no gas actually behaves as an ideal gas, simple gases such as O_2 , N_2 , H_2 , and the noble gases exhibit behavior that is very nearly ideal for normal temperatures and pressures [7].

2.2 Instruments and Sensors

2.2.1 Thermocouple

A thermocouple is a device used to measure temperature. It consists of two dissimilar metals joined at one end and electrically open at the other. When the temperature at the junction of the dissimilar metals is different from the temperature at the electrically open end, a voltage differential is induced. This voltage differential is caused by the Seebeck effect, which describes the voltage gradient along a wire resulting from a temperature gradient as seen in the following equation

$$\nabla V = S \cdot \nabla T \quad (2.23)$$

where ∇V is the voltage gradient at any point on the wire, S is the Seebeck coefficient with units of Volts per Kelvin (V/K), and ∇T is the temperature gradient at the same point with units of Kelvin (K). If we assume the electrically open end of the circuit is at a constant temperature T_1 and the temperature at the junction is at temperature T_2 , we can calculate the voltage at the electrically open junction using the following equation

$$\Delta V = \int_{T_1}^{T_2} S_A(T) - S_B(T) dT \quad (2.24)$$

where $S_A(T)$ is the Seebeck coefficient for metal A as a function of temperature, and $S_B(T)$ is the Seebeck coefficient for metal B as a function of temperature. The Seebeck coefficient for metals varies greatly from metal to metal, and for temperature ranges. In order to generate a sufficient voltage difference at the electrically open end of the circuit, specific metal pairs are chosen based on the range of temperatures to be measured.

2.2.2 Microbolometer Infrared Camera System

A microbolometer infrared camera system consists of an array of microbolometers and is typically used for imaging in the thermal infrared (8 to 15 μm). A microbolometer infrared camera system does not need to be actively cooled like most types of thermal infrared sensors, making them easier to use and less expensive. A microbolometer, as its name suggests is a small bolometer. Where a bolometer is a thermal infrared detector which uses a resistance thermometer to measure the temperature of a radiation absorber [8]. A typical bolometer consists of a radiation absorber, resistive thermometer, supporting substrate, reflector, and a read out circuit as shown in Figure 2.2. During operation, thermal infrared radiation is absorbed by the absorber causing its temperature and the temperature of the resistive thermometer (which is thermally coupled to the radiation absorber) to increase. The change in temperature of the resistive thermometer is measured by an electrical circuit [8]. For absorbers that are not optically thick, a reflector can be placed on the backside to reflect radiation that is not absorbed or reflected by the absorber. For the bolometer to respond quickly to changes in incident radiation, the components must have a low heat capacity, and the absorber and thermometer must be thermally coupled to a heat sink (so the resistive thermometer quickly returns to an ambient temperature when the thermal infrared radiation is not present).




Figure 2.2: Structure of a typical bolometer

2.2.3 LiDAR Sensing Theory

Light detection and ranging, or LiDAR, is a sensing technique for measuring the distance between the sensor and a target. LiDAR works by illuminating the target with a laser and precisely measuring the time it takes for the photons from the laser to reflect off of the target and return to the LiDAR system. A typical scanning LiDAR system consists of a laser that is scanned across the target using a moving mirror. The light leaves the laser in the form of a very narrow beam with small angular divergence, reflects off of the target, and returns to the LiDAR system. Since the speed of light in air is known (2.99×10^8 m/s) the distance between the target and the sensor can be calculated from the round-trip travel time of the beam. The process is repeated after the mirror has been rotated. The data collected by the LiDAR system can be used to create three-dimensional point clouds by combining the known three-dimensional location and orientation of the LiDAR sensor at the moment of each measurement, the measured distance from the LiDAR system to the target, and trigonometry.

2.3 Internal Combustion Engine

The internal combustion engine is the primary mode of power for nearly all over-the-road vehicles in use today. There are two main types of engines in widespread use today; two-stroke, and four-stroke (each type can be gasoline or diesel powered).

2.3.1 Two-Stroke

A two-stroke internal combustion engine, as its name suggests, takes two strokes to complete a single power stroke as seen in Figure 2.3. The intake of air and expulsion of exhaust take place during the compression and power strokes, respectively [1]. When the piston is at the bottom of its stroke, a blower forces fresh air into the cylinder from a port near the bottom of the cylinder. The fresh air being forced into the cylinder, forces the exhaust out through an open exhaust valve. When the piston rises above the inlet ports, the exhaust valve closes and the fresh air is compressed. Before the piston reaches the top of the cylinder, fuel is injected into the cylinder, which is quickly ignited (either by a spark plug in gasoline engines, or by heat of compression in diesel engines). The pressure from the combusting fuel pushes the piston down. When the piston is near the bottom of the cylinder, the exhaust valve opens and the inlet air port is exposed. The blower then begins to blow fresh air into the cylinder, pushing the exhaust out and the cycle begins again.




Figure 2.3: Diagram outlining the operation of a two-stroke internal combustion engine [1]

2.3.2 Four-Stroke

A four-stroke internal combustion engine, takes four strokes to complete a single power stroke as seen in Figure 2.4. During the intake stroke, the piston moves down, drawing a fuel-air mixture into the cylinder through the open intake valve. Next, during the compression stroke, the intake

valve is closed and the piston is drawn upward, compressing the air-fuel mixture. Next, ignition occurs, either through the firing of a spark plug in gasoline-powered engines, or through ignition caused by the heat of compression in diesel engine. The power stroke follows the ignition, as the combusting air-fuel mixture forces the piston downward. The final stroke, the exhaust stroke occurs when the exhaust valve opens and the piston is drawn upward expelling the exhaust from the cylinder.




Figure 2.4: Diagram outlining the operation of a typical four-stroke internal combustion engine [2]

2.4 Mathematics

2.4.1 Euclidean Distance

When dealing with n -band spectral measurements, it is often convenient to consider each measurement as a point in n -dimensional space. Points that close together are similar, while points that are far apart are dissimilar. Euclidean distance is a simple method of quantifying the distance between two points in n -dimensional space. In two-dimensions, the distance between two points is calculated using the following equation

$$d(\mathbf{a}, \mathbf{b}) = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2} \quad (2.25)$$

where (a_1, a_2) and (b_1, b_2) are points in two-dimensional space, and $d(\mathbf{a}, \mathbf{b})$ is the distance between them. For n-dimensions, the Euclidean distance is calculated by simply expanding the two dimensional case as seen in Equation 2.26.

$$d(\mathbf{a}, \mathbf{b}) = \sqrt{\sum_{i=1}^n (a_i - b_i)^2} \quad (2.26)$$

2.4.2 Mahalanobis Distance

The Mahalanobis distance is similar to the Euclidean distance, in that they both describe distance in n-dimensional space. The main difference between Euclidean distance and Mahalanobis distance, is that Mahalanobis distance quantifies the distance between a single point in n-dimensional space and a set of points, not a single point as in Euclidean distance. The Mahalanobis is calculated using the following equation

$$MD(\mathbf{a}) = \sqrt{(\mathbf{a} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{a} - \boldsymbol{\mu})} \quad (2.27)$$

where, $MD(\mathbf{a})$ is the Mahalanobis distance from point \mathbf{a} to a set of points with a mean of $\boldsymbol{\mu}$, and a covariance matrix of $\boldsymbol{\Sigma}$. By incorporating the covariance matrix, the Mahalanobis distance quantifies how well a point fits into a given set of points. Consider a set of points in two-dimensions, normally distributed about the origin. The data has a standard deviation of 25 along the x-axis and a standard deviation of 5 along the y-axis. Points located at (0,30) and (30,0) are the same Euclidean distance from the data set mean, but clearly (30,0) fits into the data set much better. The Mahalanobis distance bears out this difference; the point (0,30) is a Mahalanobis distance of 35.6 from the data set mean, while (30,0) is a Mahalanobis distance of 1.5 from the data set mean. This clearly shows that the point located at (30,0) fits into the data set much better than the point located at (0,30). This example can be seen in Figure 2.5.




Figure 2.5: A set of data, normally distributed about the origin with a standard deviation of 25 along the x-axis and 5 along the y-axis (blue). The point at (0,30) in green is a Euclidean distance of 30 and a Mahalanobis distance of 35.6 from the data set mean. The point at (30,0) in red is a Euclidean distance of 30 and a Mahalanobis distance of 1.5 from the data set mean.

2.4.3 Homogeneous Coordinate System

Homogeneous coordinates systems are used in projective geometry as a way of representing infinity using finite coordinates [9]. Homogeneous coordinates work by adding another dimension to the traditional Euclidean coordinate system. Using a two-dimensional example, a Euclidean coordinate system describes position using coordinate pairs. In a homogeneous coordinate system, the two-dimensional plane will be considered a plane in three-dimensions and points on the plane can be described as the intersection point between the plane $z = 1$ and infinite lines extending from the origin (0,0,0) in both directions. For example, the point (5,10) in Euclidean space could be represented by (5,10,1) in a homogeneous coordinate system. Since we are only interested in the point where the line intersects with the plane located at $z = 1$. Homogeneous coordinates are invariant to scaling, meaning that (5,10,1) and (20,40,4) represent the same line in the three-dimensional

homogeneous coordinate system and the same point on the plane located at $z = 1$. Clearly, every Homogeneous coordinate with a z -component not equal to zero will map to a point on the infinite plane located at $z = 1$. However, when the z -component is equal to zero the line does not intersect with the plane at $z = 0$. This is considered a point at infinity. Points at infinity can be used in image rectification. Consider an image of a parallel line, due to projective distortion the lines will not appear parallel (if extended they would intersect at some point). If the location of this intersection (called the vanishing point) is calculated or measured, one can rectify the image by mapping this intersection to a point at infinity $(x,y,0)$. This allows us to mathematically deal with the concept of infinity using only finite values.

2.4.4 Hessian Matrix

The Hessian matrix is a matrix of second-order partial derivatives of a function [9]. Consider a function $z = f(x, y)$, which describes a surface in three-dimensions. The Hessian matrix would be calculated using Equation 2.28 [9].

$$H(f(x, y)) = \begin{bmatrix} \frac{\partial^2 f}{\partial x^2} & \frac{\partial^2 f}{\partial x \partial y} \\ \frac{\partial^2 f}{\partial y \partial x} & \frac{\partial^2 f}{\partial y^2} \end{bmatrix} \quad (2.28)$$

From the above equation, we can see that Hessian matrix is square, and comprised of the second order partial derivatives of a function. The size of the Hessian matrix is determined by the dimensionality of the function. An n -dimensional function will have an n -by- n Hessian matrix. The Hessian matrix is often used to find local minima and maxima in multi-dimensional optimization or minimization problems.

2.4.5 Haar Transformation Matrix

The Haar transformation matrix is used in multiresolution analysis and is made up of a set of Haar wavelets. The Haar transformation matrix is applied by multiplying an intensity image by the Haar transformation matrix and the inverse of the Haar transformation matrix as seen below [10]

$$\mathbf{T} = \mathbf{H}\mathbf{F}\mathbf{H}^T \quad (2.29)$$

where \mathbf{T} is the resulting transformed image (size N-by-N), \mathbf{H} is the Haar transformation matrix (size N-by-N), and \mathbf{F} is the input image (size N-by-N). The Haar transformation matrix is comprised of N basis functions $h_k(z)$ for $k = 0, 1, 2, \dots, N$ and $z = \frac{0}{N}, \frac{1}{N}, \frac{2}{N} \dots, \frac{N-1}{N}$ [10]. The basis functions are shown in the following expressions [10].

$$h_0(z) = \frac{1}{\sqrt{N}} \quad (2.30)$$

$$h_k(z) = \frac{1}{N} \begin{cases} 2^{\frac{p}{2}}, & \frac{(q-1)}{2^p} \leq z < \frac{(q-0.5)}{2^p} \\ 2^{\frac{p}{2}}, & \frac{(q-0.5)}{2^p} \leq z < \frac{q}{2^p} \\ 0, & otherwise \end{cases}$$

$$k = 2^p + q - 1 \quad (2.31)$$

$$0 \leq p \leq n - 1 \quad (2.32)$$

$$1 \leq q \leq 2^p, p \neq 0 \quad (2.33)$$

$$N = 2^n \quad (2.34)$$

For example, for $N = 4$ the above equations force the values of k , p , and q to be as shown below.

Table 2.1: Values for k , p , and q when $N = 4$

k	p	q
0	0	0
1	0	1
2	1	1
3	1	2

Plugging these values into the above equations, we get the following Haar transformation matrix.

$$\mathbf{H} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ \sqrt{2} & -\sqrt{2} & 0 & 0 \\ 0 & 0 & \sqrt{2} & -\sqrt{2} \end{bmatrix} \quad (2.35)$$

The above process can be completed for matrices of any size to generate the appropriate Haar transformation matrix. Once the Haar transformation matrix has been populated for a given image size the Haar transformation matrix can be used to generate a multiresolution image. The multiresolution image can be used as an aid in image compression, spatial frequency analysis, and feature detection.

2.4.6 RANdom SAmple Consensus


The RANdom SAmple Consensus (RANSAC) algorithm is designed to fit a model to a dataset that contains outliers [11]. The RANSAC algorithm is carried out by first randomly selecting the minimum number of correspondence pairs required to generate the model and calculate the model from this subset [11]. Next, the calculated model is applied to the entire data set, and all of the point correspondence pairs that fit this model to within a threshold are identified and compiled into a subset [11]. If the number of points in the subset exceeds a threshold, a new model is calculated from the subset data and the algorithm terminates [11]. If the number of points in the subset does not exceed the threshold, the process starts over with new randomly selected correspondence pairs [11]. If after a certain number of iterations occurs a model has not been selected, the model with the largest subset is chosen [11]. For example, to fit a line to a set of points using RANSAC, you

would start by randomly selecting two points and calculating the line that passes through them. Next you would find all the other points in the set that lie on or near that line. If the number of points is sufficient (satisfies a user defined threshold) then a least squares regression would be applied to find the line that best fit this new (outlier free dataset). If the number of points did not exceed the threshold a new random pair of points would be selected and the process would repeat. This would continue until a solution is found.

2.5 Imaging Science

2.5.1 Morphological Processing

Morphological processing is typically done on binary images to extract spatial components such as region boundaries and the convex hull [10]. When operating on binary images, we can describe an image, or objects in an image, using only the coordinates of the non-zero pixels (or the zero pixels). For example, the binary images shown in Figure 2.6 can be expressed using only the coordinates of the non-zero (white) region and the size of the image grids. The first image contains set A and the second image contains set B .



(a) Binary image of set A

(b) Binary image of set B

Figure 2.6: Example of two binary images

Once the data is in the form of sets, any number of set operations can be executed. The union operator, \cup , combines two sets, forming a new set that contains all of the pixels present in each

individual set. The intersection operator, \cap , combines two sets, forming a new set that contains only the pixels the individual sets have in common. An example of the union and intersection operator can be seen in Figure 2.7. Other set operations, such as reflections and translations can be executed on single sets of data, and as the names suggest, reflect the set about a line or move the set.




Figure 2.7: Example of the union and intersection operator

The morphological erosion operation is a process that incorporates a structuring element. A structuring element is a set of pixels, defined by the user, which interacts with the set being operated on. An erosion of set B with structuring element C is defined by the following equation [10]


$$B \ominus C = \{z \mid (C)_z \cap B^c = \emptyset\} \quad (2.36)$$

where B^c , is the complement of B and \emptyset is the empty set. This equation means that $B \ominus C$ is comprised of the set of all points such that C translated by z is contained by B [10]. This can be visualized as the structuring element sliding around the set B and turning on any pixel in which the structuring element is completely within set B . An example of the erosion operation can be seen in Figure 2.8.

The morphological dilation operation is a very similar to erosion. A dilation of set B with structuring element C is defined by the following equation [10]

$$B \oplus C = \left\{ z \mid \left(\hat{C} \right)_z \cap B = \emptyset \right\} \quad (2.37)$$

where \hat{C} is the reflection of C about its origin. This equation means that $B \oplus C$ is comprised of the set of all points such that \hat{C} translated by z is overlapped by B [10]. This can be visualized as the structuring element sliding around the set B and turning on any pixel covered by the structuring element (while still centered on the original set B). An example of the dilation operation can be seen in Figure 2.8.



(a) Set B eroded by a square structuring element of size 10-by-10 pixels
(b) Set B dilated by a square structuring element of size 10-by-10 pixels

Figure 2.8: Example of the union and intersection operators

The morphological closing operation is a combination of erosion and dilation. The closing operation can be expressed by Equation 2.38.

$$A \bullet C = (A \oplus C) \ominus C \quad (2.38)$$

This equation means that $A \bullet C$ is calculated by first dilating set A with C and eroding the result by C . An example of the closing operation can be seen in Figure 2.9.

The morphological opening operation is a combination of erosion and dilation. The opening operation can be expressed by Equation 2.39.

$$A \bullet C = (A \ominus C) \oplus C \quad (2.39)$$

This equation means that $A \bullet C$ is calculated by first eroding set A with C and dilating the result by C . An example of the opening operation can be seen in Figure 2.9.




Figure 2.9: Example of the closing and opening operators

2.5.2 Cross-Correlation

Two dimensional cross-correlation is a common method of performing template matching in imagery. Consider two intensity images, one large image, and one small image as seen in Figure 2.10. For this example, the small image is actually a subset of the larger image selected at random. The goal of template matching is to locate where in the large image the smaller image, or template, best matches. Visually, a cross-correlation can be thought of as sliding the template over the larger image, multiplying them together (pixel-by-pixel) summing the result for every location on the larger image. Mathematically, the cross-correlation is executed using Equation 2.40 [10]

$$c(x, y) = \sum_s \sum_t I(s, t) \cdot T(x + s, y + t) \quad (2.40)$$




Figure 2.10: An image and a template image

The cross-correlation of the example image and template can be seen in Figure 2.11.

Figure 2.11: Cross-correlation of image I and template T

The result does not clearly show the point where the template best matches the image. This is because cross-correlation is heavily influenced by the local intensities of both the large image, I , and the template, T . To account for the effect of the local intensities a normalized correlation coefficient can be calculated using the following equation [10]

$$\gamma(x, y) = \frac{\sum_s \sum_t (T(s, t) - \bar{T}) \sum_s \sum_t (I(x + s, y + t) - \bar{I}(x + s, y + t))}{\sqrt{\sum_s \sum_t (T(s, t) - \bar{T})^2 \sum_s \sum_t (I(x + s, y + t) - \bar{I}(x + s, y + t))^2}} \quad (2.41)$$

where \bar{T} is the mean of the template and $\bar{I}(x + s, y + t)$ is the mean value of the large image coincident with the template [10]. The resulting normalized correlation coefficient clearly identifies the location where the template best matches the large image as shown in Figure 2.12.



Figure 2.12: Normalized correlation coefficient for image I and the template T

2.5.3 One-Dimensional Special Functions

In order to solve certain problems it can be helpful to use one-dimensional functions that can only be described piecewise and cannot be expressed as a single algebraic expression. These one-dimensional special functions are defined with x as the dependent variable, and y as the independent variable. We define x_0 as the position of the function on the x-axis and b as a scaling factor [12]. A zero value for x_0 means that the function is centered at zero on the x-axis, a positive value for x_0 means the center of the function is shifted by $|x_0|$ to the right, and a negative value for x_0 means the center of the function is shifted by $|x_0|$ to the left [12]. A sign change for b will result in a reflection about the line $x = x_0$ [12].

2.5.4 Rectangular Function

The rectangular function is a one-dimensional special function described by the following equation [12].

$$RECT\left(\frac{x-x_0}{b}\right) = \begin{cases} 1, & \left|\frac{x-x_0}{b}\right| > \frac{1}{2} \\ \frac{1}{2}, & \left|\frac{x-x_0}{b}\right| = \frac{1}{2} \\ 0, & \left|\frac{x-x_0}{b}\right| < \frac{1}{2} \end{cases}$$

The above equation describes the value of the RECT function under three conditions. Looking carefully at the expressions, it becomes clear that the RECT is a rectangular function of width b centered at x_0 , as seen in Figure 2.13.




Figure 2.13: The rectangle function $RECT\left(\frac{x-x_0}{b}\right)$

2.5.5 Comb Function

The comb function is a one-dimensional special function described by Equation 2.42 [12].

$$COMB\left(\frac{x-x_0}{b}\right) = |b| \sum_{n=-\infty}^{\infty} \delta(x-x_0-n \cdot b) \quad (2.42)$$

The above equation describes the value of the COMB function as a function of x , x_0 , and b . This describes an infinite number of Dirac delta functions of area $|b|$, spaced by $|b|$, and shifted by x_0 .

A Dirac delta function, $\delta(x)$, is an impulse function with an area of one, that is zero everywhere except $x = 0$. Since the Dirac delta function has a finite area (equal to one) but no span (it is only non-zero at $x = 0$) it can be thought of as having an infinite magnitude at $x = 0$. The Dirac delta function is useful in approximating quantities that are highly localized in a given coordinate system [12]. An example of a COMB function can be seen in Figure 2.14.




Figure 2.14: The comb function $COMB\left(\frac{x-x_0}{b}\right)$ where the vertical axis denotes the area of the Dirac delta function

Chapter 3

Phenomenologies

As stated in Chapter 1, the objective of the work presented here is to develop a multi-system, physics-based approach towards remote vehicle mass estimation. This chapter will describe, in detail, all of the vehicle phenomenologies we considered as being potentially correlated to vehicle mass and observable with a remote sensing system. Acceleration and deceleration rate, brake temperature, engine acoustics, engine heat generation, exhaust plume, frame twist, suspension response, tire deformation, tire temperature, and vibration response are the ten phenomenologies whose relationship to vehicle mass will be explored. Where possible, vehicle mass will be directly related to a remotely observable characteristic of the phenomenology. This remotely observable characteristic of the phenomenology will then be evaluated and potential measurement techniques will be listed and detailed.

3.1 Acceleration and Deceleration Rate

The acceleration and deceleration rate of a vehicle is affected by the total mass of the vehicle. The acceleration rate of a fully loaded vehicle will be less than that of a vehicle carrying a lesser load for the same transmission gear and engine output. Similarly, the magnitude of the deceleration rate of a fully loaded vehicle will be less than that of a vehicle carrying a partial load, given the same braking system and application method. Since deceleration is simply negative acceleration,

we will henceforth refer to only acceleration, which can be either positive or negative. The relationship between vehicle load and acceleration rates can be described with the equation for kinetic energy from classical mechanics

$$E_k = \frac{1}{2} m \cdot \Delta v^2 \quad (3.1)$$

where E_k is the kinetic energy required to increase or decrease the velocity of an object with mass m by an amount equal to Δv . The change in velocity of the vehicle can be calculated from the acceleration rate of the vehicle using the following equation

$$\Delta v = \int_{t_i}^{t_f} a(t) dt \quad (3.2)$$

where t_i and t_f are the initial time and final time, respectively, and $a(t)$ is the acceleration rate of the vehicle as a function of time t . Rearranging Equation 3.1 and substituting using Equation 3.2 we can express mass, m in terms of acceleration.

$$m = 2 \cdot \frac{E_k}{\left(\int_{t_i}^{t_f} a(t) dt \right)^2} \quad (3.3)$$

This simple relationship between acceleration, vehicle mass, and kinetic energy is crucial for any conservation of energy based approaches to vehicle mass estimation. This means that accurate measurement of a vehicle's acceleration rate over a period of time coupled with accurate measurement of the amount of energy applied to the vehicle over the same period of time will result in an accurate estimation of the vehicle's mass.

3.1.1 Measurement Techniques

The acceleration and velocity of a vehicle as a function of time can be calculated from position measurements using the following equations.

$$a(t) = \frac{d^2}{dt^2} x(t) \quad (3.4)$$

$$v(t) = \frac{d}{dt}x(t) \quad (3.5)$$

By utilizing a video camera system that collects imagery as a function of time, the position versus time measurement can be made by extracting accurate position measurements in three-dimensions from imagery data (assuming the video camera system collects imagery at a known rate). In the subsequent sections, we will review some of the methods of extracting position information from imagery.

Position Measurements with an Uncalibrated Camera

One of the simplest methods of measuring the position of a vehicle remotely is to acquire a sequence of images from a downward tilted video camera system positioned on a bridge or overpass that overlooks a region of planar, straight roadway with at least one known ground distance along the road axis [3]. In order to extract meaningful spatial data from the imagery in this scenario, first the projective distortion must be removed. This can be done for any plane in an image by applying the proper transform (assuming minimal lens distortion). This transform can be calculated by finding two of the plane's vanishing points. A vanishing point is the intersection point in the image plane of lines that are parallel in the world coordinate system. The lane demarcations can be exploited to calculate a vanishing point in the road plane. If we assume that the lane demarcations are parallel we can simply extend two lane demarcation lines until they intersect, this intersection is the vanishing point. The coordinates on the image plane of this intersection is the vanishing point for the road plane for lines running parallel to the direction of traffic. An example of this can be seen in Figure 3.1.



Figure 3.1: Imagery taken using an uncalibrated camera system, highlighted in the image are the lines used to determine the vanishing point and the location of the vanishing point [3]

A second vanishing point is assumed to be at infinity since the camera's horizontal axis is oriented so that any lines on the road surface that are perpendicular to the direction of traffic would be imaged parallel to the x-axis [3]. The affine transform H required to remove the projective distortion for the road plane can then be calculated from the vanishing point using the following equation [3]

$$H = \begin{bmatrix} 1 & \frac{-f_1}{f_2} & 0 \\ 0 & 1 & 0 \\ 0 & \frac{-1}{f_2} & 1 \end{bmatrix} \quad (3.6)$$

where f_1 and f_2 are the x-coordinate and y-coordinate, respectively, of the road plane vanishing point for lines parallel to the direction of traffic. The transform, H , will map each coordinate pair on the image plane to the road plane. Now that the projective distortion is removed, the next step is to use the known ground distance along the road axis to calculate the proper scale factor to convert from arbitrary units to units of length along the direction of the road plane (the cross-road

scale factor will be incorrect). Using this calculated affine transform and scale factor, the change in position of a vehicle as a function of time can be calculated simply by tracking the vehicle from frame to frame and transforming its image coordinates to world coordinates (assuming the frame rate of the camera is known). Grammatikopoulos et al. produced results with an estimated accuracy of approximately ± 3 km/h [3]. This method imposes some rigid constraints as to where the video camera system is located and how it is oriented relative to the road surface. We can relax some of these restrictions and potentially improve upon the accuracy if we know the camera matrix of the video camera system.

Camera Matrix

When we obtain an image with a typical camera, we are creating a two-dimensional representation of a three-dimensional scene. The camera matrix is a 3-by-4 matrix, which maps three-dimensional homogeneous world coordinates to two-dimensional coordinates on a camera's image plane. The camera equation has eleven degrees of freedom; five of which are internal camera parameters and 6 which are external camera parameters (the twelfth term is a uniform scale factor). The internal camera parameters are; the scale factor in the x-coordinate direction (α_x), the scale factor in the y-coordinate direction (α_y), the skew (s), and the coordinates of the principal point (x_0, y_0) [9]. The scale factors refer to the product of the number of pixels per unit length along an axis and the focal length of the camera [9]. The skew parameter is only non-zero when the pixel elements of the camera are skewed causing the x-axis and the y-axis to be not perpendicular (this is very uncommon [9]). The external camera parameters have six degrees of freedom, which describe the position and orientation of the imaging system in the three-dimensional world coordinate system [9]. Equation 3.7 shows that by combining the internal and external camera parameters we can calculate the 3-by-4 camera matrix, P

$$P = K \cdot R \cdot C \quad (3.7)$$

where K is a 3-by-3 matrix that contains all of the internal camera parameters, R is the product of three 3-by-3 matrices, each of which describes the camera's rotation in one of the three axes of the three-dimensional world coordinate frame, and C contains the camera center location in the

world coordinate frame

$$K = \begin{bmatrix} fm_x & s & x_0 \\ 0 & fm_y & y_0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3.8)$$

$$R = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_x & -\sin \theta_x \\ 0 & \sin \theta_x & \cos \theta_x \end{bmatrix} \begin{bmatrix} \cos \theta_y & 0 & -\sin \theta_y \\ 0 & 1 & 0 \\ \sin \theta_y & 0 & \cos \theta_y \end{bmatrix} \begin{bmatrix} \cos \theta_z & -\sin \theta_z & 0 \\ \sin \theta_z & \cos \theta_z & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3.9)$$

$$C = \begin{bmatrix} 1 & 0 & 0 & -x_C \\ 0 & 1 & 0 & -y_C \\ 0 & 0 & 1 & -z_C \end{bmatrix} \quad (3.10)$$

where f is the focal length of the camera, m_x and m_y are the number of pixels per unit length in image coordinates in the x and y directions, respectively, s is the skew parameter, (x_0, y_0) are the coordinates of the principal point in terms of pixel dimensions, θ_x , θ_y , and θ_z are the angles of rotation about the x, y, and z axes, respectively, and (x_C, y_C, z_C) are the coordinates of the camera center in the world coordinate system [9]. These parameters are combined to create the camera matrix which is used to map the three-dimensional world onto the two-dimensional image plane of the imaging system. This dimension reduction from the three-dimensional world coordinate system to the two-dimensional coordinate system of the camera's image plane results in an unrecoverable loss of information about the scene. This means that while there is a single set of two-dimensional coordinates that can result from an image of a three-dimensional scene, there are an infinite number of three-dimensional scenes that could have produced any single image. Specifically, if the camera matrix is known, a point in three-dimensions can be mapped to a single point on the camera's image plane and a point on the camera's image plane can be mapped to a ray in three-dimensions (a ray has an infinite number of points).

Calculating the Camera Matrix

The preferred method of calculating the camera matrix for a given imaging system involves a full camera calibration. This calibration can be done using the RIT - Digital Imaging and Remote

Sensing (DIRS) calibration facility. This calibration facility consists of a frame that houses reflective targets positioned on three planes as seen in Figure 3.2. To calibrate an imaging system, images are collected of the three-plane frame from nine locations, each with a different orientation. The position of each reflector is identified in each of the images. The three-dimensional location of each reflector is known for the calibration frame, and the camera's location and orientation relative to the camera frame is precisely measured. From this, the internal camera parameters are calculated. Once the internal camera parameters are known for a given imaging system, the full camera matrix can be calculated for any camera position and orientation.




Figure 3.2: RIT DIRS laboratory camera calibration facility

Position Measurements with a Calibrated Camera

When working with a fully calibrated imaging system with a known camera matrix, we can relax some of the imposed restrictions required for extracting spatial information from an uncalibrated system. Specifically, when measuring a target with known dimensions (e.g. license plate) it is possible to calculate the position of the target in the three-dimensional world coordinate system

regardless of its orientation relative to the imaging system (as long as the entire target is visible and well resolved). This can be done by identifying the coordinates on the edges of the target in the image coordinate system, mapping these coordinates to rays in the world coordinate system, and solving for the position and orientation for which the target's true dimensions match. This means that the three-dimensional position coordinates of a target with known dimensions can be measured as a function of time using a calibrated video camera system positioned in an off-normal orientation relative to the target. Garibotto et al. developed a method for measuring the velocity of a vehicle by exploiting the known spacing between characters on a license plate and were able to measure the speed of a vehicle with accuracy better than 4 % [13]. However, for a measurement of this type to be an accurate method of measuring vehicle position, the imaging system must be positioned with an unobstructed view of the target and the dimensions of target must be known. For an unobstructed view of the license plate, the video camera system would need to be positioned so as to image the front or the rear of the vehicle, requiring the video camera system to be in the path of the vehicle, or above it. A more desirable position for the video camera system from an operational standpoint would be on the side of the road. With a camera looking across a road, perpendicular to the flow of traffic, the motion of the vehicle would be captured largely as translational motion in consecutive frames of a video as opposed to a head-on camera view, in which vehicle motion is captured as changes in object depth. As a result, a cross-road method could utilize more of the camera's field-of-view than a head-on method, which should result in better accuracy. In order to move to the cross-traffic mounted camera system, we need to relax the restriction of knowing our target's dimensions. To do this we can move from a single camera system to a two-camera system.

Position Measurements with Two Calibrated Cameras

Measurements of the position of a vehicle as a function of time with a system of two calibrated video camera systems can be done with no a priori knowledge of the vehicle. By positioning the calibrated camera systems a precisely known distance apart, each with a precisely known orientation, both observing the same general area, three-dimensional spatial data can be obtained about the scene. Consider a single calibrated video camera system observing a vehicle from the side of

the road. Given the image coordinate location of any single pixel object, we can use the camera equation to calculate the three-dimensional ray on which the object must be located in the three-dimensional world coordinate system. If a second calibrated camera system was to observe the same single pixel object at exactly the same time, we could calculate a second three-dimensional ray on which the object also must be located in the three-dimensional world coordinate system. Since two rays can only intersect in one location we only need to calculate the intersection point of the two rays to determine the coordinates of the object in the three-dimensional world coordinate system. This simplified description only works under ideal conditions absent of noise, since real-world measurements are not perfect and always contain some level of noise, a more complicated technique is required. To obtain the best solution for a given measurement configuration one must find the object coordinates in the world coordinate system that best fits the measured data [9]. This requires defining a cost or error function and then minimizing the function, typically with a direct linear transformation or analogous method [9]. Zhu and Zheng developed a method of measuring the speed of a vehicle using a two-camera video system and tested it using a camera orientation that was a hybrid of the cross-traffic and head-on camera orientations. They achieved accuracy better than 1 % [14].

Video Detection Algorithms

Each of the methods of measuring the position of a vehicle as a function of time discussed above require some type of detection and tracking algorithm to operate autonomously. The single camera measurement techniques require the detection and tracking of a predefined target, for example a license plate, throughout multiple video frames. The two camera method requires unique points, or features, on the vehicle to be identified in two simultaneously collected images and then tracked throughout multiple video frames.

For any autonomous method of measuring vehicle position as a function of time, the system must be able to detect the presence of a vehicle in a scene. Consider a video camera system observing a stretch of road over a period of time, most of the collected images would be of an empty road with a relatively static background. When a vehicle enters the scene, the system must detect its presence and begin tracking the vehicle, or some portion of the vehicle through

subsequent frames. Zielke et al. developed a method for vehicle detection that is based on the general symmetry of most vehicles [15]. Their method considers each row and column from an image to be one-dimensional intensity functions. They split each of these one-dimensional functions into their constituent even and odd functions and calculate a degree of symmetry value as a function of the location of the axis of symmetry and the width of the observation window [15]. They then determine the most significant axis of symmetry and observation interval by imposing restrictions on the observation interval and maximizing the degree of symmetry [15]. This method showed good promise, however, it is designed for an application viewing vehicles from the front or the rear, and vehicles have much less symmetry when viewed from the side making this method less desirable for a side-view camera orientation.

Another method, this one employed by Grammatikopoulos et al., analyzes multiple video frames and calculates the mean and standard deviation of the intensity values at each pixel location and considers any intensity value within one standard deviation of the range to be background [3]. This is continuously done to account for changes in illumination and temporal changes in the background. When a pixel falls outside the range designated as background it is considered as originating from an object of interest [3]. The resulting binary image indicating objects of interest and background, is filtered using morphological operators to eliminate objects that are too small and to fill in holes [3]. The result is a binary image, representing the entire area of the object of interest. This method will work well regardless of the camera orientation, as long as the approximate size of the vehicle in pixels is known.

Current Video Tracking Algorithms

Once an object of interest has been detected, it must be tracked throughout subsequent, and possibly antecedent, video frames. If the target is known, it can be tracked using a cross-correlation-based search method. This is a very common method of locating known targets in a scene and is used by Grammatikopoulos and as Bertozzi in their tracking algorithms [3, 16]. However, this cross-correlation method is not robust to changes in scale or rotation of the target. As a result, when searching for a known target of unknown scale and orientation, one must generate multiple targets, each scaled and oriented differently which exponentially increases the computational

cost of such a search. To mitigate this, most implementations of this method impose restrictions on how much the target can change. Restrictions on changes in translation, orientation and scale for successive frames are imposed once the target has been initially detected. This significantly decreases the computational cost of the algorithm and reduces false alarms.

Another method of tracking an object throughout a sequence of video frames is called feature tracking. Features in this context, are defined as parts of an object that can be detected easily, such as corners or lines [17]. The key aspect of feature tracking is automatically identifying the same feature (e.g. corner, edge) in a pair of images. A common algorithm for identifying the same feature in a pair of images is the Speeded-Up Robust Features (SURF) algorithm. The SURF algorithm is divided into three main parts; the first part is the detection of points of interest in the image (corners, edges, blobs), next the area directly surrounding each point of interest is characterized and represented by a feature vector, and last the feature vectors are matched between different images [18]. The detection of points of interest needs to be invariant to changing camera-target geometry and must be computationally cost-effective. For the SURF algorithm, a Fast-Hessian Detector is used to identify points of interest. The Fast-Hessian Detector calculates a discretized version of the determinant of the Hessian matrix, where the Hessian matrix $\mathcal{H}(\mathbf{x}, \sigma)$ is defined as [18]

$$\mathcal{H}(\mathbf{x}, \sigma) = \begin{bmatrix} L_{xx}(\mathbf{x}, \sigma) & L_{xy}(\mathbf{x}, \sigma) \\ L_{xy}(\mathbf{x}, \sigma) & L_{yy}(\mathbf{x}, \sigma) \end{bmatrix} \quad (3.11)$$

where $L_{xx}(\mathbf{x}, \sigma)$, $L_{yy}(\mathbf{x}, \sigma)$, $L_{xy}(\mathbf{x}, \sigma)$, and $L_{yx}(\mathbf{x}, \sigma)$ are the convolutions of the Gaussian second order partial derivatives, $\frac{\partial^2}{\partial x^2}g(\sigma)$, $\frac{\partial^2}{\partial y^2}g(\sigma)$, $\frac{\partial}{\partial x}\frac{\partial}{\partial y}g(\sigma)$, and $\frac{\partial}{\partial y}\frac{\partial}{\partial x}g(\sigma)$ individually applied to the image and σ is the standard deviation of the Gaussian [18]. The Gaussian second order derivatives are discretized and the determinant of the Hessian is calculated using these discretized kernels for multiple values of σ at each pixel location in the image [18]. The interest points are located at the extrema of the Laplacian of this three-dimensional matrix [18]. The characterization of the area directly surrounding each interest point is determined by first calculating the Haar-wavelet response in the x and y direction for a circular neighborhood immediately surrounding each interest point and using these wavelet responses to generate an orientation vector [18]. Next, a square region centered on the interest point, and orientated along the orientation vector is consid-

ered, and the Haar-wavelet response is calculated along the orientation vector and perpendicular to the orientation vector for several subregions within the square regions [18]. The resulting wavelet responses are used to create a 4-dimensional descriptor vector associated with each subregion. The final output for each interest point is a 4-dimensional orientation vector, and a 4-dimensional descriptor vector for each subregion (typically 16 subregions) which are concatenated into a single vector [18]. After locating the interest points and characterizing their surrounding area, the final step is to match interest points for image pairs. This matching can be done based on the Euclidean or Mahalanobis distance of these descriptor vector [18].

3.2 Brake Temperature

A vehicle's braking system is designed to convert the vehicle's kinetic energy to a different form. A vehicle being operated with a high load will have more kinetic energy than the same vehicle carrying a standard load traveling at the same velocity. Nearly all of the braking systems utilized by over-the-road vehicles today are friction brakes. Friction brakes convert a vehicle's kinetic energy into thermal energy by pressing together two surfaces moving at different velocities (typically one surface is moving and one surface is stationary). If a friction style braking system were employed, more thermal energy would be created for the same change in velocity by a vehicle being operated with a high load than would be generated by a vehicle carrying a lesser load. Assuming the vehicle's kinetic energy is converted directly to heat by the braking system, the thermal energy generated can be described by the relationship [2]

$$E_b = \frac{m}{2} (V_1^2 - V_2^2) + \frac{I}{2} (\omega_1^2 - \omega_2^2) \quad (3.12)$$

where E_b denotes the thermal energy generated by the brakes, V_1 the velocity of the vehicle prior to braking, V_2 the velocity of the vehicle after braking, I the moment of inertia of the vehicle's drive-train, ω_1 the angular velocity of the drive-train prior to braking, and ω_2 the angular velocity of the drive-train after braking. Solving the equation for mass we get the following equation.

$$m = \frac{2 \cdot E_b - I (\omega_1^2 - \omega_2^2)}{V_1^2 - V_2^2} \quad (3.13)$$

If we restrict this to only cases where the vehicle comes to a complete stop and assume that the drive-train is disengaged (application of the clutch in a manual vehicle) the equation simplifies to:

$$m = \frac{2 \cdot E_b}{V_1^2} \quad (3.14)$$

This simplified equation has three unknowns, the mass of the vehicle (m), the thermal energy generated by the braking system (E_b), and the initial velocity of the vehicle (V_1). The initial velocity of the vehicle can be measured accurately using a number of remote sensing methods that were discussed in Section 3.1.1. The thermal energy generated by the braking system over a period of time can be calculated from the following equation [7]

$$E_b = ((T_f - T_i) \cdot c_{Brakes} \cdot m_{Brakes}) + E_{Cooling} \quad (3.15)$$

where T_i and T_f are the initial temperature of the braking system and the final temperature of the braking system, respectively (Kelvin), c_{brakes} is the specific heat capacity of the material making up the braking system (Joules per Kelvin kilograms), m_{brakes} is the mass of the braking system (kilograms), and $E_{Cooling}$ is the thermal energy lost to cooling over the measurement period (Joules). This relationship suggests that the total amount of kinetic energy converted to thermal energy by the braking system can be calculated from measurements of the braking system temperature during a braking period and knowledge of the braking system mass, braking system material composition, and the amount of thermal energy lost to cooling over the measurement period. Combining Equations 3.14 and 3.15 we get the following equation for calculating the mass of a vehicle.

$$m = \frac{2}{V_1^2} (((T_f - T_i) \cdot c_{Brakes} \cdot m_{Brakes}) + E_{Cooling}) \quad (3.16)$$

From this equation, the mass of a vehicle can be calculated from measurements of the vehicle's velocity and brake temperature immediately prior to a period of braking, and measurements of the vehicle's brake temperature at the completion of the braking period (when the vehicle has reached a complete stop). In addition, knowledge of the mass and material composition of the braking system must be known, as well as the amount of thermal energy lost to cooling during the braking

period. These parameters may be able to be estimated based on environmental conditions, vehicle type and braking system type.

The two main style of friction brakes used today are disc-style and drum-style. Disc brakes are most common in passenger vehicles and consist of a disc that rotates with the wheel and a stationary caliper that uses hydraulics to pinch the rotating disc between two brake pads. A diagram illustrating the main parts of a typical disc brake system can be seen in Figure 3.3. Drum brakes are more common in larger vehicles and consist of a drum that rotates with the wheel and two brake shoes that are pressed into the inside of the rotating drum. A diagram illustrating the main parts of a typical drum brake system can be seen in Figure 3.4.




Figure 3.3: Diagram of a typical disc brake system




Figure 3.4: Diagram of a typical drum brake system

If we assume that the thermal energy lost to cooling, mass of the braking system, and the specific heat capacity of the braking system can be estimated based on vehicle type, style of braking system, vehicle speed and environmental conditions, then our ability to calculate the mass of a vehicle using Equation 3.16 is dependent on our ability to accurately measure the initial velocity of the vehicle (discussed in Section 3.1.1) and our ability to accurately measure the temperature of the braking system.

