

Traffic measurement and vehicle classification with a single magnetic sensor

Sing Yiu Cheung, Sinem Coleri, Baris Dunder, Sumitra Ganesh, Chin-Woo Tan and Pravin Varaiya

Traffic Measurement Group

Department of Electrical Engineering and Computer Science
University of California, Berkeley, CA 94720

May 20, 2004

Abstract

Wireless magnetic sensor networks offer a very attractive, low-cost alternative to inductive loops for traffic measurement in freeways and at intersections. In addition to vehicle count, occupancy and speed, the sensors yield traffic information (such as vehicle classification) that cannot be obtained from loop data. Because such networks can be deployed in a very short time, they can also be used (and reused) for temporary traffic measurement.

This paper reports the detection capabilities of magnetic sensors, based on two field experiments. The first experiment collected a two-hour trace of measurements on Hearst Avenue in Berkeley. The vehicle detection rate is better than 99 percent (100 percent for vehicles other than motorcycles); and estimates of vehicle length and speed appear to be better than 90 percent. Moreover, the measurements also give inter-vehicle spacing or headways, which reveal such interesting phenomena as platoon formation downstream of a traffic signal.

Results of the second experiment are preliminary. Sensor data from 37 passing vehicles at the same site are processed and classified into 6 types. Sixty percent of the vehicles are classified correctly, when length is *not* used as a feature. The classification algorithm can be implemented in real time by the sensor node itself, in contrast to other methods based on high scan-rate inductive loop signals, which require extensive offline computation. We believe that when length is used as a feature, 80-90 percent of vehicles will be correctly classified.

Keywords: automatic vehicle classification, sensors, traffic surveillance, vehicle detectors, advanced traffic management systems

Executive summary

Wireless sensor networks bring together three technological innovations: micro-sensors, low-power radios, and low-power microprocessors, all powered by batteries. This report is part of a series that explores the use of wireless sensor networks for traffic monitoring.

These networks consist of a set of sensor nodes (SN) and one access point (AP). Each SN comprises a magnetic sensor, a microprocessor, a radio, and a battery. Each SN is encased in a 5"-diameter 'smart stud' that is glued to the center of a lane, where traffic is to be measured. Sensor measurements are processed within the SN and the results are transmitted via radio to the AP, which is located on the side of the road (figure 1).

Although the work is still in its experimental stage, the results indicate that these sensors are more accurate than inductive loops, and yield additional traffic information, such as headways and vehicle classification, not available with standard traffic measurement technology such as inductive loops, radar, and video.

The technology is sufficiently robust that the first small-scale field trials are scheduled for the summer of 2004. A private corporation, Sensys Networks, Inc., has attracted venture capital to develop vehicle detection systems based on this technology.¹

The current cost of these networks is about one-third that of inductive loops, and expected to rapidly decline with volume. Because such networks can be deployed in a very short time, they can also be used (and reused) for temporary traffic measurement. Because the sensor nodes are 'smart' they can assist automatic diagnosis, greatly facilitating maintenance.

A reliable and extensive traffic surveillance system is essential to measuring and improving the performance of traffic operations. Surveillance systems today rely on inductive loops. They are expensive to deploy and maintain, which has limited the extent of coverage and reliability. Alternatives, including radar and video, have not yet replaced loops. Wireless sensor networks offer a very low-cost and accurate solution that could serve as the foundation for an extensive, dense traffic monitoring system for both freeways and urban streets.

The research has overcome two major challenges to the use of wireless sensor networks for traffic monitoring. First, a communication protocol must be designed that provides extremely low power operation and guarantees delay bounds: For intersection traffic control, the presence of a vehicle must be reported within 0.1 sec; the battery-powered SNs must last 5-10 years. This challenge is met by the invention of the Pedamacs protocol, which operates the network in a time-synchronous manner and places the radios in a 'sleep' mode whenever a node is not in use. The Pedamacs protocol and calculations of its power requirements are presented in [1].

Second, appropriate sensors and low-power data processing algorithms must be developed to give the accuracy of loop detectors. We have explored the use of acoustic and magnetic sensors. The results of acoustic sensors are reported in [10].

This report presents the detection capabilities of magnetic sensors, based on two field experiments. The first experiment collected a two-hour trace of measurements on Hearst Avenue in Berkeley. The vehicle detection rate is better than 99 percent (100 percent for

¹ Professor Pravin Varaiya is a founder of Sensys.

vehicles other than motorcycles); and estimates of vehicle length and speed appear to be better than 90 percent. Moreover, the measurements also give inter-vehicle spacing or headways, which reveal such interesting phenomena as platoon formation downstream of a traffic signal.

Results of the second experiment are preliminary. Sensor data from 37 passing vehicles at the same site are processed and classified into 6 types. Sixty percent of the vehicles are classified correctly, when length is *not* used as a feature. The classification algorithm can be implemented in real time by the sensor node itself, in contrast to other methods based on high scan-rate inductive loop signals, which require extensive computation [2]. We believe that when length is used as an additional feature, 80-90 percent of vehicles will be correctly classified.

1. Introduction

Wireless magnetic sensor networks offer a very attractive alternative to inductive loops for traffic measurement in freeways and at intersections, in terms of cost, ease of deployment and maintenance, and enhanced measurement capabilities. Because such networks can be deployed in a very short time, they can also be used (and reused) for temporary traffic measurement.

