



pennsylvania

DEPARTMENT OF TRANSPORTATION

Road Condition Reporting

FINAL REPORT

1/27/2015

Christoph Mertz

Carnegie Mellon University

Carnegie Mellon University

COMMONWEALTH OF PENNSYLVANIA
DEPARTMENT OF TRANSPORTATION

CONTRACT # CMUIGA2012
WORK ORDER # 002



1. Report No. FHWA-PA-2015-003-CMU WO 002		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle Road Condition Reporting				5. Report Date 1/27/2015	
				6. Performing Organization Code	
7. Authors Christoph Mertz, Courtney Ehrlichman, John Kozar, Srivatsan Varadharajan				8. Performing Organization Report No.	
9. Performing Organization Name and Address Carnegie Mellon University 5000 Forbes Avenue Pittsburgh, PA 15213				10. Work Unit No. (TRAIS)	
				11. Contract or Grant No. CMUIGA2012 – CMU WO 002	
12. Sponsoring Agency Name and Address The Pennsylvania Department of Transportation Bureau of Planning and Research Commonwealth Keystone Building 400 North Street, 6 th Floor Harrisburg, PA 17120-0064				13. Type of Report and Period Covered Final Report, Jan. 16 2014 to Jan 15 2015	
				14. Sponsoring Agency Code	
15. Supplementary Notes					
16. Abstract This report documents the findings of the road condition reporting project where the feasibility of live reporting of the road conditions with an Android camera and computer vision algorithms was tested. An app was developed that can collect videos or images and tag them with GPS and other information. During the 2013/14 snow season the system was mounted on two snow plows and data from several snow events was collected. With this data a computer vision algorithm was developed and tested that is able to detect the percentages of snow, slush or normal road in the images. From these percentages the road condition can be calculated. Overall this project demonstrated that it is feasible to develop a road condition monitoring system.					
17. Key Words snow plow, road condition, smartphone, computer vision				18. Distribution Statement No restrictions. This document is available from the National Technical Information Service, Springfield, VA 22161	
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 16	22. Price N/A

Notice

This work was sponsored by the Pennsylvania Department of Transportation and the U.S. Department of Transportation, Federal Highway Administration. The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the Federal Highway Administration, U.S. Department of Transportation, or the Commonwealth of Pennsylvania at the time of publication. This report does not constitute a standard, specification, or regulation.

Executive Summary

For the public and the agencies responsible for snow removal it is important to have up-to-date information about the current road conditions. A candidate for this task is a smartphone based system that captures live videos of the roads from a snow plow, sends it to a central location and analyzes the data to determine the road condition. In this work order we wanted to determine if this system was feasible for the desired purpose.

During the snow 2013/14 season two snow plows were fitted with Android cameras on their windshields. A data collection app was run during each shift to collect videos, GPS and other information. After the snow event the data was downloaded and complemented with Road Condition Reporting System (RCRS) road condition history logs. Additionally, data was labeled by PennDOT experts and by CMU. This data was used to develop and test algorithms that can infer the snow condition from images.

We looked at two broad categories of algorithms. In the first category the algorithms try to mimic human judgment, whenever a human judges an image to represent road condition X, the algorithm should classify the image as condition X. Practically this means that whatever labels the human provides is accepted as the ground truth. These labels are used to train a classifier. The output of the classifier (i.e. the road condition in an image) should then be similar to what a human would have judged. In the second category the algorithms try to make objective measurements, it tries to determine the percentages of clear road, slush, or snow that is seen in the image. We showed that both approaches are feasible. After some deliberation, PennDOT decided that we should follow the objective measure approach.

The outline of the analysis method is as follows. In a preprocessing step images at difficult locations (e.g. T-intersections) or with artifacts (e.g. windshield wiper in view) are discarded. Images which have other vehicles or artificial illumination from street lamps might also be discarded or might receive special treatment. For the remaining images the percentages of snow, slush and normal road in front of the snow plow are determined by analyzing texture, color and other features. A user chosen threshold on the percentages distinguishes between the road conditions. The conditions from all the images taken on one road segment are then summarized into one final road condition for that particular segment. We demonstrated the feasibility of this approach by training and testing on day and night images.

The most accurate method to determine lane-level localization is to use differential GPS (DGPS). If this is not practical, one can instead include the road condition of the adjacent lanes in the results.