3.2.1 Measurement Techniques

The temperature of a vehicle's braking system can be measured remotely using a thermal infrared sensor. A thermal infrared sensor observing a target quantifies the total amount of thermal infrared radiance incident on the sensor from the location of the target. This radiance is a combination of target self-emission, radiance reflected off of the target, radiance scattered into the field-of-view of the sensor, and self-emission of the atmosphere between the target and the sensor. Additionally, this radiance passes through the atmosphere, which absorbs certain wavelengths of radiance. In the subsequent section we will review a method of extracting temperature information from measured thermal infrared radiance.

Temperature Measurements with a Calibrated Spectral Sensor

The total radiance leaving an opaque object $L(T, \lambda)$ is a combination of self-emitted radiance and reflected radiance, which can be calculated using the following equation (atmospheric contributions and scattering effects can be considered insignificant for small target-to-sensor distances)

$$L(T, \lambda) = \epsilon(\lambda) \cdot L_{BB}(T, \lambda) + (1 - \epsilon(\lambda)) \cdot L_{BG}(\lambda) \quad (3.17)$$

where the $L_{BB}(T, \lambda)$ is the spectral radiance emitted from a blackbody at temperature T , and $L_{BG}(\lambda)$ is the radiance of the background incident onto the surface of the object.

Kahle and Alley developed a method called Reference Channel, or REFCHAN, to calculate temperature and spectral emissivity when both the object leaving spectral radiance and the background spectral radiance are known [19]. In this method, it is assumed that the object being measured has an emissivity of one at some known wavelength, λ_0 , in the measurement window. Based on this assumption, the target could be treated as a perfect blackbody at wavelength λ_0 . The radiance emitted by a perfect blackbody is described by Equation 3.18, and by solving this equation for surface temperature we can calculate the temperature of a perfect blackbody given its measured self-emission at wavelength λ_0 as seen in Equation 3.19 [20].

$$L_{BB}(T, \lambda) = 2hc^2\lambda^{-5} \left(e^{\frac{hc}{\lambda kT}} - 1 \right) \quad (3.18)$$

$$T = \frac{hc}{k \cdot \lambda_0 \cdot \ln \left(\frac{2 \cdot c^2 \cdot h}{L(\lambda_0) \cdot \lambda_0^5} + 1 \right)} \quad (3.19)$$

Once the surface temperature of the object has been calculated, the emissivity of the object can be calculated by rearranging Equation 3.17 and solving for spectral emissivity as shown in Equation 3.20

$$\epsilon(\lambda) = \frac{L(T, \lambda) - L_{BG}(\lambda)}{L_{BB}(T, \lambda) - L_{BG}(\lambda)} \quad (3.20)$$

Kahle and Alley's method is fairly simple and easy to implement, but is only possible if the material being measured has a emissivity of one at a known wavelength. For many objects being

measured there will not be a wavelength for which it has an emissivity of one (or even approaching one), specifically for man-made materials [20]. This method is only one of several temperature emissivity separation techniques, it is important that the technique employed is properly suited for the material of interest and test conditions.

3.3 Engine Acoustics

The vast majority of vehicles today are powered by 4-stroke internal combustion engines fueled by gasoline or diesel. An internal combustion engine converts the chemical energy stored in a fuel into rotational energy used to supply power to the drive wheels of the vehicle. The operation of a typical 4-stroke gasoline fueled internal combustion engine is described in Figure 3.5. Diesel fuel powered 4-stroke internal combustion engines operate in a similar manner, however spark plugs are not used, instead the heat generated while compressing the diesel fuel initiates the ignition of the fuel.




Figure 3.5: Operation of a 4-stroke gasoline powered internal combustion engine [2]

The acoustic signal generated by a 4-stroke internal combustion engine is caused mainly by the combustion of fuel and the interaction of mechanical parts. The combustion noise is caused

by a rapid increase in pressure in the combustion chamber, which incites structural vibration of the engine, and resonance in the gas [21]. There are many mechanical parts within the engine that contribute to the acoustic signal of the operating engine to varying degrees. Fuel injectors are a reciprocating part that contributes to the acoustic signal of the engine through vibrations caused by the injector needle impacting the backstop. Intake and exhaust valves contribute to the acoustic signal of the engine through vibrations imparted to the engine through impacts during the sealing and unsealing of the combustion chambers. When an engine is operating properly, each of these contributions to the engine's acoustic signal occur in a distinct pattern. For a single cylinder, the pattern would occur as seen in Figure 3.5. Typically, vehicles utilize internal combustion engines with between 4 and 8 cylinders. Each piston is linked to a single crankshaft that along with the cam shafts, dictates exactly the order in which each cylinder proceeds through the 4-stroke cycle. This results in a repeating acoustic pattern with a period related to the angular speed of the engine by a factor of two as seen in the following expression

$$\omega_{Engine} = \frac{2}{P} \quad (3.21)$$

where ω_{Engine} is the angular frequency at which the engine is operating (Hertz) and P is the period of the repeating acoustic pattern. It should be noted that the factor of two in this equation is necessary because a 4-stroke engine requires two revolutions to fire all of the cylinders.

Once the angular frequency of the engine has been calculated, the number of cylinders can be determined by identifying how many combustions occur during each cycle. It may also be possible to estimate the size, or total displacement of the engine and the cylinder pressure through analysis of its remotely observed acoustic signal. To calculate the mechanical energy generated by the cylinders of an engine over a period of time, the following equation is used [22]

$$E_{Cylinders} = 2 \cdot \Delta t \cdot MEP \cdot D \cdot \omega_{Engine} \quad (3.22)$$

where $E_{Cylinders}$ is the mechanical energy generated by the cylinders of an engine, Δt is the period of time over which the measurement was made, MEP is the mean effective pressure, or the average cylinder pressure over an engine cycle (Pascals), D is the total displacement of the engine (cubic meters), and ω_{Engine} is the angular frequency at which the engine is operating.

If the change in velocity of the vehicle is measured over the same time period, the mass of the vehicle can be calculated by combining Equation 3.1 and Equation 3.22 and rearranging as shown in Equation 3.23.

$$m = \frac{4 \cdot \Delta t \cdot MEP \cdot D \cdot \omega_{Engine}}{\Delta v^2} \quad (3.23)$$

This suggests that engine speed, size, and cylinder pressure, when combined with the change in vehicle velocity could provide accurate vehicle mass estimations for a positively accelerating vehicle. Admittedly, measurement of engine displacement and cylinder pressure acoustically is a difficult proposition and may not be possible due to vehicle variation, the effect of mufflers, background acoustic noise, echoes, measurement geometry, and environmental conditions. However, even if engine speed is the only measurement that can be reliably made acoustically it may still be possible to calculate mass by estimating the other parameters based on vehicle type. In the subsequent section we will discuss the type of measurements that are needed to calculate the mass of a vehicle using this phenomenology.

3.3.1 Measurement Techniques

Measuring the engine speed, size, and cylinder pressure of an internal combustion engine from the acoustic noise it emits is a challenging proposition. Zheng and Leung et al. developed a signal processing procedure to decompose the acoustic noise signal of an internal combustion engine [21]. Their method showed success in identifying the individual contributions of several noise sources, however they were not able to quantify the parameters of interest to our phenomenology. It is unlikely that any measurement other than engine RPM will be possible using acoustic monitoring (and RPM may only be possible if the number of cylinders and engine type is known).

3.4 Engine Heat Generation

Internal combustion engines are designed to convert the chemical energy stored in a fuel into rotational energy that can be used to propel a vehicle. However only 14 to 26 % of the chemical energy ends up being transformed to energy that can be used to move the vehicle, the remaining

energy is converted to thermal energy during the combustion of the fuel and as a result of internal friction. The total amount of thermal energy generated by the vehicle's engine is affected by the amount of load the vehicle is carrying. The engine will need to operate at a higher speed (RPM) to achieve the same acceleration rates when the vehicle is carrying a full load compared to when the vehicle is carrying a partial load. As an engine is operated at a higher frequency the engine's combustion rate increases, and the engine generates more thermal energy. A portion of this thermal energy is absorbed by the engine itself, however the engine's cooling system actively regulates the temperature of the engine. The bulk of the thermal energy generated by the engine takes the form of heated exhaust gases, which flow from the engine through the exhaust system. The exhaust system, which is not actively cooled in most cases, absorbs a fraction of the thermal energy as the exhaust gas flows through it. The change in temperature of the exhaust system over a given time period can be expressed with the following equation

$$\Delta T_{Exhaust} = \int_{t_i}^{t_f} \frac{(P_{Total}(t) \cdot (1 - \eta) \cdot R) - P_{Cooling}(t)}{c_{Exhaust} \cdot m_{Exhaust}} dt \quad (3.24)$$

where $\Delta T_{Exhaust}$ is the overall change in temperature of the exhaust system over the time period from t_i to t_f , $P_{Total}(t)$ is the total power output of the engine as a function of time (Watts), η is the efficiency of the engine, or the ratio of useful energy produced by the engine to total energy produced by the engine (unitless), $R_{Exhaust}$ is the percent of the thermal energy produced by the engine that is absorbed by the exhaust system (unitless), $P_{Cooling}(t)$ is the thermal energy lost by the exhaust system to cooling, $c_{Exhaust}$ is the specific heat capacity of the exhaust system (Joules per Kelvin kilograms), and $m_{Exhaust}$ is the mass of the exhaust system. Combining and rearranging equations 3.24 and 3.1 we can solve for the mass of the vehicle as shown in Equation 3.25.

$$m = \frac{2\eta \left(\Delta T_{Exhaust} \cdot c_{Exhaust} \cdot m_{Exhaust} + \left(\int_{t_i}^{t_f} P_{Cooling}(t) dt \right) \right)}{\left(\int_{t_i}^{t_f} a(t) dt \right)^2 (1 - \eta) \cdot R_{Exhaust}} \quad (3.25)$$

This expression relates the overall mass of the vehicle to changes in the temperature of the exhaust system and the velocity of the vehicle, which can be measured. Measuring the temperature of the exhaust system remotely would require a MWIR or LWIR sensor with an unobstructed view


of the exhaust system to measure the surface leaving radiance. Additionally, the emissivity of the exhaust system and the radiance incident upon the exhaust system would need to be known in order to calculate the surface temperature of the exhaust system from the measured radiance. The remaining parameters in the equation would need to be estimated based on vehicle type and environmental conditions. It is unclear exactly how much vehicle-to-vehicle variation there is in exhaust system mass, specific heat capacity, and emissivity but it is likely significant enough to make the accurate measurement of mass very difficult using this method.

3.4.1 Measurement Techniques

Remotely measuring the temperature of a vehicle's exhaust system is possible due to self-emission. Section 3.2.1 discusses the mechanism by which this measurement is possible and a method of extracting the temperature data from the measured radiance.

3.5 Exhaust Plume

During the operation of internal combustion engines, air is brought into the engine, combined with fuel and ignited. A simplified balanced equation for the complete combustion of a hydrocarbon fuel and air mixture can be described with the following expression [4].



The above relationship describes the simplified combustion of a hydrocarbon fuel (CH_2), which is a good approximation of average gasoline and diesel fuel, and air ($O_2 + 4N_2$) mixture [4]. When a mixture of hydrocarbon and air undergoes complete combustion, all of the carbon present in the hydrocarbon is converted to carbon dioxide (CO_2) and all of the hydrogen in the hydrocarbon is converted to water (H_2O) leaving the nitrogen (N_2) unchanged. Using this equation it can be determined that completely combusted hydrocarbon fuel and air mixtures will have an exhaust composition that is 12.5 % carbon dioxide, 12.5 % water vapor and 75 % nitrogen gas. It also can be determined that one molecule of CO_2 is output for every one molecule of CH_2 that is input. Given that CH_2 has a molar mass of 14.02 g/mol, CO_2 has a molar mass of 44.01 g/mol, water

has a molar mass of 18.01 g/mol, and gasoline has a density of approximately 0.74 kg/L it can be determined that 2.3 kg of CO₂ and 0.95 kg of water is output for every liter (0.74 kg) of gasoline combusted. This means that the rate at which a vehicle is consuming fuel can be calculated from the rate at which the vehicle is emitting CO₂ or water vapor. Typical internal combustion engines utilized by vehicles today only convert between 14 % and 26 % of the energy stored in hydrocarbon fuels into the vehicle's kinetic energy (more accurate efficiency estimations can be made if more information about the vehicle is known, e.g. make, model, age, mileage, condition). If we assume the exhaust behaves as an ideal gas we can express the amount of fuel consumed over a period of time as a function of the volume, pressure and exhaust gas temperature as a function of time as shown in Equation 3.27

$$m_{Fuel} = \frac{M \cdot X_{Fuel}}{R} \int_{t_i}^{t_f} \frac{P_{Exhaust}(t) \cdot V_{Exhaust}(t)}{T_{Exhaust}(t)} dt \quad (3.27)$$

where m_{Fuel} is the mass of the fuel consumed by the engine over the period of time t_i to t_f , $P_{Exhaust}(t)$ is the pressure of the exhaust gas as a function of time, $V_{Exhaust}(t)$ is the volume of the exhaust gas as a function of time, R is the Ideal Gas Constant (8.31 J/(K mol)), $T_{Exhaust}(t)$ is the temperature of the exhaust gas as a function of time, M is the molar mass of the fuel (grams per mole) and X_{Fuel} is the mole fraction of fuel to exhaust (for the gas described in Equation 3.26 the mole fraction of fuel to exhaust would be 1/8 meaning that eight moles of exhaust are produced for every mole of fuel consumed). Accounting for the engine's efficiency (η), combining Equations 3.27 and 3.1 and rearranging we can solve for the mass of the vehicle as shown in Equation 3.28

$$m = \frac{2 \cdot \eta \cdot M \cdot X_{Fuel}}{R \int_{t_i}^{t_f} a(t) dt} \int_{t_i}^{t_f} \frac{P_{Exhaust}(t) \cdot V_{Exhaust}(t)}{T_{Exhaust}(t)} dt \quad (3.28)$$

This suggests that measurements of a vehicle's exhaust volume, pressure, and temperature coupled with measurements of the vehicle's acceleration over a period of time could be used to calculate the vehicle's mass. In the subsequent section we will discuss some of the challenges associated with this measurement and some of the methods that have been developed to measure exhaust gases.

3.5.1 Measurement Techniques

Remote measurement of the volume, pressure, and temperature of the exhaust gas leaving a moving vehicle is challenging mainly due to the composition of the gas. Exhaust gas is composed of CO_2 , H_2O , and N_2 , all of which are present in the Earth's atmosphere in significant amounts. This prohibits exploiting passive detection using wavebands where these molecules absorb energy since very little energy in these wavebands make it to the Earth's surface. This leaves active illumination as the most likely path for successfully characterizing the amount of exhaust gas a vehicle is emitting.

Burgard and Bishop et al. developed a method of using active illumination to measure the concentrations of exhaust gases emitted from moving vehicles [4]. Their method uses a collimated infrared (IR) source, oriented so it travels across the road to a set of detectors as seen in Figure 3.6. The IR source and detector are positioned at the height of a typical exhaust pipe of a passenger vehicle so when a vehicle passes between them the exhaust plume is directly in the path of the collimated IR beam.

The detectors are filtered to only respond to energy at wavelengths that are specific to the absorption features associated with exhaust gases of interest [4]. Gas cylinders of known concentrations are used to calibrate the system. From this calibration, the output voltage of the detector can be directly converted to an estimate of gas optical depth [4]. The goal of this device is to measure the ratio of carbon monoxide (CO) and hydrocarbons (HC) to CO_2 . CO and HC are pollutants associated with combustion whose emissions are heavily regulated in some areas. However, this method could be used to simply measure the amount of CO_2 being emitted by the vehicle as it passes by. If the exhaust plume is assumed to be symmetric, the CO_2 optical depth measured when the exhaust plume is between the source and detector could be used to calculate the amount of exhaust (molar amount not volume) in that region of the plume. This is possible from a measurement of the optical depth of the CO_2 gas if the exhaust composition can be approximated using Equation 3.26. Combine this with the speed of the vehicle and it may be possible to estimate the amount of exhaust a vehicle is emitting as a function of time.




Figure 3.6: A schematic diagram of the test setup used by Burgard and Bishop et al. [4].

3.6 Frame Twist

A vehicle is accelerated forward by its drive wheels. The rate of acceleration of the vehicle is dependent on the torque being applied to the drive wheels. Assuming no losses, the acceleration rate of the vehicle (x''), can be defined in terms of the total torque on the drive wheels (τ_w), the radius of the drive wheels (r_w), and the mass of the vehicle (m), as shown in Equation 3.29. Rearranging this expression we can solve for the mass of the vehicle in terms of the other parameters as shown in Equation 3.30. This means that the total mass of the vehicle can be calculated from its acceleration, the radius of the drive wheels, and the total torque being applied to the drive wheels.

$$x''(t) = \frac{\tau_w(t)}{r_w \cdot m} \quad (3.29)$$

$$m = \frac{\tau_w(t)}{r_w \cdot x''(t)} \quad (3.30)$$

The torque being applied to the drive wheels can be calculated as a function of the torque being generated by the engine and the gear ratio of the transmission. A typical internal combustion engine needs to operate at angular velocities much higher than the angular velocities appropriate for the drive wheels (especially during slow travel). The vehicle's transmission converts the high angular velocity rotational energy from the engine into the slower, more appropriate, angular velocities required by the drive wheels (while increasing the torque). Assuming no transmission losses, a transmission can reduce the angular velocity, by increasing the torque and vice versa. The amount by which the transmission scales the angular velocity and torque is defined by the gear ratio (G). For a transmission with a gear ratio of G , the output torque (τ_{out}), and the output angular velocity (ω_{out}) can be described as a function of the input torque (τ_{in}) and the input angular velocity (ω_{in}) as seen in Equations 3.31 and 3.32. This relationship, coupled with Equation 3.30 shows that the total mass of the vehicle can be calculated from the acceleration of the vehicle, the radius of the drive wheels, the gear ratio, and the total torque being generated by the engine.

$$\tau_{out} = \tau_{in} \cdot G \quad (3.31)$$

$$\omega_{out} = \frac{\omega_{in}}{G} \quad (3.32)$$

To measure the torque being generated by the engine, we need to consider the interaction between the engine and the vehicle frame. The vehicle's engine is semi-rigidly coupled to the frame of the vehicle, and as a result when the vehicle's engine is generating torque, an equal torque must be supplied by the vehicle's suspension to prevent the engine and frame from rotating. This torque causes the frame to twist relative to the vehicle's axles. The magnitude of this frame twist should be equal to the torque being generated by the engine multiplied by a scale factor. This scale factor should be dependent on the suspension parameters of the vehicle. As a result; if the suspension parameters, drive wheel radius, and gear ratio are known, the total mass of the vehicle can be calculated from the frame twist and vehicle position as a function of time.

3.6.1 Measurement Technique

To measure the mass of a vehicle by exploiting the frame twist phenomenology, the acceleration of the vehicle as a function of time will need to be precisely measured in a front-on orientation (the vehicle traveling towards the measurement system). To measure the gear ratio of the vehicle we will need to measure the engine RPMs and compare this to the angular velocity of the vehicle's drive wheels. To measure RPM we can exploit the engine acoustics using the methods described in Section 3.3. Additionally, the angular twist of the vehicle frame will need to be precisely measured as a function of time. Both of these measurements are possible using the techniques discussed in Section 3.1.1. However, the measurement of the angular twist of the vehicle frame will require slight modification to the feature tracking algorithm.

Angular Rotation from Feature Tracking

From Section 3.1.1, we are familiar with the SURF method of tracking features throughout video frames. To calculate the amount of rotation the features undergo from frame to frame, we can use the points selected by the SURF method in conjunction with the outlier rejection technique RANSAC. We can use correspondence pairs identified using the SURF algorithm in consecutive video frames as the input to the RANSAC algorithm that calculates the model, or transform, that best fits the data. We can then decompose the transform to determine the amount of frame rotation the vehicle is experiencing in each image.

3.7 Suspension Response

The suspension system for most vehicles acts as a damped spring coupler between the frame of the vehicle, and each of the vehicle's axles. Each axle can be approximated as a damped harmonic oscillator, which can be described with the following balance of forces equation.

$$\frac{d^2y}{dt^2} - k_d \frac{dy}{dt} - k_s (y(t) - y_{eq}) + g = 0 \quad (3.33)$$

The above balance of force equation describes the relationship between the position of a mass as a function of time $y(t)$, and the forces acting on it in a damped harmonic oscillator system in

one-dimension. Where k_d represents the damping coefficient of the system with units of Newton seconds per meter (N s/m), k_s represents the spring constant of the system with units of Newtons per meter (N/m), y_{eq} is the equilibrium position of the mass, and g is the acceleration due to gravity (-9.8 m/s^2). This balanced force equation allows the position, velocity, and acceleration of the mass to be calculated as a function of time for a damped harmonic oscillator system with known parameters. Inversely, if the position of the mass was known as a function of time, as well as the spring constant and damping coefficient, the mass can be calculated.

This means that if the suspension parameters of the vehicle are known (k_s and k_d), the mass supported by the axle can be calculated from the position of the mass as a function of time. This is a somewhat oversimplified approach that ignores the forces caused by the frame's rotation. This rotational force can be accounted for with a balance of torques equation which in conjunction with the balance of forces equation (Equation 3.33) can fully describe a two-axle vehicle. This approach will be described in detail in Chapter 5.

3.7.1 Measurement Techniques

To measure the mass of a vehicle by exploiting the suspension response phenomenology, at least two positions on the frame of the vehicle, and the vehicle's wheels will need to be tracked. The positions of the features will need to be precisely measured in a side-on orientation (the sensor located on the road-side, oriented perpendicular to the path of the vehicle) as the vehicle travels over a speed-bump. These measurements require a level of precision that may not be possible using the techniques discussed in Section 3.1.1. As a result a custom video tracking algorithm was developed.

Custom Video Tracking Algorithm

The existing video tracking algorithms discussed in Section 3.1.1 were found to be inappropriate for some of the position tracking requirements of the suspension response phenomenology. A custom video tracking algorithm was developed that had improved accuracy and precision for this application at the expense of speed and autonomy. The custom tracking algorithm was designed to track a localized region on a vehicle, instead of tracking the entire vehicle like the methods

described in Section 3.1.1. This difference is critical for the suspension response phenomenology where the movement of the vehicle's frame relative to the vehicle's axle is the observable of interest. The video tracking algorithm developed for the suspension response application is optimized to accurately track points on a vehicle traveling through the field-of-view of a stationary camera. For this application, it is assumed that the frame-to-frame movement of the vehicle is small enough to consider the difference in projective distortion between consecutive video frames to be negligible. A block diagram of the processing steps used in this tracking algorithm can be seen in Figure 3.7.

The algorithm must be provided a sequence of image frames, and the location, in the first video frame, of the region to track. The region to track is extracted from the first image, and a contrast stretch is applied to remove low intensity edges. Next, the image is converted into a binary edge-image using a basic Laplacian of Gaussian (LoG) edge detector. The LoG edge detector works by first convolving the input image with a Gaussian kernel to remove high-frequency noise, and then convolving the resulting image with a Laplacian of the Gaussian kernel as shown in Equations 3.34, 3.35, and 3.36

$$G(x, y) = \frac{1}{\sqrt{2\pi} \cdot \sigma^2} \cdot \exp\left(-\frac{x^2 + y^2}{2 \cdot \sigma^2}\right) \quad (3.34)$$

$$\Delta G(x, y) = \frac{\partial^2 G}{\partial x^2} + \frac{\partial^2 G}{\partial y^2} \quad (3.35)$$

$$I_1(x, y) = (I_0(x, y) * G(x, y)) * \Delta G(x, y) \quad (3.36)$$

where $G(x, y)$ is a symmetric Gaussian kernel, σ is the standard deviation of the Gaussian kernel, $\Delta G(x, y)$ is the Laplacian of the Gaussian kernel, $I_0(x, y)$ is the input image, and $I_1(x, y)$ is the Laplacian of Gaussian filtered image. The filtered image ($I_1(x, y)$) is converted to an edge image by finding all of the zero-crossings and setting them equal to one (all other points are zero). An example of this can be seen in Figure 3.8. Edge-based images were found to be much more reliable for tracking than intensity images, as they produced much sharper responses from the normalized cross-correlation, resulting in reduced false alarms. This binary, edge image becomes the initial template. The next image frame is loaded and a search area is defined as the previous search

location plus twenty pixels in all four directions. An array of binary edge images is generated using a range of contrast stretches to account for any lighting change. An example of this array is shown in Figure 3.9. The normalized cross-correlation of the template image and each of the images in the array are calculated using Equation 3.37 [23]

$$\gamma(u, v) = \frac{\sum_{x,y} [I(x, y) - \overline{I_{u,v}}] [t(x - u, y - v) - \bar{t}]}{\left(\sum_{x,y} [I(x, y) - \overline{I_{u,v}}]^2 \sum_{x,y} [t(x - u, y - v) - \bar{t}]^2 \right)^{0.5}} \quad (3.37)$$

where $I(x, y)$ is the input image, $t(x, y)$ is the template image, $\overline{I_{u,v}}$ is the mean of the input image in the region under the template, and \bar{t} is the mean of the template image. The normalized cross-correlation of the array of images results in an array of intensity images. An example of one of the normalized cross-correlation intensity images is shown in Figure 3.10. The magnitude of these cross-correlation surfaces is going to be highest at the location where the two images align, and the highest value would represent the best match. The highest peak is selected and recorded as the center of the tracked region in the current frame. If no acceptable match is found, the search area is increased by twenty pixels in each of the four directions, and the process is repeated. This continues until an acceptable match is found or the search area becomes too large (the search area was deemed too large when it grew beyond 60 pixels in each direction). The maximum size of an acceptable search will depend on the expected speed of the vehicle and the camera parameters. Once a match is found, it will be used as the new template for the next frame. This is done because the shape of the tracked region will change due to projective distortion as the vehicle travels through the field-of-view of the camera. By continuously updating the template, it was determined that one does not need to account for this projective distortion (assuming the frame-to-frame movement of the vehicle is small). However, as a result of continuously updating the template, we need to track and account for the sub-pixel misalignment that accumulates each time a new template is chosen. If we do not account for this misalignment, the region being tracked can drift away from the true location defined in the first frame. This potential misalignment is tracked by computing the normalized cross-correlation of the new template, and the previous template and then fitting a sub-pixel, two-dimensional Gaussian to the peak. The deviation of the center of the calculated Gaussian from the exact center of the result is the amount of misalignment (it

should be less than one in each axis). This misalignment is cumulatively summed from frame-to-frame for each axis, and is compensated for each time the sum reaches an integer pixel deviation by shifting the current template. The process is repeated for each subsequent frame using the template selected from the previous frame and continues until no acceptable match is found or the edge of the image is reached.

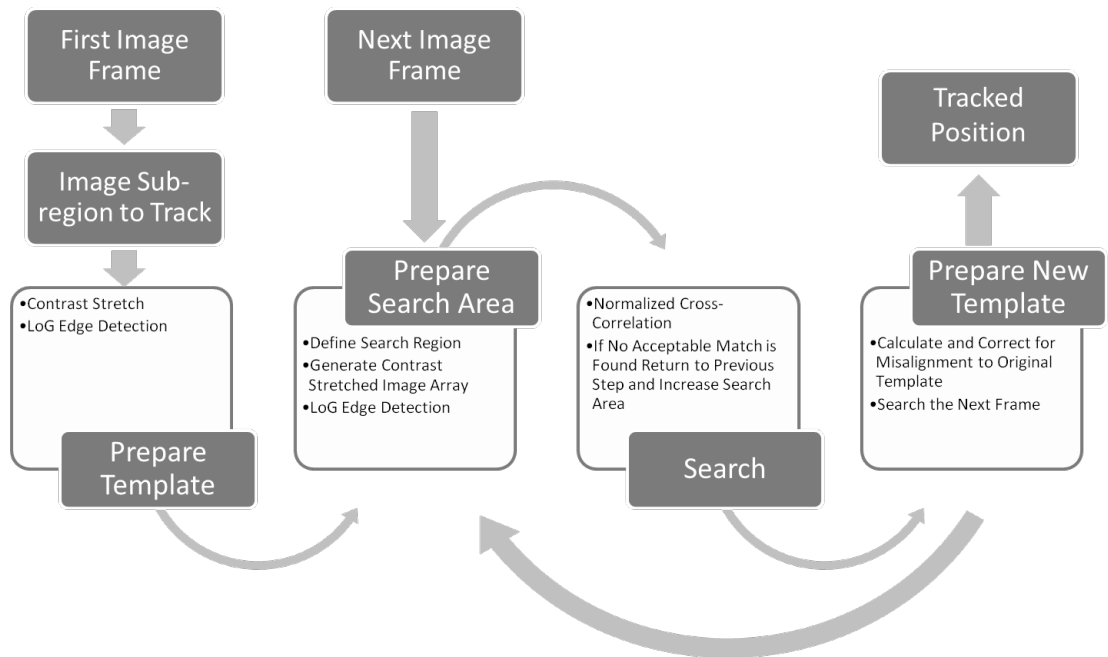


Figure 3.7: A block diagram of the data processing steps used to track position from video using the custom video tracking algorithm



Figure 3.8: The region to track (left) and a edge image of the region to track (right)



Figure 3.9: An array of binary, edge images generated by varying the contrast stretch (mild contrast stretch on the left to severe contrast stretch on the right)




Figure 3.10: An example of the intensity image created from the normalized cross-correlation of a search area image and the template image (both binary, edge images)

3.8 Tire Deformation

The entire mass of a vehicle and its contents is supported by its tires and nearly every over-the-road vehicle in use today is equipped with pneumatic tires. Pneumatic tires consist of an airtight body, called the carcass, filled with pressurized gas, typically air or nitrogen (N_2). Newton's first law of motion states that a stationary object must have a net force acting on it equal to zero. This means that for a stationary object, all forces must be perfectly balanced. For a stationary vehicle, gravity is exerting a force that is equal to the product of the vehicle's mass and the acceleration due to gravity (9.8 m/s^2 towards the center of the Earth). Since the vehicle is stationary, there must be an equal force opposing the force of gravity. This means that the tires are exerting a total force on the ground that is equal and opposite to the force gravity is exerting on the vehicle. Application of Newton's first law to the tire to ground interface shows that the tire's internal pressure and the force the tire is exerting on the ground must be perfectly balanced. If we ignore the effect of the stiffness of the tire's carcass, the total force exerted by the tires on the ground can be described with the following expression.

$$F_{Tires} = \sum_{i=1}^N P_i \cdot A_i \quad (3.38)$$

The above expression describes the total force being exerted on the ground by all of the vehicle's tires (F_{Tires}). Where N is the number of tires on the vehicle, P_i is the internal pressure of a specific tire, and A_i is the area in contact with the ground of a specific tire. Since F_{Tire} must be equal and opposite to the force of gravity acting on the vehicle, we can solve the equation for the mass of the vehicle as shown in Equation 3.39

$$m = \frac{\sum_{i=1}^N P_i \cdot A_i}{-g} \quad (3.39)$$

This relationship shows that the mass of a vehicle can be calculated by measuring the contact area of each of its tires, given that the internal pressure of each tire is known. Although, remote measurement of internal tire pressure would likely not be possible, assuming proper tire inflation may produce a good estimate of vehicle mass. In the subsequent section we will discuss the type of measurements that are needed to calculate the mass of a vehicle using this phenomenology.

3.8.1 Measurement Techniques

To measure the mass of a vehicle by exploiting the tire deformation phenomenology, the size of the contact area of each tire must be accurately measured. The contact area can be calculated from measurements of the contact area width and length. These measurements are possible using the techniques discussed in Section 3.1.1. However, due to the high level of precision required for this measurement, other methods should be investigated.

Dimension Measurements with Light Detection And Ranging (LiDAR)

A stationary LiDAR system positioned on the side of the road may be a better option for measuring the tire contact area than the methods described in Section 3.1.1. The high level of precision required, coupled with the dark tire material and shadows in the measurement area may prove too daunting a task for image-based measurements alone. A typical one-dimensional scanning LiDAR system emits a beam of light in the direction of the target, the light interacts with the target, and

a portion of the beam returns to the LiDAR system. A sensor in the LiDAR system measures the time it takes for the beam to travel to the target and back to the LiDAR system, and calculates the sensor to target distance based on this travel time measurement. A rotating mirror is used to change the direction of the beam. A one-dimensional LiDAR system could be deployed on the side of the road and image passing vehicles using the configuration seen in Figure 3.11.




Figure 3.11: Proposed configuration for one-dimensional scanning LiDAR system to measure tire contact area

3.9 Tire Temperature

As described in Section 3.8, all of the mass of the vehicle is supported by the tires. At the interface of the tire and the road surface the material of the tire is pressed into contact with the material of the road by the gravitational force acting on the vehicle. This applied stress to the tire material causes it to deform, or strain. The stress-strain relationship for an ideal elastic material is linear, meaning that an applied stress always results in the same strain. The outer surface of nearly all

tires used on over-the-road vehicles is composed of rubber, which is not an ideal elastic material but a viscoelastic material. A viscoelastic material is a substance having both viscous and elastic components. This means that unlike an ideal elastic material, rubber has a strain rate that is dependent on time [5]. For example, an ideal elastic material under a finite stress will strain a finite amount regardless of how long the strain is applied or the accumulation of previous stresses. A viscoelastic material under a finite stress will experience strain that changes as a function of time. This results in a stress-strain curve in which the strain that results from an applied stress lags in its response to a given stress, this phenomenon is known as hysteresis [5]. An example of stress-strain curves for an ideal elastic material and a viscoelastic material showing hysteresis can be seen in Figure 3.12 where the red region represents energy that is converted to heat during the loading and unloading cycle.




Figure 3.12: Stress-strain plots for ideal elastic material (left) and a viscoelastic material (right). [5]

As a vehicle is traveling down the road, the tread surface of its tires is subjected to periodic stresses as each region comes into contact with the road surface. The stresses can be broken down into two types; normal stresses, which are a result of the mass of the vehicle being supported by the tires, and shear stresses, which are a result of the rotational energy being applied to the tire to accelerate the vehicle. This results in a repeating cycle of deformation and relaxation of the tire

tread material as different parts of the tread surface come into contact with the ground as the tire rotates. This can be thought of as energy being stored in the tire material as it is deformed and then released as the material returns to its equilibrium state. During each cycle, a percentage of this energy is converted to heat due to the viscoelastic nature of the tire material. This heat could be observed as an elevated tire surface temperature. The amount of heat generated by a tire per unit of surface area for a single rotation can be calculated using the following equation

$$Q_{Tire} = \oint \sigma d\epsilon \quad (3.40)$$

where σ is the stress acting on the surface of the tire in units of Newtons per square meter (N/m^2), and ϵ is the strain, or amount of deformation, experienced by the surface of the tire in units of meters. The contour integral is simply calculating the area inside the stress-strain plot as highlighted in red in Figure 3.12. If we consider non-drive wheels traveling at a constant velocity we can ignore the effects of shear stress, leaving the normal stress as the only stress acting on the tire material. At any given moment of the tire's rotation, the portion of the tire in contact with the ground (contact patch) is being acted upon by the road with a force equal to the internal pressure of the tire. As the tire rotates, every position on the surface of the tire will experience periods of loading and unloading. If we assume, for simplicity, that the stress from the ground is equal for the entire area of the contact patch, we can express the stress on the tread of the tire as a rectangular pulse function with a frequency equal to the angular frequency of the rotating tire and a duty cycle linearly related to the downward force acting on the tire as shown in Equation 3.41

$$\sigma(t) = P_{Tire} \cdot COMB\left(\frac{t \cdot v}{2\pi \cdot r_{Tire}}\right) * RECT\left(\frac{t \cdot w_{Tire} \cdot v \cdot P_{Tire}}{4\pi^2 r_{Tire}^2 \cdot F_{Tire}}\right) \quad (3.41)$$

where, P_{Tire} is the pressure of the gas in the tire, v is the linear velocity of the vehicle, F_{Tire} is the force acting on the tire perpendicular to the ground, r_{Tire} is the rolling radius of the tire, and w_{Tire} is the width of the tire's tread. This means, that for two identical vehicles traveling at the same velocity, one unloaded and one loaded, the only difference in the stress acting upon the tire tread as a function of time would be the duration of the stress application (the width of the RECT function). An example of these stress functions can be seen in Figure 3.13.




Figure 3.13: Predicted force function for a position on the tire tread for a loaded and unloaded tire traveling at the same velocity

By combining equations 3.40 and 3.41 we see that the heat generated by the tire is dependent on the tire strain function. The tire strain as a function of time would be extremely difficult to measure and will need to be simulated. The simulation would utilize either the well documented properties of rubber, or laboratory testing to determine the strain curve that would result from the stress curve calculated using Equation 3.41 for a number of vehicle velocities, applied pressure durations and tire parameters. These curves would be combined using Equation 3.40 and used to generate a look-up-table relating the normal force being applied to the tire to; heat generation per square meter of tire tread surface per tire rotation, tire pressure, rolling radius, vehicle velocity, and tread width as seen below:

$$F_{Tire} = LUT_{Force}(Q_{Tire}, P_{Tire}, r_{Tire}, v_{Tire}, w_{Tire}) \quad (3.42)$$

The amount of thermal energy generated per rotation of the tire can be calculated from measurable parameters using the following equation

$$H_{TireTotal} = (T_f - T_i) \cdot c_{Tires} \cdot m_{Tires} + H_{Cooling} \quad (3.43)$$

where $H_{TireTotal}$ is the total heat generated by the tire during a single rotation, T_i and T_f are

the tire surface temperatures measured at the beginning and end of a tire rotation, respectively, c_{Tires} is the specific heat capacity of the material making up the tire surface, m_{Tires} is the mass of the tire surface, and $H_{Cooling}$ is the amount of thermal energy lost to cooling during the rotation. From the calculation of total heat generated by the tire during a single rotation, the amount of heat generated by a tire per unit of surface area for a single rotation can be calculated using the following equation

$$Q_{Tire_M} = \frac{H_{Tire_{Total}}}{w_{Tire} \cdot l_{Tire}} \quad (3.44)$$

where Q_{Tire_M} is the measured heat generated per unit area by the tire surface during a single rotation with units of Joules per square meter (J/m^2). Combining Equations 3.41, and 3.40 we can relate measured heat generated during a single rotation to the mass being supported by the tire using the following equation

$$m = \frac{\sum_{i=1}^N LUT_{Force} \left(Q_{Tire_{Measured_i}}, P_{Tire_i}, r_{Tire_i}, v, w_{Tire_i} \right)}{-g} \quad (3.45)$$

where $-g$ is the acceleration due to gravity. This means that the mass of a vehicle could be calculated from the knowledge of the internal pressure of the tires, tire material properties, rate of cooling of the tire surface, vehicle velocity, rolling radius of the tires, tread width of the tires, and the surface temperature of the tire tread at the beginning and end of a rotation. Tire pressure would need to be estimated based on the recommended inflation pressure of similarly sized tires. The tire material properties (specific heat capacity and tread mass) will need to be estimated based on the size of the tire and the known properties of rubber. The rate of cooling of the tire surface will have to be estimated based on environmental conditions and vehicle speed. If we assume that the internal pressure of the tires, material properties of the tire, and the cooling rate of the tire surface can be estimated based on vehicle type, tire size, vehicle speed and environmental conditions, then our ability to calculate the mass of a vehicle using Equation 3.45 is dependent on our ability to accurately measure the vehicle velocity, tire dimensions, and tire surface temperature.

3.9.1 Measurement Techniques

Remotely measuring the speed of a vehicle as a function of time is possible using several methods. Section 3.1.1 discusses the mechanism by which this measurement is possible and identifies methods by which this measurement is made. Remotely measuring the temperature of a vehicle's tire is possible due to self-emission. Section 3.2.1 discusses the mechanism by which this measurement is possible and identifies methods by which this measurement is made.

Calculating vehicle mass from changes in surface temperature of the tire tread is a difficult proposition. The calculation would involve several highly complex measurements to be made requiring precise aiming, tracking, and timing. Additionally, multiple assumptions and estimations will need to be used for parameters that cannot be measured directly. As a result, it is unlikely that this method will result in a useful estimate of vehicle mass.

3.10 Vibration Response

The vibration induced by a vehicle onto the medium over which it is traveling is affected by the total mass of the vehicle. When a vehicle is traveling over a road surface, vibrations are induced by its travel. Particularly, when a vehicle is traversing an object such as a speed bump or a railroad track, the vibrational response induced onto the object will be larger for a vehicle carrying a full load than that of the same vehicle carrying a partial load. When the tire of a vehicle comes into contact with a raised object, like a railroad track, a combination of forces are rapidly imparted to the railroad track. This sudden application of force is called, mechanical shock, which is described as a vibration where the excitation is non-periodic, such as a pulse or a step [24]. The magnitude and duration of the mechanical shock experienced by the railroad track as the vehicle travels over is dependent on the mass and the velocity of the vehicle. Without finite element analysis (FEA) modeling of the railroad track system it would be quite difficult to predict the relationship between vehicle velocity, mass, and the mechanical shock experienced by the railroad track. However, we can assume that the mechanical shock experienced by the railroad track will be correlated in some way to the mass and velocity of the vehicle. If this correlation, in fact does exist, and is measurable, we can quantify the relationship and justify more extensive FEA modeling. In the

subsequent section we will discuss a method of measuring mechanical shock.

3.10.1 Measurement Techniques

To measure the vibration induced by a vehicle onto a railroad track, the acceleration of the railroad track as a function of time must be accurately measured. This measurement can be made using calibrated accelerometers rigidly mounted to the railroad track. In addition, the distance between the accelerometer and the point of contact between the railroad track and the tire will need to be measured, as will the velocity of the test vehicle. Both of these measurements are possible using the techniques discussed in Section 3.1.1.

Vibration Response Measurements with a Calibrated Accelerometer

The vibration response of an object can be measured using precision accelerometers, designed for mechanical shock and impact testing. These accelerometers are rigidly mounted to the object-under-test and output an electrical signal that describes the acceleration experienced by the object. Typically, the accelerometers are only sensitive in a single dimension and are highly sensitive to orientation. Three of these one-dimensional accelerometers can be combined perpendicularly to form a triaxial accelerometer, which measures acceleration in three-dimensions.

Chapter 4

Collection Campaign

After the identification of the phenomenologies presented in Chapter 4, testing was conducted to ascertain how well, if at all, each phenomenology could be exploited to estimate the mass of a vehicle. This chapter will describe in detail the various experiments that were conducted to quantify the correlation between phenomenology and vehicle mass.

4.1 Wyoming County International Speedway Full-Scale Field Test

The first full-scale field test was designed to test the brake temperature, suspension response, engine acoustics, and tire deformation phenomenologies. This test was carried out at the Wyoming County International Speedway, located in Perry, New York.