These networks consist of a set of sensor nodes (SN) and one access point (AP). Each SN comprises a magnetic sensor, a microprocessor, a radio, and a battery. Each SN is encased in a 5"-diameter 'smart stud' that is glued to the center of a lane. Sensor measurements are processed within the SN and the results are transmitted via radio to the AP, which is located on the side of the road. The AP, housed in a 3"x5"x1" box attached to a pole or the roadside controller, comprises a radio, a more powerful processor, and a GPS receiver. The AP processes the data received from the SNs. The results are either sent to a local controller or to the TMC. The AP is either line- or solar-powered.

Figure 1 shows how such a network might be deployed. The little circles are the SNs and the square is the AP.




Figure 1 Deploying sensor networks at intersection and freeway

A sensor network conducts two operations: detection and traffic measurement, and communication. Communication is discussed in another report [1]. This report discusses how well a magnetic sensor can detect vehicles and estimate various traffic-related parameters.

The results of two experiments are presented. The first experiment involves a two-hour trace of measurements on Hearst Avenue, near the UC Berkeley campus. In all 332 vehicles were observed. The results are excellent: vehicle detection rate of 99% (100% if motorcycles are excluded), and estimates of vehicle length and speed that appear better than 90%. Because the measurements detect individual vehicles (rather than, say, counts and occupancies averaged over 30 sec), their analysis yields additional information, such as inter-vehicle spacing or headways and the formation of platoons downstream of a signal.

The second experiment is more limited and the results are preliminary. Measurements of the magnetic 'signatures' from 37 vehicles are processed to classify them into six types. The algorithm achieves a correct classification rate of 60 percent, in *real time*. The current algorithm does *not* use vehicle length as a feature in the classification scheme. Since it is an important feature [2], the correct classification rate should increase to 80-90 percent after incorporating length.

The report is organized as follows. Section 2 presents the results of the first experiment on vehicle detection and section 3 presents the results of the second experiment on vehicle classification. Section 4 provides some background of magnetic sensors to better explain why they are reliable and why the signature extracted from these sensors is better (more sensitive) than the one obtained from high scan-rate inductive loops. Section 5 presents some conclusions and outlines planned future work.

2. Vehicle detection

Measurements are from one sensor node placed in one lane of Hearst Avenue, indicated in Figure 2, on February 23, 2004, 8-9 pm. The node is glued to the middle of the lane. Ground truth was established by visual counting. Traffic was light, so the visual count is accurate.




Figure 2 Site of experiment with one sensor node

The 332 observed vehicles are classified as follows:

Type	Passenger vehicle	SUV	Van	Mini-truck	Bus	Truck	Motorcycle	Total
Number	248	48	18	9	4	0	5	332

2.1 Detection rate The sensor node detected 330 (99%) vehicles. The two undetected vehicles are believed to be motorcycles. Thus all non-motorcycle vehicles are detected. Interestingly, a motorcycle that passes near a node *is* detected. At locations where motorcycle detection is important, placing two nodes in a lane should ensure their detection.

The magnetic sensor measures $\text{mag}(z)$, the magnetic field in the vertical direction. The sensor's output is sampled at 128Hz, i.e. 128 times per second. The samples are compared with a threshold, resulting in a sequence of 1's and 0's. If 10 successive values are 1 (above the threshold), vehicle detection is declared. When the sample value subsequently falls below the threshold, the vehicle is declared to have passed the sensor. Thus the state machine coded in the SN processor sets a detection flag whose value is 1 for the time during which a vehicle is above the sensor, and whose value is 0 otherwise. Figure 3 displays the raw samples (left) from the passage of a single vehicle and the corresponding detection flag (right).




Figure 3 Raw samples and detection flag

2.2 Vehicle arrivals As a sequence of vehicles passes over the sensor, the detection flag produces a corresponding sequence of pulses. We call the time instant when the flag switches from 0 to 1 its *uptime*, and the instant when it switches from 1 to 0 its *downtime*. The interval between an uptime and the subsequent downtime is the *ontime*.

At a finer scale, figure 3 shows that the uptime occurs within 10 samples, i.e. in less than 0.1 s immediately after the front of the vehicle just crosses the sensor. Thus, the presence of a vehicle can be reported within 0.1 s to the controller—which is the requirement for a vehicle presence detector at an intersection. Second, the sum of the ontimes over a 30 s interval divided by 30 is the occupancy of a loop detector.

Thus a single sensor node produces measurements obtained from a single inductive loop in signal control and freeway traffic monitoring applications. Moreover, these measurements are made without a detector card used by a loop detector.

As will be seen below much more information can be extracted, because each vehicle is measured individually.

Each uptime indicates the arrival of a vehicle at the sensor. Figure 4 is a plot of vehicle arrivals at the sensor during the first 10 minutes of the experiment. (The entire two hour-long trace is not plotted, because the scale would be too small.)




Figure 4 Vehicle arrivals during the first 10 minutes

Each cross in the figure is the time when a vehicle just crosses the sensor node. Observe that the arrivals are clearly bunched together in platoons, formed by the clearing of the queue behind the signal during each green phase (see figure 2). Successive platoons are one minute apart, corresponding to the cycle time. The variability in platoon size implies that traffic is not saturated—the queue is cleared during each green phase. Typical signal control detection systems do *not* measure traffic downstream of a signal. However, the figure shows that such measurements can reveal how well the signal plan is adapted to the traffic demand. Note also that the information in the figure is produced in *real time* by the sensor node itself, with no additional processing.

It is also important to note that queue lengths at intersections can be measured by placing SNs upstream of the signal. This information would be of great value in setting signal control.