Additional information from the internal snow plow data is useful for the road condition reporting system. Particular interesting are the outside air and road temperature. This internal data could be sent to the smartphone and from there to the central location, or alternatively it could be sent via an Automatic Vehicle Localization (AVL) system.

It should be mentioned that besides smartphones there are other suitable platforms for road condition reporting system as long as they capture and transfer images and location information. Overall this project demonstrated that it is feasible to develop a road condition monitoring system.

Table of Contents

- Introduction 3
- Task 1 and 2 summary 3
 - Data overview 3
 - Images 3
 - Other data 4
 - Labels 4
 - Choice of algorithm 5
 - Road Condition Analysis Method 5
- Localization improvements 9
- Connection to internal plow data 9
- AVL option 10
- Conclusion 10
- Appendix A: ATS data capture 11

Introduction

In order to best address deteriorating road conditions in inclement weather and report the information to the traveling public through traveler information services like 511PA, it is important to have up-to-date information about the road conditions. Currently, snow plow drivers are required to report this information every 2 hours during winter months when there is precipitation and temperatures are at or near freezing. Operators then capture the information in Road Condition Reporting System (RCRS) for further dissemination. It is desirable to automate this process so the information can be captured in real-time and the drivers can focus on their primary task of plowing. A candidate for this task is an Android camera-based road condition monitoring system. In this work order we wanted to determine if this system was feasible for the desired purpose.

In the first task of this work order we installed Android cameras on snow plows and collected videos and sensor data during snow events and added supplemental data after the collection. This data forms the basis for developing a full system that can be mounted in snow plows and reports the road conditions live. The results were reported in the “Task 1 report: Data Collection and Support.” In the second task of this work order we investigated various candidate methods to analyze the collected data. In the report “Task 2 report: Image Analysis Research and Development” we described the various methods and their performance.

In this final report we will summarize the findings of the task 1 and task 2 reports, and draw conclusions.

Task 1 and 2 summary

Data overview

Detailed description of the collected data can be found in the task 1 report. Here we discuss the data relevant for the image analysis, i.e. the images themselves and the labels for the images.

Images

Images were collected at all hours of the day. As one can expect, the day-time images look quite different than the night-time images. Three day-time images are shown in Figure 1. They show different road conditions and road types and one image has a car in its view.



Figure 1 Three day-time images.

Night images are much more challenging than images taken during the day because the illumination can vary greatly. Nine images with different illumination and weather conditions are shown in Figure 2.



Figure 2 Various night time images.

Dawn and dusk images are in-between the day and the night images. Three examples are shown in Figure 3.

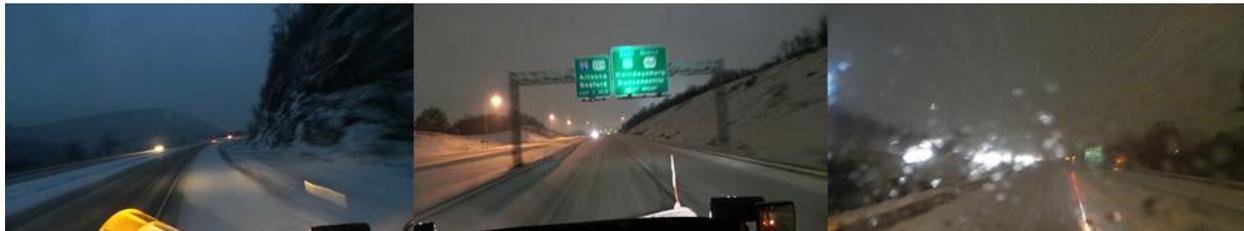


Figure 3 Three examples of images taken during dusk or dawn.

Other data

For each image we also captured a time stamp and the GPS location. Acceleration, gyroscope, and compass information were also recorded, but at this point they were not used in the analysis.

Labels

In order to train and test a classifier one needs labeled data. For our purposes, the labels are the RCRS road conditions 1 through 5 (with condition 5 being a manual entry) or the coverage of the viewed area (“road”, “slush”, or “snow”). The road condition label can be applied to a whole period (order of 1 hour), a video sequence (order of a minute), or a single image. The coverage label is applied to areas within an image.

We had different groups of people label different sets of data.