4.1.1 Test Vehicle

A 1992 Ford F600 flatbed dump-truck was chosen as the test vehicle for these experiments (Figure 4.1). This test vehicle has a gross vehicle weight rating (GVWR) of 23,000 pounds. The gross vehicle weight rating is the maximum operating mass of the vehicle, and is specified by the manufacturer. The truck is powered by a 5.9 L Cummins 4-stroke, 6-cylinder diesel engine and has a manual transmission with a two-speed rear axle. A two-speed rear axle refers to a rear axle that has a high and a low speed, this allows for so-called half gears, whereas each gear can be operated

with the rear axle in low (which would be considered half) and high (which would be considered full). A two-speed rear axle is used in medium to heavy trucks mainly to assist the truck when beginning to move while carrying a heavy load (the rear axle in low gear is designed for low speeds and provides more power). The braking system of this test vehicle consists of disc-brakes for the front axle, and drum brakes for the rear axle. The flatbed could be raised and lowered using hydraulics, which provided easy access to the frame for sensor mounting. This test vehicle was selected because of its varied braking methods (drum and disc), manual transmission with split rear end (similar to the transmission found in large trucks), functioning dump flatbed, and a high GVWR (just under the 26,000 pound limit that would have required a commercial driver's license to operate).



Figure 4.1: The 1992 Ford F600 chosen for as a test vehicle

4.1.2 Test Site

The Wyoming County International Speedway, located in Perry, NY was selected as the site of the data collect. This site was chosen because of the flexibility of the management, the proximity to the campus, and physical layout. The management allowed unfettered access to the track, and provided several large concrete blocks weighing approximately 4,000 pounds each to be used to load the test vehicle. Additionally, the management agreed to load and unload the blocks from the test vehicle during the field tests using track equipment. The track is located less than 40 miles from the RIT campus allowing for easy travel. The physical layout of the racetrack is a third of a mile asphalt oval track with a mild to moderate bank. An aerial photograph of the track can be seen in Figure 4.2.



Figure 4.2: The Wyoming County International Speedway

4.1.3 Test Description

During this full-scale field test, we conducted experiments to determine the validity of the brake temperature, suspension response, engine acoustics, and tire pressure phenomenologies as they pertain to predicting vehicle mass. Each test was executed at least three times for each of four different vehicle load amounts. The load amounts spanned the full range of acceptable loads for the test vehicle, from unloaded to nearly the full GVWR.

Brake Temperature Test

To test the brake temperature phenomenology, an experiment was designed and executed to measure the temperature change induced by the braking system during a deceleration to a stop from a known velocity over a predetermined distance with the transmission disengaged (clutch in) and with the transmission engaged (clutch out). The test vehicle was driven around the test track at a predetermined speed with no braking. Immediately prior to the vehicle entering the field-of-view of the imaging system, at a location marked with cones, the vehicle began aggressively braking and came to a complete stop at a position outside the field-of-view of the imaging system. This was repeated with the transmission engaged and transmission disengaged. Temperature data was collected using on-board and off-board sensors during the entirety of each test.

Suspension Response Test

To test the suspension response phenomenology, an experiment was designed and executed to measure the response of the test vehicle's suspension as it traveled over a standard-sized speed bump. A pair of portable speed bumps was acquired for this testing and can be seen in Figure 4.3. The test vehicle was driven over a speed bump at a constant velocity while position data of the test vehicle's frame and wheels was collected using on-board and off-board sensors.



Figure 4.3: Portable speed bumps acquired for the suspension response test

Engine Acoustics Test

To test the engine acoustics phenomenology, an experiment was designed and executed to measure the test vehicle's acoustic emission as it accelerated to a predetermined velocity. The test vehicle accelerated from a stop, to a predetermined velocity while acoustic data was collected using on-board and off-board sensors.

Tire Deformation Test

As a preliminary test for the tire deformation phenomenology, an experiment was designed and executed to measure the test vehicle's tire pressure as it accelerated and decelerated under various loading conditions. Additionally, the effect of environmental heating caused by changes in air temperature and solar loading was studied. The purpose of the test was to quantify the changes in internal tire pressure experienced by a tire during driving. The main cause for changes in the

internal tire pressure is temperature fluctuations. The test vehicle was accelerated from a stop to a predetermined speed, and then decelerated quickly to a stop while tire pressure data was collected using on-board sensors.

4.1.4 On-Board Sensor System

Multiple on-board sensors were installed on the test vehicle to collect temperature, acoustic, vibration, and tire pressure data in real time during the field collection.

Temperature Measurements

On-board temperature measurements were collected to ascertain the temperature range experienced by various vehicle systems and how they relate to remote surface temperature measurements. Eleven K-Type thermocouples were installed on the truck to measure the bulk material temperature of various system components. K-Type thermocouples are suitable for an approximate temperature range of -200 to 1350°C with an accuracy of $\pm 2.2^{\circ}\text{C}$ for temperatures greater than 0°C . Six of the thermocouples were installed on the vehicle's braking system. For the rear brakes (drum-style), one thermocouple was bolted to each brake shoe (4 total) as seen in Figure 4.4. For the front brakes, one thermocouple was bolted to each of the inside front brake pads (2 total) as seen in Figure 4.5. One thermocouple was bolted to the inside of the exhaust pipe as seen in Figure 4.6. Two thermocouples were affixed to the rear differential using adhesive thermal pads after thoroughly scouring the area with a grinder to expose bare metal as seen in Figure 4.7. One thermocouple each was installed on the oil pan and the inner cooler using adhesive thermal pads.

Two Micro-Epsilon TIM-160 microbolometer thermal imagers were installed to measure the surface temperature of rotating components; this includes the brake discs, brake drums, and tires. The Micro-Epsilon TIM 160 microbolometer has an optical resolution of 160-by-120 pixels, operates in a spectral range of 7.5 to $13\text{ }\mu\text{m}$, has a sensitivity of 0.1 K , a 64° field-of-view, and a frame rate of 10 Hz . One thermal imager was mounted in the front driver side wheel-well directed at the driver side front brake and tire as seen in Figure 4.8. The second thermal imager was mounted on the truck frame and directed towards the rear passenger side brake and tires as seen in Figure 4.9. In both cases, the thermal imagers were installed using adjustable fixtures that allowed nearly

360° of rotation.




Figure 4.4: The location of the thermocouples installed on the rear brakes



Figure 4.5: The location of the thermocouples installed on the front brakes




Figure 4.6: Location of the thermocouple installed on the exhaust




Figure 4.7: Location of the thermocouples installed on the rear differential



Figure 4.8: Location of the thermal imager installed on the rear of the truck, directed at the rear brakes



Figure 4.9: Location of the thermal imagers installed in the front, driver side wheel-well of the truck, directed at the front brakes

Vibration Measurements

On-board vibration measurements were collected to serve as truth data for the acceleration, velocity, and position of the vehicle's frame and axle. To collect this data, two PCB Piezotronics triaxial accelerometers were installed. A triaxial accelerometer measures acceleration as a function of time along three orthogonal vectors. The PCB Piezotronics triaxial accelerometers used here have a sensitivity of 10 mV per 9.8 m/s^2 (acceleration due to gravity), and can measure frequencies between 2 and 4000 Hz. One triaxial accelerometer was mounted on a steel plate that was welded to the vehicle's frame, while the other was mounted on a steel plate welded to the vehicle's rear axle. The location of both triaxial accelerometers can be seen in Figure 4.10.




Figure 4.10: Location of the triaxial accelerometers that were installed on the rear axle and the frame

Acoustic Measurements

On-board acoustic measurements were collected to serve as truth data. To collect this data, two calibrated PCB Piezotronics free-field condenser microphones were used. These microphones have a sensitivity of 50 mV/Pa, and can measure frequencies between 3.15 and 20,000 Hz. A

microphone was installed in the engine compartment directed at the valve train as seen in Figure 4.11. A second calibrated microphone was installed on the frame of the test vehicle, directed at the exhaust pipe as seen in Figure 4.12. The microphones were installed using custom mounting fixtures that allow the position and direction of the microphones to be adjusted easily.



Figure 4.11: Location of the microphone installed in the engine compartment directed at the valve train




Figure 4.12: Location of the microphone installed on the frame of the test vehicle directed at the exhaust

Tire Pressure Measurements

On-board tire pressure measurements were collected to monitor the change in tire pressure during the collection. Six Pressure Pro wireless pressure sensors were installed, one on each of the test vehicle's tires. These wireless pressure sensors have an accuracy of $\pm 3\%$, can measure pressures ranging from 55 to 2060 kPa, and update every five minutes. Figure 4.13 shows one of the wireless pressure sensors mounted on the test vehicle. The wireless receiver was installed in the cab of the vehicle.



Figure 4.13: One of the six wireless tire pressure monitors that were used to monitor tire pressure

Synchronization

It was important to be able to synchronize the measurements made with the on-board system with the measurements made with the off-board system. To do so, a global positioning system (GPS) receiver was used. The GPS receiver provided an extremely accurate time signal at 1 second intervals which were recorded with both the on-board and the off-board systems. Additionally, the GPS provided coarse position and ground speed information at 1 second intervals. To synchronize off-board imaging systems that could not interface with the GPS receiver, a pulsing light emitting diode (LED) array was mounted on the test vehicle and designed to pulse each time a GPS time signal was received. The LED array installed on the bed of the truck can be seen in Figure 4.14.



Figure 4.14: One of two LEDs installed on the bed of the truck which were used to synchronize the data collected by the off-board system with the data collected by the on-board system

Data Acquisition System

The on-board data acquisition system monitored and recorded data from each of the on-board sensors and controlled the synchronization LEDs. The on-board data acquisition system consisted of a DC-powered computer fitted with two 30 GB solid state drives, an Intel Atom processor, and a National Instruments CompactDAQ 8-slot chassis, equipped with three National Instruments Sound and Vibration DAQ modules, a thermocouple module, and a digital I/O module. The computer directly monitored and recorded data from the tire pressure monitors via their USB connected receiver, and the two microbolometers via direct USB connections. The National Instruments CompactDAQ chassis, and the installed modules, were used to control and monitor the accelerometers, microphones, thermocouples, and synchronization LEDs. After being digitized, data was passed from the National Instruments CompactDAQ chassis to the computer for storage on the solid state drives. The on-board data acquisition system was installed in the cab of the truck as seen in Figure 4.15.



Figure 4.15: The computer and data acquisition system installed in the cab of the truck

4.1.5 Off-Board Sensor System

The purpose of the off-board sensing system was to collect surface temperature, acoustic, and position data in real time during the field collection. The off-board system includes the Wildfire Airborne Sensor Program (WASP) infrared sensor system, two calibrated microphones, and a high-definition video camera.

RIT WASP System Measurements

Off-board measurements were made with RIT's WASP sensor to remotely measure the surface temperature of the test vehicle's tires, brakes, and exhaust system. RIT's WASP instrument is made up of four imaging systems, each mapping a different spectral range. Three of the imagers are 640-by-512 pixel infrared cameras each operating in a different spectral range. The short-wave infrared (SWIR) imager operates from 0.9 to 1.7 μm , the mid-wave infrared (MWIR) imager operates from 3.0 to 5.0 μm , and the long-wave infrared (LWIR) imager operates from 8.0 to 9.2 μm . Each of

these imagers has a pixel pitch of $25\text{ }\mu\text{m}$ and is equipped with a 25 mm focal length lens. The fourth camera of the WASP instrument is a high-resolution RGB camera. This camera has a pixel pitch of $25\text{ }\mu\text{m}$, a resolution of 4000-by-2672 pixels, and a 50 mm lens. The WASP system was positioned near the road and oriented so that it could observe the entire test vehicle for multiple frames as it drove past (Figure 4.16).



Figure 4.16: The WASP system setup to monitor the test vehicle

Acoustic Measurements

Off-board measurements were made with calibrated microphones to determine if quality acoustic measurements could be collected of a vehicle traveling past a stationary sensor, and if vehicle parameters could be determined (e.g. engine RPM, number of cylinders, exhaust brake applied). Additionally, the difference in signal for an off-board system as compared to an on-board system is of interest to investigate the impact of a changing sensor to target distance as it relates to signal quality. The microphones were mounted near the road and used to gather acoustic pass-by data.

Position Measurements

Off-board position measurements were made with a high-definition video camera system to remotely measure the acceleration, velocity, and position of the vehicle as a function of time. The video camera system used for this test was the Kodak Zi6, which captures video in a progressive high-definition format (720p) at 60 fps. The video camera system was positioned on the side of the track, oriented perpendicular to the path of the vehicle for all testing.

Data Acquisition System

The off-board data acquisition system monitored and recorded data from the microphones. The system consisted of a laptop and a National Instruments USB DAQ module. The National Instruments USB DAQ module was used to control and monitor the off-board microphones. After being digitized, data was passed from the National Instruments USB DAQ module to the laptop for storage. The video camera system is self-contained and stored all data internally. RIT's WASP system was controlled, monitored, and its data stored, using its computer system and custom hardware.

4.1.6 Test Execution

This field test was carried out on November 20th and 21st in 2010 at the Wyoming County International Speedway in Perry, NY. Prior to the start of testing, the test vehicle was taken to a calibrated truck scale located at a nearby farm. The test vehicle was weighed unloaded, loaded with one concrete block, and loaded with two concrete blocks (Table 4.1). This averaged out to 4,010 pounds

per concrete block. A third block was added to the truck bringing the estimated weight of the truck up to 23,550 lbs. The vehicle loaded with three blocks could not be weighed due to road weight limit restrictions en route to the truck scale. Figure 4.17 shows the test vehicle being loaded, and Figure 4.18 shows the test vehicle being weighed.

Table 4.1: The test vehicle loading conditions

Test Vehicle Loading	Total Vehicle Mass (lbs.)
Unloaded	11,520
1 Block	15,530
2 Blocks	19,540
3 Blocks	23,550 (Estimated)



Figure 4.17: The test vehicle being loaded at the track



Figure 4.18: The test vehicle being weighed

Brake Temperature Test

Prior to the start of the brake temperature test, the top speed and required braking distance was determined by preliminary testing of the fully-loaded test vehicle (total weight of 23,550 lbs). The top speed that could be achieved in the allotted distance (approximately a quarter of a mile) was found to be approximately 30 miles per hour (mph), and was used for all of the testing, regardless of loading. The braking distance required for the fully-loaded vehicle was used to identify the location where braking should begin for all test runs. The driver was instructed to begin braking at the designated location and come to a complete stop at a second designated location. There were eight combinations of vehicle loading and transmission engagement (engaged or disengaged), the test was executed for each combination three times for twenty-four total test runs. Both the on-board and off-board sensing systems operated properly for all twenty-four test runs.

Suspension Response Test

Prior to the start of the suspension response test, fiducials were added to the test vehicle to improve our ability to track specific positions on the test vehicle in video frames. Fiducial markers were added to the wheels of the test vehicle and a location on the frame directly above the wheels.

Figure 4.19 shows the fiducial markers being added to the test vehicle. The driver was instructed to traverse the speed bump at the maximum safe velocity for the fully-loaded vehicle. After several preliminary runs, it was determined that approximately 5 mph was the maximum safe velocity to traverse the speed bump with the fully-loaded vehicle. The test was then repeated for all four test vehicle loads, 3 times each, for a total of 12 test runs. On-board triaxial accelerometers measured the frame and rear axle acceleration response of the test vehicle, while off-board, a high-definition video system recorded the vehicle at a rate of approximately 30 fps. Figure 4.20 shows the test in progress. Both the on-board and off-board sensing systems operated properly for all twenty-four test runs.



Figure 4.19: A fiducial marker being added to the frame of the test vehicle



Figure 4.20: The suspension response test in progress for the test vehicle loaded with 3 concrete blocks

Engine Acoustic Test

The engine acoustic test was incorporated into the brake temperature test. At the beginning of each brake temperature test run, the test vehicle was accelerated rapidly to the predetermined velocity of 30 mph. Only this acceleration phase of the testing is relevant to the engine acoustic test. The acoustic emission of the test vehicle was measured by the on-board and off-board calibrated microphones during each test run. There were a total of twenty-four test runs (6 runs for each vehicle load, since transmission engaged/disengaged braking does not factor into this test). Both the on-board and off-board sensing systems operated properly for all twenty-four test runs.

Tire Pressure Test

The tire pressure test was executed continuously throughout all of the test runs. The wireless tire pressure monitors operated properly for the duration of the tests.

4.2 Savannah River Site Full-Scale Field Test I

The second full-scale field test was designed to test the brake temperature, frame twist, engine acoustic, exhaust braking acoustic, vibration response, tire deformation, and tire temperature phenomenologies. This test was carried out at the Savannah River Site, located in Aiken, South Carolina.

4.2.1 Test Vehicle

A Western-Star tractor pulling a 40-foot long, three-axle trailer was chosen as the test vehicle and can be seen in Figure 4.21. This test vehicle is powered by a 15.2 L Caterpillar diesel engine model C15. The C15 is an in-line six cylinder diesel engine with a bore diameter of 5.4 inches, a stroke of 6.75 inches, an average output of 550 horsepower at 2,100 rpm, and a weight of nearly 3,000 pounds. The tractor has a manual transmission and is equipped with a Jacobs-style exhaust brake. The 40-foot long, three-axle trailer has a GVWR of 90,628 pounds for 60 mph travel, and 116,000 pounds for 20 mph travel. This vehicle was selected because it represents a typical over-the-road truck.



Figure 4.21: The tractor-trailer chosen for Savannah River Site field tests

4.2.2 Test Site

The Savannah River Site, located in Aiken, SC was selected as the site of the data collect. This site was chosen because of the access to a tractor-trailer test vehicle, site support for loading and

unloading, and the physical layout. The test vehicle, a full size tractor-trailer, was provided for this testing as well as a licensed driver. Heavy equipment and licensed operators were provided for loading and unloading of the tractor-trailer test vehicle. The test location was a relatively flat, paved road, closed to most traffic, and crossed by railroad tracks in multiple locations. The railroad tracks were used as impulse bumps for the vibration response testing. The physical layout of the test location is a rectangle with approximately 2.4 km of road. An aerial photograph of the track can be seen in Figure 4.22.



Figure 4.22: The Savannah River Site testing location

4.2.3 Test Description

During this full-scale field test, experiments were conducted to determine the validity of the brake temperature, frame twist, engine acoustic, exhaust braking acoustic, vibration response, suspension response, tire deformation, and tire temperature phenomenologies as they pertain to predicting

vehicle mass. Each test was executed at least three times for each of three different vehicle load amounts. The load amounts spanned the full range of acceptable loads for the test vehicle, from unloaded to nearly the full GVWR.

Brake Temperature

To test the brake temperature phenomenology, an experiment was designed and executed to measure the temperature change induced by the braking system during a deceleration to a stop from a known velocity, over a predetermined distance, with the transmission disengaged (clutch in). The test vehicle was driven around the test track at a predetermined velocity with minimal braking. Immediately prior to the vehicle entering the field-of-view of the WASP system, at a location marked with cones, the vehicle began aggressively braking and came to a complete stop at a position outside the field-of-view of the WASP system. Temperature data was collected using on-board and off-board sensors during the entirety of each test.

Frame Twist

To test the frame twist phenomenology, an experiment was designed and executed to measure the frame twist of the test vehicle's suspension as it accelerated aggressively from a stop. The test vehicle was accelerated aggressively from a stop for a short distance while position data of the test vehicle's frame was collected using on-board and off-board sensors.

Engine Acoustic

To test the engine acoustics phenomenology, an experiment was designed and executed to measure the test vehicle's acoustic emission as it accelerated to a predetermined velocity. The test vehicle was accelerated from a stop to a predetermined velocity while acoustic data was collected using on-board and off-board sensors.

Exhaust Braking Acoustic

To determine if the acoustic signature of a vehicle's exhaust brake was identifiable, an experiment was designed and executed to measure the test vehicle's acoustic emission as it applied an exhaust brake. The test vehicle was accelerated to a predetermined velocity, maintained the velocity until it reached the designated location, began braking using the both exhaust braking and friction braking, and came to a complete stop at the marked location. Acoustic data was collected using on-board and off-board sensors during the entirety of each test.

Vibration Response

To test the vibration response phenomenology, an experiment was designed and executed to measure the vibration response induced by the test vehicle onto a set of railroad tracks as the test vehicle traveled over them. The test vehicle was driven over the railroad tracks at a constant, predetermined velocity while vibration response data was collected using off-board sensors mounted to the railroad tracks.

Suspension Response

To test the suspension response phenomenology, an experiment was designed and executed to measure the response of the test vehicle's suspension as it traveled over a standard-sized speed bump. A pair of portable speed bumps was used for this testing and can be seen in Figure 4.3. The test vehicle was driven over the speed bump at a constant, predetermined speed while position data was collected using off-board sensors.

Tire Deformation

To test the tire deformation phenomenology, an experiment was designed and executed to measure the amount of deformation experienced by the test vehicle's tires as it traveled at a constant velocity. The test vehicle was driven past a calibrated camera system at a predetermined velocity. The high-definition camera system was setup to capture images of the test vehicle's tires as it passed. Additionally, control data was taken of each of the test vehicle's tires while the vehicle was sta-

tionary. Tire deformation data was collected using off-board sensors during the entirety of each test.

Tire Temperature

To test the tire temperature phenomenology, an experiment was designed and executed to measure the temperature change of the test vehicle's tires as it accelerated, decelerated and maintained a constant velocity under various loading conditions. The test vehicle was driven around the test track at a predetermined velocity with minimal braking. Immediately prior to the vehicle entering the field-of-view of the WASP system, at a location marked with cones, the vehicle began aggressively braking and came to a complete stop at a position outside the field-of-view of the WASP system. Temperature data was collected using on-board and off-board sensors during each test.

4.2.4 On-Board Sensor System

Multiple on-board sensors were installed on the test vehicle to collect temperature, and acoustic data in real time during the field collection.

Temperature Measurements

On-board temperature measurements were collected of the test vehicle using two Micro-Epsilon TIM-160 microbolometer thermal imagers. The Micro-Epsilon TIM 160 microbolometer has an optical resolution of 160-by-120 pixels, operates in a spectral range of 7.5 to 13 μm , has a sensitivity of 0.1 K, a 64° field-of-view, and a frame rate of 10 Hz. One thermal imager was mounted on the truck frame, on the driver side, directed at the passenger side front brake and tire as seen in Figure 4.23. The second thermal imager was mounted on the truck frame, on the passenger side of the test vehicle, directed at the tractor's drive wheels on the driver side as seen in Figure 4.24. In both cases, the thermal imagers were installed using adjustable fixtures that allowed nearly 360° of rotation. A thermocouple was installed in the field-of-view of each thermal imager to provide an in-scene temperature reference.



Figure 4.23: Location of the thermal imager installed on the front driver's side of the tractor directed at the front passenger's side brakes and tire




Figure 4.24: Location of the thermal imager installed on the front passenger's side of the tractor directed at the rear driver's side brakes and tire

Acoustic Measurements

On-board acoustic measurements were collected to serve as truth data. To collect this data, two calibrated PCB Piezotronics Free-Field Condenser microphones were used. These microphones have a sensitivity of 50 mV/Pa, and can measure frequencies between 3.15 and 20,000 Hz. Both calibrated microphones were installed on the trailer, directed towards the tractor as seen in Figure 4.25. The microphones were installed using custom mounting fixtures that allow the position and direction of the microphones to be adjusted easily.




Figure 4.25: Location of the microphones installed on the test vehicle trailer directed at the tractor

Position Measurements

Coarse on-board position and velocity measurements were collected using a GPS receiver. The GPS receiver provides position and ground speed information at a rate of 5 Hz.

Data Acquisition System

The on-board data acquisition system monitored and recorded data from each of the on-board sensors. The on-board data acquisition system consisted of a DC-powered computer fitted with two 30 GB solid state drives and an Intel Atom Processor, and a National Instruments, CompactDAQ 8 Slot Chassis, equipped with one National Instruments Sound and Vibration DAQ module, a thermocouple module, and a Digital I/O module. The computer directly monitored and recorded data from the two microbolometers via direct USB connections. The National Instruments CompactDAQ Chassis and installed modules were used to control and monitor the microphones and thermocouples. After being digitized, data was passed from the National Instruments Compact-

DAQ Chassis to the computer for storage on the solid state drives. The on-board data acquisition system was installed on the test vehicle's trailer as seen in Figure 4.26.



Figure 4.26: The computer and data acquisition system installed on the test vehicle's trailer

4.2.5 Off-Board Sensor System

The purpose of the off-board sensing system was to collect surface temperature, acoustic, vibration, and position data in real time during the field collection. The off-board system includes the Wildfire Airborne Sensor Program (WASP) infrared sensor system, two calibrated microphones, two calibrated triaxial accelerometers, and a high-definition video camera.

RIT WASP System Measurements

Off-board measurements were made with RIT's WASP sensor to remotely observe the surface temperature of the test vehicle's tires, brakes, and exhaust system. The WASP system was posi-

tioned near the road and oriented so as to observe the entire test vehicle for multiple frames as it drove past as seen in Figure 4.27.



Figure 4.27: The location of RIT's WASP sensor during testing

Acoustic Measurements

Off-board measurements were made with calibrated PCB Piezotronics Free-Field Condenser microphones. The purpose of these measurements was to determine if high quality acoustic measurements can be made of a vehicle traveling past a stationary sensor, that vehicle parameters can be determined (e.g. engine RPM, number of cylinders, exhaust brake applied). The microphones have a sensitivity of 50 mV/Pa, and can measure frequencies between 3.15 and 20,000 Hz. The microphones were mounted near the road and used to gather acoustic pass-by data.

Vibration Measurements

Off-board measurements were made with calibrated triaxial accelerometers to characterize the vibration response induced by the vehicle onto railroad tracks as the vehicle traveled over them. To collect this data, two PCB Piezotronics triaxial accelerometers were mounted on the railroad tracks near the road shoulder as seen in Figure 4.28. The PCB Piezotronics triaxial accelerometers

used here have a sensitivity of 10 mV per 9.8 m/s^2 (acceleration due to gravity) and can measure frequencies between 2 and 4000 Hz.



Figure 4.28: Location of the triaxial accelerometers that were mounted on the railroad tracks

Position Measurements

Off-board position measurements were made with a high-definition video camera system to remotely measure the acceleration, velocity, and position of the vehicle as a function of time. The video camera system used for this test was the 18-megapixel Canon T2i. The video camera system was positioned in various, experiment dependent, locations throughout the testing. For the brake temperature, suspension response, tire deformation, and tire temperature experiments the video camera system was located on the road shoulder, oriented perpendicular to the path of the test vehicle. For the vibration response, engine acoustic, and frame twist experiments, the video camera was located directly in the path of the test vehicle. For the tire deformation test, a laser trigger was used to initiate the collection of multiple $1/4000$ of a second exposure images. The laser trigger and camera were positioned and adjusted so as to collect images of the vehicle's tires when traveling at a specific velocity. Figure 4.29 shows the camera and laser trigger setup.



Figure 4.29: The camera and laser trigger used to capture images of the test vehicle's tires while moving

Data Acquisition System

The off-board data acquisition system monitored and recorded data from the microphones and accelerometers. The system consisted of a laptop, a National Instruments USB DAQ module, National Instruments cDAQ chassis, and two pressure and vibration modules. The National Instruments modules were used to control and monitor the two off-board microphones, and two accelerometers. After being digitized, data was passed from the National Instruments modules to the laptop for storage via USB. The video camera system is self-contained and stored all data internally. RIT's WASP system was controlled, monitored and its data stored, using its computer system and custom hardware.

4.2.6 Test Execution

This field test was carried out on May 24th and 25th in 2011 at the Savannah River Site in Aiken, South Carolina. The test vehicle was operated in three conditions, empty trailer, 20 tons loaded on the trailer, and 40 tons loaded on the trailer. The brake temperature, frame twist, engine acoustic, exhaust braking acoustic, vibration response, suspension response, tire deformation, and tire

temperature tests were all completed during this period.

Brake Temperature

Prior to the start of the brake temperature test, the maximum velocity and required braking distance was determined by preliminary testing of the fully-loaded vehicle (40 tons loaded on the trailer). The maximum velocity that could be achieved for the fully-loaded test vehicle on the test road was found to be approximately 35 mph. This velocity was used for all of the brake testing runs, regardless of loading. The braking distance required for the fully-loaded test vehicle was used to identify the location at which braking should begin for all test runs. The driver was instructed to accelerate to 35 mph, maintain a constant velocity until he reached the first set of cones where he would begin braking, and come to a complete stop at the second set of cones. The driver was instructed to brake with the transmission disengaged for all testing. Five runs were completed for each vehicle loading for a total of 15 runs. Both the on-board and off-board sensing systems operated properly for all fifteen test runs.

Tire Temperature

The tire temperature test was conducted simultaneously with the brake temperature test. The on-board sensors were directed at both the test vehicle's brakes and tires so data could be gathered for both simultaneously. Periods of positive acceleration, and periods of constant velocity are of particular interest for the tire temperature testing. Five runs were completed for each vehicle loading for a total of 15 runs. Both the on-board and off-board sensing systems operated properly for the entirety of all fifteen test runs.

Exhaust Braking Acoustic

The exhaust braking test was incorporated into the brake temperature test. Two additional brake temperature test runs were made for each vehicle loading condition. During these runs, the driver was instructed to brake using the exhaust brake. Six total runs were completed using the exhaust brake. Only the on-board sensing system was operated for this test.

Frame Twist

Prior to the start of the frame twist test, the distance required by the unloaded test vehicle to accelerate rapidly from a stop to the point where the driver would typically shift, was determined through preliminary testing of the unloaded test vehicle. The high-definition camera system was placed in the path of the test vehicle leaving room for the test vehicle to accelerate towards it and safely come to a stop. The driver was instructed to accelerate rapidly from a stop, at a location marked by cones, towards the camera system until he typically would begin shifting. At the point where he would typically begin shifting, the driver would apply the brakes and decelerate to a stop. Four runs were completed for each vehicle loading for a total of 12 runs. Both the on-board and off-board sensing systems operated properly for all twelve test runs.

Engine Acoustic

The engine acoustic test was incorporated into the frame twist test. The acoustic emission of the test vehicle was measured by the on-board and off-board calibrated microphones during each frame twist test run. Four runs were completed for each vehicle loading for a total of 12 runs. Both the on-board and off-board sensing systems operated properly for all twelve test runs.

Suspension Response

Prior to the start of the suspension response test, a speed bump was placed across a flat and straight portion of road. The driver was instructed to traverse the speed bump at the maximum velocity he felt was safe for the fully-loaded vehicle. After several preliminary runs, it was determined that approximately 2 to 3 mph was the maximum velocity at which the driver felt comfortable traveling over the speed bump with the fully-loaded vehicle. The test was then conducted for all three test vehicle loads five times each, for a total of fifteen test runs. The on-board GPS receiver recorded coarse vehicle position and velocity data, while off-board, a high-definition video system recorded the vehicle at a rate of approximately 30 fps. Figure 4.30 shows the test in progress. Both the on-board and off-board sensing systems operated properly for all fifteen test runs.



Figure 4.30: The suspension response test in progress for the test vehicle loaded with 40 tons, the arrow marks the location of the speed bump

Vibration Response

Prior to the start of the vibration response test, the driver was instructed to traverse the railroad tracks at the maximum velocity he felt was safe for the fully-loaded vehicle. After several preliminary runs, it was determined that approximately 2 to 3 mph was the maximum velocity at which the driver felt comfortable traveling over the railroad tracks with the fully-loaded vehicle. The test was then conducted for all three vehicle loads five times each, for a total of fifteen test runs. The on-board GPS receiver recorded coarse vehicle position and velocity data, while off-board a high-definition video system recorded the vehicle frame at a rate of approximately 30 fps, and the accelerometers mounted to the railroad tracks recorded the vibrational response the test vehicle induced. Figure 4.31 shows the test in progress. Both the on-board and off-board sensing systems operated properly for all fifteen test runs.



Figure 4.31: The vibration response test in progress for the test vehicle loaded with 40 tons, the circles mark the locations of the accelerometers

Tire Deformation

Prior to the tire deformation test, the laser trigger and high-definition imaging system were setup along the path of the test vehicle and adjusted for a test vehicle velocity of approximately 30 mph. The test vehicle was driven past the laser trigger and imaging system at a constant speed of approximately 30 mph. The laser trigger and high-definition imaging system captured short exposure images of each of the test vehicle's tires as the vehicle passed. The test was conducted for all three test vehicle loads five times each, for a total of fifteen test runs. In addition, control data was collected of the test vehicle's tires while the vehicle was stationary for all three test vehicle loads.

4.3 Savannah River Site Full-Scale Field Test II

The third full-scale field test was designed to test the frame twist, tire deformation, and suspension response phenomenologies. This test was carried out at the Savannah River Site, located in Aiken, South Carolina.

4.3.1 Test Vehicles

The Western-Star tractor pulling a 40-foot long, three-axle trailer described in Section 4.2.1 was chosen as a test vehicle for these experiments.

A 2006 Honda Accord was chosen as a test vehicle and can be seen in Figure 4.32. This test vehicle is a 4-door model powered by a 2.4 L I4 engine. The vehicle has a 5-speed manual transmission. Unloaded (with two passengers) the test vehicle weighed approximately 3,340 pounds. This vehicle was selected because it represents a typical passenger vehicle.



Figure 4.32: The Honda Accord used for this field test

4.3.2 Test Site

The Savannah River Site, located in Aiken, SC was selected as the site of the data collect. This site was chosen because of the access to a tractor-trailer test vehicle, site support for loading and unloading, and the physical layout. The primary test vehicle, a full size tractor-trailer, was provided for this testing as well as a licensed driver. Heavy equipment and licensed operators were provided for loading and unloading of the tractor-trailer test vehicle. The test location was a relatively flat, paved road, closed to most traffic. The physical layout of the test location is a rectangle with approximately 2.4 km of road. An aerial photograph of the track can be seen in Figure 4.22.

4.3.3 Test Description

During this full-scale field test, experiments were conducted to further quantify the correlation between vehicle mass and the frame twist, tire deformation, and suspension response phenomenologies. Each test was executed at least five times for each of three different vehicle load amounts. The load amounts spanned the full range of acceptable loads for the test vehicle, from unloaded to nearly the full GVWR.

Frame Twist

For this test of the frame twist phenomenology, an experiment was designed and executed to measure the test vehicles' frame twist as they accelerated from a stop. This test was designed to simulate a real-world scenario. All of the measurements were made remotely from a ground based high-speed video camera system positioned at an off-road location. The test vehicles were accelerated aggressively from a stop for a short distance while position data of the test vehicles' frame was collected using the high-speed video camera system positioned at an off-road location.

Tire Deformation

For this test of the tire deformation phenomenology, an experiment was designed and executed to measure the test vehicles' tire deformation as they traveled at low velocity. This test was designed to simulate a real-world scenario. All of the measurements were made remotely, from a ground based LiDAR system positioned at an off-road location. The test vehicles traveled at a constant, low velocity while three-dimensional position measurements were collected using the off-board LiDAR system.

Suspension Response

For this test of the suspension response phenomenology, an experiment was designed and executed to measure the test vehicles' suspension response as they traveled over a speed bump. This test was designed to simulate a real-world scenario. All of the measurements were made remotely, from a ground based high-definition video camera system positioned at an off-road location. The test

vehicles were driven over the speed bump at a constant velocity while position data was collected using the off-board high-definition camera system.

4.3.4 On-Board Sensor System

Since the purpose of this field test was to simulate true, real-world test conditions, the only on-board sensor that was installed on either test vehicle was for the collection of acoustic data. The on-board acoustic data was collected because off-board acoustic data was unavailable for this testing.

Acoustic Measurements

On-board acoustic measurements were collected during this field test for the tractor trailer test vehicle only. To collect this data, a calibrated PCB Piezotronics Free-Field Condenser microphone was used. The microphone has a sensitivity of 50 mV/Pa, and can measure frequencies between 3.15 and 20,000 Hz. The calibrated microphone was installed on the rear windshield of the tractor, directed towards the exhaust. The microphone was installed using a suction cup that allowed for easy removal and adjustment.

4.3.5 Off-Board Sensor System

Off-board position measurements were made for each experiment during this field collection. Due to the unique requirements associated with each experiment, a different instrument was used for each. The frame twist test required an instrument capable of collecting imagery at a high frame rate, the suspension response test required an instrument capable of collecting imagery with high spatial resolution, and the tire deformation test required an instrument capable of high-resolution spatial data in three dimensions.

Position Measurements

The Vision Research Phantom v5.1 was selected as the instrument for the frame twist testing. The Vision Research Phantom v5.1 has an 8-bit, 1024-by-1024 CMOS detector capable of recording

1,200 fps at full resolution. Data is directly recorded to flash memory on the camera, and transferred to a support computer after the triggered collection has occurred in RAW CINE format. The camera has a Continuous Adjustable Resolution (CAR) feature that allows for a trade-off between image dimension and frames per second; from 1,200 fps at 1024-by-1024 pixels to 95,000 fps at 64-by-32 pixels. The camera was used in the 1024-by-1024 pixel, 1,200 fps mode for this test. The camera was placed in the path of the test vehicle at an offset distance of approximately 45 m and utilized a 400 mm focal length Nikon optic to acquire imagery of the vehicles that filled the field-of-view.

The ARRIFLEX D-21 was selected as the camera for the suspension response testing. The ARRIFLEX D-21 is a professional, cinematic-style digital camera that has a spinning, electronically adjustable reflex mirror shutter with a shutter angle that can be set between 11.2° and 180° and records at a resolution of 2880-by-1620 pixels at 30 fps in uncompressed 12-bit per pixel Bayer mode. The camera was used with a shutter angle of 11.2° for this collection to minimize motion blur of the moving vehicle in each individually recorded frame. The camera was placed approximately 10 m from the road, oriented perpendicular to the path of the vehicle to acquire imagery of the vehicle while it traverses the speed bump. The offset distance of 10 m was chosen to ensure that imagery of the entirety of the vehicle's suspension response was acquired.

The LMS400-2000 LiDAR system was selected as the instrument for the tire deformation testing. The LMS400-2000 is a one-dimensional scanning LiDAR system with a scanning angle of 70° , an angular resolution of 0.1° to 1.0° , a scanning frequency of 270 to 500 Hz, and an operating range between 0.7 to 3.0 m assuming a minimum 10 % reflectance target at 650 nm. The LiDAR system was placed approximately 0.5 m from the path of the test vehicle. The system was oriented so the one-dimensional scanner was operating vertically, capitalizing on the motion of the test vehicle to create a three-dimensional spatial map of the test vehicle as it passed.

4.3.6 Test Execution

This field test was carried out on December 10th and 11th in 2012 at the Savannah River Site in Aiken, South Carolina. The tractor trailer test vehicle was operated in the following three conditions, empty trailer, 20 tons loaded on the trailer, and 40 tons loaded on the trailer. The

Honda Accord test vehicle was operated in the following three conditions; no load (except for the driver and passenger) 3,340 pounds, 480 pounds loaded in the trunk for a total weight of 3,820 pounds, and 980 pounds loaded in the trunk for a total weight of 4,320 pounds. The frame twist, suspension response, and tire deformation tests were all completed during this period.

Frame Twist

The frame twist test was setup along a stretch of straight road at the testing site. The Vision Research Phantom v5.1 high-speed video camera system was set up so as to capture the test vehicles' acceleration from a stop, as shown in Figure 4.33. The high-speed camera was placed in the path of the vehicle at an offset distance of approximately 45 m and utilized a 210 mm focal length Nikon optic to achieve imagery of the vehicle that filled the field-of-view. The driver was instructed to accelerate from a stop towards the camera system as he typically would accelerate from a stop sign. The high-speed camera system began collecting data immediately prior to the test vehicle beginning its acceleration, and stopped collecting approximately two seconds after the test vehicle began its acceleration. Five runs were completed for each vehicle loading and for each test vehicle for a total of 30 runs.




Figure 4.33: The instrument configuration for the full-scale field test conducted at The Savannah River Site in December of 2012

Tire Deformation

The tire deformation test was run concurrently with the frame twist test. The frame twist testing monitored the initial acceleration of the test vehicle and the tire deformation testing monitored the test vehicle once it reached its desired velocity. The LMS400-2000 LiDAR system was setup on the edge of the road, oriented towards the path of the test vehicle and set to scan vertically. The LiDAR system was positioned less than 0.5 m from the designated path of the test vehicle to maximize the resolution of the collected data (while keeping the LiDAR system relatively safe from accidentally being run-over). The driver was instructed to travel past the LiDAR system at a constant velocity of approximately 3 to 5 mph. The LiDAR system collected data continuously for the duration of the each test run. Five runs were completed for each vehicle loading and for each test vehicle for a total of 30 runs.

Suspension Response

The suspension response test was run concurrently with the frame twist test and the tire deformation test. The frame twist testing monitored the initial acceleration of the test vehicle, the tire deformation testing monitored the test vehicle once it reached its desired velocity, and the suspension response test monitored the test vehicle as it continued down the stretch of road and traveled over a speed bump. The ARRIFLEX D-21 high-definition video camera system was setup off the road, oriented perpendicular to the path of the test vehicle. The high-definition camera system was positioned back from the road with the speed bump at the edge of its field-of-view to ensure all of the test vehicles' suspension response to the speed bump was captured. The driver was instructed to travel over the speed bump at a velocity he felt was safe for the fully-loaded vehicle. The high-definition camera system began collecting data immediately prior to the test vehicle entering its field-of-view, and stopped collecting after the test vehicle exited its field-of-view. Five runs were completed for each vehicle loading and for each test vehicle for a total of 30 runs.

4.4 Small-Scale Testing

To better understand the measurement requirements related to the tire deformation phenomenology, two small-scale experiments were conducted. The first experiment was conducted at Rochester Institute of Technology (RIT) in Henrietta, New York; the second experiment was conducted at the Wyoming County International Speedway in Perry, New York.

4.4.1 Test Vehicle

The 1992 Ford F600 flatbed dump-truck described in Section 4.1.1 was chosen as the test vehicle for both of these experiments.

4.4.2 Test Site

A parking lot on the campus of RIT was chosen as the site for the first small-scale test. This site was chosen for convenience and for easy access to students which were required for the experiment (as weight for the test vehicle). The Wyoming County International Speedway (described in

Section 4.1.2) was chosen as the site for the second small-scale test. This site was chosen because a closed course was required for this small-scale test. Additionally, at the Wyoming County International Speedway we were given access to weights that could remain on the test vehicle while it is moving (unlike students).

4.4.3 Test Plan

Both of the small-scale experiments were conducted to better understand the correlation between vehicle mass and the tire deformation phenomenology.

Tire Deformation I

For this test of the tire deformation phenomenology, an experiment was designed and executed to measure the tire deformation of a stationary test vehicle. The test was designed to measure the contact area and deformation of the stationary test vehicle's rear tires under different loading conditions. Foam-board sheets were placed under the rear tires of the test vehicle under various loading conditions and the contact area of the tires were mapped using spray paint. In addition, contact area and deformation measurements of the test vehicle's rear tires were made remotely under different loading conditions using a LiDAR scanner.

Tire Deformation II

For this test of the tire deformation phenomenology, an experiment was designed and executed to measure the tire deformation of both a static test vehicle and a moving test vehicle. The test was designed to measure the contact area and deformation of the test vehicle's tires under different loading conditions and tire inflation levels. Sheets of paper were placed under the tires of the test vehicle and the contact area of the tires was mapped using spray paint for various combinations of vehicle loading and tire inflation. In addition, contact area and deformation measurements of the test vehicle's tires were made remotely under different loading conditions using a LiDAR scanner as the test vehicle moved at a low velocity.

4.4.4 On-Board Sensor System

No on-board sensors were utilized for either of the small-scale tests.