2.2. Vehicle length and speed The *ontime* or interval between successive uptime and downtime of the detection flag is the time during which a vehicle is above the sensor. Consider n successive vehicles, with measured ontimes t_1, t_2, \dots, t_n . To fix ideas, take $n = 11$, for reasons argued by Coifman et al [3]. Suppose the unknown lengths of these vehicles (in meters) are l_1, l_2, \dots, l_n . Suppose the (assumed) common unknown speed of these vehicles is v (m/s). Then the $(n + 1)$ unknowns l_1, l_2, \dots, l_n and v satisfy n equations,

$$(1) \quad t_i = l_i \times v, \quad i = 1, \dots, n.$$

Suppose we know the distribution $p(l)$ of vehicle lengths. Then we can obtain a maximum likelihood estimate of the vehicle speed, \hat{v} , and the vehicle lengths

$$(2) \quad \hat{l}_i = \frac{t_i}{\hat{v}}, i = 1, \dots, n.$$

As shown in [3], a robust estimate of the speed v is given by

$$(3) \quad \bar{v} = \frac{\bar{l}}{\bar{t}},$$

in which \bar{l} is the *median* vehicle length and \bar{t} is the *median* of the n observed ontimes t_1, t_2, \dots, t_n . We adopt this procedure, choosing $n = 11$ (as suggested in [3]) and also $n = 5$, for purposes of comparison. We take the median vehicle length as $\bar{l} = 5$ meters.

Figure 5 displays the individual vehicle speed estimates for a 5-point and an 11-point median. The latter estimate is smoother, as expected. With traffic flowing at the rate of 330 vehicles/hour, the passage of 11 vehicles takes about 2 minutes, so the 11-point estimate corresponds to a 2-min average. Under a heavier traffic flow of, say, 2,000 v/hr, this would be a 20-second average.

The headway is obtained if we subtract from the uptime (arrival) of a vehicle the downtime (departure) of the preceding vehicle. Figure 6 plots the headway in seconds for the first 10 minutes. Observe how the signal light creates departures in platoons, separated by large headways.

Using the speed estimate \bar{v} for \hat{v} in (2) the vehicle length estimates are

$$(4) \quad \hat{l}_i = \frac{t_i}{\bar{v}}, i = 1, \dots, n.$$

Figure 7 below is the histogram of vehicle lengths for the experiment.

2.3 Accuracy of estimates For freeway data, Coifman et al [3] find the standard deviation (σ) of a 10-point median speed estimate to be 2.5 mph. So we expect that with a probability of 0.95, the estimates differ from the true average speed by less than 2σ or 5 mph. We can measure the speed of an *individual* vehicle by placing two sensor nodes at a known distance, or by measuring the slew rate of an individual sensor's signal. These measurements will be carried out in the future.

As mentioned before, a better estimate of speed and length is provided by a maximum likelihood estimate, which requires knowledge of the length distribution $p(l)$. Alternatively, we can estimate this distribution, together with the speed. We will pursue this analysis in the future.




Figure 5 Vehicle speed using 5-point and 11-point median




Figure 6 Headway vs arrival time




Figure 7 Histogram of vehicle length (m)

3. Vehicle classification

One can estimate volumes of long vehicles (trucks) and short vehicles (cars) from 30-second average single-loop measurements of occupancy and counts [5][6]. However, classification of *individual* vehicles requires finer measurement.

This section reports results of a simple classification scheme based on a single dual-axis magnetic sensor, which measures the earth's magnetic field in both the vertical direction ($\text{mag}(z)$) and along the direction of the lane ($\text{mag}(x)$). Each of these measurements is sampled at 64Hz, i.e. 64 times per second. A vehicle's samples are processed and two pieces of information are extracted. First, the slope or rate of change of consecutive samples is compared with a threshold and declared to be +1 (-1) if it is positive and larger than (or negative with magnitude larger than) the threshold, or 0 if the magnitude of the slope is smaller than the threshold. The result is a 'hill pattern' that reveals the 'peaks' and 'valleys' in the vehicle's $\text{mag}(z)$ and $\text{mag}(x)$ signatures. The second piece of information is the largest value of the samples. Although in the results reported below, this information was generated off-line, it should be clear that it could be extracted in real time by the sensor node itself. Note that information about vehicle length is *not* used.

A simple algorithm uses this information to classify the vehicle into six types: passenger vehicles (1), SUV (2), Van (3), Bus (4), mini-truck (5), truck (6), and other (7).

Figure 8 below displays the raw samples and the hill pattern from four passenger vehicles.