1. The RCRS history log gives the road conditions that were reported by the plow operators approximately every two hours. These records exist for all the data.

2. An unofficial “greatest hits” list of 50 videos were selected and labeled by experts from different county offices. They gave road condition labels for the start, middle and end of the video.
3. CMU gave road condition labels to images and gave coverage labels to areas of images.

Whenever more than one person labeled a data set we looked at the consistency between the labelers to estimate the accuracy of the labels.

Choice of algorithm

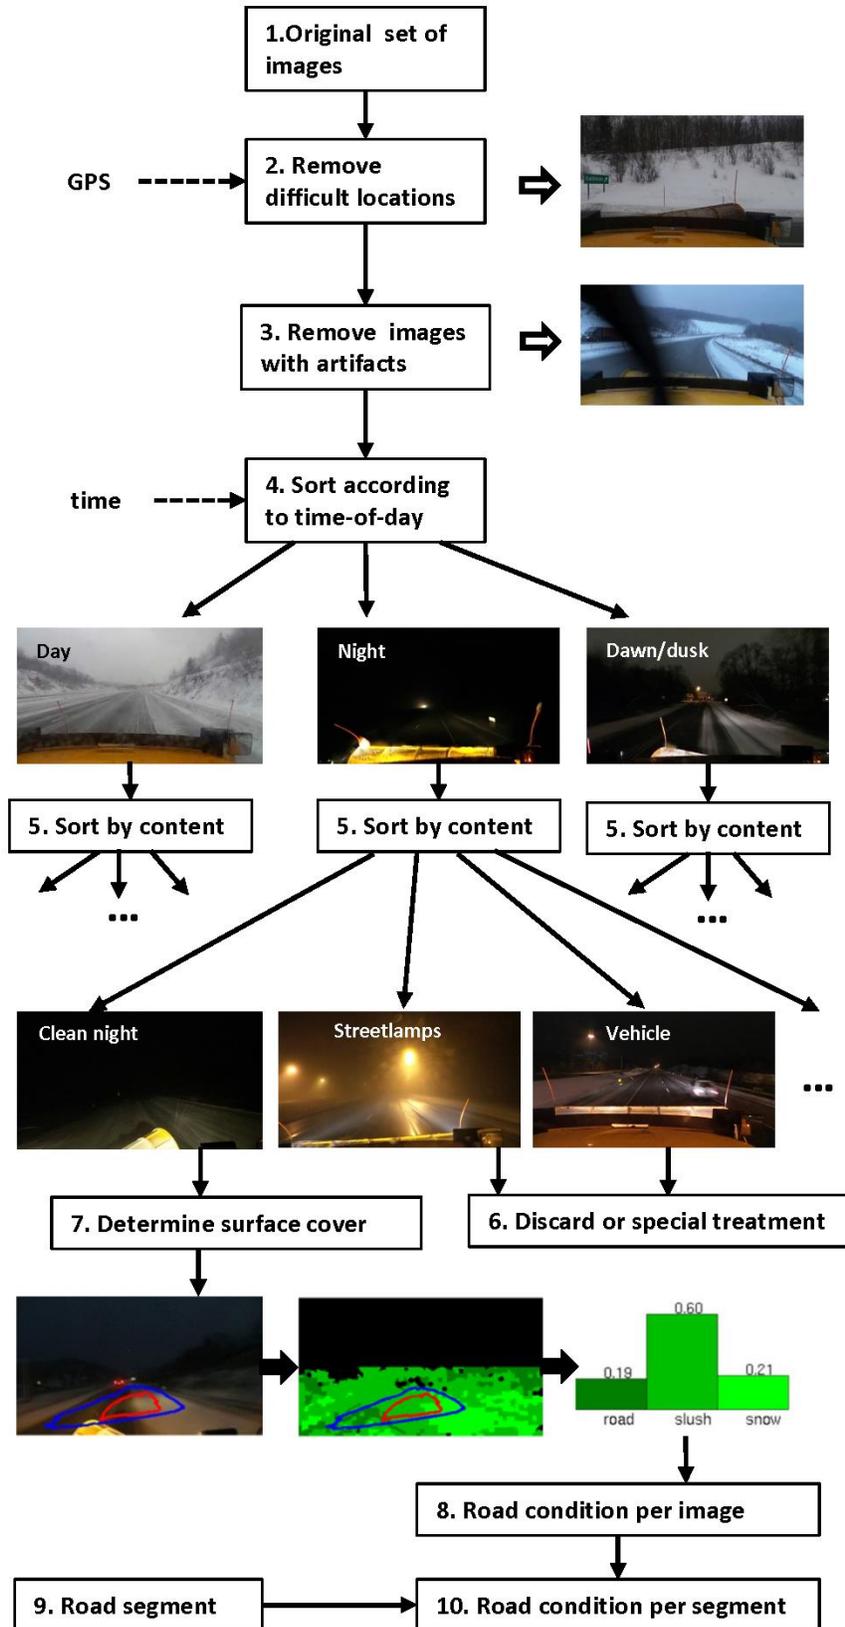
We looked at two broad categories of algorithms. In the first category the algorithms try to **mimic human judgment**, whenever a human judges an image to represent road condition X, the algorithm should classify the image as condition X. Practically this means that whatever labels the human provides is accepted as the ground truth. These labels are used to train a classifier. The output of the classifier (i.e. the road condition in an image) should then be similar to what a human would have judged. In the second category the algorithms try to **make objective measurements**, it tries to determine the percentages of clear road, slush, or snow that is seen in the image.