4.4.5 Off-Board Sensor System

Both of the small-scale experiments utilized paper or foam-board and spray paint to map the contact area of the test vehicle's tires under different loading conditions. Both of the small-scale experiments also used LiDAR systems to remotely map the contact area and deformation of the test vehicle's tires.

Position Measurements

The LMS100 LiDAR system was selected as the instrument for the first small-scale test. The LMS100 is a one-dimensional scanning LiDAR system with a scanning angle of 270° , an angular resolution of 0.25° to 0.5° , a scanning frequency of 25 to 50 Hz, and a maximum operating distance of approximately 20 m assuming a minimum 10 % reflectance target at 650 nm. The LiDAR system was placed approximately 1.5 m from the stationary test vehicle and was oriented so the one-dimensional scanner was operating vertically as seen in Figure 4.34. A rotating platform was used to create a three-dimensional spatial map of the test vehicle.



Figure 4.34: The LiDAR system position and orientation used for the first small-scale experiment

The LMS400-2000 LiDAR system was selected as an instrument for the second small-scale test. The LMS400-2000 is a one-dimensional scanning LiDAR system with a scanning angle of 70° , an angular resolution of 0.1° to 1.0° , a scanning frequency of 270 to 500 Hz, and an operating range between 0.7 and 3.0 m assuming a minimum 10 % reflectance target at 650 nm. The LiDAR system was placed approximately 0.5 m from the path of the test vehicle. The system was oriented so the one-dimensional scanner was operating vertically, capitalizing on the motion of the test vehicle to create a three-dimensional spatial map of the test vehicle as it passed.

4.4.6 Test Execution

The first small-scale test took place on February 3rd in 2012 at a parking lot on the campus of The Rochester Institute of Technology. The second small-scale test took place on October 8th in 2012 at the Perry International Speedway. Both of these tests pertained to the tire deformation phenomenology.

Tire Deformation I

During this test, forty-one students and faculty members volunteered to support the experiment as load for the test vehicle. Each volunteer had their weight recorded and was assigned a number. The test was conducted by first positioning the rear tires of the unloaded test vehicle onto a sheet of foam-board as seen in Figure 4.35. Next, the volunteers climbed onto the test vehicle and positioned themselves uniformly on the flatbed as seen in Figure 4.36. Spray paint was then used to mark the contact area of the rear tire. This was repeated for the loading conditions seen in Table 4.2.



Figure 4.35: The rear tires of the test vehicle positioned on foam-board prior to measuring contact area



Figure 4.36: Several volunteers loaded on the flatbed of the test vehicle

Table 4.2: The test vehicle loading conditions for the small-scale test

Test Vehicle Loading	Total Vehicle Mass (lbs)
Unloaded	11,520
13 Volunteers	13,881
27 Volunteers	16,458
41 Volunteers	18,693

Tire Deformation II

This test was setup along a straightaway on the track at the Wyoming County International Speedway. The test vehicle was operated in four loading conditions and four tire inflation conditions as listed in Tables 4.3 and 4.4. The LMS400-2000 LiDAR system was positioned near the path of

the test vehicle oriented perpendicular to the path of the test vehicle, and set to scan vertically. For each tire inflation and loading condition, the test vehicle was driven past the LiDAR system at a constant velocity of approximately 3 to 5 mph. The LiDAR system collected data continuously for the duration of each test run. After the test vehicle had passed the LiDAR system, it was stopped with each tire positioned on a sheet of paper. Spray paint was used to mark the outline of the contact area of the test vehicle's tires.

Table 4.3: The test vehicle loading conditions

Test Vehicle Loading	Total Vehicle Mass (lbs)
Unloaded	11,520
1 Block	15,530
2 Blocks	19,540
3 Blocks	23,550 (Estimated)

Table 4.4: The test vehicle tire inflation pressures

% Inflation	Tire 1 (psi)	Tire 2 (psi)	Tire 3 (psi)	Tire 4 (psi)	Tire 5 (psi)	Tire 6 (psi)
100	100	95.5	89.0	88.0	88.5	89.0
90	91.5	85.5	79.0	78.5	79.5	79.5
80	81.5	74.5	68.0	70.5	69.0	69.0
70	66.5	61.5	59.0	55.0	60.0	60.0

Chapter 5

Experimental Results and Modeling

Upon the completion of each test described in Chapter 4 the collected data was processed and analyzed. The results of each test were used to better understand the validity of the phenomenologies as a method of estimating vehicle mass, and the lessons learned during each test were used to refine subsequent testing methodologies. In some cases, phenomenologies were deemed unsuitable and research into them was ceased to concentrate efforts on the more promising phenomenologies. Specifically, the suspension response and frame twist phenomenologies were selected as the best suited phenomenologies based on the data collected and the analysis conducted. It should be noted that the tire deformation phenomenology was found to be inconclusive, with more data needed to fully determine its suitability. Physics-based models were constructed for the suspension response and frame twist phenomenologies. The models were validated using the data collected during a final full-scale field test. The results of this field test showed that both phenomenologies could accurately estimate the mass of a vehicle using remote measurements, if accurate vehicle parameters were known. Specifically, the frame twist phenomenology requires the combined spring constants of the vehicle's suspension, and the position of the engine's axis of rotation relative to the suspension. The suspension response phenomenology requires the spring constants and damping coefficients of each of the vehicle's suspensions (front, rear, trailer), the unloaded mass of the vehicle, the unloaded moment of inertia of the vehicle, and the unloaded center of gravity of the vehicle. If all of these parameters are known with an uncertainty of 3 to 5 percent, the mass

estimation could have an uncertainty of approximately 10 percent for each method. As the uncertainty in the input parameters grow, so does the uncertainty in the mass estimation. A method was proposed to combine the results of the suspension response and frame twist phenomenologies into a single mass estimation with reduced uncertainty.

5.1 Brake Temperature

Experiments were conducted at the Wyoming County International Speedway and the Savannah River Site to better understand the brake temperature phenomenology. The experiments were designed to quantify the correlation between the temperature signature of a vehicle during a braking period and the mass of the vehicle and also to assess the ability of a remote sensing system to measure this signature.

5.1.1 Wyoming County International Speedway Full-Scale Field Test

During the full-scale field test conducted at the Wyoming County International Speedway on November 20th and 21st of 2010, data pertaining to the brake temperature phenomenology were collected using on-board and off-board instruments. These data were processed and analyzed to determine if correlation existed between brake temperature and vehicle mass, and if that phenomenology was observable from a remote sensing system. Due to a hard drive failure, no usable data was collected by the WASP system for this test.

During the brake temperature testing, the driver's side front brake (disc) and the passenger side rear brake (drum) of the test vehicle were monitored as the vehicle decelerated to a complete stop. This monitoring was done by two microbolometer array thermal imagers, multiple thermocouples, and RIT's WASP system mounted near the road. Three runs with the clutch engaged and three without the clutch engaged were conducted for each of the four vehicle loads for a total of twenty-four runs. An example of the on-board IR imagery data can be seen in Figures 5.1 and 5.2.




Figure 5.1: Gray-scale LWIR image of the front driver's side disk brake collected by the on-board imager and the tracking ROIs




Figure 5.2: Gray-scale LWIR image of the rear passenger's side drum brake collected by the on-board imager and the tracking ROIs

Data Processing

The two LWIR imagers mounted to the test vehicle stored each video frame as a radiometric audio video interleave (AVI) file, or .ravi. Custom code was written to ingest the .ravi movie files into MATLAB. Each video frame was read in as a 120-by-160 array of 16-bit signed integers. The integer values are in uncalibrated digital counts that should be linearly dependent on the sensor reaching radiance in the spectral range of 7.5 to 13 μm . Visual analysis of this data identified that reflected energy was significantly influencing the measured radiance. Figure 5.1 shows a radiance image of the front driver's side brake rotor, where the impact of reflected energy is clear. To avoid the influence of the reflected energy a basic tracking algorithm was written to extract the minimum radiance (in uncalibrated digital counts) of a region of the brake rotor in each frame. This should minimize the effect of reflected radiance if we assume that the rotor has a roughly uniform surface temperature, and emissivity. The tracking algorithm required the user to identify a region-of-interest (ROI) in the first image in the sequence (this ROI should be large and contain unique geometric features). Next, the user must identify a region within the ROI over which to take the minimum temperature (this ROI should be small, roughly uniform, and outside the influence of reflected radiance). Figures 5.1 and 5.2 show radiance images, each with two ROIs identified. Once the two ROIs have been identified, the location of the large ROI is found in each image frame by taking the two-dimensional normalized cross-correlation of the large ROI and each image frame. The two-dimensional normalized cross-correlation is analogous to taking the dot product of two normalized vectors in one dimension. The result of the two-dimensional normalized cross-correlation is an intensity array with a peak value at the location where the ROI and the image frame are most similar. If the ROI has unique features, and these features are present within the image (in roughly the same scale and orientation), this method will identify the location of the ROI within the image. The final step is locating the small ROI using its relative location within the large ROI and finding its minimum and maximum radiance. The minimum radiance should be used for analysis of the low emissivity brake rotor as it should be least impacted by reflected radiance, the maximum radiance should be used for analysis of the high emissivity (cavity-like) brake drum to minimize the influence of the surrounding, cooler, areas. This tracking algorithm was used to track the radiance of the brake rotor and the brake drum over the course of each test

run.

The thermocouple data did not require data processing. The output from the data acquisition module was calibrated temperature in degrees Celsius as a function of time.

Results

The measured radiance values in uncalibrated digital counts (DC) were extracted from the on-board IR video for the red ROIs shown in Figures 5.1 and 5.2 using the tracking algorithm discussed in the previous section. This was completed for all of the brake test runs (24 runs total). Without knowing the emissivity of the regions being observed, or the radiance incident upon them from other sources, it is not possible to convert the measured radiance to surface temperature. However, in the spectral region in which the IR cameras operate we can expect the radiance to have an approximately linear correlation to surface temperature. To see how well the thermocouple measurements track with the radiance measured from the ROIs, the measurements were compared for identical periods. Figure 5.3 shows the measured radiance and the measured temperature for the same time period plotted on different scales. The measured radiance was for a region on the brake rotor, and the measured temperature was for the back of the corresponding brake pad. The comparison clearly shows that the self-emission of the rotor surface, changes much more rapidly than the temperature of the thermocouple. This is a result of the low thermal inertia of the brake rotor surface and the high thermal inertia of the braking system. As a result, the temperature measurements change much more slowly than the radiance measurements. This means that the thermocouple measurements are not well suited to validate the radiance measurements.




Figure 5.3: Comparison of on-board temperature and radiance measurements.

Despite our inability to derive surface temperature from the self-emitted radiance of the test vehicle's brakes, the radiance data was analyzed for correlation to vehicle mass. Each radiance signal was plotted versus time and analyzed. The disc brake radiance signals were decomposed into two values, which described the curve; the difference between peak radiance and initial radiance and the difference between final radiance and initial radiance. An example of this can be seen in Figure 5.4. The drum brake temperature signals were decomposed into a single value that described the curve; the difference between the final radiance and initial radiance. An example of this can be seen in Figure 5.5. For both measurements, the final radiance was chosen to be 3 seconds after the vehicle came to a complete stop.




Figure 5.4: The characteristic values of the driver side front disc radiance signal in digital counts. Note that the final radiance was chosen to be 3 seconds after the vehicle was stopped.




Figure 5.5: The characteristic values of the passenger side rear drum radiance signal in digital counts. Note that the final radiance was chosen to be 3 seconds after the vehicle was stopped.

The velocity of the test vehicle was measured using a GPS module throughout each test. The test vehicle's velocity immediately prior to the start of braking was used to calculate the kinetic energy of the vehicle using the equation we described in Section 3.2 (rotational kinetic energy was disregarded for this analysis). The results can be seen in Tables 5.1 and 5.2. Due to a GPS malfunction, velocity data was unavailable for some test runs. For the remaining data, the normalized correlation coefficient was calculated to quantify the correlation between each signal parameter and the kinetic energy of the test vehicle at the onset of braking and can be seen in Table 5.3.

Table 5.1: Results from the Brake Test with transmission disengaged; where Peak is the peak radiance, Peak - Start is the difference between the peak radiance and the initial radiance, Final is the radiance 3 seconds after the vehicle stops, Final - Start is the difference between the radiance 3 seconds after the vehicle stops and the initial radiance, and Min is the minimum radiance

Speed (m/s)	Kinetic Energy (kJ)	Braking Rate (kJ/s)	Disc Brake					Drum Brake		
			Peak	Start	Peak - Start	Final	Final - Start	Final	Min	Final - Start
15.07	593.61	118.72	167	-210	377	14	224	-166	-240	74
16.82	739.37	147.87	550	-30	580	325	355	7	-110	117
12.55	411.67	68.61	430	90	340	275	185	80	30	50
13.17	610.89	101.82	153	-218	371	-42	176	-145	-232	87
14.04	694.72	138.94	519	-60	579	241	301	-40	-106	66
13.58	649.67	129.93	560	17	543	385	368	51	-7	58
12.40	681.19	136.24	400	-60	460	343	403			

Table 5.2: Results from the Brake Test with transmission engaged; where Peak is the peak radiance, Peak - Start is the difference between the peak radiance and the initial radiance, Final is the radiance 3 seconds after the vehicle stops, Final - Start is the difference between the radiance 3 seconds after the vehicle stops and the initial radiance, and Min is the minimum radiance

Speed (m/s)	Kinetic Energy (kJ)	Braking Rate (kJ/s)	Disc Brake					Drum Brake		
			Peak	Start	Peak - Start	Final	Final - Start	Final	Min	Final - Start
13.07	446.10	74.35	445	95	350	270	175	77	-11	88
14.35	538.24	76.89						97	33	64
16.51	712.48	142.50	960	208	752	726	518	194	-14	208
13.58	649.67	92.81	606	82	524	505	423	112	26	86
13.79	669.50	133.90	723	172	551	513	341	138	65	73
13.68	659.55	131.91	835	219	616	636	417	228	119	109
13.22	774.65	154.93	730	221	509	704	483	191	26	165
13.43	798.95	133.16	798	246	552	786	540	241	129	112
13.74	836.10	167.22	1054	399	655	996	597	334	182	152
12.04	773.66	110.52	286	-80	366	254	334	-48	-136	88
12.45	827.47	165.49	494	-21	515	462	483	99	-86	185
12.60	848.11	141.35	648	101	547	617	516	128	28	100

Table 5.3: Results from the normalized correlation coefficient calculations that show the correlation between each signal parameter and the kinetic energy of the test vehicle

	Normalized Correlation Coefficient of Parameter to Vehicle Mass	
Parameters	Kinetic Energy (Test A)	Kinetic Energy (Test B)
Peak - Start (Disc)	0.81	0.35
Final - Start (Disc)	0.72	0.83
Final - Start (Drum)	0.65	0.51

Conclusions

The effect of the vehicle's transmission on the thermal energy generated during braking was determined to be not significant. The measured differences in self-emitted radiance for the braking system between the transmission engaged (Test B) and the transmission disengaged (Test A) test runs fell within the noise of the measurements.

The correlation analysis did not show a strong correlation between the vehicle's kinetic energy and any signature associated with the emitted radiance of the rear passenger's side drum brake. Analysis of the radiance measurements show a much slower increase in emitted radiance than the disc brake during a period of braking. The disc brake's emitted radiance increases, peaks and begins to decrease over the period of measurement while the drum brake's emitted radiance increases, begins to plateau, but does not peak or begin to decrease over the period of measurement. This difference is likely caused by differences in the surfaces being observed; for the disc brake, the actual braking surface is being measured but for the drum brake, the actual braking surface is obscured.

The correlation analysis showed the strongest correlation between the vehicle's kinetic energy and the difference between the final radiance and the initial radiance for the front disk brakes. This was true for testing with the transmission engaged and with the transmission disengaged. The difference between the final radiance and initial radiance emitted by the disc brake does not appear to be affected by the braking method (transmission engaged, or transmission disengaged).


Due to a malfunction of RIT's WASP sensor, we were unable to determine if any of the brake temperature signatures measured could be observed from an off-board system.

5.1.2 Savannah River Site Full-Scale Field Test I


During the full-scale field test conducted at the Savannah River Site on May 24th and 25th of 2011, data pertaining to the brake temperature phenomenology was collected using on-board and off-board instruments. The data was processed and analyzed to determine if there was a correlation between each phenomenology and vehicle mass, and if that phenomenology was observable from a remote sensing system.

During the brake temperature testing, the driver side tractor drive wheels, and the passenger

side front wheel were monitored with two microbolometer array thermal imagers. An example of the on-board IR imagery data can be seen in Figure 5.6. RIT's WASP system was mounted near the road and monitored the test vehicle during the braking period. An example of the WASP imagery can be seen in Figure 5.7.



(a) Passenger side front tractor brake



(b) Driver side rear tractor brake

Figure 5.6: Gray-scale LWIR images of the tractor's brakes collected by the on-board imager




Figure 5.7: Gray-scale LWIR image of the test vehicle immediately upon coming to a complete stop collected with WASP

Data Processing


The two LWIR imagers mounted to the test vehicle stored each video frame as a radiometric audio video interleave (AVI) file, or .ravi. Custom code was written to ingest the .ravi movie files into MATLAB. Each video frame was read in as a 120-by-160 array of 16-bit signed integers. The integer values are in uncalibrated digital counts that should be linearly dependent on the sensor reaching radiance in the spectral range of 7.5 to 13 μm . Visual analysis of the data did not show any significant influence by reflected energy as was seen in some of the previous testing (of disk brakes). This was expected since the test vehicle was equipped with drum-style brakes only, and the reflectance issue was only observed on the highly reflective surface of the rotor of disk-style brakes. The visual analysis did show that the visible region of the drum brake that heats up was fairly small, only about 10 pixels, as seen in Figure 5.8. Since the area-of-interest was the hottest in its general area, tracking its location only required a simple algorithm that tracked the location of the pixel with the highest intensity within a user-defined area. The areas-of-interest selected for each camera can be seen in Figure 5.8. This tracking algorithm was used to track the maximum radiance of the test vehicle's brakes as a function of time for each test run.

The WASP sensor located near the road, stored each video frame as an ENVI image file or

.img. Only LWIR data was collected during this testing. Custom processing code was used to apply a flat-field correction and the data was saved as 32-bit double-precision floating-point TIFF images. Each video frame was stored as a 512-by-640 array of double-precision floating-point values. The values are in calibrated sensor-reaching radiance ($W/(m^2 sr)$) in the spectral range of 8.0 to 9.2 μm . Visual analysis of the data showed that only one of the test vehicle's brakes was consistently visible during the collection period. A basic tracking algorithm was written to extract the maximum radiance of a user-selected region of the brake in each frame. The tracking algorithm required the user to identify a region-of-interest (ROI) in the first image in the sequence. The region of interest selected was the passenger-side rear tractor drive wheel as seen in Figure 5.9. The location of the ROI is found in each image frame by taking the two-dimensional normalized cross-correlation of the large ROI and each image frame. The two-dimensional normalized cross-correlation is analogous to taking the dot product of two normalized vectors in one dimension. The result of the two-dimensional normalized cross-correlation is an intensity array with a peak value at the location where the ROI and the image frame are most similar. If the ROI has unique features, and these features are present within the image (in roughly the same scale and orientation), this method will identify the location of the ROI within the image. The final step is finding the maximum radiance within the tracked ROI. This tracking algorithm was used to track the radiance of the brake drum of the passenger side rear tractor drive wheel over the course of each test run.



(a) Passenger side front tractor brake with two areas-of-interest highlighted



(b) Driver side rear tractor brake with one area-of-interest highlighted

Figure 5.8: Location of the areas-of-interest used to track the location and radiance of the hottest regions of the test vehicle's brakes.



Figure 5.9: Location of the area-of-interest used in the WASP imagery to track the location and radiance of the hottest region of the test vehicle's brake.

Results

The measured radiance values in uncalibrated digital counts (DC) were extracted from the on-board IR video for the ROIs shown in Figure 5.8 using the tracking algorithm discussed in the previous section. Without knowing the emissivity of the regions being observed, or the radiance incident upon them from other sources, it is not possible to convert the measured radiance to surface temperature. However, in the spectral region in which the IR cameras operate we can expect the radiance to have an approximately linear correlation to surface temperature. An example of the measured uncalibrated radiance data can be seen in Figure 5.10.




Figure 5.10: The characteristic values of the passenger side front brake radiance signal in digital counts as measured by the on-board IR sensor.

The measured radiance values in uncalibrated radiance were extracted from the off-board IR video for the ROI shown in Figure 5.9 using the tracking algorithm discussed in the previous section. This was completed for all of the brake test runs (24 runs total). An example of the measured uncalibrated radiance data can be seen in Figure 5.11.




Figure 5.11: The characteristic values of the passenger side front brake radiance signal in digital counts as measured by the off-board IR sensor (WASP)

Conclusions

The data showed that only a small area of the test vehicle's brakes could be observed remotely, and only briefly for a given orientation. Only one of the test vehicle's brakes was visible by the off-board sensor, and only during the final few seconds of the braking process. Additionally, the on-board data showed a very large range of observed radiance values. This range is due to thermal energy stored in the brake material from previous instances of braking, as the initial radiance increased with each test run. This suggests that the driving habits immediately (and perhaps not so immediately) preceding an instance of braking have a very large impact on the initial temperature of the braking system. The analysis of the radiance measurements made with the on-board and off-board sensors led to the conclusion that the remotely-sensed brake temperature would not be a

viable method of estimating vehicle mass.

5.2 Engine Acoustics

Experiments were conducted at the Wyoming County International Speedway and the Savannah River Site to better understand the engine acoustics phenomenology. The experiments were designed to quantify the correlation between the engine acoustic signature of a vehicle during an acceleration period and the mass of the vehicle, and also to assess the ability of a remote sensing system to measure this signature.

5.2.1 Wyoming County International Speedway Full-Scale Field Test

During the engine acoustic testing, the acoustic emission of the test vehicle was monitored as the vehicle accelerated from a stopped position. This monitoring was done by two on-board calibrated PCB Piezotronics free-field condenser microphones (one mounted in the engine compartment and one mounted near the exhaust), and two off-board PCB Piezotronics free-field condenser microphones mounted near the road. Six test runs were conducted for each of the four vehicle loads for a total of twenty-four runs.

Data Processing

The two PCB Piezotronics free-field condenser microphones mounted on the test vehicle monitored the acoustic emission of the test vehicle inside the engine compartment and near the exhaust pipe. The two PCB Piezotronics free-field condenser microphones mounted near the road measured the total acoustic emission of the test vehicle as observed by a stationary sensor. The data was collected at a rate of 51.2 kHz and stored in a binary proprietary format (National Instruments). The data was ingested using National Instruments DIAdem software and exported to a universal format. The data was then ingested by MATLAB, and low pass filtered to remove the effects of frequencies higher than 300 Hz. This filtering was done to reduce noise in the data signals. Next, a spectrogram was calculated for each test run from the filtered data. A spectrogram, or spectral waterfall plot, is a plot that shows the spectral energy of a signal as it varies with time. A

Hamming window of approximately 0.75 seconds was used as the spectrogram window to reduce the effects of aliasing. An example of spectrograms calculated from the on-board microphones over identical time periods can be seen in Figure 5.12. The spectrograms calculated from the data collected by the on-board microphone located inside the engine compartment showed a much lower signal-to-noise ratio than the data collected by the on-board microphone located near the exhaust. The decreased signal-to-noise ratio is likely caused by the reflected energy (or echoes) resulting from being in an enclosed area. Due to the lowered signal-to-noise ratio, the data collected by the on-board microphone located in the engine compartment was disregarded. This left three spectrogram plots for each test run calculated from the on-board microphone located near the output of the exhaust pipe, and the two off-board microphones located near the road.

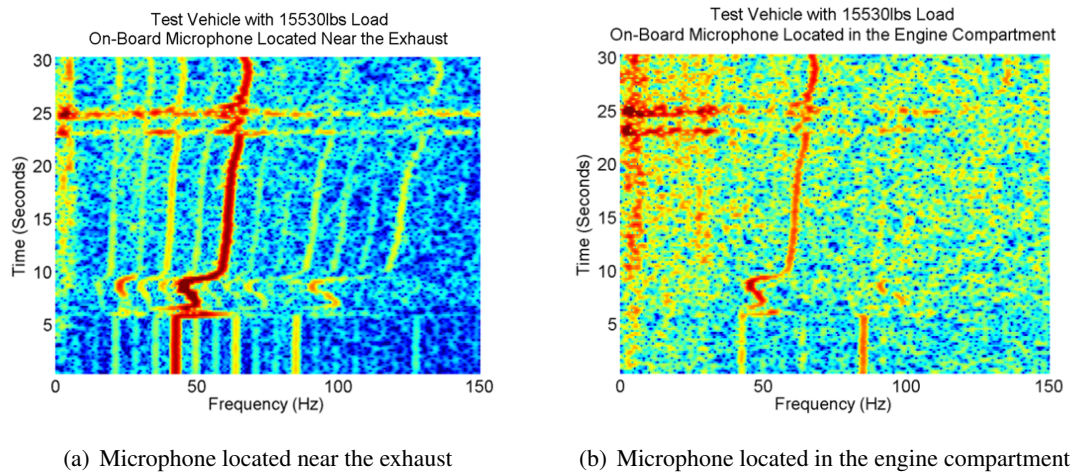


Figure 5.12: Spectrogram from the on-board microphones for the test vehicle loaded with one block accelerating from a stop

Results

The spectrograms calculated from the on-board data were analyzed to determine how the engine operating frequency manifested in the spectrogram. The results suggest that the first harmonic of the signal is half of the operating frequency of the engine. For example, an engine operating at 1000 RPM (or 16.67 Hz) would emit an acoustic signal with a first harmonic of 500 RPM (or 8.33

Hz). This is consistent with the operation of a 4-stroke engine because it takes two full revolutions for all of the engine's cylinders to fire. The first harmonic of a signal can be calculated from the spectral energy plot (or spectrogram) by measuring the distance between adjacent harmonics. Figure 5.13 shows an example of how the fundamental frequency is calculated and identifies the measured harmonics.




Figure 5.13: A spectrogram from the on-board microphone located near the exhaust of the test vehicle, collected while idling. The top line illustrates how the fundamental frequency was calculated from the average distance between local maxima (harmonics) in the frequency domain. The bottom 9 lines show the location of the first nine harmonics.

The results show that the first harmonic is difficult to directly identify from the spectrogram, but the subsequent harmonics can be easily observed. Since each consecutive harmonic is separated by a frequency equal to the first harmonic, it is possible to calculate the first harmonic simply by measuring the average separation between adjacent harmonics. In Figure 5.13 it takes 28.5 Hz to span four harmonics, which averages to 7.13 Hz per harmonic. This means the first harmonic is located at 7.13 Hz, which means the engine was idling at approximately twice that, 14.26 Hz or 855.6 RPM. Further analysis of the test data showed that the harmonic with the most energy was consistently the sixth, and the test vehicle has six cylinders. This relationship between the number of cylinders and the harmonic with the most energy makes sense, as the most significant component of the first harmonic would be the combustion events, which occur at almost equal spacing. The fact that the combustion events are not exactly evenly spaced is likely the reason why the sixth harmonic appears wider than the other in the spectrogram.

Comparison between the data collected by the on-board microphone (located near the output port of the exhaust) and the off-board microphones shows a significant reduction in signal for the off-board microphones (Figure 5.14). The increased distance between the off-board microphone and test vehicle has caused the magnitude of the observed engine signal to be reduced. Additionally, the data collected by the off-board microphone shows four strong vertical lines, which were caused by being in close proximity to the WASP sensor, and the generator used to power it.

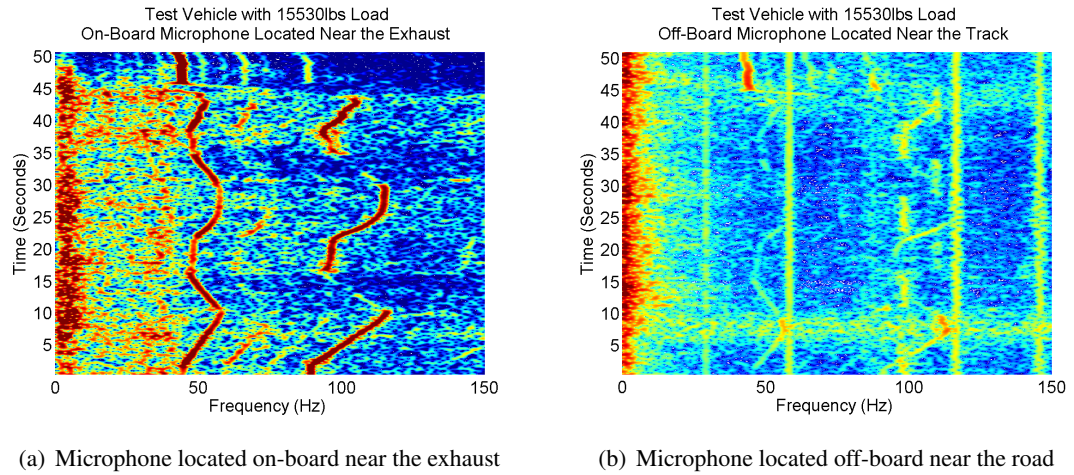


Figure 5.14: Spectrogram from the on-board and off-board microphones for the test vehicle loaded with two blocks accelerating from a stop

Conclusions

From analysis of the data collected by the on-board microphone located near the output of the exhaust system it appears that it may be possible to extract the operating frequency of an engine and the number of cylinders it has through acoustic analysis. However, even with the data collected by the on-board sensor, there are periods of time where only one or two of the engine harmonics are discernible. To measure the operating frequency of the engine, at least two adjacent harmonics must be able to be identified (assuming our analysis is correct).

5.2.2 Savannah River Site Full-Scale Field Test I

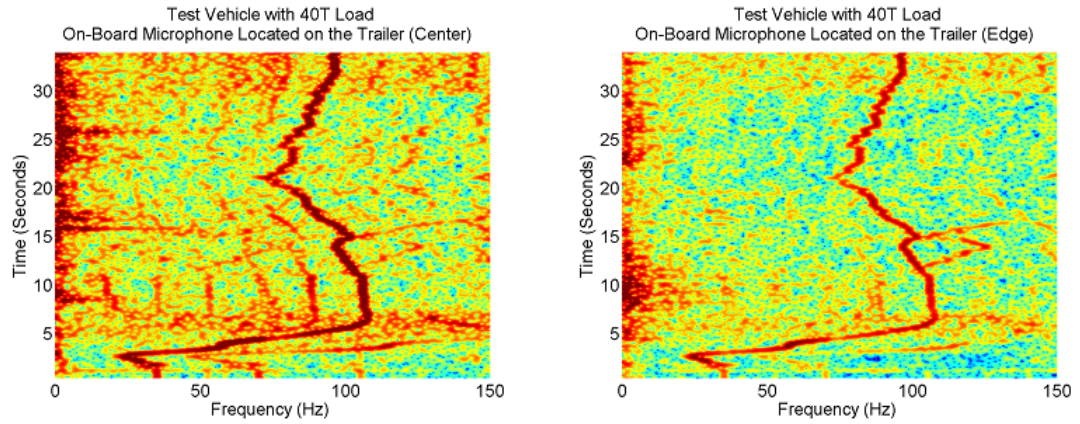
During the engine acoustic testing, the acoustic emission of the test vehicle was monitored as the vehicle accelerated from a stopped position. This monitoring was done by two on-board calibrated PCB Piezotronics free-field condenser microphones mounted on the tractor end of the trailer and orientated towards the output of the exhaust (one mounted near the center of the trailer, the other towards the edge), and two off-board PCB Piezotronics free-field condenser microphones mounted near the road. Ten test runs were conducted for each of the three vehicle loads for a total of thirty runs.

Data Processing

The two PCB Piezotronics free-field condenser microphones mounted on the test vehicle monitored the acoustic emission of the test vehicle near the exhaust pipe. The two PCB Piezotronics free-field condenser microphones mounted near the road measured the total acoustic emission of the test vehicle as observed by a stationary sensor. The data was collected at a rate of 51.2 kHz and stored in a binary proprietary format (National Instruments). The data was ingested using National Instruments DIAdem software and exported to a universal format. The data was then ingested by MATLAB, and low pass filtered to remove the effects of frequencies higher than 300 Hz. This filtering was done to reduce noise in the data signals. Next, a spectrogram was calculated for each test run from the filtered data. A spectrogram, or spectral waterfall plot, is a plot that shows the spectral energy of a signal as it varies with time. A Hamming window of approximately 0.75 seconds was used as the spectrogram window to reduce the effects of aliasing. An example of spectrograms calculated from the on-board microphones over identical time periods can be seen in Figure 5.15. The spectrograms calculated from the data collected by the on-board microphone located near the edge of the trailer showed a lower signal-to-noise ratio than the data collected by the on-board microphone located near the center of the trailer. The decreased signal-to-noise ratio for the microphone located near the edge of the trailer is likely caused by air traveling past as the vehicle travels, the microphone located near the center of the trailer is less affected because it is shielded by the tractor. One of the off-board microphones malfunctioned and was replaced with the on-board microphone located on the edge of the trailer. As a result, three acoustic measurements were collected for each test run, and spectrograms were generated for each measurement.

Results

The spectrograms calculated from the on-board data were analyzed to determine how the engine operating frequency manifested in the spectrogram. Previous testing suggested that the first harmonic of the signal is half of the operating frequency of the engine, and that the number of cylinders the engine had determines which harmonic has the highest power (the sixth harmonic will have the highest power for a six cylinder engine). The on-board data exhibited the expected harmonics as seen in Figure 5.16.



(a) Microphone located near the center of the trailer oriented towards the output of the exhaust system (b) Microphone located near the edge of the trailer oriented towards the output of the exhaust system

Figure 5.15: Spectrogram from the on-board microphones for the test vehicle loaded with one block accelerating from a stop




Figure 5.16: Location of the fifth, sixth and seventh harmonic of the test vehicle. The harmonics took place at 29.2, 35.1 and 40.9 Hz, respectively, meaning that the first harmonic occurs at 5.9 Hz and the test vehicle's engine is operating at 700 RPM.

The frequency with the highest power was tracked as a function of time and converted to the operating frequency of the engine using the following equation (for a 4-stroke engine)

$$\omega_{Engine} = \frac{2 \cdot \omega_{Peak}}{N_{Cylinders}} \quad (5.1)$$

where ω_{Engine} is the operating frequency of the engine (Hz), ω_{Peak} is the measured frequency with the highest power (Hz), and $N_{Cylinders}$ is the number of cylinders in the engine (unitless). Four-stroke engines require two full revolutions for all of the engine's cylinders to fire. This is accounted for by the factor of two in Equation 5.1. Figure 5.17 shows the engine RPM, calculated by tracking the frequency with the highest power from Figure 5.16.


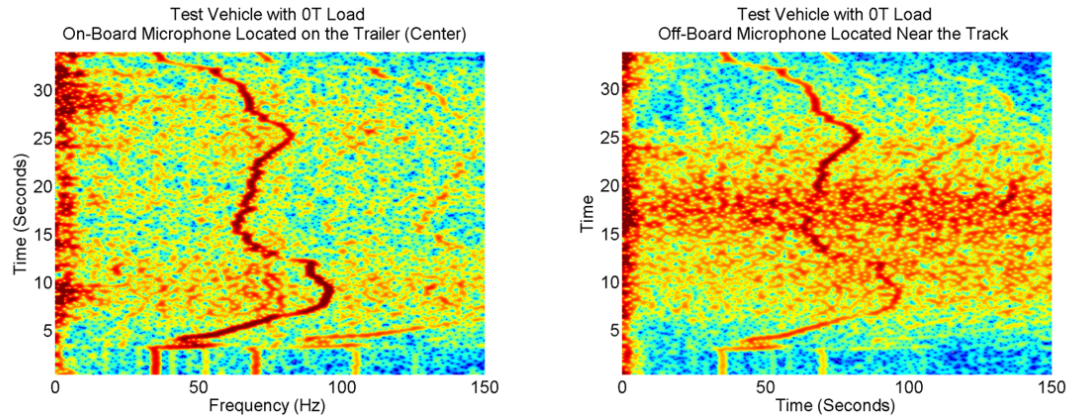


Figure 5.17: The calculated engine RPM for the test vehicle as a function of time

Comparison between the data collected by the on-board microphone (located near the center of the trailer) and the off-board microphone shows a significant reduction in signal for the off-board microphones (Figure 5.18). The increased distance between the off-board microphone and test vehicle has caused the magnitude of the observed engine signal to be reduced. However, it appears that the pertinent signals are still observable.



(a) Microphone located on-board near the center of the trailer

(b) Microphone located off-board near the road

Figure 5.18: Spectrogram from the on-board and off-board microphones for the unloaded test vehicle accelerating from a stop

Conclusions

The analysis of the data collected during this field test agreed well with the data collected during the previous field test. Extracting the operating frequency and number of cylinders of an engine through acoustic analysis does appear to be possible. However, to measure the operating frequency of the engine, at least two adjacent harmonics must be able to be identified. From the data collected to date, that does not always appear possible due to signal noise. Noise mitigation may be achieved using directional microphones, to increase signal and reduce background noise, or through post processing. Additional testing is required to build confidence in our assumption that the harmonic with the highest power correlates to the number of cylinders in the engine (all testing so far has been on six cylinder engines).

5.3 Exhaust Braking

An experiment was conducted at the Savannah River Site to better understand the exhaust braking acoustic phenomenology. The experiment was designed to quantify the correlation between the exhaust braking acoustic signature of a vehicle during an exhaust braking period and the mass of the vehicle, and also to assess the ability of a remote sensing system to measure this signature.

5.3.1 Savannah River Site Full-Scale Field Test I

During the exhaust braking acoustic testing, the acoustic emission of the test vehicle was monitored as the vehicle decelerated with and without an exhaust brake. This monitoring was done by two on-board calibrated PCB Piezotronics free-field condenser microphones mounted on the tractor end of the trailer and orientated towards the output of the exhaust (one mounted near the center of the trailer, the other towards the edge). No off-board microphones were used during this testing. Eight test runs were conducted; four with the exhaust brake activated, and four without.

Data Processing

The two PCB Piezotronics free-field condenser microphones mounted on the test vehicle monitored the acoustic emission of the test vehicle near the exhaust pipe. The two PCB Piezotronics free-field condenser microphones mounted near the road measured the total acoustic emission of the test vehicle as observed by a stationary sensor. The data was collected at a rate of 51.2 kHz and stored in a binary proprietary format (National Instruments). The data was ingested using National Instruments DIAdem software and exported to a universal format. The data was then ingested by MATLAB, and low pass filtered to remove the effects of frequencies higher than 300 Hz. This filtering was done to reduce noise in the data signals. Next, a spectrogram was calculated for each test run from the filtered data. A spectrogram, or spectral waterfall plot, is a plot that shows the spectral energy of a signal as it varies with time. A Hamming window of approximately 0.75 seconds was used as the spectrogram window to reduce the effects of aliasing. Only the data from the on-board microphone located near the center of the trailer and oriented towards the output of the exhaust system was used.

Results

The spectrograms were analyzed to determine how the acoustic signal changed when the exhaust brake was applied. The acoustic signal of the test vehicle changes significantly when the exhaust brake is applied as seen in Figure 5.19. The results clearly show a difference between braking methods. For the spectrogram generated from the measurements taken during the standard braking (no exhaust brake), only the sixth, twelfth, and eighteenth harmonic of the engine are clearly

visible, and even those are reduced in power during the braking period. For the spectrogram generated from the measurements taken during the exhaust braking, nearly all of the harmonics show an increase in power. This observation agrees well with the known operation of the vehicle and exhaust braking. When the driver takes his foot off of the accelerator, the amount of fuel being injected into the cylinders is reduced, resulting in less energetic (or no) combustion. This explains why the harmonics appear to have less power during the standard braking measurements. When a Jacobs-style exhaust brake (the type installed on the test vehicle) is activated, no combustion takes place at all, but instead a valve is opened to release all of the high pressure gas within the cylinder. This release of high pressure gas occurs at the same frequency the combustion would occur at relative to the engine operating frequency and generates significant acoustic energy.

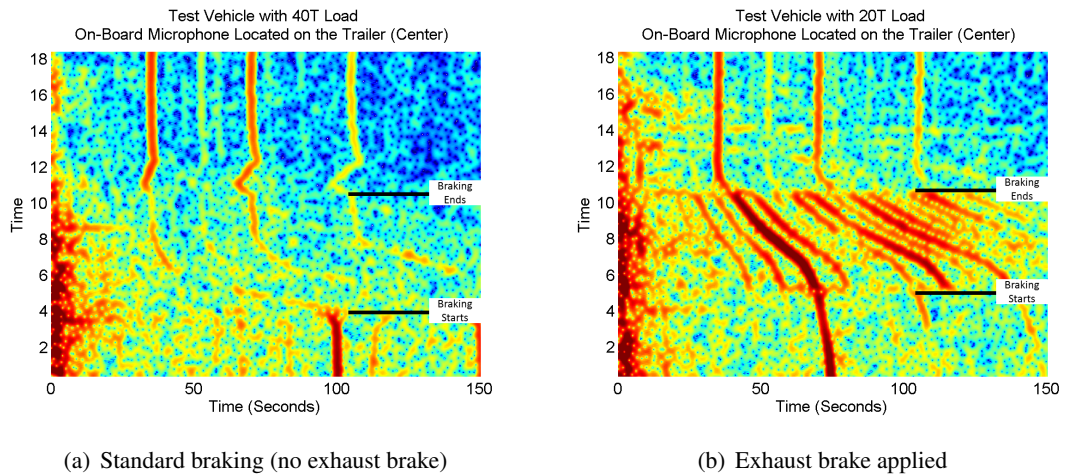


Figure 5.19: Spectrogram from the test vehicle decelerating to a stop with and without the exhaust brake

Conclusions

The analysis of the data collected during this field test showed that it is possible to determine between standard braking, and Jacobs-style exhaust braking using acoustic signals. This result is not particularly surprising, as the human ear can easily distinguish between standard braking Jacobs-style exhaust braking due to its unique staccato sound. The spectral signature of the Jacobs-style

exhaust brake showed nearly identical spectral content to the normal vehicle operation (except with more energy), which was surprising given the very different perception by the human ear. These results suggest that determining if a Jacobs-style exhaust brake is being applied may be possible using an automated system. However, off-board measurements are needed to build confidence.

5.4 Tire Temperature

An experiment was conducted at the Savannah River Site to better understand the tire temperature phenomenology. The experiment was designed to quantify the correlation between the tire temperature signature of a vehicle during operation and the mass of the vehicle, and also to assess the ability of a remote sensing system to measure this signature.

5.4.1 Savannah River Site Full-Scale Field Test I

During the tire temperature testing, the driver side tractor drive wheels, and the passenger side front wheel were monitored with two microbolometer array thermal imagers. RIT's WASP system was mounted near the road and monitored the test vehicle during the braking period. Since the main focus of this experiment is to monitor the temperature of the tires when the test vehicle is operating a roughly constant (non-zero) velocity, the WASP data will only be used to determine how well the tire radiance can be measured from an off-board location. An on-board GPS module monitored the test vehicle's velocity as a function of time for the duration of each test run (no data processing was required).