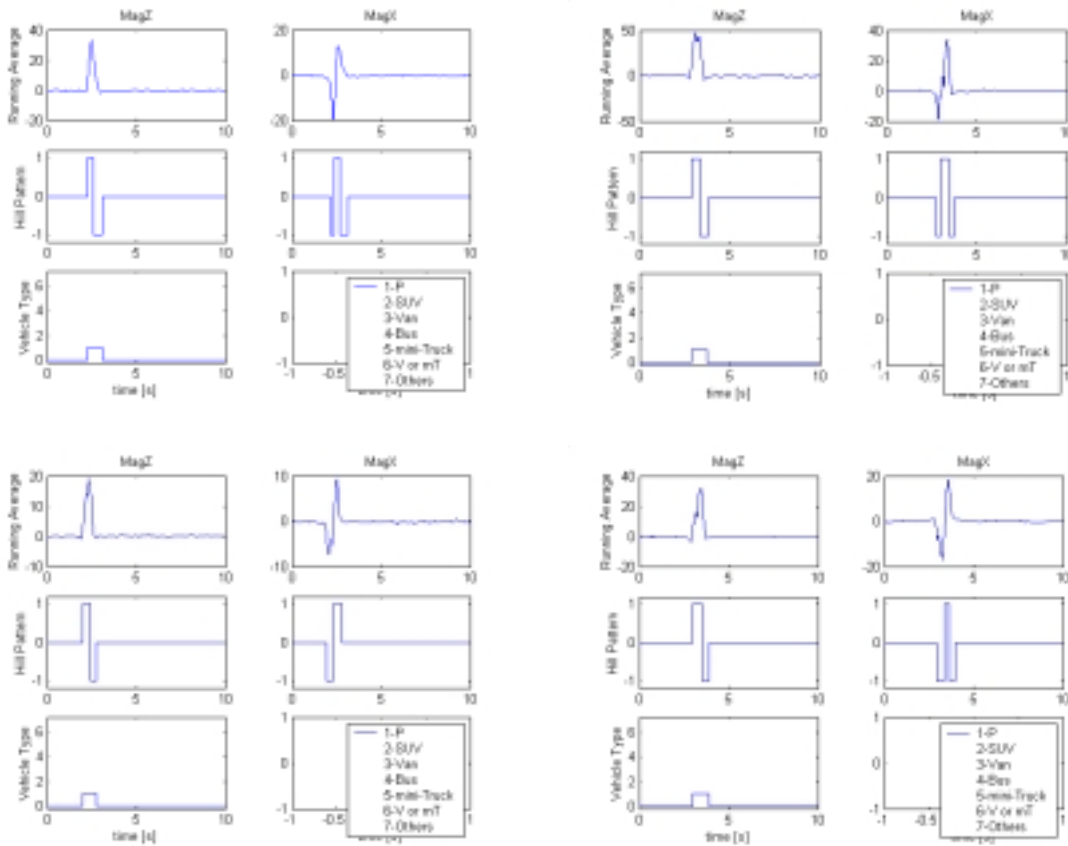


Figure 8 Raw samples and hill pattern from four passenger vehicles

The figure shows six plots for each vehicle. The top row shows the raw samples of $\text{mag}(z)$ and $\text{mag}(x)$. The second row shows the hill pattern. In each case the $\text{mag}(z)$ profile shows a single 'peak' revealed by the hill pattern (+1,-1) or one positive slope followed by one negative slope. In each case the $\text{mag}(x)$ profile shows one 'valley' followed by one 'peak' revealed by the hill pattern (-1,+1,-1). (The reason for the pattern is explained in the next section.) The plot on the third row gives the outcome of the algorithm, which classifies the signature into seven types. In each of these four cases, the algorithm decides that it is indeed a passenger vehicle (type 1).