We showed that both approaches are feasible. After some deliberation, PennDOT decided that we should follow the objective measure approach. The reasons to use this approach are as follows:

1. The strength of computers is to measure things, e.g. it can accurately measure areas in an image and calculate percentages of the different coverage (road/slush/snow).
2. Human judgments are inconsistent. Even experts do not agree on the road conditions all the time (90% agreement on clear images, 80% agreement otherwise). Inconsistent labels will make it difficult to get a good classifier.
3. Different PennDOT Districts might want to have different thresholds for road conditions. The same amount of snow on the ground might be considered a more severe condition in southern areas where snow is much less common.
4. The information the public wants to get through 511 might be different than what the snow removal operators want. The public wants to know how long and how dangerous their trips will be. The snow removal operators need to plan the removal schedule, best road treatment, etc.
5. If someone wants to change how the road conditions are calculated it is straightforward to change the ROI or the thresholds on the percentages. In contrast, if the human judgment is changed a whole set of data needs to be relabeled and the classifier needs to be retrained.

Road Condition Analysis Method

The figure below shows the flow of the final road condition analysis method:



1. Original set of images

The images are taken by the camera. They might be subsampled or taken at certain time intervals to reduce the amount of data that needs to be transferred.

2. Remove difficult locations

At some locations the road is not in the view of the camera. The figure shows an example of a T-intersection. Once the difficult locations are identified, it is straightforward to exclude images that were taken in their vicinity by comparing GPS values.

3. Remove images with artifacts

An image where the windshield wiper is in the field of view (as shown in the Figure) or similar artefacts should be discarded. We already have a vision algorithm that discards images with windshield wipers.

4. Sort according to time-of-day

We found that images taken at different times of the day (day, night, dawn/dusk) need to be analyzed differently. At this stage in the process we will sort the images accordingly. Since we have for each image the time it was taken, this sorting is straightforward. We will have to decide if dawn/dusk will be one bin or if there needs to be a finer sorting. We also need to determine the best time thresholds (times before and after sunrise/sunset).

5. Sort by content

We need to further sort the images by their content. The most significant ones for night images are street lamps and other artificial illumination. One set of objects that might be present at any time is other vehicles. We do need to decide which contents to sort for. Certainly we will need to sort for illumination in night images. Maybe we need to sort for vehicles, but it might be that vehicles are treated implicitly in step 7 where they are classified as “unknown” and therefore ignored in the coverage calculation.

6. Discard or special treatment of certain images

The “clean” images will be analyzed as described in the next steps (7.-9.). Images with special content that do not occur frequently (e.g. large truck obstructing the view) or only intermittently can be discarded. Images with extra content that occur for extended periods of time need special treatment. We will first have to determine what these are; the most obvious candidates are long stretches with artificial illumination in populated areas and dense traffic during rush hour. The special treatment might be pre-processing images before moving to the next steps (7.-9.), training a classifier with different features, or creating masks for objects (e.g. masking out vehicles).