Data Processing

The two LWIR imagers mounted to the test vehicle stored each video frame as a radiometric audio video interleave (AVI) file, or .ravi. Custom code was written to ingest the .ravi movie files into MATLAB. Each video frame was read in as a 120-by-160 array of 16-bit signed integers. The integer values are in uncalibrated digital counts that should be linearly dependent on the sensor reaching radiance in the spectral range of 7.5 to 13 μm . A tracking algorithm was written to extract the minimum, maximum, and median radiance (in uncalibrated digital counts) of a region on the

tread surface of the tire in each frame. The tracking algorithm required the user to identify a region-of-interest (ROI) in the first image in the sequence (this ROI should be large and contain unique geometric features). Next, the user must identify a region within the ROI over which to record the minimum, maximum, and median radiance (this ROI should be small and roughly uniform). Figure 5.20 shows the radiance images, each with two ROIs identified. Once the two ROIs have been identified, the location of the large ROI is found in each image frame by taking the two-dimensional normalized cross-correlation of the large ROI and each image frame. The result of the two-dimensional normalized cross-correlation is an intensity array with a peak value at the location where the ROI and the image frame are most similar. If the ROI has unique features, and these features are present within the image (in roughly the same scale and orientation), this method will identify the location of the ROI within the image. The final step is locating the small ROI using its relative location within the large ROI and finding its minimum, maximum, and median radiance. The minimum, maximum, and median radiance signals (Figure 5.21) were analyzed and it was determined that the median radiance signal consistently had the best signal-to-noise ratio. The median radiance signal was then smoothed using a boxcar filter that averaged 9 samples (1 second). This tracking algorithm was used to track the radiance of the two tires (passenger side front tractor tire, and driver side rear tractor tire) over the course of each test run.

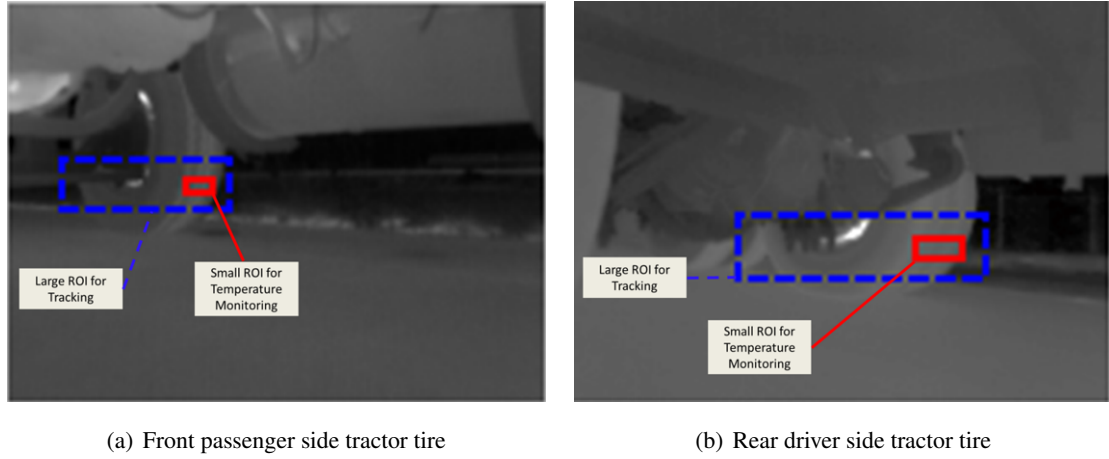


Figure 5.20: Gray-scale LWIR images collected by the on-board imagers and the corresponding tracking ROIs




Figure 5.21: Minimum, maximum, and median tire radiance for a fully loaded (40 ton) test vehicle.

The WASP sensor located near the road, stored each video frame as an ENVI image file or .img. Only LWIR data was collected during this testing. Custom processing code was used to apply a flat-field correction and the data was saved as 32-bit double-precision floating-point TIFF images. Each video frame was stored as a 512-by-640 array of double-precision floating-point values. The values are in calibrated sensor-reaching radiance ($\text{W}/(\text{m}^2 \text{sr})$) in the spectral range of 8.0 to 9.2 μm . Visual analysis of the data showed that only one of the test vehicle's tires was consistently visible during the collection period. A basic tracking algorithm was written to extract the maximum radiance of a user-selected region of the tire in each frame. The tracking algorithm required the user to identify a region-of-interest (ROI) in the first image in the sequence. The region of interest selected was the driver-side rear tractor drive wheel as seen in Figure 5.22. The location of the ROI is found in each image frame by taking the two-dimensional normalized

cross-correlation of the large ROI and each image frame. The two-dimensional normalized cross-correlation is analogous to taking the dot product of two normalized vectors in one dimension. The result of the two-dimensional normalized cross-correlation is an intensity array with a peak value at the location where the ROI and the image frame are most similar. If the ROI has unique features, and these features are present within the image (in roughly the same scale and orientation), this method will identify the location of the ROI within the image. The final step is finding the maximum radiance within the tracked ROI. This tracking algorithm was used to track the radiance of the tire tread of the driver-side rear tractor drive wheel over the course of each test run.



Figure 5.22: Location of the area-of-interest used in the WASP imagery to track the location and calibrated radiance of the hottest region of the test vehicle's tire tread.

The measured radiance values ($\text{W}/(\text{m}^2 \text{sr})$) were extracted from the off-board IR video for the ROI shown in Figure 5.22 using the tracking algorithm discussed in the previous section. This was completed for all of the brake test runs (24 runs total). An example of the measured radiance data can be seen in Figure 5.23.




Figure 5.23: The characteristic values of the diver side front tire tread radiance signal in digital counts as measured by the off-board IR sensor (WASP).

Results

The measured radiance values in uncalibrated digital counts (DC) were extracted from the on-board IR video for the ROIs shown in Figure 5.20 using the tracking algorithm discussed in the previous section. The brake test runs were chosen because they were the only tests that where the test vehicle was operated at typical real-world speeds (30 - 35 mph). Without knowing the emissivity of the regions being observed, or the radiance incident upon them from other sources, it is not possible to convert the measured radiance to surface temperature. However, in the spectral region in which the IR cameras operate we can expect the radiance to have an approximately linear correlation to surface temperature. The velocity of the test vehicle (measured with the on-board GPS module) was plotted on the same chart to better identify any trends between tire tread

temperature and vehicle velocity. An example of the measured uncalibrated radiance data for each vehicle load can be seen in Figure 5.24. The plot shows the measured uncalibrated radiance for each tire tread (left axis) and the velocity of the test vehicle (right axis). The results show a very large range of starting radiance values, likely caused by differences in road temperature. In Figure 5.24; the 40 ton test run took place on May 24th, 2011 at 10:53AM, the 20 ton test run took place on May 25th, 2011 at 8:38AM, and the 0 ton test run took place on May 25th, 2011 at 11:06AM. The temperature of the road surface increased sharply during the morning hours during this collection, and this change in temperature most likely caused the large temperature range observed in Figure 5.24. In addition, visual analysis of the on-board IR video showed that the tire tread surface and the road surface always had very nearly the same emitted radiance. Further analysis of the data could not discern any trends linking vehicle load, vehicle velocity, and tire tread temperature.

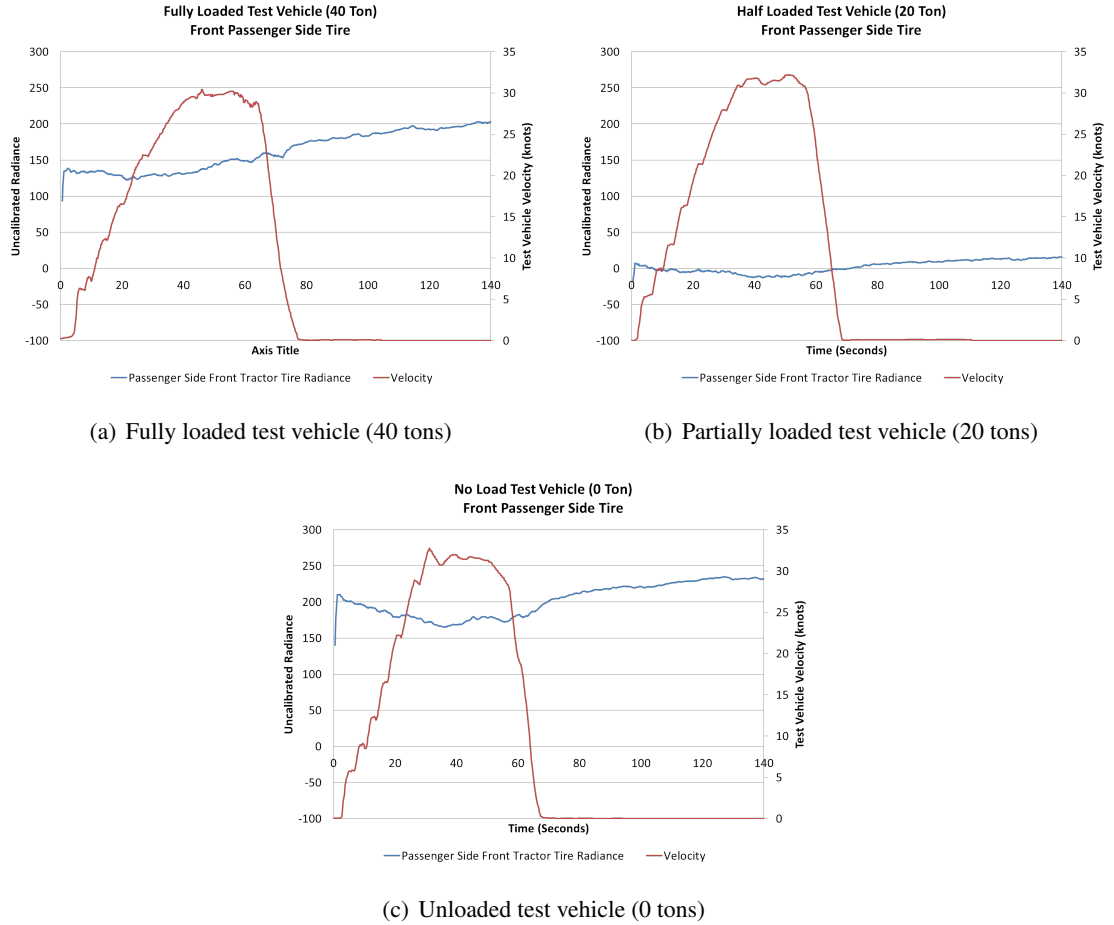


Figure 5.24: Plots of the test vehicle's velocity, and the test vehicle's tire tread uncalibrated radiance as measured by the on-board sensors.

The measured radiance values were extracted from the off-board IR video for the ROI shown in Figure 5.22 using the tracking algorithm discussed in the previous section. An example of the measured radiance data can be seen in Figure 5.23. If we assume that the tire tread has an emissivity of one, the measured radiance corresponds to a surface temperature of approximately 22.1 °C (71.8 F). This temperature is consistent with the probable road temperature at the time of the measurement as seen in Figure 5.25 (the road temperature was likely cooler than the air at this point in the morning). This data suggests that it is possible to accurately measure the surface temperature of a moving vehicle using an off-board IR sensor.




Figure 5.25: Air temperature measurements during the field testing (from Aiken, SC weather station)

Conclusions

The data showed that it is possible to measure the surface-leaving-radiance of a moving vehicle's tire tread using an off-board IR imager. If the emissivity of the surface material is known (and it is sufficiently high as to disregard reflected radiance) it is possible to convert the surface-leaving-radiance to surface temperature. We assumed an emissivity of one for the surface of the tire, which resulted in temperature values that appeared reasonable. While the results of the off-board sensor were encouraging, analysis of the on-board measurements were less so. The data suggests that the surface temperature of the tire tread is significantly affected by the temperature of the road surface. When the road surface temperature is significantly higher than the air temperature, the tire surface gains thermal energy when it is in contact with the road surface (through conduction),

and then loses thermal energy to the cooler air (mainly through convection). In order for the tire temperature phenomenology to accurately estimate the mass of the test vehicle, we must be able to quantify how much thermal energy is gained by the tire through contact with the road surface (from the tire's internal deformation, not conductive heat transfer). The measured data suggests that the thermal energy gained by the tire surface as a result of internal viscoelastic deformation is insignificant compared to the amount of thermal energy gained by the tire surface as a result of conductive heat transfer with the road surface. As a result it was concluded that remotely-sensed tire temperature is not a viable method of estimating vehicle mass.

5.5 Vibration Response

An experiment was conducted at the Savannah River Site to better understand the vibration response phenomenology. The experiment was designed to quantify the correlation between the vibration response signature induced by a vehicle onto an impediment (railroad track) and the mass of the vehicle, and also to assess the ability of a remote sensing system to measure this signature.

5.5.1 Savannah River Site Full-Scale Field Test I

During the vibration response testing, vibration data was collected from accelerometers mounted on railroad tracks as the test vehicle was driven over them. A Canon T2i Camera was used to capture high-resolution video of the test vehicle as it traveled over the railroad tracks. The camera was mounted on the side of the road, oriented perpendicular to the path of the test vehicle. Five test runs were conducted for each of the three vehicle loads for a total of fifteen runs. An example of the off-board video data can be seen in Figure 5.26.



Figure 5.26: Off-board video frame collected of the test vehicle traveling over the railroad tracks

Data Processing

The off-board Canon T2i 18-megapixel camera system collected high-definition video of the test vehicle and stored it in QuickTime file format. These files each have a resolution of 1920-by-1088 and were collected at 30 frames-per-second. Before the vibration data could be processed the distance from the test vehicle's tires to the accelerometers needed to be extracted from the video data. This was done by first ingesting the movie files into MATLAB and converting them to individual frames. Twelve frames were extracted, each of the six vehicle tires centered on each rail. Lines were drawn between the railroad tracks at the point where the test vehicle contacted the rail L_1 , at the location of the accelerometer L_2 , and between lines L_1 and L_2 (perpendicular to them both) L_3 , as seen in Figure 5.27. The distance, in meters, between the test vehicle's tire and the accelerometer was calculated using the following equation

$$d = \frac{L_0 \cdot f}{p} \sqrt{\frac{1}{L_1^2} + \frac{1}{L_2^2} - \frac{2 \cdot \cos\left(2 \cdot \arctan\left(\frac{p \cdot L_3}{2 \cdot f}\right)\right)}{L_1 \cdot L_2}} \quad (5.2)$$

where f is the focal length of the camera (m), p is the pixel pitch of the camera (m), and L_0 is the distance between the rails (m).




Figure 5.27: Image frame of the tractor-trailer traveling over the railroad tracks. The distances L_1 , L_2 , L_3 were manually measured (in pixels)

The two off-board PCB triaxial accelerometers measured the acceleration of each railroad track as a function of time along three orthogonal vectors with a sensitivity of 10 mV per 9.8 m/s^2 (acceleration due to gravity), for the frequencies between 2 and 4000 Hz. The data was collected at a rate of 51.2 kHz and stored in a binary proprietary format (National Instruments). The data was ingested using National Instruments DIAdem software and exported to a universal format. In MATLAB, the data was band-pass filtered to remove information outside the range of sensitivity of the instrument (2 to 4000 Hz). The time it takes the test vehicle to travel from one railroad track to the next is calculated from the time between acceleration impulses on each railroad track. For example, if the test vehicle is traveling from right to left, the acceleration spike from the tire's contact with the rail would first occur on the rightmost rail and then the leftmost rail. The time between the spikes is the time it took the test vehicle to travel from one rail to the next. Since distance between the rails is known, we can calculate the velocity of the test vehicle from the time. Next, the triaxial bipolar acceleration data was combined into acceleration magnitude value using the following equation

$$A_{\text{Magnitude}} = \sqrt{A_x^2 + A_y^2 + A_z^2} \quad (5.3)$$

where $A_{\text{Magnitude}}$ is the magnitude of the total acceleration acting on the sensor (m/s^2), A_x is the acceleration in the x direction (m/s^2), A_y is the acceleration in the y direction (m/s^2), and

A_z is the acceleration in the z direction (m/s^2). The acceleration magnitude data, as a function of time was broken up into sections. Each section was associated with one of the test vehicle's axles. An example of a section of data can be seen in Figure 5.28. For each section, the peak acceleration magnitude, area under the curve of the acceleration magnitude, and the location of the peak frequency was recorded. The resulting data can be seen in Table 5.4.




Figure 5.28: The acceleration magnitude experienced by the railroad track when the first wheel of the tractor-trailer crosses it

Table 5.4: Data calculated from each of the railroad track test runs; where Peak is the maximum total acceleration magnitude, AUC is the area under the curve of the total acceleration magnitude plot, and FFT Peak is the location of the peak frequency (frequency with the highest energy) in Hertz

Load (T)	Run	Avg Calculated Distance (m)	Speed (m/s)	Energy (J)	Peak (g)	AUC (g*s)	FFT Peak
40	2	0.36	2.14	207471	2.23	0.06	0.05
40	3	0.37	1.73	135147	1.64	0.06	0.03
40	4	0.34	1.85	154458	1.78	0.06	0.04
40	5	0.36	1.99	178863	1.69	0.06	0.04
20	2	0.40	1.77	85232	1.70	0.06	0.04
20	3	0.38	2.35	149947	2.35	0.06	0.06
20	4	0.42	2.10	120282	2.00	0.06	0.05
0	1	0.72	1.73	27174	1.09	0.05	0.03
0	2	0.69	1.59	23045	1.33	0.05	0.03
0	3	0.71	1.69	25828	1.19	0.05	0.03
0	4	0.74	1.56	21967	0.82	0.04	0.02
0	5	0.71	1.65	24792	0.94	0.04	0.03

Results

The data calculated above was analyzed to determine how well they correlated with vehicle load (or total vehicle kinetic energy). The normalized correlation coefficient was calculated for the matrix seen in Table 5.4 and the results can be seen in Table 5.5. The results show a strong correlation between many of the calculated values from the railroad test runs. Unfortunately, the

data also shows a strong negative correlation between the sensor-to-tire distance and the kinetic energy of the vehicle. This tells us that the driver consistently drove closer to the sensors when the tractor-trailer was fully loaded than he did when the truck was unloaded. This may be driving some of the correlation we observed and more data will be needed to quantify the significance.

Table 5.5: The results of the data correlation analysis, the highlighted cells show substantial correlation to vehicle kinetic energy

	Load (T)	Run	Avg Calculated Distance (m)	Speed (m/s)	Energy (J)	Peak (g)	AUC (g*s)	FFT Peak
Load (T)	1.00							
Run	0.17	1.00						
Avg Calculated Distance (m)	-0.93	-0.12	1.00					
Speed (m/s)	0.54	0.04	-0.70	1.00				
Energy (J)	0.94	0.15	-0.93	0.78	1.00			
Peak (g)	0.72	-0.10	-0.87	0.90	0.86	1.00		
AUC (g*s)	0.79	-0.03	-0.91	0.76	0.85	0.91	1.00	
FFT Peak	0.63	0.05	-0.79	0.99	0.83	0.94	0.84	1.00

Conclusions

The analysis of the data collected during this field test showed a significant correlation between the kinetic energy of the test vehicle and the peak acceleration magnitude, area under the curve of the acceleration magnitude, and the frequency with the highest energy. This correlation may be a result of the unintended inverse correlation between sensor-to-vehicle distance and vehicle kinetic energy. More data is needed to determine if the correlation observed is real.

5.6 Tire Deformation

Experiments were conducted at Rochester Institute of Technology (RIT), the Wyoming County International Speedway, and the Savannah River Site to better understand the tire deformation phenomenology. The experiments were designed to quantify the correlation between the tire deformation of a vehicle (both moving and stationary) and the mass of the vehicle, and also to assess the ability of a remote sensing system to measure this deformation.

5.6.1 Small-Scale Testing I

During the small-scale field tests conducted at Rochester Institute of Technology (RIT) in Henrietta, New York and Wyoming County International Speedway in Perry, New York data pertaining to the tire deformation phenomenology was collected using off-board instruments. The data was processed and analyzed to quantify the correlation between the tire deformation phenomenology and vehicle mass, and determine how well a remote sensing system could observe the phenomenology.

During this tire deformation testing, an outline of the rear two driver-side tire footprints was obtained, as well as a three-dimensional LiDAR point cloud of the outer tire, for the stationary test vehicle under four loading conditions. The footprint outlines were collected by placing foam-board between the tires and the road surface and spray painting around the perimeter (Figure 5.29). The LiDAR point cloud was collected using a LMS100 one-dimensional LiDAR system. The LiDAR system was placed approximately 1.5 meters from the stationary test vehicle and was oriented so the scanner was operating vertically. A rotating platform was used to create a three-dimensional spatial map of the test vehicles' outer driver side rear tire. Additionally, the internal pressure of tires being observed was recorded for each test load.



Figure 5.29: An example of a tire footprint outline, measured by spray painting around the perimeter of the tire

Data Processing

The footprint outlines were photographed and analyzed manually to determine the contact area of the two tires for each of the vehicle loads. A plot of measured contact area versus added weight can be seen in Figure 5.30




Figure 5.30: A plot of measured contact area of the inner and outer rear driver side tires versus added vehicle load

The data collected by the LiDAR system was stored as ENVI files. Custom code was written in MATLAB to ingest the ENVI files and store them as an array of distance values (in mm) and also an array of intensity values (dependent on the surface reflectance of the material for 650 nm light). The distance array was mapped to a Cartesian coordinate system using knowledge of the angular step size used by the LiDAR system (both vertical and horizontal angular step size). The resulting three-dimensional point cloud was manually manipulated with custom code to align, rotate, scale, and translate the point cloud so the tire and wheel were centered and oriented normal to the viewing orientation (Figure 5.31).




Figure 5.31: A three-dimensional point cloud aligned, rotated, scaled, and translated to best show the test vehicle's tire

Results

The total area of both tires' footprints (summed) was plotted versus total vehicle load. This result was compared to the predicted total footprint area of the two tires. The predicted total footprint area was calculated based on the total vehicle load, and the internal pressure of each tire, and assumed that the two measured tires would support one-third of the total vehicle load (the test vehicle has six tires in total). It should be noted that only the slope of the simulated footprint area versus total vehicle load was calculated, the y-intercept was chosen to match the measured data. This was done to account for the zero-load footprint area of the tire. Theoretically, the footprint area of the unloaded tire would be zero; however in reality the unloaded tire has a footprint with a non-zero area. Plots of the measured and simulated footprint area data versus total vehicle load can be seen in Figure 5.32.




Figure 5.32: Plot of the measured total footprint area of the rear driver side tires (measured from the footprint outlines), and the predicted total footprint area versus total vehicle load

The length and width of the footprints were plotted versus total vehicle weight as well and can be seen in Figure 5.33.




Figure 5.33: Plot of each tire’s footprint length and width as a function of total vehicle load (for the rear driver side two tires)

Conclusions

The results of the footprint outline measurements show that the total tire contact area behaves as expected for various vehicle loading conditions. Extrapolation from the measured data shows a zero-load footprint area that is larger than expected. This could be a result of the estimation that the rear two tires support one-third of the total vehicle load. Future measurement should include measurements of all of the tires (or at least all of the tires on one side of the test vehicle) so an estimation will not need to be made. The results of the footprint length and width measurements show that the tire footprint primarily expands only in the length direction (length being the direction the tire would roll) while the width of the tire footprint remains approximately constant. This means that measurement of only the tire length may be sufficient to estimate the contact area (if the width of the tire is known).

5.6.2 Small-Scale Test II

During this tire deformation testing, an outline of all of the tire footprints were obtained, as well as a three-dimensional LiDAR point cloud of the outer passenger side tires, for the test vehicle under

four loading conditions and four tire pressure conditions. The footprint outlines were collected by placing paper between the tires and the road surface and spray painting around the perimeter (Figure 5.34). The LiDAR point cloud was collected using a SICK LMS400-2000 one-dimensional LiDAR system. The test was setup so the LiDAR was scanning vertically as the test vehicle was driven past. This allowed both passenger side tires to be measured with the LiDAR system during each test run. The internal pressure of all of the tires was measured and recorded for each test load. The vehicle loads measured during this test were 11,520 pounds (unloaded), 15,530 pounds (1 concrete block), 19,540 pounds (2 concrete blocks) and 23,550 pounds (3 concrete blocks). For each vehicle load, four tire pressures were used; 100 percent (full inflation), 90 percent, 80 percent and 70 percent. This resulted in a total of 16 test runs.




Figure 5.34: An example of a tire footprint outline, measured by spray painting around the perimeter of the tire

Data Processing

The footprint outlines were photographed and analyzed manually to determine the contact area of each of the test vehicle's tires for all of the loading, and tire inflation conditions. The results can be seen in Table 5.6. The tires are referred to as follows; front driver side (FD), front passenger side (FP), rear driver side outer (RDO), rear driver side inner (RDI), rear passenger side outer (RPO), and rear passenger side inner (RPI).

Table 5.6: Table of the internal tire pressure, vehicle load, and contact area for all of the test vehicle tires.

Pressure (%)	Added Weight (lbs)	Total Weight (lbs)	Area (cm ²)						Pressure (psi)					
			FD	FP	RDO	RDI	RPO	RPI	FD	FP	RDO	RDI	RPO	RPI
100	0	11500	417	441	397	387	390	349	100	97	89	89	90	91
100	4000	15500	448	485	436	418	407	378	100	96	89	88	89	89
100	8000	19500	514	544	483	470	448	424	100	96	89	88	89	89
100	12000	23500	547	588	542	542	502	512	100	96	89	88	89	89
90	0	11500	517	538	494	439	486	442	90	85	79	78	78	78
90	4000	15500	557	568	537	513	519	531	92	86	80	79	80	80
90	8000	19500	582	627	592	577	553	543	92	86	80	79	80	80
90	12000	23500	607	647	620	592	568	559	92	86	80	78	80	80
80	0	11500	555	574	514	489	510	503	81	75	71	68	70	69
80	4000	15500	634	598	542	524	554	553	82	75	72	68	71	70
80	8000	19500	662	690	667	649	651	616	81	75	71	68	71	69
80	12000	23500	694	705	701	693	676	643	82	75	72	69	72	70
70	0	11500	570	570	502	456	540	492	66	60	59	55	59	59
70	4000	15500	604	674	529	480	608	537	67	62	60	56	60	60
70	8000	19500	672	711	644	581	689	637	67	62	59	55	60	60
70	12000	23500	756	776	704	663	728	671	66	62	59	55	60	60

The data collected by the LiDAR system was stored as ENVI files. Custom code was written in MATLAB to ingest the ENVI files and store them as an array of distance values (in mm) and also an array of intensity values (dependent on the surface reflectance of the material for 650 nm light). The distance array was mapped to a Cartesian coordinate system using knowledge of the angular step size used by the LiDAR system and the speed of the test vehicle. The resulting three-dimensional point cloud was manually manipulated with custom code to align, rotate, scale, and translate the point cloud so the tire and wheel were centered and oriented normal to the viewing orientation (Figure 5.35). Unfortunately, due to misalignment of the LiDAR system, the top of the

wheel and tire was not captured for any of the test runs.




Figure 5.35: A three-dimensional point cloud aligned, rotated, scaled, and translated to best show the test vehicle's tire

Results

The area data derived from the footprint outlines and the measured pressure data were combined to calculate the amount of mass each tire was supporting using the following equation

$$M_T = P_T \cdot A_T \quad (5.4)$$

where M_T is the mass the tire is supporting (lbs.), P_T is the internal pressure of the tire (psi), and A_T is the area of the tire's contact footprint (in^2). The results were summed for all of the tires to calculate the total mass of the vehicle. The calculated test vehicle mass versus actual test vehicle mass can be seen in Figure 5.36. The data shows that the measured vehicle mass is significantly over-estimated for all test runs. This large discrepancy is likely a result of error in the tire footprint area measurement. The area measurement does not take into account the voids in the tire area caused by the tread pattern on the surface of tire. Additionally, due to physical restrictions and wind the spray paint may not have accurately characterized the length of the tire footprints.




Figure 5.36: Plot comparing the calculated test vehicle mass (derived from the footprint outlines) versus the actual test vehicle mass

The length of the contact area of the outer-most passenger side tires was extracted from the rectified LiDAR data. This was done by first identifying the uppermost horizontal line of data that did not capture any of the passing tire. In Figure 5.35, this is the uppermost red line of data. This line of data was then plotted, along with the two lines of data directly above (Figure 5.37). This data was used to then extrapolate precisely where the tire footprint began and ended. The solid lines in Figure 5.37 represent the length of the tire footprint as calculated by each line (blue is the uppermost, red is the middle), and the extrapolated length is purple. The measured length for each test run can be seen in Figure 5.38. The results show a fairly linear relationship between percent inflation and measured footprint length.




Figure 5.37: Plot showing three adjacent horizontal lines of LiDAR data at the road/tire interface as the tire travels through the field-of-view




Figure 5.38: Plot showing three adjacent horizontal lines of LiDAR data at the road/tire interface as the tire travels through the field-of-view

Conclusions

The results of the footprint outline measurements show significant disagreement between the total vehicle mass calculated from the measured contact area and internal tire pressure and the true mass of the test vehicle. The data suggests that there is an offset not being accounted for. This offset, or correction factor, is most likely a function of the tire tread surface and how it interacts with the road. Our calculation assumes that the entire contact area is experiencing a uniform force equal to the internal tire pressure. In reality, the edges of the contact area are experiencing significantly less force than the area in the center of the contact area. Modeling of the tire system may help to better understand the tire-road interface and identify the cause of this observed discrepancy.

The LiDAR measurements showed that it is possible to extract the length of the tire's footprint using a ground based LiDAR system. The measured data showed tire footprint lengths that increased inversely with tire pressure at a rate that was consistent with the expected result. Anal-

ysis of the data showed that the angular resolution of the LiDAR system was too coarse for this application. Future testing will attempt to increase the angular resolution through changes to the hardware settings and by positioning the LiDAR system closer to the path of the test vehicle.

5.6.3 Wyoming County International Speedway Full-Scale Test

During the tire deformation testing, the pressure of all of the test vehicle's tires was monitored. This monitoring was done by six on-board wireless pressure sensors, one mounted on each of the test vehicle's six tires. The tire pressure was monitored during the brake testing (24 runs) and the suspension testing (12 runs) for a total of 36 test runs.

Data Processing

The six Pressure Pro 16 Wheel Monitor wireless pressure sensors were installed, one on each of the test vehicle's tires and measured the internal pressure of the tires in psi. The data was collected at an update rate of 5 minutes by the wireless receiver located in cab of the truck. The wireless receiver was connected via RS232 to the National Instruments I/O module. The data was converted by the on-board computer system and stored as a CSV file.

Results

The data was plotted as a function of time for each tire as seen in Figure 5.39. A linear trend was fit to each signal and plotted; the results can be seen in Figure 5.40.




Figure 5.39: Measured tire pressure for each of the test vehicle's tires as a function of time




Figure 5.40: Linear trends of the measured tire pressure for each of the test vehicle's tires as a function of time

Conclusions

The results of the tire deformation test showed small temperature variations. This suggests that the vehicle's load and operation style (heavy acceleration and braking) do not significantly affect tire pressure. The linear trends show that over the 3 hours monitored, the pressure increased for all tires by an amount less than 4%. Most of this increase was likely caused by the increasing ambient temperature.

5.6.4 Savannah River Site Full-Scale Field Test I

During the tire deformation testing, high resolution images were obtained for the outermost passenger side tire for each axle as the test vehicle traveled at an approximately constant speed. The

images were obtained with a Canon T2i 18-megapixel camera, mounted on the side of the road oriented so its optical axis was approximately perpendicular to the path of travel of the test vehicle. A laser trigger was used to initiate the collection of a $250\mu\text{s}$ exposure image when the test vehicle's tire was centered in the field-of-view. Images were collected for six test runs at each of three vehicle loads (unloaded, 20 tons, and 40 tons) for a total of eighteen runs. An example of the imagery collected can be seen in Figure 5.39.



Figure 5.41: Off-board image collected of the front passenger side test vehicle wheel while the test vehicle was traveling at a constant velocity


Data Processing

The data collected by the Canon T2i camera was stored as JPEG files. Custom code was written in MATLAB to ingest the image file, automatically identify the edge of the wheel (rim), and remove the projective distortion in the plane of the rim. Each image was read in as a 5184-by-3456 pixel 8-bit color image, and converted to an 8-bit intensity image. Next, a high pass filter was applied and the image was thresholded (Figure 5.42). This was done to highlight the edges in the image and allow large objects in the image to be identified. All but the largest ten connected regions in the


image were then removed and the remaining regions were closed using the 'close' morphological operator (Figure 5.43). This left us with ten large regions, one of which should be the wheel. Next, a 'flood fill' function was used to fill in each of the remaining regions and the most elliptical was selected (Figure 5.44). To determine how elliptical a region is, the following equation is used

$$E = \left| 1 - \frac{N}{4\pi\sqrt{|\mathbf{S}|}} \right| \quad (5.5)$$

where N is the number of points in the region, and $|\mathbf{S}|$ is the determinant of the covariance matrix of the points in the region. A value of zero denotes a perfect ellipse, the larger the value, the less elliptical the region is. This expression works because $4\pi\sqrt{|\mathbf{S}|}$ is the area of a solid ellipse, and N is the number of points in the region (the area of the region). The more these two numbers deviate, the less elliptical the region is.




(a) After the high-pass filter was applied



(b) After the high-pass filter and thresholding was applied

Figure 5.42: An example of the tire image after the high-pass filter was applied (left) and after the thresholding was applied (right)



(a) After all but the ten largest regions in the image were removed (b) After the 'close' morphological operator was applied

Figure 5.43: An example of the tire image after the ten largest regions in the image were removed (left) and after the morphological closing was applied (right)



(a) After the flood fill operator was applied

(b) The most elliptical region

Figure 5.44: An example of the tire image after the flood fill operator was applied (left) and the most elliptical region (right)

After identifying the most elliptical region in the image (which we assume to be the rim of the tire), the next step is to remove the projective distortion in the plane of the rim. To remove the projective distortion, first the ellipse must be characterized. This is done by identifying the perimeter of the ellipse, finding its center, and its major and minor axis (Figure 5.45).




Figure 5.45: The perimeter of the detected ellipse (blue), the major axis (green), and the minor axis (red)

The characterization of the ellipse is comprised of five parameters; the length of its major radius, the length of its minor radius; the location of the ellipse center along the x-axis, the location of the ellipse center along the y-axis, and the rotation (in radians) of the major axis from being parallel to the x-axis. The projective transform, required to rectify the ellipse is calculated from the following equation

$$\mathbf{T} = \mathbf{T}_{Scale} \mathbf{T}_{Rotation} \mathbf{T}_{Translation} \quad (5.6)$$

where $\mathbf{T}_{Translation}$ is the transform required to translate the image so the ellipse center lies at coordinates $(0, 0)$, $\mathbf{T}_{Rotation}$ is the transform required to rotate the image so the ellipse major axis is parallel with the x-axis, and \mathbf{T}_{Scale} is the transform required to scale (along the x-axis and y-axis) the ellipse to make it circular. Figure 5.46 shows the processed image where the tire rim has been rectified.




Figure 5.46: The processed image, where the tire rim has been rectified

Results

The processed images were then used to measure the rim to tire edge distances at multiple locations around the perimeter of the tire. This was done by manually identifying the rim and the tire edge in the processed images and measuring the length in pixels. The length in pixels was converted to real-world units (inches) using the known rim diameter. An example of the results can be seen in Figure 5.47. The results did not agree with our knowledge of the rim to tire edge distances. We

know that the tire deforms from the supported weight of the test vehicle, resulting in the reduction in rim to tire edge length at the tire to road interface. The measurements derived from the rectified images, consistently show the opposite, which we know to be erroneous. This error is likely caused by rectification error caused by the bulge of the tire sidewall away from the rectified plane of the rim edge. Other likely contributors to this error are; lens distortion, and self-obscuration.




Figure 5.47: The processed image, with the measured rim to tire edge distances displayed

Conclusions

The analysis of the data collected during this field test showed that single-view images do not work well for measuring the dimensions of three-dimensional objects. Further work could be done to attempt to recover usable measurement data from the single-view images (camera calibration); however the more practical solution would be to employ a camera array, or a LiDAR system for this measurement.

5.6.5 Savannah River Site Full-Scale Field Test II

The tire deformation test used a LiDAR system to measure the tire deformation of a moving vehicle. The test was setup so the vehicle's driver side tires passed within 0.5 m of the LiDAR system to ensure that the highest possible resolution was achieved without risking damaging the LiDAR system. The LiDAR system scanned vertically as the vehicle passed, using the movement of the vehicle to fully measure each of the wheels on the driver's side of the vehicle. LiDAR data was collected using the Sick LMS400-2000; this is a one-dimensional scanning LiDAR system that collects data at 0.2° increments at a rate of approximately $10\mu\text{s}$ per point.

Data Processing

Extensive work was done to determine the most accurate way to extract deformed and undeformed tire shape and thickness from the LiDAR data sets collected during this test. The initial approach converted the raw LiDAR distance data points into a three-dimensional point cloud using trigonometry and knowledge of the LiDAR collection cycle. Next the road surface was isolated using LiDAR scans collected prior to the vehicle entering the field-of-view. The angle of the road surface was determined and the data was rotated and translated so the Z-plane was located on the road surface. The resulting three-dimensional point cloud can be seen in Figure 5.48. The point cloud has real-world units in the Z and Y direction (mm); however the X-axis units are dependent on the speed of the vehicle. To correct the X-axis units, knowledge of the wheels, specifically the circular nature of the wheel rims, is exploited. We know that the wheel rim is circular, so when the X-axis scaling factor is chosen correctly the three-dimensional point cloud representation of the rim will be circular. This means that the X-axis scaling factor can be determined by selecting the factor that results in the most circular wheel rim. Since we are only adjusting the X-axis scaling factor and not the Y-axis (which is already in units of mm) the end result will be an X-axis correctly scaled and in units of mm. Figure 5.49 shows some visual examples of how the X axis is adjusted until the rim is circular. Next, horizontal and vertical slices of the scaled and oriented point cloud are extracted and filtered to remove noise. The results can be seen in Figure 5.50 which show the horizontal slice of the tire (we assume to be undeformed), and the vertical slice of the tire (we assume is deformed). From Figure 5.50 we can quantify the differences between a

deformed and undeformed tire.




Figure 5.48: Three dimensional point cloud rotated so the Z plane is on the road surface




Figure 5.49: Visual representation of how the X axis is adjusted from raw unscaled units on the left to units of mm on the right by fine adjustments until the rim appears round




Figure 5.50: Comparison of undeformed horizontal slices (red & blue) and a deformed vertical slice (green)

Results

The full-scale data set was processed using the methods described above. The results showed that the inherent noise of the LiDAR system when measuring dark, non-reflective surfaces such as tires was too large and precluded accurate characterization of the tire shape. Additionally, slight misalignment between the LiDAR system and the vehicle's path coupled with the rounded edges of most tires made measuring the undeformed tire size too inaccurate on the scale needed to estimate vehicle weight. The raw data sets were re-evaluated and a new processing chain was developed to precisely calculate the vertical tire deflection (deformed tire), using a combination the measured road profile and the intensity data collected by the LiDAR system. The first step in the new processing chain was to convert the raw distance data into a 3D point cloud as in the previous method. The next step was to accurately characterize the road surface as a function of distance from the LiDAR system. This was done by analyzing thousands of LiDAR scans taken prior to and after the vehicle passing through the field-of-view. The road scans were averaged, excluding outliers, to create a very accurate two-dimensional profile of the slice of road surface directly in front of the LiDAR system. The point cloud and road scan were then rotated so that the road profile was as flat as possible (linear fit). Next, a look up table (LUT) was generated from this rotated road profile that described the height of the road surface at a range of distances from the LiDAR system. Next, the three-dimensional point cloud was overlaid with the intensity data collected by the LiDAR system to create a hybrid image as seen in Figure 5.51. From this hybrid image, the interface of the wheel rim and the tire nearest to the road surface is easily discerned. This gives us the 3-dimensional location of the interface, which includes the distance from the LiDAR system along the road surface. This distance can be fed into the road surface LUT to determine precisely how far the wheel rim and tire interface is from the road surface, thus giving an accurate tire thickness measurement (± 2 mm). Figure 5.52 shows a false color image of a vehicle tire with the wheel rim and tire interface, and road surface marked.




Figure 5.51: Point cloud with intensity data overlaid




Figure 5.52: False color image with wheel/tire interface and road surface identified

Conclusions

The deformed tire thickness measurements described above was completed for all of the test runs and the results showed excellent repeatability as can be seen in Figure 5.53. These results show three different vehicle loads (unloaded, 50 % of capacity and 100 % of capacity, denoted by color) each measured during different vehicle passes in front of the LiDAR system. The correlation between tire thickness and vehicle weight was calculated for each tire for both the tractor-trailer and the passenger vehicle and can be seen in Tables 5.7 and 5.8. The results show a strong correlation between vehicle load and tire deformation for the rear tractor tires, the trailer tires, and the rear passenger vehicle tires. Additionally, the results show essentially no correlation between vehicle load and the front tire deformation for either vehicle. This means that a comparison between the tire deformation of the front tire and all of the other tires should be a very good way to determine vehicle load. In order for this method to be used in a real-world application, where no a priori knowledge of the vehicle is known, the undeformed tire thickness must be known. We have proposed a method of estimating the undeformed tire thickness using the tire specifications written on the side wall of the tire. The specifications (seen in Figure 5.54) can be used to calculate the

undeformed tire thickness using the following equation

$$UndeformedTireThickness(mm) = \frac{(100SectionWidth)}{AspectRatio} - FlangeThickness \quad (5.7)$$

where flange thickness is equal to 17.5 mm for modern passenger tires. Finally, using the ratio of deformed to undeformed tire thicknesses for the front tire and one or all of the rear tires, an estimation of the total vehicle load can be made.




Figure 5.53: Deformed tire thickness from wheel 6 (last trailer wheel) of the tractor-trailer tested during the full-scale field test at the Savannah River Site

Table 5.7: Correlation coefficient between vehicle weight and tire deformation for the tractor-trailer

		Tire Deformation					
Correlation Coefficient	Weight (lbs)	Wheel 1	Wheel 2	Wheel 3	Wheel 4	Wheel 5	Wheel 6
Weight (lbs)	1.000						
Wheel 1	-0.210	1.000					
Wheel 2	-0.881	0.317	1.000				
Wheel 3	-0.880	0.385	0.984	1.000			
Wheel 4	-0.993	0.235	0.891	0.881	1.000		
Wheel 5	-0.986	0.210	0.822	0.819	0.989	1.000	
Wheel 6	-0.991	0.220	0.837	0.832	0.991	0.992	1.000

Table 5.8: Correlation coefficient between vehicle weight and tire deformation for the passenger vehicle measured

		Tire Deformation					
Correlation Coefficient	Weight (lbs)	Wheel 1	Wheel 2	Wheel 3	Wheel 4	Wheel 5	Wheel 6
Weight (lbs)	1.000						
Wheel 1	-0.225	1.000					
Wheel 2	-0.995	0.229	1.000				




Figure 5.54: The tire specifications printed on the sidewall

5.7 Suspension Response

Experiments were conducted at the Wyoming County International Speedway and the Savannah River Site to better understand the suspension response phenomenology. The experiments were designed to quantify the correlation between the suspension response signature of a vehicle as it travels over a speed bump and the mass of the vehicle, and also to assess the ability of a remote sensing system to measure this signature. A physics-based model was developed to predict the suspension response signature of the vehicle, and the model was validated using data from a full-scale field test.