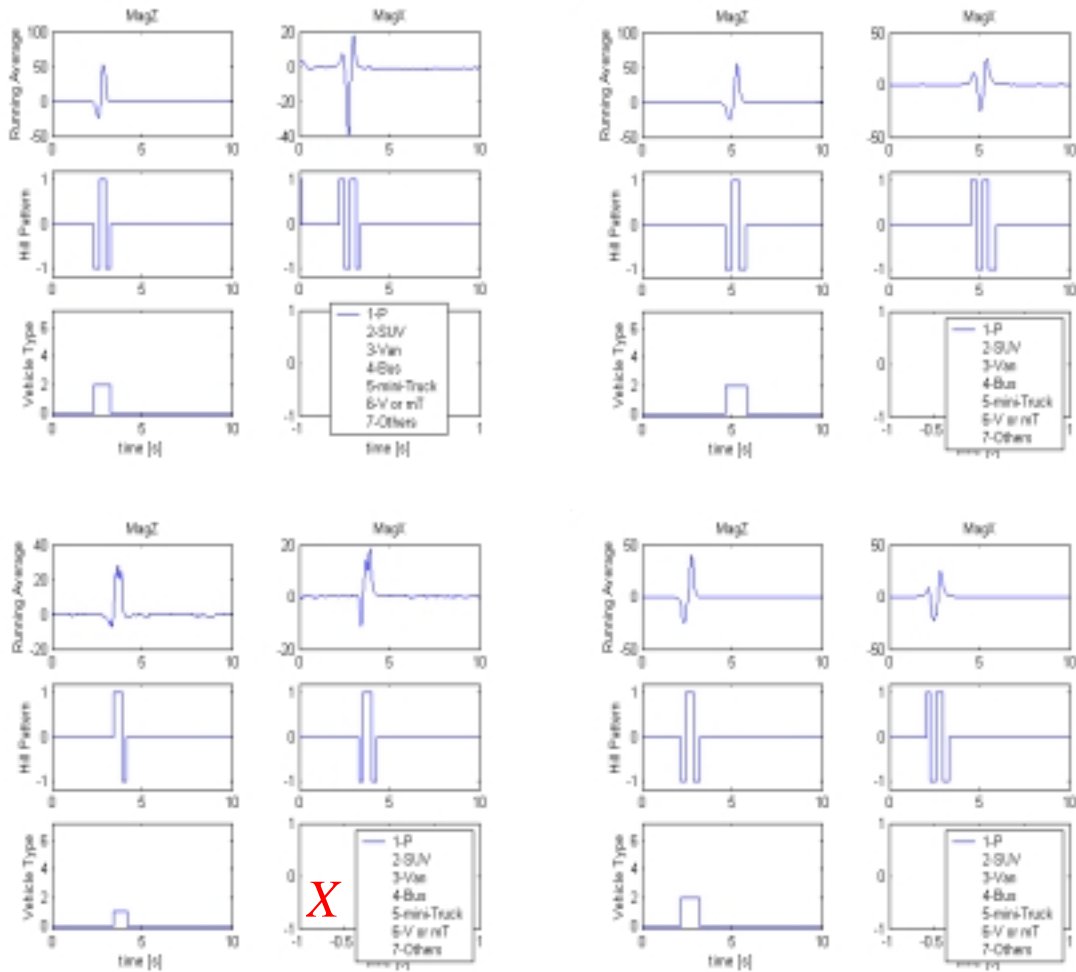


Figure 9 Raw samples and hill pattern from four SUVs

Figure 9 displays the raw samples, the hill pattern and the classification of four SUVs. In three cases (top two and bottom right) the patterns of both $\text{mag}(z)$ and $\text{mag}(x)$ are different from the passenger vehicle pattern of Figure 8. These patterns are $(-1, +1, -1)$ and $(+1, -1, +1, -1)$ respectively, both of which are different from the type-1 pattern.

The vehicle on the bottom left, however, is misclassified as a passenger vehicle, because the initial negative slope of $\text{mag}(z)$ is too small in magnitude to cross the threshold, and $\text{mag}(x)$ does not show the initial positive slope of the three other signatures.

The misclassification may be due to several reasons. The two peaks in $\text{mag}(x)$ in the three correctly classified SUVs are due to two masses of steel (ferrous material) separated by a significant gap. Such a distribution is not detected in the misclassified SUV. It is possible that a higher sampling rate might reveal the 'missing' peak. It is also possible that this SUV is built differently from the others, making its signature similar to that of a passenger vehicle. Also a lower $\text{mag}(z)$ threshold would reveal the small valley that is visible in the raw samples. Lastly, if SUVs are longer than passenger vehicles, then taking length as a feature, might lead to a correct classification.