7. Determine surface cover

In this step each area of the image is classified into one of four classes: “road”, “slush”, “snow”, or “unknown”. The algorithm we developed for this first segments the image into superpixels and then classifies each superpixel as one of the four classes. The classifier uses >100 features and a SVM as the learning method. For a selected ROI the percentages of “road”, “slush”, and “snow” can then be

calculated. The ROI can be chosen by the user, e.g. it can be the region in front of the snow plow or it could include adjacent lanes.

8. Road condition per image

The road condition per image can now be determined by thresholds on the road/slush/snow percentages. This threshold can be chosen by the user.

9. Road segment

The boundaries of road segments can be chosen by the user or they can be determined from the road conditions. Boundaries could be exits, intersections or county limits. It could also be the locations where the road condition changes from one value to another.

10. Road condition per segment

The final output is the road condition per road segment. For each image we will find the closest road segment according to its GPS location. The road condition of the segment is then the average of the road conditions of the images that belong to the segment. Besides a single road condition number one could also report things like the variation of road conditions between the images.

The following table shows the different modules, their current status, and the amount of effort needed for building a full system in a 3-step approach. Step-2 is building a base system for implementation on PennDOT computers and Step-3 is building a full system.

Module	status	task 2	task 3	effort needed
2. Remove difficult locations	missing	x,c		1
3. Remove images with artifacts	exists for windshield wiper	c		2
4. Sort according to time-of-day	exists	c		1
5. Sort by content	missing		x	3
6. Discard or special treatment	missing		z	3
7a. Determine surface cover (day)	basic exists	y,c		2
7b. Determine surface cover (dawn/dusk/night)	basic exists		y	3
8. Road condition per image	basic exists	y,c		1
9. Road segment	missing	x,c		2
10. Road condition per segment	missing	x,c		2

x = develop from scratch, y = improve current version, z = develop exploratory code*, c = convert code for use on PennDOT computer

1-3 = some to lots of effort needed

*As we do not know all of the special cases and how to best deal with them we will first have to investigate the cases and test various methods.

Localization improvements

The standard GPS from smartphones has an accuracy of 10m-20m. That is sufficient to determine the road segment the snow plow is driving on. However, it is not sufficient to determine the lane. A typical lane is about 4m wide. It is necessary to know the lane because different lanes can have different road conditions. In the extreme case the main lanes of the interstate might be cleared while the snow plow is finishing up by removing the snow from the shoulders. The images would show heavy snow cover when in fact the traffic is flowing freely.

One can improve the GPS location by using other recorded data, mainly from accelerometer and gyroscope, and map data. This can smooth out outliers and ensure that the location is on the street, but it is not enough to determine the lane.

A different approach would be to **analyze the images to determine the lane** the truck is driving in. This would involve developing a whole new computer vision algorithm. The algorithm would probably work for some images, but not for all situations. Figure 4 shows example images with varying difficulties to determine the lane position from visual clues alone. Even for a human it is not possible to determine the lane position in the right image of Figure 4.



Figure 4 From left to right are images with increasing difficulty to determine the lane location from visual clues only.

Adding **differential corrections to the GPS** (DGPS) improves the accuracy to about 1m, enough to determine the lane. This correction is not available inside a smartphone. One would have to use a separate unit and somehow connect it to the smartphone. It might also be an option for Automatic Vehicle Localization (AVL) systems.

Another possibility is to not determine the lane the plow is driving in and instead **determine the condition of adjacent lanes**. First we find out if the road has more than one lane in the travel direction by looking up the current road on a map with the help of the GPS location. If it is only one lane, we assume that this is the lane we are traveling in. If there are more lanes, then we determine the road condition of the two adjacent lanes. To the traveling public we will report the lane with the best road condition. For the snow plow operations we can report all three conditions. Notice that when the truck is driving in the left-most (or right-most lane) there actually is no adjacent lane to the left (right).

Connection to internal plow data

There is data available on standard snow plows that are helpful in determining the road condition. Particularly useful are air and road temperatures. Knowing the speed of the vehicle can be used to

smooth out GPS or interpolate GPS where it is not available and thereby improve localization. Plow up/down and the status of the vehicle lights can help in the analysis of the captured images by priming the algorithm to expect an additional object in the field-of-view (the plow) and to expect certain illumination.