5.7.1 Wyoming County International Speedway Full-Scale Field Test

During the full-scale field test conducted at the Wyoming County International Speedway on November 20th and 21st of 2010, data pertaining to the suspension response phenomenology was collected using on-board and off-board instruments. These data were processed and analyzed to determine if correlation existed between the suspension response phenomenology and vehicle mass, and if the suspension response phenomenology was observable from a remote sensing system.

During the testing, the position of the test vehicle was monitored as the vehicle traveled over a speed bump. This monitoring was done by two on-board PCB triaxial accelerometers, and an off-board video camera mounted near the road. Three test runs were conducted for each of the

four vehicle loads for a total of twelve runs. An example of the off-board video data can be seen in Figures 5.55.



Figure 5.55: Off-board video collected of the test vehicle traveling over the speed bump; the fiducials added to the frame of the test vehicle are circled in red. Fiducials can also be seen on the wheels of the test vehicle.

Data Processing

The two PCB triaxial accelerometers mounted to the test vehicle stored acceleration as a function of time for three orthogonal axes (X,Y and Z). The data was collected at a rate of 51.2 kHz and stored in a binary proprietary format (National Instruments). The data was ingested using National Instruments DIAdem software and exported to a universal format. Preliminary processing of the data showed that the accelerometers were not sensitive to the very low frequency acceleration signal that was of interest to calculating the test vehicle's position. As a result, this data could not be used as on-board acceleration, velocity, or position data.

The off-board Kodak Zi6 video camera system collected high-definition video of the test vehicle and stored it in QuickTime .mov file format. These files each have a resolution of 1280-by-720 and were collected at 30 frames-per-second. The parameter of interest for this test is the response

of the test vehicle's suspension response. Since the suspension is simply a damped spring that couples the vehicle's frame to its axle, its response can be measured by tracking the difference between the center of the vehicle's wheel and a position on the vehicle's frame as shown in Figure 5.56. Figure 5.57 shows the basic steps used to calculate the suspension response from the measured video files. To extract the suspension response from the measured video files, the movie files were first ingested and converted into images. Next, the custom video tracking algorithm described in Section 3.7.1 was used to track center of the fiducals and the algorithm described in Section 5.6.4 was used to track the center of the vehicle's wheels. The results from these tracking algorithms were the image coordinates of the centers of each wheel and their corresponding frame fiducial as a function of time. Each signal was then rotated to align the x-axis with the axis of vehicle travel and low-pass filtered to reduce noise. The data was then converted from pixels to physical units based on the known diameter of the test vehicle's wheel. The distance between the frame and the wheel was then calculated for both the front and rear axle. The equilibrium distance (the distance between the fiducial and the wheel prior to the speed bump) was subtracted from each signal resulting in a plot describing the suspension's deviation from equilibrium as a function of time. This resulted in two signals for each run, a signal for the front suspension and for the rear suspension, each plotting the suspension's deviation from equilibrium as a function of time. Each signal was then split into two non-overlapping signals to differentiate between when the front axle and the rear axle traveled over the speed bump. After all of the data sets had been processed and evaluated it was determined that the unloaded (zero block) data sets were unusable due to low signal-to-noise which can be attributed to being collected at a greater sensor-to-camera distance than the other data sets.



Figure 5.56: Off-board video collected of the test vehicle traveling over the speed bump, the suspension response (the distance between the frame and the fiducial) is marked in red.

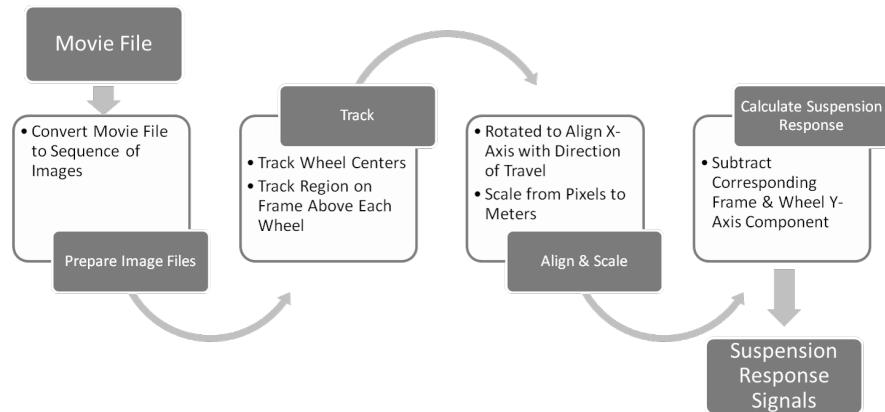


Figure 5.57: A block diagram of the processing steps used to convert the movie files into suspension response signals

Results

The suspension displacement data as function of time was smoothed using a small Gaussian kernel. Figure 5.58 shows an example of the smoothed suspension displacement from equilibrium of the

rear suspension as it traveled over a speed bump (the gray area represents the time when the vehicle's wheel was on the speed bump). The numerical derivative and double derivative of the suspension displacement from equilibrium were taken and plotted as a function of time (Figures 5.59 and 5.60, respectively).




Figure 5.58: Smoothed rear suspension displacement from static equilibrium as the vehicle's rear axle traveled over a speed bump, gray area represents the time the tire was on the speed bump




Figure 5.59: First derivative (velocity) of the rear suspension displacement as the vehicle's rear axle traveled over a speed bump, gray area represents the time the tire was on the speed bump




Figure 5.60: Second derivative (acceleration) of the rear suspension displacement as the vehicle's rear axle traveled over a speed bump, gray area represents the time the tire was on the speed bump

The suspension displacement from equilibrium as a function of time was extracted from the off-board video for the front and rear suspensions of the test vehicle using the process discussed in the previous section. To describe the suspension response, we used the basic balance of forces equation for a single damped spring system with no external forces, as defined in Equation 5.8, where $x(t)$ is the length of the spring, k_d is the damping coefficient, k_s is the spring constant, g is the acceleration due to gravity, and m is the supported mass. Supported mass being defined as the mass being supported by the suspension, which will be different for the front and rear suspensions depending on the size, mass, and position of the vehicle's load. Prior to the test vehicle reaching the speed bump, we can assume that the system is in static equilibrium with the spring force counteracting the gravitational force resulting in a compressed spring with a length x_{eq} defined in Equation 5.9.

$$x''(t) + \left(\frac{k_d}{m}\right) x'(t) + \frac{k_s}{m} x(t) - g = 0 \quad (5.8)$$

$$x_{eq} = \frac{m \cdot g}{k_s} \quad (5.9)$$

Equation 5.8 can be rewritten in terms of the deviation from the equilibrium spring length $x_M(t)$ (defined in Equation 5.10) resulting in Equation 5.11. This subtle change of variables from the actual spring length $x(t)$ to the deviation from the equilibrium spring length $x_M(t)$ is critical because the actual spring length cannot be easily measured remotely, while the deviation from the equilibrium spring length can be measured using the tracking methods described in the previous section.

$$x_M(t) = x(t) - x_{eq} \quad (5.10)$$

$$x_M''(t) + \left(\frac{k_d}{m}\right) x_M'(t) + \frac{k_s}{m} x_M(t) = 0 \quad (5.11)$$

We do not know the spring constant (k_s) or damping coefficient (k_d) of the test vehicle, but if we consider them to be constant we can use the ratios m/k_d and m/k_s as scaled versions of the desired parameter, supported mass (m). This still leaves Equation 5.11 with two unknown ratios;

k_d/m and k_s/m and three measured values; $x_M(t)$, $x'_M(t)$, and $x''_M(t)$. We can reduce the number of unknown ratios in Equation 5.11 by only considering points where $x'_M(t) = 0$ or $x_M(t) = 0$. Solving for m/k_d or m/k_s at these specific points results in the following equations:

$$\frac{m}{k_s} = -\frac{x_M(t)}{x''_M(t)} \quad (5.12)$$

$$\frac{m}{k_d} = -\frac{x'_M(t)}{x''_M(t)} \quad (5.13)$$

In order to solve Equation 5.12 we need to identify the points where $x'_M(t) = 0$, and the magnitude of $x''_M(t)$ is as large as possible (when the magnitude of $x'_M(t)$ approaches zero, Equations 5.12 and 5.13 become unstable). In addition, because we have assumed no external forces in the balance of forces equation (Equation 5.11), the point must not occur while the corresponding tire is on the speed bump (traveling over the speed bump exerts an external force). The red Xs in figures 5.61 and 5.62 show the relevant locations where $x'_M(t) = 0$, the magnitude of $x''_M(t)$ is large, and no external forces are present for the front and rear suspensions, respectively. Similarly, in order to solve Equation 5.13 we need to identify the points where $x_M(t) = 0$, the magnitude of $x''_M(t)$ is as large as possible, and the corresponding tire is not on the speed bump. The green Xs in Figures 5.61 and 5.62 represent these locations for the front and rear suspensions, respectively. Fewer points were selected for the front suspension than the rear suspension because the front suspension response damped out more quickly (the acceleration returned to zero more rapidly).




Figure 5.61: Points on the front suspension displacement plot where the supported mass ratios were calculated, green Xs denotes points where spring is in the equilibrium position, red Xs denote points where the first derivative will be zero




Figure 5.62: Points on the rear suspension displacement plot where the supported mass ratios were calculated, green Xs denotes points where spring is in the equilibrium position, red Xs denote points where the first derivative will be zero

Using Equation 5.12, the ratio, m/k_s , was calculated at the locations of the red Xs in Figures 5.61 and 5.62 using the tracked data from all usable data sets and the results can be seen in Figures

5.63 and 5.64. Using Equation 5.13 the ratio, m/k_d , was calculated at the location of the green Xs in Figures 5.61 and 5.62 using the tracked data from all usable data sets and the results can be seen in Figures 5.65 and 5.66.






Figure 5.63: Front suspension m/k_d calculations for three different vehicle loads




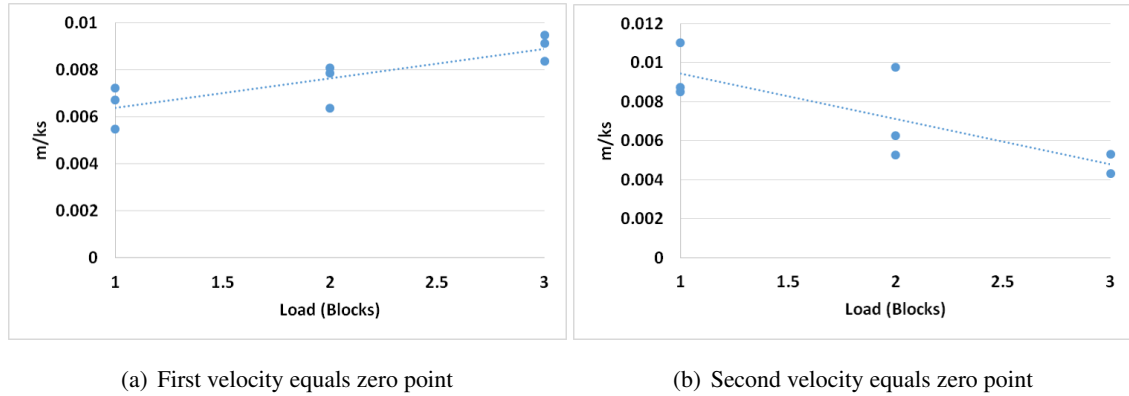
(a) First displacement equals equilibrium point

(b) Second displacement equals equilibrium point



(c) Third displacement equals equilibrium point

Figure 5.64: Rear suspension m/k_d calculations for three different vehicle loadsFigure 5.65: Front suspension m/k_s calculations for three different vehicle loads

Figure 5.66: Rear suspension m/k_s calculations for three different vehicle loads

For each of the plots seen in Figures 5.63, 5.64, 5.65, and 5.66 the normalized correlation coefficient was calculated between the ratios (m/k_d and m/k_s) and vehicle mass. The results can be seen in Tables 5.9 and 5.10. A normalized correlation coefficient of one signifies a perfect positive correlation, a normalized correlation coefficient of negative one signifies a perfect negative correlation and a normalized correlation coefficient of zero signifies no correlation. Each row in Tables 5.9 and 5.10 represents a different X from Figures 5.61 and 5.62 where the ratio (m/k_d or m/k_s) was calculated. The normalized correlation coefficient value highlighted in green denotes the highest level of correlation in each table. Note that the highest levels of correlation to vehicle mass occur in the ratios calculated from the rear suspension immediately after the vehicle's tire moves past the speed bump.

Table 5.9: Normalized correlation coefficient for the calculated ratio m/k_d to vehicle mass

Point of Calculation	Normalized Correlation Coefficient
Front Displacement = Zero (1)	0.06
Rear Displacement = Zero (1)	0.75
Rear Displacement = Zero (2)	0.60
Rear Displacement = Zero (3)	-0.26

Table 5.10: Normalized correlation coefficient for the calculated ratio m/k_s to vehicle mass

Point of Calculation	Normalized Correlation Coefficient
Front Velocity = Zero (1)	0.32
Rear Velocity = Zero (1)	0.83
Rear Velocity = Zero (2)	-0.79

Conclusions

Due to the limitations of the on-board accelerometers we were not able to validate the off-board vehicle position measurements obtained through the video tracking process described in the Data Processing Section. However, through analysis of the tracked video it became apparent that remotely sensing the position of a vehicle's frame and axles using video is sufficiently effective. Improvements to the effectiveness of this method could be achieved by increasing the resolution and sampling rate of the video. This would reduce noise and allow for more accurate calculation of the first and second derivative, which is critical to this phenomenology.

The simple balance of forces model used to estimate the supported mass (to within a scale factor) was found to be adequate for establishing correlation between the suspension response phenomenology and vehicle mass. The correlation analysis showed strong correlation between the calculated supported mass (within a scale factor) and the actual mass of the test vehicle for the rear suspension response. Specifically, strong correlation resulted when the supported mass was calculated using data immediately after the vehicle moved off of the speed bump. This was expected, as acceleration of the suspension is required for the supported mass calculation and the more time passes after the introduction of energy to the system by the speed bump, the less the suspension is accelerating. The suspension oscillation is being damped out, and the system quickly returns to static equilibrium after passing over the speed bump. Once the suspension returns to static equilibrium, the acceleration of the suspension is zero which makes Equations 5.12 and 5.13 unsolvable (divide by zero).

The correlation analysis did not show any significant correlation between the front suspension response of the vehicle and the vehicle's mass. This may be because of how quickly the front suspension oscillations damp out after traveling over the speed bump. This quick return to static equilibrium resulted in the supported mass calculation being made when the suspension was barely accelerating, which amplified the noise, which could be masking any correlation. Additionally, the

blocks used as load were all positioned on middle and rear of the vehicle, which would influence the front suspension less than the rear.

5.7.2 Savannah River Site Full-Scale Field Test I

During the full-scale field test conducted at the Savannah River Site on May 24th and 25th of 2011, data pertaining to the suspension response phenomenology was collected with on-board and off-board instruments. These data were processed and analyzed to determine if correlation existed between the suspension response phenomenology and vehicle mass, and if the suspension response phenomenology was observable from a remote sensing system. During the suspension response testing, the position of several areas of the test vehicle were monitored as the vehicle traveled over a speed bump. This monitoring was done by an on-board GPS module, and an off-board video camera mounted near the road. Five test runs were conducted for each of the three vehicle loads for a total of fifteen runs. An example of the off-board video data can be seen in Figure 5.67.



Figure 5.67: Off-board video frame collected of the test vehicle traveling over the speed bump

Data Processing

The on-board GPS measurements were used to document the velocity of the test vehicle as a function of time, but were not useful for position measurements due to the coarse precision (approximately 1 meter) of the module.

The off-board Canon T2i 18-megapixel camera system collected high-definition video of the test vehicle and stored it in QuickTime file format. These files each have a resolution of 1920-by-1088 and were collected at 30 frames-per-second. The parameter of interest for this test is the response of the test vehicle's suspension. Since the suspension is simply a damped spring that

couples the vehicle's frame to its axle, its response can be measured by tracking the difference between the center of the vehicle's wheel and a position on the vehicle's frame. Figure 5.68 shows the basic steps used to calculate the suspension response from the the measured video files.

To extract the suspension response from the measured video files, the algorithm described in Section 5.6.4 was used to track the center of the vehicle's wheels, and the custom video tracking algorithm described in Section 3.7.1 was used to track the corresponding frame locations throughout each video frame. Figure 5.69 shows the locations on the test vehicle that were selected for tracking. Note we only selected wheels that did not have another wheel closely behind. This was because we wanted to observe the suspension response fully without being influenced by the next wheel traveling over the speed bump. The results from this tracking algorithm were the image coordinates of the centers of each wheel and their corresponding frame location as a function of time. Each signal was then rotated to align the x-axis with the axis of vehicle travel. The data was then converted from pixels to physical units based on the known diameters of the test vehicle's wheels. The y-axis distance between the center of the tracked frame location and the corresponding wheel (defined in Figure 5.69) was then calculated for the three pairs. The equilibrium distance (the distance between the center of the frame location and the wheel nearest it, prior to the speed bump) was subtracted from each signal resulting in a plot describing the suspension's deviation from equilibrium as a function of time. This resulted in three signals for each run, plotting the suspension's deviation from equilibrium as a function of time.

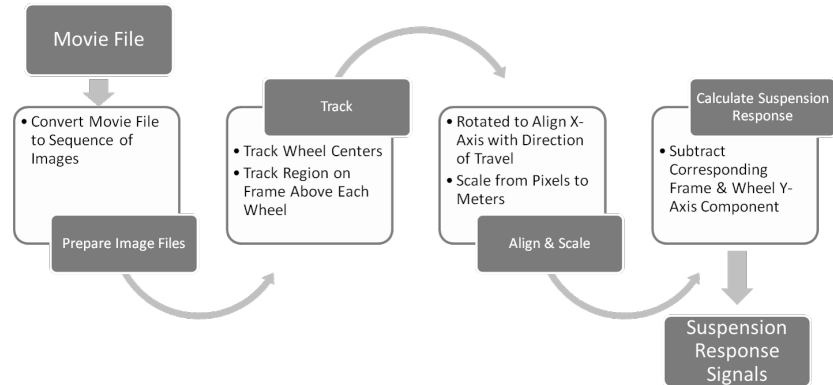


Figure 5.68: A block diagram of the processing steps used to convert the movie files into suspension response signals




Figure 5.69: The locations on the test vehicle that were tracked throughout the video frames; A & B, C & D, and E & F each make up a suspension pair

Results

The suspension displacement from equilibrium as a function of time was extracted from the off-board video for three positions on the test vehicle using the process described in the previous section. Figure 5.70 shows an example of the suspension displacement from equilibrium of the

rear end of the trailer as it traveled over a speed bump (the gray area represents the time when the vehicle's wheel was on the speed bump). The numerical derivative and double derivative of the suspension displacement from equilibrium were taken and plotted as a function of time (Figures 5.71 and 5.72, respectively).




Figure 5.70: Rear of the trailer suspension displacement from static equilibrium as the trailer's rearmost axle traveled over a speed bump, gray area represents the time the tire was on the speed bump




Figure 5.71: First derivative (velocity) of the rear of the trailer suspension displacement from static equilibrium as the trailer's rearmost axle traveled over a speed bump, gray area represents the time the tire was on the speed bump




Figure 5.72: Second derivative (acceleration) of the rear of the trailer suspension displacement from static equilibrium as the trailer's rearmost axle traveled over a speed bump, gray area represents the time the tire was on the speed bump

To describe the suspension response, we again used the basic balance of forces equation for a single damped spring system with no external forces, which we defined previously (repeated in Equation 5.14). In section 5.7.1 we calculated the supported mass of the suspension system by considering only the points where $x'_M(t) = 0$ or $x_M(t) = 0$, however this method was strongly impacted by noise since only a single point in time was considered. To improve upon this method we can solve Equation 5.14 using two points where $x'_M(t)$ or $x_M(t)$ are equal. For example, Equation 5.15 describes the relationship between suspension position, velocity, and acceleration at time, t_1 . If we find another point in time, t_2 which has the same velocity, we can set the equations equal to each other as shown in Equation 5.16, and solve for the ratio m/k_s as shown in Equation 5.17. Similarly, we can solve for the ratio m/k_d by finding two points in time with the same position value, as shown in Equation 5.18.

$$x''_M(t) \cdot m + k_d \cdot x'_M(t) + k_s \cdot x_M(t) = 0 \quad (5.14)$$

$$k_d \cdot x'_M(t_1) = -x''_M(t_1) \cdot m - k_s \cdot x_M(t_1) \quad (5.15)$$

$$m \cdot (x''_M(t_1) - x''_M(t_2)) = k_s (x_M(t_2) - x_M(t_1)) \quad (5.16)$$

$$\frac{m}{k_s} = \frac{x_M(t_2) - x_M(t_1)}{x''_M(t_1) - x''_M(t_2)} \quad (5.17)$$

$$\frac{m}{k_d} = \frac{x'_M(t_2) - x'_M(t_1)}{x''_M(t_1) - x''_M(t_2)} \quad (5.18)$$

In order to solve Equation 5.17 we need to identify two points which have the same $x'_M(t)$ value, and where the magnitude of $x''_M(t)$ is as large as possible (when subtracting two small magnitude $x''_M(t)$ values Equations 5.17 and 5.18 become unstable). In addition, because we have assumed no external forces, the points must not occur while the corresponding tire is on the speed bump. The red Xs in Figure 5.73 shows examples of the relevant points in time which have the same $x'_M(t)$ value, the magnitude of $x''_M(t)$ is large, and no external forces. Similarly, in order to solve Equation 5.18 we need to identify two points which have the same $x_M(t)$ value, the magnitude of $x''_M(t)$ is large, and no external forces. The green Xs in Figures 5.74 represents these locations for one of the three suspensions considered during this test.




Figure 5.73: Location of the pair of points with values of 0.01 m/s on the velocity versus time plot for the rearmost trailer suspension loaded with 40 tons






Figure 5.74: Location of the pair of points with values of 0.008 m on the suspension position versus time plot for the rearmost trailer suspension loaded with 40 tons

Using Equation 5.17, the ratio m/k_s , was calculated from the tracked data using pairs at similar locations to the one denoted with red Xs in Figure 5.73. This was done for the front tractor suspension, the front trailer suspension, and the rear trailer suspension for all of the test runs and the results can be seen in Figure 5.75. Similarly, the ratio m/k_d was calculated using Equation 5.18 from the tracked data using pairs at similar locations to the ones denoted with green Xs in Figure 5.74. This was done for the front tractor suspension, the front trailer suspension, and the rear trailer suspension for all of the test runs and the results can be seen in Figure 5.76.




(a) Tractor front suspension m/k_s calculations for three different vehicle loads

(b) Trailer front suspension m/k_s calculations for three different vehicle loads




(c) Trailer rear suspension m/k_s calculations for three different vehicle loads

Figure 5.75: Suspension m/k_s calculations versus total vehicle load for three different vehicle suspensions



(a) Tractor front suspension m/k_d calculations for three different vehicle loads

(b) Trailer front suspension m/k_d calculations for three different vehicle loads



(c) Trailer rear suspension m/k_d calculations for three different vehicle loads

Figure 5.76: Suspension m/k_d calculations versus total vehicle load for three different vehicle suspensions

For each of the plots seen in Figures 5.75 and 5.76 the normalized correlation coefficient was calculated between the ratios (m/k_d and m/k_s) and vehicle load. The results can be seen in Tables 5.11 and 5.12.

A normalized correlation coefficient of one signifies a perfect positive correlation, a normalized correlation coefficient of negative one signifies a perfect negative correlation and a normalized correlation coefficient or zero signifies no correlation. Each row in Tables 5.9 and 5.10 represents a different X from Figures 5.61 and 5.62 where the ratio (m/k_d or m/k_s) was calculated. The normalized correlation coefficient value highlighted in green denotes the highest level of correlation in each table. Note that the highest levels of correlation to vehicle mass occur in the ratios calculated

from the rear suspension immediately after the vehicle's tire moves past the speed bump.

Table 5.11: Normalized correlation coefficient for the calculated ratio m/k_s to vehicle mass

Suspension	Normalized Correlation Coefficient
Tractor Front	-0.230
Trailer Front	0.905
Trailer Rear	0.844

Table 5.12: Normalized correlation coefficient for the calculated ratio m/k_d to vehicle mass

Suspension	Normalized Correlation Coefficient
Tractor Front	0.0639
Trailer Front	0.806
Trailer Rear	0.886

Conclusions

The two-point method of calculating the supported mass (to within a scale factor) described in the previous section was developed as an incremental improvement on the single-point method described in Section 5.7.1. The two-point method is marginally more robust to noise than the one-point method as it utilizes twice as many points. Additionally, the calculation of m/k_d using the two-point method is invariant to any erroneous shift in suspension displacement. This invariance is important as calculating the true equilibrium position was found to be more challenging than expected due to lens distortion and tracking error. As with the one-point method, the two-point method requires that no external forces are acting on the suspension system and that the suspension is oscillating (the more dynamic the oscillations, the more robust the calculation is to noise). This means that the two-point method (like the one-point method) can only be used in the period immediately after the wheel moves past the speed bump but before the oscillations damp out.

The correlation analysis showed strong correlation between the calculated supported mass and the actual mass of the vehicle load for both the front and rear trailer suspension responses,

while the front tractor suspension, did not. The strong correlation between the trailer suspension responses and the mass of the vehicle load was expected, as the load was added to the trailer causing both the front and the rear of the trailer to experience an increase in supported load. It should be noted that the front trailer suspension is actually the suspension of the rear two tractor axles, as we are assuming that the fifth-wheel coupling between the trailer and the tractor is rigid. These results, combined with the results from Section 5.7.1 provide sufficient validation of this phenomenology's correlation to vehicle mass. The next step is to construct a physics-based model of this phenomenology.

5.7.3 Physics-Based Model

In Sections 5.7.1 and 5.7.2 it was shown that the response of a vehicle's suspension strongly correlates to the mass it is supporting. This was done using two similar methods of calculating the supported mass of the vehicle's suspension, a one-point method and a two-point method, each derived from the basic balance of forces equation. While both the one-point and the two-point method proved adequate for determining correlation, both methods have inherent limitations. Both methods rely on only a very small number of data points to calculate the supported mass, both methods require that no external forces are acting on the system, and both methods work best when the suspension is moving significantly. Additionally, the assumption that each suspension acts independent of the others is incorrect, as each of a vehicle's suspensions is coupled to the others in some way (dependent on vehicle type). A much improved method would utilize a large number of data points, including data points taken while external forces are applied and also incorporate the effects of the other suspensions. This would require a model of multiple damped oscillators coupled by the vehicle frame that predicts the response to a given input signal. Such a model would allow for the calculation of the actual vehicle load, not the supported mass of a given suspension.

To generate such a model, we should first explicitly define the phenomenology being modeled. The objective of the suspension response phenomenology is to estimate the mass of a vehicle by observing the response of its suspension while it travels over an impediment of some kind (speed bump). This impediment excites the vehicle's suspension and induces a response. The response

is manifested as movement of the vehicle frame relative to the vehicle's axles. By fitting this response to a model of a suspension system with known parameters, supporting an unknown mass, the mass of the vehicle frame can be calculated. For this model, the vehicle will be approximated as a rigid body, supported by two damped springs (a simple passenger vehicle). Excitation will be introduced to the model as the vertical position of each of the vehicle's axles. The model will be governed by a balance of forces equation, and a balance of torques equation which will be combined to define the position and orientation of the vehicle frame at all times. It should be noted that this model excludes the mass of the vehicle's axles.

A diagram of the suspension response model for a typical two-axle vehicle can be seen in Figure 5.77. In this diagram, y_{A1} and y_{A2} are the vertical positions of the vehicle axles. These vertical positions are defined to equal zero when the vehicle is traveling on flat road. A positive value for y_{A1} or y_{A2} would indicate a rise in the road as shown in Figure 5.78. A negative value for y_{A1} or y_{A2} would indicate an indent in the road as shown in Figure 5.79. The variables y_{F1} and y_{F2} are defined as the vertical positions of the vehicle frame at the points where the suspension is joined to the frame (front and rear, respectively). For simplicity, these vertical positions are defined to equal zero when the vehicle is traveling on flat road and the suspension spring is in an uncompressed state as shown in Figure 5.80. When any load is added to the system, the springs compress, resulting in negative values for y_{F1} and y_{F2} as seen in Figure 5.81. The deviation of the front suspension spring from its uncompressed state can be calculated by subtracting y_{A1} from y_{F1} . Likewise, the deviation of the rear suspension spring from its uncompressed state can be calculated by subtracting y_{A2} from y_{F2} . A positive spring length represents a stretched spring, while a negative spring length represents a compressed spring. The variable y_{cg} is defined as the vertical position of the vehicle frame at its center of gravity, and can be thought of as lying on a line connecting points y_{F1} and y_{F2} at a distance x_{cg} from y_{F1} . The variable d is defined as the distance between the point where the front and rear suspensions are joined to the frame (this constant value can be assumed to be the distance between the center of the front and rear axle). The variable θ is the angle from horizontal of the line connecting points y_{F1} and y_{F2} . The variables F_1 , F_{cg} , and F_2 are the forces acting on the vehicle's frame due to the front suspension, due to gravity, and due to the rear suspension, respectively. In this simplified model it is assumed that forces F_1 , F_{cg} , and

F_2 all act perpendicular to the flat ground.




Figure 5.77: Simplified suspension diagram for a standard two-axle vehicle




Figure 5.78: Example of the front axle traveling over a rise in the road




Figure 5.79: Example of the rear axle traveling over an indent in the road




Figure 5.80: Example of the completely unloaded vehicle traveling over flat road (both springs are in the uncompressed state)




Figure 5.81: Example of the loaded vehicle traveling over a flat road (both springs are in a compressed state)

The vehicle frame position can be described completely with only the vertical position of the point on the frame where the front suspension is joined, y_{F1} , and the angle of the frame, θ . The vertical position of any point on the line connecting y_{F1} and y_{F2} can be described as a function of y_{F1} and θ :

$$y_{F2} = y_{F1} + d \cdot \sin(\theta) \quad (5.19)$$

$$y_{cg} = y_{F1} + x_{cg} \cdot \sin(\theta) \quad (5.20)$$

For simplification, the small angle approximation ($\sin(\theta) = \theta$) will be used in Equations 5.19 and 5.20.

For this model, forces can only be applied by the front suspension, the rear suspension, and by gravity. The position of these applied forces are labeled in red in Figure 5.77 as F_1 , F_2 , and F_g , respectively.

$$F_1 = k_{s1} (y_{A1} - y_{F1}) + k_{d1} (\dot{y}_{A1} - \dot{y}_{F1}) \quad (5.21)$$

$$F_2 = k_{s2} (y_{A2} - y_{F2}) + k_{d2} (\dot{y}_{A2} - \dot{y}_{F2}) \quad (5.22)$$

$$F_g = m \cdot g \quad (5.23)$$

Where k_{s1} and k_{s2} are the spring constants of the front and rear suspension, respectively, k_{d1} and k_{d2} are the damping coefficients of the front and rear suspension, respectively, m is the total mass of the vehicle frame and load, and g is the acceleration due to gravity.

Each force defined above induces a torque about every position on the vehicle frame. However, for this model, torque will only be considered about the center of gravity of the vehicle frame. The torques are defined as, τ_1 , the torque applied by the front suspension, τ_2 , the torque applied by the rear suspension, and τ_g , the torque applied by the force of gravity.

$$\tau_1 = -F_1 \cdot x_{cg} \quad (5.24)$$

$$\tau_2 = F_2 \cdot (d - x_{cg}) \quad (5.25)$$

$$\tau_g = F_g \cdot 0 = 0 \quad (5.26)$$

Based on Equation 5.26, the torque applied by the force of gravity can be disregarded. This is the reason the center of gravity of the frame was chosen as the point at which all torque were calculated about in this model.

Now that the forces and torques have been defined, the balance of forces and balance of torques equations can be expressed in terms of system parameters:

$$\ddot{y}_{cg} \cdot m = F_1 + F_2 + F_g \quad (5.27)$$

$$\ddot{\theta} \cdot J = \tau_1 + \tau_2 \quad (5.28)$$

Equation 5.27 defines the vertical acceleration of the center of gravity of the frame as a function of the total mass of the vehicle frame and load, m , and the forces acting on it. Equation 5.28

defines the angular acceleration of the vehicle frame about its center of gravity as a function of the rotational inertia of the vehicle frame, J , and the torques acting on it.

By combining Equations 5.19, 5.20, 5.21, 5.22, 5.23, and 5.27, the balance of forces equation can be expressed in terms of the suspension parameters (k_{s1} , k_{s2} , k_{d1} , k_{d2}), the vertical position of the vehicle frame at the point where the front suspension is joined (y_{F1}), the vertical position of the front and rear axles (y_{A1} , and y_{A2}), the angle of the frame (θ), the distance of the center of gravity of the vehicle from the front axle (x_{cg}), the mass of the vehicle (m), and the acceleration due to gravity (g).

Similarly, by combining Equations 5.19, 5.21, 5.22, 5.24, 5.25, and 5.28, the balance of torques equation can be expressed in terms of the suspension parameters (k_{s1} , k_{s2} , k_{d1} , k_{d2}), the vertical position of the vehicle frame at the point where the front suspension is joined (y_{F1}), the vertical position of the front and rear axles (y_{A1} , and y_{A2}), the angle of the frame (θ), the acceleration due to gravity (g), the distance of the center of gravity of the vehicle from the front axle (x_{cg}), and the rotational inertia of the vehicle frame (J).

The balance of forces and balance of torques equations taken together, make up a system of differential equations with two-dependent variables ($y_{F1}(t)$ and $\theta(t)$) and one independent variable, time (t). If the system parameters, and initial conditions are known, this system of differential equations can be solved for any given road profile. For example, given the system parameters and road profile below, the response of the vehicle frame can be calculated as a function of time by solving the system of differential equations. The system of differential equations was solved using a differential equation solver in MATLAB (ode45.m) and the results can be seen below:

$$k_{s1} = 30,000 \text{ N/m}$$

$$k_{d1} = 1000 \text{ N s/m}$$

$$k_{s2} = 40,000 \text{ N/m}$$

$$k_{d2} = 2000 \text{ N s/m}$$

$$m = 1200 \text{ kg}$$

$$d = 3 \text{ m}$$

$$x_{cg} = 2 \text{ m}$$

$$J = 2000 \text{ kg m}^2$$

$$g = -9.8 \text{ m/s}^2$$




Figure 5.82: The vertical position of the front axle (y_{A1}), and rear axle (y_{A2}) traveling over a simulated speed bump




Figure 5.83: The vertical position of the vehicle frame at the front suspension and rear suspension attachment points (y_{F1} and y_{F2}) as the vehicle travels over a simulated speed bump

The example above shows that the model, as it has been described, can predict the position of the vehicle frame as a function of time, given that the vehicle parameters and the road profile are known. This model can be used to solve for the mass of the vehicle frame and load if the suspension parameters, road profile, and frame position are known. The suspension parameters can be estimated by vehicle type (with the exception of rotational inertia which may be estimated as a function of the frame and load mass). The road profile can be measured by tracking the vertical position of the center of the vehicle axles relative to their initial positions (assuming the vehicle is traveling on flat road at the onset of the tracking). Measuring the frame position and orientation is more complicated. Even if the vehicle is traveling over flat road and experiencing no suspension bounce from previous impediments, the suspensions are still in equilibrium with the force of gravity. This means that the vehicle frame and orientation can only be measured relative to the static equilibrium position and orientation (static equilibrium does not mean the vehicle is stationary, only that the suspension is not bouncing). Measuring the vehicle frame and orientation relative to the static equilibrium position and orientation can be done by tracking the vertical positions of points on the frame directly above each axle relative to their initial positions (assuming the vehicle is at static equilibrium at the onset of the tracking). This measurement must

be corrected by an offset equal to the static equilibrium position and orientation. Fortunately, the static equilibrium position, y_{F1_0} , and orientation, θ_0 , can be calculated as a function of the system parameters:

$$y_{F1_0} = \frac{g \cdot m (d - x_{cg})}{d \cdot k_{s1}} \quad (5.29)$$

$$\theta_0 = \frac{g \cdot m (x_{cg} (k_{s1} + k_{s2}) - d \cdot k_{s2})}{d^2 \cdot k_{s1} \cdot k_{s2}} \quad (5.30)$$

Using the above definitions for the static equilibrium position and orientation of the frame, the tracked position and orientation can be corrected, resulting in the required frame position and orientation $y_{F1}(t)$ and $\theta(t)$

$$y_{F1}(t) = y_{M1}(t) + y_{F1_0} \quad (5.31)$$

$$\theta(t) = \theta_M(t) + \theta_0 \quad (5.32)$$

where $y_{M1}(t)$ is the measured position of the frame above the front axle relative to the static equilibrium position and $\theta_M(t)$ is the angular orientation of the frame relative to the static equilibrium orientation.

This allows the system of differential equations to be expressed in terms of measurable values and suspension parameters:

$$\begin{aligned} 0 = & m \left(x_{cg} \cdot \ddot{\theta}_M + \ddot{y}_{M1} \right) + k_{s1} (y_{M1} - y_{A1}) + k_{s2} (y_{M1} - y_{A2} + d \cdot \theta_M) \\ & + k_{d1} (\dot{y}_{M1} - \dot{y}_{A1}) + k_{d2} (\dot{y}_{M1} - \dot{y}_{A2} + d \cdot \dot{\theta}_M) \end{aligned} \quad (5.33)$$

$$\begin{aligned} 0 = & J \cdot \ddot{\theta}_M + k_{s1} \cdot x_{cg} (y_{A1} - y_{M1}) \\ & + k_{s2} (y_{A2} \cdot x_{cg} - y_{M1} \cdot x_{cg} - d \cdot y_{A2} + d \cdot y_{M1} + d^2 \cdot \theta_M - d \cdot \theta_M \cdot x_{cg}) \\ & + k_{d1} (x_{cg} \cdot \dot{y}_{A1} - x_{cg} \cdot \dot{y}_{M1}) \\ & + k_{d2} (x_{cg} \cdot \dot{y}_{A2} - x_{cg} \cdot \dot{y}_{M1} - d \cdot \dot{y}_{A2} + d \cdot \dot{y}_{M1} + d^2 \cdot \dot{\theta}_M - d \cdot x_{cg} \cdot \dot{\theta}_M) \end{aligned} \quad (5.34)$$

Solving Equation 5.34 for the angular acceleration ($\ddot{\theta}_M$) results in one of the required ordinary differential equations. Substituting the result into Equation 5.33 and solving it for the vertical acceleration of the front of the frame (\ddot{y}_{M1}) results in the other. These equations for $\ddot{\theta}_M$ and \ddot{y}_{M1} make up the system of differential equations that can be used to model the position and orientation of the frame as a function of time for a given set of inputs. The measured, estimated, and unknown parameters are listed below:

$$\begin{aligned} \text{Measured Parameters: } & y_{M1}(t); \theta_M(t); y_{A1}(t); y_{A2}(t); d \\ \text{Estimated Parameters: } & k_{s1}; k_{d1}; k_{s2}; k_{d2} \\ \text{Unknown Parameters: } & m; x_{cg}; J \end{aligned} \tag{5.35}$$

However, the unknown values defined above; m , x_{cg} , and J are not precisely the terms of interest. The main value of interest is the mass of the vehicle's load (m_L), not the total mass of the vehicle (m). Assuming that the unloaded mass of the vehicle, m_0 , and its center of gravity, x_0 , are known; the unknown parameters can be rewritten in terms of the unloaded vehicle parameters (known) and the load parameters (unknown). The total mass of the vehicle, m , is simply the sum of the mass of the unloaded vehicle, m_0 , and the mass of the load, m_L :

$$m = m_0 + m_L \tag{5.36}$$

If both the unloaded vehicle and the load are approximated as rigidly coupled point masses, the position of the combined center of gravity, x_{cg} , can be calculated using the following equation

$$x_{cg} = \frac{m_0 \cdot x_0 + x_L \cdot m_L}{m_0 + m_L} \tag{5.37}$$

where x_0 is the x-axis center of gravity location of the unloaded vehicle, and x_L is the x-axis center of gravity location of the vehicle's load.

The rotational inertia of the loaded vehicle about its center of gravity, J , can be approximated as simply the sum of the rotational inertias of the unloaded vehicle and the load about the center of gravity of the total (x_{cg}).

$$J = m_0 (x_{cg} - x_0)^2 + m_L (x_{cg} - x_L)^2 \quad (5.38)$$

Substituting Equation 5.37 into Equation 5.38 results in Equation 5.39.

$$J = \frac{m_0 \cdot m_L (x_0 - x_L)^2}{m_0 + m_L} \quad (5.39)$$

Now, as shown in Equations 5.36, 5.37, and 5.39 the three unknown values; m , x_{cg} and J are expressed in terms of only two unknown values; m_L and x_L . Substituting the system of differential equations derived from Equations 5.33, and 5.34 results in the final system of differential equations. A more detailed mathematical derivation of the system of differential equations and the resulting final equations can be found in Appendix A. The final measured, estimated, and unknown parameters are listed below:

$$\begin{aligned} \text{Measured Values: } & y_{M1}(t); \theta_M(t); y_{A1}(t); y_{A2}(t); d \\ \text{Estimated Values: } & k_{s1}; k_{d1}; k_{s2}; k_{d2}; m_0; x_0 \\ \text{Unknown Values: } & m_L; x_L \end{aligned} \quad (5.40)$$

To use the differential equations to solve for the unknown parameters, the position and orientation ($y_{M1}(t)$ and $\theta_M(t)$) will be modeled for various combinations of the unknown values (m_L , and x_L) and the results will be compared to the measured position and orientation. The combination of unknown values that produces the modeled position and orientation that agrees best with the measured position and orientation will be chosen as the best estimates for the unknown values. The error in the position estimation will be defined by taking the root mean square of the difference between the measured position ($y_{M1}(t)$) and the estimated position ($\widehat{y_{M1}(t)}$) of the vehicle frame above the front axle.

$$Error_{Position} = \sqrt{\sum \left(\left(y_{M1}(t) - \widehat{y_{M1}(t)} \right)^2 \right)} \quad (5.41)$$

The error in the orientation estimation will be defined by taking the root mean square of the difference between the measured orientation ($\theta_M(t)$) and the estimated orientation ($\widehat{\theta_M(t)}$) of the vehicle frame.

$$Error_{Orientation} = \sqrt{\sum \left(\left(\theta_M(t) - \widehat{\theta_M(t)} \right)^2 \right)} \quad (5.42)$$

To calculate the total error in the estimation, the position error must be combined with the orientation error. Since the units are different (the position error is in units of meters and the orientation error is in units of radians) the orientation error must be converted from radians into the equivalent error in meters. This can be done by simply scaling the orientation error by the distance between the axles (d). This was determined by calculating the change in position of $y_{M1}(t)$ by a change in orientation $\theta_M(t)$ and use of the small angle approximation. Therefore the total error is defined as:

$$Error_{Average} = \frac{(Error_{Position} + d \cdot Error_{Orientation})}{2} \quad (5.43)$$

Using simulated data, Figures 5.84 and 5.85 show what would be considered the 'truth' data, which is the measured position and orientation of the frame relative to the equilibrium position for the system values listed below:

$$k_{s1} = 30,000 \text{ N/m}$$

$$k_{d1} = 1000 \text{ N s/m}$$

$$k_{s2} = 40,000 \text{ N/m}$$

$$k_{d2} = 2000 \text{ N s/m}$$

$$d = 3 \text{ m}$$

$$x_0 = 0.1 \text{ m}$$

$$x_L = 2.9 \text{ m}$$

$$m_0 = 300 \text{ kg}$$

$$m_L = 900 \text{ kg}$$




Figure 5.84: The vertical position of the vehicle frame at the front suspension relative to the static equilibrium position (y_{M1}) as the vehicle travels over a simulated speed bump




Figure 5.85: The orientation of the vehicle frame relative to the static equilibrium orientation (θ_M) as the vehicle travels over a simulated speed bump

While keeping all of the known values constant, the position and orientation of the frame as a function of time is modeled for a range of unknown values (m_L and x_L) using the system of differential equations. For each combination of unknown values, the position error, orientation error, and average error is calculated. A two-dimensional heat-map of the resulting position, orientation, and average errors can be seen in Figures 5.86, 5.87, and 5.88, respectively.