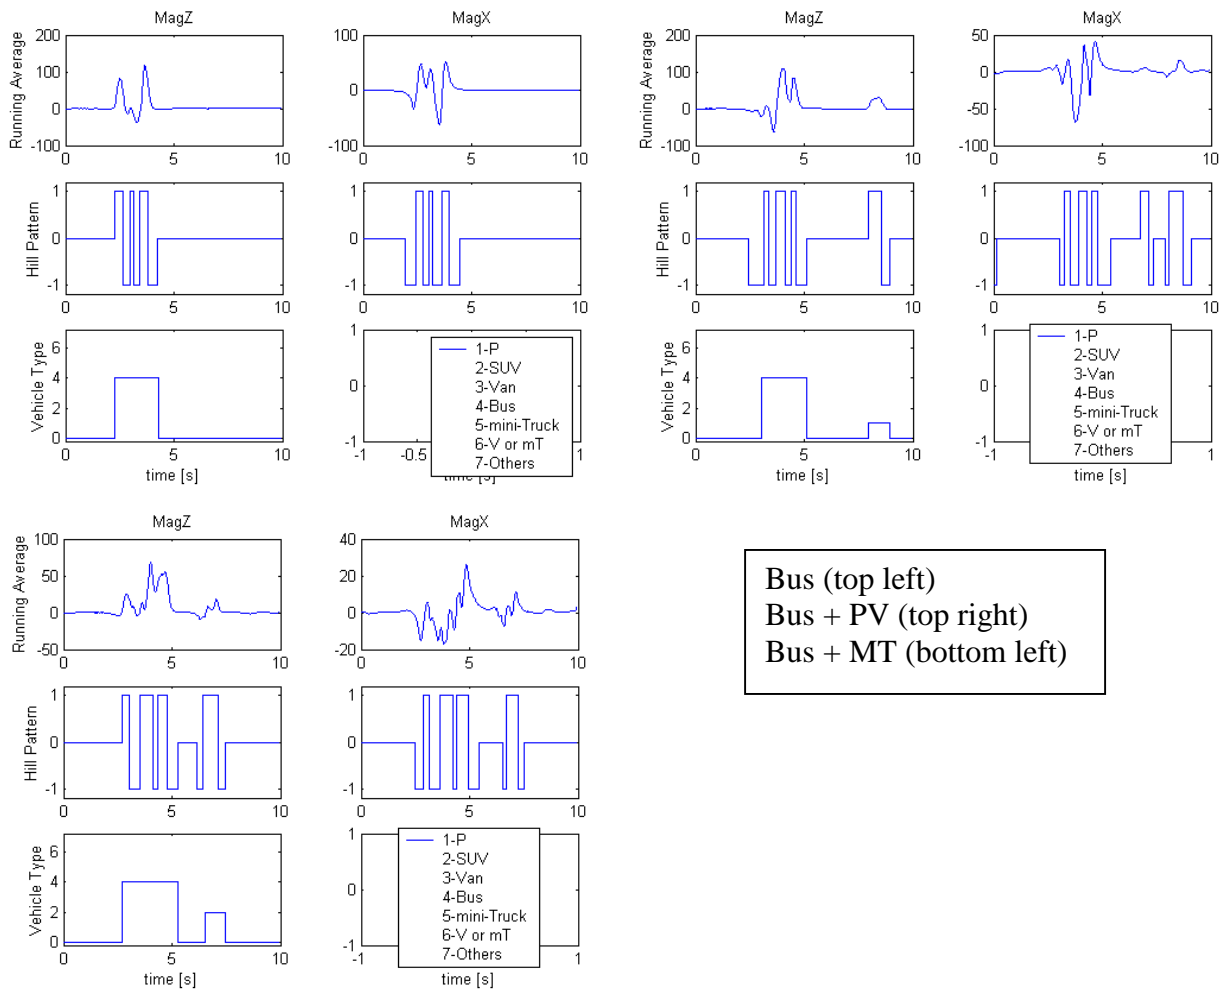


Figure 10 Raw samples and hill pattern from four vans

Figure 10 displays the raw samples and the hill pattern for four vans. Once again three of the four vehicles are correctly classified as Type 6, and one is misclassified. The van hill pattern for $\text{mag}(z)$ is $(+1, -1, +1, -1)$ and for $\text{mag}(x)$ is $(-1, +1, -1)$ (the same as that of a passenger vehicle). The $\text{mag}(z)$ pattern for the misclassified vehicle does have the two peaks, but the first peak is too small to pass the threshold test. Again, incorporation of length might have helped.

Figure 11 displays the raw samples and hill pattern from three buses. In two cases, the bus is followed by another vehicle. The $\text{mag}(z)$ pattern of $(+1, -1, +1, -1, +1, -1)$ or $(-1, +1, -1, +1, -1, +1, -1)$ has three peaks; moreover the raw $\text{mag}(z)$ samples have a much higher maximum value than those of the other vehicle types. The $\text{mag}(x)$ pattern in all three cases is $(-1, +1, -1, +1, -1, +1)$. Their very distinct patterns readily distinguish buses from the other vehicle types.

Figure 12 displays the raw samples and hill pattern from four mini-trucks (pickups). The distinguishing $\text{mag}(z)$ pattern is $(+1, -1, +1)$ (different from the preceding patterns) and the $\text{mag}(x)$ pattern is $(-1, +1, -1)$. Moreover the maximum value of $\text{mag}(z)$ is larger than those of passenger vehicles. However, in two of the four cases, the mini-truck is classified as a passenger vehicle. Within the data set of the experiment, classification of mini-trucks is the most challenging.



Bus (top left)
 Bus + PV (top right)
 Bus + MT (bottom left)

Figure 11 Raw samples and hill pattern from three buses

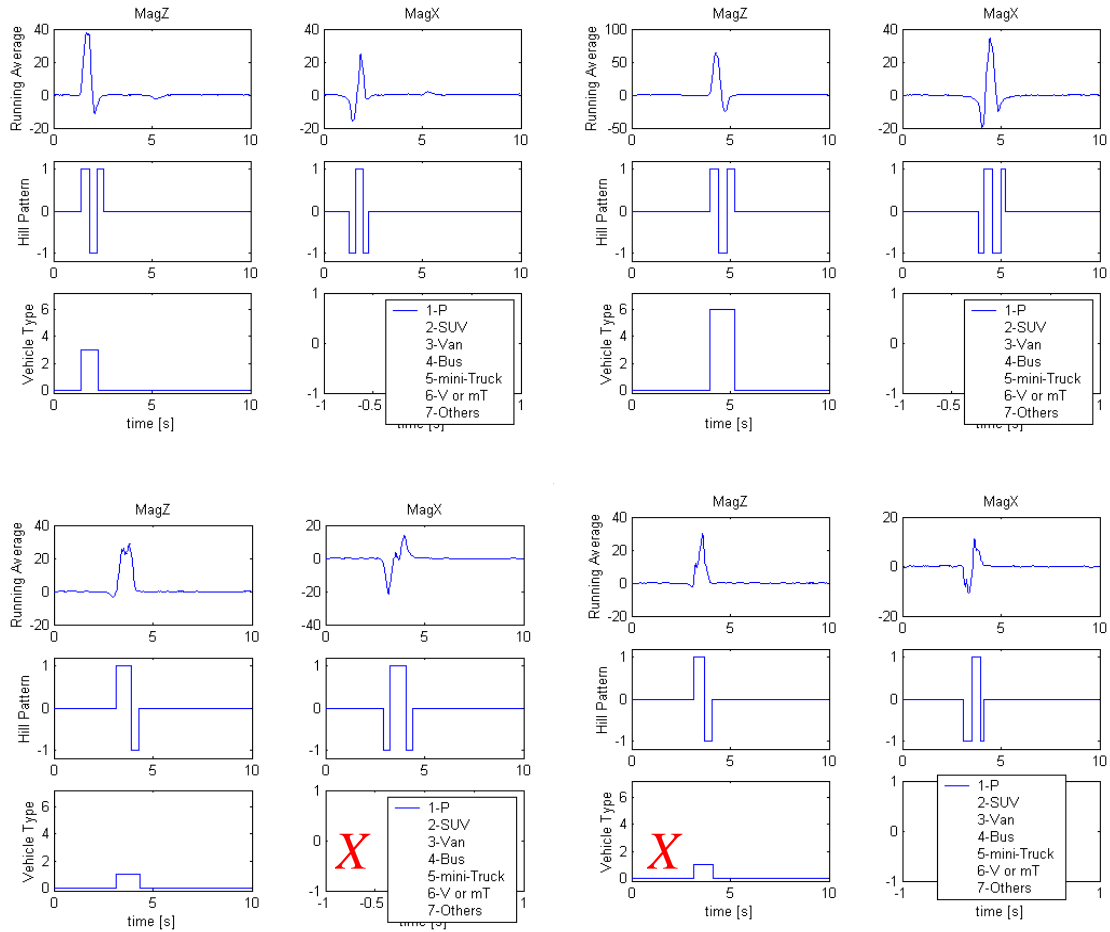


Figure 12 Raw samples and hill pattern from four mini-trucks

Total	PV	SUV	Van	Bus	MT	Van, MT
PV 15	11	4				
SUV 7	3	4				
Van 5	1			1		3
Bus 3				3		
MT 7	4	1			2	

Table 1 Classification using magnetic sensor

Table 1 indicates that 24 out of 37 vehicles (63 percent) were correctly classified. If mini-trucks are combined with passenger cars, 28 vehicles (75 percent) are correctly

classified. Obviously, the sample size is too small to make any firm judgment. Nevertheless, the technique is very promising, and some tentative conclusions appear warranted.

First, buses, vans and passenger vehicles are all correctly classified. The troublesome vehicles are SUVs and mini-trucks. The data did not contain any trucks, and further experiments will tell how well their signatures are distinguishable.

Second, and very importantly, the classification is based on measurements from a *single* sensor, without using length as a feature. Moreover, it should be clear that the algorithm can be implemented in real time. All loop-detector signature based classification schemes reviewed next require two loops (to extract speed and hence length), and the computations involved can only be carried off-line.

3.1 Comparison with loop signature-based classification

We compare the classification results based on the magnetic sensor signatures with three studies that use signatures from inductive loop signals scanned at a high rate (about 140 Hz). A survey of other studies that use inductive-loop signatures is provided in [2].

The studies use somewhat different categories, as seen in Table 2.

The loop-based studies are all based on pattern recognition methods. The computation requirements are such that a real-time implementation would be expensive. All these studies use length as an important feature.

Study	Type 1	Type 2	Type 3	Type 4	Type 3	Type 6	Type 7
[7]	PV, MV, MT, SUV	Small bus, van, pickup		Big bus, one-unit truck		Vehicle with trailer	
[2]	PV,	SUV, pickup,	Van	Bus	2-axle truck,	>:2-axle truck	Limousine
[4]	PV	SUV	Van	Truck	Pickup	Trailer	
Magnetic sensor	PV	SUV	Van	Bus	Mini-truck	Van, mini-truck	

Table 2 Classification studies: PV = passenger vehicle, MV=minivan, MT=mini-truck.

We pause to compare our scheme with the detailed discussion of pattern recognition schemes in Sun [2]. Two sets of schemes are tested. The first set, called ‘decision-theoretic’ methods, uses these features: length, largest magnitude of the measured inductance, variance and skewness (defined as the third or fourth central moment of the signal).

Significant pre-processing of the raw data is needed before the features can be extracted. As explained in [2], [9], (1) the signal magnitude must be normalized; (2) the ontime

must be multiplied by speed to convert the time axis into length; (3) a spline function must be interpolated through the raw samples and the result must be re-sampled so that each signature has an equal number of sample points. This pre-processing requires significant computational resources.

The decision-theoretic scheme uses a heuristic ‘decision tree’ like the one shown in Figure 13.




Figure 13 Decision tree for classification into 7 types of vehicles. Source [2].

The procedure works as follows. The six thresholds values b_1, \dots, b_6 are selected so as to give the best results for a ‘training set’ of vehicles. From Figure 13 we see that the length feature is compared with four thresholds b_1, \dots, b_4 to give five regions, namely, $(0, b_1), (b_1, b_2), (b_2, b_3), (b_3, b_4), (b_4, \infty)$. If the length falls within the middle three intervals, the vehicle is immediately classified as being of Type 3, 4 or 5. If the length falls into the first interval, it is classified as being of Type 1 or 2, and then the magnitude of the signal is compared with b_5 to determine its type (1 or 2). Similarly, if the length is larger than b_4 , the vehicle is declared to be of Type 6 or 7, and the ambiguity is subsequently resolved by comparing its skewness with the threshold b_6 . Observe the major role played by length in the classification.

As discussed in [2], different orders in which features are examined (first length, then magnitude in Figure 13) give different decision procedures. There is no theoretical reason to prefer one order to another.

The second set of classification schemes examined in [2] uses a neural network with many other features, including the discrete Fourier transform and Karhunen-Loeve transform coefficients, in addition to those listed earlier. The final scheme uses a total of 51 features. Once again, extraction of these features requires a great deal of computation.

The correct classification rate obtained by these schemes (for the seven classes listed in Table 2) is about 80 percent.

In summary, the following tentative conclusion appears warranted: magnetic signatures are better than inductive-loop signatures in terms of (1) computational burden, (2) improved sensitivity (speed and length are not used), and (3) implementation. The improved sensitivity is explained in the next section.

4. Magnetic sensor vs inductive loop measurements

The magnetic sensor used in these experiments is the Honeywell chip, HMC 1001/1002. These magneto-resistive sensors convert the magnetic field to a differential output voltage, capable of sensing magnetic fields as low as 30 μ gauss [11]. The earth's magnetic field is between 250 and 650 mGauss. A ferromagnetic material, such as iron, with a large permeability, changes the earth's magnetic field. The resulting change in the voltage is sampled at 128 Hz. The magnetic field is a three-dimensional vector.