We tested if it is viable to feed this data into the smart camera. We got an Advanced Tracking System (ATS) unit that can connect to the snow plow data box and transmit the data via Ethernet or serial port. In “Appendix A: ATS data capture” we describe how we connected to the box via Wifi peer-to-peer and the data fields we could retrieve. The manufacturer’s software was used to receive and display the data on a Window’s PC. To our knowledge there is no such program for smartphones. We believe that it will take a considerable amount of work and collaboration with the manufacturer to make this data available to our app on the smart camera.

We also realized that this data is collected and sent to central operation by AVL systems. AVLs are commercially available and they are more robust and user-friendly than what we could develop ourselves. This option will be discussed in the next section.

AVL option

The three essential data fields that are needed for a road condition reporting system are image (or video), time, and location (GPS). For live reporting one also needs a communication link. The smart camera we have been using is an inexpensive, flexible solution. However, there are other platforms that are also suitable. Our algorithms and methods might need some tuning if the hardware changes, but in general it is agnostic about how the data is captured. A particular interesting option is Automatic Vehicle Localization (AVL). They are reliable, robust, tightly integrated into the vehicle and commercially available. Their localization capabilities should be better, because they can mount the GPS antenna outside the vehicle and they probably use better GPS units. Differential GPS (DGPS) might be available for some AVL systems, which would give around 1m position accuracy. They already have access to other data on the truck like outside and road temperature. What an AVL probably does not have compared to a smart phone/camera are acceleration, gyroscope, and compass information. But those are not essential. A crucial question is what cameras are available for AVL. At the moment we do not know if they are better or worse. The relevant properties are sensitivity, dynamic range, noise and resolution.

Conclusion

Overall this project demonstrated that it is feasible to develop a road condition monitoring system. The central part of such a system is the analysis of images captured from the operating snow plow. We believe that the most promising method is to first select out clean (i.e. no artificial lights or objects like vehicles) day, night and dawn/dusk image. For each time-of-day we will then train classifiers that calculate the percentages of road, slush, snow, or unknowns. The road conditions are then determined

by thresholds on the percentages. Images with artificial light, vehicles or other objects could be discarded or specially treated.

Additional data like air and road temperature will be very useful. Either our current app needs to be modified to receive this data or one could use an AVL. We have outlined the work that needs to be done to build a basic and a full system.

Appendix A: ATS data capture

Below are sample screen shots of a product from Certified Power Inc. called the Freedom ATS. It is a device used to capture various parameters from a snow plow i.e. vehicle speed, salt usage (spread rate), outside and road temperatures plus various other vehicle and plow information. The ATS control box is shown in Figure 5.



Figure 5 Top and side views of the Freedom ATS box.

The ATS Server software was installed and setup on a laptop for the collection of the data files from the ATS unit. Configuration for the network settings is as follows:

Network Settings	Database Settings
Host IP: 192.168.1.200	Log In: SgSIS_Admin
Protocol: UDP	Password: *****
Listen Port: 20000	Server Name: \SgSIS
Relay Port: 6650	Server IP: 0.0.0.0
Relay IP: 192.168.11.120	Server Port: 11015
<input checked="" type="checkbox"/> Acknowledge ATS	DataBase: SgSISV1
<input type="checkbox"/> Listen on Startup	Stored Procedure: LoadSpreaderReport
<input type="checkbox"/> Relay ACK Required	Database Timeout: 15
Data Destination	SQL Server Protocol: TCP/IP
<input checked="" type="checkbox"/> CSV File	Test Connection
CSV Output File Name: C:\Users\jkozar\ATSDAT	
<input type="checkbox"/> SQL Server	Display Data Filter
<input type="checkbox"/> IP Relay	Truck ID to Capture: Undefined
	<input type="checkbox"/> Apply Filter

The ATS unit has a “Digital Input Configuration” screen for various control inputs that lets you setup and assign a name to which the “Event log” will record the status:

Digital Input Configuration

Welcome to the digital input configuration web page.

From here, you can assign a function and configure the active input voltage level for any/all of the available digital inputs.