Figure 5.86: Heat-map showing the position error as a function of the mass of the load (kg) and the center of gravity of the load (m) (correct value denoted by red dot)




Figure 5.87: Heat-map showing the orientation error as a function of the mass of the load (kg) and the center of gravity of the load (m) (correct value denoted by red dot)




Figure 5.88: Heat-map showing the averaged error as a function of the mass of the load (kg) and the center of gravity of the load (m) (correct value denoted by red dot)

These results clearly show that this method converges to the correct solution when using the average error metric (Figure 5.88). However, this simulation was done using perfect data with no measurement error. Using simulated data with ± 1 mm of random noise added, Figures 5.89 and 5.90 show what would be considered the 'truth' data, which is the measured position and orientation of the frame relative to the equilibrium position for the same system parameters listed previously, but now with noise applied:




Figure 5.89: The vertical position of the vehicle frame at the front suspension relative to the static equilibrium position (y_{M1}) as the vehicle travels over a simulated speed bump with simulated measurement error (± 1 mm)




Figure 5.90: The angular orientation of the vehicle frame relative to the static equilibrium position (θ_M) as the vehicle travels over a simulated speed bump with simulated measurement error (± 1 mm)

A three-dimensional representation of the resulting position, orientation, and total errors for the simulated data with measurement error can be seen in Figures 5.91, 5.92, and 5.93, respectively.




Figure 5.91: Heat-map showing the position error as a function of the mass of the load (kg) and the center of gravity of the load (m) for data with simulated measurement error (correct value denoted by red dot)




Figure 5.92: Heat-map showing the orientation error as a function of the mass of the load (kg) and the center of gravity of the load (m) for data with simulated measurement error (correct value denoted by red dot)




Figure 5.93: Heat-map showing the averaged error as a function of the mass of the load (kg) and the center of gravity of the load (m) for data with simulated measurement error (correct value denoted by red dot)

These results show that the model is capable of accurately estimating the mass and center of gravity of the vehicle's load to within a reasonable margin of error, even in the presence of measurement error when using simulated data. The next step is to validate the model using actual field measurements.

5.7.4 Savannah River Site Full-Scale Field Test II

During the full-scale field test conducted at the Savannah River Site on December 10th and 11th of 2012, data pertaining to the suspension response phenomenology was collected using off-board instruments only. These data were processed and used to test the effectiveness of the physics-based model described in Section 5.7.3 in calculating the mass of the test vehicle based on the remotely-observed suspension response of the vehicle as it travels over an impediment (speed bump).

During this testing, the position of several areas of the test vehicle were monitored as the vehicle traveled over a speed bump. This monitoring was done by an off-board video camera mounted near the road. Two test vehicles were used during this field test, a 2006 Honda Accord,

and a Western-Star tractor pulling a 40-foot long trailer. Each test vehicle was tested unloaded, at half of the recommended maximum load, and at the recommended maximum load. Five test runs were conducted for each vehicle load for a total of fifteen runs per vehicle. Examples of the off-board video data can be seen in Figures 5.55 and 5.67.



Figure 5.94: Off-board video frame collected of the passenger vehicle traveling over the speed bump



Figure 5.95: Off-board video frame collected of the tractor-trailer traveling over the speed bump

Data Processing

The off-board ARRIFLEX D-21 camera system collected high-definition video of the test vehicles and stored each frame in ‘Digital Picture Exchange’ file format. These files each have a resolution of 2880-by-1620 and were collected at 30 frames-per-second in uncompressed 12-bit per pixel Bayer mode. The camera was placed approximately 10 m from the road and oriented perpendicular to the path of the vehicle. The parameters of interest for this test are the position and orientation of the vehicle frame, and the location of the wheel centers as a function of time. To extract the

desired data from the video frames, we employed the tracking algorithm described in Section 3.7.1 to track the position of several regions on the vehicle as a function of time. Figures 5.96 and 5.97 show the locations on the test vehicles that were selected for tracking.



Figure 5.96: Off-board video frame collected of the tractor-trailer traveling over the speed bump



Figure 5.97: Off-board video frame collected of the tractor-trailer traveling over the speed bump

After the positions of interest have been tracked, the data is converted from pixels to units of distance (meters) using known vehicle dimensions. The tracked wheel diameters, and distance between wheel centers are used to calculate the proper scale factor for this conversion. Next the data is rotated so that the x-axis is aligned parallel to the path of the vehicle and the y-axis is aligned parallel to the force of gravity. This is done automatically by rotating the tracked position data so that their endpoints align. This simple method is adequate as we have required the test vehicle to be in static equilibrium as it enters the field-of-view of the camera system, and we have determined that the suspension system has returned to static equilibrium prior to leaving the field-of-view of the camera system. Next, the data is both low-pass filtered to remove noise, and

high-pass filtered to remove the lowest frequency components (demeans the data and also removes any lens distortion effects like barrel distortion or pincushion distortion that caused the data to tail upwards or downwards near the edges).

To aid in computation, we fit a Gaussian-like function to the tracked wheel data to allow for direct calculation of the first derivative and avoid having to rely on numeric derivatives. The form used for this function can be seen in Equation 5.44. Where H is the height of the speed bump, t is time, t_0 is the point in time when the wheel is at the peak of the speed bump, W is the width of the speed bump, and V is the velocity of the test vehicle. The derivative of this function can be seen in Equation 5.45. An example of the modeled versus measured wheel data can be seen in Figure 5.98. The model does not account for the response of the tire, which causes some disagreement in the period directly after the speed bump.

$$f(t) = H \cdot \exp\left(-\left(\frac{1.3 \cdot V \cdot (t - t_0)}{W}\right)^4\right) \quad (5.44)$$

$$f'(t) = 11.54 \cdot \left(\frac{V}{W}\right)^4 \cdot (t_0 - t)^3 \cdot f(t) \quad (5.45)$$




Figure 5.98: Modeled data fit to the measured y-axis wheel position

Results

For the passenger vehicle, position data for the center of each wheel, and a frame position directly above each wheel was extracted from the off-board video and processed using the methods described in the previous section. An example of this processed data can be seen in Figures 5.99 and 5.100. For the tractor-trailer, position data for the center of each trailer wheel, and a location on the trailer were extracted from the off-board video and processed using the methods described in the previous section. An example of the processed data for the tractor-trailer can be seen in Figure 5.101.




Figure 5.99: Example data set for the front of the passenger car




Figure 5.100: Example data set for the rear of the passenger car





Figure 5.101: Example data set for the front of the tractor-trailer

For the passenger vehicle data sets, the physics-based model described in Section 5.7.3 was used to generate nearly two-hundred and fifty thousand simulations for each test run. Each simulation used a different combination of vehicle load mass, and vehicle load center of gravity. Each simulation used the same best-fit Gaussian-like approximation of the tracked wheel centers as inputs, and used the vehicle and suspension parameters seen in Table 5.13. Where m_0 is the unloaded mass of the vehicle, x_0 is the x-axis distance of the vehicle's center of gravity from the front wheel center (positive means towards the rear of the vehicle), J_0 is the moment of inertia of the unloaded vehicle, k_{s1} is the front suspension spring constant, k_{d1} is the front suspension damping coefficient, k_{s2} is the rear suspension spring constant, and k_{d2} is the rear suspension damping coefficient. All of these parameters are required to be known for this method of analysis to be possible. The only parameters allowed to vary during the simulations were the mass of the load, and the center of gravity of the load. The load mass array ranged from 0 to 600 kg at step sizes of 2.7 kg and the center of gravity array ranged from 0 to 5 m at step sizes of 0.01 m. Each generated frame response was compared to the measured frame response, and three error values were calculated for each simulation based on the error equations from Section 5.7.3. The frame position error, which characterized how well the model fit the y-axis position of the front frame


position, the orientation error, which characterized how well the model fit the frame rotation, and the averaged error, which combined the position and orientation errors. Heat map representations of the averaged error can be seen in Figure 5.102. It should be noted that it took extensive research and brute force calculations to find valid suspension parameters for this test vehicle. Future work should employ field testing to measure the suspension parameters, or enlist the support of manufacturers to generate the required database of suspension response values.

Table 5.13: Table of passenger vehicle parameters


Vehicle Parameter	Value
m_0	1430 kg
x_0	1.4 m
J_0	2261 kg m ²
k_{s1}	53,306 N/m
k_{d1}	3555 N s/m
k_{s2}	104,640 N/m
k_{d2}	4465 N s/m




(a) Average error heat map for the unloaded car




(b) Average error heat map for the unloaded car




(c) Average error heat map for the half loaded car



(d) Average error heat map for the half loaded car



(e) Average error heat map for the fully loaded car




(f) Average error heat map for the fully loaded car

Figure 5.102: Suspension m/k_d calculations versus total vehicle load for three different vehicle suspensions


For the tractor-trailer data sets, we were unable to use the full model described in Section 5.7.3. The full model would require that a position be tracked at the front and rear of the trailer simultaneously, which was not possible due to the limited field-of-view of the imaging system (by the time the rear of the trailer got into frame, the front of the trailer had moved out of frame). To account for this limitation, the model was modified so as to restrict movement of the front end of the trailer (essentially the front of the trailer was treated as a pin style hinge) and allow normal movement of the rear of the trailer. This modified model was used to generate nearly two-hundred and fifty thousand simulations for each test run. Each simulation used a different combination of vehicle load mass, and vehicle load center of gravity. Each simulation used the same best-fit Gaussian-like approximation of the tracked trailer wheel centers as inputs, and used the vehicle and suspension parameters seen in Table 5.14. Where k_{s2} is the rear trailer suspension spring constant (for each of the rear trailer wheels), and k_{d2} is the rear trailer suspension damping coefficient (for each of the rear trailer wheels). As we are using a simplified model, the moment of inertia of the trailer is not required. The only parameters allowed to vary during the simulations were the mass of the load, and the center of gravity of the load. The load mass array ranged from 0 to 50,000 kg at step sizes of 225 kg and the center of gravity array ranged from 0 to 20 m at step sizes of 0.045 m. Each generated frame response was compared to the measured frame response, and three error values were calculated for each simulation based on the error equations from Section 5.7.3. The frame position error, which characterized how well the model fit the y-axis position of the front frame position, the orientation error, which characterized how well the model fit the frame rotation, and the averaged error, which combined the position and orientation errors. Heat map representations of the averaged error can be seen in Figure 5.103.

Table 5.14: Table of tractor-trailer parameters


Vehicle Parameter	Value
m_0	16,000 kg
x_0	6 m
k_{s2}	3,898,903 N/m
k_{d2}	63,744 N s/m



(a) Average error heat map for the unloaded tractor-trailer (b) Average error heat map for the unloaded tractor-trailer



(c) Average error heat map for the half loaded tractor-trailer (d) Average error heat map for the half loaded tractor-trailer



(e) Average error heat map for the fully loaded tractor-trailer (f) Average error heat map for the fully loaded tractor-trailer

Figure 5.103: Suspension m/k_d calculations versus total vehicle load for three different vehicle suspensions

Conclusions

The physics-based method of calculating the mass of a vehicle's load described in Section 5.7.3 was tested in the previous section using data collected remotely from a passenger vehicle and a tractor-trailer. The results from the passenger car test utilized the unmodified physics-based model, and showed excellent results. The calculated and predicted mass and center of mass values agreed very well in all cases (the disagreement in center of mass location observed in the unloaded condition is irrelevant as a load with a mass of zero does not have a center of mass). The tractor-trailer results were derived from a simplified version of the physics-based model, which did not consider the suspension response of the front end of the trailer. This was necessary because the field-of-view of the imaging system was not sufficient to capture the entirety of the trailer. Since the interaction between the suspensions could not be leveraged, the results of the tractor-trailer do not converge to a point but to a line. This was expected, because the results are completely dependent on the load that the rear suspension is experiencing, and the rear suspension can experience the same load from any combination of load mass and center of gravity location. For example, a very heavy load with a center of gravity located near the front of the trailer would exert the same force on the rear suspension as a much lighter load with a center of gravity located near the rear of the trailer. Despite being forced to use the simplified model, the results of the test were quite good. With the exception of the unloaded conditions, all of the estimated values fell on or very close to the line that contained the true value. If the test was rerun with either two synchronized imaging systems that allowed for high resolution measurements to be made of both the front and rear ends of the trailer, the data could be processed with the full physics-based model and would likely get results that converge to the correct solution as the passenger car data did.

5.8 Frame Twist

A physics-based model was developed to predict the relationship between the frame twist signature of a vehicle accelerating from a stop and the mass of the vehicle. Two experiments were conducted at the Savannah River Site to validate the model, and also to assess the ability of a remote sensing system to measure the appropriate signatures.

5.8.1 Physics-Based Model

In order for a vehicle to accelerate, rotational energy must be transferred from the vehicle's engine to its drive wheels. During this transfer of energy, a torque is applied by the engine to the vehicle's drive shaft. An equal torque must be supplied by the vehicle's frame to prevent the engine from rotating; this torque causes the frame to twist relative to the vehicle's axles. For an engine mounted so the crankshaft is perpendicular to the drive-shaft, known as a transverse engine, the frame twist would manifest as a front to back rotation. For an engine mounted so the crankshaft is parallel to the drive-shaft, known as a longitudinal engine, the frame twist would manifest as a left-to-right rotation (when facing the vehicle). When the engine is applying torque to the vehicle frame, τ_E , the frame twists relative to the axles and the vehicle's suspension springs are compressed and stretched as seen in Figure 5.104. This compression and stretching of the suspension springs causes pushing and pulling forces to be applied at the points where the frame and suspension are joined. The amount of torque generated by these forces, τ_F , is dependent on their distance from the axis of rotation, r , and the amount of frame twist, θ , as seen in Equation 5.46 (the θ and θ_0 terms accounts for the force applied non-perpendicularly to the radial axis). The forces, F_1 and F_2 are generated by the vehicle's suspension and can be described in terms of the frame twist and the total suspension spring constant, k_s , as seen in Equations 5.47 and 5.48. By combining all the equations we can describe the torque generated by the suspension as a function of frame twist (Equation 5.49). Since we know that the torque the engine is applying to the frame is equal in magnitude and opposite in direction to the torque the suspension is applying on the frame we can describe the engine torque as a function of the frame twist as seen in Equation 5.50.

$$\tau_F(t) = r \cdot \cos(\theta(t) + \theta_0) \cdot (F_2 - F_1) \quad (5.46)$$

$$F_1(t) = -\frac{k_s \cdot r \cdot \tan(\theta(t))}{2} \quad (5.47)$$

$$F_2(t) = \frac{k_s \cdot r \cdot \tan(\theta(t))}{2} \quad (5.48)$$

$$\tau_F(t) = k_s \cdot r^2 \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) \quad (5.49)$$

$$\tau_E(t) = -k_s \cdot r^2 \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) \quad (5.50)$$




Figure 5.104: Longitudinal engine imparting torque, τ_E , to the vehicle frame causing it to twist until the torque generated by the suspension forces (F_1 and F_2) are sufficient to match

Now that we can relate the torque being generated by the engine to the amount of frame twist, we need to describe how much power the engine is outputting. Rotational power is defined by the torque and angular velocity of the rotating shaft (Equation 5.51). So to calculate the power output of the engine we need to measure the angular velocity of the engine, $\omega_E(t)$. The angular velocity of the engine can be measured directly from the acoustic emissions of the engine. Substituting Equation 5.50 into Equation 5.51 we can express the power output of the engine in terms of the frame twist and the engine frequency (and several scalars).

$$P_E(t) = \tau_E(t) \cdot \omega_E(t) \quad (5.51)$$

$$P_E(t) = -\omega_E(t) \cdot k_s \cdot r^2 \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) \quad (5.52)$$

Assuming that all of the rotational power being generated by the engine is converted into linear vehicle motion, the definite integral of the power output of the engine would equal the change in kinetic energy of the vehicle as shown in Equation 5.53. Where $KE(t)$ is the vehicle's kinetic energy at time t . Plugging in our expression for engine power and the standard equation for kinetic energy we get Equation 5.54. Assuming the vehicle is accelerating from a stopped position, solving for the mass of the vehicle m , and taking the absolute value of the definite integral, we are left with Equation 5.55. Note that we take the absolute value because frame twist may be positive or negative depending on the rotation direction of the engine.

$$KE(t_f) - KE(t_0) = \int_{t_0}^{t_f} P_E(t) dt \quad (5.53)$$

$$\frac{1}{2} \cdot m \cdot (x'(t_f)^2 - x'(t_0)^2) = -k_s \cdot r^2 \int_{t_0}^{t_f} \omega_E(t) \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) dt \quad (5.54)$$

$$m = \frac{2 \cdot k_s \cdot r^2 \left| \int_{t_0}^{t_f} \omega_E(t) \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) dt \right|}{x'(t_f)^2} \quad (5.55)$$

From Equation 5.55, we can solve for the mass of the vehicle in terms of three dynamic values; the engine angular velocity ($\omega_E(t)$), frame twist ($\theta(t)$), and the velocity of the vehicle ($x'(t)$) and three scalar values; the suspension spring constant (k_s), distance from the engine axis of rotation to the suspension spring (r), and the angle of the engine force relative to the suspension spring (θ_0). The dynamic values can be measured remotely, but the scalar values will need to be estimated based on vehicle type. To model this phenomenology, the following values will be used to simulate the tractor trailer (note that the acceleration and engine speed are assumed to be constant).

$$k_s = 180,000 \text{ N/m}$$

$$r = 0.8 \text{ m}$$

$$\theta_0 = 32^\circ$$

$$x''(t) = 0.5 \text{ N/m}^2$$

$$\omega_E = 33 \text{ Hz}$$

Using the scalar values above we can model the expected frame twist value using Equation 5.55. The acceleration rate will be assumed to be a constant 0.5 m/s^2 , and the engine speed will be assumed to be a constant 33 Hz (1980 RPM). The movement of the vehicle in this simulation is shown in Figure 5.105




Figure 5.105: The modeled position and velocity of the test vehicle versus time used for all simulated data

Using Equation 5.55 and the simulated and estimated values described above, we were able to calculate the predicted frame twist values for the unloaded, partially loaded, and fully loaded tractor trailer. The results can be seen in Figure 5.106.




Figure 5.106: The modeled position and velocity of the test vehicle versus time used for all simulated data

The simulated results show the expected peak frame twist values to range from 0.4 degrees, for an unloaded tractor trailer, to 1.1 degrees for a full loaded tractor trailer accelerating at a constant (and conservative) 0.5 m/s^2 . This frame twist magnitude, while small is not prohibitively small as to disqualify remote measurement. This model predicts that the frame twist signature should be large enough to measure remotely. The next step is to validate the model using actual field measurements.

5.8.2 Savannah River Site Full-Scale Field Test I

During the full-scale field test conducted at the Savannah River Site on May 24th and 25th of 2011, data pertaining to the frame twist phenomenology was collected using off-board instruments. These data were processed and analyzed to validate the methodology detailed in the previous section and determine if the frame twist phenomenology was observable from a remote sensing system. During the testing, the front of the tractor was monitored with a Canon T2i 18-megapixel video camera as the vehicle accelerated from a stop. The monitoring was done with the video camera mounted in the center of the road directed at the center of the tractor's grill. Acoustic data was collected with the T2i camera microphone. Four test runs were conducted for each of the three vehicle loads for a total of twelve runs. An example of an image from the off-board video

data can be seen in Figure 5.107.



Figure 5.107: Off-board video frame collected of the test vehicle accelerating from a stop

Data Processing

The video data collected for each test run by the Canon T2i video camera was stored as a Quick-Time multimedia (MOV) file. Each movie file was ingested into MATLAB and processed using a series of custom algorithms as shown in Figure 5.108. First, each frame of the MOV file was extracted and converted into a 1920-by-1088 pixel gray-scale image. Next, a correlation-based video tracking algorithm was used to track and isolate the same region of the vehicle's grill for each frame. The tracking method was designed to track the location of a predefined region, and to be tolerant of small frame-to-frame changes in scale and rotation. An example of the results of the tracking algorithm can be seen in Figure 5.109.

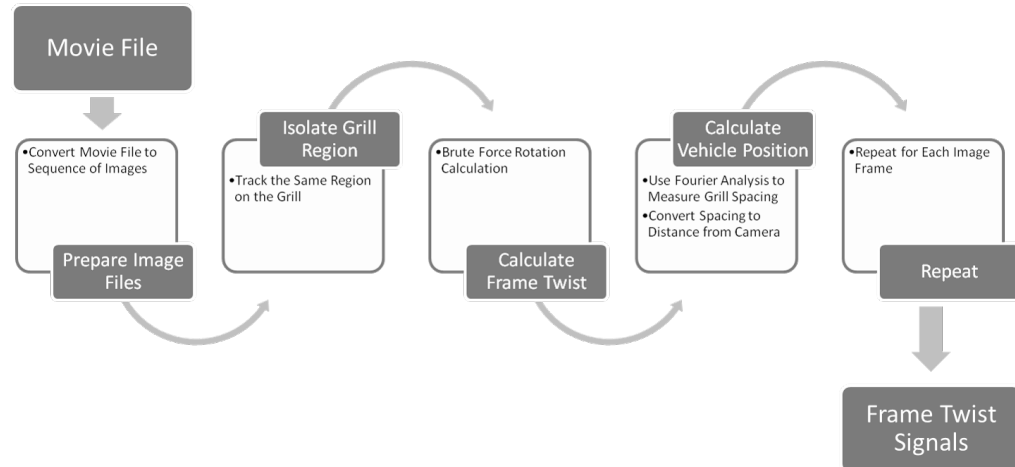


Figure 5.108: A block diagram of the processing steps used to convert the movie files into frame twist signals




Figure 5.109: An image frame with the extracted grill region identified

To measure the angle of the test vehicle's frame as a function of time, a simple brute force approach was employed. Each isolated grill region extracted by the tracking algorithm was rotated using a bilinear interpolation method of image rotation. The rotated image was cropped to remove the edge artifacts caused by the rotation, and the standard deviation was calculated for each vertical column of pixels and summed. This resulted in a scalar value that described the vertical variability of the image (a value of zero would occur if each row of the image was identical). This was repeated for a range of rotation values and the rotation that resulted in the lowest value was selected as the rotation of the grill (relative to the vertical image axis). The process was repeated for each image frame. See Figure 5.110 for an example of the rotation calculation method.




Figure 5.110: The left image is the unrotated grill image, the middle plot shows the vertical variability values for a range of image rotations (-5° to 5°), and the right image shows the grill image rotated so the grill runs perpendicular to the horizontal image axis.

To measure the camera to vehicle distance as a function of time, the spacing of the grill was calculated using Fourier analysis. After the grill images were rotated so the grill runs perpendicular to the horizontal image axis, the average of each vertical column was taken resulting in a grill profile. A Hamming window was applied to the profile to reduce instances of aliasing. The absolute value of the squared Fourier transform was then taken and the peak frequency was identified as the frequency of the grill, and the inverse of the frequency as the period (in pixels). See Figure

5.111 for an example of this process.




Figure 5.111: Left plot shows the grill intensity profile, middle plot shows the demeaned intensity profile with a Hamming window applied, and the right plot shows the power spectra of the middle plot

The grill spacing in pixels was calculated for each video frame using this method. The camera to vehicle distance for each frame was then calculated using the equation below. Where d is the vehicle to camera distance, S is the actual grill spacing in meters, f is the focal length of the camera in meters, ρ is the pixel pitch of the camera's detector array in meters, and N is the grill spacing in pixels (the period).

$$d = \frac{S \cdot f}{\rho \cdot N} \quad (5.56)$$

The audio data collected for each test run was used to calculate the engine speed using Fourier analysis. The audio data was first extracted from the video file and converted into a pressure versus time signal. A 0.3 s window of data was selected starting at time zero, a Hamming window was applied, and the absolute value of the squared Fourier transform was taken resulting in a power spectra. The operating frequency of the engine was calculated by locating the peaks in the power spectra associated with its 3rd, 6th, and 9th harmonic and then dividing the frequency by the number of the harmonic. For example, if the 3rd harmonic was found at 30 Hz the engine frequency would be 10 Hz. Three harmonics were used because no single harmonic was found to be reliably persis-

tent in the audio signal. The initial search locations were restricted to frequencies corresponding to the typical idling frequency of a large diesel engine, approximately 420 - 960 RPM (or 7 - 16 Hz). Once the initial engine frequency was identified, the window of data was moved forward in time by 1 ms and the process was repeated, with the search area for each harmonic being substantially restricted to within 5 Hz of the previous harmonic. If the three harmonics agree well, the average of the three calculated engine frequency values is used, if one of the three harmonics does not agree it is discarded, if none of the harmonics agree well, the lowest (3rd harmonics) is selected and used to calculate the engine frequency. This continues until the entire pressure signal has been processed, resulting in an engine frequency versus time signal. An example of this process can be seen in Figure 5.112.




Figure 5.112: From left-to-right the raw pressure signal, the pressure signal with a Hamming window applied, and the power spectra (during idle)

Using the processes described in this section the angular frame twist, position, and engine speed of the test vehicle was calculated as a function of time for each of the test runs. The results were analyzed to establish their correlation to vehicle mass.

Results

The data processing described in the previous section resulted in angular frame twist, position, and engine speed as a function of time for each test run. An example dataset can be seen in Figure 5.113.




Figure 5.113: From left-to-right; the angular frame twist, position, and engine speed of the unloaded test vehicle accelerating from a stop

Equation 5.55, from the previous section was modified slightly to account for the unknown spring constant, k_s , to give Equation 5.60 which we used to calculate the mass to spring constant ratio of the test vehicle for each test run. The final time, t_f was identified as the time where the test vehicle reached its peak velocity and the values for r and θ_0 were estimated to be 0.8 m and 33° , respectively. The results can be seen in Figure 5.114.

$$\frac{m}{k_s} = \frac{2 \cdot r^2 \left| \int_{t_0}^{t_f} \omega_E(t) \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) dt \right|}{x'(t_f)^2} \quad (5.57)$$




Figure 5.114: The calculated mass to spring constant ratio for each test run versus the true test vehicle mass

Since the actual vehicle mass was known for each test run, we can use the mass to spring constant ratio to calculate the unknown spring constant k_s . The spring constant k_s was calculated to be 182,395 N/m. Scaling the data in Figure 5.114 by the calculated value for k_s we can get the calculated mass versus true mass plot seen in Figure 5.115. The results show an average error of approximately 4000 kg.




Figure 5.115: The calculated mass versus the true test vehicle mass

Conclusions

Multiple image processing techniques were employed to successfully extract the frame twist, vehicle position, and engine angular velocity data as a function of time from the raw image files. Analysis of the extracted data showed low to moderate noise. This confirmed that it was possible to remotely measure the parameters pertinent to the frame twist phenomenology. Improvements to the effectiveness of this method could be achieved by increasing the resolution and sampling rate of the video. This would reduce noise and allow for more accurate calculation of the vehicle velocity, which is critical to this phenomenology. Additionally, a higher quality acoustic measurement system should allow for more accurate measurement of the engine angular velocity.

The results seen in Figure 5.115 show an extremely strong correlation between the calculated vehicle mass and the true vehicle load. The correlation coefficient was calculated to be 97.7%. The calculated value for the spring constant k_s was found to be 182,395 N/m which is 3-4 times larger than the spring constant calculated for the tractor suspension for the suspension response

phenomenology. However it is inappropriate to compare these spring constants directly as the frame twist phenomenology is impacted by the combined stiffness of the front and rear tractor suspensions, while the suspension response is only impacted by one or the other.

While these results are very encouraging, the magnitude of the error in the calculated mass is fairly large (approximately 10 %). This error is likely caused by noise in the engine RPM measurement, and to a lesser extent, noise in the vehicle velocity measurement. This noise could be reduced by utilization of a video camera system that collects data at a faster rate, and by an improved acoustic measurement system.

5.8.3 Savannah River Site Full-Scale Field Test II

During the full-scale field test conducted at the Savannah River Site on December 10th and 11th of 2012, data pertaining to the frame twist phenomenology was collected using on-board and off-board instruments for the tractor trailer vehicle. These data were processed and used to test the effectiveness of the physics-based model described in Section 5.8.1 in calculating the mass of the test vehicle based on the remotely-observed frame twist of the vehicle as it accelerates from a stop. Improved measurement equipment was employed to try and improve upon the results generated during Savannah River Site Full-Scale Field Test I.

During this testing, the front of the tractor was monitored with the Vision Research Phantom v5.1 high-speed video camera as the vehicle accelerated from a stop. The monitoring was done with the video camera mounted in the center of the road directed at the center of the tractor's grill. Acoustic data was collected with a calibrated PCB Piezotronics Free-Field Condenser microphone mounted on the outside of the rear windshield of the tractor cab.

The same test vehicle was used during this test as in the previous SRS frame twist test; a Western-Star tractor pulling a 40-foot long trailer. The test vehicle was tested unloaded, at half of the recommended maximum load, and at the recommended maximum load. Four test runs were conducted for each vehicle load for a total of twelve runs. An examples of the off-board video data can be seen in Figure 5.116.



Figure 5.116: Example frame from the Phantom Video camera

Data Processing

The high-speed video data collected for each test run by the Phantom video camera was stored uncompressed in a proprietary format (.CINE) developed by Vision Research. The ImageJ software package was used to convert each from of the .CINE video files into individual images. These images were ingested into MATLAB and processed using a series of custom algorithms as shown in Figure 5.117. First, each image frame was ingested into MATLAB as a 1024-by-1024 pixel gray-scale image. Next, a correlation-based tracking algorithm was used to track and isolate the same region of the vehicle's grill for each frame. In order to track the extremely small frame-to-frame changes in position points on the bumper were also tracked and isolated. The tracking algorithm was designed to be robust to small frame-to-frame changes in scale and rotation. The regions tracked by the algorithm are identified in Figure 5.118.

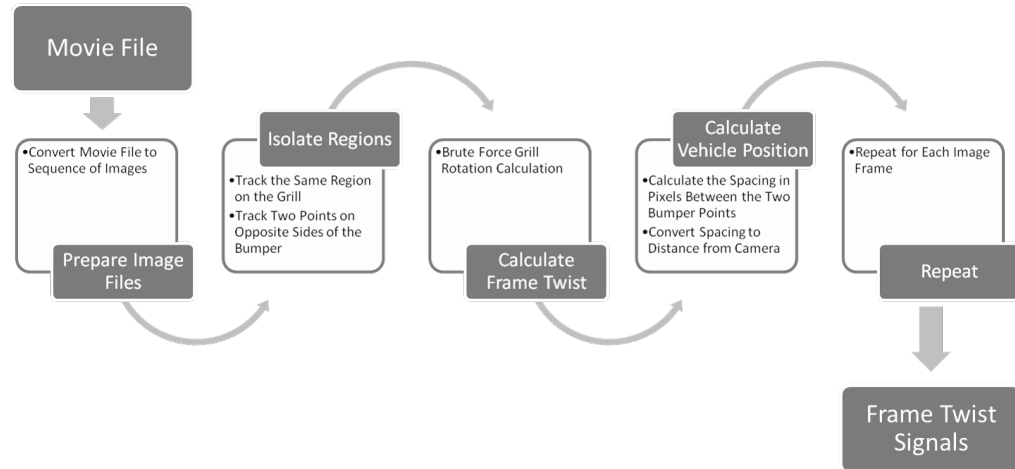


Figure 5.117: A block diagram of the processing steps used to convert the movie files into frame twist signals




Figure 5.118: The regions on the test vehicle tracked by the algorithm

To measure the angle of the test vehicle's frame and the distance of the test vehicle from the camera as a function of time using the grill data, the same approach described in Section 5.8.2 was employed. In addition, the position of two points on the bumper (Identified in Figure 5.119) were tracked and function of time. These tracked points were used to calculate the vehicle to camera distance as a function of time using the Equation below

$$d = \frac{S \cdot f}{\rho \cdot N} \quad (5.58)$$

where S is the true distance between the points (known), f is the focal length of the camera, ρ is the pixel pitch of the camera, and N is the measured distance between the points in pixels. The camera to vehicle distance from this method was averaged with the camera to vehicle distance calculated from the grill spacing method to reduce noise.




Figure 5.119: The tracked points on the bumper used to calculate the vehicle to camera distance

The acoustic data collected for each test run was used to calculate the engine speed using Fourier analysis. The audio data was collected by the PCB Piezotronics Free-Field Condenser microphone mounted on the truck's rear windshield at a rate of 52 kHz. This data was ingested into MATLAB as a pressure versus time signal and processed using the same method described in Section 5.8.2. The higher quality microphone resulted in a much cleaner signal the data collected by the Canon T2i camera. An example of the processed data can be seen in Figure 5.120.




Figure 5.120: From left-to-right the raw pressure signal, the pressure signal with a Hamming window applied, and the power spectra (during idle)

After the acoustic data was processed into an engine frequency versus time signal it needed to be synchronized with the data signals extracted from the high-speed video (frame twist, and position). This synchronization was achieved through frame by frame analysis of on-board video taken during each test run which captured the engine RPM and the vehicle speed from the vehicle's instrument panel. An example from of this video can be seen in Figure 5.121. This method of synchronization was both challenging and imprecise, it is highly recommended that any future testing utilizes a synchronized data acquisition system to mitigate this.



Figure 5.121: An example of the on-board video of the vehicle's instrument panel captured during each test run

Results

The data processing described in the previous section resulted in angular frame twist, position, and engine speed as a function of time for each test run. An example dataset can be seen in Figure 5.122.




Figure 5.122: From left-to-right; the angular frame twist, position, and engine speed of the unloaded test vehicle accelerating from a stop.

Prior to processing, all of the data sets were analyzed and it was determined that the high-speed

video did not capture the initial frame twist during several test runs. The frame of the vehicle begins twisting between 0.5 s and 1 s prior to any perceivable forward motion of the vehicle. This can be seen in the left-most two plots in Figure 5.122. As a result, several of the data sets were unusable. This left us with three datasets for the unloaded vehicle, two for the half-loaded vehicle, and one for the fully loaded vehicle.

The remaining suitable data was plugged into Equation 5.59 to solve for the mass to spring constant ratio of the test vehicle. Previously, we had solved for the ratio only at the point where velocity was at its peak. For this dataset, to increase accuracy we now solve for the ratio at each point in time and then use a least squares fit to determine the ratio. The values for r and θ_0 were estimated to be 0.8 m and 33° , respectively. The results can be seen in Figure 5.123.

$$\frac{m}{ks} = \frac{2 \cdot r^2 \left| \int_{t_0}^{t_f} \omega_E(t) \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) dt \right|}{x'(t_f)^2} \quad (5.59)$$




Figure 5.123: The calculated mass to spring constant ratio for each suitable test run versus the true test vehicle mass

Since the actual vehicle mass was known for each test run, we can use the mass to spring constant ratio to calculate the unknown spring constant k_s . The spring constant k_s was calculated to be 574,842 N/m. Scaling the data in Figure 5.123 by the calculated value for k_s we can get the calculated mass versus true mass plot seen in Figure 5.124.




Figure 5.124: The calculated mass for each suitable test run versus the true test vehicle mass

Conclusions

High-speed video data was used to measure the position and frame twist of the test vehicle at high spatial and temporal resolution using various image processing techniques. Acoustic data was used to accurately measure the engine speed of the test vehicle using Fourier based signal processing techniques. Synchronizing these three signals from two different data sources proved to be challenging and likely introduced significant error. The difficulty arose from the inconsistent relationship between the engine speed and the vehicle response stemming from the operator's use of the clutch and accelerator. The style in which the operator manipulates the clutch and accelerometer resulted in very different responses of frame twist, vehicle position, and engine speed. This makes it very difficult to identify persistent patterns for synchronization of the signals. To mitigate this, both the acoustic data and the video data should be taken using the same data acquisition system to ensure they are reliably synchronized.

Several of the data sets were deemed unusable because the high-speed video was started too

late and did not capture the initial frame twist of the vehicle. It is critical that the vehicle is measured while the frame is in a relaxed state; the frame twist at this time correlates to the zero torque condition, and all subsequent measurements must be made relative to it. For many of the data sets from this collect, the high-speed video was initiated after frame began to twist, but before the vehicle began moving forward (there is typically a 0.5 s to 1 s offset between when the frame first twist and the first measurable forward movement of the vehicle occurs). Due to this late initiation of the high-speed video, it was impossible to determine the amount of frame twist associated with zero torque. To mitigate this, the vehicle should be observed for a considerable time (2 s to 3 s) prior to acceleration in order to identify the zero torque frame twist. It should be noted, that the vehicle does not need to be completely stopped in order to identify the zero torque condition, simply not accelerating, or decelerating with an engine brake (standard braking would be acceptable).

The results seen in Figure 5.123 showed a significant reduction as compared to the processed data from the SRS 1 collect (results shown in Figure 5.114). The reduction was approximately a factor of three which suggests three times as much energy was required during the SRS 1 test as compared to the SRS 2 test to change the velocity of the test vehicle by the same amount. This discrepancy was investigated and it was determined that the SRS 1 test data was collected while the test vehicle was driving slightly uphill. Even a mild to moderate difference in incline of 3° - 5° would account for the discrepancy we observed between the two test sets. To account for this, future testing would need to be done on roads with an incline of zero degrees, or the incline would need to be factored into the calculation; this was done in the following equation

$$m = \frac{k_s \cdot r^2 \left| \int_{t_0}^{t_f} \omega_E(t) \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) dt \right|}{\left(\frac{1}{2} \cdot x'(t_f)^2 \right) + (-g \cdot (h(t_f) - h(t_0)))} \quad (5.60)$$

where g is the force due to gravity, $h(t_f)$ is the vehicle's elevation at time t_f , and $h(t_0)$ is the vehicle's elevation at time t_0 . This equation shows that a small change in elevation can require a significant amount of energy. For example, an increase in elevation of 5 cm would require the same amount of energy as a change in velocity of 1 m/s. This means that the spring constant value of 574,842 N/m is more accurate than the spring constant value calculated previously

(182,395 N/m). This is about ten times larger than the spring constant calculated for the tractor suspension for the suspension response phenomenology. However it is inappropriate to compare these spring constants directly as the frame twist phenomenology is impacted by the combined stiffness of the front and rear tractor suspensions, while the suspension response is only impacted by one or the other.

The results shown in Figure 5.124 show an excellent correlation between the calculated vehicle mass and the true vehicle load. The correlation coefficient was calculated to be 99.7%. This reflects a nearly perfect correlation between the measured and the true mass data and shows good agreement for test runs with the same vehicle loads.

The results of this test showed strong correlation and good repeatability. However there are multiple areas in which improvements could be made. First, it is critical that all of the measured signals are synchronized. The most accurate method of achieving this synchronization is by utilizing a single data acquisition system for all measurements (audio and video). Second, the data acquisition system must begin measuring the vehicle at least 2 s prior to any forward motion. This will ensure that the zero torque condition is captured. Finally, observation of the test vehicle should be extended. Increasing the measurement window to 10 s - 15 s should improve the overall effectiveness of the system and make the measurements more robust to noise and measurement error. This could be achieved by reducing the frame rate of the video acquisition system, as the 1200 frames-per-second rate of the Phantom camera proved to be faster than is required. A more reasonable 100 - 300 frames-per-second should provide more than adequate temporal resolution.

5.9 Combined Phenomenology and Error Analysis

The work in the previous sections has shown that the vehicle's suspension response to a speed bump, and the vehicle's frame twist response during acceleration were found to be the most effective methods to calculating the mass of the vehicle. In this section we will examine the effect of parameter estimation and measurement error on the suspension response and frame twist phenomenology and then combine the mass estimations from these two methods into a mass estimate that is more accurate than the mass estimate from either method on their own.

5.9.1 Suspension Response Error Analysis

The suspension response phenomenology estimates the mass and center of gravity of a vehicle's load by modeling the response of the vehicle's suspension as it is disturbed from equilibrium as a result of traveling over a speed-bump. This estimation is made by modifying the load parameters for the model until the simulated data best matches the measured suspension response of the vehicle. The required inputs to the model are; the spring constants and damping coefficients of the suspension, and the moment of inertia and center of mass of the unloaded vehicle. To analyze the impact of each of these parameters, data was simulated using the model for the passenger car. Data was simulated for a typical passenger sedan being driven over a speed bump with 200 kg in the trunk. The suspension and frame parameters used for the model can be seen below and the simulated data can be seen in Figure 5.125. From top-to-bottom the parameters are; front spring constant, front damping coefficient, rear spring constant, rear damping coefficient, unloaded vehicle mass, wheel base, center of gravity of the unloaded vehicle, moment of inertia of the unloaded vehicle, and velocity of the vehicle.

$$k_{s1} = 60,000 \text{ N/m}$$

$$k_{d1} = 3500 \text{ N s/m}$$

$$k_{s2} = 30,000 \text{ N/m}$$

$$k_{d2} = 3000 \text{ N s/m}$$

$$m_0 = 1430 \text{ kg}$$

$$d = 2.74 \text{ m}$$

$$x_{cg} = 1.4 \text{ m}$$

$$J = 2000 \text{ kg m}^2$$

$$V = 1.1 \text{ m/s}$$




Figure 5.125: The simulated suspension response data used to characterize the effect of error on the mass estimation

Using this simulated data, the mass of the load and the location of its center of gravity was calculated using the suspension response method described previously but with various degrees of erroneous parameters. Figure 5.126 shows the results of the suspension response calculation using the correct system parameters, with the best 0.1 % highlighted in green. This green region represents the suspension response estimation for the mass and center of gravity of the vehicle load. If the true value falls within this region the calculation would be considered correct.




Figure 5.126: Average error heat map for the simulated data calculated with the true parameter values

The amount of error was varied to find the maximum error for the front spring constant (k_{s1}) that would still allow the calculated mass and center of gravity window to contain the true value. We found that the total front spring constant error could exceed $\pm 70\%$ of the true value and still allow the truth data to fall within the 0.1% window as seen in 5.127. This robustness to error for the front spring constant parameter is because the load is being added to the trunk (far from the front suspension). This would not be the case for a load being added to the front of the vehicle.




Figure 5.127: Average error heat map for the front spring constant k_{s1} with -70% error (left) and 70% error (right)

The amount of error was varied to find the maximum error for the front damping coefficient (k_{d1}) that would still allow the calculated mass and center of gravity window to contain the true value. We found that if the total front damping coefficient error did not exceed $\pm 28\%$ of the true value, the 0.1% window would contain the true mass and center of gravity as seen in 5.128. This robustness to error for the front damping coefficient parameter is because the load is being added to the trunk (far from the front suspension). This would not be the case for a load being added to the front of the vehicle.




Figure 5.128: Average error heat map for the front spring constant k_{d1} with -28% error (left) and 28% error (right)

The amount of error was varied to find the maximum allowable error for the rear spring constant (k_{s2}) that would still allow the calculated mass and center of gravity window to contain the true value. We found that if the total rear spring constant error did not exceed $\pm 5\%$ of the true value, the 0.1% window would contain the true mass and center of gravity as seen in 5.129. This sensitivity was expected, and should be applied to the front spring constant as well to allow for accurate calculation of masses near the front of the vehicle. The data also shows the expected relationship between the error and the estimated mass and center of gravity window. When the rear spring constant value was too low, the window moved towards lower mass and the center of gravity more towards the front of the vehicle. When the rear spring constant value was too high the reverse was true.




Figure 5.129: Average error heat map for the rear spring constant (k_{s2}) with -5% error (left) and 5% error (right)

The amount of error was varied to find the maximum error for the rear damping coefficient (k_{d2}) that would still allow the calculated mass and center of gravity window to contain the true value. We found that if the total rear damping coefficient error did not exceed $\pm 5\%$ of the true value, the 0.1% window would contain the true mass and center of gravity as seen in 5.130. This sensitivity was expected, and should be applied to the front damping coefficient as well to allow for accurate calculation of masses near the front of the vehicle. When the rear damping coefficient value was too low, the window moved towards lower mass and the center of gravity remained nearly constant. When the rear damping coefficient value was too high the reverse was true.





Figure 5.130: Average error heat map for the rear damping coefficient (k_{d2}) with -5% error (left) and 5% error (right)

The amount of error was varied to find the maximum error for the moment of inertia of the unloaded vehicle (J_0) that would still allow the calculated mass and center of gravity window to contain the true value. We found that if the total error did not exceed $\pm 3\%$ of the true value, the 0.1% window would contain the true mass and center of gravity as seen in 5.131. When the moment of inertia was too low, the window moved towards the lower mass and the center of gravity moved towards the rear of the vehicle. When the moment of inertia was too high the reverse was true.

Table 5.15: Recommended maximum error values for the suspension response estimated parameters

	Acceptable Error (%)	Acceptable Error
ks1	± 5	± 3000 N/m
kd1	± 5	± 175 N s/m
ks2	± 5	± 1500 N/m
kd2	± 5	± 225 N s/m
J0	± 3	± 150 kg m ²

Figure 5.131: Average error heat map for the moment of inertia of the unloaded vehicle (J_0) with -3% error (left) and 3% error (right)

Based on these simulated results we can provide guidelines for acceptable error in the estimation of the parameters required for the suspension response phenomenology. Table 5.15 shows the acceptable error (as a percent of the true value) and the actual range of values (based on the simulated passenger car data).

The estimated parameters are used to generate simulated suspension response data, which must be compared to measured suspension response data in order to estimate the mass and center of gravity of the vehicle's load. The measured suspension response data, describes the position of the vehicle's frame relative to its axles. To measure this, we track the position of the wheel centers, and a position on the frame directly above each wheel. For all of the test collects, this

data was tracked using an image-based measurement system. Now that substantial modeling has been done, we will try to define the minimum collection requirements of the measurement system needed to accurately calculate the mass and center of gravity of a vehicle. Based on the two vehicles modeled during this study, a typical passenger vehicle (Honda Accord) and a tractor-trailer, the measurement system needs to observe the front-most and rear-most axles of the vehicle simultaneously from a time just prior to the vehicle reaching the speed-bump until approximately 3 s after the last axle travels over the speed bump (in both models most of the suspension motion is damped out in less than 3 s). For a vehicle traveling 2 m/s this works out to a field-of-view of 6 m plus twice the vehicle length. So for the Honda Accord, the ideal field-of-view would be approximately 11.5 m. For the tractor-trailer, the ideal field-of-view would be approximately 41 m for the entire tractor-trailer, or 32 meter for the trailer only. The minimum measurement resolution required (based on the vehicles modeled during this study) is approximately 2 mm for the passenger vehicle and approximately 4 mm for the tractor-trailer and the minimum sampling rate is approximately 30 frames-per-second. Since the suspension response mass and center of gravity calculation relies on a comparison of modeled and measured data using a least square fit method it is fairly robust to moderate bipolar noise (on the order of ± 0.5 mm to 1 mm). This means that the minimum requirements for the passenger vehicle could be met with a single, high resolution camera. For example; a single Canon EOS 5D Mark III with a fixed 24 mm focal length positioned approximately 7.7 m from the path of the vehicle would have a field-of-view of 11.5 m and a spatial resolution on the vehicle of 2 mm. A diagram of the configuration can be seen in Figure 5.132. If the same camera was used, the tractor-trailer would require two synchronized cameras, positioned approximately 15.3 m from the path of the vehicle to achieve the desired field-of-view and spatial resolution of 41 m and 4 mm, respectively. The cameras would need to be precisely aligned to ensure parallel optical axes and an equal road-to-camera distance, or the system would need to be calibrated in place. A diagram of this configuration can be seen in Figure 5.133.




Figure 5.132: Diagram of the observation geometry meeting the minimum spatial resolution and field-of-view requirements for a passenger vehicle




Figure 5.133: Diagram of the observation geometry meeting the minimum spatial resolution and field-of-view requirements for a tractor-trailer

5.9.2 Frame Twist Error Analysis

The frame twist phenomenology works by measuring the power generated by the engine, and the position of the vehicle and uses these measurements to calculate the mass of the vehicle by assuming all of the power generated by the engine is converted into the vehicle's kinetic energy. To calculate the amount of power generated by the engine, several characteristics of the vehicle must be known. Specifically, the conversion factor from frame twist angle to engine torque (Equation 5.61). This conversion factor is linearly dependent on the total spring constant (k_s), and non-linearly dependent on the geometry parameters that describe the position of the engine's axis of rotation relative to the suspension springs (r and θ_0). First, we will examine how uncertainty in each individual parameter effects the uncertainty of the calculated torque, and then we will combine these uncertainties to into a single expression for the total uncertainty in the calculate torque. Uncertainty in the spring constant (Δk_s) results in the uncertainty in the engine torque ($\Delta \tau_E$) as shown in Equation 5.62. This is a linear relationship, as a 5 % error in spring constant would result in a 5 % error in calculated torque. Similarly, we can express the uncertainty in the calculated engine torque ($\Delta \tau_E$) caused by uncertainty in the geometry parameters (Δr and $\Delta \theta_0$) using Equations 5.63 and 5.64, respectively. Assuming non-correlated error, we can combine these expression to find the total uncertainty in the calculated engine torque from all three parameters (Δk_s , Δr , and $\Delta \theta_0$) using the standard propagation of error equation as shown in Equation 5.65.

$$\tau_E(t) = -k_s \cdot r^2 \cdot \cos(\theta(t) + \theta_0) \cdot \tan \theta(t) \quad (5.61)$$

$$\Delta \tau_E = r^2 \cdot \cos(\theta(t) + \theta_0) \cdot \tan(\theta(t)) \cdot \Delta k_s \quad (5.62)$$

$$\Delta \tau_E = 2 \cdot k_s \cdot r \cdot \cos(\theta(t) + \theta_0) \cdot \tan(\theta(t)) \cdot \Delta r \quad (5.63)$$

$$\Delta \tau_E = k_s \cdot r^2 \cdot \sin(\theta(t) + \theta_0) \cdot \tan(\theta(t)) \cdot \Delta \theta_0 \quad (5.64)$$

$$\begin{aligned}
\Delta\tau_E^2 = & \left(r^2 \cdot \cos(\theta(t) + \theta_0) \cdot \tan(\theta(t)) \cdot \Delta k_s \right)^2 + \\
& \left(2 \cdot k_s \cdot r \cdot \cos(\theta(t) + \theta_0) \cdot \tan(\theta(t)) \cdot \Delta r \right)^2 + \\
& \left(k_s \cdot r^2 \cdot \sin(\theta(t) + \theta_0) \cdot \tan(\theta(t)) \cdot \Delta\theta_0 \right)^2
\end{aligned} \tag{5.65}$$

The to calculate the mass of a vehicle using the frame twist phenomenology the position, angular frame twist, and engine speed of the vehicle need to be measured precisely. Now that substantial amounts of data have been collected and analyzed, we will try to define the minimum collection requirements for the measurement system to accurately calculate the vehicle mass. The frame twist measurement of a tractor-trailer requires a minimum resolution of approximately 0.05° and a sampling rate of approximately 300 frames-per-second. These measurements need to be made over a time period starting at least 3 seconds before acceleration begins and ending after the vehicle has traveled approximately 6 m. The measurement of the distance between the camera and the vehicle requires a minimum resolution of approximately 5 mm. This precision is required to allow for accurate calculation of the numerical derivative needed to calculate the velocity of the vehicle. An imaging system can be used to measure the vehicle's frame twist and position by tracking points on the vehicle and calculating the angle and distance between them as a function of time. To achieve the required level of resolution using points on opposite sides of the vehicle (typical vehicle width of approximately 2 m) the spatial resolution would need to be at least 0.7 mm. The diagram of an acceptable configuration can be seen in Figure 5.134. Note that the vehicle fully fills the field-of-view after traveling the 6 m. This configuration would have a spatial resolution of 0.7 mm when the vehicle is 11.8 m away and 0.34 mm when the vehicle is 5.8 m from the camera. This translates to an angular frame twist resolution of 0.02° when the vehicle is 11.8 m away and 0.01° when the vehicle is 5.8 m away. The camera to vehicle distance measurement would have a resolution of 4.2 mm when the vehicle is 11.8 m away and 1 mm when the vehicle is 5.8 m from the camera. This imaging system would require an aperture diameter of no greater than 7.8 mm to ensure a depth of field that would span the 6 m of vehicle travel at the distances of stated.




Figure 5.134: Diagram of the observation geometry meeting the minimum spatial resolution and field-of-view for a tractor-trailer. The left diagram shows the vehicle prior to acceleration and the right diagram shows the vehicle after traveling 6 m.

Acoustic measurements are needed to calculate the engine speed as a function of time. All of the data analyzed suggests that a standard off-the-shelf microphone sampling at the typical 52 kHz rate is sufficient for the engine speed calculation.

5.9.3 Combined Phenomenology

The suspension response and the frame twist phenomenologies proved best suited for the remote measurement of vehicle mass. Separately they have been shown to accurately estimate the mass of a vehicle when provided with sufficient characteristics of the vehicle, and accurate measurements of the pertinent vehicle signatures. However, uncertainty in the mass estimation results from uncertainty in the estimation of the vehicle characteristics. Combining the results from the two phenomenologies should result in a smaller uncertainty.

The suspension response phenomenology generates a two-dimensional heat map with vehicle load mass on the x-axis, center of gravity of the load on the y-axis and model fit error is represented by the color scale. An example of this is shown in Figure 5.135. The green region in the center shows the area where we have high confidence that the true vehicle mass and center of gravity lies and the **X** in the center is the true value.




Figure 5.135: The simulated suspension response data used to characterize the effect of error on the mass estimation

The frame twist phenomenology estimates the total vehicle mass. From this, the mass of the load can be calculated by simply subtracting the unloaded vehicle mass from the estimated total mass. This estimation for the mass of the load (with uncertainty) can be used to enhance the results from the suspension response test by excluding masses that do not agree with the frame

twist results. Figure 5.136 shows the suspension response data with vertical lines added which represent the frame twist estimate of the vehicle mass. Figure 5.137 shows the same plots with the green region modified to only include the load estimations that agree with both the suspension response and the frame twist phenomenology.




Figure 5.136: The simulated suspension response data used to characterize the effect of error on the mass estimation with the frame twist error bars added




Figure 5.137: The simulated suspension response data used to characterize the effect of error on the mass estimation with regions excluded based on the frame twist analysis

In conclusion, the suspension response and frame twist phenomenologies have been shown to both be viable methods of measuring the mass of a vehicle's load. Each method calculates a vehicle load mass with an uncertainty, giving a range of values in which the true vehicle mass is likely to lie. Combining the results from the suspension response and frame twist phenomenologies can reduce the size of the range, giving a narrower mass estimation window.

Chapter 6

Conclusions

The research presented here has thoroughly investigated several phenomenologies as methods of estimating the mass of a vehicle remotely. Brake temperature, engine acoustics, exhaust output, tire temperature, vehicle-induced ground vibration, tire deformation, suspension response, and engine torque-induced frame twist were all investigated. In this Chapter we will summarize the work done for each phenomenology, list the major accomplishments of this research, and describe recommendations for future research.

6.1 Brake Temperature

The brake temperature phenomenology was considered as a potential method of remotely estimating vehicle mass. When a vehicle decelerates to a stop using a typical braking system, all of the vehicle's kinetic energy is converted into thermal energy. This thermal energy is localized to the braking system, causing a dramatic increase in temperature. Research into the brake temperature phenomenology took the form of two field tests, each employing on-board LWIR imagers precisely oriented to observe the brake components. The results of these tests showed that latent thermal energy from previous instances of braking, coupled with a very small and typically obscured area of interest, made the brake temperature phenomenology non-viable for estimating vehicle mass.

6.2 Engine Acoustics

The engine acoustic phenomenology was considered as a potential method of remotely estimating vehicle mass. The acoustic signature of an engine is dependent on, among other things, its size and its speed of operation. If the size and average cylinder pressure of an engine are known, and the speed of the engine measured, it could be possible to estimate the power output of the engine as a function of time (if the vehicle's transmission is fully engaged). Research into the engine acoustic phenomenology took the form of two field tests, each employing microphones positioned to observe the engine acoustic signature as the vehicle was operating. The results of these tests showed that it was possible to remotely measure the speed of the engine, but no method was found to estimate the engine size or average cylinder pressure. As a result, the engine acoustic method alone was deemed non-viable for estimating vehicle mass. However, the ability to measure the engine speed remotely proved critically important to the engine torque-induced frame twist phenomenology's positive results.

6.3 Engine Exhaust

The engine exhaust output phenomenology was considered as a potential method of remotely estimating vehicle mass. The temperature of the exhaust system, the plume generated by the exhaust system, and the impact of exhaust braking were all investigated as potential methods of estimating vehicle mass, or methods of providing supplementary information to a different mass estimation technique. The temperature of the exhaust system should be correlated with the amount of heat the engine is generating, and the heat production of the engine is correlated to the power output of the engine (typical combustion engines are typically only 11 - 26 percent efficient). Research into the exhaust temperature phenomenology took the form of small tests where LWIR images were captured of vehicles during operation. The results of these tests showed a huge variation in exhaust system emissivity, and styles, as well as a very slow rate of heat transfer from the engine to the exhaust system. As a result, the engine/exhaust temperature phenomenology was deemed non-viable for estimating vehicle mass. To exploit exhaust plume for mass estimation, the volume of the plume would need to be measured, and then compared to changes in the kinetic

energy of the vehicle. Due to the composition of the exhaust plume (all gases are present in large quantities in the atmosphere) measuring the volume of exhaust being produced by a vehicle did not seem plausible. Research into the exhaust braking phenomenology identified a method of determining if a Jacob's-style exhaust brake was actively being used by the vehicle, which could be useful as a supplementary signal in a conservation of energy based mass estimations.

6.4 Tire Temperature

The tire temperature phenomenology was considered as a potential method of remotely estimating vehicle mass. When the vehicle is moving, the continuous loading and unloading of the tire material results in the generation of thermal energy. Since the amount of tire surface in contact with the ground at any given moment is directly related to the mass of the vehicle, the temperature of the tires of a moving vehicle may be correlated to its mass. Research into the tire temperature phenomenology took the form of one field test, which employed on-board LWIR imagers precisely oriented to observe the tires. The results of this test showed that the amount of thermal energy generated by the tires is too small to be observed above the other sources of thermal energy experienced by the tires. As a result, the tire temperature phenomenology was deemed non-viable for estimating vehicle mass.

6.5 Vehicle-Induced Ground Vibration

The vehicle-induced ground vibration phenomenology was considered as a potential method of remotely estimating vehicle mass. When a vehicle travels over a road surface, vibrations are induced by its travel. The amount of vibrational energy imparted by the vehicle to the ground should be correlated to the mass of the vehicle. Research into the vehicle-induced ground vibration phenomenology took the form of one field test, which employed triaxial accelerometers mounted to the top surface of a set of railroad tracks. The results of the testing were inconclusive, a strong correlation was measured between the peak energy observed, and the kinetic energy of the vehicle, but this correlation may have been caused by the unintentional correlation between vehicle mass and distance from the accelerometer at the point of impact. It was also observed that the position

of the accelerometer relative to the railroad track support structure (ties and spikes) had a strong impact on the observed induced acceleration. More research into this phenomenology is required to determine if it is a viable method of estimating vehicle mass.

6.6 Tire Deformation

The tire deformation phenomenology was considered as a potential method of remotely estimating vehicle mass. The entire mass of a vehicle is supported by the tires. To support the load of the vehicle, the tires deform to increase the area of the region in contact with the ground. If the total area of the tires in contact with the road is measured, and the internal pressure of the tires is known, the mass of the vehicle can be calculated. Research into the tire deformation phenomenology took the form of multiple field tests employing, individually, LiDAR, and high-definition cameras. The results of this testing showed a strong correlation between tire deformation and the mass of the vehicle. In order to advance this phenomenology to the point where an estimation of mass could be made from remotely measured tire deformation, a tire model needs to be developed. Due to the complex structure of the tire carcass and internal belts, this model will likely need to be generated using finite element analysis (FEA). More research into this phenomenology is required to determine if it is a viable method of estimating vehicle mass.

6.7 Suspension Response

The suspension response phenomenology was considered as a potential method of remotely estimating vehicle mass. The suspension system for most vehicles can be thought of as a set of damped springs which couple the vehicle's frame to its axles. When the vehicle is traveling over flat ground, the suspension remains in a state of static equilibrium. However, when the vehicle travels over an impediment, like a speed bump, the suspension responds. This response is dependent on the mass and moment of inertia of the vehicle, as well as the suspension parameters. Research into the suspension response phenomenology took the form of two field tests, each employing video cameras which captured the position of the vehicle as it traveled over a speed bump. The results of these tests showed a strong correlation of approximately 90 percent between the

measured suspension response and vehicle mass. Because of these positive results, a physics-based model was built to simulate the suspension response of a vehicle given an input signal (the road profile) and a set of vehicle parameters; spring constants, damping coefficients, vehicle mass, vehicle center of gravity, and the vehicle's moment of inertia. This physics-based model was used to develop an algorithm that, when given a suspension response and road profile, can solve for the mass and center of gravity of the vehicle, given that the suspension parameters (spring constants and damping coefficients), and unloaded vehicle parameters (mass, moment of inertia, and center of gravity) are known. A third field test was executed to validate this model. The results of this field test were very positive, and showed that the method was accurate to within 10 percent, if all the required parameters were known accurately. Finally, an uncertainty analysis was carried out to quantify how the expected error in the input parameters and suspension response measurement affected the uncertainty in the vehicle mass estimation. Expected error in each of the vehicle parameters had roughly the same relative impact on the uncertainty in the vehicle mass estimation, with the exception of the moment of inertia of the vehicle which was found to have slightly greater impact than the others.

6.8 Engine Torque-Induced Frame Twist

The engine torque-induced frame twist phenomenology was considered as a potential method of remotely estimating vehicle mass. For a vehicle to accelerate, rotational energy must be transferred from the vehicle's engine to its drive wheels. The amount of power transferred from the engine to the vehicle's drive-train can be calculated from the engine torque applied to the drive shaft, and the engine's speed. If the power applied by the engine can be measured during a period of vehicle acceleration, and the velocity of the vehicle can be measured over the same period, conservation of energy can be invoked to calculate the mass of the vehicle. A physics-based model was built to describe how the vehicle's suspension responds to the engine torque causing the vehicle's frame to twist relative to its axles. Using this model, the engine torque could be calculated from the amount of frame twist, and knowledge of the vehicle's suspension parameters and the engine's position relative to the suspension. A method was developed that utilized this model to calculate vehicle mass from measured frame twist, engine speed, and vehicle position over a

period of vehicle acceleration. A field test was carried out to validate the model. The results of this field test were very positive and showed a correlation to vehicle mass of 97.7 percent, and an agreement with the true vehicle mass to within 10 percent. A second field test was carried out with improved instrumentation; a Vision Research Phantom v5.1 high-speed video camera, and a PCB Piezotronics Free-Field Condenser microphone. The results of this test were inconclusive. Many of the collected datasets had to be discarded due to late triggering of the high-speed video camera leaving only five usable datasets. The five usable datasets showed an excellent correlation to vehicle mass of 99.7 percent and an agreement with the true vehicle mass that was better than 10 percent. However, more datasets are needed to confirm that the error was truly reduced. Finally, an error analysis was completed to quantify how uncertainty in the input parameters, frame twist, engine speed and vehicle velocity measurements affected the uncertainty in the vehicle mass estimation. Expected error in each of the vehicle parameters had roughly the same relative impact on the uncertainty in the vehicle mass estimation, with the exception of the distance between the engine axis of rotation and the vehicle's suspension which was found to have slightly greater impact than the others.

6.9 Combined Phenomenology

The suspension response and frame twist phenomenologies were both shown to be viable methods of measuring the mass of a vehicle. Each method produces an estimation for the mass of the vehicle with an uncertainty which is dependent on the expected error in the parameter estimations, as well as the measurement error. Combining the estimated mass and uncertainties from each phenomenology can result in a mass estimation with a reduced uncertainty. A method was described that uses the mass estimation uncertainty value from the frame twist phenomenology to restrict the vehicle mass and center of gravity estimation from the suspension response phenomenology. This restricted mass estimation was shown to have less uncertainty than either of the phenomenologies produced alone.

6.10 Major Accomplishments

Multiple original contributions to science emerged during the course of this research. A unique video tracking algorithm was developed to meet a specific requirement of the suspension response phenomenology. The algorithm was designed to accurately track the position of a specific region on a moving vehicle as it traveled through the field-of-view of a video camera, and be robust to changing illumination conditions and slowly varying changes in scale and rotation. Also arising from the suspension response phenomenology was an algorithm to fit the measured suspension response data from the front and rear suspensions of a vehicle to a unique system of second-order differential equations which described the response of a coupled damped-harmonic oscillator to an input impulse signal (speed bump). The output of this algorithm is the mass and center of gravity of the load being carried by the vehicle, given that the suspension parameters, unloaded mass, unloaded center of gravity, and unloaded moment of inertia of the vehicle are known. Research into the frame twist phenomenology resulted in a unique method of calculating the speed of an engine using Fourier analysis of acoustic data measured on or near the vehicle. The method exploited the strong harmonics present in the engine's acoustic signal, which allows it to be robust to background noise, echoes, and other issues associated with acoustic measurements. From the suspension response and the frame twist phenomenology research, a set of optimal collection requirements were defined that provide guidance into accurate measurement of both phenomenologies.

6.11 Future Work

This research investigated several phenomenologies as potential methods of remotely calculating the mass of a vehicle. The majority of the phenomenologies investigated were found to be poorly suited for remote vehicle mass estimation and research into them was ceased to focus on the more promising phenomenologies. For the phenomenologies deemed promising for remote vehicle mass estimation, there are improvements that can be made. In the following subsections we will briefly provide recommendations for future research into the tire deformation, suspension response, and frame twist phenomenologies.

6.11.1 Tire Deformation

The tire deformation phenomenology was shown to be strongly correlated to vehicle mass. The next step in continuing this research would be to develop a tire model that can accurately predict the shape of a tire under different loading conditions and internal pressures. Due to the complex construction and shape of a typical tire, an FEA model is recommended. Once an accurate model has been created, a measured tire's shape could be compared to a set of simulated tire shapes to determine the loading and internal pressure that most closely matches. LiDAR is recommended as the best method for remotely measuring the shape of the tire, as it has been shown to be effective for measuring a tire in profile.

6.11.2 Suspension Response

The suspension response phenomenology was shown to be an effective method of remotely measuring the mass of a vehicle. The next step in continuing this research would be to measure all of the suspension response related parameters (spring constants, damping coefficients, unloaded frame mass, and unloaded frame center of gravity) for a vehicle, and then repeat the suspension response field test using the measurement setup described in Section 5.9.1. Positive results from this test would validate the suspension response physics-based model. Additional work for the suspension response phenomenology would be to study the variability in the suspension response related parameters. Of particular interest would be the variability present in vehicles of the same class (e.g. mid-size sedans). Small in-class variability would allow all vehicles of that class to be characterized by a single parameter value (instead of different values for each make and model).

6.11.3 Frame Twist

The frame twist phenomenology was shown to be an effective method of remotely measuring the mass of a vehicle. The next step in continuing this research would be to measure all of the frame twist-related parameters (spring constants, damping coefficients, and engine geometry). The results of these measurements, combined with the physics-based model described in Section 5.8.1 could be used to predict the amount of frame twist that would result from a given engine torque.

This could then be verified by applying a known torque to the frame, and measuring the frame twist response. Additional work for the frame twist phenomenology would be to study the variability in the frame twist-related parameters. Of particular interest would be the variability present in vehicles of the same class. Small in-class variability would allow all vehicles of that class to be characterized by a single parameter value.

Appendix A

Suspension Response Model Derivation

```
[ reset() :
```

The position and orientation of the vehicle frame can be described completely with two terms:

-> The vertical position of the frame where the front suspension joins the frame ($y_{F1}(t)$)

-> The angle, relative to horizontal, formed by a line passing through $y_{F1}(t)$ and the position on the frame where the rear suspension joins the frame $y_{F2}(t)$ ($\theta(t)$)

Define the vertical position of the frame's center of gravity, $y_{cg}(t)$, and the position on the frame where the rear suspension joins the frame, $y_{F2}(t)$, in terms of $y_{F1}(t)$, $\theta(t)$, the position of the center of gravity, x_{cg} , and the distance between the axles, d :

```
[ Eq1 := y[cg](t) = y[F1](t) + x[cg]*Symbol::theta(t);
  Eq2 := y[F2](t) = y[F1](t) +(Symbol::theta(t)*d);
   $y_{cg}(t) = y_{F1}(t) + \theta(t) x_{cg}$ 
   $y_{F2}(t) = y_{F1}(t) + d \theta(t)$ 
```

Define the three forces acting on the vehicle frame F_1 , F_2 , and F_g in terms of the suspension parameters (k_{s1} , k_{s2} , k_{d1} , k_{d2}):

-> F_1 is force acting on the frame from the front suspension

-> F_2 is force acting on the frame from the rear suspension

-> F_g is force acting on the frame from gravity

-> k_{s1} and k_{d1} are the spring constant and damping coefficient of the front suspension respectively

-> k_{s2} and k_{d2} are the spring constant and damping coefficient of the rear suspension respectively

```
[ Eq3 := F[1] = k[s1]*(y[A1](t)-y[F1](t)) + k[d1]*(diff(y[A1](t),t) -
  diff(y[F1](t),t));
  Eq4 := F[2] = k[s2]*(y[A2](t)-y[F2](t)) + k[d2]*(diff(y[A2](t),t) -
  diff(y[F2](t),t));
  Eq5 := F[g] = m*g;
   $F_1 = k_{d1} \left( \frac{\partial}{\partial t} y_{A1}(t) - \frac{\partial}{\partial t} y_{F1}(t) \right) + k_{s1} (y_{A1}(t) - y_{F1}(t))$ 
   $F_2 = k_{d2} \left( \frac{\partial}{\partial t} y_{A2}(t) - \frac{\partial}{\partial t} y_{F2}(t) \right) + k_{s2} (y_{A2}(t) - y_{F2}(t))$ 
   $F_g = g m$ 
```

Define the three torques acting on the vehicle frame τ_1 , τ_2 , and τ_g in terms of the suspension parameters (k_{s1} , k_{s2} , k_{d1} , k_{d2}):

-> τ_1 is torque acting on the frame from the front suspension

-> τ_2 is torque acting on the frame from the rear suspension

-> τ_g is torque acting on the frame from gravity

```
[ Eq6 := Symbol::tau[1] = -F[1]*x[cg];
  Eq7 := Symbol::tau[2] = F[2]*(d-x[cg]);
  Eq8 := Symbol::tau[g] = 0;
   $\tau_1 = -x_{cg} F_1$ 
   $\tau_2 = (d - x_{cg}) F_2$ 
   $\tau_g = 0$ 
```

Define the balance of forces equation in terms of the vertical acceleration of the frame's center of gravity

```
[ Eq9 := diff(diff(y[cg](t),t),t)*m = F[1] + F[2] + F[g];
```

$$m \frac{\partial^2}{\partial t^2} y_{cg}(t) = F_g + F_1 + F_2$$

Express the balance of forces equation in terms of the system parameters by substituting Eq1, Eq3, Eq4, and, Eq5 into Eq 9

```
Eq10 := collect(simplify(subs(Eq9, Eq1, Eq3, Eq4, Eq5, Eq2)),
[k[s1], k[s2], k[d1], k[d2]]);
```

$$m \frac{\partial^2}{\partial t^2} (y_{F1}(t) + \theta(t) x_{cg}) = (y_{A1}(t) - y_{F1}(t)) k_{s1} + (y_{A2}(t) - y_{F1}(t) - d \theta(t)) k_{s2} \\ + \left(\frac{\partial}{\partial t} y_{A1}(t) - \frac{\partial}{\partial t} y_{F1}(t) \right) k_{d1} + \left(\frac{\partial}{\partial t} y_{A2}(t) - \frac{\partial}{\partial t} (y_{F1}(t) + d \theta(t)) \right) k_{d2} + g m$$

Define the balance of torques equation in terms of the angular acceleration about the frame's center of gravity

```
Eq11 := diff(diff(Symbol::theta(t), t), t) * J =
Symbol::tau[1] + Symbol::tau[2];
```

$$J \frac{\partial^2}{\partial t^2} \theta(t) = \tau_1 + \tau_2$$

Express the balance of torques equation in terms of the system parameters by substituting Eq1, Eq3, Eq4, and, Eq5 into Eq 11

```
Eq12 := simplify(subs(Eq11, Eq6, Eq7, Eq4, Eq3, Eq2));
```

$$J \frac{\partial^2}{\partial t^2} \theta(t) = -x_{cg} \left(k_{d1} \left(\frac{\partial}{\partial t} y_{A1}(t) - \frac{\partial}{\partial t} y_{F1}(t) \right) + k_{s1} (y_{A1}(t) - y_{F1}(t)) \right) \\ - \left(k_{s2} (y_{F1}(t) - y_{A2}(t) + d \theta(t)) + k_{d2} \left(\frac{\partial}{\partial t} (y_{F1}(t) + d \theta(t)) - \frac{\partial}{\partial t} y_{A2}(t) \right) \right) (d - x_{cg})$$

Find the static equilibrium position and orientation

```
Eq13 := subs(Eq10, diff(diff(y[F1]
(t), t), t)=0, diff(diff(Symbol::theta(t), t), t)=0, diff(y[F1]
(t), t)=0, diff(Symbol::theta(t), t)=0, diff(y[A2](t), t)=0, diff(y[A1]
(t), t)=0, y[A1](t)=0, y[A2](t)=0):
Eq14 := subs(Eq12, diff(diff(y[F1]
(t), t), t)=0, diff(diff(Symbol::theta(t), t), t)=0, diff(y[F1]
(t), t)=0, diff(Symbol::theta(t), t)=0, diff(y[A2](t), t)=0, diff(y[A1]
(t), t)=0, y[A1](t)=0, y[A2](t)=0):
Eq15 := solve([Eq13, Eq14], [y[F1]
(t), Symbol::theta(t)], IgnoreSpecialCases):
Eq16 := y[F1Static] = Eq15[1][1][2];
Eq17 := Symbol::theta[Static] = Eq15[1][2][2];
```

$$y_{F1Static} = \frac{g m (d - x_{cg})}{d k_{s1}}$$

$$\theta_{Static} = \frac{g m (x_{cg} k_{s1} - d k_{s2} + x_{cg} k_{s2})}{d^2 k_{s1} k_{s2}}$$

Redefine $y_{F1}(t)$, and $\theta(t)$ in terms of the measurable position and orientation relative to static

equilibrium ($y_{M1}(t)$ and $\theta_M(t)$):

```
Eq18 := y[F1](t) = y[M1](t) + y[F1Static];
Eq19 := Symbol::theta(t) = Symbol::theta[M](t) + Symbol::theta[Static];
y_F1(t) = y_M1(t) + y_F1Static
theta(t) = theta_M(t) + theta_Static
```

Express the full balance of forces and balance of torques equations terms of the measurable position and orientation relative to static equilibrium ($y_{M1}(t)$ and $\theta_M(t)$):

```
Eq20 := collect(simplify(subs(Eq10, Eq18, Eq19, Eq16, Eq17)),
[k[s1], k[s2], k[d1], k[d2]]);
Eq21 := simplify(expand(subs(Eq12, Eq18, Eq19, Eq16, Eq17)));
```

$$m \left(x_{cg} \frac{\partial^2}{\partial t^2} \theta_M(t) + \frac{\partial^2}{\partial t^2} y_{M1}(t) \right) = (y_{A1}(t) - y_{M1}(t)) k_{s1} + (y_{A2}(t) - y_{M1}(t) - d \theta_M(t)) k_{s2}$$

$$+ \left(\frac{\partial}{\partial t} y_{A1}(t) - \frac{\partial}{\partial t} y_{M1}(t) \right) k_{d1} + \left(\frac{\partial}{\partial t} y_{A2}(t) - \frac{\partial}{\partial t} y_{M1}(t) - d \frac{\partial}{\partial t} \theta_M(t) \right) k_{d2}$$

$$J \frac{\partial^2}{\partial t^2} \theta_M(t) = k_{d1} x_{cg} \frac{\partial}{\partial t} y_{M1}(t) - k_{d2} x_{cg} \frac{\partial}{\partial t} y_{A2}(t) - k_{d1} x_{cg} \frac{\partial}{\partial t} y_{A1}(t) + k_{d2} x_{cg} \frac{\partial}{\partial t} y_{M1}(t)$$

$$+ d k_{d2} \frac{\partial}{\partial t} y_{A2}(t) - d k_{d2} \frac{\partial}{\partial t} y_{M1}(t) - y_{A1}(t) x_{cg} k_{s1} - y_{A2}(t) x_{cg} k_{s2} - d^2 k_{d2} \frac{\partial}{\partial t} \theta_M(t) + y_{M1}(t) x_{cg} k_{s1}$$

$$+ y_{M1}(t) x_{cg} k_{s2} + d y_{A2}(t) k_{s2} - d y_{M1}(t) k_{s2} - d^2 \theta_M(t) k_{s2} + d k_{d2} x_{cg} \frac{\partial}{\partial t} \theta_M(t) + d \theta_M(t) x_{cg} k_{s2}$$

Solve the full balance of torques equation (Eq21) for the angular acceleration relative to the static equilibrium (double derivative of $\theta_M(t)$)

```
Eq22 := simplify(solve(Eq21, diff(diff(Symbol::theta[M]
(t), t), t), IgnoreSpecialCases)):
Eq23 := diff(diff(Symbol::theta[M](t), t), t) = simplify(expand(Eq22[1]));
```

$$\frac{\partial^2}{\partial t^2} \theta_M(t) = \frac{d k_{d2} \frac{\partial}{\partial t} y_{A2}(t)}{J} - \frac{d k_{d2} \frac{\partial}{\partial t} y_{M1}(t)}{J} - \frac{y_{A1}(t) x_{cg} k_{s1}}{J} - \frac{y_{A2}(t) x_{cg} k_{s2}}{J} - \frac{d^2 k_{d2} \frac{\partial}{\partial t} \theta_M(t)}{J}$$

$$+ \frac{y_{M1}(t) x_{cg} k_{s1}}{J} + \frac{y_{M1}(t) x_{cg} k_{s2}}{J} + \frac{d y_{A2}(t) k_{s2}}{J} - \frac{d y_{M1}(t) k_{s2}}{J} - \frac{d^2 \theta_M(t) k_{s2}}{J} - \frac{k_{d1} x_{cg} \frac{\partial}{\partial t} y_{A1}(t)}{J}$$

$$- \frac{k_{d2} x_{cg} \frac{\partial}{\partial t} y_{A2}(t)}{J} + \frac{k_{d1} x_{cg} \frac{\partial}{\partial t} y_{M1}(t)}{J} + \frac{k_{d2} x_{cg} \frac{\partial}{\partial t} y_{M1}(t)}{J} + \frac{d \theta_M(t) x_{cg} k_{s2}}{J} + \frac{d k_{d2} x_{cg} \frac{\partial}{\partial t} \theta_M(t)}{J}$$

Substitute the expression for the angular acceleration relative to the static equilibrium (Eq23) into the balance of forces equation (Eq20) and solve for the vertical acceleration relative to the static equilibrium (double derivative of $y_{M1}(t)$)

```
Eq24 := simplify(expand((subs(Eq20, Eq23)))):
Eq25 := solve(Eq24, diff(diff(y[M1](t), t), t), IgnoreSpecialCases):
Eq26 := diff(diff(y[M1](t), t), t) = Eq25[1];
```

$$\begin{aligned}
\frac{\partial^2}{\partial t^2} y_{M1}(t) = & - \left(J k_{d1} \frac{\partial}{\partial t} y_{M1}(t) - J k_{d2} \frac{\partial}{\partial t} y_{A2}(t) - J k_{d1} \frac{\partial}{\partial t} y_{A1}(t) + J k_{d2} \frac{\partial}{\partial t} y_{M1}(t) - J y_{A1}(t) k_{s1} \right. \\
& - J y_{A2}(t) k_{s2} + J y_{M1}(t) k_{s1} + J y_{M1}(t) k_{s2} + m k_{d1} x_{cg}^2 \frac{\partial}{\partial t} y_{M1}(t) + m k_{d2} x_{cg}^2 \frac{\partial}{\partial t} y_{M1}(t) \\
& - m y_{A1}(t) x_{cg}^2 k_{s1} - m y_{A2}(t) x_{cg}^2 k_{s2} + m y_{M1}(t) x_{cg}^2 k_{s1} + m y_{M1}(t) x_{cg}^2 k_{s2} + J d k_{d2} \frac{\partial}{\partial t} \theta_M(t) \\
& + J d \theta_M(t) k_{s2} - m k_{d1} x_{cg}^2 \frac{\partial}{\partial t} y_{A1}(t) - m k_{d2} x_{cg}^2 \frac{\partial}{\partial t} y_{A2}(t) - d m k_{d2} x_{cg} \frac{\partial}{\partial t} y_{M1}(t) \\
& + d m k_{d2} x_{cg}^2 \frac{\partial}{\partial t} \theta_M(t) - d^2 m k_{d2} x_{cg} \frac{\partial}{\partial t} \theta_M(t) + d m y_{A2}(t) x_{cg} k_{s2} - d m y_{M1}(t) x_{cg} k_{s2} \\
& \left. + d m \theta_M(t) x_{cg}^2 k_{s2} - d^2 m \theta_M(t) x_{cg} k_{s2} + d m k_{d2} x_{cg} \frac{\partial}{\partial t} y_{A2}(t) \right) / (J m)
\end{aligned}$$

Eq 23 and Eq26 are the final ordinary differential equations used to model the position of the vehicle frame

Bibliography

- [1] *Detroit Diesel, Series 53 Service Manual*. Detroit Diesel Corporation, 1990.
- [2] Jearl Walker, David Halliday, Robert Resnick, and J Walker. *Fundamentals of physics*. Wiley New York, 2008.
- [3] Lazaros Grammatikopoulos, George Karras, and Elli Petsa. Automatic estimation of vehicle speed from uncalibrated video sequences. In *Proceedings of International Symposium on Modern Technologies, Education and Professional Practice in Geodesy and Related Fields*, pages 332–338, 2005.
- [4] Daniel A Burgard, Gary A Bishop, Ryan S Stadtmuller, Thomas R Dalton, and Donald H Stedman. Spectroscopy applied to on-road mobile source emissions. *Applied spectroscopy*, 60(5):135A, 2006.
- [5] Marc A Meyers and Krishan Kumar Chawla. *Mechanical behavior of materials*, volume 547. Cambridge University Press Cambridge, 2009.
- [6] John R Schott. *Remote sensing*. Oxford University Press, 2007.
- [7] M. Silberberg. *Chemistry: The Molecular Nature of Matter and Change*. McGraw-Hill, 2000.
- [8] PL Richards. Bolometers for infrared and millimeter waves. *Journal of Applied Physics*, 76(1):1–24, 1994.

- [9] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [10] R.C. Gonzalez and R.E. Woods. *Digital Image Processing*. Pearson/Prentice Hall, 2008.
- [11] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [12] Jack D Gaskill. *Linear systems, Fourier transforms, and optics*. Wiley New York, 1978.
- [13] Giovanni Garibotto, Paolo Castello, Enrico Del Ninno, Pietro Pedrazzi, and Gianfranco Zan. Speed-vision: speed measurement by license plate reading and tracking. In *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*, pages 585–590. IEEE, 2001.
- [14] Junda Zhu, Liang Yuan, Yuan F Zheng, and Robert L Ewing. Stereo visual tracking within structured environments for measuring vehicle speed. *Circuits and Systems for Video Technology, IEEE Transactions on*, 22(10):1471–1484, 2012.
- [15] Thomas Zielke, Michael Brauckmann, and Werner Vonseelen. Intensity and edge-based symmetry detection with an application to car-following. *CVGIP: Image Understanding*, 58(2):177–190, 1993.
- [16] Massimo Bertozzi, Alberto Broggi, Alessandra Fascioli, and Stefano Nichele. Stereo vision-based vehicle detection. In *IEEE Intelligent Vehicles Symposium*, pages 39–44. Citeseer, 2000.
- [17] Emanuele Trucco and Konstantinos Plakas. Video tracking: a concise survey. *Oceanic Engineering, IEEE Journal of*, 31(2):520–529, 2006.
- [18] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Computer Vision–ECCV 2006*, pages 404–417. Springer, 2006.
- [19] Anne B Kahle and Ronald E Alley. Separation of temperature and emittance in remotely sensed radiance measurements. *Remote Sensing of Environment*, 42(2):107–111, 1992.

- [20] Carl Salvaggio and Craig J Miller. Comparison of field-and laboratory-collected midwave and longwave infrared emissivity spectra/data reduction techniques. In *Aerospace/Defense Sensing, Simulation, and Controls*, pages 549–558. International Society for Optics and Photonics, 2001.
- [21] GT Zheng and AYT Leung. Internal combustion engine noise analysis with time-frequency distribution. *Journal of Engineering for Gas Turbines and Power*, 124:645, 2002.
- [22] John Baechtel. *Performance Automotive Engine Math*, volume 204. CarTech Inc, 2011.
- [23] Jae-Chern Yoo and Tae Hee Han. Fast normalized cross-correlation. *Circuits, systems and signal processing*, 28(6):819–843, 2009.
- [24] Cyril M Harris. *Shock and vibration handbook*. McGraw-Hill, 1996.