In the vehicle detection experiment, we only measure the field in the vertical direction, $\text{mag}(z)$; for vehicle classification, we use both $\text{mag}(z)$ and $\text{mag}(x)$ —the change along the direction of the vehicle's motion. Figures 8-12 show the resulting waveforms.

The magnetic sensor is passive, and energy is consumed in the electronic circuit that measures the change in the resistance and the A/D conversion. By contrast, the inductive loop is an active device. A 6' by 6' copper loop is excited by a 20kHz voltage, creating a magnetic field. When a conducting material passes over the loop, the inductance is lowered. The loop detector card measures the change in the inductance. Special high scan-rate detector cards used for vehicle classification sample the inductance at about 140Hz.

The tiny magnetic sensor measures a highly localized change. As the vehicle travels over the sensor, it records the changes in the fields caused by different parts of the vehicle. By contrast, the 6' by 6' standard loop geometry results in the "integration of the inductive signature over the traversal distance ... which can remove distinctive features from the inductive signature [4]." So the standard loop is not ideal for vehicle classification. Figure 14 reproduces the inductive loop signatures of a pickup truck and a passenger car. Comparison with the magnetic signatures clearly shows that the latter provides much more detail.




Figure 14 Inductive loop signature from pickup truck (left) and passenger car (right). Source [7]

5. Conclusions and future plans

The limited experiments reported here strongly suggest that magnetic sensors provide traffic measurements that are both more accurate and more informative than from loop detector measurements. A single sensor provides count accuracy exceeding 99 percent, and average vehicle speed and length estimates better than 90 percent. Moreover, a single sensor can classify six types of vehicles with accuracy better than 60 percent. We believe that if two sensors are used, individual vehicle speeds can be very accurately measured, and vehicles can be classified with accuracy better than 80 percent. Significantly, all these estimates can be carried out in real time.

An earlier study [1] described a communication protocol that consumes so little power that a sensor node can be supplied by energy from two AA batteries for more than three years. More careful designs by Sensys indicate a lifetime exceeding seven years. The low-cost, ease of deployment and maintenance, and greater information of these sensor networks, suggest that they can serve as a foundation for an accurate, extensive, and dense traffic surveillance system.

For the immediate future, we plan work in several directions. First, we will conduct experiments with two nodes to measure individual vehicle speeds and improve vehicle classification. Of great interest is the classification of different kinds of trucks. We intend to carry out these measurements at a weigh-in-motion station.

We also believe that magnetic sensors can be placed on bridges and overpasses, where it is difficult to cut the pavement to install loop detectors. We hope to conduct such tests. The absence of detectors at these locations (where congestion often occurs) is a significant gap in freeway traffic monitoring.

Second, we will conduct more extensive experiments on urban streets with multiple lanes and higher volumes. In collaboration with Sensys, we will compare the measurements made by magnetic sensors with those reported by loops at an intersection.

Over the longer term, we will explore other sensing modalities, including temperature and fog sensors, and accelerometers. The interesting thing about the sensor network (figure 1) is that the same communication and node architecture can be used to process and communicate measurements from different sensors.

The PeMS project [12] has shown the value of traffic data for measuring and improving freeway performance. The project also shows how difficult it is to maintain California's loop detector system. Wireless sensor networks may provide the ideal low-cost, accurate traffic surveillance system needed to improve our transportation system.

6. Acknowledgement

This work was performed as part of the California PATH Program of the University of California, in cooperation with the California Department of Transportation. The contents of this paper 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 State of California.

7. References

- [1] S. Coleri and P. Varaiya. PEDAMACS: Power efficient and delay aware medium access protocol for sensor networks. PATH working paper. To appear.
- [2] C. Sun. An investigation in the use of inductive loop signatures for vehicle classification. California PATH Research Report UCB-ITS-PRR-2002-4, 2000.
- [3] B. Coifman, S. Dhoorjaty and Z.-H. Lee. Estimating median velocity instead of mean velocity at single loop detectors. *Transportation Research, Part C*, vol 11C, nos. 3-4, pp. 211-222, June-August 2003.
- [4] C. Oh, S.G. Ritchie and S.-T. Jeng. Vehicle re-identification using heterogeneous detection systems. 83rd TRB Annual Meeting, January 2004, Washington, D.C.
- [5] J. Kwon, A. Skabardonis and P. Varaiya. Estimation of truck traffic volume from single loop detector using lane-to-lane speed correlation. 82nd TRB Annual Meeting, January 2003, Washington, D.C.
- [6] X. Zhang, Y. Wang and N.L. Nihan. Monitoring a freeway network in real-time using single-loop detectors: System design and implementation. 83rd TRB Annual Meeting, January 2004, Washington, D.C.
- [7] S. Oh, S.G. Ritchie and C. Oh. Real time traffic measurement from single loop inductive signatures. 81st TRB Annual Meeting, January 2002, Washington, D.C.
- [8] A.N. Knaian. A wireless sensor network for smart roadbeds and intelligent transportation systems. M.S. thesis. Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 2000.
- [9] B. Abdulbahai and S.M. Tabib. Spatio-temporal inductive-pattern recognition for vehicle re-identification. *Transportation Research, Part C*, vol 11C, nos. 3-4, pp. 223-240, June-August 2003.
- [10] J. Ding. Vehicle detection by sensor network nodes. MS thesis. Department of Electrical Engineering and Computer Science. University of California, Berkeley, CA, Fall 2003.
- [11] <http://www.honeywell.com/magnetic/datasheet/magsen.pdf>
- [12] <http://pems.eecs.berkeley.edu/>