Input 1	Input 2	Input 3	Input 4
Function : Undefined	Function : Undefined	Function : Undefined	Function : Undefined
Name : <input type="text"/>			
Active State Low : <input type="radio"/> High : <input type="radio"/>	Active State Low : <input type="radio"/> High : <input type="radio"/>	Active State Low : <input type="radio"/> High : <input type="radio"/>	Active State Low : <input type="radio"/> High : <input type="radio"/>
Active Text : <input type="text"/>			
InActive Text : <input type="text"/>			
Severity Error : <input type="radio"/> Status : <input type="radio"/>	Severity Error : <input type="radio"/> Status : <input type="radio"/>	Severity Error : <input type="radio"/> Status : <input type="radio"/>	Severity Error : <input type="radio"/> Status : <input type="radio"/>
Counts : 0 <input type="text"/> Per : Undefined	Counts : 0 <input type="text"/> Per : Undefined	Counts : 0 <input type="text"/> Per : Undefined	Counts : 0 <input type="text"/> Per : Undefined

Input 5	Input 6	Input 7	Input 8
Function : Undefined	Function : Undefined	Function : User Defined	Function : User Defined
Name : <input type="text"/>	Name : <input type="text"/>	Name : Oil Satus	Name : Plow position
Active State Low : <input type="radio"/> High : <input type="radio"/>	Active State Low : <input type="radio"/> High : <input type="radio"/>	Active State Low : <input checked="" type="radio"/> High : <input type="radio"/>	Active State Low : <input checked="" type="radio"/> High : <input type="radio"/>
Active Text : <input type="text"/>	Active Text : <input type="text"/>	Active Text : CHECK OIL	Active Text : Up
InActive Text : <input type="text"/>	InActive Text : <input type="text"/>	InActive Text : Norm	InActive Text : Down
Severity Error : <input type="radio"/> Status : <input type="radio"/>	Severity Error : <input type="radio"/> Status : <input type="radio"/>	Severity Error : <input checked="" type="radio"/> Status : <input type="radio"/>	Severity Error : <input type="radio"/> Status : <input checked="" type="radio"/>
Counts : 0 <input type="text"/> Per : Undefined	Counts : 0 <input type="text"/> Per : Undefined	Counts : 0 <input type="text"/> Per : Undefined	Counts : 0 <input type="text"/> Per : Undefined

Record data on input state change :

Event Log

Welcome to the controller Event Log.

The Event Log contains lists of events broken down into status or general events and errors or more serious events. Should either log fill, the oldest data in the log will be overwritten.

Select one of the buttons to move through the event log. Refreshing or reloading the page will only refresh the currently displayed page.

Plow position Up [06/17/2014, 09:50:14]

Oil Satus CHECK OIL [06/17/2014, 09:50:17]

Oil Satus Norm [06/17/2014, 09:50:18]

Oil Satus CHECK OIL [06/17/2014, 09:50:18]

Also the “Installation Options” screen needs to be setup for the connection from the ATS unit and server (laptop):

Installation Options

Welcome to installation options.

This page allows the user to view and modify installation options. Typically these options would be modified only by a key operator or Service personnel.

Language : English	Stack Warning Threshold : 50 %	Keypad Port Number : 11016
Session Timeout : 2 minutes	Verbose Event Logging : <input type="checkbox"/>	Keypad Port Type : TCP
Customer Username : admin	Customer Password : <input type="text"/>	
Spreader Type : ACS	Breadcrumb Rate : 3 seconds	Resident Data Archive Size : 512 Kb
Record spreader status changes : <input checked="" type="checkbox"/>	Gps Installed : <input checked="" type="checkbox"/>	Removable Data Archive Size : 256 Mb
Default Truck Id : 7A9CE4		
Popup Help : <input checked="" type="checkbox"/>		

SERVER INTERFACE

Logical Interface : String Format 1 Include MAC Address : Data Transmit Rate : 0 seconds

Port Type : UDP Auto Server Selection : 10 seconds

ACK Timeout : 2 seconds

Primary	Secondary	Serial Port Settings
Physical Interface : Wired Ethernet	Wireless	Baud Rate : 19200
Target IP Address : 192.168.11.117	192.168.1.118	Databits : 8 Databits
Port Number : 11014	20000	Parity : None
		Flow Control : None

NOTE - You may have to reboot for some changes to take effect.

The “Wireless configuration” screen shows the setup for the wireless control of this unit through an “ADHOC” network with the name of “snowplow”:

